EDP Industry Rpt. 7/1/85 - View from the Valley
EDP Industry Rpt 7/31/85 Rollercoaster Revenue
Electronics 4/19/84 New System...
Electronics 4/14/86 - Tandem...
Financial Times Bs Info Ltd. 1/84 Computers in Banking
IBM Jrnl. 1987 - Structure of System/88
IEEE Computer Soc. Tech Comm Database Engr 6/83 - Quarterly Bulletin

Chilton I+Cs 10/87 - Reliability, Availability + fault tolerance
Chilton I+Cs 10/87 - Applying Fault Tolerant System Architecture
Codd, E.F. Articles
~~Compnies~~ Coming Computer Indust. Shakeout Book 1984
Cost of Network Ownership
Data Communications 10/83 - Solving the Problem
Data Communications 8/85 + 9/85 articles by Madsen + Foley
Datamation 1/85 + 11/85 IBM articles
Datamation 10/3/86 - Cowen - response by Tangney
DataPro 1984 - User Rating of Comp. Sys.
DEC World 9/21/87 Product Announcements
Drexel Burnham Lambert 2/4/85 T.Co Rpt
Drexel Burnham Lambert 12/9/83 - Entering new growth Phase

Articles about Tandem 1983-1987

~~Tandem Authors~~

Computer Language Mag. 1/85 - Exotic Lang. of Month - TAL
Computer Tech Review Summer '83 - Arch Determines Cost . . .
Computer World 1978 - Library Net . . .
Computer World 6/8/87 - 12 Rules for a Distributed Data base
The Commercial Computer Systems Indust - Alex Brown + Sons
Mail Message re articles 5/87
ACM Commun 4/70 - Nucleus of Multiprog. Per Brinch Hansen
ACM Comunications 8/85 - Case study: Cirrus Network
Auerbach 8/83 - Reprint on Prototyping
The Big Score - Book section on Tandem
Business Week 7/14/80 - What makes T. Run
Business Week 11/84 - Who's Excellent Now?
Business Week 4/21/86 - Computers: when will the slump End?

ARTICLES x AUTHOR
file - needs work

Article & Author
Files Index

CONTENTS - ARTICLE FILES
at the Corporate Information Center

December 27, 1989

This listing includes articles, reprints, reports, surveys, & large articles on Tandem and its staff, customers or competitors, fault-tolerance, venture capital, etc that are kept in the ARTICLE FILE. Included are papers in the AUTHOR FILE written by Tandem Staff members, excluding Tandem Technical Reports and Specifications. There is some cross referencing to items in the CLIPPING FILES.

NOTE: Entries on this list are crossed when possible by source, title and author.
In the AUTHOR FILE papers are filed by author's last name.
In the ARTICLE FILE papers are filed by source.
In the CLIPPING FILES the papers are filed chronologically.

Tandem Computers, Inc - Proprietary and Confidential

ACM, copyright 1980
"XRAY: Instrumentation for Multiple Computers", by Russ Blake
- SEE AUTHOR FILE

ACM, Communications of - April 1970
"The Nucleus of a Multiprogramming System", by Per Brinch Hansen
SEE ARTICLE FILE

ACM, Communications of - February 1978
"B-Trees Re-Examined", by Gerald Held
- SEE AUTHOR FILE

ACM, Communications of - February 1982
"Relational Database: A Practical Foundation for Productivity", by E F Codd
- SEE ARTICLE FILE

ACM, Communications of - August 1985
"Case Study: The CIRRUS Banking Network", edited by David Gifford and Alfred Spector
- SEE ARTICLE FILE

ACM Software Engineering Notes - July 1989
"The Three Questions", by Tom Van Vleck
- SEE AUTHOR FILE

"Advanced Program-to-Program Communciation in SNA", by Pete Homan
IBM Systems Journal - v22, n4, 1983
- SEE AUTHOR FILE

Alex. Brown & Sons
"The Commercial Computer Systems Industry" by Alex. Brown & Sons, Inc.
August 1989
- SEE ARTICLE FILE

"Applying Fault Tolerant System Architectures", by Jonathan A. Humphry
Chilton's - I&CS - Oct. 1987
- SEE ARTICLE FILE

"Architecture for High Volume Transaction Processing", by Robert W W Horst and Timothy C K Chou - Proceedings of the 12th Annual International Symposium on Computer Architecture (IEEE) - June 1985
- SEE AUTHOR FILE

Auerbach - August 1983 - "Prototyping"

Auerbach Publishers Inc
"An Overview of Fault-Tolerant Systems", By Patricia A Becker - Auerbach Publishers Inc - copyright 1986
- SEE AUTHOR FILE

Ball, Brian
Fault Tolerant Computing Exhibition, Technical Papers of (proceedings - held April 1987 - sponsored by MITRE - "Evolution of Tandem's Appro

to Fault Tolerance", by Brian Ball - paper dated March 1987
- SEE AUTHOR FILE

Bartlett, Joel F
"A NonStop Operating System" - copyright 1977
- SEE AUTHOR FILE

Bartlett, Wendy
"Fault Tolerance: Which Scheme is Best for You?" by Wendy Bartlett and Michael Treese - EDN July 13, 1989
- SEE AUTHOR FILE

Bate, Ruth
NCR Corporation Company Profile and Strategies - Competitive Analysis Report - Ruth Bates - Sept 25, 1987
- SEE ARTICLE FILE (under NCR)

Becker, Patricia A
"An Overview of Fault-Tolerant Systems", By Patricia A Becker - Auerbach Publishers Inc - copyright 1986
- SEE AUTHOR FILE

Ben-Ari, M
"Principles of Concurrent Programming" - book, by M Ben-Ari and Eiichi Watanabe - in Japanese - copyright 1982
- SEE AUTHOR FILE

"Big Score, The The Billion-Dollar Story of Silicon Valley" (book)
- have section on Tandem only
- SEE ARTICLE FILE

Bit - 1976 (v16)
"Tuning Garwick's Algorithm for Repacking Sequential Storage", by David S Wise and Dan C Watson
- SEE AUTHOR FILE

"Beyond OSI: Open Systems Networking", by David Kirk
- Telecommunications - October 1985
- SEE AUTHOR FILE

Blake, Russ
"Tailor: A Simple Model That Works" - 1979 Conference on Simulation, Measurement and Modeling of Computer Systems
- SEE AUTHOR FILE

Blake, Russ
"XRAY: Instrumentation for Multiple Computers" - copyright 1980 by the Association for Computing Machinery
- SEE AUTHOR FILE

Blake, Russ
"Optimal Control of Thrashing" - copyright 1982 by Tandem Computers
- SEE AUTHOR FILE

Borr, Andrea
"Transaction Monitoring in ENCOMPASS: Reliable Distributed Transaction Processing" - from the Proceedings of the Seventh International Conference on Very Large Databases, (IEEE) Sept 1981 - SEE AUTHOR FILE; SEE ALSO Tandem Technical Report 81.2

Borr, Andrea
"Robustness to Crash in a Distributed Database: A Non Shared-Memory Multi-Processor Approach" - from the Proceedings of the Tenth International Conference on Very Large Data Bases, Aug 1984 - - SEE AUTHOR FILE; SEE ALSO Tandem Technical Report 84.2

Brinch, Hansen
"The Nucleus of a Multiprogramming System" - ACM, Communications of - April 1970
- SEE ARTICLE FILE

"B-Trees Re-Examined", by Gerald Held
- Communications of the ACM - February 1978
- SEE AUTHOR FILE

BusinessWeek - July 14, 1980
"What Makes Tandem Run?"
- SEE ARTICLE FILE

BusinessWeek - November 1984
"Who's Excellent Now?"; cover story - followup on companies listed in book "In Search of Excellence" two years prior

BusinessWeek - April 21, 1986
"Computers: When Will the Slump End?" cover story by John W Wilson
- SEE ARTICLE FILE

BusinessWeek - February 23, 1987
"How Jim Treybig Whipped Tandem Back Into Shape", by Jonathan B Levine - SEE CLIPPING FILE

Carlyle, R Emmett
"IBM's Soothing Balm", by R Emmett Carlyle - Datamation - November 15, 1985
- SEE ARTICLE FILE

"Checkpointing Techniques for Fault-Tolerant Process-Pairs", by Gary Francis Tom - Submitted to the Department of Electrical Engineering and Computer Science in Partial Fulfillment of the Requirements for the Degree of Bachelor of Science at the Massachusetts Institute of Technology, June 1981"
- SEE AUTHOR FILE

Chester, Michael
"Fault-Tolerant Computers Mature", by Michael Chester - Systems & Software - March 1985
- SEE ARTICLE FILE

Chou, Timothy C K
"An Architecture for High Volume Transaction Processing" -
Proceedings of the 12th Annual International Symposium on Computer Architecture (IEEE) - June 1985
- SEE AUTHOR FILE

CIRRUS Banking Network
ACM, Communications of - August 1985
"Case Study: The CIRRUS Banking Network", edited by David Gifford and Alfred Spector
- SEE ARTICLE FILE

Codd, E F
"Does Your DMBS Run by the Rules?" , by E.F. Codd
- Computerworld - October 21, 1985
- SEE ARTICLE FILE

"Extending the Database Relational Model to Capture More Meaning", by E.F. Codd
- ACM Trans. on Database Systems - December 1979
- SEE ARTICLE FILE

"Inside IBM's Relational 'Strategy'", by E.F. Codd
- Computerworld - December 1986
- SEE ARTICLE FILE

"Is Your DBMS Really Relational?" by E.F. Codd
- Computerworld - October 14, 1985
- SEE ARTICLE FILE

"1987: The Beginning of a New Era in Data Processing" by E.F. Codd
- Information Week - May 4, 1987 - on NonStop SQL
- SEE CLIPPING FILES AND ARTICLE FILE

"The Real Strenghts of Relational Systems" by E.F. Codd
- Computerworld - February 6, 1984
- SEE ARTICLE FILE

"Relational Database: A Practical Foundation for Productivity",
- by E.F. Codd
- Communications of ACM - February 1982 -
- SEE ARTICLE FILE

"A Relational Model of Data for Large Shared Data Banks" by E.F. Cod
- Communications of the ACM - January, 1983
- SEE ARTICLE FILE

"Coming Computer Industry Shakeout" (book), by Stephen T McClellan
- 1984 - have pages on Tandem only - SEE ARTICLE FILE
- SEE ARICLE FILE

"The Commercial Computer Systems Industry" by Alex.Brown & Sons, Inc.
August 1989
- SEE ARTICLE FILE

Compcon - Spring 1977
"System Architecture for Nonstop Computing", by James Katzman
- SEE AUTHOR FILE

Computer Architecture, Proceedings from the 12th Annual Symposium (IEEE) - "Architecture for High Volume Transaction Processing", by Robert W Horst and Timothy C K Chou - - June 1985
- SEE AUTHOR FILE

Computer Language (magazine) - Jan 1985
"Exotic Language of the Month Club - TAL", by Serg Koren
- SEE ARTICLE FILE

Computer Technology Review - Summer 1983
"Architecture Determines Cost & Reliability of Fault-Tolerant Design", by Dennis McEvoy and Sandra Metz - SEE ARTICLE FILE

ComputerWorld - copyright 1978
"Library Net Puts Minis at Front and Back Ends" - deals with OCLC - SEE ARTICLE FILE

ComputerWorld - June 8, 1987
"Twelve Rules for a Distributed Data Base", by C J Date - ComputerWorld - June 8, 1987 - SEE ARTICLE FILE

"Computers in Banking and Finance The Impact of Fault Tolerant Systems" by Chris Summers - Financial Times Business Information Ltd - received Jan 1984 - SEE ARTICLE FILE

Computing Surveys - March 1977
"Multiprocessor Organization -A Survey", by Philip H Enslow Jr
- SEE AUTHOR FILE

"Concurrent SP/k A Language Supporting Concurrent Processes", by Reinhard G Menzel - thesis, University of Toronto - July 1976
- SEE AUTHOR FILE

Corporate Competitive Newsletter
-Monthly newsletter producced by Ralph Chiarella, Neil Johnson, Ken Bradley, Bruce Wood, Ruth Bates - started Oct. 2, 1987 - SEE ARTICLE FILE (under title)

"The Costs of Network Ownership" by Dr. Michael E. Treacy and Index Group, Inc. 1989
- SEE ARTICLE FILE

[Cowen Survey - listed under Datamation-Cowen]

Data Communications - Oct 1983
"Solving the Problems of Distributed Databases", by Jim Gray and Sandy Metz - SEE ARTICLE FILE; SEE ALSO AUTHOR FILE

Data Communications - Aug 1985
"When the Shoe's on the Vendor's Foot: a Look at Tandem's Corporate

Network", by Kent Madsen and David Foley - SEE ARTICLE FILE

Data Communications - Sept 1985
"How Multiprocessor Nodes Can Become More Sociable", by Kent Madsen & David Foley - SEE ARTICLE FILE

Data Communications - Feb 1986
"Time-Staged Delivery Networks Save Time, Enhance Productivity", by Jeri Edwards - SEE AUTHOR FILE

Datamation - January 1, 1985
"IBM: Mainframes in 1990", by Norman Weizer and Frederic Withington - SEE ARTICLE FILE

Datamation - April 1, 1985
"A Measure of Transaction Processing Power", by Anon et al
- SEE AUTHOR FILE (Jim Gray) - SEE ALSO Tandem Technical Report 85.2

Datamation - November 15, 1985
"IBM's Soothing Balm", by R Emmett Carlyle - SEE ARTICLE FILE

Datamation-Cowen 1983/1984 Mini/Microcomputer Conference - section on customer loyalty. Audrey Fricke has the complete conference/survey - SEE ARTICLE FILE

Datamation-Cowen 1984/85 Mini/Microcomputer Conference (complete) - SEE ARTICLE FILE

Datamation-Cowen 1985 Computer/Telecommunications Industry Conference - SEE ARTICLE FILE

Datamation-Cowen 1985/1986 Mini/Microcomputer Conference - SEE ARTICLE FILE

Datamation-Cowen Survey 1988/89 Mini/Microcomputer conference - SEE ARTICLE FILE

Datamation-Cowen - mail memo from Cacey Tangney, "More on Customer Satisfaction" - Oct 3, 1986 - in response to "unfavorable" survey - SEE ARTICLE FILE under Datamation-Cowen

Datapro 1984 User Rating of Computer Systems (mini & mainframe sections only) - Marketing Info Center has complete 1984 survey plus surveys from previous years - SEE ARTICLE FILE

Date, C J
"Twelve Rules for a Distributed Data Base",
by C J Date - ComputerWorld - June 8, 1987 - SEE ARTICLE FILE

Dec Professional Jan. 1988
"OLTP on the Vaxcluster" by Berl Hartman.

" Online Transaction Processing - What is it? And should I care?"
by Phillip A Naecker

DECWorld Product Announcements
"This publication describes the new product announcements at the recent DECWorld show in Boston." From Neil Johnston, Sept 21, 1987 - SEE ARTICLE FILE

"Designer's Overview of Transaction Processing", by Lloyd Smith, copyright 1979 - SEE AUTHOR FILE

Drexel Burnham Lambert, Inc - Dec 9, 1983
"Tandem Computers Entering New Growth Phase" - SEE ARTICLE FILE

Drexel Burnham Lambert, Inc - Feb 4, 1985
"Tandem Computers Inc - Company Report - Outlook Improving" - SEE ARTICLE FILE

Duke, Chris and Sandy Metz
"Message-Transport Software Does the Whole Job" - Electronic Design - July 8, 1982 - SEE AUTHOR FILE

EDN - July 13, 1989 - "Fault Tolerance: Which Scheme is Best for You?" by Wendy Bartlett and Michael Treese - SEE AUTHOR FILE

EDN - April 28, 1988 - "Silicon compilers Tame 10,000-Gate-Plus ASICs, Gate Arrays" by William J Stenzel - SEE AUTHOR FILE

EDP Industry Report - July 1, 1985 - Review & Forecast Part I - "View from the Valley: Large-Scale Systems Coming off Slower Growth, Facing Lower Income Levels in 1985 Despite Continuing Demand" - SEE ARTICLE FILE

EDP Industry Report - July 31, 1985 - Review & Forecast Part II - "Rollercoaster Revenue Climbs and Slides Unsettle Stomachs for Medium-, Small- Scale System Vendors" - SEE ARTICLE FILE

Edwards, Jeri
"Time-Staged Delivery Networks Save Time, Enhance Productivity" - Data Communications - Feb 1986 - SEE AUTHOR FILE

Electronic Business - May 1, 1984
"A Faultless Market Draws Many New Contenders"

This is in the folder with:
Mini-Micro Systems - Jan 1984
"Competition Heats Up.", by Susan Glazer - SEE ARTICLE FILE
They are essentially the same article.

Electronic Business - February 15, 1987
"Tandem: Growth in a Niche That Never Fails", by Anne Knowles - SEE CLIPPING FILE

Electronic Design - July 8, 1982
"Message-Transport Software Does the Whole Job", by Chris Duke and

Sandy Metz - SEE AUTHOR FILE

Electronics - April 19, 1984
"New System Manages Hundres of Transactions Per Second", by Robert Horst and Sandra Metz - SEE ARTICLE FILE

Electronics - April 14, 1986 - cover story
"Tandem Makes a Good Thing Better", by Jonah McLeod and "Tandem's Old Architecture Pays Off In New Markets", by Clifford Barney (articles are attached) - SEE ARTICLE FILE

Enslow, Philip H Jr
"Multiprocessor Organization-A Survey" - Computing Surverys - March 1977 - SEE AUTHOR FILE

"Evolution of Tandem's Approach to Fault Tolerance" - Fault Tolerant Computing Exhibition, Technical Papers of (proceedings) - held April 1987 - sponsored by MITRE - "Evolution of Tandem's Approach to Fault Tolerance", by Brian Ball - paper dated March 1987 - SEE AUTHOR FILE

"Fault-Tolerant Computer Systems Market Analysis and Forecast United States Market (1984-1990)" - Frost and Sullivan - Jan 1986 - SEE ARTICLE FILE

"Fault Tolerant Computers Market in Europe" (also titled "Fault Tolerant Computer Systems European Market Analysis and Forecast 1983-1990") - Frost and Sullivan - Aug 1985 - SEE ARTICLE FILE

"Fault-Tolerant Computers Mature", by Michael Chester - Systems & Software - March 1985 - SEE ARTICLE FILE

"Fault-Tolerant Computers US Market Analysis and Forecast (1983 - 1988)" - Frost and Sullivan - Aug 1984 - SEE ARTICLE FILE

Fault Tolerant Computing Exhibition, Technical Papers of (proceedings) - held April 1987 - sponsored by MITRE - "Evolution of Tandem's Approach to Fault Tolerance", by Brian Ball - paper dated March 1987 - SEE AUTHOR FILE

"Fault-Tolerant Computing System", by James A Katzman - copyright 1977 - first revision January 1979 - SEE AUTHOR FILE

"Fault Tolerant Control for Intelligent Manufacturing", by Glen B. Alleman - IMS October 1985 Proc. - SEE ARTICLE FILE

"Fault-Tolerant System" - book, by Eiichi Watanabe - in Japanese. This is a colection of technical papers/articles published in the past, with addition information such as a draft proposal of Japanese Industry Standard of Highly Reliable Computing System. Papers were written by various Tandem and non-Tandem people. Book received in Oct 1986. - SEE AUTHOR FILE

"Fault-tolerant System Architecture of Tandem NonStop System", presentation by Eiichi Watanabe - Jan 1984 - abstract in English, paper in Japanese - SEE AUTHOR FILE

"Fault-Tolerant Systems"
- report from ITOM International Co - copyright 1982
- SEE ARTICLE FILE

Financial Times Business Information Ltd - received Jan 1984
"Computers in Banking and Finance The Impact of Fault Tolerant Systems" by Chris Summers - SEE ARTICLE FILE

Frost and Sullivan - Aug 1984
"Fault-Tolerant Computers US Market Analysis and Forecast (1983 - 1988)" - SEE ARTICLE FILE

Frost and Sullivan - Aug 1985
"Fault Tolerant Computers Market in Europe" (also titled "Fault Tolerant Computer Systems European Market Analysis and Forecast 1983-1990") - SEE ARTICLE FILE

Frost and Sullivan - July 1988
"Fault Tolerant Computer Systems in Europe 1986-1992" - SEE ARTICLE F

Frost and Sullivan - Aug 1988
"Fault Tolerant Computer Systems in the U.S." - SEE ARTICLE FILE

Frost and Sullivan - Jan 1986
"Fault-Tolerant Computer Systems Market Analysis and Forecast United States Market (1984-1990)" - SEE ARTICLE FILE

Gray, Jim
"A Transaction Model" - paper - Feb 1980 - SEE AUTHOR FILE

Gray, Jim
"Solving the Problems of Distributed Databases" by Jim Gray and Sandy Metz - Data Communications - Oct 1983 - SEE AUTHOR FILE

Gray, Jim
"A Measure of Transaction Processing Power", by Anon et al - Datamation - April 1, 1985 - SEE AUTHOR FILE - SEE ALSO Tandem Technical Report 85.2

Gray, Jim
"Transparency in its Place" by Jim Gray - Unix Review - May 1987 - SEE ARTICLE FILE; SEE ALSO AUTHOR FILE

"Guardian/Expand: A NonStop Networ", by James A Katzman and H Robert Taylor - copyright 1978 -SEE AUTHOR FILE

Hamilton, Paul
"Relialbility, Availability, and Fault Tolerance" - Chilton's I&CS - Oct 1987 - SEE ARTICLE FILE

Haskell, George
Preliminary Analysis of Stratus Product Announcements April 18, 1984
- SEE CLIPPING FILES

Helland, Pat
Transaction Monitoring Facility (TMF) - Quarterly Bulletin of the IEE Computer Society Technical Committee on Database Engineering - June 1985
- SEE AUTHOR FILES

Held, Gerald
"B-Trees Re-Examined" - Communications of the ACM - Feb 1978
- SEE AUTHOR FILE

Holbrook, Robert
"New RDBMS Dispel Doubts, Perform OLTP Applications" - Computer Technology Review - Spring 1988
- SEE AUTHOR FILE

Homan, Pete
"Advanced Program-to-Program Communciation in SNA" - IBM Systems Journal - vol 22, n 4, 1983 - SEE AUTHOR FILE

Horst, Robert W
"An Architecture for High Volume Transaction Processing", by Robert W Horst - IEEE Proceedings of the 12th Annual International Symposium on Computer Architecture - June 17-19, 1985

Humphry, Jonathan
"Applying Fault Tolerant System Architectures" - Chilton's I&CS Oct 1987 - SEE ARTICLE FILE

"IBM's Soothing Balm", by R Emmett Carlyle
- Datamation - November 15, 1985 - SEE ARTICLE FILE

IBM Systems Journal - v22, n4, 1983
"Advanced Program-to-Program Communciation in SNA", by Pete Homan
- SEE AUTHOR FILE

IBM Systems Journal - v26, n3, 1987
"Structure of System/88, a Fault-Tolerant computer" by E.S. Harrison
- SEE ARTICLE FILE

IEEE Quarterly Bulletin of the IEEE Computer Society Technical Committee on Database Engineering - June 1983 ("This issue of Database Engineering focuses on highly available systems. The first five articles are from commercial vendors who market fault-tolerant computer systems.") - SEE ARTICLE FILE

Industry Week - October 15, 1984
"IW's Excellence in Management Awards", includes Jim Treybig.
- SEE ARTICLE FILE

Information Week - May 4, 1987
"1987: The Beginning of a New Era in Data Processing", by E F Codd - Information Week - May 4, 1987 - on NonStop SQL - SEE CLIPPING FILES

Information Week - November 9, 1987
"Tandem's Treybig: No More Mister Niche Guy", by Paul E. Schindler - Information Week - November 9, 1987 - SEE CLIPPING FILES

ITOM International Co
- "Fault-Tolerant Systems - copyright 1982 - SEE ARTICLE FILE

Japan Electronic Industry Newsletter
- monthly newsletter sent by email by Eiichi Watanabe beginning April 1987 - SEE ARTICLE FILE (under title)

Japan Technology and Industry Report - quarterly report sent by email by Eiichi Watanabe beginning June 30, 1987 - SEE ARTICLE FILE (under title)

Johnston, Neil
DECWorld Product Announcements - "This publication describes the new product announcements at the recent DECWorld show in Boston." From Neil Johnston, Sept 21, 1987 - SEE ARTICLE FILE

Katzman, James
"System Architecture for Nonstop Computing", by James A Katzman - Compcon Spring 1977 - SEE AUTHOR FILE

Katzman, James A
"Fault-Tolerant Computing System", by James A Katzman - copyright 1977 - first revision January 1979 - SEE AUTHOR FILE

Katzman, James A
"Guardian/Expand: A NonStop Network", by James A Katzman and H Robert Taylor - copyright 1978 - SEE AUTHOR FILE

Kirk, David
"Beyond OSI: Open Systems Networking", by David Kirk - Telecommunications - October 1985 - SEE AUTHOR FILE

Laurich, Larry
"Product Development Cycle" - video tape, by Larry Laurich - Nov 11, 1980 - 35 min - SEE AUTHOR FILE

Levering, Robert
100 Best Companies to Work for In America (book) - Have section on Tandem only - SEE ARTICLE FILE

"Library Net Puts Minis at Front and Back Ends"
- ComputerWorld - copyright 1978 - deals with OCLC - SEE ARTICLE FILE

Mackie, Dave
"The Tandem 16 NonStop System" - no source, no date - SEE AUTHOR FILE

Madsen, Kent
"NonStop Transaction Processing" - by Lloyd Smith and Kent Madsen - Datamation March 1983 - SEE AUTHOR FILE UNDER SMITH AND MADSEN

Martin Simpson - July 25, 1983
Tandem Computers End-User Survey - SEE ARTICLE FILE

McClellan, Stephen T
Coming Computer Industry Shakeout (book) - 1984 - have pages on Tandem only - SEE ARTICLE FILE

McEvoy, Dennis - 1983 speech at NCC - SEE ARTICLE FILE

McSweeney, Tim - SEE Stratus - August 1985 - "Competitive Analysis" SEE CLIPPING FILE

"Measure of Transaction Processing Power", by Anon et al
- Datamation - April 1, 1985 - see Jim Gray in AUTHOR FILE
- SEE ALSO Tandem Technical Report 85.2

Menzel, Reinhard G
"Concurrent SP/k A Language Supporting Concurrent Processes", by Reinhard G Menzel - thesis, University of Toronto - July 1976 - SEE AUTHOR FILE

"Message-Transport Software Does the Whole Job", by Chris Duke and Sandy Metz - Electronic Design - July 8, 1982 - SEE AUTHOR FILE

Miller, Donald V
"Cryptanalysis of a Two Round Version of Des Using Index Implications
Cryptologia - October 1988
- SEE AUTHOR FILE

Mini-Micro Systems - Jan 1984
"Competition Heats Up.", by Susan Glazer - SEE ARTICLE FILE

[Also in this folder: Electronic Business - May 1, 1984
"A Faultless Market Draws Many New Contenders" - this is essentially the same article as above.]

M.I.T. - Sloan School of Management - February 1988
"The Realities of Electronic Data Interchange: How much Competitive Advantage?" by Robert I. Benjamin et al.

MITRE
Fault Tolerant Computing Exhibition, Technical Papers of (proceedings) - held April 1987 - sponsored by MITRE - "Evolution of Tandem's Approach to Fault Tolerance", by Brian Ball - paper dated March 1987 - SEE AUTHOR FILE

Morgan Stanley - Nov 28, 1983
"Stratus Fault-Tolerant Systems Facing a Bright Future"
- SEE CLIPPING FILES

MSRA Inc - April 1984
"Stratus and the Fault-Tolerant Computer Market" - SEE CLIPPING FILES

"Multiprocessor Organization-A Survey", by Philip H Enslow Jr
- Computing Surveys - March 1977 - SEE AUTHOR FILE

NCR Corporation Company Profile and Strategies - Competitive Analysis Report - Ruth Bates - Sept 25, 1987 - SEE ARTICLE FILE (under NCR)

"1987: The Beginning of a New Era in Data Processing"
"1987: The Beginning of a New Era in Data Processing", by E F Codd - Information Week - May 4, 1987 - on NonStop SQL - SEE CLIPPING FILES

Nikkei Computer - v6 n10 - 1985
"Tandem & Other Vendors" - Nikkei Computer - v6 n10 - 1985, by Eiiche Watanabe - in Japanese - SEE AUTHOR FILE

"NonStop Operating System", by Joel F Bartlett, copyright 1977
- SEE AUTHOR FILE

"NonStop Transaction Processing" by Lloyd Smith and Kent Madsen
- Datamation March 1983
- SEE AUTHOR FILE

"OLTP on the Vaxcluster" by Berl Hartman - Dec Professional Jan.88

"100 Best Companies to Work for In America" (book), by Robert Levering - Have section on Tandem only - SEE ARTICLE FILE

["On the Reliability of the IBM MVS/XA" - see entry under Stratus Updates]

"Online Transaction Processing - What is it? And should I care?" by Philip A. Naecker - Dec Professional Jan.88

"Optimal Control of Thrashing", by Russ Blake
- copyright 1982 by Tandem Computers - SEE AUTHOR FILE

[Otto, Richard - Stratus-System/88 Competitive Updates, Comparisons, etc - see entry under Stratus Updates]

"Overview of Fault-Tolerant Systems, An", By Patricia A Becker - Auerbach Publishers Inc - copyright 1986 - SEE AUTHOR FILE

"Overview of the Implementation of Fault-Tolerant Systems", by Eiichi Watanabe (in Japanese) - July 8, 1988 - SEE AUTHOR FILE

PaineWebber - May 24, 1984
"Fault-Tolerant Computers: A Progress Report" - SEE ARTICLE FILE

PaineWebber - Nov 17, 1985
Tandem Computer Status Report - SEE ARTICLE FILE

"Principles of Concurrent Programming" - book, by M Ben-Ari and Eiichi Watanabe - in Japanese - copyright 1982 - (note: Ben-Ari is not a Tandem employee) - SEE AUTHOR FILE

"Product Development Cycle" - video tape, by Larry Laurich - Nov 11, 1980 - 35 min - SEE AUTHOR FILE

Prudential-Bache Securities - June 18, 1985
"Tandem Computers Inc Software House Survey" - SEE ARTICLE FILE

Prudential-Bache Securities - January 20, 1986
"Company Repositing for Better Profitability and Growth Nearly Complete" - SEE ARTICLE FILE

"The Realities of Electronic Data Interchange: How Much Competitive Advantage?" by Robert I. Benjamin, David W. De Long, Michael S. Scott Morton. - February 1988
SEE ARTICLE FILE

"Reliability, Availability, and Fault Tolerance", by Paul Hamilton - Chilton's I&CS Oct. 1987 - SEE ARTICLE FILE

"Robustness to Crash in a Distributed Database: A Non Shared-Memory Multi-Processor Approach", by Andrea Borr - from the Proceedings of the Tenth International Conference on Very Large Data Bases, Aug 1984 - SEE AUTHOR FILE - SEE ALSO Tandem Technical Report 84.2

"Rollercoaster Revenue Climbs & Slides..."
- EDP Industry Report - July 1, 1985 - Review & Forecast Part I - "View from the Valley: Large-Scale Systems Coming off Slower Growth, Facing Lower Income Levels in 1985 Despite Continuing Demand"
- SEE ARTICLE FILE

July 31, 1985 - Review & Forecast Part II - "Rollercoaster Revenue Climbs and Slides Unsettle Stomachs for Medium-, Small- Scale System Vendors" - SEE ARTICLE FILE

"Running Late in Fail-Safe Computers", by Ira Sager - Venture - Dec 1983

Sager, Ira
"Running Late in Fail-Safe Computers", by Ira Sager - Venture - Dec 1983

Silicon Compilers Tame 10,000-Gate-Plus ASICs, Gate Arrays, by William J. Stenzel - EDN - April 28, 1988
- SEE AUTHOR FILE

Simulation, Measurement and Modeling of Computer Systems, 1979 Conference on - "Tailor: A Simple Model That Works", by Russ Blake
- SEE AUTHOR FILE

Smith, Lloyd
"Designer's Overview of Transaction Processing", by Lloyd Smith, copyright 1979 - SEE AUTHOR FILE

"NonStop Transaction Processing", by Lloyd Smith and Kent Madsen - Datamation March 1983 -SEE AUTHOR FILE UNDER SMITH AND MADSEN

Software House Survey -
Prudential-Bache Securities - June 18, 1985 -
"Tandem Computers Inc Software House Survey" - SEE ARTICLE FILE

"Solving the Problems of Distributed Databases", by Jim Gray and Sandy Metz - Data Communciations - Oct 1983 - SEE AUTHOR FILE

Standard & Poor's Industry Surveys - June 6, 1985
"Computers & Office Equipment Basic Analysis" - SEE ARTICLE FILE

Stenzel, William J, EDN - April 28, 1988
"Silicon Compilers Tame 10,000-Gate-Plus ASICs, Gate Arrays - SEE AUTHOR FILE

Stratus - Nov 23, 1983
"Fault-Tolerant Systems Facing a Bright Future" - Morgan Stanley - SEE CLIPPING FILES

Stratus - April 18, 1984
Preliminary Analysis of Stratus Product Announcements, by George Haskel - SEE CLIPPING FILES

Stratus - April 1984
"Stratus and the Fault-Tolerant Computer Market" - MSRA, Inc - SEE CLIPPING FILES

Stratus - August 1985
"Competitive Analysis", by Tim McSweeney - SEE CLIPPING FILES

Stratus - March 20, 1986
Tandem VS Stratus Comparison, by Rich Otto - SEE CLIPPING FILES

Stratus - Stratus-System/88 Competitive Updates
Update #1, Dec 6, 1985, by Richard Otto
Update #2, Mar 12, 1986, by Richard Otto
Note: we also have copies for distribution of Stanford paper "On the Reliability of the IBM MVS/XA", cited in this update.
Update #3, May 1, 1986, by Richard Otto
Update #4, May 13, 1986, by Elias Zelkha
Update #5, June 6, 1986, by Pete Vankuran
- SEE CLIPPING FILES FOR ALL OF THE ABOVE

"Structure of System/88, a Fault-Tolerant computer" , by E.S. Harrison IBM Systems Journal, V26, n3, 1987 - SEE ARTICLE FILE

Summers, Chris
"Computers in Banking and Finance The Impact of Fault Tolerant

Systems", by Chris Summers - Financial Times Business Information Ltd - received Jan 1984 - SEE ARTICLE FILE

"System Architecture for Nonstop Computing", by James A Katzman - Compcon Spring 1977 - SEE AUTHOR FILE

Systems & Software - March 1985
"Fault-Tolerant Computers Mature", by Michael Chester - Systems & Software - March 1985 - SEE ARTICLE FILE

Tangney, Cacey
Datamation-Cowen - mail memo from Cacey Tangney, "More on Customer Satisfaction" - Oct 3, 1986 - in response to "unfavorable" survey - SEE ARTICLE FILE under Datamation-Cowen

"Tailor: A Simple Model That Works", by Russ Blake - 1979 Conference on Simulation, Measurement and Modeling of Computer Systems - SEE AUTHOR FILE

"Tandem & Other Vendors" - Nikkei Computer - v6 n10 - 1985, by Eiiche Watanabe - in Japanese - SEE AUTHOR FILE

"Tandem Computers will have it next" - Yankee Conveyor - Aug. 1987 - SEE CLIPPING FILE

"Tandem 16 NonStop System, The", by Dave Mackie - no source, no date - SEE AUTHOR FILE

"Tandem's Treybig: No More Mister Niche Guy", by Paul E. Schindler - Information Week - November 9, 1987 - SEE CLIPPING FILES

Taylor, H Robert
"Guardian/Expand: A NonStop Network", by James A Katzman and H Robert Taylor - copyright 1978 - SEE AUTHOR FILE
(note: it is unclear whether or not Taylor was a Tandem employee)

Telecommunications - October 1985
"Beyond OSI: Open Systems Networking", by David Kirk - SEE AUTHOR FILE

"Time-Staged Delivery Networks Save Time, Enhance Productivity", by Jeri Edwards - Data Communications - Feb 1986 - SEE AUTHOR FILE

Tom, Gary Francis
"Checkpointing Techniques for Fault-Tolerant Process-Pairs", by Gary Francis Tom - Submitted to the Department of Electrical Engineering and Computer Science in Partial Fulfillment of the Requirements for the Degree of Bachelor of Science at the Massachusetts Institute of Technology, June 1981" - SEE AUTHOR FILE

"Transaction Model", by Jim Gray - paper - Feb 1980 - SEE AUTHOR FILE

"Transaction Monitoring in ENCOMPASS: Reliable Distributed Transaction Processing", by Andrea Borr - from the Proceedings of the

Seventh International Conference on Very Large Databases, (IEEE) Sept 1981 - SEE AUTHOR FILE - SEE ALSO Tandem Technical Report 81.2

"Transparency in its Place", by Jim Gray - Unix Review - May 1987 - SEE ARTICLE FILE; SEE ALSO AUTHOR FILE

Treacy, Michael
"The Costs of Network Ownership" - the Index Group,Inc.1989 - SEE ARTICLE FILE

Treybig, Jim - September 30, 1987 Cowen Conference

"Tuning Garwick's Algorithm for Repacking Sequential Storage", by David S Wise and Dan C Watson - Bit 1976 (v16) - SEE AUTHOR FILE (under Watson - it is unclear if Wise was a Tandem employee)

"Twelve Rules for a Distributed Data Base", by C J Date - ComputerWorld - June 8, 1987 - SEE ARTICLE FILE

Van Vleck, Tom
"The three Questions" - ACM Software Engineering Notes, July 1989

Venture - Dec 1983
"Running Late in Fail-Safe Computers", by Ira Sager - SEE ARTICLE FILE

Very Large Databases, Proceedings of the Seventh International Conference on, Sept 1981 - "Transaction Monitoring in ENCOMPASS: Reliable Distributed Transaction Processing", by Andrea Borr - SEE AUTHOR FILE - SEE ALSO Tandem Technical Report 81.2

Very Large Data Bases, Proceedings of the Tenth International Conference on, Aug 1984 - "Robustness to Crash in a Distributed Database: A Non Shared-Memory Multi-Processor Approach", by Andrea Borr - SEE AUTHOR FILE - SEE ALSO Tandem Technical Report 84.2

"View from the Valley: Large-Scale Systems Coming off Slower Growth, Facing Lower Income Levels in 1985 Despite Continuing Demand" EDP Industry Report - July 1, 1985 - Review & Forecast Part I - - SEE ARTICLE FILE

Watson, Dan C
"Tuning Garwick's Algorithm for Repacking Sequential Storage", by David S Wise and Dan C Watson - SEE AUTHOR FILE (under Watson - it is unclear if Wise was a Tandem employee)

Watanabe, Eiichi
"Fault-Tolerant System" - book, by Eiichi Watanabe - in Japanese. This is a colection of technical papers/articles published in the past, with addition information such as a draft proposal of Japanese Industry Standard of Highly Reliable Computing System. Papers were written by various Tandem and non-Tandem people. Book received Oct 1986. - SEE AUTHOR FILE

"Fault-tolerant System Architecture of Tandem NonStop System", presentation by Eiichi Watanabe - Jan 1984 - abstract in English, paper in Japanese - SEE AUTHOR FILE

Japan Electronic Industry Newsletter - monthly newsletter sent by email by Eiichi Watanabe beginning April 1987 - SEE ARTICLE FILE (under title)

"Overview of the Implementation of Fault-Tolerant Systems", by Eiichi Watanabe (in Japanese) - July 8, 1988 - SEE AUTHOR FILE

Japan Technology and Industry Report - quarterly report sent by email by Eiichi Watanabe beginning June 30, 1987 - SEE ARTICLE FILE (under title)

"Principles of Concurrent Programming" - book, by M Ben-Ari and Eiichi Watanabe - in Japanese - copyright 1982 - (note: Ben-Ari is not a Tandem employee) - SEE AUTHOR FILE

"Tandem & Other Vendors" - Nikkei Computer v6 n10 1985, by Eiiche Watanabe - in Japanese - SEE AUTHOR FILE

"What Makes Tandem Run?"
BusinessWeek - July 14, 1980 - SEE ARTICLE FILE

Wise, David S
"Tuning Garwick's Algorithm for Repacking Sequential Storage", by David S Wise and Dan C Watson - SEE AUTHOR FILE (under Watson - it is unclear if Wise was a Tandem employee)

"XRAY: Instrumentation for Multiple Computers", by Russ Blake - copyright 1980 by the Association for Computing Machinery - SEE AUTHOR FILE

Yankee Group Market Watch - June 5, 1986
"OLTP Market Update: NCR, Tandem's High-End Strategies" - SEE ARTICLE FILE

EXOTIC LANGUAGE OF THE MONTH CLUB

Transaction Application Language

By Serg Koren and Pierre Provencher

A relatively recent advance in computer systems has been the development of fault-tolerant computers. One such system is the TANDEM machine—a multiprocessor, multiuser transaction-oriented minicomputer, which the manufacturer claims experiences minimal downtime. To take advantage of the NON-STOP (TANDEM trademark) capabilities of the system, TANDEM developed and implemented a high-level, block-structured language known as TAL (Transaction Application Language).

TAL is used by TANDEM to develop all of its system software and by programmers to write applications where performance is a consideration or where system-level access is needed. TAL is proof that communication and system module applications do not need to be developed at the machine-code level. This article will touch briefly on TAL's main points.

TAL has many similarities to Pascal and ALGOL. It has procedures, subprocedures, and functions and is strongly-typed— although not as strongly as Pascal. TAL also has a standard I/O and function library as well as the ability to access the bit level.

TAL comments are delimited by !. Like Pascal, TAL permits local and global variables and has the same *BEGIN END* structure. A TAL identifier is up to 31 characters long and starts with a letter. The ^ is considered a letter permitting readable identifiers:

```
int line^buffer^of^terminal[0:39];
```

TAL's assignment operator is := and = represents a logical test. Assignment can be used after a variable declaration, making initialization easy. (Strangely, literals use = for assignment). Using apostrophes around the assignment statement transforms assignment into a move of multiple elements. Assignment can also be chained.

```
array1':='array2 for 4;
a:=b:=c:=d;
```

One of the unique features of the language allows the user to interface non-standard hardware or to temporarily override the original sysgen configuration of a piece of hardware.

TAL is compiled in one pass and generates true TANDEM machine code (a custom processor board, not a CPU). The language allows calls to COBOL, FORTRAN, and library routines. In addition, TAL routines are callable by both COBOL and FORTRAN.

Procedures (known as PROCs) in TAL have their own local data area and can be called recursively. Like most other procedure-oriented languages, TAL's PROCs can pass or receive parameters. Parameters can be passed as a reference or a value although arrays are always passed by reference. A procedure or a function can also be passed as a parameter. A typical PROC shell is presented in Listing 1, with a typical call being,

```
CALL Sampleproc(A,B);
```

In addition, PROCs can return a value which makes the PROC equivalent to functions in other languages. To do this the procedure is given a type, and a *RETURN* statement is included in the body of the procedure which will return a value of that type. The procedure in Listing 2 can be called with

```
Result := Typedproc;
```

Procedures may also be assigned attributes. For example, the attribute RESIDENT indicates that the portion of code for this procedure should not be swapped out of memory. The procedure where execution begins is known as the main PROC

```
PROC Sampleproc(Param1,Param2);

   INT Param1,
       Param2;
BEGIN

   INT Localvariable1,
       Localvariable2;

     <code here>

END;      ! end of proc
```

Listing 1.

```
INT PROC Typedproc;      !without parameters
BEGIN
     <code>
     RETURN Intvalue;    !where Intvalue is global here
     <more code>
     RETURN Intvalue;    !you can have more than 1 RETURN
END;
```

Listing 2.

Computer Language Jan 1985

and is designated by the attribute *MAIN* following the PROC name:

```
PROC Starthere MAIN;
```

In addition to one procedure calling another, TAL allows PROCs to contain and call SUBPROCs. Just as with PROCs, SUBPROCs can pass parameters and be typed. Similarly, SUBPROCs have their own local data space and attributes.

TAL has the standard loop constructs found in other languages: the *FOR*, *DO UNTIL*, and *WHILE DO*. The *FOR* loop can be ascending (*TO*) or descending (*DOWNTO*). A primitive *CASE* is also provided. The *CASE* statement works with an index variable that starts at zero. *OTHERWISE* can be used to assure that a value greater than the last case index specified will have predictable results.

```
CASE indexvariable of
    begin
    a:=a+1; !indexvariable = 0
    b:=b+1; !indexvariable = 1
    otherwise c:=c+1;
              !indexvariable > 1
    end;
```

Although it is confusing and not structured, TAL allows the use of *DEFINE*. The *DEFINE* assigns a portion of program source text to an identifier.

```
DEFINE total= count:=count+1 #;
    .....
    if new^customer then total;
```

One unique feature of the TAL *IF* statement is that it can be used in an assignment statement:

```
Result := IF Thisistrue THEN 3 ELSE 0;
```

which will assign a value of 3 if the variable *Thisistrue* is logically true. If *Thisistrue* is logically false, a value of 0 will be assigned. As long as an expression can be evaluated as logically true or false, it can be imbedded within an *IF*.

```
IF (X := Y - Z) THEN CALL Myproc;
```

will evaluate the expression within parentheses. If the resultant value of X is greater than zero then the expression is true and the *CALL* will be executed.

TAL handles six data types, single-word (16 bit), double-word, strings (8-bit bytes), fixed point (64 bit), and reals (64 or 32 bit). Individual data types can be converted to any other through built-in type transfer functions. Data declaration can be specified as direct or indirect addressing (Listing 3). A TANDEM word is 16 bits.

A dot before an identifier indicates standard indirect addressing. Indirect addressing was introduced because of a TANDEM restriction on memory addressing by a program. Without using extended addressing, 64K work is available for the data area of a TAL program. Only the first 32K is byte addressable. Standard indirect addressing provides access to the 64K while extended addressing (increasing the address to 32 bits from 16 bits) permits the user to define a larger memory segment.

For extended addressing it is the programmer's responsibility to manage memory through *ALLOCATESEGMENT* and *DEALLOCATESEGMENT* calls. The individual address of a variable can be accessed by prefixing the @ operator to the identifier.

Any data variable (except literals) can be specified as a one-dimensional array by specifying bounds. Multiple dimensional arrays are not permitted. Any variable can be indexed as an array. A pointer is an indirect addressing variable with no dimension. The compiler does not build run-time boundary checks even if bounds are specified.

```
int .pointer; !is a standard pointer
int .ext top; !is an extended
              !address pointer
```

Variables and arrays can be defined as read only and then are coded in-line instead of in the data area.

```
string
  .three^letters[0:2]='P':=["abc"];
```

TAL also allows the user to perform block moves of data through the use of the move statement. For example, consider the code segment presented in Listing 4.

This example is a bit strange but none the less sometimes useful. Again, you can only move data of the same type although, as in the last example, the system does not check variable bounds.

Numerous data conversion functions are available to transform one data type to another. The built-in function *CODE* enables entry of machine assembler in the program and can be useful to write such things as a trap handling routine.

After an I/O or a GUARDIAN call, the condition code of the operation can easily be checked using an *IF* statement with no test variable.

```
if < then call error;
!the condition code is less than 0
```

Bit manipulations can only be done on integers and are useful for bit maps, etc. The examples presented in Listing 5 show the major statement types. TANDEM numbers the bits from the high-order bit (bit 0 is the high-order bit). A test on bit

```
literal
    true = 1,                          !definition of
    false = 0;                         ! constants
int
    counter:=0,                        !direct integer
   .array^of^ten^elements[0:9];        !indirect array
int(32)
    dblword;                           !double integer
string
    .a^to^c:=["abc"];                  !string declaration
```

Listing 3.

```
INT A,
    C;
STRING B[0:3];  !a 4 byte string array

C := 1;
B ':=' "TEST";  !would "move" the 4 bytes into the array
A ':=' C FOR 2; !would "move" the 2 words starting at variable
                !C into the memory starting at variable A.
                !A would then have the value 1 and C would have
                !the value of "TE".
```

Listing 4.

15 is an easy way to determine if the value is even or odd.

```
word1.<8:9>:=word2.<10:11>;
if word.<15> then ...
```

TAL allows the software developer to use address pointers to either byte or word data as the following shows:

```
INT .Indirectvariable[0:5];
              !six word array
STRING .Stringpointer :=
    @Indirectvariable '<<' 1;
```

The @ indicates "address of" and '<<' 1 does a logical left bit shift of 1 bit. This effectively multiplies the integer address by 2 in order to get the address of the string. This is one way of dealing with integers as if they were strings. The statement

```
Stringpointer[1] := "A";
```

would set the rightmost byte of the first word of *Indirectvariable* to the letter A.

Since TAL is used to write all of TANDEM's compilers and interpreters, a string search facility is provided in the form of a *SCAN* statement. *SCAN* tests a string of any length until either a target character or an ASCII null is encountered. The address of the target character is returned as a result of the *SCAN*. Whereas *SCAN* searches strings from left to right, a similar command, *RSCAN*, does so from right to left. The form of the syntax for *SCAN* is:

```
SCAN <string array> WHILE (or
    UNTIL) <target character> ->
    <next address>
```

In all TAL has 19 different types of statements, none of these dealing with I/O. All I/O is handled through a standard library of PROCs. A typical I/O statement is *READ*, and appears as a *CALL* in the body of a user PROC. For example:

```
CALL READ(filenum,filebuf,
    readcount,countread);
```

is typical. Filenum is the file to be read, filebuf is an integer array into which data is to be read, readcount is the number of bytes to read, and countread is returned by the *READ PROC* and indicates the actual number of bytes read. All other I/O is handled similarly.

TAL also allows the user to code applications that communicate with each other by means of inter-process messages. Library routines are provided to handle these. *AWAITIO* is a PROC that either awaits an I/O completion, specifies a time-out period for an I/O, or checks for a completion of an I/O. This PROC is used to code applications where processing need not wait for I/Os, where I/Os must be funneled, where I/O timing is critical, etc.

Library routines are provided for creation, deletion, and system management

```
A.<1> := 1;  ! set bit one of variable A
A := A.<3:6>;! set the value of A to the value of bits 3 thru 6
             ! bit extraction
IF A.<4> THEN RETURN; ! bit test
```

Listing 5.

RP/M T.M.

By the author of Hayden's "CP/M Revealed."

New resident console processor RCP and new resident disk operating system RDOS replace CCP and BDOS without TPA size change.

User 0 files common to all users; user number visible in system prompt; file first extent size and user assignment displayed by DIR; cross-drive command file search; paged TYPE display with selectable page size. SUBMIT runs on any drive with multiple command files conditionally invoked by CALL. Automatic disk flaw processing isolates unuseable sectors. For high capacity disk systems RDOS can provide instantaneous directory access and delete redundant nondismountable disk logins. RPMPIP utility copies files, optionally prompts for confirmation during copy-all, compares files, archives large files to multiple floppy disks. RPMGEN and GETRPM self-install RP/M on any computer currently running CP/M®2.2. Source program assembly listings of RCP and RDOS appear in the RP/M user's manual.

RP/M manual with RPMGEN.COM and GETRPM.COM plus our RPMPIP.COM and other RP/M utilities on 8" SSSD $75. Shipping $5 ($10 nonUS). MC, VISA.

118 SW First St. - Box C
Warrenton, OR. 97146

Micro Methods, Inc.

(503) 861-1765

CIRCLE 52 ON READER SERVICE CARD

OPT-TECH SORT™

SORT/MERGE program for IBM-PC & XT

Now also sorts dBASE II files!

- Written in assembly language for **high performance**
  Example 4,000 records of 128 bytes sorted to give key & pointer file in 30 seconds **COMPARE!**
- Sort ascending or descending on up to nine fields
- Ten input files may be sorted or merged at one time
- Handles variable and fixed length records
- Supports all common data types
- Filesize limited only by your disk space
- Dynamically allocates memory and work files
- Output file can be full records, keys or pointers
- Can be run from keyboard or as a batch command
- Can be called as a subroutine to many languages
- Easy to use, includes on-line help feature
- Full documentation — sized like your PC manuals
- **$99** —VISA, M/C, Check, Money Order, COD, or PO
  Quantity discounts and OEM licensing available

To order or to receive additional information write or call:

OPT-TECH DATA PROCESSING
P.O. Box 2167 Humble, Texas 77347
(713) 454-7428
Requires DOS, 64K and One Disk Drive

CIRCLE 67 ON READER SERVICE CARD

of other processes, files, etc. If additional efficiency or capability is necessary, TAL allows for actual TANDEM assembly code to reside within TAL PROCs. Device control is another major TANDEM function—everything from laboratory equipment to automatic tellers in banking (a common application).

SETMODE is the TAL library PROC that allows the user to temporarily override the device characteristics originally specified during sysgen. This procedure allows the user to code an application such that any device used by the application becomes virtually port-independent.

For example, if a data communications line is normally sysgened as being 9600 baud with 7 data bits, an application could alter this to 300 baud with 8 data bits for special communications needs. *SETMODE* can handle printers, terminals, communications lines, etc. and lets the user alter a wide range of characteristics.

Tal becomes a really powerful tool when associated with the GUARDIAN operating system. File access, process control and checkpointing are not defined in the language itself, but the calls needed are sourced in the TAL program. This gives serious advantages: only the needed procedures are sourced in by the programmer, the guardian routines are not reserved TAL identifiers, and the program stays structured since those calls are treated as procedures.

As is the case with most languages, TAL is still evolving to accommodate new uses, hardware, etc. Although TAL is usable only on TANDEMs, it is a highly structured and powerful language with a wide following. Part of the strength of TAL lies in the simplicity of its statement types and part in the wide range of standard library procedures provided. ■

Serg Koren is a professional programmer with Smith-Kline Clinical Laboratories and has been using TAL for six years.

Pierre Provencher graduated with a degree in computer science from the Univ. of Montreal, Que. He has been consulting for on-line full tolerance systems for five years.

DSD 80

The Dynamic Screen Debugger for CP/M

DSD-80 is the most advanced debugging program available for CP/M-80 and compatible operating systems. Spend less time debugging and more time programming because DSD is the only utility to offer all of the most important features in one package.

Satisfaction is guaranteed or a full refund!

Full featured simulator provides for
- Execution only within boundaries
- Write protected memory
- Stack overflow protection
- Stack underflow protection
- Real time subroutines
- Breakpoint on specified memory or register values

Full screen display includes
- 2 memory displays
- Register display
- Stack display
- Instruction display

Fully upward compatible with DDT

Single keystroke commands for
- Stepping instructions
- Subroutine execution
- Scrolling memory displays

Full symbol support including
- Loading symbol files
- Interactive symbol definitions

Endorsed by Leor Zolman — author of BDS C

Uses less than 15k of memory

On line help information

Clearly written 50 page manual

Free updates for one year

Only $195.00

SOFTADVANCES

P.O. Box 49473 Austin, TX 78765 (512) 478-4763

Available on 8" SSSD & various 5.25" formats
IBM PC version available soon!
Visa & MC accepted

CIRCLE 77 ON READER SERVICE CARD

UNPARALLELED **PERFORMANCE** and **PORTABILITY** in an **ISAM** PACKAGE at an **UNBEATABLE** PRICE

The company that introduced micros to B-Trees in 1979 and created ACCESS MANAGER™ for Digital Research, now redefines the market for high performance, B-Tree based file handlers. With c-tree™ you get:

- complete C source code written to K & R standards of portability
- high level, multi-key ISAM routines and low level B-Tree functions
- routines that work with single-user and network systems
- no royalties on application programs

$395 COMPLETE

PC-DOS

c-tree™ BY FAIRCOM

2606 Johnson Drive
Columbia MO 65203

CIRCLE 29 ON READER SERVICE CARD

EXOTIC LANGUAGE OF THE MONTH CLUB

Transaction Application Language

By Serg Koren and Pierre Provencher

A relatively recent advance in computer systems has been the development of fault-tolerant computers. One such system is the TANDEM machine—a multiprocessor, multiuser transaction-oriented minicomputer, which the manufacturer claims experiences minimal downtime. To take advantage of the NON-STOP (TANDEM trademark) capabilities of the system, TANDEM developed and implemented a high-level, block-structured language known as TAL (Transaction Application Language).

TAL is used by TANDEM to develop all of its system software and by programmers to write applications where performance is a consideration or where system-level access is needed. TAL is proof that communication and system module applications do not need to be developed at the machine-code level. This article will touch briefly on TAL's main points.

TAL has many similarities to Pascal and ALGOL. It has procedures, subprocedures, and functions and is strongly-typed— although not as strongly as Pascal. TAL also has a standard I/O and function library as well as the ability to access the bit level.

TAL comments are delimited by !. Like Pascal, TAL permits local and global variables and has the same *BEGIN END* structure. A TAL identifier is up to 31 characters long and starts with a letter. The ^ is considered a letter permitting readable identifiers:

```
int line^buffer^of^terminal[0:39];
```

TAL's assignment operator is := and = represents a logical test. Assignment can be used after a variable declaration, making initialization easy. (Strangely, literals use = for assignment). Using apostrophes around the assignment statement transforms assignment into a move of multiple elements. Assignment can also be chained.

```
array1':=' array2 for 4;
a:=b:=c:=d;
```

One of the unique features of the language allows the user to interface nonstandard hardware or to temporarily override the original sysgen configuration of a piece of hardware.

TAL is compiled in one pass and generates true TANDEM machine code (a custom processor board, not a CPU). The language allows calls to COBOL, FORTRAN, and library routines. In addition, TAL routines are callable by both COBOL and FORTRAN.

Procedures (known as PROCs) in TAL have their own local data area and can be called recursively. Like most other procedure-oriented languages, TAL's PROCs can pass or receive parameters. Parameters can be passed as a reference or a value although arrays are always passed by reference. A procedure or a function can also be passed as a parameter. A typical PROC shell is presented in Listing 1, with a typical call being,

```
CALL Sampleproc(A,B);
```

In addition, PROCs can return a value which makes the PROC equivalent to functions in other languages. To do this the procedure is given a type, and a *RETURN* statement is included in the body of the procedure which will return a value of that type. The procedure in Listing 2 can be called with

```
Result := Typedproc;
```

Procedures may also be assigned attributes. For example, the attribute RESIDENT indicates that the portion of code for this procedure should not be swapped out of memory. The procedure where execution begins is known as the main PROC

```
PROC Sampleproc(Param1,Param2);

   INT Param1,
       Param2;
BEGIN

   INT Localvariable1,
       Localvariable2;

     <code here>

END;      ! end of proc
```

Listing 1.

```
INT PROC Typedproc;       !without parameters
BEGIN
     <code>
     RETURN Intvalue;     !where Intvalue is global here
     <more code>
     RETURN Intvalue;     !you can have more than 1 RETURN
END;
```

Listing 2.

Computer Language Jan 1985

and is designated by the attribute *MAIN* following the PROC name:

```
PROC Starthere MAIN;
```

In addition to one procedure calling another, TAL allows PROCs to contain and call SUBPROCs. Just as with PROCs, SUBPROCs can pass parameters and be typed. Similarly, SUBPROCs have their own local data space and attributes.

TAL has the standard loop constructs found in other languages: the *FOR*, *DO UNTIL*, and *WHILE DO*. The *FOR* loop can be ascending (*TO*) or descending (*DOWNTO*). A primitive *CASE* is also provided. The *CASE* statement works with an index variable that starts at zero. *OTHERWISE* can be used to assure that a value greater than the last case index specified will have predictable results.

```
CASE indexvariable of
    begin
    a:=a+1; !indexvariable = 0
    b:=b+1; !indexvariable = 1
    otherwise c:=c+1;
              !indexvariable > 1
    end;
```

Although it is confusing and not structured, TAL allows the use of *DEFINE*. The *DEFINE* assigns a portion of program source text to an identifier.

```
DEFINE total= count:=count+1 #;
     ......
     if new^customer then total;
```

One unique feature of the TAL *IF* statement is that it can be used in an assignment statement:

```
Result := IF Thisistrue THEN 3 ELSE 0;
```

which will assign a value of 3 if the variable *Thisistrue* is logically true. If *Thisistrue* is logically false, a value of 0 will be assigned. As long as an expression can be evaluated as logically true or false, it can be imbedded within an *IF*.

```
IF (X := Y - Z) THEN CALL Myproc;
```

will evaluate the expression within parentheses. If the resultant value of X is greater than zero then the expression is true and the *CALL* will be executed.

TAL handles six data types, single-word (16 bit), double-word, strings (8-bit bytes), fixed point (64 bit), and reals (64 or 32 bit). Individual data types can be converted to any other through built-in type transfer functions. Data declaration can be specified as direct or indirect addressing (Listing 3). A TANDEM word is 16 bits.

A dot before an identifier indicates standard indirect addressing. Indirect addressing was introduced because of a TANDEM restriction on memory addressing by a program. Without using extended addressing, 64K work is available for the data area of a TAL program. Only the first 32K is byte addressable. Standard indirect addressing provides access to the 64K while extended addressing (increasing the address to 32 bits from 16 bits) permits the user to define a larger memory segment.

For extended addressing it is the programmer's responsibility to manage memory through *ALLOCATESEGMENT* and *DEALLOCATESEGMENT* calls. The individual address of a variable can be accessed by prefixing the @ operator to the identifier.

Any data variable (except literals) can be specified as a one-dimensional array by specifying bounds. Multiple dimensional arrays are not permitted. Any variable can be indexed as an array. A pointer is an indirect addressing variable with no dimension. The compiler does not build run-time boundary checks even if bounds are specified.

```
int .pointer; !is a standard pointer
int .ext top; !is an extended
               !address pointer
```

Variables and arrays can be defined as read only and then are coded in-line instead of in the data area.

```
string
  .three^letters[0:2] = 'P' := ["abc"];
```

TAL also allows the user to perform block moves of data through the use of the move statement. For example, consider the code segment presented in Listing 4.

This example is a bit strange but none the less sometimes useful. Again, you can only move data of the same type although, as in the last example, the system does not check variable bounds.

Numerous data conversion functions are available to transform one data type to another. The built-in function *CODE* enables entry of machine assembler in the program and can be useful to write such things as a trap handling routine.

After an I/O or a GUARDIAN call, the condition code of the operation can easily be checked using an *IF* statement with no test variable.

```
if < then call error;
!the condition code is less than 0
```

Bit manipulations can only be done on integers and are useful for bit maps, etc. The examples presented in Listing 5 show the major statement types. TANDEM numbers the bits from the high-order bit (bit 0 is the high-order bit). A test on bit

```
literal
    true = 1,                          !definition of
    false = 0;                         ! constants
int
    counter:=0,                        !direct integer
   .array^of^ten^elements[0:9];        !indirect array
int(32)
    dblword;                           !double integer
string
    .a^to^c:=["abc"];                  !string declaration
```

Listing 3.

```
INT A,
    C;
STRING B[0:3];  !a 4 byte string array

C := 1;
B ':=' "TEST";  !would "move" the 4 bytes into the array
A ':=' C FOR 2; !would "move" the 2 words starting at variable
                !C into the memory starting at variable A.
                !A would then have the value 1 and C would have
                !the value of "TE".
```

Listing 4.

15 is an easy way to determine if the value is even or odd.

```
word1.<8:9>:=word2.<10:11>;
if word.<15> then ...
```

TAL allows the software developer to use address pointers to either byte or word data as the following shows:

```
INT .Indirectvariable[0:5];
            !six word array
STRING .Stringpointer :=
    @Indirectvariable '<<' 1;
```

The @ indicates "address of" and '<<' 1 does a logical left bit shift of 1 bit. This effectively multiplies the integer address by 2 in order to get the address of the string. This is one way of dealing with integers as if they were strings. The statement

```
Stringpointer[1] := "A";
```

would set the rightmost byte of the first word of *Indirectvariable* to the letter A.

Since TAL is used to write all of TANDEM's compilers and interpreters, a string search facility is provided in the form of a *SCAN* statement. *SCAN* tests a string of any length until either a target character or an ASCII null is encountered. The address of the target character is returned as a result of the *SCAN*. Whereas *SCAN* searches strings from left to right, a similar command, *RSCAN*, does so from right to left. The form of the syntax for *SCAN* is:

```
SCAN <string array> WHILE (or
    UNTIL) <target character> ->
    <next address>
```

In all TAL has 19 different types of statements, none of these dealing with I/O. All I/O is handled through a standard library of PROCs. A typical I/O statement is *READ*, and appears as a *CALL* in the body of a user PROC. For example:

```
CALL READ(filenum,filebuf,
    readcount,countread);
```

is typical. Filenum is the file to be read, filebuf is an integer array into which data is to be read, readcount is the number of bytes to read, and countread is returned by the *READ PROC* and indicates the actual number of bytes read. All other I/O is handled similarly.

TAL also allows the user to code applications that communicate with each other by means of inter-process messages. Library routines are provided to handle these. *AWAITIO* is a PROC that either awaits an I/O completion, specifies a time-out period for an I/O, or checks for a completion of an I/O. This PROC is used to code applications where processing need not wait for I/Os, where I/Os must be funneled, where I/O timing is critical, etc.

Library routines are provided for creation, deletion, and system management

```
A.<1> := 1;  ! set bit one of variable A
A := A.<3:6>;! set the value of A to the value of bits 3 thru 6
             ! bit extraction
IF A.<4> THEN RETURN; ! bit test
```

Listing 5.

RP/M T.M.

By the author of Hayden's "CP/M Revealed."

New resident console processor RCP and new resident disk operating system RDOS replace CCP and BDOS without TPA size change.

User 0 files common to all users; user number visible in system prompt; file first extent size and user assignment displayed by DIR; cross-drive command file search; paged TYPE display with selectable page size. SUBMIT runs on any drive with multiple command files conditionally invoked by CALL. Automatic disk flaw processing isolates unuseable sectors. For high capacity disk systems RDOS can provide instantaneous directory access and delete redundant nondismountable disk logins. RPMPIP utility copies files, optionally prompts for confirmation during copy-all, compares files, archives large files to multiple floppy disks. RPMGEN and GETRPM self-install RP/M on any computer currently running CP/M®2.2. Source program assembly listings of RCP and RDOS appear in the RP/M user's manual.

RP/M manual with RPMGEN.COM and GETRPM.COM plus our RPMPIP.COM and other RP/M utilities on 8" SSSD $75. Shipping $5 ($10 nonUS). MC, VISA.

118 SW First St. - Box G
Warrenton, OR. 97146

Micro Methods, Inc.

(503) 861-1765

CIRCLE 52 ON READER SERVICE CARD

OPT-TECH SORT™

SORT/MERGE program for IBM-PC & XT

Now also sorts dBASE II files!

- Written in assembly language for **high performance**
  Example: 4,000 records of 128 bytes sorted to give key & pointer file in 30 seconds **COMPARE!**
- Sort ascending or descending on up to nine fields
- Ten input files may be sorted or merged at one time
- Handles variable and fixed length records
- Supports all common data types
- Filesize limited only by your disk space
- Dynamically allocates memory and work files
- Output file can be full records, keys or pointers
- Can be run from keyboard or as a batch command
- Can be called as a subroutine to many languages
- Easy to use, includes on-line help feature
- Full documentation — sized like your PC manuals
- $99 —VISA, M/C, Check, Money Order, COD, or PO
  Quantity discounts and OEM licensing available

To order or to receive additional information write or call:

OPT-TECH DATA PROCESSING

P.O. Box 2167 Humble, Texas 77347

(713) 454-7428

Requires DOS, 64K and One Disk Drive

CIRCLE 67 ON READER SERVICE CARD

of other processes, files, etc. If additional efficiency or capability is necessary, TAL allows for actual TANDEM assembly code to reside within TAL PROCs. Device control is another major TANDEM function—everything from laboratory equipment to automatic tellers in banking (a common application).

SETMODE is the TAL library PROC that allows the user to temporarily override the device characteristics originally specified during sysgen. This procedure allows the user to code an application such that any device used by the application becomes virtually port-independent.

For example, if a data communications line is normally sysgened as being 9600 baud with 7 data bits, an application could alter this to 300 baud with 8 data bits for special communications needs. *SETMODE* can handle printers, terminals, communications lines, etc. and lets the user alter a wide range of characteristics.

Tal becomes a really powerful tool when associated with the GUARDIAN operating system.

File access, process control and checkpointing are not defined in the language itself, but the calls needed are sourced in the TAL program. This gives serious advantages: only the needed procedures are sourced in by the programmer, the guardian routines are not reserved TAL identifiers, and the program stays structured since those calls are treated as procedures.

As is the case with most languages, TAL is still evolving to accommodate new uses, hardware, etc. Although TAL is usable only on TANDEMs, it is a highly structured and powerful language with a wide following. Part of the strength of TAL lies in the simplicity of its statement types and part in the wide range of standard library procedures provided. ■

Serg Koren is a professional programmer with Smith-Kline Clinical Laboratories and has been using TAL for six years.

Pierre Provencher graduated with a degree in computer science from the Univ. of Montreal, Que. He has been consulting for on-line full tolerance systems for five years.

DSD 80

The Dynamic Screen Debugger for CP/M

DSD-80 is the most advanced debugging program available for CP/M-80 and compatible operating systems. Spend less time debugging and more time programming because DSD is the only utility to offer all of the most important features in one package.

Satisfaction is guaranteed or a full refund!

Full featured simulator provides for:
Execution only within boundaries
Write protected memory
Stack overflow protection
Stack underflow protection
Real time subroutines
Breakpoint on specified memory or register values
Full screen display includes:
2 memory displays
Register display
Stack display
Instruction display
Fully upward compatible with DDT

Single keystroke commands for:
Stepping instructions
Subroutine execution
Scrolling memory displays
Full symbol support including:
Loading symbol files
Interactive symbol definitions
Endorsed by Leor Zolman — author of BDS C
Uses less than 15k of memory
On line help information
Clearly written 50 page manual
Free updates for one year

Only $195.00

SOFTADVANCES
P.O. Box 49473 Austin, TX 78765 (512) 478-4763

Available on 8" SSSD & various 5.25" formats
IBM PC version available soon!
Visa & MC accepted

CIRCLE 77 ON READER SERVICE CARD

UNPARALLELED
PERFORMANCE
and
PORTABILITY
in an
ISAM PACKAGE
at an
UNBEATABLE
PRICE

c-tree
BY FARCOM
2606 Johnson Drive
Columbia MO 65203

The company that introduced micros to B-Trees in 1979 and created ACCESS MANAGER™ for Digital Research, now redefines the market for high performance, B-Tree based file handlers. With c-tree™ you get:

- complete C source code written to K & R standards of portability
- high level, multi-key ISAM routines and low level B-Tree functions
- routines that work with single-user and network systems
- no royalties on application programs

$395 COMPLETE

[illegible]

Access Manager and CP/M are trademarks of Digital Research, Inc.
[illegible]

CIRCLE 29 ON READER SERVICE CARD

EXOTIC LANGUAGE OF THE MONTH CLUB

Transaction Application Language

By Serg Koren and Pierre Provencher

A relatively recent advance in computer systems has been the development of fault-tolerant computers. One such system is the TANDEM machine—a multiprocessor, multiuser transaction-oriented minicomputer, which the manufacturer claims experiences minimal downtime. To take advantage of the NON-STOP (TANDEM trademark) capabilities of the system, TANDEM developed and implemented a high-level, block-structured language known as TAL (Transaction Application Language).

TAL is used by TANDEM to develop all of its system software and by programmers to write applications where performance is a consideration or where system-level access is needed. TAL is proof that communication and system module applications do not need to be developed at the machine-code level. This article will touch briefly on TAL's main points.

TAL has many similarities to Pascal and ALGOL. It has procedures, subprocedures, and functions and is strongly-typed— although not as strongly as Pascal. TAL also has a standard I/O and function library as well as the ability to access the bit level.

TAL comments are delimited by !. Like Pascal, TAL permits local and global variables and has the same *BEGIN END* structure. A TAL identifier is up to 31 characters long and starts with a letter. The ^ is considered a letter permitting readable identifiers:

```
int line^buffer^of^terminal[0:39];
```

TAL's assignment operator is := and = represents a logical test. Assignment can be used after a variable declaration, making initialization easy. (Strangely, literals use = for assignment). Using apostrophes around the assignment statement transforms assignment into a move of multiple elements. Assignment can also be chained.

```
array1':=' array2 for 4;
a:=b:=c:=d;
```

One of the unique features of the language allows the user to interface nonstandard hardware or to temporarily override the original sysgen configuration of a piece of hardware.

TAL is compiled in one pass and generates true TANDEM machine code (a custom processor board, not a CPU). The language allows calls to COBOL, FORTRAN, and library routines. In addition, TAL routines are callable by both COBOL and FORTRAN.

Procedures (known as PROCs) in TAL have their own local data area and can be called recursively. Like most other procedure-oriented languages, TAL's PROCs can pass or receive parameters. Parameters can be passed as a reference or a value although arrays are always passed by reference. A procedure or a function can also be passed as a parameter. A typical PROC shell is presented in Listing 1, with a typical call being,

```
CALL Sampleproc(A,B);
```

In addition, PROCs can return a value which makes the PROC equivalent to functions in other languages. To do this the procedure is given a type, and a *RETURN* statement is included in the body of the procedure which will return a value of that type. The procedure in Listing 2 can be called with

```
Result := Typedproc;
```

Procedures may also be assigned attributes. For example, the attribute RESIDENT indicates that the portion of code for this procedure should not be swapped out of memory. The procedure where execution begins is known as the main PROC

```
PROC Sampleproc(Param1,Param2);

   INT Param1,
       Param2;
BEGIN

   INT Localvariable1,
       Localvariable2;

     <code here>

END;       ! end of proc
```

Listing 1.

```
INT PROC Typedproc;        !without parameters
BEGIN
    <code>
    RETURN Intvalue;   !where Intvalue is global here
    <more code>
    RETURN Intvalue;   !you can have more than 1 RETURN
END;
```

Listing 2.

Computer Language Jan 1985

and is designated by the attribute *MAIN* following the PROC name:

```
PROC Starthere MAIN;
```

In addition to one procedure calling another, TAL allows PROCs to contain and call SUBPROCs. Just as with PROCs, SUBPROCs can pass parameters and be typed. Similarly, SUBPROCs have their own local data space and attributes.

TAL has the standard loop constructs found in other languages: the *FOR*, *DO UNTIL*, and *WHILE DO*. The *FOR* loop can be ascending (*TO*) or descending (*DOWNTO*). A primitive *CASE* is also provided. The *CASE* statement works with an index variable that starts at zero. *OTHERWISE* can be used to assure that a value greater than the last case index specified will have predictable results.

```
CASE indexvariable of
    begin
    a:=a+1; !indexvariable = 0
    b:=b+1; !indexvariable = 1
    otherwise c:=c+1;
              !indexvariable > 1
    end;
```

Although it is confusing and not structured, TAL allows the use of *DEFINE*. The *DEFINE* assigns a portion of program source text to an identifier.

```
DEFINE total= count:=count+1 #;
    .....
    if new^customer then total;
```

One unique feature of the TAL *IF* statement is that it can be used in an assignment statement:

```
Result:= IF Thisistrue THEN 3 ELSE 0;
```

which will assign a value of 3 if the variable *Thisistrue* is logically true. If *Thisistrue* is logically false, a value of 0 will be assigned. As long as an expression can be evaluated as logically true or false, it can be imbedded within an *IF*.

```
IF (X:= Y - Z) THEN CALL Myproc;
```

will evaluate the expression within parentheses. If the resultant value of X is greater than zero then the expression is true and the *CALL* will be executed.

TAL handles six data types, single-word (16 bit), double-word, strings (8-bit bytes), fixed point (64 bit), and reals (64 or 32 bit). Individual data types can be converted to any other through built-in type transfer functions. Data declaration can be specified as direct or indirect addressing (Listing 3). A TANDEM word is 16 bits.

A dot before an identifier indicates standard indirect addressing. Indirect addressing was introduced because of a TANDEM restriction on memory addressing by a program. Without using extended addressing, 64K work is available for the data area of a TAL program. Only the first 32K is byte addressable. Standard indirect addressing provides access to the 64K while extended addressing (increasing the address to 32 bits from 16 bits) permits the user to define a larger memory segment.

For extended addressing it is the programmer's responsibility to manage memory through *ALLOCATESEGMENT* and *DEALLOCATESEGMENT* calls. The individual address of a variable can be accessed by prefixing the @ operator to the identifier.

Any data variable (except literals) can be specified as a one-dimensional array by specifying bounds. Multiple dimensional arrays are not permitted. Any variable can be indexed as an array. A pointer is an indirect addressing variable with no dimension. The compiler does not build run-time boundary checks even if bounds are specified.

```
int .pointer; !is a standard pointer
int .ext top; !is an extended
              !address pointer
```

Variables and arrays can be defined as read only and then are coded in-line instead of in the data area.

```
string
  .three^letters[0:2]='P':=["abc"];
```

TAL also allows the user to perform block moves of data through the use of the move statement. For example, consider the code segment presented in Listing 4.

This example is a bit strange but none the less sometimes useful. Again, you can only move data of the same type although, as in the last example, the system does not check variable bounds.

Numerous data conversion functions are available to transform one data type to another. The built-in function *CODE* enables entry of machine assembler in the program and can be useful to write such things as a trap handling routine.

After an I/O or a GUARDIAN call, the condition code of the operation can easily be checked using an *IF* statement with no test variable.

```
if < then call error;
!the condition code is less than 0
```

Bit manipulations can only be done on integers and are useful for bit maps, etc. The examples presented in Listing 5 show the major statement types. TANDEM numbers the bits from the high-order bit (bit 0 is the high-order bit). A test on bit

```
literal
    true = 1,                          !definition of
    false = 0;                         ! constants
int
    counter:=0,                        !direct integer
    .array^of^ten^elements[0:9];       !indirect array
int(32)
    dblword;                           !double integer
string
    .a^to^c:=["abc"];                  !string declaration
```

Listing 3.

```
INT A,
    C;
STRING B[0:3]; !a 4 byte string array

C := 1;
B ':=' "TEST";  !would "move" the 4 bytes into the array
A ':=' C FOR 2; !would "move" the 2 words starting at variable
                !C into the memory starting at variable A.
                !A would then have the value 1 and C would have
                !the value of "TE".
```

Listing 4.

15 is an easy way to determine if the value is even or odd.

```
word1.<8:9> := word2.<10:11>;
if word.<15> then ...
```

TAL allows the software developer to use address pointers to either byte or word data as the following shows:

```
INT .Indirectvariable[0:5];
            !six word array
STRING .Stringpointer :=
    @Indirectvariable '<<' 1;
```

The @ indicates "address of" and '<<' 1 does a logical left bit shift of 1 bit. This effectively multiplies the integer address by 2 in order to get the address of the string. This is one way of dealing with integers as if they were strings. The statement

```
Stringpointer[1] := "A";
```

would set the rightmost byte of the first word of *Indirectvariable* to the letter A.

Since TAL is used to write all of TANDEM's compilers and interpreters, a string search facility is provided in the form of a *SCAN* statement. *SCAN* tests a string of any length until either a target character or an ASCII null is encountered. The address of the target character is returned as a result of the *SCAN*. Whereas *SCAN* searches strings from left to right, a similar command, *RSCAN*, does so from right to left. The form of the syntax for *SCAN* is:

```
SCAN <string array> WHILE (or
    UNTIL) <target character> ->
    <next address>
```

In all TAL has 19 different types of statements, none of these dealing with I/O. All I/O is handled through a standard library of PROCs. A typical I/O statement is *READ*, and appears as a *CALL* in the body of a user PROC. For example:

```
CALL READ(filenum,filebuf,
    readcount,countread);
```

is typical. Filenum is the file to be read, filebuf is an integer array into which data is to be read, readcount is the number of bytes to read, and countread is returned by the *READ PROC* and indicates the actual number of bytes read. All other I/O is handled similarly.

TAL also allows the user to code applications that communicate with each other by means of inter-process messages. Library routines are provided to handle these. *AWAITIO* is a PROC that either awaits an I/O completion, specifies a time-out period for an I/O, or checks for a completion of an I/O. This PROC is used to code applications where processing need not wait for I/Os, where I/Os must be funneled, where I/O timing is critical, etc.

Library routines are provided for creation, deletion, and system management

```
A.<1> := 1;  ! set bit one of variable A
A := A.<3:6>;! set the value of A to the value of bits 3 thru 6
             ! bit extraction
IF A.<4> THEN RETURN; ! bit test
```

Listing 5.

RP/M T.M.

By the author of Hayden's "CP/M Revealed."

New resident console processor RCP and new resident disk operating system RDOS replace CCP and BDOS without TPA size change.

User 0 files common to all users; user number visible in system prompt; file first extent size and user assignment displayed by DIR; cross-drive command file search; paged TYPE display with selectable page size. SUBMIT runs on any drive with multiple command files conditionally invoked by CALL. Automatic disk flaw processing isolates unuseable sectors. For high capacity disk systems RDOS can provide instantaneous directory access and delete redundant nondismountable disk logins. RPMPIP utility copies files, optionally prompts for confirmation during copy-all, compares files, archives large files to multiple floppy disks. RPMGEN and GETRPM self-install RP/M on any computer currently running CP/M®2.2. Source program assembly listings of RCP and RDOS appear in the RP/M user's manual.

RP/M manual with RPMGEN.COM and GETRPM.COM plus our RPMPIP.COM and other RP/M utilities on 8" SSSD $75. Shipping $5 ($10 nonUS). MC, VISA.

118 SW First St. - Box G
Warrenton, OR. 97146

Micro Methods, Inc.

(503) 861-1765

CIRCLE 52 ON READER SERVICE CARD

OPT-TECH SORT™

SORT/MERGE program for IBM-PC & XT

Now also sorts dBASE II files!

- Written in assembly language for **high performance**
  Example: 4,000 records of 128 bytes sorted to give key & pointer file in 30 seconds. **COMPARE!**
- Sort ascending or descending on up to nine fields
- Ten input files may be sorted or merged at one time
- Handles variable and fixed length records
- Supports all common data types
- Filesize limited only by your disk space
- Dynamically allocates memory and work files
- Output file can be full records, keys or pointers
- Can be run from keyboard or as a batch command
- Can be called as a subroutine to many languages
- Easy to use, includes on-line help feature
- Full documentation — sized like your PC manuals
- $99 — VISA, M/C, Check, Money Order, COD, or PO
  Quantity discounts and OEM licensing available

To order or to receive additional information write or call:

OPT-TECH DATA PROCESSING
P.O. Box 2167 Humble, Texas 77347
(713) 454-7428

Requires DOS, 64K and One Disk Drive

CIRCLE 67 ON READER SERVICE CARD

other processes, files, etc. If additional efficiency or capability is necessary, TAL allows for actual TANDEM assembly code to reside within TAL PROCs. Device control is another major TANDEM function—everything from laboratory equipment to automatic tellers in banking (a common application).

SETMODE is the TAL library PROC that allows the user to temporarily override the device characteristics originally specified during sysgen. This procedure allows the user to code an application such that any device used by the application becomes virtually port-independent.

For example, if a data communications line is normally sysgened as being 9600 baud with 7 data bits, an application could alter this to 300 baud with 8 data bits for special communications needs. *SETMODE* can handle printers, terminals, communications lines, etc. and lets the user alter a wide range of characteristics.

Tal becomes a really powerful tool when associated with the GUARDIAN operating system. File access, process control and checkpointing are not defined in the language itself, but the calls needed are sourced in the TAL program. This gives serious advantages: only the needed procedures are sourced in by the programmer, the guardian routines are not reserved TAL identifiers, and the program stays structured since those calls are treated as procedures.

As is the case with most languages, TAL is still evolving to accommodate new uses, hardware, etc. Although TAL is usable only on TANDEMs, it is a highly structured and powerful language with a wide following. Part of the strength of TAL lies in the simplicity of its statement types and part in the wide range of standard library procedures provided. ■

Serg Koren is a professional programmer with Smith-Kline Clinical Laboratories and has been using TAL for six years.

Pierre Provencher graduated with a degree in computer science from the Univ. of Montreal, Que. He has been consulting for on-line full tolerance systems for five years.

DSD 80

The Dynamic Screen Debugger for CP/M

DSD-80 is the most advanced debugging program available for CP/M-80 and compatible operating systems. Spend less time debugging and more time programming because DSD is the only utility to offer all of the most important features in one package.

Satisfaction is guaranteed or a full refund!

Full featured simulator provides for:
Execution only within boundaries
Write protected memory
Stack overflow protection
Stack underflow protection
Real time subroutines
Breakpoint on specified memory or register values
Full screen display includes:
2 memory displays
Register display
Stack display
Instruction display
Fully upward compatible with DDT

Single keystroke commands for:
Stepping instructions
Subroutine execution
Scrolling memory displays
Full symbol support including:
Loading symbol files
Interactive symbol definitions
Endorsed by Leor Zolman — author of BDS C
Uses less than 15k of memory
On line help information
Clearly written 50 page manual
Free updates for one year

Only $195.00

SOFTADVANCES
P.O. Box 49473 Austin, TX 78765 (512) 478-4763

Available on 8" SSSD & various 5.25" formats
IBM PC version available soon!
Visa & MC accepted

CIRCLE 77 ON READER SERVICE CARD

UNPARALLELED PERFORMANCE and PORTABILITY in an ISAM PACKAGE at an UNBEATABLE PRICE

The company that introduced micros to B-Trees in 1979 and created ACCESS MANAGER™ for Digital Research, now redefines the market for high performance, B-Tree based file handlers. With c-tree™ you get:

- complete C source code written to K & R standards of portability
- high level, multi-key ISAM routines and low level B-Tree functions
- routines that work with single-user and network systems
- no royalties on application programs

$395 COMPLETE

c-tree™ BY FAIRCOM
2606 Johnson Drive
Columbia MO 65203

CIRCLE 29 ON READER SERVICE CARD

EXOTIC LANGUAGE OF THE MONTH CLUB

Transaction Application Language

By Serg Koren and Pierre Provencher

A relatively recent advance in computer systems has been the development of fault-tolerant computers. One such system is the TANDEM machine—a multiprocessor, multi-user transaction-oriented minicomputer, which the manufacturer claims experiences minimal downtime. To take advantage of the NON-STOP (TANDEM trademark) capabilities of the system, TANDEM developed and implemented a high-level, block-structured language known as TAL (Transaction Application Language).

TAL is used by TANDEM to develop all of its system software and by programmers to write applications where performance is a consideration or where system-level access is needed. TAL is proof that communication and system module applications do not need to be developed at the machine-code level. This article will touch briefly on TAL's main points.

TAL has many similarities to Pascal and ALGOL. It has procedures, subprocedures, and functions and is strongly-typed— although not as strongly as Pascal. TAL also has a standard I/O and function library as well as the ability to access the bit level.

TAL comments are delimited by !. Like Pascal, TAL permits local and global variables and has the same *BEGIN END* structure. A TAL identifier is up to 31 characters long and starts with a letter. The ^ is considered a letter permitting readable identifiers:

```
int line^buffer^of^terminal[0:39];
```

TAL's assignment operator is := and = represents a logical test. Assignment can be used after a variable declaration, making initialization easy. (Strangely, literals use = for assignment). Using apostrophes around the assignment statement transforms assignment into a move of multiple elements. Assignment can also be chained.

```
array1':=' array2 for 4;
a:=b:=c:=d;
```

One of the unique features of the language allows the user to interface non-standard hardware or to temporarily override the original sysgen configuration of a piece of hardware.

TAL is compiled in one pass and generates true TANDEM machine code (a custom processor board, not a CPU). The language allows calls to COBOL, FORTRAN, and library routines. In addition, TAL routines are callable by both COBOL and FORTRAN.

Procedures (known as PROCs) in TAL have their own local data area and can be called recursively. Like most other procedure-oriented languages, TAL's PROCs can pass or receive parameters. Parameters can be passed as a reference or a value although arrays are always passed by reference. A procedure or a function can also be passed as a parameter. A typical PROC shell is presented in Listing 1, with a typical call being,

```
CALL Sampleproc(A,B);
```

In addition, PROCs can return a value which makes the PROC equivalent to functions in other languages. To do this the procedure is given a type, and a *RETURN* statement is included in the body of the procedure which will return a value of that type. The procedure in Listing 2 can be called with

```
Result := Typedproc;
```

Procedures may also be assigned attributes. For example, the attribute RESIDENT indicates that the portion of code for this procedure should not be swapped out of memory. The procedure where execution begins is known as the main PROC

```
PROC Sampleproc(Param1,Param2);

   INT Param1,
       Param2;
BEGIN

   INT Localvariable1,
       Localvariable2;

     <code here>

END;      ! end of proc
```

Listing 1.

```
INT PROC Typedproc;        !without parameters
BEGIN
     <code>
     RETURN Intvalue;      !where Intvalue is global here
     <more code>
     RETURN Intvalue;      !you can have more than 1 RETURN
END;
```

Listing 2.

and is designated by the attribute *MAIN* following the PROC name:

```
PROC Storthere MAIN;
```

In addition to one procedure calling another, TAL allows PROCs to contain and call SUBPROCs. Just as with PROCs, SUBPROCs can pass parameters and be typed. Similarly, SUBPROCs have their own local data space and attributes.

TAL has the standard loop constructs found in other languages: the *FOR*, *DO UNTIL*, and *WHILE DO*. The *FOR* loop can be ascending (*TO*) or descending (*DOWNTO*). A primitive *CASE* is also provided. The *CASE* statement works with an index variable that starts at zero. *OTHERWISE* can be used to assure that a value greater than the last case index specified will have predictable results.

```
CASE indexvariable of
    begin
    a:=a+1; !indexvariable = 0
    b:=b+1; !indexvariable = 1
    otherwise c:=c+1;
              !indexvariable > 1
    end;
```

Although it is confusing and not structured, TAL allows the use of *DEFINE*. The *DEFINE* assigns a portion of program source text to an identifier.

```
DEFINE total= count:=count+1 #;
    .....
    if new^customer then total;
```

One unique feature of the TAL *IF* statement is that it can be used in an assignment statement:

```
Result := IF Thisistrue THEN 3 ELSE 0;
```

which will assign a value of 3 if the variable *Thisistrue* is logically true. If *Thisistrue* is logically false, a value of 0 will be assigned. As long as an expression can be evaluated as logically true or false, it can be imbedded within an *IF*.

```
IF (X := Y - Z) THEN CALL Myproc;
```

will evaluate the expression within parentheses. If the resultant value of X is greater than zero then the expression is true and the *CALL* will be executed.

TAL handles six data types, single-word (16 bit), double-word, strings (8-bit bytes), fixed point (64 bit), and reals (64 or 32 bit). Individual data types can be converted to any other through built-in type transfer functions. Data declaration can be specified as direct or indirect addressing (Listing 3). A TANDEM word is 16 bits.

A dot before an identifier indicates standard indirect addressing. Indirect addressing was introduced because of a TANDEM restriction on memory addressing by a program. Without using extended addressing, 64K work is available for the data area of a TAL program. Only the first 32K is byte addressable. Standard indirect addressing provides access to the 64K while extended addressing (increasing the address to 32 bits from 16 bits) permits the user to define a larger memory segment.

For extended addressing it is the programmer's responsibility to manage memory through *ALLOCATESEGMENT* and *DEALLOCATESEGMENT* calls. The individual address of a variable can be accessed by prefixing the @ operator to the identifier.

Any data variable (except literals) can be specified as a one-dimensional array by specifying bounds. Multiple dimensional arrays are not permitted. Any variable can be indexed as an array. A pointer is an indirect addressing variable with no dimension. The compiler does not build run-time boundary checks even if bounds are specified.

```
int .pointer; !is a standard pointer
int .ext top; !is an extended
              !address pointer
```

Variables and arrays can be defined as read only and then are coded in-line instead of in the data area.

```
string
  .three^letters[0:2]='P':=["abc"];
```

TAL also allows the user to perform block moves of data through the use of the move statement. For example, consider the code segment presented in Listing 4.

This example is a bit strange but none the less sometimes useful. Again, you can only move data of the same type although, as in the last example, the system does not check variable bounds.

Numerous data conversion functions are available to transform one data type to another. The built-in function *CODE* enables entry of machine assembler in the program and can be useful to write such things as a trap handling routine.

After an I/O or a GUARDIAN call, the condition code of the operation can easily be checked using an *IF* statement with no test variable.

```
if < then call error;
!the condition code is less than 0
```

Bit manipulations can only be done on integers and are useful for bit maps, etc. The examples presented in Listing 5 show the major statement types. TANDEM numbers the bits from the high-order bit (bit 0 is the high-order bit). A test on bit

```
literal
    true = 1,                     !definition of
    false = 0;                    ! constants
int
    counter:=0,                   !direct integer
   .array^of^ten^elements[0:9];   !indirect array
int(32)
    dblword;                      !double integer
string
   .a^to^c:=["abc"];              !string declaration
```

Listing 3.

```
INT A,
    C;
STRING B[0:3];  !a 4 byte string array

C := 1;
B ':=' "TEST";  !would "move" the 4 bytes into the array
A ':=' C FOR 2; !would "move" the 2 words starting at variable
                !C into the memory starting at variable A.
                !A would then have the value 1 and C would have
                !the value of "TE".
```

Listing 4.

15 is an easy way to determine if the value is even or odd.

```
word1.<8:9>:=word2.<10:11>;
if word.<15> then ...
```

TAL allows the software developer to use address pointers to either byte or word data as the following shows:

```
INT .Indirectvariable[0:5];
            !six word array
STRING .Stringpointer :=
    @Indirectvariable '<<' 1;
```

The @ indicates "address of" and '<<' 1 does a logical left bit shift of 1 bit. This effectively multiplies the integer address by 2 in order to get the address of the string. This is one way of dealing with integers as if they were strings. The statement

```
Stringpointer[1] := "A";
```

would set the rightmost byte of the first word of *Indirectvariable* to the letter A.

Since TAL is used to write all of TANDEM's compilers and interpreters, a string search facility is provided in the form of a *SCAN* statement. *SCAN* tests a string of any length until either a target character or an ASCII null is encountered. The address of the target character is returned as a result of the *SCAN*. Whereas *SCAN* searches strings from left to right, a similar command, *RSCAN*, does so from right to left. The form of the syntax for *SCAN* is:

```
SCAN <string array> WHILE (or
    UNTIL) <target character> ->
    <next address>
```

In all TAL has 19 different types of statements, none of these dealing with I/O. All I/O is handled through a standard library of PROCs. A typical I/O statement is *READ*, and appears as a *CALL* in the body of a user PROC. For example:

```
CALL READ(filenum,filebuf,
    readcount,countread);
```

is typical. Filenum is the file to be read, filebuf is an integer array into which data is to be read, readcount is the number of bytes to read, and countread is returned by the *READ PROC* and indicates the actual number of bytes read. All other I/O is handled similarly.

TAL also allows the user to code applications that communicate with each other by means of inter-process messages. Library routines are provided to handle these. *AWAITIO* is a PROC that either awaits an I/O completion, specifies a time-out period for an I/O, or checks for a completion of an I/O. This PROC is used to code applications where processing need not wait for I/Os, where I/Os must be funneled, where I/O timing is critical, etc.

Library routines are provided for creation, deletion, and system management

```
A.<1> := 1;  ! set bit one of variable A
A := A.<3:6>;! set the value of A to the value of bits 3 thru 6
             ! bit extraction
IF A.<4> THEN RETURN; ! bit test
```

Listing 5.

RP/M T.M.

By the author of Hayden's "CP/M Revealed."

New resident console processor RCP and new resident disk operating system RDOS replace CCP and BDOS without TPA size change.

User 0 files common to all users; user number visible in system prompt; file first extent size and user assignment displayed by DIR; cross-drive command file search; paged TYPE display with selectable page size. SUBMIT runs on any drive with multiple command files conditionally invoked by CALL. Automatic disk flaw processing isolates unuseable sectors. For high capacity disk systems RDOS can provide instantaneous directory access and delete redundant nondismountable disk logins. RPMPIP utility copies files, optionally prompts for confirmation during copy-all, compares files, archives large files to multiple floppy disks. RPMGEN and GETRPM self-install RP/M on any computer currently running CP/M®2.2. Source program assembly listings of RCP and RDOS appear in the RP/M user's manual.

RP/M manual with RPMGEN.COM and GETRPM.COM plus our RPMPIP.COM and other RP/M utilities on 8" SSSD $75. Shipping $5 ($10 nonUS). MC,VISA.

118 SW First St. - Box G
Warrenton, OR. 97146

Micro Methods, Inc.

(503) 861-1765

CIRCLE 52 ON READER SERVICE CARD

OPT-TECH SORT™

SORT/MERGE program for IBM-PC & XT

Now also sorts dBASE II files!

- Written in assembly language for **high performance**
  Example: 4,000 records of 128 bytes sorted to give key & pointer file in 30 seconds. **COMPARE!**
- Sort ascending or descending on up to nine fields
- Ten input files may be sorted or merged at one time
- Handles variable and fixed length records
- Supports all common data types
- Filesize limited only by your disk space
- Dynamically allocates memory and work files
- Output file can be full records, keys or pointers
- Can be run from keyboard or as a batch command
- Can be called as a subroutine to many languages
- Easy to use, includes on-line help feature
- Full documentation — sized like your PC manuals
- $99 — VISA, M/C, Check, Money Order, COD, or PO
  Quantity discounts and OEM licensing available

To order or to receive additional information write or call:

OPT-TECH DATA PROCESSING
P.O. Box 2167 Humble, Texas 77347
(713) 454-7428
Requires DOS, 64K and One Disk Drive

CIRCLE 67 ON READER SERVICE CARD

of other processes, files, etc. If additional efficiency or capability is necessary, TAL allows for actual TANDEM assembly code to reside within TAL PROCs.

Device control is another major TANDEM function—everything from laboratory equipment to automatic tellers in banking (a common application).

SETMODE is the TAL library PROC that allows the user to temporarily override the device characteristics originally specified during sysgen. This procedure allows the user to code an application such that any device used by the application becomes virtually port-independent.

For example, if a data communications line is normally sysgened as being 9600 baud with 7 data bits, an application could alter this to 300 baud with 8 data bits for special communications needs. *SETMODE* can handle printers, terminals, communications lines, etc. and lets the user alter a wide range of characteristics.

Tal becomes a really powerful tool when associated with the GUARDIAN operating system. File access, process control and checkpointing are not defined in the language itself, but the calls needed are sourced in the TAL program. This gives serious advantages: only the needed procedures are sourced in by the programmer, the guardian routines are not reserved TAL identifiers, and the program stays structured since those calls are treated as procedures.

As is the case with most languages, TAL is still evolving to accommodate new uses, hardware, etc. Although TAL is usable only on TANDEMs, it is a highly structured and powerful language with a wide following. Part of the strength of TAL lies in the simplicity of its statement types and part in the wide range of standard library procedures provided. ■

Serg Koren is a professional programmer with Smith-Kline Clinical Laboratories and has been using TAL for six years.

Pierre Provencher graduated with a degree in computer science from the Univ. of Montreal, Que. He has been consulting for on-line full tolerance systems for five years.

DSD 80

The Dynamic Screen Debugger for CP/M

DSD-80 is the most advanced debugging program available for CP/M-80 and compatible operating systems. Spend less time debugging and more time programming because DSD is the only utility to offer all of the most important features in one package.

Satisfaction is guaranteed or a full refund!

Full featured simulator provides for
- Execution only within boundaries
- Write protected memory
- Stack overflow protection
- Stack underflow protection
- Real time subroutines
- Breakpoint on specified memory or register values

Full screen display includes
- 2 memory displays
- Register display
- Stack display
- Instruction display

Fully upward compatible with DDT

Single keystroke commands for
- Stepping instructions
- Subroutine execution
- Scrolling memory displays

Full symbol support including
- Loading symbol files
- Interactive symbol definitions

Endorsed by Leor Zolman — author of BDS C

Uses less than 15k of memory

On line help information

Clearly written 50 page manual

Free updates for one year

Only $195.00

SOFTADVANCES

P.O. Box 49473 Austin, TX 78765 (512) 478-4763

Available on 8" SSSD & various 5.25" formats

IBM PC version available soon!

Visa & MC accepted

CIRCLE 77 ON READER SERVICE CARD

UNPARALLELED PERFORMANCE and PORTABILITY in an ISAM PACKAGE at an UNBEATABLE PRICE

The company that introduced micros to B-Trees in 1979 and created ACCESS MANAGER™ for Digital Research, now redefines the market for high performance, B-Tree based file handlers. With c-tree™ you get:

- complete C source code written to K & R standards of portability
- high level, multi-key ISAM routines and low level B-Tree functions
- routines that work with single-user and network systems
- no royalties on application programs

$395 COMPLETE

c-tree BY FAIRCOM

2606 Johnson Drive
Columbia MO 65203

CIRCLE 29 ON READER SERVICE CARD

EXOTIC LANGUAGE OF THE MONTH CLUB

Transaction Application Language

By Serg Koren and Pierre Provencher

A relatively recent advance in computer systems has been the development of fault-tolerant computers. One such system is the TANDEM machine—a multiprocessor, multiuser transaction-oriented minicomputer, which the manufacturer claims experiences minimal downtime. To take advantage of the NON-STOP (TANDEM trademark) capabilities of the system, TANDEM developed and implemented a high-level, block-structured language known as TAL (Transaction Application Language).

TAL is used by TANDEM to develop all of its system software and by programmers to write applications where performance is a consideration or where system-level access is needed. TAL is proof that communication and system module applications do not need to be developed at the machine-code level. This article will touch briefly on TAL's main points.

TAL has many similarities to Pascal and ALGOL. It has procedures, subprocedures, and functions and is strongly-typed— although not as strongly as Pascal. TAL also has a standard I/O and function library as well as the ability to access the bit level.

TAL comments are delimited by !. Like Pascal, TAL permits local and global variables and has the same *BEGIN END* structure. A TAL identifier is up to 31 characters long and starts with a letter. The ^ is considered a letter permitting readable identifiers:

```
int line^buffer^of^terminal[0:39];
```

TAL's assignment operator is := and = represents a logical test. Assignment can be used after a variable declaration, making initialization easy. (Strangely, literals use = for assignment). Using apostrophes around the assignment statement transforms assignment into a move of multiple elements. Assignment can also be chained.

```
array1':=' array2 for 4;
a:=b:=c:=d;
```

One of the unique features of the language allows the user to interface non-standard hardware or to temporarily override the original sysgen configuration of a piece of hardware.

TAL is compiled in one pass and generates true TANDEM machine code (a custom processor board, not a CPU). The language allows calls to COBOL, FORTRAN, and library routines. In addition, TAL routines are callable by both COBOL and FORTRAN.

Procedures (known as PROCs) in TAL have their own local data area and can be called recursively. Like most other procedure-oriented languages, TAL's PROCs can pass or receive parameters. Parameters can be passed as a reference or a value although arrays are always passed by reference. A procedure or a function can also be passed as a parameter. A typical PROC shell is presented in Listing 1, with a typical call being,

```
CALL Sampleproc(A,B);
```

In addition, PROCs can return a value which makes the PROC equivalent to functions in other languages. To do this the procedure is given a type, and a *RETURN* statement is included in the body of the procedure which will return a value of that type. The procedure in Listing 2 can be called with

```
Result := Typedproc;
```

Procedures may also be assigned attributes. For example, the attribute RESIDENT indicates that the portion of code for this procedure should not be swapped out of memory. The procedure where execution begins is known as the main PROC

```
PROC Sampleproc(Param1,Param2);

   INT Param1,
       Param2;
BEGIN

   INT Localvariable1,
       Localvariable2;

     <code here>

END;     ! end of proc
```

Listing 1.

```
INT PROC Typedproc;        !without parameters
BEGIN
     <code>
     RETURN Intvalue;      !where Intvalue is global here
     <more code>
     RETURN Intvalue;      !you can have more than 1 RETURN
END;
```

Listing 2.

Computer Language Jan 1985

$49.95
ata.
S™
Y!
. $175
Mark
. $150
. $160
NC.
6-8641

and is designated by the attribute *MAIN* following the PROC name:

```
PROC Starthere MAIN;
```

In addition to one procedure calling another, TAL allows PROCs to contain and call SUBPROCs. Just as with PROCs, SUBPROCs can pass parameters and be typed. Similarly, SUBPROCs have their own local data space and attributes.

TAL has the standard loop constructs found in other languages: the *FOR*, *DO UNTIL*, and *WHILE DO*. The *FOR* loop can be ascending (*TO*) or descending (*DOWNTO*). A primitive *CASE* is also provided. The *CASE* statement works with an index variable that starts at zero. *OTHERWISE* can be used to assure that a value greater than the last case index specified will have predictable results.

```
CASE indexvariable of
    begin
    a:=a+1; !indexvariable = 0
    b:=b+1; !indexvariable = 1
    otherwise c:=c+1;
              !indexvariable > 1
    end;
```

Although it is confusing and not structured, TAL allows the use of *DEFINE*. The *DEFINE* assigns a portion of program source text to an identifier.

```
DEFINE total= count:=count+1 #;
    .....
    if new^customer then total;
```

One unique feature of the TAL *IF* statement is that it can be used in an assignment statement:

```
Result := IF Thisistrue THEN 3 ELSE 0;
```

which will assign a value of 3 if the variable *Thisistrue* is logically true. If *Thisistrue* is logically false, a value of 0 will be assigned. As long as an expression can be evaluated as logically true or false, it can be imbedded within an *IF*.

```
IF (X:= Y - Z) THEN CALL Myproc;
```

will evaluate the expression within parentheses. If the resultant value of X is greater than zero then the expression is true and the *CALL* will be executed.

TAL handles six data types, single-word (16 bit), double-word, strings (8-bit bytes), fixed point (64 bit), and reals (64 or 32 bit). Individual data types can be converted to any other through built-in type transfer functions. Data declaration can be specified as direct or indirect addressing (Listing 3). A TANDEM word is 16 bits.

A dot before an identifier indicates standard indirect addressing. Indirect addressing was introduced because of a TANDEM restriction on memory addressing by a program. Without using extended addressing, 64K work is available for the data area of a TAL program. Only the first 32K is byte addressable. Standard indirect addressing provides access to the 64K while extended addressing (increasing the address to 32 bits from 16 bits) permits the user to define a larger memory segment.

For extended addressing it is the programmer's responsibility to manage memory through *ALLOCATESEGMENT* and *DEALLOCATESEGMENT* calls. The individual address of a variable can be accessed by prefixing the @ operator to the identifier.

Any data variable (except literals) can be specified as a one-dimensional array by specifying bounds. Multiple dimensional arrays are not permitted. Any variable can be indexed as an array. A pointer is an indirect addressing variable with no dimension. The compiler does not build run-time boundary checks even if bounds are specified.

```
int .pointer; !is a standard pointer
int .ext top; !is an extended
              !address pointer
```

Variables and arrays can be defined as read only and then are coded in-line instead of in the data area.

```
string
  .three^letters[0:2]='P':=["abc"];
```

TAL also allows the user to perform block moves of data through the use of the move statement. For example, consider the code segment presented in Listing 4.

This example is a bit strange but none the less sometimes useful. Again, you can only move data of the same type although, as in the last example, the system does not check variable bounds.

Numerous data conversion functions are available to transform one data type to another. The built-in function *CODE* enables entry of machine assembler in the program and can be useful to write such things as a trap handling routine.

After an I/O or a GUARDIAN call, the condition code of the operation can easily be checked using an *IF* statement with no test variable.

```
if < then call error;
!the condition code is less than 0
```

Bit manipulations can only be done on integers and are useful for bit maps, etc. The examples presented in Listing 5 show the major statement types. TANDEM numbers the bits from the high-order bit (bit 0 is the high-order bit). A test on bit

```
literal
    true = 1,                        !definition of
    false = 0;                       ! constants
int
    counter:=0,                      !direct integer
   .array^of^ten^elements[0:9];      !indirect array
int(32)
    dblword;                         !double integer
string
    .a^to^c:=["abc"];                !string declaration
```

Listing 3.

```
INT A,
    C;
STRING B[0:3];  !a 4 byte string array

C := 1;
B ':=' "TEST";  !would "move" the 4 bytes into the array
A ':=' C FOR 2; !would "move" the 2 words starting at variable
                !C into the memory starting at variable A.
                !A would then have the value 1 and C would have
                !the value of "TE".
```

Listing 4.

15 is an easy way to determine if the value is even or odd.

```
word1.<8:9>:=word2.<10:11>;
if word.<15> then ...
```

TAL allows the software developer to use address pointers to either byte or word data as the following shows:

```
INT .Indirectvariable[0:5];
            !six word array
STRING .Stringpointer :=
    @Indirectvariable '<<' 1;
```

The @ indicates "address of" and '<<' 1 does a logical left bit shift of 1 bit. This effectively multiplies the integer address by 2 in order to get the address of the string. This is one way of dealing with integers as if they were strings. The statement

```
Stringpointer[1] := "A";
```

would set the rightmost byte of the first word of *Indirectvariable* to the letter A.

Since TAL is used to write all of TANDEM's compilers and interpreters, a string search facility is provided in the form of a *SCAN* statement. *SCAN* tests a string of any length until either a target character or an ASCII null is encountered. The address of the target character is returned as a result of the *SCAN*. Whereas *SCAN* searches strings from left to right, a similar command, *RSCAN*, does so from right to left. The form of the syntax for *SCAN* is:

```
SCAN <string array> WHILE (or
    UNTIL) <target character> ->
    <next address>
```

In all TAL has 19 different types of statements, none of these dealing with I/O. All I/O is handled through a standard library of PROCs. A typical I/O statement is *READ*, and appears as a *CALL* in the body of a user PROC. For example:

```
CALL READ(filenum,filebuf,
    readcount,countread);
```

is typical. Filenum is the file to be read, filebuf is an integer array into which data is to be read, readcount is the number of bytes to read, and countread is returned by the *READ PROC* and indicates the actual number of bytes read. All other I/O is handled similarly.

TAL also allows the user to code applications that communicate with each other by means of inter-process messages. Library routines are provided to handle these. *AWAITIO* is a PROC that either awaits an I/O completion, specifies a time-out period for an I/O, or checks for a completion of an I/O. This PROC is used to code applications where processing need not wait for I/Os, where I/Os must be funneled, where I/O timing is critical, etc.

Library routines are provided for creation, deletion, and system management

```
A.<1> := 1;  ! set bit one of variable A
A := A.<3:6>;! set the value of A to the value of bits 3 thru 6
             ! bit extraction
IF A.<4> THEN RETURN; ! bit test
```

Listing 5.

RP/M T.M.

By the author of Hayden's "CP/M Revealed."

New resident console processor RCP and new resident disk operating system RDOS replace CCP and BDOS without TPA size change.

User 0 files common to all users; user number visible in system prompt; file first extent size and user assignment displayed by DIR; cross-drive command file search; paged TYPE display with selectable page size. SUBMIT runs on any drive with multiple command files conditionally invoked by CALL. Automatic disk flaw processing isolates unuseable sectors. For high capacity disk systems RDOS can provide instantaneous directory access and delete redundant nondismountable disk logins. RPMPIP utility copies files, optionally prompts for confirmation during copy-all, compares files, archives large files to multiple floppy disks. RPMGEN and GETRPM self-install RP/M on any computer currently running CP/M® 2.2. Source program assembly listings of RCP and RDOS appear in the RP/M user's manual.

RP/M manual with RPMGEN.COM and GETRPM.COM plus our RPMPIP.COM and other RP/M utilities on 8" SSSD $75. Shipping $5 ($10 nonUS). MC, VISA.

118 SW First St. - Box G
Warrenton, OR. 97146

Micro Methods, Inc.

(503) 861-1765

CIRCLE 52 ON READER SERVICE CARD

OPT-TECH SORT™

SORT/MERGE program for IBM-PC & XT

Now also sorts dBASE II files!

- Written in assembly language for **high performance**
  Example: 4,000 records of 128 bytes sorted to give key & pointer file in 30 seconds. **COMPARE!**
- Sort ascending or descending on up to nine fields
- Ten input files may be sorted or merged at one time
- Handles variable and fixed length records
- Supports all common data types
- Filesize limited only by your disk space
- Dynamically allocates memory and work files
- Output file can be full records, keys or pointers
- Can be run from keyboard or as a batch command
- Can be called as a subroutine to many languages
- Easy to use, includes on-line help feature
- Full documentation — sized like your PC manuals
- **$99** —VISA, M/C, Check, Money Order, COD, or PO
  Quantity discounts and OEM licensing available

To order or to receive additional information write or call:

OPT-TECH DATA PROCESSING
P.O. Box 2167 Humble, Texas 77347
(713) 454-7428
Requires DOS, 64K and One Disk Drive

CIRCLE 67 ON READER SERVICE CARD

of other processes, files, etc. If additional efficiency or capability is necessary, TAL allows for actual TANDEM assembly code to reside within TAL PROCs. Device control is another major TANDEM function—everything from laboratory equipment to automatic tellers in banking (a common application).

SETMODE is the TAL library PROC that allows the user to temporarily override the device characteristics originally specified during sysgen. This procedure allows the user to code an application such that any device used by the application becomes virtually port-independent.

For example, if a data communications line is normally sysgened as being 9600 baud with 7 data bits, an application could alter this to 300 baud with 8 data bits for special communications needs. *SETMODE* can handle printers, terminals, communications lines, etc. and lets the user alter a wide range of characteristics.

Tal becomes a really powerful tool when associated with the GUARDIAN operating system. File access, process control and checkpointing are not defined in the language itself, but the calls needed are sourced in the TAL program. This gives serious advantages: only the needed procedures are sourced in by the programmer, the guardian routines are not reserved TAL identifiers, and the program stays structured since those calls are treated as procedures.

As is the case with most languages, TAL is still evolving to accommodate new uses, hardware, etc. Although TAL is usable only on TANDEMs, it is a highly structured and powerful language with a wide following. Part of the strength of TAL lies in the simplicity of its statement types and part in the wide range of standard library procedures provided. ■

Serg Koren is a professional programmer with Smith-Kline Clinical Laboratories and has been using TAL for six years.

Pierre Provencher graduated with a degree in computer science from the Univ. of Montreal, Que. He has been consulting for online full tolerance systems for five years.

DSD 80

The Dynamic Screen Debugger for CP/M

DSD-80 is the most advanced debugging program available for CP/M-80 and compatible operating systems. Spend less time debugging and more time programming because DSD is the only utility to offer all of the most important features in one package.

Satisfaction is guaranteed or a full refund!

Full featured simulator provides for
- Execution only within boundaries
- Write protected memory
- Stack overflow protection
- Stack underflow protection
- Real time subroutines
- Breakpoint on specified memory or register values

Full screen display includes
- 2 memory displays
- Register display
- Stack display
- Instruction display

Fully upward compatible with DDT

Single keystroke commands for
- Stepping instructions
- Subroutine execution
- Scrolling memory displays

Full symbol support including
- Loading symbol files
- Interactive symbol definitions

Endorsed by Leor Zolman — author of BDS C

Uses less than 15k of memory

On line help information

Clearly written 50 page manual

Free updates for one year

Only $195.00

SOFTADVANCES

P.O. Box 49473 Austin, TX 78765 (512) 478-4763

Available on 8" SSSD & various 5.25" formats
IBM PC version available soon!
Visa & MC accepted

CIRCLE 77 ON READER SERVICE CARD

UNPARALLELED **PERFORMANCE** and **PORTABILITY** in an **ISAM** PACKAGE at an **UNBEATABLE** PRICE

The company that introduced micros to B-Trees in 1979 and created ACCESS MANAGER™ for Digital Research, now redefines the market for high performance, B-Tree based file handlers. With c-tree™ you get:

- complete C source code written to K & R standards of portability
- high level, multi-key ISAM routines and low level B-Tree functions
- routines that work with single-user and network systems
- no royalties on application programs

$395 COMPLETE

Specify format:
8" CP/M® 5¼" PC-DOS 8" RT-11

for VISA, MC or COD orders, call
1-314-445-6833

c-tree BY FAIRCOM

2606 Johnson Drive
Columbia MO 65203

Access Manager and CP/M are trademarks of Digital Research, Inc.
c-tree and the circular disc logo are trademarks of FairCom
©1984 FairCom

CIRCLE 29 ON READER SERVICE CARD

E4-V2

Summer 1983

COMPUTER TECHNOLOGY REVIEW

The Systems Integration Sourcebook

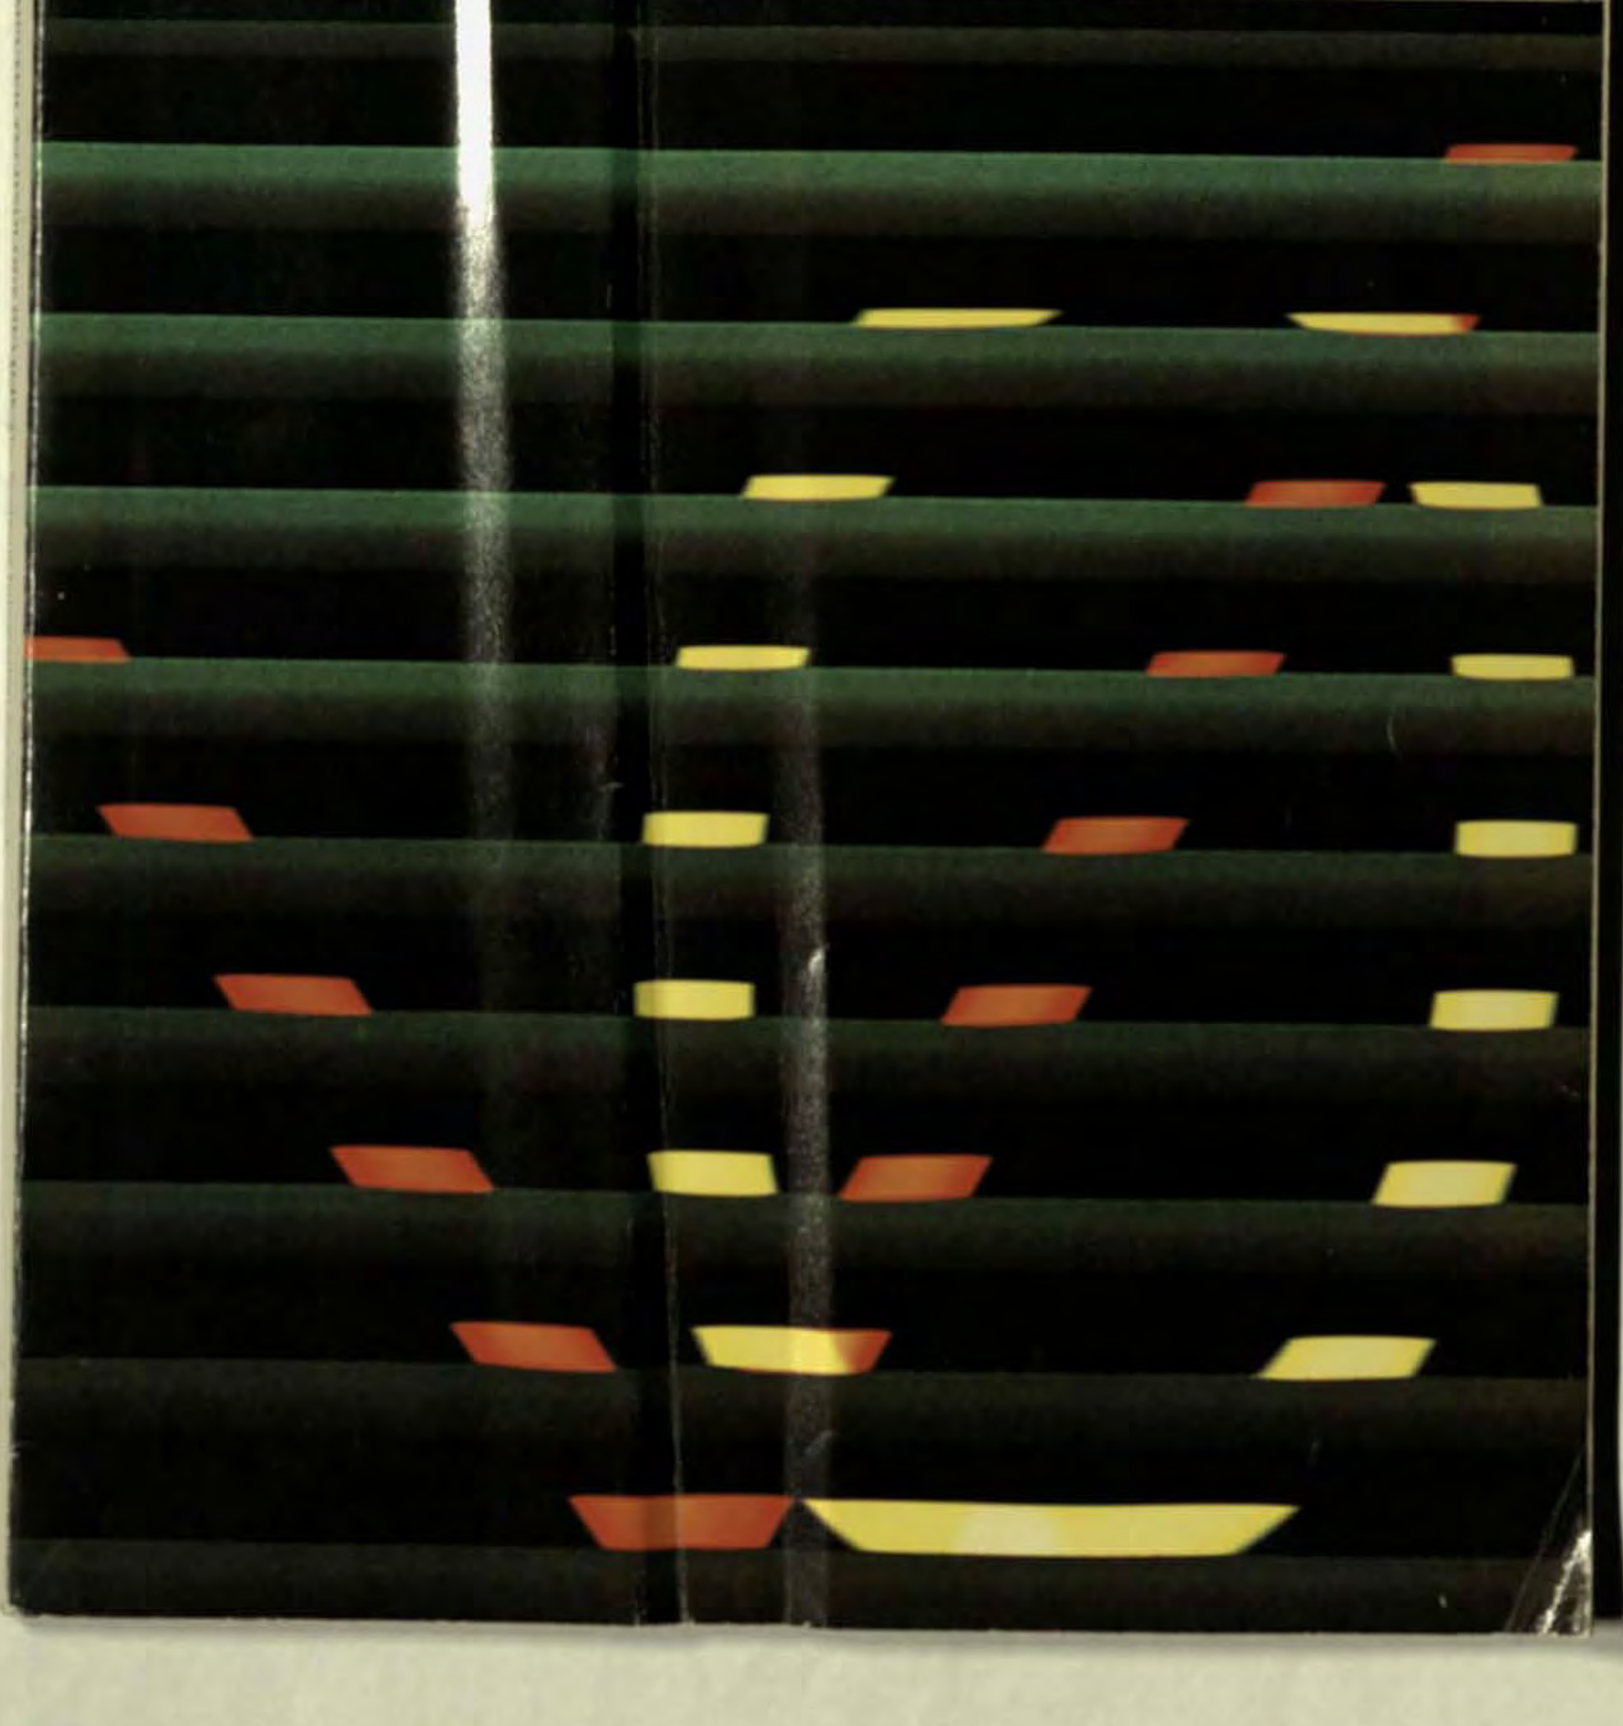

Feature For Feature We Out-perform Everyone

And With a Pricing Structure That Will Pleasantly Surprise You.

We at Falco Data Products are seriously committed to offer the DP profession the absolute best in emulations with emphasis on *quality* and *end-user benefits* as well as real price/performance. All terminals come with standard P31 Green screen. Call *NOW* for full information regarding any of these Brilliant™ terminals.

TS-2624

- HP 2624A* Emulation
- Software Controlled Setup
- 4 Pages of Screen Memory
- 64 Character Line Drawing Set
- 8 Preprogrammed or User Definable Screen Labeled Function Keys
- Local Printer Port
- Modified Data Tags

TS-100SP

- DEC VT-100† Emulation
- DEC VT-52† Emulation
- Conversation and Block Modes
- 14 User Programmable Function Keys
- Local Printer Port
- Split Screen
- Smooth Scroll
- 100 Special Graphic Characters
- 16K of Screen Memory

TS-132

- VT-132† Emulation
- VT-100 Emulation
- VT-52 Emulation
- 80 or 132 Characters Per Line
- 14 User Programmable Function Keys
- Advanced Video Feature Standard

TS-2

- ADM-2‡Emulation
- ADM-2 Style Keyboard
- ADM-1‡Mode
- Software Controlled Setup
- 32 Preprogrammed Function Keys
- 12 User Programmable Non-Volatile Function Keys
- Printer Port

TS-42

- ADM-42‡Emulation
- 8 Pages of Screen Memory Expandable to 12 Pages
- Code Compatible with ADM-2
- 32 ADM-42 Preprogrammed or User Definable Function Keys
- 12 Non-Volatile Function Keys
- Local Printer Port

Nationwide Service by Dow Jones

*Hewlett-Packard
†Digital Equipment Corp.
‡Lear Siegler Inc.

Our new TS-212^PLUS INTEGRAL MODEM provides cost effective, high performance 300 or 1200 BPS; fits all Brilliant Terminals. Just some features: • Microprocessor Controlled • Stored Number Dialing • Keyboard and Computer Dialing • Dial Until Connection

Call now

(408) 745-7123

1286 Lawrence Station Road
Sunnyvale, CA 94086

Falco
DATA PRODUCTS

Our Monitor Line Starts Here.

Mitsubishi's unsurpassed selection includes monitors for computer graphics, imaging, CAD/CAM, management systems, personal computers, video games, and more. In monochrome. In color with sizes from 12″ through 26″ diagonal. With your choice of slot or dot mask, a wide array of phosphors, and resolution ranging from an ultra-high 2,000-line version to MIS models. And if that still isn't enough variety for you, we'll custom-manufacture to your specs on orders of a thousand or more.

Our Line Doesn't Stop Short On Advanced Features, Either.

Mitsubishi's super-high contrast glass, for instance, delivers a 35% contrast enhancement when compared to CRTs with conventional clear faceplates. Further, an extremely dense black combines with our extended color range to give a deeper tone to screen images. All monitors can incorporate a special anti-reflective coating to cut glare. And all are designed to assure high-fidelity reproduction with exceptional reliability.

Packaging? Compact, light, and very easy to interface. And as you'd expect, we offer a choice of physical dimensions, plus rack-mount and lowboy configurations.

We'll Also Go Out Of Our Way To Help You.

With a stocking policy that assures you of a high product flow, as well as improving your ROI and cash flow positions. Plus application assistance, pricing specifics, and descriptive literature.

Write or call for details. Mitsubishi Electronics America, Inc., Display Products Group, 991 Knox Street, Torrance, CA 90502. Call (213) 515-3993; Eastern U.S. (617) 938-1220.

MITSUBISHI ELECTRONICS
ADVANCED AND EVER ADVANCING

COMPUTER TECHNOLOGY REVIEW

The Systems Integration Sourcebook

From the Publisher

More Valuable Than Ever Before!

It's becoming abundantly evident that what is reawakening the economy from its severe depression is the renaissance of American technology. Productivity is increasing. Quality is improving. Design is getting more innovative. Hardware is working harder, software is working smarter in goods and services for factory, office, and home.

At the core is technology. And at the core of computer technology is COMPUTER TECHNOLOGY REVIEW—the one publication that puts it all together for the computer professional, effectively bridging the gaps of the computer market segments previously serviced only by vertical publications. We cover the minicomputer, micro, and mainframe markets so thoroughly that even if the computer professional read all the pertinent vertical publications he still wouldn't get the depth or breadth of information he receives in each single issue of COMPUTER TECHNOLOGY REVIEW; information that's essential for him to make intelligent buying decisions when integrating and designing computer-based systems.

COMPUTER TECHNOLOGY REVIEW is no ephemeral buyer's guide nor unwieldy directory of goods and services. Rather, it is the unique, comprehensive Systems Integration Sourcebook. As such, its value for systems and engineering management as well as corporate management has grown dramatically from issue to issue. This is our latest issue—the largest and finest we have ever produced. It is my sincere hope that it becomes the most valuable yet for you in your career.

Best personal regards,

Yuri R. Spiro
Publisher

P.S. To ensure that you're on the list to receive future issues, please fill out, date, and sign the reader qualification card attached to the front cover of this issue.

We Print Better Bottom Lines

SSI printers and computer/printer interfaces work better and faster

and they work for less

The more you know about SSI printers and interfaces, the easier your printing problems look.

We have answers. Answers to interfacing mainframes and minis to the printer you're using. Answers to plug-and-go compatibility with IBM. Answers to controllers (including long-line) for DEC, Data General, Perkin-Elmer, Texas Instruments. Serial interfaces, synchronous and asynchronous, with the protocols you need.

And we make printers. Printers with eye-opening performance, at eye-opening prices, from 140 cps up…to five times higher than their counterparts. They drive further distances…cost as little as **half.**

For state-of-the-art, we offer you our Mercurion 1 printer system—using non-impact ion deposition imaging and giving you speeds to 5,280 LPM (a page per second). Resolution is 240x240 pixels, landscape or portrait.

It's really time you knew us. That's the bottom line. Call or mail the coupon now.

I want to print better… with a better bottom line

☐ We have a __________-type computer.
☐ Have rep contact me.
☐ Send me literature. Our major interest is __________

Name ____________________

Title ____________________

Company ____________________

Address ____________________

City, State, Zip ____________________

Southern Systems, Inc.

2841 Cypress Creek Road
Ft. Lauderdale, FL 33309
(305) 979-1000 (800) 327-5602
Telex 522135

Southern Systems Canada:
(416) 229-2589

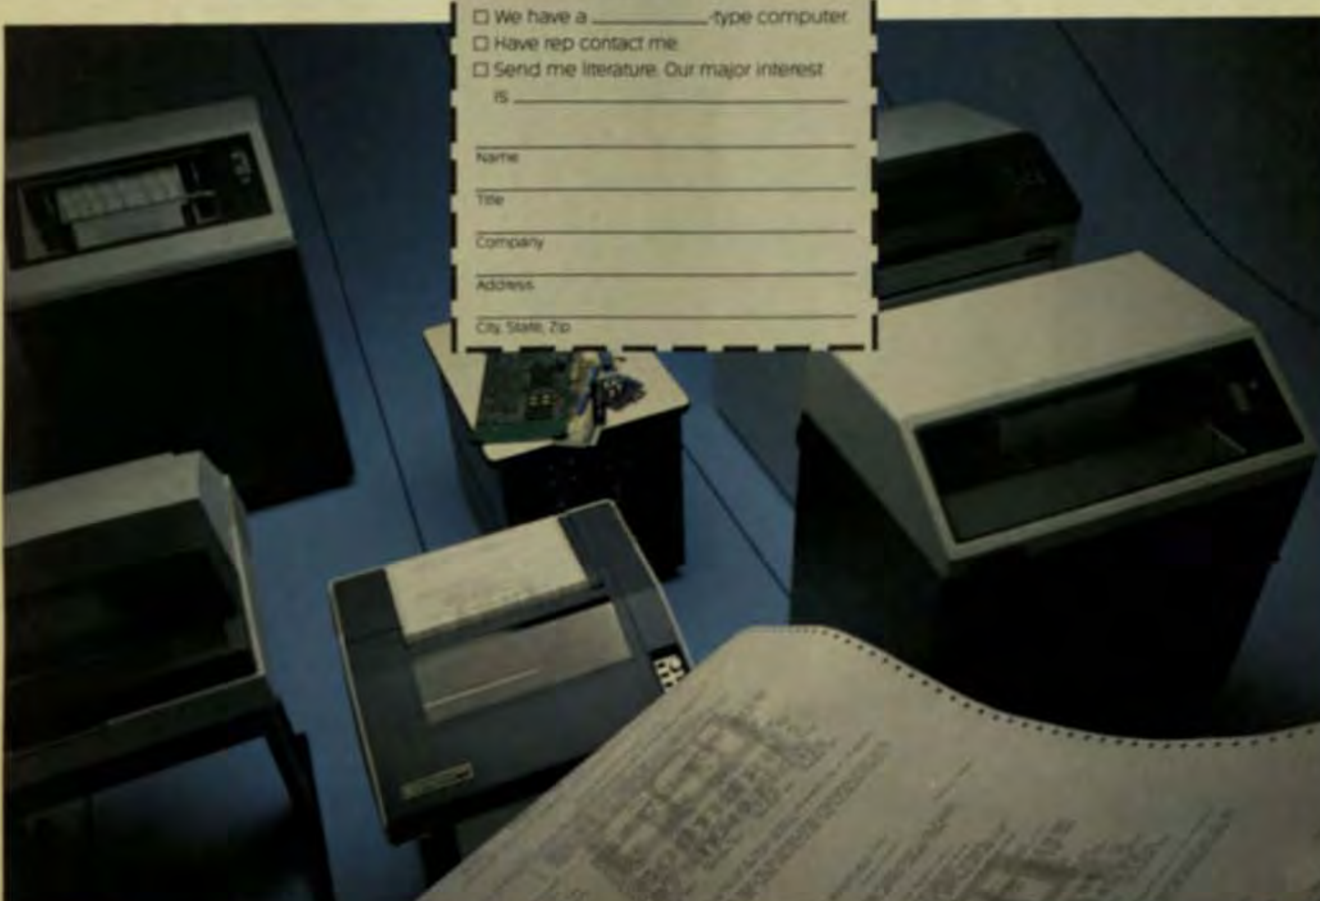

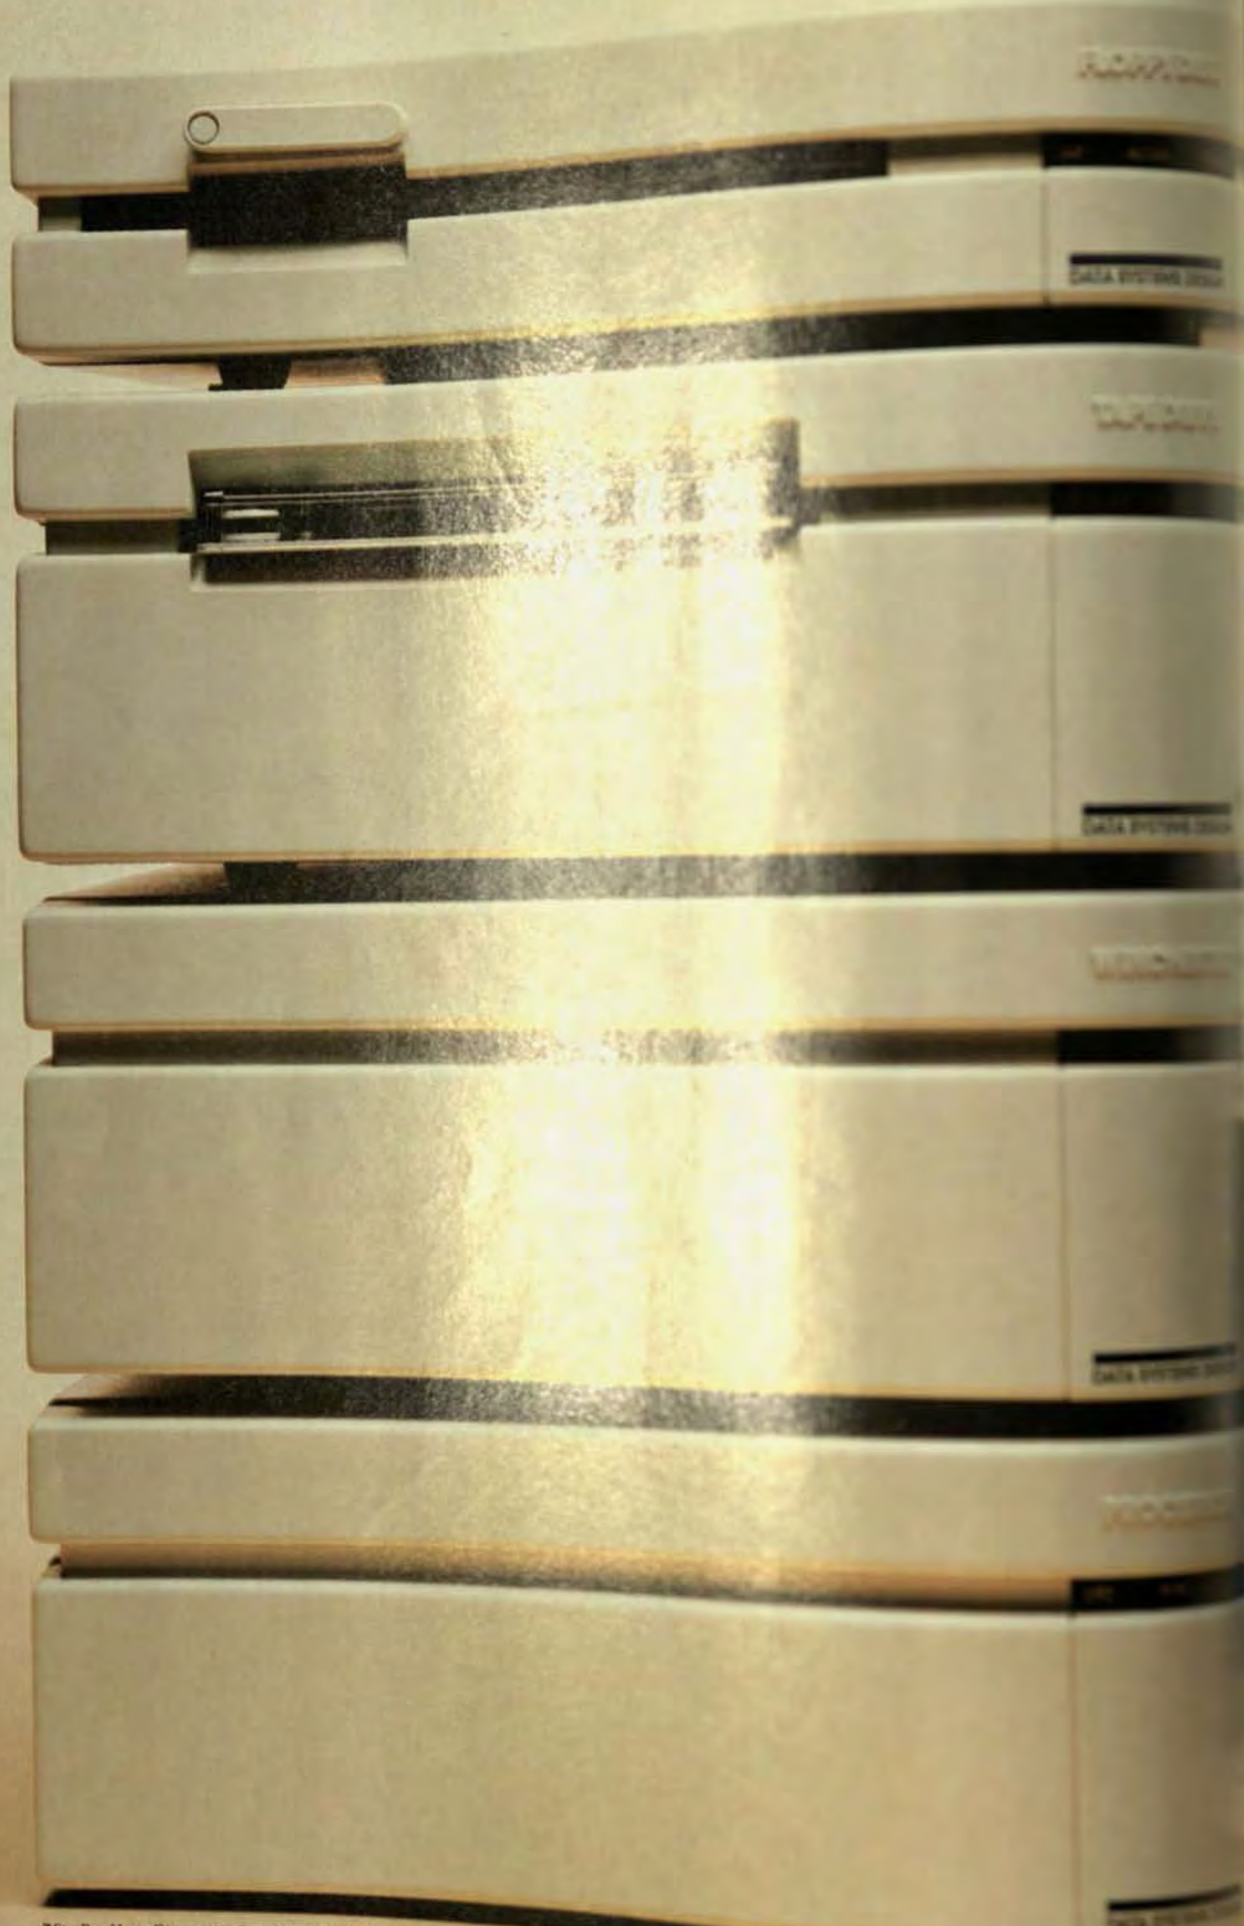

™StacPac, HyperDiagnostics, Rapid Module Exchange and HyperService are trademarks of Data Systems Design, Inc.
®DEC is a registered trademark of Digital Equipment Corporation. ®Multibus is a registered trademark of Intel Corporation.

DSD STACKS THE DEC. AND THE MULTIBUS, TOO.

Presenting StacPac™ systems and modules.

This is a brand new deal for systems integrators.

StacPac systems and modules. Modular storage and backplane units that stand alone or stack together to form the basis of a small but powerful tabletop computer.

They're available in either DEC®- or Multibus®-compatible versions. And there are lots of advantages to both.

On the DEC side, you can put a StacPac system in places a rack-mount just doesn't fit.

On the Multibus side, you can use our StacPac system to deliver serious minicomputer performance with microcomputer economy.

Either way you get greater flexibility to configure just the right system. And the easy upgradability to expand it whenever you like.

But perhaps one of the most important advantages is that you can protect your software investment—and your customers'—because inside, you're still selling the same system. You'd just never know it to look at it.

Which brings up another point.

Underneath that slick, compact exterior is the very latest in storage technology. 8" slimline floppy drives. High capacity, highly reliable 8" Winchesters. Compact ¼" cartridge tape drives. And the best high performance controllers in the business. (Packaged systems are also available without controllers at your option.)

And if you're thinking about 5¼" storage, there's our 5¼" Winchester/floppy module. (Available in the Spring of 1983.)

Configuring an elegant system solution is a simple matter of picking out the storage option you want and adding your own CPU boards, I/O and memory to our uncommonly accommodating backplane units.

Naturally, we back it all up with HyperDiagnostics,™ Rapid Module Exchange™ and HyperService.™ Some of the most economical and intelligent service features ever devised.

And you can have more information about StacPac modules just by writing for our brochure.

But you just watch. You start marketing your systems in our StacPac modules and your competition will say you have an unfair advantage.

And you know something? They'll be right.

Corporate Headquarters: 2241 Lundy Avenue, San Jose, CA 95131. Eastern Region Sales and Service: Norwood, MA, (617) 769-7620. Central Region Sales and Service: Dallas, TX, (214) 980-4884. Western Region Sales: Santa Clara, CA, (408) 727-3163.

DATA SYSTEMS DESIGN

INTERNATIONAL SALES: Australia 03/544 3444; Belgium and Luxembourg 02/7209038; Canada 416/625 1907; Denmark 02/63 22 33; Finland 90/88 50 11; France 03/411 5454; Hong Kong and Peoples Republic of China 03/696231; Israel 52-52444; Italy 02/4047648; Japan, Osaka 06/323 1707, Tokyo 03/345 1411; Netherlands 02977-22456; New Zealand 04/693 008; Norway 02/78 94 60; Singapore, Malaysia, and Indonesia 2241077; Spain 01/433 2412; Sweden 08/38 03 70; Switzerland 01/741 41 11; United Kingdom 7073/34774; West Germany and Austria 089/1204-0; Yugoslavia 61/263 261

★ ★ ★ E X T R A ★ ★ ★

32-Bit Computer Breaks $10K!

Universe 68/05 First to Smash Price Barrier

The new Universe 68/05 is the first true 32-bit computer priced under $10,000 (OEM quantity one). "True" because, unlike other 68000-based systems, the Universe 68/05 handles 32 bits in parallel on its VERSAbus.

Outperforms VAX*

Its *price* is even more impressive when you look at Universe 68/05 *performance* versus that of 32-bit "superminis" several times more expensive, like the VAX-11/750.

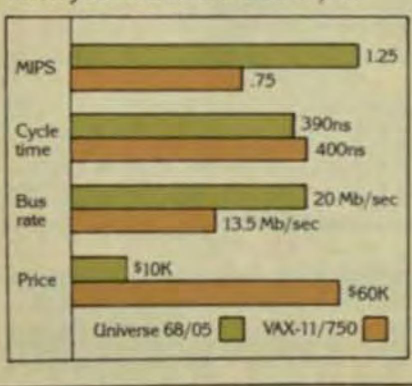

High-Speed 68000, 4Kb Cache, 32-Bit Bus

The key to that performance is a 4Kb cache that eliminates processor wait-states and takes full advantage of a 12.5MHz 68000 processor. Also included are a separate 68000 I/O processor, four serial I/O ports (expandable to 64), 256Kb RAM (expandable to 3Mb), 20Mb/sec, 32-bit VERSAbus, 10Mb Winchester, 1.25Mb floppy disk, and 5-slot card cage. All in a 7-inch enclosure.

UNIX-Compatible Real-Time OS, Too

UNOS*, our UNIX* Rev7-compatible operating system with real-time features, runs Pascal, Fortran, C, BASIC, DBMS, and third party application programs.

For more information, just attach your business card to this ad and mail to Charles River Data Systems, 983 Concord St., Framingham, MA 01701. Or call us at (617) 655-1800. We'll send you a copy of "The Insider's Guide to the Universe," a detailed discussion of the technical concepts behind this remarkable new computer.

CHARLES RIVER DATA SYSTEMS

*VAX is a trademark of Digital Equipment Corporation. UNIX is a trademark of Bell Laboratories. UNOS is a trademark of Charles River Data Systems.

Architecture Determines Cost and Reliability of Fault-Tolerant Design

Different architectures are available to achieve fault-tolerant computing. Cost-effectiveness and reliability are now available for commercial environments.

by **Dennis McEvoy and Sandra Metz,**
Tandem Computers Inc.

Fault tolerance, or the ability of computers to recover automatically from failures and continue processing, is becoming a critical element in systems used in an on-line, interactive mode. There are several fault-tolerant system architectures and configurations available today, each optimized for the needs of different high-reliability markets. This article will discuss several ways to achieve fault tolerance in computing systems, and the relative merits of each.

A classical system architecture typically consists of a centralized processor, an I/O channel, and controllers to manage such peripherals as terminals and disks. Because a failure in any one of these components can take the entire system off-line, this configuration is acceptable only for applications that don't require continuous system availability.

Applications that demand continuous availability typically require some kind of fault-tolerant computing system. Such a design must strategically duplicate components in an efficient and cost-effective manner. The methods of duplication—and the resulting levels of fault tolerance—vary dramatically and are suitable for a wide range of applications.

SWITCHED BACKUP

A very basic approach to fault-tolerant computing is to duplicate the central-processing unit. Fig 1 shows a dual-processor system in which one processor acts as a backup for the other. The processors are connected through a switch, and either one can communicate with the other components to control the entire system.

This standby configuration is known as switched backup and provides the benefit of fault tolerance in the case of a single processor failure. Switched backup is one of the most commonly used configurations for large, on-line mainframe applications.

Several problems are associated with switched backup. One involves every company's biggest concern—dollars. Since the second processor is a backup, it remains idle until the first processor fails. If the backup happens to be a $2 million mainframe, a fairly expensive piece of hardware will be idle until a failure occurs in the on-line system.

However, to gain better return on your investment, you can use the backup system for program development while the on-line system is still functional. In that case, however, when the on-line processor fails, the development work on the backup processor must be halted so that the backup system can be switched on-line.

Another problem with a switched backup is that it provides fault tolerance for only one situation—a single processor failure. An I/O-channel failure between the bus switch and disk controller, for example, would quickly turn both $2 million mainframes castors up. The bus switch itself is another single point of system failure.

OEM Discounts Available.

Introducing the Auragen of the computer species.

Now a higher form of computer technology has evolved. It combines fault tolerance, high performance and ease of use. The natural selection for your business transaction needs.

Now you can have fault tolerant applications at price/performance levels never available before. For under $138,000, our entry level 2-Cluster system provides up to twice the performance of today's fault tolerant systems.

AURAGEN™ begins with an advanced 32-bit distributed architecture. Both your professional and end-user staff will find it easy to learn and use. AUROS,™ our operating system based on UNIX,™ is enhanced with performance, security, privacy and recovery capabilities specifically tailored for the transaction market. We provide ANSI 74 High Level COBOL, FORTRAN 77, Pascal, C and even business BASIC that run fully fault tolerant. No special instructions are necessary. No reprogramming. No retraining. No added expense.

We also provide a full range of interactive productivity tools for programmers, including AURAGEN's full-function relational DBMS and Transaction Processing Management System. Then we've added menus, help facilities, and an English-like query language to make your staff more productive. The extensive communications packages that we provide include SNA, BSC, X.25 and more.

The concept that sets us above and apart from all other fault tolerant systems is simple. Our fault tolerance is totally transparent and is so efficient you'll never want to turn it off. Our backup hardware is available to do productive work during normal operations. No downtime. No manual backup procedures. No added expense.

With AURAGEN's distributed design, you can grow in precisely controlled increments. The basic fault tolerant system contains two Clusters. Each Cluster contains several tightly coupled MC68000 processors. Separate processors handle fault tolerance, user programs, communications, and database functions. The system can be expanded to 32 Clusters.

AURAGEN is so advanced, you'd select our system even if we didn't provide fault tolerance. But we do. In today's competitive business environment, you need every survival factor working for you. We provide them all. And we want to tell you more about them. That's why we invite you to ask for more information.

Auragen
Systems Corp.
THE NATURAL SELECTION

AURAGEN Systems Corp. **CTR/583**
Two Executive Drive, Fort Lee, N.J. 07024
Call toll-free 800-221-1741. In New Jersey (201) 461-3400.

I'm interested in the AURAGEN System. Tell me more.

Name ____________________
Title ____________________
Company ____________________
Address ____________________
City __________ State __________ Zip __________
Telephone ____________________

*AURAGEN and AUROS are registered trademarks of AURAGEN Systems Corp.
*UNIX is a registered trademark of Bell Laboratories.
© 1983 AURAGEN Systems Corp.

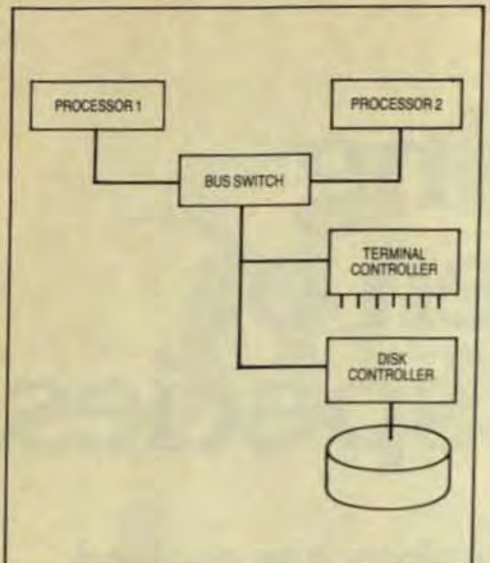

Fig 1 A very basic approach to fault-tolerant computing is switched backup, which duplicates the central-processing unit. A switch connects two processors so that either can communicate with the other system components and control the computer as a whole.

TRIPLE MODULAR REDUNDANCY

A higher level of fault tolerance is achieved with a configuration known as triple modular redundancy. Like switched backup, triple modular redundancy is a hardware-only approach, but it goes a step further by providing three copies of every piece of system hardware (Fig 2).

The three processors are lock-stepped together and simultaneously run the same instruction stream. A piece of hardware called a voter receives and compares the output from all three processors. If one output doesn't match the others, the voting hardware accepts the output from the two matching processors and passes it along to the appropriate interface.

This approach has been used successfully for many years in such areas as space research and nuclear power plants, where computer failure is unacceptable. In the commercial environment, however, triple modular redundancy poses two major problems: tripled cost and complex design. A single clock must run the entire system, and the clock must be implemented so that it can't possibly fail. This is a technically complex requirement.

These limitations weren't a stumbling block in applications like the U.S. space program, where literally billions of dollars were riding on the fact that the computer would be available 100% of the time.

COMPARISON LOGIC

A similar approach is to have four processors running the same instruction stream (Fig 3). The processor outputs are paired with each other and compared. If a mismatch occurs, the pair containing the mismatch is shut down. The system continues running with the remaining pair of processors.

Although similar to triple modular redundancy, such comparison logic has the advantage of being easier to implement than voter logic, and therefore is less expensive to design and maintain. However, the financial problems of redundant hardware all running a single instruction stream and performing identical operations still remain.

Also, comparison logic causes a good processor to be taken off-line along with the problem unit, thus wasting a resource that could be working with the system. And, again, the clock is the weakest link in the system, a single point at which the entire system could fail.

A MULTICOMPUTER SYSTEM

A different approach to fault-tolerant computing, involving a unique, integrated hardware/software architecture, was taken by Tandem Computers. Tandem's NonStop system was designed specifically for the on-line transaction-processing marketplace. This arena requires a system architecture that's not only fault-tolerant and easily expandable, but also maintains a price/performance ratio suitable for commercial transaction-processing applications.

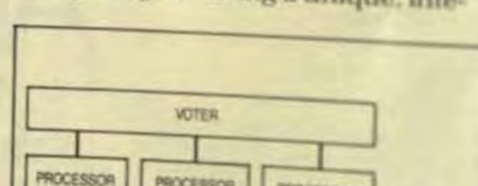

Fig 2 A higher level of fault tolerance is achieved with triple modular redundancy, which provides three copies of every piece of system hardware. The three processors are lock-stepped together and voter logic receives and compares the output from all three.

This architecture combines hardware and software so that there's no single point of failure (Fig 4). Components are duplicated but are not redundant. Each processor runs its own instruction stream, so no backup components sit idle until a failure occurs.

Fault tolerance is achieved by keeping each processor in constant communication with the others in the system via a dual high-speed interprocessor system bus. If one processor fails, its workload will automatically be absorbed by the remaining processors.

Each processor has its own I/O channel and each I/O controller is dual ported. Thus, if one processor fails, another can still use that controller. Since disks are physically duplicated, with identical data stored on each, even a disk crash can't bring down the system. The duplicate disks pay for themselves by ensuring complete system fault tolerance and improved performance. One disk performs seeks on the inner portion of its surface while the other disk performs the same seeks on its outer portion. This provides a significant performance improvement in seek time.

The software architecture is based on independent processors and independent programs within them, all of which communicate through a message system that is an integral part of the operating system. This operating system sees all programs and data transfers as communications distributed over several processors.

Programs can access any device anywhere in the system, even those not physically connected to the processor running the program. Conversely, each program is unaffected by the processor on which it runs. The operating system sees all physical resources as logical files. Only the message-routing part of the operating

SOMETIMES, THE BEST PLACE TO GET A DEC PDP-11 SYSTEM ISN'T DEC

When you buy DEC PDP-11 based System 58 from Cambridge Digital, it's ready to get to work as soon as you hook it up. Because at Cambridge Digital, we sell enhanced-performance DEC systems that are fully integrated with DEC-compatible products. Products you can't buy from DEC. So you spend your time working on your system, instead of just trying to make it work. It's part of the edge you get when you buy from Cambridge Digital.

Take our ready-to-run System 58. It features a DEC LSI-11/02, 21, or 23 with internal RAM memory of 32Kb, expandable to 4Mb. The 11/23 based system directly addresses up to 4Mb of memory. Plus you get an RLV12 compatible Winchester disk with you choice of 10, 20 or 40Mb, and a floppy with usable capacity of 1.02Mb.

And, you get a full range of development tools from CDS, too. Like a fully-supported and enhanced operating system based upon UNIX* System III. Also TSX-Plus, RT-11, μ-power PASCAL, and RSX-11M.

It's all ready to run. Available for as little as $7,350. Available in as little as 10 days. But its not available from DEC, so call Cambridge Digital, instead, and get The Edge.

For more information on our System 58 and a copy of our Building Blocks For Systems Catalog, including a description of the seven guarantees you get when you get The Edge, call or write.

Main Office Dept. 7509, P.O. Box 568, 65 Bent Street, Cambridge, Massachusetts 02139. Telex 92-1401/ COMPUMART CAM 800-343-5504 In Massachusetts 617-491-2700

New York District Office 516-935-3111

*UNIX is a trademark of Bell Laboratories.

I want The Edge:

Name_____ Title_____

Organization/Company_____

Address_____

City_____ State/Province_____

Zip/Postal Code_____ Country_____

Phone No. (   )_____

My main product interests are: DEC ☐ NCR ☐

(7509) 58

Cambridge Digital

DIVISION OF COMPUMART

The Edge in System Integration

800-343-5504

In Massachusetts call 617-491-2700

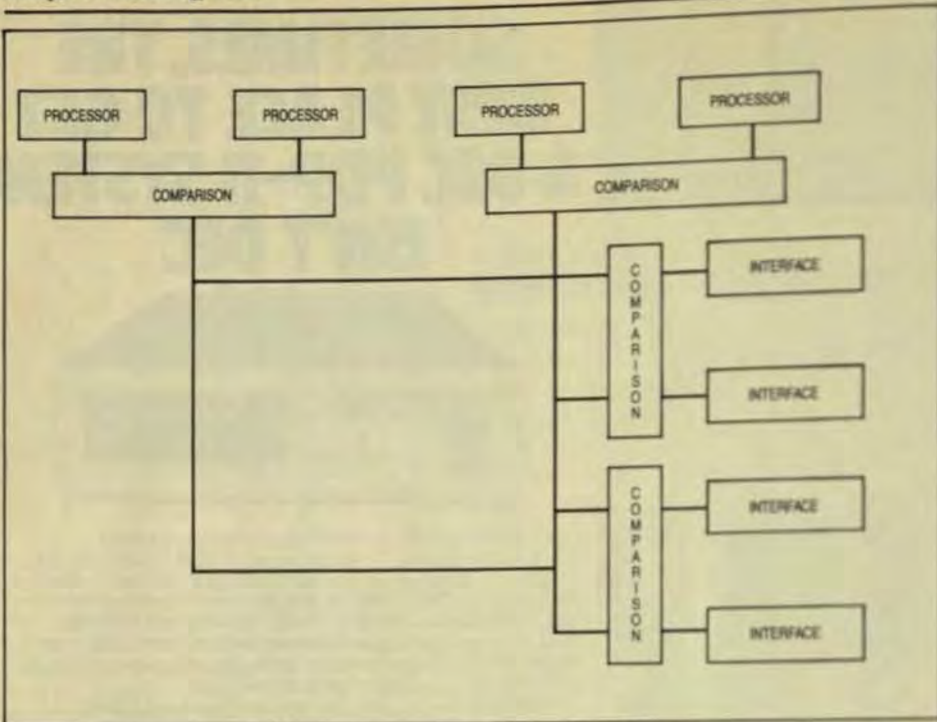

Fig 3 An approach similar to triple modular redundancy is to have four processors running the same instruction stream, with their output paired and compared. If a mismatch occurs, the pair containing the mismatch is shut down and the remaining pair continues to run the system.

system knows the geographic locations of resources, so data can be rerouted and resources dynamically reallocated during a failure.

Software fault tolerance is improved by distributing the operating system across all the processors in the system. If one processor fails, its job will be picked up by those remaining. In a system where three or four copies of the hardware are lock-stepped together, a severe software bug could crash all the processors simultaneously. Past experience shows that software bugs severe enough to crash a system are typically related either to timing or to the placement of data in memory. In a NonStop system, however, a software failure will probably not crash more than one processor because the timing circumstances and exact placement of data in memory are unique to each processor.

SYSTEM EXPANSION

On-line transaction processing is characterized by the need for easy system expansion. For example, if a banking application initially involves control of 50 automated teller machines and the application is successful, the bank will typically wish to expand the automated-teller network.

The interprocessor bus system is the key to easy hardware expansion because each processor module has its own memory, I/O channel, and bus controller. To expand this system, additional processor modules are simply added to the bus. The message system keeps track of all system resources, so the system software load is automatically allocated to the appropriate hardware. The system can then be fine tuned for optimum performance.

It's possible to have a modularly expandable multiprocessor system where the processors are linked through a shared memory. A shared-memory architecture provides fast data access to all the processors and provides additional processing power. But the additional processing power obtained doesn't increase linearly.

For example, a second processor doubles your investment and provides only 1.7 to 1.8 times the power of a single CPU because of contention for the shared memory and associated shared resources. A third unit provides approximately 2.4 times the power for three times the cost. By the time the fourth processor is added, it's possible to reach a state of diminishing returns.

One way to get around the contention problem is to locate a cache memory in front of each CPU. Each

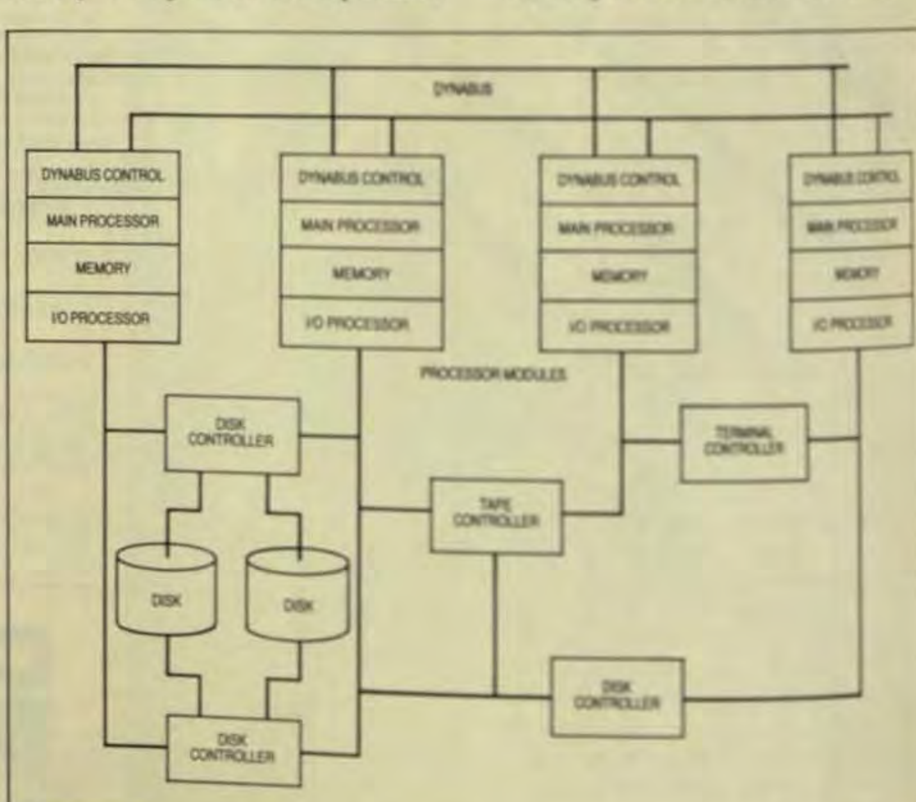

Fig 4 In a multicomputer system architecture, components are duplicated but not redundant. Each processor runs its own instruction stream and none of the backup components sits idle waiting for a failure.

cache still has to fetch instructions from memory though, so that while the problem may be minimized, the limits of the shared memory will be reached eventually.

A multicomputer software architecture eliminates such system bottlenecks as memory contention because the operating system runs in each processor, and the only communication between CPUs occurs when a user program requests I/O from another processor. This communication is done over the system bus, which runs at an aggregate rate of 26 Mbytes/s. Since the processor memory runs at 5 Mbytes/s, bus contention doesn't cause a problem.

Like the shared-memory approach, multicomputer architecture is modularly expandable. But because interprocessor communication occurs in the system bus, memory contention is eliminated along with the potential single-point failure of a shared memory. Therefore, a dual-processor system yields a full two processors' worth of computing power, three units yield three processors' worth of power, and so forth. There's virtually no performance performance degradation associated with expandability.

Another advantage of multicomputer architecture is the networking scheme inherent in a single system. Since the message system already keeps track of the physical location of all resources, a configuration can be expanded into a network of systems without any changes in existing software. The message system keeps track of each system as easily as it keeps track of each processor, and the user doesn't require any additional programming.

This networking scheme also contains a distributed database capability, so data files can be distributed among a network of systems, while the message system automatically keeps track of their locations. The user doesn't need to know where a specific piece of data resides in order to access it.

CHOOSING A FAULT-TOLERANT SYSTEM

The choice of a fault-tolerant system should depend on the application. Some require 100% availability, but only during critical time periods. An automated betting system at a race track, for instance, must be available on race days, but could be serviced at night. In contrast, a system that monitors medical equipment in a hospital must be available continuously, even during periods of maintenance and repair.

You can determine the degree of fault tolerance required by analyzing the cost of failure for a given application. For example, suppose a system advertised a guaranteed up-time of over 99%. That may sound fairly reliable, but 1% downtime translates into approximately a third of a day each month. If the application involved is control of automated teller terminals and the downtime occurs during a Friday lunch hour, the 99% uptime won't soothe the angry customers. The cost of failure in this case requires 100% system availability.

Other considerations include the intial system cost, ease and cost of expansion and modification, and the vendor-supplied system software. Also, if a database management system is needed for the application, you should choose a vendor who can supply one because it's difficult to add features like networking and database management to a fault-tolerant system unless they're designed into it at the start. ■

Dennis L. McEvoy *is vice president of software development at Tandem. He joined in 1974 as a member of the original software development team and project manager for the development of the firm's NonStop operating system.*

Sandra Metz *is a technical public relations specialist at Tandem. She joined the firm in 1982 after working as a writer for Hewlett-Packard and holds a BS in technical writing from Carnegie-Mellon Univ.*

ALL IN ONE. ALL FOR $50

All of the significant state-of-the-art advances for systems integrators are covered in-depth in separately edited sections in every issue of Computer Technology Review.

It's like receiving 16 separate publications in each single issue. Order today and receive CTR and it s updated supplements for an entire year for only $50.00.

COMPUTER TECHNOLOGY REVIEW

The Systems Integration Sourcebook

924 Westwood Boulevard, Suite 650
Los Angeles, California 90024

Ferrite vs. thin film may not be the either/or situation you think it is.

In fact, this free booklet from Memorex explores a number of ways to increase your current storage capacity. All without making a "forever" commitment one way or the other. Or backtracking. Or disrupting your entire operation.

It covers the pros and cons of ferrite head and thin film technology. (And since we build both, we can be pretty objective.)

Plus short term additions. Long term planning. Leasing. Buying. Bridging.

Bridging? Sure—the best of both worlds. It's a great way to solve your immediate problem without creating an even bigger one later.

For really fast action, call right now and your free copy will go out in tonight's mail. Or you can send us the coupon if you can stand a few more sleepless nights.

MEMOREX

A Burroughs Company

Call now, toll free. 800-538-9303

Memorex Corporation
Attn: Marketing Communications MS 1217
San Tomas at Central Expressway
Santa Clara, CA 95052

Send my survival manual.

NAME ______

TITLE ______

FIRM ______

ADDRESS ______

CITY ______ STATE ______ ZIP ______

PHONE ______

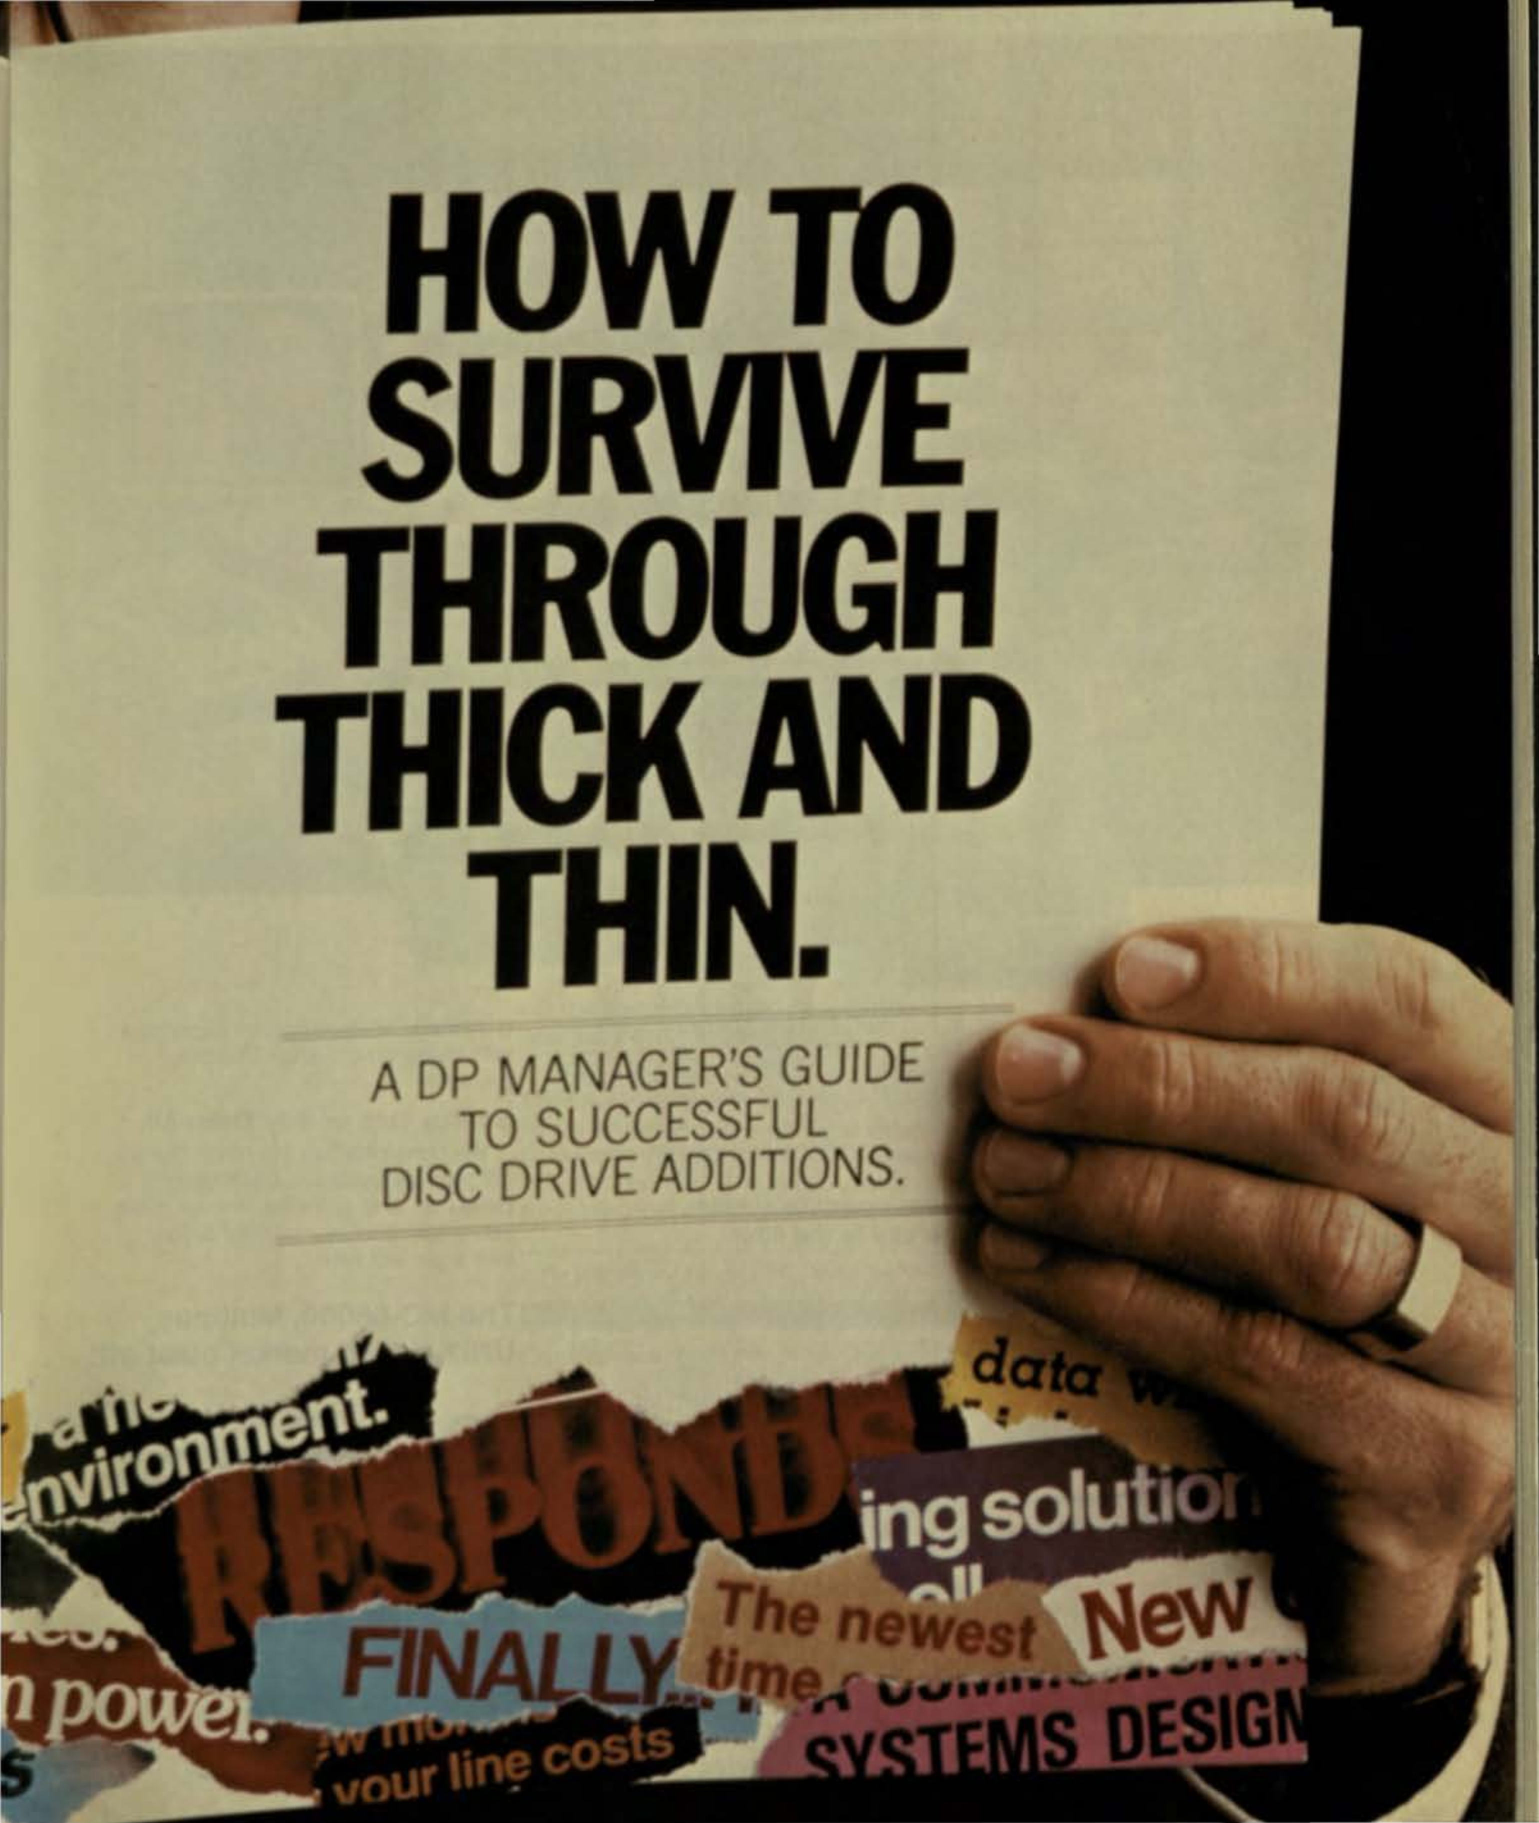

HOW TO SURVIVE THROUGH THICK AND THIN.
A DP MANAGER'S GUIDE TO SUCCESSFUL DISC DRIVE ADDITIONS.
nvironment.
data
ing solution
The newest
New
FINALLY...
power.
your line costs
SYSTEMS DESIGN

SYSTEMS INTEGRATORS GUIDE TO THE UNIVERSE

One stop shopping for the Systems Integrator and OEM.

Imagine being able to buy everything you need for the product you are developing from ONE supplier. Multibus boards? Yes. Fully assembled systems? Yes. A real-time multitasking coprocessor for UNIX™? Yes. . . . and yes; Bell Labs System V UNIX™, the supported UNIX ported by Unisoft. And COSMOS has software to get you one step closer to your final image processing product or office automation or graphics product.

Our products are compatible . . . part of the same solar system.

For example COSMOS rt/UNX is a real-time, multitasking coprocessor which provides real-time capabilities to UNIX systems without modification to the standard UNIX kernel. This means that all UNIX facilities including the UNIX file system and debugging facilities are available to the real-time application. And the UNIX processes can even monitor the application.

We've also designed a powerful single board computer for you to use in a stand-alone or multi-user environment to run our real-time operating system and your application. Or use our real-time facilities, our graphics capability, and our board-level products as key elements in your process control system.

From board to system, the sky is the limit.

You can start with COSMOS board-level components or take advantage of our STARFIELD Series—a universe of system configurations which includes ANTARES for major development projects, the ORION system for developing and running scientific applications, or LYRA, the perfect target system. These configurations can be applied to a galaxy of processing solutions: For example, ORION includes 1 MEG ECC memory, serial and parallel I/O, 474 M of mass storage, a 1/2" streamer tape plus floating-point or array processor and an optional Ethernet interface and runs UNIX.

All the STARFIELD system products run UNIX and support C, COBOL, Fortran, PASCAL and BASIC. The STARFIELD Series can be expanded with many off-the-shelf Multibus products.

So Buy One or Buy Them All.

We don't promise the moon but we do promise systems, board level products, and software that will make you a star and launch your products into a perfect orbit.

The MC-68000, Multibus UNIX-based market blast-off!

Winning Combinations!

COSMOS SYSTEMS, INC.
A Subsidiary of C.M. Technologies, Inc.

430 Toyama Drive
Sunnyvale, CA 94086 (408) 744-0721
TELEX 334417 CABLE CHIMINOCO

Eastern Regional Office
Walt Waldron—Eastern Regional Sales Mgr.
569 River Road
Fair Haven, NJ 07701 (201) 530-7788

Western Regional Office
Tom McKee—Western Regional Sales Mgr.
30100 Crown Valley Parkway, Suite 32
Laguna Niguel, CA 92677 (714) 495-2661

™ Trademark of Bell Labs.

Integrated DP Support Systems Benefit From Advance Site Planning

Total integration includes the computer-support systems that keep the computers running. Pre-installation planning can increase reliability and cut costs.

by **David C. Holscher,**
Liebert Corp.,
and **Leslie D. Davidson,**
Merrill Lynch & Co., Inc.

The conventional view of systems integration generally covers a wide range of data-processing equipment, from central-processing units and mass-storage systems to printers, drives, and other peripherals. Total integration, however, includes the computer-support systems that keep the data-processing installation up and running. Although critical to overall system performance, these support considerations are usually left until last on the list of priorities. Four fundamental factors influence data-processing performance—speed, capacity, applications, and availability. Of these factors, availability is the one that underlies optimization of the other three. And maximum availability is the goal of computer support.

COMPUTER-SUPPORT SYSTEMS

Computer-support systems should offer four primary features:

- Precision to meet the unique needs of data-processing installations.
- Reliability to ensure maximum uptime.
- Cost-effectiveness, from purchase and installation to operation and maintenance.
- Flexibility in adapting to change and growth.

A packaged approach to a dedicated support system offers several advantages in providing these features—system compatibility, factory testing of the total system, potential investment-tax credits, single-source responsibility from one supplier, ease of planning and installation, and a modular approach for future needs.

Because of the critical nature of the relationship between the data-processing and computer-support systems, it's essential to incorporate support within the overall integration process. Three particular support areas are especially important because of their direct impact on data-processing operations. Those areas are:

- Precision environmental control.
- Electrical power protection.
- Centralized monitoring and control.

Other, noncritical support areas include fire suppression and access security.

A CASE STUDY

Merrill Lynch & Co. is a worldwide financial-services company that operates an extensive data-processing network through its operations systems group. The firm is vitally dependent on its information-management resources on a daily basis. Corporate revenues in 1982 exceeded $5 billion. The company handled 13% of publicly listed equity volume and over 30% of odd-lot New York Stock Exchange volume last year. That translates into an average processing volume of some 80,000 transactions per day.

Applications of this system include stock transactions, research-information retrieval, stock quotaan interior space of the building, away from outside walls, roof, and basement. It should not be near sources of water (such as kitchens and bathrooms), areas with flammable or explosive materials, or electrical rooms.

For effective environmental control, a total vapor seal enclosing the entire data-processing area will lock out unwanted moisture and heat. Proper insulation will also greatly improve the efficiency of environmental-control systems.

Four basic types of environmental-control systems are available: air-cooled, water-cooled, glycol-cooled, and chilled water. Selection is a function of existing facilities, operating costs, and energy-conservation goals. Liebert environmental control systems include monitoring and control components to maintain the environment within very tight tolerances and minimize operating costs by reducing humidifier and compressor operation. Dual refrigeration circuits provide standby capacity. Semihermetic compressors are energy-efficient, durable, and field serviceable.

Merrill Lynch has standardized on a 15T environmental-control system that can handle approximately 1000 ft² of typical data-center area. The units are placed regularly around the perimeter of the data center, with air flow directed under the raised floor. In a large data center, it's also advisable to place several units in the middle of the room. With the amount of cabling under the raised floor, this configuration will ensure even air distribution throughout the room. A redundancy factor of one unit for every five is built in initially to provide for system-operating problems.

Environmental-control considerations also extend beyond the data-processing hardware itself and should not be ignored. For example, paper stock must be kept at precise temperature and humidity to avoid printer jams. Tape storage has similar requirements. At the main facility in Manhattan, 35% of the entire building is related to data processing and requires precise environmental control.

Integration of an effective environmental-control system can provide substantial energy savings. In contrast to the general office environment, where lights and equipment can be turned on and off to manage energy consumption, the data center offers relatively few sources for energy management because the entire facility must be operational at all times.

The main data center is supported by 130 environmental-control systems. The three smaller centers in the New York area (Fig 2) use an average of 15 units, and the nine regional centers employ three each. Low-end environmental-control systems are installed at all 500 branch offices. In 1979, Merrill Lynch initiated an energy-conservation program that included an effort to minimize environmental-control energy costs at the main data center. The first step was to add a strainer-type cycle in the cooling towers that are used to release the heat accumulated by the environmental-control systems.

The second step was to specify for all new units, the use of cooling systems that combine traditional refrigeration methods with a second circuit that uses a glycol solution. The nature of glycol and its heat-rejection properties allow the system to minimize or eliminate compressor operation for much of the year. Since compressor operation comprises a major portion of overall system-operating costs, the savings can be substantial.

Nearly $1 million has been saved per year by such measures. The payback period for these environmental-control-system improvements was a short two years—an excellent return on investment.

Another energy-management option is to install heat-recovery modules to apply waste heat from the data center to other parts of the building. In some cases, an entire building can be heated by use of such a system. Climate conditions and the relative size of the data center in comparison with the overall building will determine whether this is a feasible option.

POWER PROTECTION

The companion to environmental con-

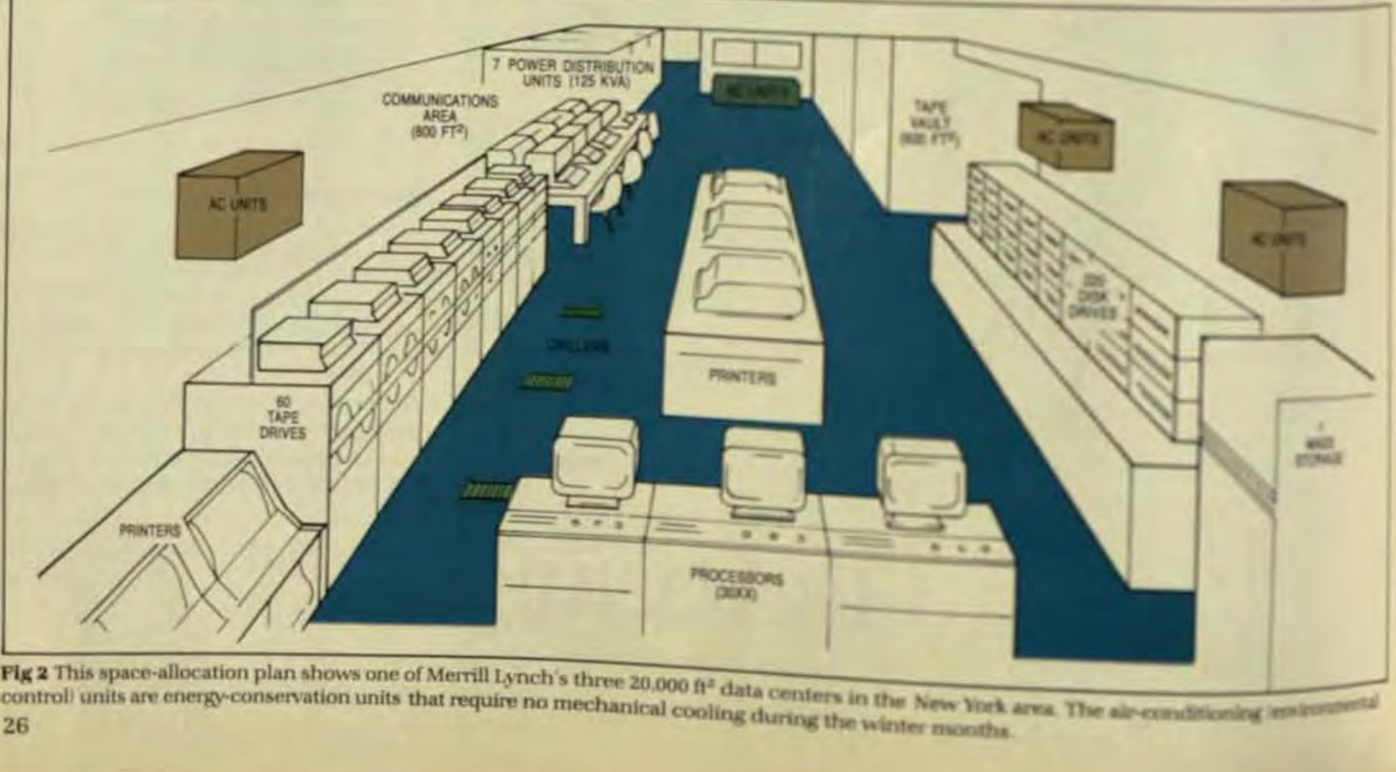

Fig 2 This space-allocation plan shows one of Merrill Lynch's three 20,000 ft² data centers in the New York area. The air-conditioning (environmental control) units are energy-conservation units that require no mechanical cooling during the winter months.

Second to None!

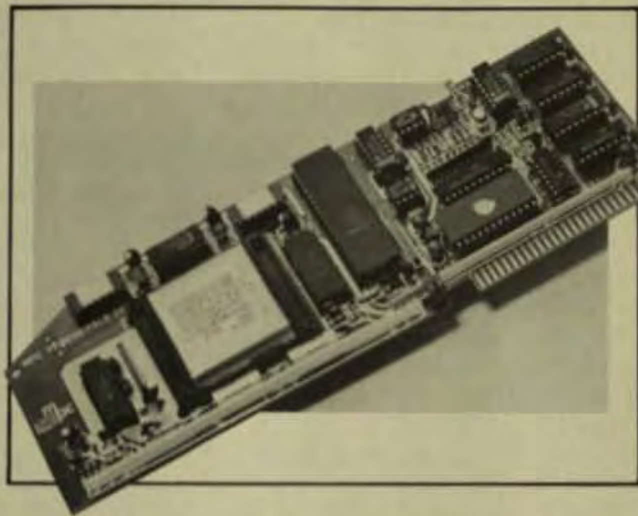

OMNI·BUFFER™
Universal Printer Buffer

Specifications:
- 64K Buffer
- Serial and parallel I/O
- Serial protocols RDY/BSY, X on/X off, ETX/ACK
- Switch selectable baud rates
- Full Centronics implementation
- Multiple copies
- Pause control
- Data compression

Omni-Buffer Supports:

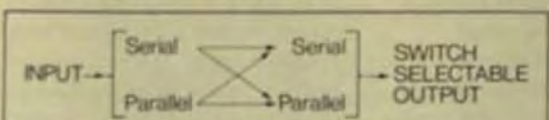

BUBDISK™ 128K
Bubble Memory for APPLE

Specifications:
- 128K non-volatile memory
- Portable
- Solid state Bubble Memory
- Not affected by heat, dust, grease, grit
- Immune to pollution and vibration
- Slot independent
- Three times faster and 1000 times more reliable than standard floppy disk

Omni-Buffer and Bubdisk are two unique products. Bubdisk, a 128K non-volatile bubble memory disk emulator, is available only from MPC. Omni-Buffer, a universal printer buffer, has features that far outpace the competition in a similar price range. Omni-Buffer and Bubdisk both carry MPC's TWO YEAR "NO HASSLE" WARRANTY. If you even "think" something is wrong with an MPC product, we'll replace it. Yes, Omni-Buffer and Bubdisk are indeed "SECOND TO NONE".

mpc

MPC Peripherals Corp.
9424 Chesapeake Drive
San Diego, CA 92123
619-278-0630
1-800-321-1028
TWX 910-335-1177

Chip-Design Systems Let Integrators Participate In Producing Custom ICs

Interactive design systems reduce the time from chip specification to prototype and increase the likelihood that the logic will work the first time around.

by **Graham Shenton**, International Microelectronic Products

Today the systems integrator has several choices about how to build his product: standard logic devices, gate arrays (semi-custom chips), and custom chips. In the past, the gap between standard logic, gate arrays, and hand-drawn custom chips was vast, resulting in high design costs and long turn-around times for full custom parts. Only high-volume manufacturers could consider full custom.

That condition no longer exists. New design methods integrate all phases of the design process from block-diagram input to full hand-drawn custom chips, with systems designers actively participating in the chip-design process. Modern design systems reduce the time from chip specification to prototype chips to as little as 14 weeks. Design time, and thus design cost, is significantly reduced.

In addition to lowering design cost, such design systems significantly increase the probability of getting a working chip the first time around because extensive chip simulation is an important part of modern chip design. Today's simulation programs are generally superior to TTL breadboard implementations of a circuit in predicting and evaluating actual performance because they model the performance of the custom chip rather than the performance of the circuit, as implemented in standard logic.

Systems integrators will never sacrifice either system functions or performance in order to use custom chips. Custom-chip functions may include logic elements as simple as a 2-input NAND gate or as complex as a microprocessor. Digital and analog circuits may be included on the same chip. And 3μ CMOS is rapidly becoming the fabrication standard for custom chips, although NMOS is readily available.

This article begins with a brief discussion of the alternatives available to the systems integrator: standard logic, gate arrays, and custom. It continues with a more detailed discussion of the current status of custom chips, how they are designed, how they are produced, their capabilities and limitations, and guidelines to their applicability. It ends with a discussion of computer-aided design (CAD) systems, their current state and future outlook, and one particular system is examined as an example.

STANDARD LOGIC

The semiconductor giants of the world produce an amazing array of logic devices for the systems designer to choose from. With the shrinking of geometries and the packing of more functions onto a single chip, systems that once occupied several cabinets can now reside on a single PC board.

The primary advantage of standard logic is the ability to design and implement a system quickly and relatively inexpensively. Parts are available from stock. There is no front-end chip design to boost total system design time and cost, and design changes can be implemented relatively easily.

Disadvantages, however, are many. Systems implemented with standard logic will always require more devices than the same system implemented with gate arrays or custom chips. More chips mean more board space, more interconnects, more possible failure points, higher power consumption, and significantly higher cost at volume production levels. In addition, reverse engineering of such systems is relatively easy, making it possible for competitors to analyze and reproduce designs.

Of course there are places where standard logic is the only reasonable way to go. The most obvious are low-volume applications where size, power consumption, and the possibility of reverse engineering are less important than front-end chip-design time and chip-design costs. Also, systems that are subject to engineering changes in their early stages are usually implemented in standard logic until the design is solid enough for customization.

GATE ARRAYS

Gate arrays, or semicustom chips, offer an attractive alternative to standard logic devices as a way to pull a variety of functions into a single chip. They provide a middle ground between standard logic and custom chips, in that a master slice is common to a large number of applications, with customization at the interconnect level.

Logic designers generally find gate arrays to be a friendly way to implement random logic within a system. They work with standard logic elements, interconnecting them just as they do standard logic chips. The main difference is that the interconnects are made at the chip level rather than at the board level.

Gate arrays offer a convenient way to collect random logic into a single chip, significantly reducing circuit-board space and power consumption. Design times are relatively short—typically 12 to 16 weeks from design input to prototype chips for moderately complex (3000- to 4000-gate chips). They are frequently cost-effective in volumes as low as 1000 chips.

Manufacturers offer gate arrays in a variety of configurations, typically containing from 300 to 5000 gates on a single chip. It is virtually guaranteed, however, that a particular design will not use all of the available gates.

Consider the case in which a design requires 525 gates, but gate arrays are available in 500 and 1000 gate sizes. What will the designer do? He can squeeze out some functions to make his design fit the smaller chip; he can design in more functions to make it fit the larger chip; or he can split his design by putting some functions in external chips. Regardless of the choice made, some gates will not be used in the gate array, resulting in wasted silicon on the chip. Add to this the need for interconnect channels, and it's clear why a gate array will invariably be larger than an equivalent custom chip.

The major advantages of gate arrays when compared to standard logic are lower production costs in moderately low volumes because of lower component costs and less PC-board space, lower power consumption, fewer interconnects, and design security. The major advantage relative to custom is that gate arrays are cost-effective in smaller volumes than full custom chips. (However, the number frequently quoted as the crossover point—100,000 chips—is no longer valid, as shall be seen later.)

Compared to full custom chips, the major disadvantages of gate arrays are: a larger chip size that results in higher per-chip production costs; some lack of flexibility due to the limited types of devices available on a particular chip; and the fixed chip configurations, which lead to wasted gates for particular applications.

FULL CUSTOM

The distinction between custom chips and gate arrays is clear: gate arrays are customized at the interconnect level only, whereas custom chips are customized at all mask levels. While gate arrays invariably will have wasted elements on the finished chip, a custom chip consists entirely of functional components.

Until recently, custom chips were designed and produced in much the same manner as standard logic devices. An experienced chip designer would manually draw the pattern of each mask layer on Mylar. The hand-drawn patterns were then digitized, stored in a computer, and used as a database for producing a pattern-generation tape. The tape was then used to generate production masks.

This procedure is extremely time-consuming, making lead times from chip specification to prototype chips very long—frequently a year or more. Further, it's reliably estimated that there are no more than 2500 experienced chip designers worldwide, and their employers prefer to use their talents to develop high-volume standard logic devices rather than relatively low-volume custom chips.

If the custom-chip industry was to survive in competition with gate arrays and VLSI standard logic, design times had to be shortened and more engineers capable of developing custom chips had to be found. CAD seemed to be the right answer to shorten the design cycle, and the pool of over 200,000 systems designers seemed to be the best source for additional chip-design talent.

THE FIRST STEP: CELL LIBRARIES

Cell-library-based systems were the first step in meeting these goals. Computerized development systems produce fully automated chip designs by using predefined elements contained in their standard library of logic elements. Library elements are individually hand-drawn by an experienced chip designer, digitized, and stored in the CAD system's database.

Library elements range from simple 2-input NAND gates to a variety of flip-flops and registers. Libraries containing programmable-logic arrays (PLAs), microprocessors, and analog elements will be available soon. Typical libraries contain elements equivalent to those offered in standard chips so that designers can implement their systems directly.

Design times rival gate arrays from design input to prototype chips, but differ from gate arrays in several important factors. First, there are no wasted elements in a chip that's designed from a cell library. The designer includes only those elements he needs, rather than interconnect prefacricated elements, so the resulting chip will invariably be smaller than an equivalent gate array.

Conversely, he loses some economy of scale because his chip is customized at all mask levels, and the result is higher volume requirements for cost-effective production. As a general rule, however, the smaller chip size will offset the cost for custom fabrication and result in a lower final per-chip cost in moderate volumes.

Although the designer realizes significantly greater flexibility and

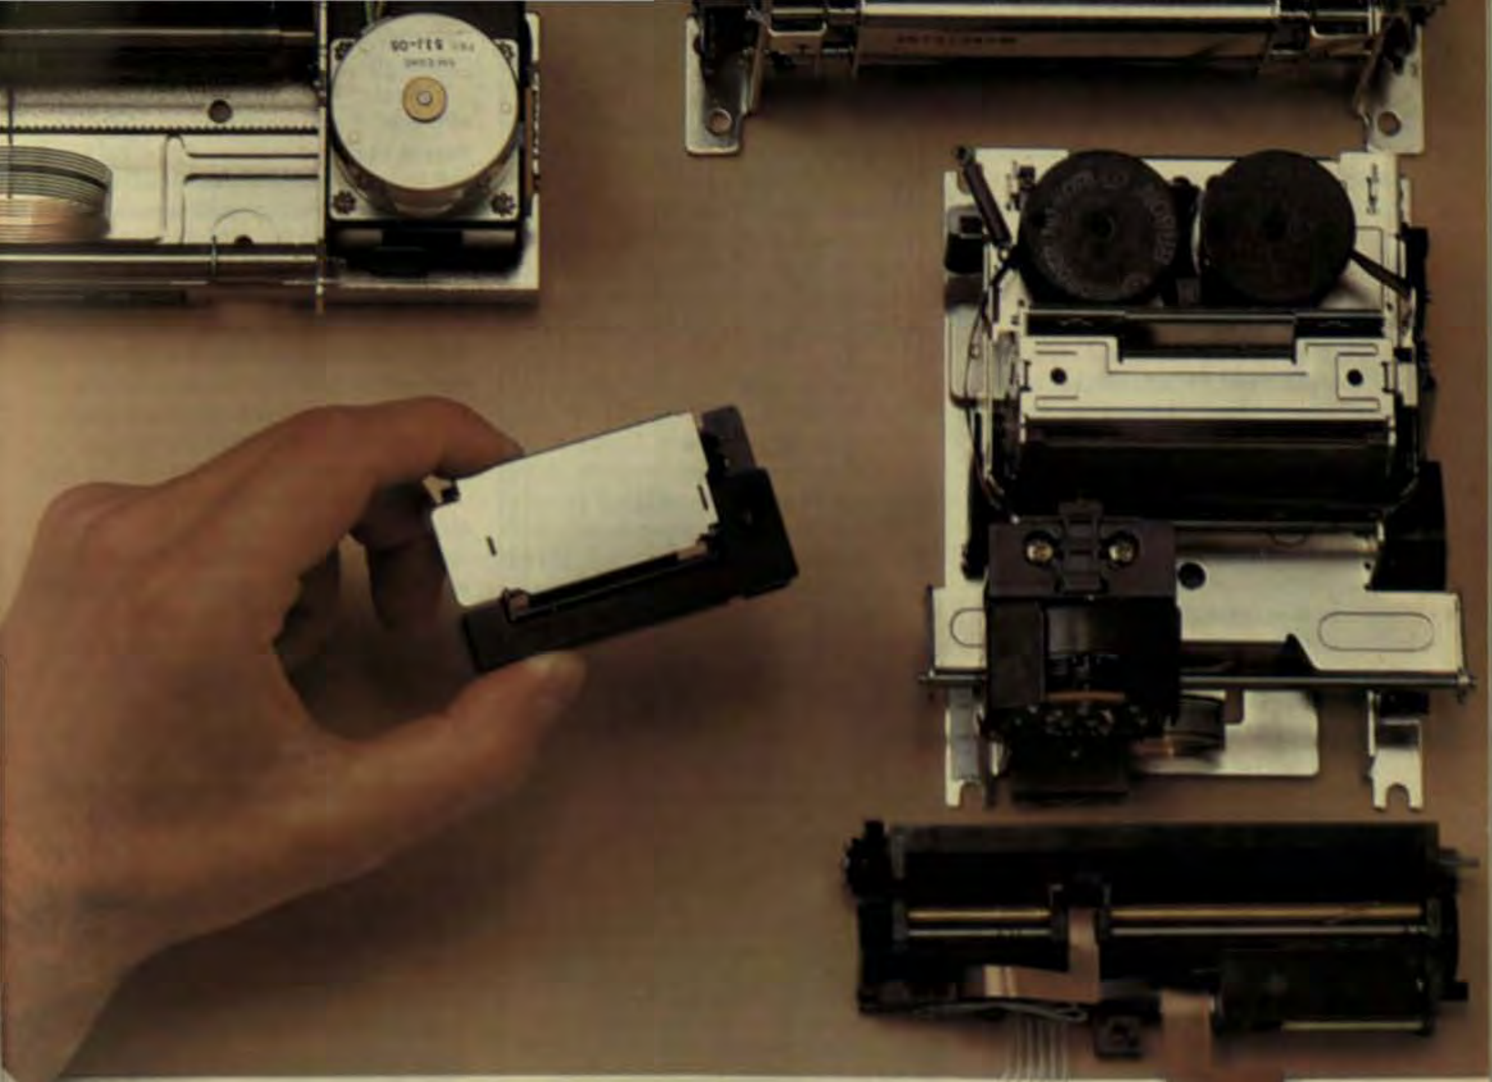

EPSON PRINTER MECHANISMS: THEY COME WITH A HELPING HAND.

PORTABLE/HAND-HELD COMPUTER PRINTERS				
Model	Head	Columns	Speed	Weight
M1XX	Impact Shuttle	16 to 40	0.4 to 0.7 LPS	2.1 oz
M2XX	7 wire impact	21 to 31	2.4 LPS	28.0 oz
M5XX	7 wire impact	40	3.0 LPS	59.0 oz
M12XX	Thermal	40	0.5 LPS	5.2 oz
M3XXX	9 wire impact	80	80 CPS	4.4 lb

You've heard that Epson delivers more printer mechanisms than anyone else. That's true. More than 600,000 shipped per month.

You've heard that our out-of-box failure rate is measured in 10ths of a percent. True again.

But maybe you're concerned that the leader in ECR and POS printers is too busy to worry about your application. *Not a chance!*

That's how we got where we are. By caring about your problems. Lending you a helping hand in selecting a printer for your application. Helping you package it as an integral part or as a stand-alone. And always delivering a cost-effective, quality solution. On schedule.

The printer mechanisms shown here can easily be designed into portable and hand-held computers. We have many more. In many variations. For many applications. Call or write us today. Let us start by giving you a helping hand.

SW Region (714) 751-1919 • **NW Region** (408) 985-8828 • **SE Region** (404) 458-9666
NE Region (617) 245-8007 • **CENTRAL Region** - (815) 455-2570

EPSON

EPSON AMERICA, INC.

OEM Products Division
Printer Group
3415 Kashiwa Street, Torrance,
CA 90505 (213) 534-0360
Telex: 182412

Six reasons why you should have an Array Processor from Floating Point Systems.

1. More computing power for the money.

At Floating Point Systems, we offer five array processors to meet a range of scientific and engineering applications. Attach one to your host computer and offload numerically-intensive calculations easily and cost-effectively... all for prices starting at $40,000 (U.S.) for 38-bit products. Our 64-bit systems start at less than $300,000 (U.S.)

2. Greater precision for greater accuracy.

For large, high-precision applications, the 64-bit FPS-164 offers 15 decimal digits of precision. Our 38-bit array processors provide up to 8 decimal digits of precision — 2 digits greater than 32-bit formats — ideal for signal, image and geophysical processing applications.

3. FORTRAN and more, for flexible programming.

For the FPS-164, our ANSI 77 FORTRAN Compiler generates code that optimizes use of the FPS-164's architecture. Our new Single Job Executive (SJE) supports complete job processing in the FPS-164. Our FPS-164 math library offers over 400 FORTRAN-callable subroutines for additional performance.

Our 38-bit array processors — the FPS-100, AP-120B, AP-180V and AP-190L — are supported by a comprehensive math library that offers over 450 FORTRAN-callable subroutines for signal, image, geophysical processing and other applications.

4. Solutions in minutes instead of hours, hours instead of days.

Because of their unique parallel pipelined architecture, our array processors provide high-speed computational throughput — up to 12-million floating point operations per second, assuring faster project turnaround time.

5. Large main data memory for bigger calculations.

The FPS-164 offers up to 58 Megabytes of main data memory (directly addressable), with a Disk Subsystem for up to 3,000 Megabytes of storage.

Our 38-bit array processors also have large main data memory — these range from a maximum of 64K words for the FPS-100 to a maximum of 448K words for AP-190L. You can add on 80- or 300-Megabyte disk storage systems, and a General Purpose Intelligent I/O Processor to control A/D and D/A equipment for real-time applications.

6. Superior reliability and worldwide support.

Our array processors have established impressive records for reliability and maintainability, and provide our customers with long-running, dependable operation. The FPS-164 features error-correcting memory, internal diagnostic system with a diagnostic microprocessor.

Behind this reliabilty stands our strong support: service facilities at key locations throughout the world, remote diagnostics, documentation, installation, training and more.

For further information, write today or call our nearest Sales Office ... or our toll free number, (800) 547-1445.

The world leader in array processors.

FLOATING POINT SYSTEMS, INC.

P.O. Box 23489
Portland, OR 97223
(503) 641-3151
TLX: 360470 FLOATPOIN BEAV

FPS Sales and Service Worldwide.
U.S.: Albuquerque (NM), Atlanta (GA), Dedham (MA), Denver (CO), Hartford (CT), Houston (TX), Laguna Hills (CA), Los Angeles (CA), New Orleans (LA), New York (NY), Orlando (FL), Palo Alto (CA), Philadelphia (PA), Rockville (MD), Schaumburg (IL), Seattle (WA). INTERNATIONAL: Canada — Calgary, Montreal, Ottawa; England, Bracknell, Berkshire; France, Rungis; Japan, Tokyo; Netherlands, Gouda; West Germany, Haar

DISTRIBUTORS: Australia and New Zealand, Milsons Point, N.S.W. (Techway PTY. LTD.); Austria, Vienna (Elektronische Bauelemente Und Geräte); Finland, Helsinki (OY Emmett AB); India, Bombay (Hinditron Computers PVT. LTD.); Israel, Tel Aviv (Eastronics LTD.); Korea, Seoul (Korea Computer Center, Inc.); Singapore (Scientek Corporation); Southern Africa, Johannesburg (Anker Data Systems); Sweden and Norway, Vasholm (Ta Konsult AB); Taiwan and Hong Kong, Taipei (Scientek Corporation)

© Floating Point Systems, Inc. 1982

Performance runs in our family.

At WICAT, our primary objective is performance. This attitude has created today's most comprehensive line of supermicro computers. In terms of breadth of product, software support, and perhaps most importantly, reliability, WICAT is a performance company.

WICAT has created a family of supermicro computers versatile enough to meet virtually any computing requirement. Our systems support from 1 to 48 users, 256K to 14 MB of main memory and from 10 MB to 474 MB of disk storage. So from the fast single user desk top S-150 to the powerful rack mounted S-220, WICAT provides the strength, speed, and flexibility to perform according to your company's unique requirements.

In the course of just two years, we have been successful in providing WICAT's customers with two operating systems (MCS and UNIX), 9 programming languages, as well as word processing, data base management systems, a virtual electronic spreadsheet, and the WISE authoring system. This is performance. Our software designers have concentrated on producing tools which formerly existed only in the mainframe or minicomputer environment.

This performance story began with the introduction of the first 68000 based microcomputer and now includes the installation and support of thousands of systems worldwide. Consider this a personal invitation to call us and get more information about the WICAT family.

WICATsystems

P.O. Box 539 · 1875 South State · Orem, UT 84057 · 801-224-6400 · Ext. 241

Cache Memory Increases Instruction Execution Speed for Superminis

Despite its high cost, a cache memory that uses separate instruction and data banks can improve system speed enough to make its price worthwhile.

by **Robert L. Hawk**,
Gould Inc.,
Computer Systems Division

Cache memory—a high-speed buffer memory physically located in the central-processing unit and functionally located between the CPU and main memory—is one of the major contributions to high-performance superminis. Because the addresses of the instructions and data to be used next by the CPU are generally the ones located near addresses already in use, obvious performance gains are possible when these items are kept in a cache rather than accessed from main memory with each fetch. Generally, the performance improvement more than justifies the relatively high cost of such a system.

A short, tight loop that reuses the same instructions—and perhaps even the same data—is a typical example of this type of program structure. The addition of cache memory to a CPU thus reduces execution time by cutting the time needed to fetch instructions and operands. In many large computers, main-memory access time is in the 600-ns range, while cache-memory cycle time is now commonly below 100 ns.

The increases in absolute execution speed made possible by cache memory were demonstrated in tests carried out at the Fort Lauderdale headquarters of Gould Inc., Computer Systems Division. The tests were conducted to document the performance of a new superminicomputer, the Concept 32/67, which executes at the rate of more than 1.6 (MIPS) (million instructions per second).

These tests measured the execution speed of single-precision Whetstone instructions with a single CPU. The speed of the computer without cache memory was used as the baseline measurement. When various parts of the cache memory were added to the CPU through privileged software instruction execution, the following increases in execution speed over the baseline measurement were obtained:

- CPU with no cache 100%
- Addition of operand cache bank only 120%
- Addition of instruction cache bank only 170%
- Addition of operand and instruction cache banks 230%

HOW CACHE MEMORY WORKS

A superminicomputer's memory system resembles a pyramid in three areas: size, access time, and price/-performance (Fig 1). At the very top of the pyramid is 32-Kbyte dual-bank memory. This unique design is a 2-way set-associative cache. Below that in the memory pyramid is a possible 16 Mbytes of MOS memory with 600-ns cycle time.

At the next and broader level of the memory hierarchy come moving-head hard disks and their faster cousins, fixed-head hard disks. At the bottom of the hierarchial memory structure is the virtually unlimited broad base of rotating or mass memory made up of magnetic tape, floppy disks, and streaming tape, all of which are relatively inexpensive—and extremely slow.

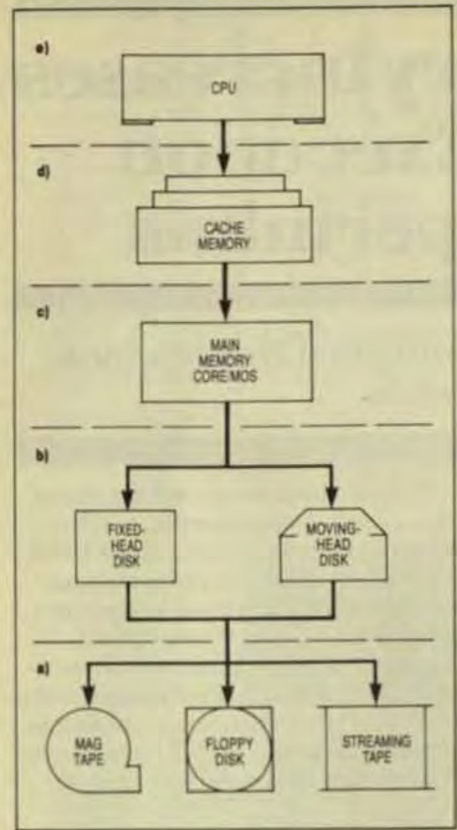

Fig 1 In a hierarchical memory system, the first level (**a**) consists of mass-storage devices with low costs and low speeds. Faster access times are available at levels 2, 3, and 4 (**b,c,d**), but at increasing costs per bit. With cache memory, an efficient CPU (**e**) can access data at high speed while retaining the low cost per bit of mass-memory devices.

While it would be prohibitive in terms of cost and physical space to provide an entire memory made up of 45-ns cache memory, a high-performance computer deals with its memory structure in such a way that the entire memory system approaches the speed of the cache memory at the apex of the hierarchy triangle. In the example system, the CPU first looks at the cache memory when it needs to fetch an instruction or operand.

If the instruction or operand is in the cache memory, the result can be returned in the same cycle. While the processor is executing a program, accessing data from its high-speed cache memory, other pieces of the program are moving from the rotating memory to the dynamic MOS RAM, and are even being prefetched into the cache.

FEATURES OF CACHE MEMORY

What happens in the cache memory, and how it helps provide the price/performance of this machine can be understood best if you look in detail at some of the features of cache memory:

- **Instruction and data banks.** The cache memory is split into two independent 16-Kbyte banks, an instruction cache, and an operand or data cache. Because the instructions are almost always present in the instruction bank, there's an almost perfect hit rate that allows the CPU as a whole to process near its maximum speed.

An example of the dual cache memory's power might be a program that manipulates a large array. The routines of the program fit quite comfortably in the 16-Kbyte instruction cache bank, and the array data itself is constantly shuffled in and out of the other side of the cache memory, yet the program functions at a 100% hit rate in the instruction bank. In contrast, virtually all other cache-memory systems mix instructions and data in a monolithic cache memory, so that instructions often are purged unnecessarily when new data is added.

- **Two-way set-associative operation.** In the two-way set-associative approach, both the instruction and operand caches are further divided into two equal banks (Fig 2). Each cache bank is provided with a RAM index array that's addressed by the least-significant address bits of the larger memory array.

As an instruction or operand is loaded from main memory into the cache memory, the most-significant address bits are stored in one of the two index arrays at the same location as addressed by the least-significant bits. For each index location, the index array also carries a valid bit that is set when the upper address bits are loaded into the index RAM.

During read operations, these RAMs are addressed by the low-order bits of the memory address; its upper-order bits are compared to the contents of the index array. If the equal compare is obtained, the valid bit is set and a cache hit occurs, indicating that the addressed memory location is present in cache and that the addressed contents of the cache array can be gated to the CPU processing logic as though that data had come from main memory.

If neither cache bank indicates a hit, the cache control logic must fetch the data from main memory and store it in one of the cache banks. Since both index and cache banks are addressed by the same least-significant address bit, cache can hold two memory locations that have the same least-significant address bits—one location in each cache bank.

And since the index arrays of the two banks have different contents, when the processor attempts to access a third memory location with the same least-significant address

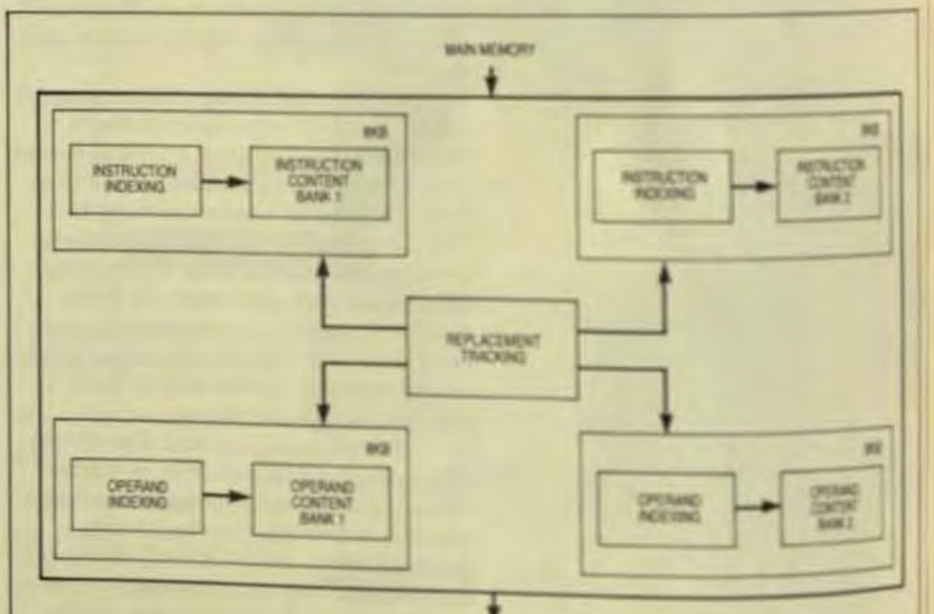

Fig 2 The 2-way set-associative cache memory with separate instruction and operand organization provides the key to cache hit rates in the mid-90% range and above.

A22 Minicomputer System

Like DEC's MICRO/PDP-11...only better!

Dataram's A22 Minicomputer System has been designed to offer features not available with DEC's popular MICRO/PDP-11.

Features like an 8" RX02 floppy, 1.0 MB quad memory, two fans in the card cage area, and space for a ¼" cartridge tape drive. Thirty amps of +5.0 volt power is provided. And the A22 is available in rack-mountable or tabletop versions.

Best of all, it's available now...only from Dataram.

- LSI-11/23
- 8" RX02 floppy
- 1.0 MB quad memory
- Two fans in card cage area
- 20 MB Winchester
- Supports RT-11, RSX-11M, RSX-11M-PLUS, UNIX, and TSX-PLUS
- Rack-mountable or tabletop
- Space for ¼" cartridge tape drive

DEC, MICRO/PDP-11, PDP, LSI-11, RSX, and RT-11 are trademarks of Digital Equipment Corporation
TSX-PLUS is a trademark of S&H computer systems, inc.
UNIX is a trademark of Bell Laboratories.

DATARAM

Dataram Corporation □ Princeton Road □ Cranbury, New Jersey 08512 □ (609) 799-0071 □ TWX: 510-685-2542

COMARK CORPORATION

DISKSTOR

The low cost OEM computer system that expands with you.

From the Busmasters

Featuring:
- 8 or 16 slots
- 5MHZ 8085 or 10MHZ 8086
- 8" Peripheral mounting
- 200/400 Power Supply
- Cast Aluminum Bezel
- Keylock Access
- Flexible I/O Backpanel
- 64K or 128KB RAM up to 1MB

Options:
- 18 or 36MB Winchesters
- 3M Tape Cartridge
- CP/M, CP/M86, MSDOS, MPM 86
- Custom I/O Panels
- Serial Communications
- Field Upgrade Kits

Prices start at $1,695 in Multibus*, $1,795 in Q-Bus**
Call for OEM Pricing and complete catalog.

COMARK CORPORATION
The BusMasters
93 West Street
Post Office Box 474
Medfield, MA 02052
(617) 359-8161

*Multibus is a trademark of Intel Corp.
**Q-Bus is a trademark of DEC

bits, a conflict exists. Then cache must decide which of the previous cache locations to replace with the third location. A cache-replacement algorithm makes this decision.

- **Lookahead, prefetching, and block moves.** A four-stage pipeline in the CPU separates the four main parts of instruction execution: fetch, decode, execute, and store. While instruction four is being fetched at the top level of the pipeline, instruction three is being decoded at the next level, instruction two is being executed, and the results of the first instruction are being stored.

The time for an individual instruction to move through the pipeline is 600 ns per instruction. However, once the pipe is primed by the fourth fetch, the process is completing an instruction at the rate of once every 150 ns—the true effective speed of the processor for most single instructions.

If an instruction to be fetched is not in cache memory, it can be fetched from main memory by the cache's own microengine. When the instruction returns from main memory it will immediately and simultaneously be inserted in a cache-memory location and in the pipeline, where it will be used by the execution unit. If a needed operand isn't in the cache memory, the cache's microengine will fetch it and place it in the operand register in the CPU and in the cache memory itself.

The internal program counter is always advanced to the next instruction beyond the one just loaded into the pipeline. Thus, even when the pipeline is full and can't accept the next instruction, the program counter's prefetch feature checks the cache for the next required instruction and initiates a fetch into cache if necessary.

Cache memory addresses are used to make intelligent decisions about data fetches from main memory. When the processor presents an address, it's checked to determine whether it's currently resident in cache memory. If there's a hit, the data is delivered to the processor; if there's a miss, the cache's intelligent control initiates access to main memory, bringing the referenced item into the data cache bank for use. At the same time, it brings into the cache more data than is immediately needed—data located adjacent to the requested data—on the lookahead theory that the processor is going to need that data soon.

- **Replacement algorithm.** Having the data present in the cache memory is the key to making that cache memory useful, so much study and refinement has gone into data-replacement algorithms. A hit rate (which can be defined as the percentage of times the needed operand or instruction is found in the cache memory) of 50% or less makes the cache virtually useless in terms of significantly enhancing the computer's speed. Having the needed data present in the cache memory is a direct function of the replacement algorithm. Four such general algorithms are used:

Random replacement. Easy to implement in hardware, very fast, and generally inefficient, this algorithm discards contents from cache memory to make room for new data with no logical decision about the near-term usefulness of the data being flushed.

First-in, first-out (FIFO). This common algorithm is particularly useful in straight-line, structured programming where the CPU rarely returns to a previously used instruction or operand, but presents very real problems in programs that employ frequently called subroutines.

Least frequently used. This algorithm, which replaces data on the basis of its average access by the running program, is difficult to implement in hardware—especially with set sizes greater than two. It also presents significant problems in counting, recordkeeping, and general housekeeping.

Least recently used. This keeps track of addresses being used by the CPU. In the event of a miss in the cache memory, the least recently used data is flushed on the general principle that if the processor hasn't called on a location recently, it probably has no further need for the data stored there. While somewhat difficult to implement, this principle prevents thrashing of data between main memory and cache.

- **Write-through storage.** With a block size of one word, storage of results achieves maximum efficiency in minimum space. By comparison, a CPU operating with a 4-word block would waste three storage locations when attempting to store one word of the block because the entire 4-word block must be stored as a unit. To avoid the waste of three locations in the cache memory, systems using 4-word blocks usually store results only in main memory.

However, because of the system's one-word block size, results can be stored directly in cache memory without waste of precious and limited cache memory. Those results thus become immediately available in cache memory to the program in execution. The result is also written through and simultaneously stored in main memory to ensure that the main memory is always consistent with the cache memory—an important consideration if another device, such as an intelligent I/O processor, is accessing the main memory concurrently with the program executing in the CPU.

- **Board layout.** The cache memory physically resides as an integral part of the CPU—three printed-circuit boards in the Concept 32/67. One PCB is the microstore, which contains the microsequencer. The second PCB in the CPU handles instruction and execution, and contains those functions as well as the bit-slice microprocessors. The third board is the cache/system-bus unit. The cache memory itself comprises static RAMs with a 4K by 4-bit organization, and 45-ns address access.

PERFORMANCE MEASUREMENT

Hit ratios in monolithic cache memories across the superminicomputer industry range from 70 to 90% for instructions, and rarely higher than 80% for operands. While there is no recognized standard equipment or method for hit-rate measurement, there is a device, a "hit speedometer," that uses a hardware counter to note fetches and monitor a signal that indicates cache-memory hits. The device measures results for every 10 million fetches (1.5 to 4 s of run time, depending on the complexity of the instructions), then averages a large series of such groups of 10 million.

Because of the various refinements in the cache system described, its hit rate is routinely in the 93 to 97% range, depending on the program in execution, and can be a perfect 100% in the cache-memory instruction bank. That represents a speed improvement of about four to one over main-memory addressing with the same program.

In terms of programming, cache memory is totally transparent to the programmer writing the application. In fact, cache memory's ability to maximize the execution of tight loops and structured iterative programs encourages good programming practice by rewarding it with maximum speed.

Cache memory is expensive, but a rough estimate suggests that about 2% of the system's price is because of implementation. The payback for that 2% increase is total system throughput of as much as 230% compared to a similar machine without cache, depending on the application, program, and system configuration. ■

Robert Hawk, *senior section manager for advanced development at Gould, is reponsible for system bus hardware development.*

COMPUTER VISION

. . . We made it affordable!

The
IVP-2
VIDEO
I/O SYSTEM

Digitizes and stores video images in real time · 480 x 512 resolution · 64 gray levels · RS-232 and parallel interface to any computer · Input from any TV camera or video recorder Composite video output of reconstructed digital image Video memory may be down-loaded from host computer Supports computer generated graphics overlays on real time video images · User programmable, microprocessor controlled for maximum flexibility

FOXVILLE
COMMUNICATIONS
CORPORATION

7741 E. Gray Road
Suite #17
Post Office Box 5419
Scottsdale, AZ 85261
(602) 948-9817
Telex 165750

In Europe, contact

FOXVILLE
EUROPE

Hundelgemsesteenweg 105
9220 Merelbeke
Belgium
Telex 12901

Analogic's AP500, the Emancipated Array Processor

Frees You and Your Host for Greater Productivity

The Emancipation

Analogic's AP500 32-bit Floating Point Array Processor is the first truly independent array processor. It offers the performance and throughput capabilities needed to meet most time-critical, complex, and I/O-bound applications.

With little or no Host involvement, the state-of-the-art Control Processor based on the 32/16-bit MC68000 supermicro and ultra-fast 40-bit Arithmetic Pipeline deliver the power and flexibility to solve most applications in real-time.

The AP500 has unprecedented I/O capabilities for COMPLETE system expansion. Now free the Host from burdensome I/O with the IEEE-796 standard Multibus, the RS-232 Serial Line Interface, and the two 6.25MHz Auxiliary I/O Ports. Only Analogic provides the necessary paths for array processor system growth.

OEM and End-User Benefits

Packaged to Fit—Contained in a 5.25" rack-mountable chassis or available as a card set with backplane, this AP500 is compact/rugged power for the systems integrator.

Easy-to-Use—Programmable in high-level or assembly language, AP500 Software is modularly designed to get most applications off the ground quickly and painlessly.

Tools to Optimize—AP500 Software provides many simple programming techniques for execution-time optimization . . . customized to meet your application demands.

Priced to Please—For comparable performance, you will have to pay from two to ten times as much for other array processors or minicomputers.

Complete System Flexibility

The AP500 supports virtually [illegible] term configuration. You can build [illegible] system from a wide range of Multibus compatible peripherals including [illegible] tapes, A/Ds, D/As, display [illegible] CPUs and LAN controllers, or [illegible] data acquisition and data display [illegible] to the high-speed Auxiliary I/O Ports.

Write for comprehensive literature [illegible]

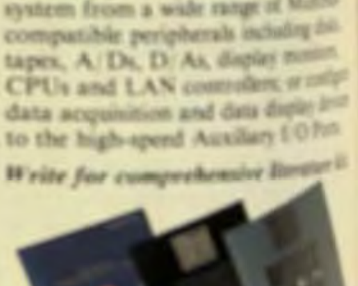

ANALOGIC
COMPUTER SYSTEMS DIVISION

ANALOGIC CORPORATION • Audubon Road • Wakefield, Massachusetts 01880
TEL: (617) 246-0300 • TLX U.S. & Canada 951-070, Int'l 681-7180

Array Processors Speed Up Complex Calculations At Low Cost

Array processors perform number crunching in scientific/ engineering applications and run at speeds up to a thousand times faster than the computer they support.

by **Ron Baldridge,** TRW LSI Products

Array processors are special-purpose peripheral processors that perform the high-speed, repetitive arithmetic operations that are often required in large scientific or engineering applications. Typically, they are 10 to 1000 times faster than the computer they support—and at substantially less cost than supercomputers.

As new, higher-speed semiconductor components are developed, both the speed and size of AP memories will increase, and their physical size and power consumption will be reduced. New parallel and pipeline architectures will further enhance the effectiveness of these units.

Today, APs are being used in the discovery of oil through seismic data processing, the enhancement of satellite photographs through image processing, the identification of manufacturing flaws in piece parts through video-camera pattern recognition, and the discovery of brain tumors through X-ray scan processing.

Solving these problems often requires special mathematical algorithms such as the fast-Fourier transform (FFT) or matrix convolution and inversion to solve simultaneous linear differential equations over 2- or 3-dimensional space. Such problems can usually be broken down into long sequences of repetitive multiplications and additions of arrays of numbers—hence the term array processor (AP).

The primary function of the array processor in these applications is to perform this type of numerical computation in the fastest time possible and at the most reasonable cost. This is generally accomplished through the use of fast semiconductor components, and an architecture that uses special techniques such as parallel circuits and pipelining to perform high-speed math operations.

Array processors aren't meant to be used for logic operations or string manipulations, but are at their best when performing the high-speed multiplication and addition of arrays of numbers, using a special mathematical representation known as floating point.

USE OF FLOATING POINT FORMAT

Scientific and engineering problems often deal with very large or very small numbers. For example, the speed of light is about 186,000 miles per second, and the line width of an interconnect on an integrated circuit can be less than 1 μ.

To represent such a wide range of numbers in a compact way, a special notation known as floating point was devised. A floating point number is simply a shorthand way of representing a number without writing all the digits. For example, 186,000 may be written .186 x 10^6. If base 10 is assumed, this number can be represented by three quantities: a sign (plus or minus), a signed exponent, and a significand or mantissa. For this example: SIGN = +
EXPONENT = +6
MANTISSA = .186

Since a number can be written with the decimal point in different places just by a change in the exponent, the point is said to float. For exam-

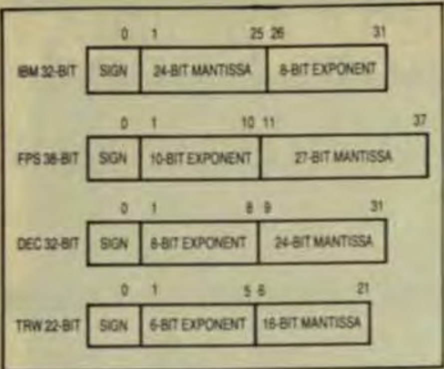

Fig 1 Floating point formats provide a variety of range and precision. For example, Floating Point Systems uses 38 binary bits to get an extended range of over $10^{\pm 153}$ and a precision of eight significant decimal digits. In contrast, the TRW short format is designed for high-speed digital signal-processing applications, where a range of $10^{\pm 10}$ and a precision of five significant digits is often adequate.

ple: $.186 \times 10^6 = 1.86 \times 10^4$. To maximize precision in long strings of computations, floating point numbers are often normalized—written with a zero to the left of the point and a nonzero digit immediately to the right of the point, as in 0.186×10^6.

Floating point notation is a compact way to represent very large or very small numbers in the array processor's memory. The size of the exponent determines the range of numbers represented, while the mantissa size determines the number of digits of precision that can be maintained. Fig 1 summarizes several popular floating point formats in use.

The speed with which an array processor computes is usually measured by millions of floating point operations per second (MFLOPS). In a single cycle, the typical AP can perform one floating point multiply, two floating point adds, and three memory accesses. This permits speeds in today's APs in excess of 24 MFLOPS.

By contrast, a VAX-11/780 general-purpose machine computes at a rate of 0.5 MFLOPS. An AP's achieveable speed, however, can be considerably less than the theoretical speed due to a variety of overhead factors between the host and the AP, as well as possible under-utilization of the AP's full parallel-computing capability. The latter could well depend on the match between the AP architecture and the structure of the problem solution.

CAT-SCAN FOR SHARP IMAGE

Computerized axial tomography (CAT) scanners were developed in the early 1970s to provide a thin cross-sectional X-ray picture or slice through a patient's body as seen from above (Fig 2). Sharp, high-quality images can be constructed from the absorption data received when an X-ray beam is rotated 360° around the patient. A set of simultaneous linear equations must be solved by the AP to reconstruct the image. In addition, the image often needs to be enhanced by a filtering out of false structures, such as shadows, introduced by the hardware. Current CAT-scan systems have scan times of 2 to 5 s, can reconstruct an image in about 30 s, and can precisely locate structures to within 1 mm. These systems are being used to identify and diagnose such things as tumors in the head or to locate structures precisely so that surgical tubes can be placed for drainage.

Faster scan times (less than 1 s) will be needed in the future to avoid the need to sedate fidgety patients, to produce 3-D images, and to visualize moving organs such as the heart. These systems will also benefit from faster APs to process the larger amounts of data required to produce these images in close to real time. The array processor is therefore integral to the development of new equipment for medical diagnostics and treatment.

TYPES OF ARRAY PROCESSORS

The most costly, but by far the highest-performance method of processing arrays of numbers is to employ a supercomputer such as Control Data Corp.'s Cyber-205 or Cray Corp.'s CRAY-1. These machines can cost up to $8 million, but often operate in the 50-to 100-MFLOP range. They may be used to solve very large problems, such as weather forecasting or animation sequences in motion pictures. In these units, the AP functions required to process arrays or vectors of numbers are built into the architecture, thereby eliminating the host/AP overhead.

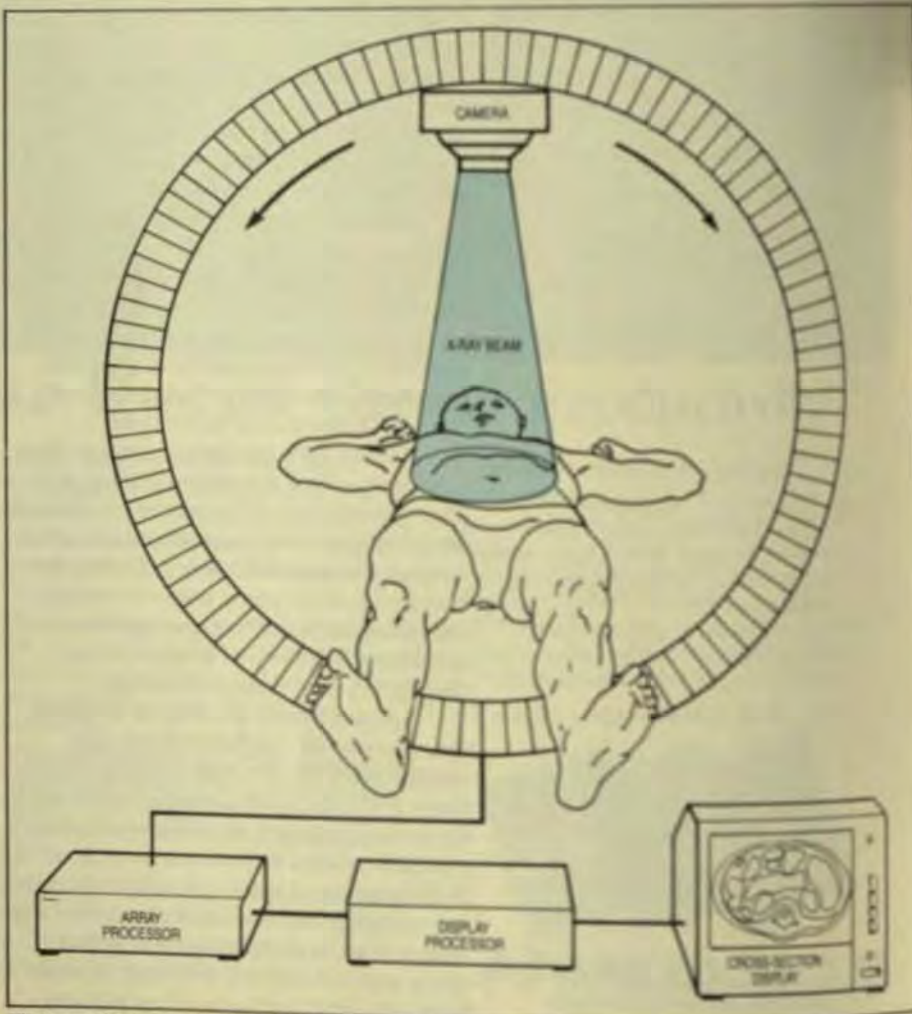

Fig 2 A CAT-scanner rotates an X-ray beam 360° around a patient to produce a thin cross-sectional slice of the patient's body. The AP must rapidly solve a series of simultaneous equations as well as digitally filter out aberrations in the image before it is displayed.

Such supercomputers use some of the fastest components available, typically emitter-coupled logic (ECL), and often require special cooling. Other types of built-in APs include the hostless AP, which has a built-in control processor that provides very high performance at considerably less than supercomputer cost. The most common type of AP is attached—a separate peripheral processor, usually one designed to communicate with a variety of host through an I/O channel. An array processor may also be an add-in—typically a low-cost, board-level product that plugs into the backplane of the host. In some configurations the host and AP may actually share a common memory through a remote-memory interface. The attached AP may cost in the $25,000 to $100,000 range, and has typical performance of 10 to 20 MFLOPS.

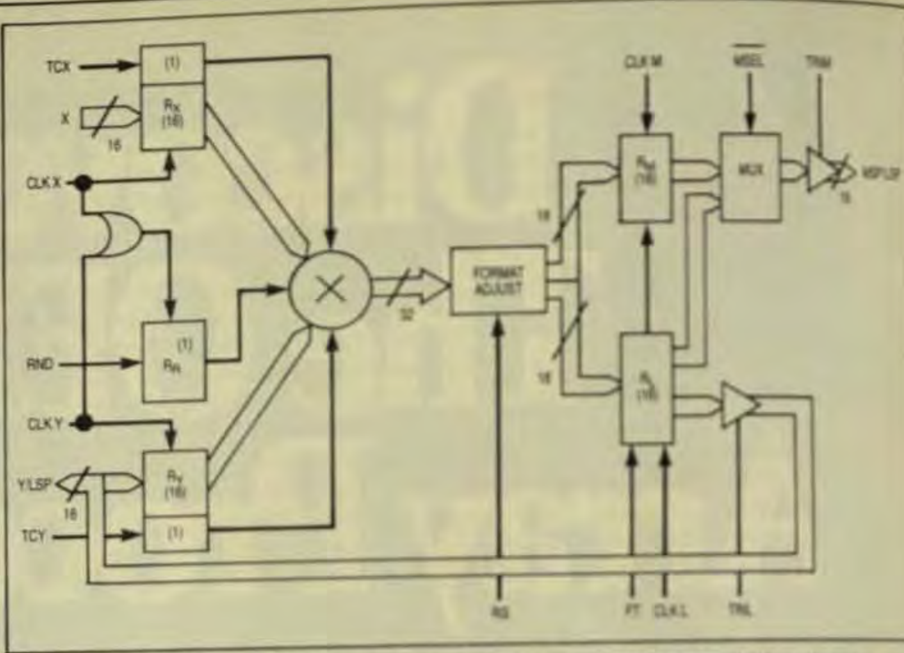

Fig 3 An example of a new high-speed components is the TRW MPY016K, which performs a 16x16-bit multiply in less than 50 ns, compared to 165 ns in the previous generation of devices and at the same power consumption and cost

TRENDS IN TECHNOLOGY

Array processors continue to increase in performance while lowering or maintaining their price. This situation is made possible through the development of new, higher-speed semiconductor components. For example, high-density, low-cost 64-Kbit MOS RAM memories will become available at access speeds under 100 ns, perhaps approaching 50 ns in the next two years.

Already available are high-density gate arrays like the Fujitsu VH series, which exceed 10,000 gates implemented in low-power CMOS and run at very high speeds. Advances in such design tools as simulation and layout will permit rapid turnaround times on complex circuits. And since the highly repetitive and parallel structure of some AP architectures lends itself well to gate-array implementation, benefits will include reduced circuit count, power, and cost.

The commercial multiplier IC is another type of high-speed component just beginning to be released. The TRW MPY016K (Fig 3), for example, is capable of performing a full 16x16-bit binary multiply in less than 45 ns—a significant increase in speed from today's standard 145-ns multiplier. Next-generation APs appearing in 1984 should see at least a 3-fold increase in computational speed by using these devices.

Several semiconductor manufacturers are also working on high-speed floating point ICs that will use a single IC to do the total floating point multiply and addition in less than 100 ns. The first such device to appear is the TRW TDC1022—a 22-bit arithmetic unit. These ICs will permit very powerful APs to be built in compact packages that resemble today's small minicomputers. In fact, powerful APs are likely to be integral in the design of next-generation engineering workstations.

TYPES OF ARCHITECTURE

Powerful new 16-bit microcomputers such as the Motorola 68000 will be built into the AP so it can handle its own peripherals separate from the host, as well as control intelligent I/O channels. For example, the new ST-100 from Star Technologies, which uses two 68000s, allows the sharing of APs by one host, the sharing of hosts by one AP, or even a hostless AP. Microprocessors may also be dedicated to handling other tasks, such as controlling the supervisory program or controlling a local-area-network interface.

As more and more functions move from the host to the AP it may make sense to integrate the host directly into the AP architecture. This, in effect, is analogous to the supercomputer approach on a smaller scale. New machines may use pipelining and other techniques employed by supercomputers, possibly even to the point of software compatibility. Thus, a small Cyber-205 could appear in an engineering workstation.

Since an AP processes data in parallel, it seems logical to assign a separate processor to each parallel data path. So far, such configurations of parallel processors have proven to be fraught with programming difficulties—such as how to synchronize the operations of the parallel processors.

However, research is being conducted on a new type of element known as a systolic processor. Arrays of such elements may be connected together so that data flows through the system much as blood is pumped through the body by the heart. These arrays permit very efficient implementations of such algorithms as the fast-Fourier transform used frequently in digital signal processing. ■

Ron Baldridge *is market development manager at TRW-LSI Products, and was previously strategic planning analyst at Mostek Corp. Ron has a BSEE from the Univ. of Florida and an MSEE from the Univ. of Tennessee.*

Boost Your Micro's Ego

This Single Board Array Processor Will Have Your Micro Thinking Like a Mainframe

It's fast enough for real-time analysis. Small enough to fit in a single slot. And flexible enough to perform almost any algorithm.

That's the APB-3000 array processor. From Marinco.

It's Multibus* compatible. With other bus versions soon to follow.

To the host, it simply looks like extra memory.

Data flows through dual ports and settles in a mapped memory that can represent integer, floating point and complex numbers. Yet both program memory and data memory are always accessible to the host.

The APB-3000 executes instructions in 100 nsecs. Which adds up to 10 million floating point operations per second. For instance, a 1024 point complex FFT is processed in 4.1 msecs. Start to finish. And extra boards can be linked for 20, 30, 50 MFLOPS or more.

Standard routines include FFTs, power spectral density, correlation, digital filtering and deconvolution.

Or create your own.

It's simple. The board is easy to talk to. Especially in FORTRAN.

Order the board with standard ROM. Or with RAM for development. If you wish, Marinco will back you up with extensive engineering and software support.

While the board may inflate your micro's ego, it won't inflate your budget. The cost: $4250. Complete. And quantity discounts are available.

So call us today.

But remember. When your micro suddenly begins acting like a mainframe, don't worry. A strong ego is healthy.

And besides—it's results that count.

MARINCO, INC.
A HELIONETICS COMPANY

11760 Sorrento Valley Road,
San Diego, CA 92121
(619) 453-5203
Telex 69-7901

*Trademark of Intel Corporation

THERE IS MORE THAN ONE GOOD REASON TO CALL HEURIKON

Make HEURIKON your source for Multibus microcomputers and sytem components.

Since 1972 Heurikon has been working hard to make life a little easier for those who "put it all together." Powerful Z-80™ and MC68000™ based Multibus™ microcomputers like MLZ-92A and HK68 have helped Heurikon customers find economical solutions for their problem applications.

MLZ-90A single board microcomputer with nine byte-wide memory sockets for use with RAM or ROM (AM9511 and floppy disk drive controller optional).

MLZ-91A single board CP/M™ system with on-card floppy disk drive controller, winchester interface, optional AM9511, streamer tape interface, two serial ports, one parallel port, 64K bytes RAM with parity, two EPROM sockets, and GPIB CONTROLLER.

MLZ-92A single board CP/M™ system with four serial ports on-card, floppy disk drive controller, winchester interface, optional AM9511, Centronics printer interface, 64K bytes of RAM with parity, and two EPROM sockets.

MLZ-93A single board CP/M™ system with 128K bytes of dual ported RAM, four EPROM sockets, floppy disk drive controller, optional AM9511 and powerful serial port features including SDLC and HDLC protocol support and modem controls.

HK68™ powerful and versatile single board UNIX™ or CP/M-68K™ system with MC68000 CPU, MMU, quad channel DMA, four serial ports, 128K or 256K bytes of on-board RAM with parity (expandable to 1M byte on-board!), and two iSBX™ connectors for I/O expansion.

MLZ-VDC intelligent 640 x 480 x 4 color graphic controller based on the NEC 7220 controller chip with on-board Z-80 CPU, DMA controller, and user definable FIFO interface to Multibus.™ Users may display up to 16 colors from a 4K palette. Up to 1024 x 1024 x 3 interlaced also available.

iSBX Expansion Models
Heurikon also stocks iSBX™ modules for I/O expansion on its HK68 and MLZ-VDC. These modules may be used on any other board complying with Intel iSBX specifications.

Presently available are the SBX-FDIO floppy disk controller module and the SBX-SIO serial port expansion module.

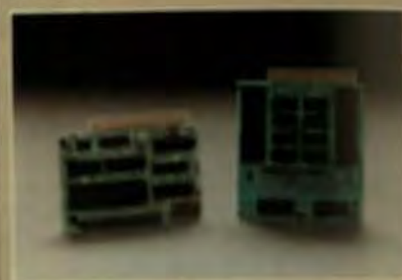

The SBX-FDIO can control up to four single or double sided, single or double density 5¼ inch or 8 inch floppy disk drives. A software toggle allows both 5¼ inch and 8 inch drives to be intermixed on the same cable!

The SBX-SIO is a quad serial port expansion module utilizing Zilog SCC controller chips with built in baud rate generators. The SBX-SIO is available in synchronous and asynchronous models.

Systems
Heurikon also supplies completely integrated systems with your choice of CP/M-80™, MP/M-80™, CP/M-68K™, and "Berkeley enhanced" Unix operating systems.

Systems are available with four, six, and fourteen slot enclosures with 5¼ inch or 8 inch dual floppy or floppy winchester drive combinations.

HEURIKON CORP.

3001 Latham Dr. Madison, WI 53713 Telex [illegible]

Call Heurikon Direct 1 800 356•9602 In Wisconsin 1 608 271•8700

Z-80 is a trademark of Zilog Corp. Multibus & iSBX are trademarks of Intel Corp. UNIX is a trademark of Bell Lab. CP/M-68K is a trademark of Digital Research. HK-68 is a trademark of Heurikon Corp.

Microprocessors Boost Performance Levels of Multiterminal Systems

A new system architecture that employs dedicated microprocessors supports multiterminal office systems without a significant loss of response time.

by **James Ferguson,**
Cado Systems Corp.

Switching to more powerful minis or mainframes to accommodate larger numbers of terminals loses the cost-effectiveness of microcomputer processing and means that the enormous standing investment in programming and stored records would have to be abandoned. What's needed is an entirely different architecture—one capable of supporting what is effectively an infinite number of terminals with a single bus that can link N processors to a common data base without limiting the type of processor that can be used.

An open-ended 64-terminal system has been designed to meet this challenge. Since computers degrade exponentially rather than linearly as terminals are added, significant performance improvements were also mandated. While this structure could have been developed conceptually by any computer manufacturer, its detailed implementation is unique. Because no company develops a product in a vacuum, there's a continuum of product development to take advantage of unique operating systems, software, and peripherals. In this case, there were four primary factors to consider:

- The architecture had to allow existing user software to be 100% upward-compatible to the new processor.
- The new system had to run the existing operating system, which employs a language-driven interpreter based on a proprietary version of BASIC.
- A high percentage of today's users employ both word-processing and management-inquiry software that require a disproportionate amount of memory.
- The system should retain the user-responsive speed of a 19.2-Kbaud terminal rather than revert to slower units. However, a number of traditional computer-design features could be ignored because the target was business-information processing, not telemetry processing, scientific research, or other specialized uses that would have required sophisticated number crunching and a considerable amount of resulting overhead.

BASIC DESIGN ISSUES

In the typical business environment, there are two primary reasons that systems degrade when terminals are added. The first is disk arm movement and the reading of data from the disk. To minimize this factor, a direct-access technique was developed to place the record key and data on the same track.

The second factor leading to system degradation is the time required to service each terminal—and the faster the terminal, the more serious the problem. For example, a 19.2-Kbaud terminal performing word processing on a byte-to-byte basis requires CPU service once every 500 μs. In a 50-terminal system, a different unit would have to be serviced every 10 μs. But since each interrupt requires 100 μs, such a large system is wholly impractical.

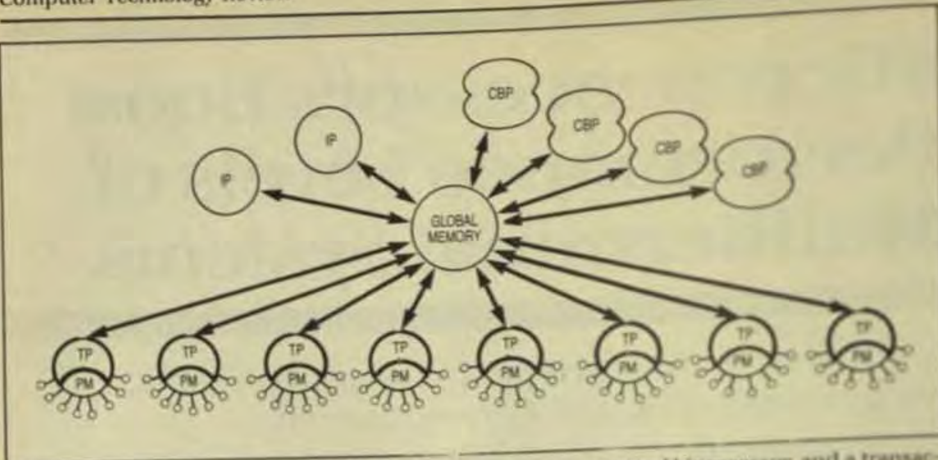

Fig 1 The basic Tiger architecture uses an intranet processor, a control biprocessor, and a transaction processor, which communicate through a shared global memory to handle as many as eight terminals. Each terminal has 8 Kbytes of private memory.

Fortunately, there are alternatives to traditional approaches. Rather than take the technical risks involved in the development of an ultra-high-speed central processor to perform all the processing, disk access, and interrupt service, consider sharing the processing load among multiple microprocessors, each dedicated to a single task or family of tasks. Such an architecture would provide a direct growth path for current microprocessor systems and allow direct transportability of existing microcomputer software.

SHARED WORK LOAD

Using this approach, Cado's Tiger architecture consists of a global memory and three types of standard microprocessor that share the total work load (Fig 1). Depending on the number of terminals to be supported, the system requires either one or two intranet processors (IPs); up to four 2-channel control biprocessors (CBPs); and one transaction processor (TP) for each eight terminals. Since each microprocessor is supplied with its own ROM and RAM memory, each can perform a substantial amount of processing without loading the memory bus.

The TP only performs tasks it can complete in less than 62.5 µs. Of necessity, these are quite simple tasks, such as "add two numbers," or "display a number"—but that covers a great many business-processing operations. Any task requiring a longer processing time is passed to one of the other microprocessors (the IP), so interrupts are not used as they would be in a traditional system architecture.

Working on a polling basis, and without interrupts, the IP handles the overload tasks in turn. When it finds a job allocated to it by one of the TPs, it takes as long as needed to complete that task, posts a message in that TP's mailbox, and moves on to the next operation.

For business processing, this approach is quite feasible, because less than 10% of the instructions involve the IP at all. Typically, the IP is only called upon to perform a few arithmetic functions such as long multiply and long divide, and to access the database.

The third processor in the system, the Intel 8089 CBP, uses one of its channels for the physical disk read and its other channel to perform short-duration jobs, such as "move a string" and horizontal queuing of disk-read requests. In this approach, the arm responds to each request to read data in its current direction of travel, reversing its motion only when no further data in that direction is requested.

All three processor types are involved in functions requiring a major disk read. A TP unable to perform such a large function itself posts the task in the IP's mailbox. After performing some calculations, the IP determines the needed track location and posts an arm-movement request in the mailbox of the CBP, then moves on to another function.

When the arm movement command is executed, the IP returns to the original chore and requests a disk transfer from the other channel of the CBP, then goes on to yet another job. The CBP transfers the requested track's data to active memory and alerts the original TP to proceed with processing the original task.

Throughout, each processor performs only its prescribed functions, interleaving them with other tasks for peak efficiency and an absolute minimum of overhead.

MEMORY ALLOCATION

To avoid using the system's memory bus more than absolutely necessary, each processor has its own 8 Kbytes of private ROM for program storage and its own private RAM for processing. The system also has a global memory shared by each of the processors. The size of the global memory varies with the size of the system, but each TP has at least 128 Kbytes available to it.

That amount of memory nominally allocates 16 Kbytes of RAM to each user, 8 Kbytes for the program being used and 8 Kbytes for a complete disk-track buffer. While every user is allocated a disk buffer that's uniquely his, it remains available to other users because it's located in global memory.

With this disk-buffer capacity available, it's feasible to access an entire track at a time, rather than read individual records or sectors. An entire 8192-byte track can be read with a single arm movement. For a 100-record track, that design eliminates almost 200 arm movements because the arm can start a read at any point on the track and continue until it reaches that same point on the next revolution.

In effect, each user in this configuration has access to up to 64 Kbytes of buffered memory, which reduces the number of disk accesses. Today's low memory costs make it practical to implement even greater amounts of memory with this design.

In a typical business application, this capability virtually eliminates disk accesses. It's rare that eight operators are performing eight difficult jobs that are totally unrelated to one another. Far more typical is a function such as order entry in which multiple operators work with two files—with no need to access the

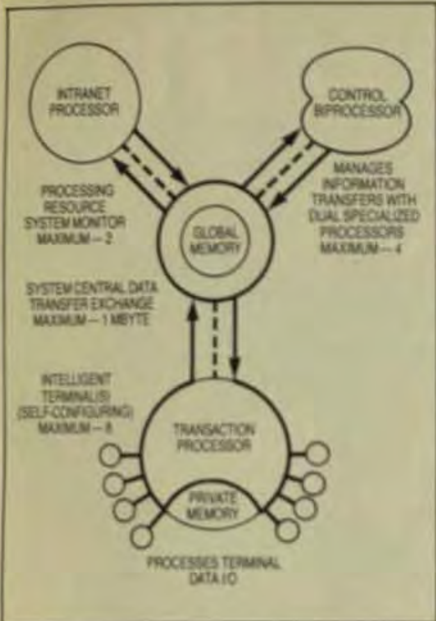

Fig 2 When expanded to handle 64 terminals, this design requires two IPs, four CBPs, and eight TPs.

manufacturing file, the payroll file, or any other unrelated file stored on the disk.

Hence, a crew performing order entry reads needed tracks into the disk buffer from the customer and parts files, leaves them there, and goes on to perform the order-entry task. As many additional tracks as needed can be read into other vacant track buffers for use by the operators until the job is complete or all track buffers are filled. In the latter case, the system determines the least-used track in the buffers, sends it back to the disk, and transfers a new track to the newly vacated buffer. That way, there's always a high probability that the needed data file is in active memory, where the crew can operate on it without waiting for a disk access.

EXPANSION METHOD

This architecture allocates a TP and 128 Kbytes of memory to each group of eight terminals. One CBP can support 16 terminals, and a single IP can support up to four clusters of eight terminals. As many as 64 terminals can be handled by the addition of more microprocessor support (Fig 2).

Conceptually, this architecture has universal application for business-computing systems. It relies on a software system with minimal overhead, and a family relationship between the microprocessors. Otherwise, however, the selection of individual system components is irrelevant. New microprocessor devices can be integrated as they are developed and tested.

This type of architecture is beyond the design stage at this point. It has gone through a complete test phase and demonstrated that it can support more than 50 terminals without significant degradation. While the system takes full advantage of its 16-bit bus for high throughput, its primary speed advantage results from the elimination of bottlenecks in disk access, interrupt processing, and bus transfers. ■

James Ferguson *is vice president of software development at Cado Systems Corp. A founder of the company, he developed the architecture and operating systems of its 8-bit microcomputer products.*

Home Sweet Home for Your Multibus*

Finally there's a system chassis that is designed and manufactured with the thoroughness and care you expect in your Multibus system. It's Electronic Solutions' new Multichassis™.

- **9 slots, 0.6" spacing—or 7 slots, 0.75" spacing**
- **Hefty 4-output 300W power supply— 40A at +5V**
- **Cool operation even with high-density boards**

There's a field-proven card cage and backplane, plus full RFI filtering, locking front panel function switch, power fail detection, and quiet dual cooling fans with quick-change filters.

And best of all, the removable front panel lets you easily customize the Multichassis to match your company color and logo.

So treat your Multibus system to an elegant but affordable new home—the Multichassis by Electronic Solutions. Call us today for full specifications and prices.

Electronic Solutions
5780 Chesapeake Ct.
P.O. Box 85244
San Diego, CA 92138
Call Toll Free
(800) 854-7086
In California (619) 292-0242
In Canada call:
Transduction Ltd.(416) 625-1907

*Multibus is a [illegible] Corp.

HOW DOES 1 = 50?

With SUPER STAR... The First S-100 Computer To Use a 5 MB Removable Cartridge.

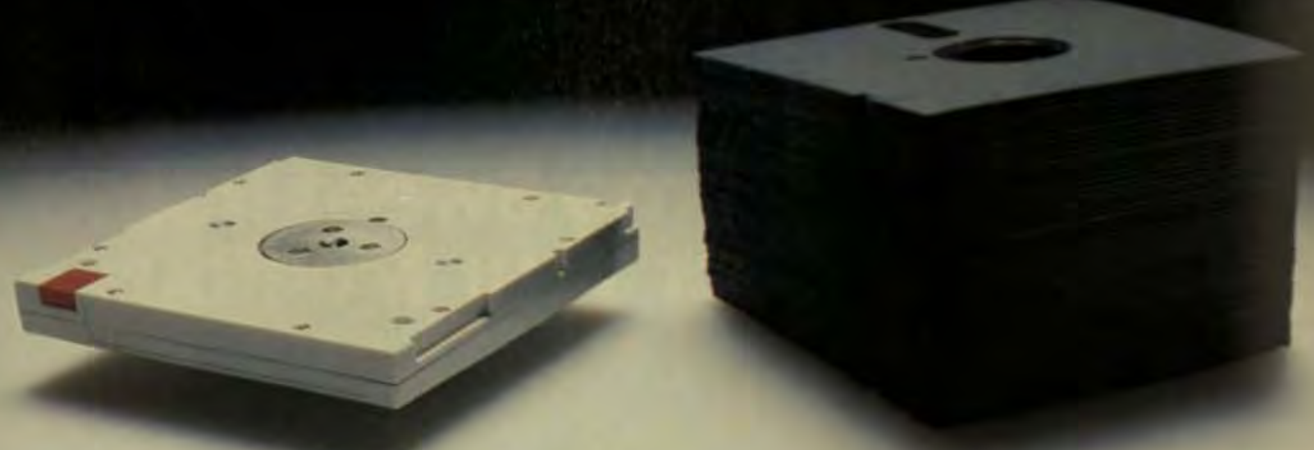

Advanced Digital's new SUPER STAR™ is the first of a family of S-100-based computer systems to be equipped with a 5-1/4" 5-MB Fixed/ 5-MB Removable Cartridge Winchester disk drive (10 MB on-line storage capacity) as standard. Now you can have the capacity of up to 50 floppy diskettes and the speed and convenience of 100% backup using the industry-standard Dysan 5-MB cartridge.

SUPER STAR is the ideal system for business, professional, and educational applications. Its CP/M* operating system (standard) gives you access to a vast library of readily available software programs. The six-slot motherboard allows plenty of room for expansion.

And as your business grows, SUPER STAR can grow with you. If you need more memory, just plug it in. SUPER STAR is based on the S-100 IEEE standard. To add more users, simply add Advanced Digital's SUPER SLAVE™ processor boards

and TurboDOS* operating system. Now you've got an extremely powerful, multi-user, multi-tasking, multiprocessing SUPER STAR computer system for up to four users.

SUPER STAR is truely a Super computer system. Look at these outstanding features:

- New Slim-Line Profile
- 5-1/4", 10-MB Winchester Disk Drive, 5-MB Fixed, 5-MB Removable Cartridge
- 5-1/4", Half-High, 48TPI Floppy Disk Drive (Osborne Format Compatable)
- Six Slot Motherboard
- Switch-selectable 110/220V Power
- Up to 4 Users With SUPER SLAVES And TurboDOS
- CP/M Operating System Standard
- One Year Warranty

SUPER STAR is the perfect system for applications requiring large data storage capacity, speed, performance, reliability, and dependable, convenient back-up capability. And at a suggested retail price of $5,000, SUPER STAR is a SUPER VALUE.

See the entire Advanced Digital product line, including the new SUPER STAR, at your local, quality computer dealer or contact:

ADVANCED DIGITAL CORPORATION

5432 Production Dr.,
Huntington Beach, CA 92649
Phone: (714) 891-4004
Telex: 4722065 ITTSCSMA

*CP/M is a trademark of Digital Research Corp. TurboDOS is a trademark of Software 2000 Inc.

SUPER STAR, SUPER SIX, SUPER SLAVE are trademarks of Advanced Digital Corp.

For European Distributor, contact: (714) [illegible]

Global Memory Systems Offer Alternatives for Increased Performance

Integrating a local bus structure with a common global memory enhances the interaction of basic system elements to maximize overall system efficiency.

by **Joseph P. Altnether,** Intel Corp.

An architecture that combines local and global memory for distributed processing increases the power and speed of microcomputer systems by optimizing the interaction among the microprocessor, memory, and peripherals. While the efficiency of any of these basic system elements can be diluted or enhanced by the performance of the others, cost-effective off-the-shelf solutions are available to meet specific system requirements.

For example, if the microprocessor is faster than the memory or the peripheral, it must be slowed down because synchronization with the memory/peripheral is necessary to transfer data among the various elements. The least disadvantageous method for synchronization is to use WAIT states to stretch the allotted memory/peripheral access-cycle time, because only one clock period is needed each time a WAIT state is used—and even that may not always be required. One WAIT state in an iAPX 86 system, for example, reduces system performance by more than 8% (Table 1), and two or more WAIT states degrade it even further. To achieve maximum performance levels, a system needs high-performance elements that can be integrated effectively.

Due to the numerous system requests for data, memory has almost as much effect on overall system operation as the microprocessor itself. Local memory is accessed more often than global memory and operates without WAIT states, enhancing system performance. Local memory is usually small—generally less than 64 Kbytes—and as such is the domain of high-speed static RAMs and integrated RAMs (iRAMs).

Global memory, on the other hand, is much larger—usually a megabyte or more—and utilizes the full addressing capabilities of the microprocessor. Because global memory is accessed less frequently, its performance can be less than that of the local memory without seriously affecting overall system performance. By combining a multiple local-bus structure with a common global bus, better memory performance can be obtained (Fig 1).

Table 2 shows four alternatives that meet the requirements of most types of global memory systems:

- High integration.
- High integration/high performance.
- Error detection and correction.
- High Performance/High Reliability.

HIGH INTEGRATION

High integration uses a single LSI controller to provide address multiplexing, refresh timing, and memory-refresh synchronization and arbitration. Until the advent of LSI controllers, the alternatives were either a discrete TTL controller, or partial solutions that provided only segments of the logic required to support a dynamic RAM system.

In either case, the system designer was left with the arbitra-

TOSHIBA HAS

OUR NEW 74HC HIGH-SPEED CMOS LOGIC FAMILY COMBINES LSTTL SPEED WITH CMOS LOW-POWER.

Get off the design tight-rope with the 74HC series, a high-speed CMOS Logic Line with an operating speed thirty times higher than the standard C^2MOS logic.

This new generation of integrated circuits makes possible new applications in high-speed portable instruments which couldn't be achieved with current LSTTL or CMOS devices.

No matter what your design parameters, Toshiba America offers the most comprehensive CMOS Logic Line available today, ranging from our standard *4000/4500 series* through our higher speed *40H series* to our newest line, the *74HC series*, the fastest yet.

ALABAMA, Components Unlimited, Inc., (205) 883-7938; **ARIZONA,** Western Microtechnology, (602) 948-4240; **CALIFORNIA,** Image Electronics, (714) 730-0303, Integrated Electronics, (213) [illegible] (408) 287-6488, (916) 424-5297, Ryno Electronics, (714) 292-6022, Western Microtechnology, (408) 725-1660; **COLORADO,** A.C.T., (303) 422-8229, Integrated Electronics, (303) 292-6121; **CONNECTICUT,** [illegible] Inc., (203) 265-3134, Future Electronics, Inc., (203) 272-1852, Milgray/Connecticut, (203) 795-0711; **FLORIDA,** Future Electronics, Inc., (305) 291-2600, Milgray/Florida, Inc., (305) 647-5747, (800) [illegible]; **GEORGIA,** Components Unlimited, Inc., (404) 441-1447, Milgray Electronics, Inc., (404) 393-9666; **ILLINOIS,** GBL Gould Electronics, (312) 593-3220, Integrated Electronics/Intercomp Inc., (312) 843-2040, Semiconductor Specialists, Inc., (312) 279-1000; **KANSAS,** Milgray Electronics, Inc., (913) 236-8800; **MARYLAND,** Milgray/Washington, Inc., (301) 468-6400, (800) 638-6656; **MASSACHUSETTS,** Cronin Electronics, Inc., [illegible] 449-5000, Future Electronics, Inc., (617) 366-2400, Milgray/New England, Inc., (617) 272-6800; **MINNESOTA,** Semiconductor Specialists, Inc., (612) 854-8841, The Joel Company, (612) 545-5660; **MISSOURI,** [illegible]

THE ANSWERS.

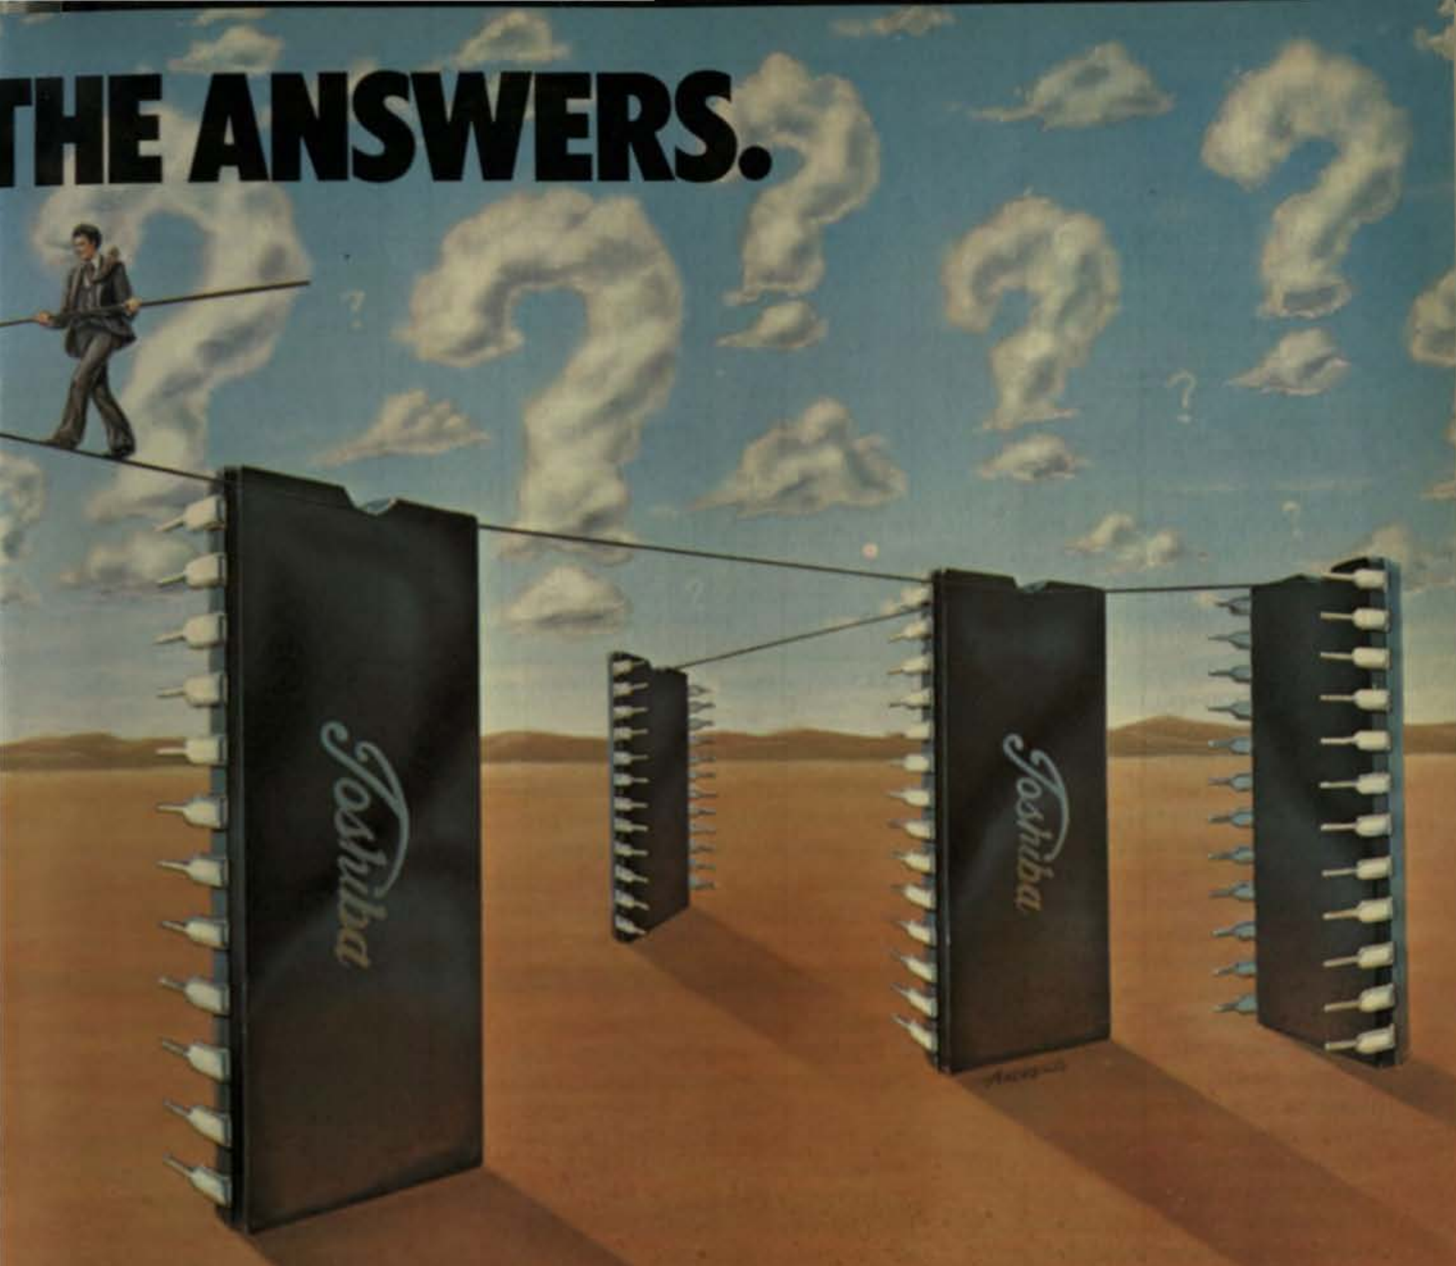

For more information about our entire logic line, write Toshiba America, Inc., 2441 Michelle Drive, Tustin, CA 92680, (714) 730-5000. Or call your local sales rep or distributor.

Wherever your imagination leads, now or in the future, Toshiba will have the answers you need.

Characteristics/ Logic Families	New HS-C^2MOS (74HC Series)	LSTTL	*HS-C^2MOS (40H Series)	Std. C^2MOS (4000/4500 Series)
Prop. Delay Time (typ) GATE (C_L= 15pF)	8ns	9ns	15ns	125ns
Max. Clock Freq. (typ) J/K F/F (C_L= 15pF)	60MHz	45MHz	20MHz	2MHz
Quiescent Power Diss. (typ) (GATE)	0.01μW	8mW	0.01μW	0.01μW
Noise Margin V_{IH} (min)/V_{IL} (max)	3.5V/1.5V	2.0V/0.8V	4.0V/1.0V	3.5V/1.5V
Output Current I_{OH} (min)/I_{OL} (min)	4mA/4mA	0.4mA/4mA	0.36mA/ 0.8mA	0.12mA/ 0.36mA
Op. Volt. Range	2-6V	4.75-5.25V	2-8V	3-18V
Op. Temp. Range	−40-85°C	0-70°C	−40-85°C	−40-85°C

*Data believed to be accurate and representative of each logic family.

TOSHIBA AMERICA, INC.

A WORLD STANDARD IN MOS.

277-2400; Semiconductor Specialists, Inc., (816) 452-3900; **NEW JERSEY,** General Components, Inc., (609) 227-1800, Milgray/Delaware Valley, Inc., (609) 983-5010, (800) 257-7808, (800) 257-7111; **NEW MEXICO,** Western Microtechnology, (602) 948-4240; **NEW YORK,** A.C.I. Electronics, (516) 293-6630, CAM/RPC Electronics, (716) 865-2080, (716) 865-2083, Future Electronics, Inc., (315) 463-6633, Milgray/Upstate New York, (716) 385-9330, Milgray Electronics, Inc., (516) 546-5600, Rome Electronics, (315) 337-5400; **NORTH CAROLINA,** Components Unlimited, Inc., (919) 782-3004; **OHIO,** Milgray/Cleveland, Inc., (216) 447-1520, (800) 321-0006, (800) 362-2808; **PENNSYLVANIA,** Semiconductor Specialists, Inc., (412) 963-7241; **TEXAS,** A.C.T., (214) 980-1888, (512) 452-5254, A.C.T. Houston, (713) 496-4000; **UTAH,** Integrated Electronics, (801) 298-1868; **VIRGINIA,** Components Unlimited, Inc., (804) 237-6291; **WASHINGTON,** Integrated Electronics, (206) 455-2727; **CANADA,** Carsten Electronics, Ltd., (613) 729-4138, (416) 751-2371, (514) 334-8321, Future Electronics, Inc., (416) 663-5563, (514) 694-7710, (613) 820-8313, (604) 438-5545, (403) 259-6408.

tion circuit, the most difficult control element to design. Resolution of high-speed asynchronous requests can cause elements of the arbiter to operate beyond specification, causing the most nefarious problem—arbitration failure. Because such failures occur so infrequently, they must be traced with

TABLE 1 — EFFECT OF WAIT STATES ON IAPX SYSTEM PERFORMANCE				
NUMBER OF WAIT STATES	0	1	2	3
SYSTEM PERFORMANCE (%)	100	91.7	83.7	73.7

an oscilloscope, which requires great tenacity. To overcome the arbitration problem, Intel's 8203 uses a two-stage flip-flop circuit to synchronize refresh and access requests. In addition, its design replaces 30 to 40 TTL packages. This high-integration solution reduces the component count to lower total system cost and provides more reliable system operation.

Because of the temporary nature of the data stored in the 64K dynamic RAM, the controller performs a refresh cycle to preserve data. Since this function is not synchronized with memory-access cycles, the memory will occasionally be involved in refresh when a memory request is made. The controller's high-speed arbiter resolves these conflicts and synchronizes the microprocessor via the system's acknowledge signal (SACK). If refresh is being performed and a memory cycle is requested, SACK requests WAIT states until the controller can service the memory access to the microprocessor.

At high microprocessor speeds, the system speed is limited by the ready response time of the controller. At higher speeds, the memory is assumed to be not ready until the controller transmits that it is ready. A WAIT state is automatically added to allow time for the controller to inform the

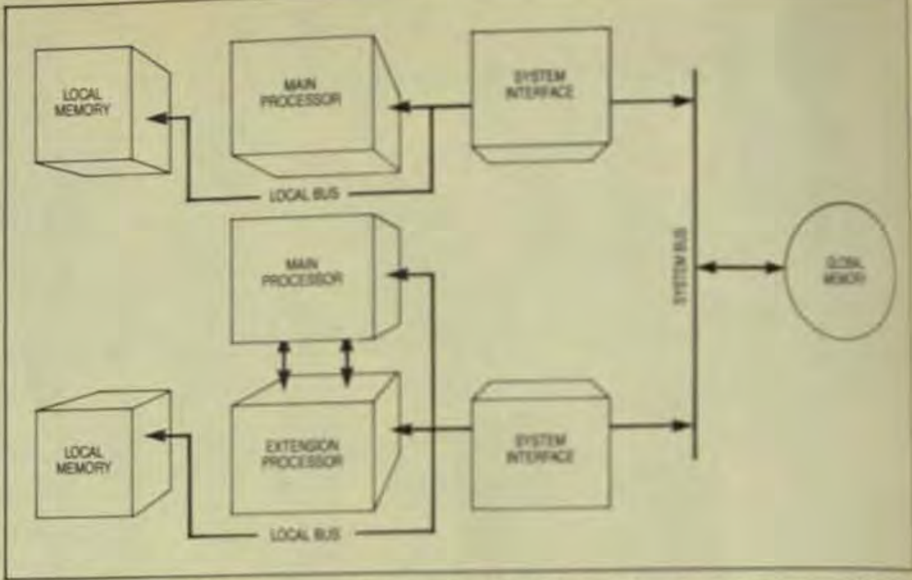

Fig 1 New system architectures combine local and global memories to enhance processing power and speed.

microprocessor of a not-ready-to-transfer condition. This system design offers the lowest cost and performance of the four basic types.

HIGH INTEGRATION/HIGH PERFORMANCE

The addition of five TTL packages to the high-integration design (Fig 2) increases system performance. This type of memory system operates at 5 MHz with no WAIT states, using 150-ns dynamic RAMs. It monitors the status bits from the microprocessor to increase overall system performance. These bits indicate the type of operation to be performed and allow the controller to begin a memory cycle earlier in the microprocessor cycle and so it will have the time to offer a not-ready indication and inject a WAIT state. WAIT states are required only 3% of the time, and only if refresh is occurring when microprocessor access is requested. This system has all the features of a high-integration system—reduced cost, lower component count, and high reliability—in addition to increased speed.

TABLE 2 — MICROSYSTEM DESIGN SOLUTIONS AND ALTERNATIVES			
SYSTEM CRITERIA	SOLUTION	MAXIMUM BLOCK SIZE	DESCRIPTION
SIMPLE, LOW-COST LOWER DESIGN	HIGH INTEGRATION	256 KBYTES	LSI-BASED SYSTEM FOR COMPLETE MEMORY CONTROL TO OVERALL COST AND FACILITATE DESIGNS
LOW-COST, HIGH-PERFORMANCE DESIGN	HIGH INTEGRATION AND HIGH PERFORMANCE	256 KBYTES	LSI-BASED SYSTEM OPTIMIZED FOR NO WAIT-STATE OPERATION AND HIGH PERFORMANCE
HIGH RELIABILITY	ERROR CORRECTION	256 KBYTES	LSI-BASED SYSTEM FOR ERROR-CORRECTION TO PROVIDE VERY HIGH RELIABILITY, FACILITATE DESIGNS, AND MINIMIZE BOARD SPACE
HIGH PERFORMANCE, RELIABILITY	SUPERIOR PERFORMANCE	512 KBYTES	LSI-BASED SYSTEM FOR ERROR-CORRECTION AND HIGH PERFORMANCE OPERATION. PROVIDES SUPERIOR RELIABILITY

ERROR CORRECTION

The error-correction solution provides increased reliability by the addition of error-correction codes (ECCs). The soft-error rate typically is 10 times greater than the hard error rate, and in some systems the natural reliability level may be insufficient for the application.

Natural reliability is defined as the mean time between failures (MTBF), an inverse function of the product of the number of devices and the device failure rate. The failure rate is the sum of both the hard-error failure rate and the soft-error failure rate. Thus, a large system comprising thousands of devices with individually low failure rates could have an MTBF that exceeds the system failure rate simply because of the sheer number of devices.

In most applications, the MTBF target is defined by such factors as system downtime, service-call costs, and system costs. In other applications, danger to life or finances demand extremely high reliability.

Error-correction codes increase the natural reliability by several orders of magnitude (Fig 3). Data for these curves is based on the Intel 2164A with a soft-error

ANNOUNCING THE WY-50.

PRETTY BABY

A NEW GENERATION OF LOW-COST TERMINALS FROM WYSE

Never before has anyone put so much into something so good looking. Or so compact. It took revolutionary design to do it. Design a lot of people couldn't accomplish for the price. But we did.

In fact, the birth of the WY-50 introduces a new standard for low cost terminals. It was conceived to give you 30% more viewing area. Developed to meet the most advanced European ergonomic standards. And best of all, delivered for a surprisingly low price.

The WY-50 sells for only $695.00.

BORN WITH ADVANCED FEATURES:

- 14" screen.
- 80/132 column format.
- Soft-set up mode.
- High resolution characters.
- Low-profile keyboard.
- Industry compatible.
- Only $695.00.

For more information on the revolutionary design, outstanding features and unique good looks of our Pretty Baby, contact WYSE TECHNOLOGY and we'll send you a brochure filled with everything you need to know.

The WY-50. It's a real beauty.

WYSE *Make the Wyse Decision.*

WYSE TECHNOLOGY, *3040 N. First St., San Jose,* CA 95134, 408/946-3075, TLX 910-338-2251. *in the east, call 516/293-5563, call toll-free 800/538-8157, ext. 932, in CA 800/672-3470, ext. 932 in the midwest, 313/227-5080, in So. CA 213/340-2013*

Inherent Pixel Memory In AC Plasma Displays Increases Screen Data

By avoiding the need to refresh, ac plasma panels can display large quantities of data without flicker. As many as four concurrent windows can be manipulated.

by P. Pleshko and Victor Tang, IBM Corp.

AC plasma displays can support large amounts of data without refresh because the inherent memory of each pixel eliminates the need to refresh the display. As a result, large quantities of information can now be displayed on flat-panel devices, bringing large-capacity information capabilities to office information processing.

Early non-CRT devices included light-emitting diodes (LEDs) and dc gas-discharge devices, both limited to applications that required only a few characters. These applications were soon expanded to include displays with 32 to 240 lines of characters for such applications as banking. Then liquid-crystal (LCD), electroluminescent, and vacuum fluorescent displays were introduced.

Each of these technologies suffered from the need for refresh. That is, the information displayed by a particular pattern of pixels has to be excited repetitively by selectively applied voltage patterns so that the information is regenerated at least 60 times per second to prevent the viewer from perceiving screen flicker.

This refresh requirement limits the maximum number of characters that can be displayed because the duty factor (and hence, the brightness or contrast ratio) decreases as the number of display lines increases. Thus, the many products based on these technologies have been incapable of displaying single or multiple pages of information.

AC plasma displays, in comparison, allow information written on the display to remain until erased and rewritten. The average brightness isn't governed by duty-cycle considerations that depend on the number of horizontal lines of pixels, and there's no perceptible flicker because the image is sustained at several tens of kilohertz—approximately two and a half orders of magnitude higher than the flicker frequency.

This article discusses IBM's 581 ac plasma flat-panel display subassembly, which is a unique, large-screen, high-information-content display with minimum screen update time, and high power efficiency for computer-output applications.

DISPLAY-PANEL CHARACTERISTICS

The 581 ac plasma-display panel has crisp characters formed by the pixels that are illuminated. The black layer of material on the rear surface of the panel is applied to eliminate rear glass-surface reflections (ghosting).

The front surface has a bronze glass sheet affixed to it to enhance the contrast by reducing glare (front-surface specular reflections) by diffusing and filtering incident light. The panel has an average brightness of 12 foot-lamberts averaged over space and time, or 30 foot-lamberts peak brightness (without the antiglare) with a contrast ratio of 18.

The 960 vertical lines and 768 horizontal lines at 71.4 lpi provide 0.75 million pixels in the approximately 1 ft^2 viewing area of the panel. This high resolution allows the displaying of complex halftones, approximately 10,000 characters of information, or


TeleVideo

Brand new terminal. Same old story.

If you're looking for the best terminal value, look to a familiar name: TeleVideo. The story of our new 924 is the same story you've heard for the last four years. We've combined the best innovations in technology, design, and quality manufacturing to give you everything you want in an advanced terminal. And we've done it at a price that can't be beaten.

The 924 is comfortable, with a tilt and swivel non-glare screen. Its lightweight, low profile keyboard fits your fingers naturally. Our 32 programmable, non-volatile function keys turn often used instructions into simple one-button commands. The 924 provides extra display features, including full screen editing, character and block graphics, definable scrolling regions for multiple window operation, plus 32 character-by-character visual attributes.

The 924 has the intelligence you need too, with logical attributes to define protected and unprotected regions. And, with standard one page or optional four page memory, there's plenty of room to work.

If you don't need all the power of the 924, try our 915. It has all the design advantages of the 924, and is the superior terminal in its class. It just costs less.

The 924 and 915 mark a great new chapter in the TeleVideo story. And, like all TeleVideo terminals, they're reliable performers, backed nationwide by General Electric Instrumentation and Computer Service Centers. Contact us today about the 924 and 915 story and you'll enjoy the happy ending. Call 800-538-8725 (in California call (408) 745-7760), or send us this coupon.

TeleVideo Systems, Inc.
Dept. 221C
1170 Morse Avenue, Sunnyvale, CA 94086

Yes, I'd like to know more about TeleVideo's new 924 and 915 terminals.

Name ____________
Address ____________
City ____________
State ____________ Zip ____________
Telephone (______) ____________

Santa Ana/California, 714-557-6095
Sunnyvale/California, 408-745-7760
Atlanta/Georgia, 404-399-6464
Dallas/Texas, 214-980-9978
Chicago/Illinois, 312-351-9350
Boston/Massachusetts, 617-668-6891
New Jersey/New York, 201-267-8805
Woking, Surrey/United Kingdom, 44-9905-6464

TeleVideo Systems, Inc.

Fig 1 The high-resolution of IBM's 581 ac plasma-display panel displays complex half-tones (**a**), two pages of information (**b**), or multiple-screen partitioning (**c**).

two pages of information, or multiple screen partitioning (1920-character formats) for viewing and interacting (Fig 1).

A block diagram of the ac plasma-display panel is shown in Fig 2. The 960 vertical lines are brought out equally, half on the top and half on the bottom of the panel, as even and odd lines. The 768 lines are also brought out equally, half on the left and half on the right side of the panel. This was done to reduce the density of interconnections to the panel and affects the architecture of the interface. The interface, as shown, consists of 10 data-addressing lines, five function lines, and two lines with signals sent to the using system. The signal lines operate at TTL logic levels, with the down level as the active state (ONE).

The 10 data-addressing lines consist of the following:

- The load-horizontal line selects the axis into which the data is loaded. When inactive, data is loaded into vertical modules. When it is activated, it directs the module-select and the even-odd data-input lines to the horizontal driver modules.
- The four module-select lines select driver modules on either axis.
- The even-data line, when active, sends data to either the left-horizontal or top-vertical drivers.
- The odd-data line, when active, sends data to either the right-horizontal or bottom-vertical drivers.
- The clock-data line enters data on the negative clock transition and can be any frequency up to 3.0 MHz.
- The reset-X data resets all latches in all vertical driver modules and is usually activated prior to the loading of new data into the vertical drivers.
- The set-panel line is used to override data latched in the vertical drivers, and will cause a 960-pixel write or erase to occur without altering the data set in the vertical latches.
- A full-panel update (erasing the panel, writing new data, and allowing 26 ms for asynchronous ...) can be performed in 200 ms.

The five function lines consist of the following:

- The **write command** causes latched data in both the vertical and horizontal drivers to be written on the panel. The maximum amount of data that can be latched at any one time is 960 pixels in the X direction and ... pixels in the Y direction.
- The **erase command** causes erasure of the data on the panel corresponding to the data latched in the vertical and horizontal drivers, and can erase as many as 960 × 16 pixels.
- The **test-input signal** operates a diagnostic facility internal to the display panel. Four patterns are generated in the sequence, which takes less than 10 s. ... logic must be issued by the using system during a power-on sequence. It resets the control logic and causes the auto-adjust circuitry to find the operating voltage automatically.
- The **oscillator line** is self-explanatory.
- The **ready line** sends a signal to the using system, indicating when the panel is ready to receive data or commands. The panel is not ready when a write, erase, reset logic, test, or power-on sequence is in progress.

Image limit is a signal sent to the using system warning of an image overload condition. This condition is reached when 31% of the pixels are on. If the number of pixels is increased more than an additional ... after the image-limit signal is generated, a hardware-controlled reset and panel erase will occur.

The signal ground consists of seven pins. The power required by the panel as a function of the percentage of pixels is shown in Fig 3. The ... on this linear relationship ...

The terminals that face up to everyone's problems.

Now Qume offers a line of three CRT terminals that help you face every one of your company's needs: The Qume QVT™ series.

The low-cost Qume QVT 102™ has all the best features of the most popular terminals and can emulate three with a keystroke! QVT 103™ is code-compatible with the DEC VT100/VT 131 and offers 80 or 132 character/line, split screen and up to 4 pages of memory. And QVT 108™ faces up to the most demanding jobs with 22 function keys, 13 editing functions and 2 pages of memory.

Qume QVT terminals all share striking ergonomic design. Tilt/swivel screens in green or amber. Detached, low-profile keyboards. Big 9 x 12 character cells. And a long list of other features to make people more comfortable and productive.

Best of all, Qume QVT terminals are very competitively priced. You can have performance, features and reliability without facing a budget problem.

Talk to your Qume sales office about filling all your terminal needs with a single purchase of QVT terminals. Or write Qume, 2350 Qume Drive, San Jose, California 95131.

One size fits all.

Qume.®
A Subsidiary of ITT

QUME CORPORATION
HEADQUARTERS
2350 Qume Drive
San Jose, CA 95131
CALIFORNIA
San Jose (408) 942-4111
Culver City (213) 410-1458
Santa Ana (714) 957-4040
COLORADO
Aurora (303) 752-3000
ILLINOIS
Palatine (312) 991-7250
OHIO
Dayton (513) 439-0469
TEXAS
Irving (214) 659-0745
NEW JERSEY
Edison (201) 225-5005
MASSACHUSETTS
Bedford (617) 275-3200
GEORGIA
Decatur (404) 284-8500
GERMANY
Düsseldorf PH: 0211743016
ENGLAND
Reading, Berkshire
PH: 734-584-646
FRANCE
Boulogne PH: (1) 6082334
CANADA
Quebec PH: (514) 695-3837

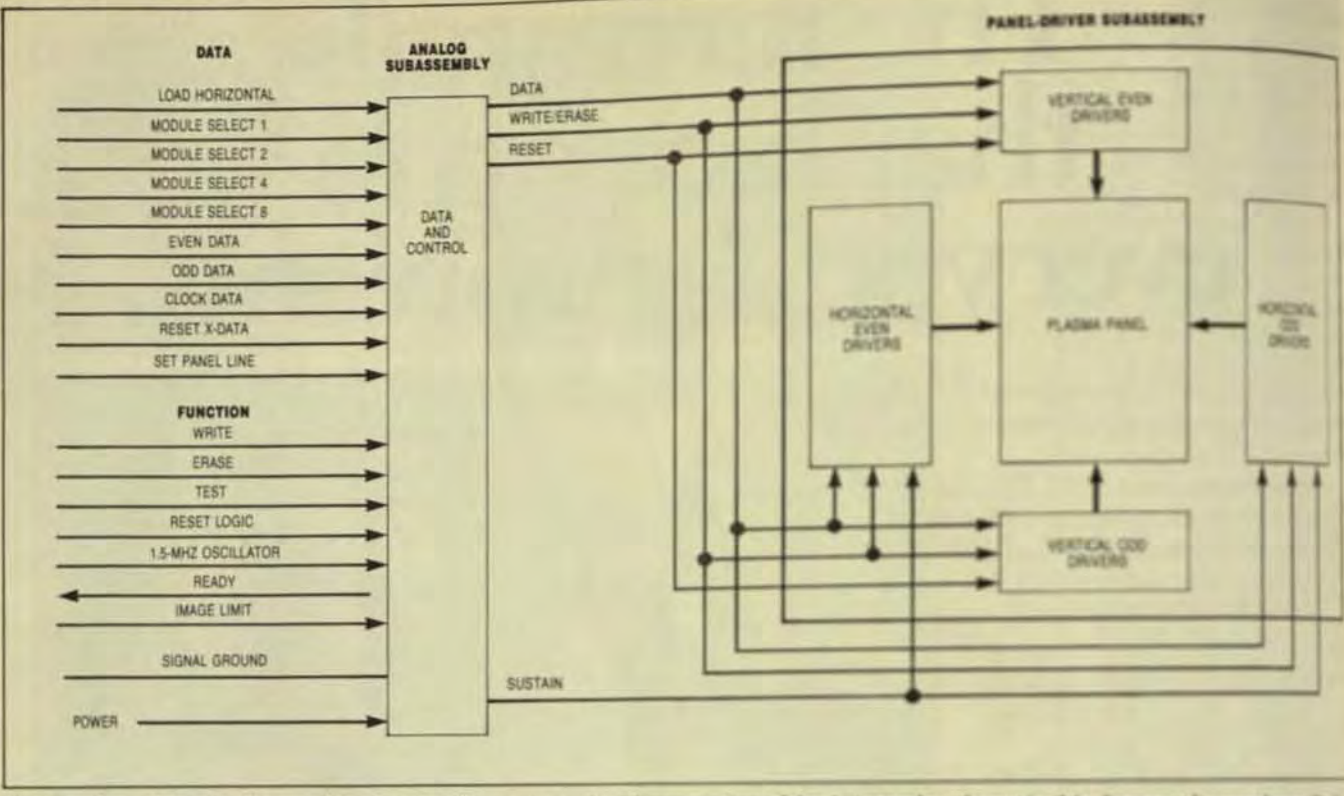

Fig 2 Ten data-addressing lines and five function lines are required for operation of the 581 panel as shown in this diagram of an ac plasma display

intercept) is due to the standby power of the circuits and the power needed to drive the capacitive loads in the panel.

One of the unique functions of the 581 technology is a built-in auto-adjust. This term refers to the microprocessor function that automatically sets the amount of the operating voltage needed to sustain an information pattern after it is written or erased on the panel. Auto-adjust is accomplished by use of lit pixel patterns and sequences to test for the voltage operating window and then referencing the sustain-operating voltage to the voltage at the bottom of the panel's voltage-operating window. This sequence eliminates any initial factory adjustment or adjustment in the field because it's initiated at each power-on sequence.

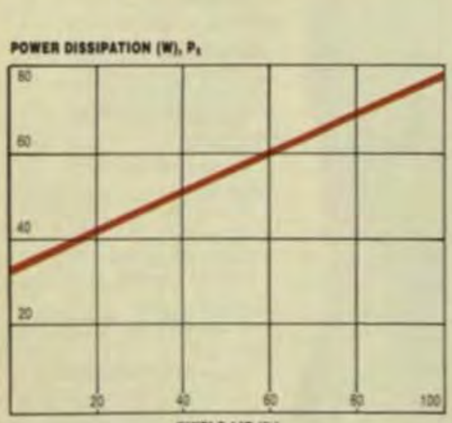

Fig 3 The power required by IBM's 581 panel is shown as a percent of pixels lit.

The price of this technology is plotted in Fig 4 and extrapolated for other display sizes. This curve shows that because of the drivers required, a matrix flat-panel technology isn't as cost-effective in smaller panel sizes. This situation is reflected in the cents/pixel curve as well. This curve can also be used to assess the price-effectiveness of alternative flat-panel technologies at various sizes. Both economies of scale in manufacturing and hardware-design effectiveness are evaluated in such comparisons.

Fig 5 shows an extrapolated luminous-efficacy curve for different panel sizes in the 581 technology with the panel condition in which all cells are lit. This curve can be used to compare the electro-optical efficacy of alternative technologies and designs. It provides an indication of the total power "efficiency" of a 581-like functional-display block. Those designs that are above the curve have a less power-efficient design, and those below are more power-efficient.

AC PLASMA-DISPLAY OPERATION

The basic 2-terminal ac plasma-display device is shown schematically in Fig 6a. Driving-point terminals are connected via capacitors to a gas-filled cavity that emits light in the visible range when a discharge is formed. Because of the capacitive coupling to the gas, an ac waveform is required for continuous emission of light pulses.

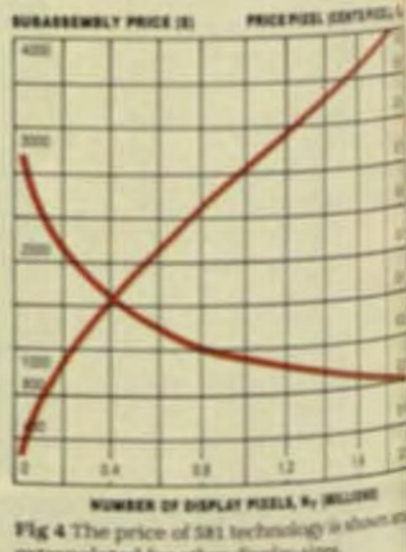

Fig 4 The price of 581 technology is shown and extrapolated for other display sizes.

Driving each cell individually in a large array of cells would result in a prohibitively high cost for drivers and connections between the drivers and the cells. One way to reduce driver and interconnection costs is to interconnect the electrodes of each cell in a matrix array.

For individual plasma cells to operate independently when interconnected as a matrix array, an impedance in series with the gas discharge is required at each intersection of X,Y lines. Capacitors provide the necessary electrical isolation and are fabricated with a continuous film of dielectric material.

As shown in Fig 6b, an ac plasma panel consists of an assembly of two substrates spaced to form a chamber containing a neon gas mixture. Each substrate holds a set of parallel conductors covered by a transparent dielectric, and the substrates are oriented so that the two sets of conductors are orthogonal. Selected intersections of these conductors become localized spots of neon-colored light when driven by suitable signals, thereby creating a display.

The most common mode of operation for an ac plasma display is the bistable or memory mode of operation. In this mode, the cell produces two levels of brightness, a bright ON level and an OFF level. The bistable operation yields cell brightness that is independent of the number of pixels in the display panel. Brightness is determined by the intersection of the dynamic-load line and plasma-discharge characteristic.

In the bistable mode, a continuous ac square-wave voltage called a "sustain" voltage is simultaneously applied to all cells. This causes the ON cells to discharge repetitively while OFF cells remain unfired. Because all of the cells are sustained in parallel, the use of a single sustain driver for a large matrix array requires relatively large sustain currents.

In Fig 7, the potential difference V_{ab} that equals the firing (ionizing voltage V_f of gas, is applied to the cell at time t_0. The gas ionizes, and electrons and ions move toward the anode and cathode respectively. This first discharge produces a short burst of neon-colored light and a voltage (from the wall charge) that opposes the external potential difference, quenching the discharge.

At time t_1, the external voltage polarity is reversed and reduced in magnitude to a point below V_f. This re-

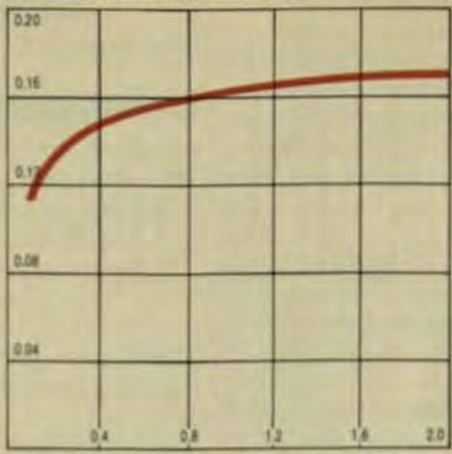

Fig 5 This extrapolated luminous-efficacy curve for the 581 panel assumes that all pixels are lit.

duced applied voltage is called the sustain voltage of the cell. Since the polarity of the applied voltage has been reversed, it now aids the wall voltage produced during the first firing. The sum of these two voltages exceeds the firing voltage and a second firing occurs.

This second discharge produces another pulse of light, and an opposing wall voltage. At time t_2 the polarity is again reversed, initiating a third firing that results in a third burst of light and an opposing wall voltage.

In summary, a high potential difference V_f is used to initiate the first gas discharge. Subsequently, a sustain drive of alternating polarity and lower amplitude is used to produce a burst of light and a sustaining wall voltage for every alternation. Thus the cell has been turned ON (written), and is being sustained in this ON state. Typical values for V_f and for the sustain operation are 150V and 90V, respectively.

The sustain voltage can't produce firings without the aid of the wall voltage, so the erase operation is simply one in which a weak firing is forced that cannot produce enough wall voltage to sustain the subsequent firings. This reduction in the wall voltage to or close to zero is generally accomplished by a low-amplitude alteration of the applied ac waveform.

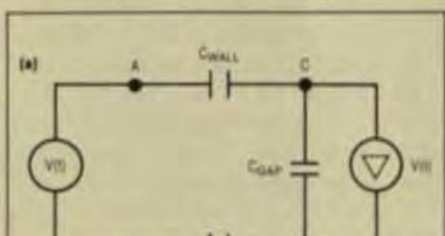

Fig 6 A simplified ac plasma-panel cell is shown as an electrical equivalent (**a**) and as a physical representation (**b**). Two substrates form a chamber containing a neon-gas mixture. Each substrate holds a set of parallel conductors covered by a transport dielectric.

The gas discharge that occurs in a cell emits light from two regions—the negative glow and the positive column. For neon with 0.1% argon dopant, the radiation from the negative glow region is at 5850 Å, while the radiation from the positive column is predominantly at 7630 Å.

Cell brightness is given by the time-averaged value of the light emissions from a cell. Thus, the peak value of light emission and its variation with time, space, and repetition rate, determine the cell brightness. The instantaneous cell-light output waveform follows the current waveform in time and varies directly with the cell current. Therefore, as cell current increases with wall capacitance and pressure, so does the brightness. In addition, an increase in the sustain frequency, which increases the number of light pulses per unit time, increases cell brightness.

AN AC PLASMA-DISPLAY OFFICE APPLICATION

The IBM 3290 information panel utilizes plasma technology in a display with functions required to meet the needs of information workers. This flat panel provides display capacity, information handling, ergonomics, and application support.

The user of the IBM 3290 is the information worker. It's been estimated that 55% of the American labor force are information workers, paid to process data or information. But whereas each industrial worker is supported by about $40,000 of capital investment, the information worker has yet to be supported at a similar level.

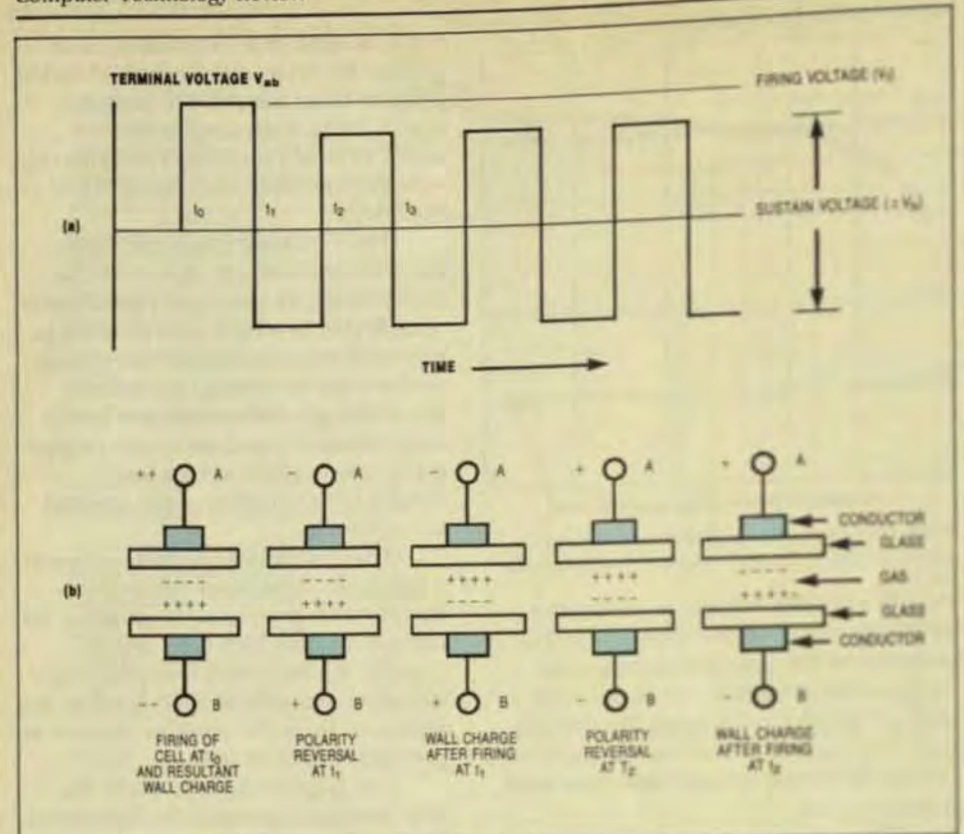

Fig 7 The sequence of events that lights a panel cell involves two firing (ionizing) voltages. The first discharge produces a short burst of neon-colored light; the second reverses the polarity and produces a second light pulse. Subsequently, a sustaining wall voltage is produced for every alternation.

For example, one study reports that the information worker spends 90% of his time collecting, distilling, organizing, and communicating information, and less than 4% in analysis and exploitation. To provide higher quality tools to information workers, close to half of the office population will have terminals by the end of the decade. Databases are transforming entire industries.

The challenge is to provide displays with capabilities that expand the effectiveness of the information worker. The 3290, a flat-panel device that's capable of displaying alphanumeric data and business graphics, has a capacity of about 730,000 pixels (almost 10,000 characters) with advanced screen management capabilities.

The display is partitioned into as many as four concurrent windows, each of which can be the size of a conventional 1920-character display. Each window can interact with a different application in a different processor, in widely dispersed geographical areas. In effect, one information panel is capable of functioning as four displays. This capability of concurrent viewing and concurrent processing permits "relational informational processing." For example, it allows the display of menus in one window, the application results in another, and the graphical representation of results in a third.

While all these capabilities provide a critical mass of functions, it is the applications that expand the effectiveness and efficiency of the information worker. There are four important applications. They are: report management distribution system (RMDS), program development, large-screen database/data communications, presentation graphics, and information-center data analysis.

RMDS provides facilities for the viewing of a full page of computer-printed output. Studies show that it costs up to $0.25 per page to deliver a printer-output page to the information worker. This application facilitates significant reduction in the use and handling of paper.

Program development exploits the large screen and its multiple concurrent viewing and execution capabilities to improve programmer productivity. The programmer can view program listings, compiler output, and test executions simultaneously while making corrections.

The database/data communications with the information panel permits users to interact with multiple applications concurrently and view the results simultaneously for correlation and checking operations. Data from the active screen can be moved to copy areas for subsequent comparisons with other screens from the same application or compared with screens from other applications.

Information-center data analysis, supported by the graphical data display manager, allows the construction of pictoral graphs and charts while other screens on the display interact with an application. In this manner, a new level of data-analysis support with concurrent data and graphics update is provided. This application offers data-analysis capability to the professional who lacks in-depth programming skills.

FUTURE TECHNOLOGY

One image-quality parameter is resolution. AC plasma technology has been demonstrated at 125-lpi resolution. The resolution required to reduce the visual perception of staircasing will be mitigated by the ability to produce gray levels, as has been demonstrated in the industry. Screen-update times will also be faster, with 1/30 of a second being projected.

Color has been demonstrated in ac plasma displays by use of a gas that emits ultraviolet light instead of visible light to excite different color phosphors. However, crosstalk among cells requires photon isolation, which complicates the internal structure of the panel. Color resolution of 50 to 60 lpi seems possible in the future. ■

Peter Pleshko *is manager of display technology development and manufacturing at IBM's Kingston, NY, facility. He received a PhD in electrical engineering from New York Univ. and an MBA from Pace Univ. He holds 11 patents.*

Victor Tang *is manager of strategy planning for plasma-display products at IBM, which he joined in 1968. He has BSEE and MS degrees from Purdue Univ. and a Master's from Columbia Business School.*

Elastomeric Switches Find Growing Keyboard Market

by **Paul Miller**, Selec-Key

Elastomeric switches (contact pads) have made a substantial penetration of the switch market. They have created a new niche in the specialty market by offering distinctive features that allow additional versatility in the design of the final electronic device. Contact pads are used in just about every type of electronic product that uses a printed-circuit board. Off-shore producers are supplying them to their domestic markets and are now penetrating the United States electronic markets.

A contact pad is a momentary switching device normally supplied with an eleastomeric contact to close a circuit. The conductive contacts are located directly over the traces of a PCB and short the traces when the pad is pressed, which eliminates the need for other switching devices and connectors.

Numerous versions currently are on the market. Some of the most popular designs are:

- A dome design, usually with a ring around the top that provides over travel. This type is used under hard keytops.
- A cone design, normally used for short-stroke applications. This type is also mounted under hard keytops.
- An integrated design that incorporates a keytop into the contact pad. Legends can be printed on the keytop.
- A nontactile design that operates in a manner similar to the membrane switches constructed with laminates of polycarbonate and polyester.

Traditionally, silicone rubber has been used in the contact pad because of its excellent environmental and electrical properties. Silicone rubber has excellent ozone resistance; it is outstanding in high and low temperatures; its electrical properties are well known; and it is resistant to most chemicals that would be around keyboards. Air vents are designed into the part to relieve pressure when the switch is closed. These vents are interconnected to each switch location to distribute pressure equally over the entire keyboard.

The resistivity of the elastomeric contact material will test 5 to 10 Ω/cm under standard conditions, however actual parts may go as high as 200 Ω in applications on PCBs. This resistance could lead to rapid temperature buildup if high current loads are drawn. You should use no higher than 30 mA for 0.5 s. Contact bounce is less than 5 ms, so the circuitry must be designed to function with silicone contact pads.

The life expectancy of silicone contact pads varies with the design and material formulation. A general rule of thumb is the greater the activating force, the shorter the functional life. Typically, a ring-dome design in a standard silicone material will function several million cycles. The ring on top of the dome relieves some of the flex stress that's normally carried by the dome walls and improves both the life and feel of the part.

In the near future, new technology will reduce the resistance of the contact and materials will be improved to increase the elastomeric pad's functional life. The use of solid-state devices is rapidly increasing, and the use of contact pads will grow with the market. ■

Intelligent Keyboards

by **Mark Tiddens**, Key Tronic Corp.

A dramatic change in the design and packaging of CRT terminals is becoming evident—an unbundling of the terminal into two separate components, the monitor and the keyboard. In the last couple of years, most terminals have changed to the use of detached keyboards because of human factors (ergonomic) concerns centered around the need for flexible physical placement of the keyboard and monitor. Keyboards have also adopted lower-profile designs with nonglare keys. Ergonomic trends for monitors include compact packaging, tilt and swivel, nonglare surfaces, and alternate display colors (such as amber).

Separation of the keyboard and display is only the first step terminals have taken to follow the same trend as personal computers and component television; it makes sense to separate the display from the rest of the system so that it can be chosen according to the application and the latest available technology.

In order to unbundle the terminal, the electronics have to be removed from the display. Although they could be placed in a separate box, the terminal electronics have been integrated to a point at which they can be enclosed within the keyboard. Detached keyboards already require a printed-circuit board and some electronics, so PCB extension and added electronics is more cost-effective.

Thus, intelligent keyboards were devised—keyboards that only need to be connected to a monitor to become a complete CRT terminal. Key Tronic, for example, recently began providing OEM customers with custom keyboards containing terminal electronics to emulate a VT-100 terminal. Zenith Data Systems is another firm that has recently introduced terminals that can be purchased without the monitor.

When the terminal electronics are placed inside the keyboard, one question concerns where the power supply should be. Another is the number of cables that will be needed. Since it's preferable to have only one coiled cable going to the keyboard, a practical design solution is a small connector box that can be placed behind or beneath the desktop. This box can hide from view all the connectors and cables for the power, modem, printer, monitor, and other peripherals.

All cables would then connect to a single box that could either be plugged directly into the ac wall socket, or fastened to the desk or wall. The desktop is then left with two components—a low-profile keyboard and compact monitor—both of which take up little room and make a desktop terminal more practical by eliminating the need for a separate table.

The unbundling of electronics equipment—including everything from hi-fi to television, personal computers, and now terminals—will have far-reaching effects in the future. You can envision the mixing and matching of many types of computer, visual, and audio components. The continued standardization of terminal features and interfaces and the unbundling of the monitor from the terminal are major steps in this direction. ■

Take A Closer Look At The MPC Alternatives

• Ergonomics • Price/Performance • Plug Compatible Emulation • UL Listed • FCC & CSA Approved • Available Now*

MPC 1100

The MPC 1100 video display terminal is a comprehensive PERKIN-ELMER™ 1251 Emulator. It is 100% plug compatible with all RELIANCE-PLUS™ operating system software and application programs. An advanced monochrome alphanumeric terminal, the 1100 offers significant advantages over the 1251 personality.

MPC 2100

The MPC 2100 series of color video display terminals emulate the full functionality of the ISC 8301™ series. They accept the software commands of PLOT-10™, IGS™ and ISSCO™ graphic software. A high resolution, self convergent monitor assures superior color graphics.

MPC 1200

The MPC 1200 series are ANSI ... display terminals. The 1200 fully emu... the DEC™ VT/131 & VT/132. The MPC ... is completely compatible with DEC™ ... architecture, accepting all REGIS™ ... commands to produce superior ... graphics.

You'll See Built-in Extras as Standard Features

MPC's Alternatives bring you the costly options of other terminals as standard features. A new generation of engineering excellence delivers an unmatched range of standard features:

- **14" Screen:** Enlarged screen assures easier reading and reduced eyestrain.
- **Detachable Keyboard**
- **24 Programmable Function Keys:** Plus, the numeric keypad adds 10 more programmable function keys!
- **40/80/132 Columns:** User selectable for individual needs.
- **Enhanced Scrolling:** Jump, Smooth, and Bi-directional scrolling for easier reading and editing.
- **8 Pages of Video Memory:**
- **Two RS-232 Ports**
- **Double High/Double Wide Character Capability**
- **Soft Set-Up Mode Plus Host Configurable:** All attributes can be changed dynamically from keyboard *or* host computer.
- **Multiple Screen Attributes**
- **Flexible Split Screens**

Color Monitor Option: All MPC terminals also available in color.

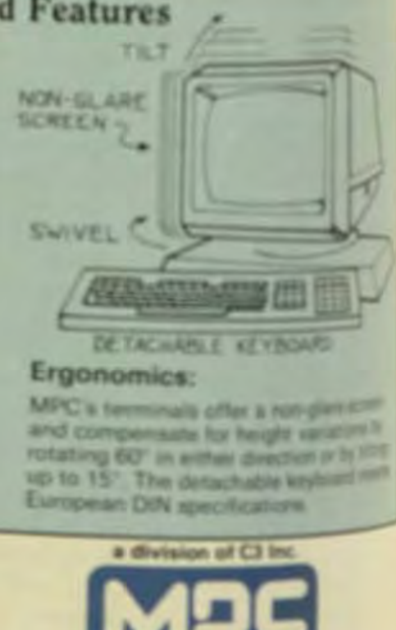

Ergonomics:

MPC's terminals offer a non-glare ... and compensate for height variations ... rotating 60° in either direction or by ... up to 15°. The detachable keyboard ... European DIN specifications.

Contact Your Nearest MPC Distributor

■ **Kierulff Electronics:** Call your Kierulff Division, (outside Calif 800-338-8811) ■ **Computer Peripherals Inc.,** Richardson, TX (214) 644-3606 ■ **Digital Solutions, Inc.,** Marietta, GA (404) 955-4488 ■ **Dayton-Forrester Associates,** Canoga Park, CA (213) 701-0127 ■ **Pro Com Sales,** Elk Grove, IL (312) 860-1028 ■ **The Tricom Group,** Hicksville, NY (516) 681-1222 ■ **MIDCOM COMMUNICATIONS, INC.,** Southfield, MI (313) 353-5696 ■ **Computer Dataco Inc.,** Kansas City, MO (816) 221-1212 ■ **Peripheral Business Systems,** Kirkland, WA (206) 823-6661 ■ **Peripheral Equipment,** Pleasant Grove, UT (801) 785-5009 ■ **Dynamic Systems North West,** Mukilteo, WA (206) 745-5311 ■ **RC Data Inc.,** San Jose, CA (408) 946-3800 ■ **Trans Alaska Data Systems, Inc.,** Anchorage, AK (907) 561-1776 **Andor Systems,** Ontario, Canada (416) 746-2775

a division of C3 Inc.

MPC

Micro Products Company
Rte. 634 & Acacia Lane
Post Office Box 198
Sterling, Virginia
Telex:
Phone: (703)

Barry Morgan-C3 Inc
Berkshire, England

Emulator Terminals Must Overcome Many Compatibility Issues

An alternative to captive markets is provided by units that duplicate performance at lower cost or with added features. Some even replace more than one product.

by **Robert J. Pryor,**
Visual Technology Incorporated

No clear set of standards permits terminals to be interchanged among vendors. That's because terminal vendors' control-code sets vary, as well as the specific ways the terminals respond to those control codes. The ANSI X3.64 standard, which was intended to alleviate the problem, was completed in 1979. Although it was praised as the coming industry standard, only a miniscule number of the terminals installed today actually follow this standard. As a result, when de facto standards develop because of a select number of commercial successes, they create a viable market for emulator terminals.

COMPATIBILITY

Compatibility issues in the computer industry seem to be all-pervasive and never-ending—which is only to be expected in a dynamic industry that's experienced staggering growth rates in a short time frame. Compatibility issues will continue to be of concern to the systems integrator, who is ultimately responsible for ensuring that all the pieces of the system function together in all their targeted environments.

The key items the systems integrator must be concerned with in the selection of an emulator CRT terminal include:

- **Compatibility.** The terminal should operate as identically as possible to the one it is emulating, for all known potential environments.
- **Cost.** Obviously the terminal should cost less than the one it is emulating, or some other compelling reasons will be needed to purchase it.
- **Packaging.** Superior ergonomics can provide a strong motivation to use an emulator terminal. On the other hand, a look-alike emulator terminal may be seen as a plus.
- **Additional Functions.** An emulator terminal may provide additional functions not on the original. These can be a plus if they are especially useful, a minus if they cause incompatibility—or, as is usually the case—insignificant.

Today a variety of choices are available to the systems integrator without a sacrifice of compatibility to existing applications. Emulator terminals provide unique advantages that allow the integrator to tailor the system to his requirements and remain competitive in an increasingly competitive environment. It's important to start with a clear goal of improving the total system—on the basis of cost, functions, ergonomics, or all of these. Then the peripheral and other components can be chosen to meet the intended goal best.

Following this approach can provide the systems integrator with a uniquely competitive product offering if he chooses wisely. Emulator products in general should be chosen carefully after thorough evaluation to ensure that the total system will not only be competitive, but successful as well.

EMULATION TERMINALS

Most commercially successful code structures were initially established by the major original vendors of CRTs, whether independents like Lear Siegler, Applied Digital Data Systems (ADDS), and Hazeltine, or mainframe vendors like IBM and Digital Equipment Corp. Generally, emulations of these companies' products could be made and sold, at a lower cost, into the same markets. Ergonomic packaging and other features not available on the vendors' products could be made available in emulations, thereby offering a choice of products to an otherwise captive market.

Another factor driving the CRT market is that it became increasingly competitive, and market growth proved to be insufficient to sustain the growing number of vendors. Since new vendors, mostly with some off-shore manufacturing connections, were too late to the market to create their own code sets, they emulated someone else's. At the same time, in order to maintain their growth rates, existing independent vendors tried to take a share of competitors' markets by emulating their products.

The logical extension of all this would be to have a multiproduct emulator in one terminal, with more features at lower cost. That would provide a vendor with the most potential market penetration with a single product. The first such device, the Visual 200, was introduced in 1979. It emulated the four major products of its time: Digital's VT52, the Lear Siegler ADM 3A, Hazeltine 1500, and the ADDS Consul 520.

The systems integrator may choose to employ an emulator product for any or all of the reasons mentioned. As a point of reference, the installed base of dumb and smart terminals is split into two major and roughly equal segments: IBM 3270-compatible (bisynchronous, SNA) and ASCII. This article will deal with the latter type, which can be further segmented into captive, independent, and graphics markets.

CAPTIVE MARKETS

By far the most successful ASCII terminal produced by a computer vendor is Digital's VT100, introduced in 1979. By virtue of its tremendous commercial success, this product is the closest the ASCII industry has come to a standard. The interesting thing about this product is that its stated goal was to create a standard, but not for a proprietary code structure.

Essentially a dumb terminal (no local processing), the VT100 included smooth scroll, split screen, and double-size characters. While it supported the older VT52 code set, Digital announced that this product would also support the new ANSI X3.64 standard.

Independent vendors saw this as an opportunity to be part of the new wave, since the specifications were available to everyone. There was only one problem: terminals compatible with ANSI X3.64 wouldn't work perfectly in DEC environments. This situation points up the difficulty of multiple vendors trying to produce compatible products from documentation alone.

Ultimately, the VT100—rather than the standard it was based on—became the standard. This allowed vendors to create look-alike products at lower prices with additional features. Two of the most successful independents have been CIE Terminals (formerly C.Itoh U.S.A.) with its CIT-101, and Visual Technology Inc. with its Visual 100.

Both these vendors had products that functioned identically to a VT100 in DEC applications where even a small difference would have rendered the products useless. Both had look-alike keyboards so that operator retraining was unnecessary. C.Itoh carried the compatibility a step further by offering a look-alike cabinet. The only obvious differences between a CIT-101 and the VT100 were extra features offered by the C.Itoh product.

The Visual 100 offered a package that was clearly distinguishable from the VT100 by virtue of its tilt and swivel display capabilities. This package also permitted a 14-in. CRT to be installed, making the 132-column display more readable (a problem with the VT100 and subsequently with the CIT-101). The success of these two emulator products shows that there are different user preferences and, therefore, alternatives for the systems integrator even in a captive market.

INDEPENDENT MARKETS

Before the VT100 became a dominant force, several hundred thousand terminals were installed, a large number of them original-vendor products and the remainder from independents like ADDS, Lear Siegler, and Hazeltine. Over the years, the installed base of these vendors' products grew because of private labeling arrangements with large-volume users in the computer industry. As a result, the Consul 520 from ADDS, the ADM 3A from Lear Siegler, and the Hazeltine 1500 became de facto standards in their own right.

As applications were written around the control-code sets of those terminals, an increasing investment was being made in them, which made it very costly or impossible for the systems integrator to change to another code set, especially for existing systems where terminals were being added. In effect, a pseudocaptive market was created, although slowly at first, that had many of the characteristics of any mainframe or minicomputer vendor's market.

The same challenges and potential opportunities were created for outside vendors. As with the captive markets, the considerations of cost, features, and packaging were central for emulator terminals. Numerous products that emulated competitors' terminals were introduced by Televideo, a newcomer; Volker Craig of Canada; Microterm; Soroc; Datamedia; and even ADDS, Lear Siegler and Hazeltine.

As a result, the systems integrator's choices appeared to be greatly expanded. Being locked into a specific terminal seemed to be a thing of the past as manufacturers emulated each other's products in a quest for market share, and a whole new set of issues had to be dealt with in selection of a terminal vendor: product quality and reliability, vendor viabil-

PENPAD™ reads your writing as easy as ABC... ABC... ABC...

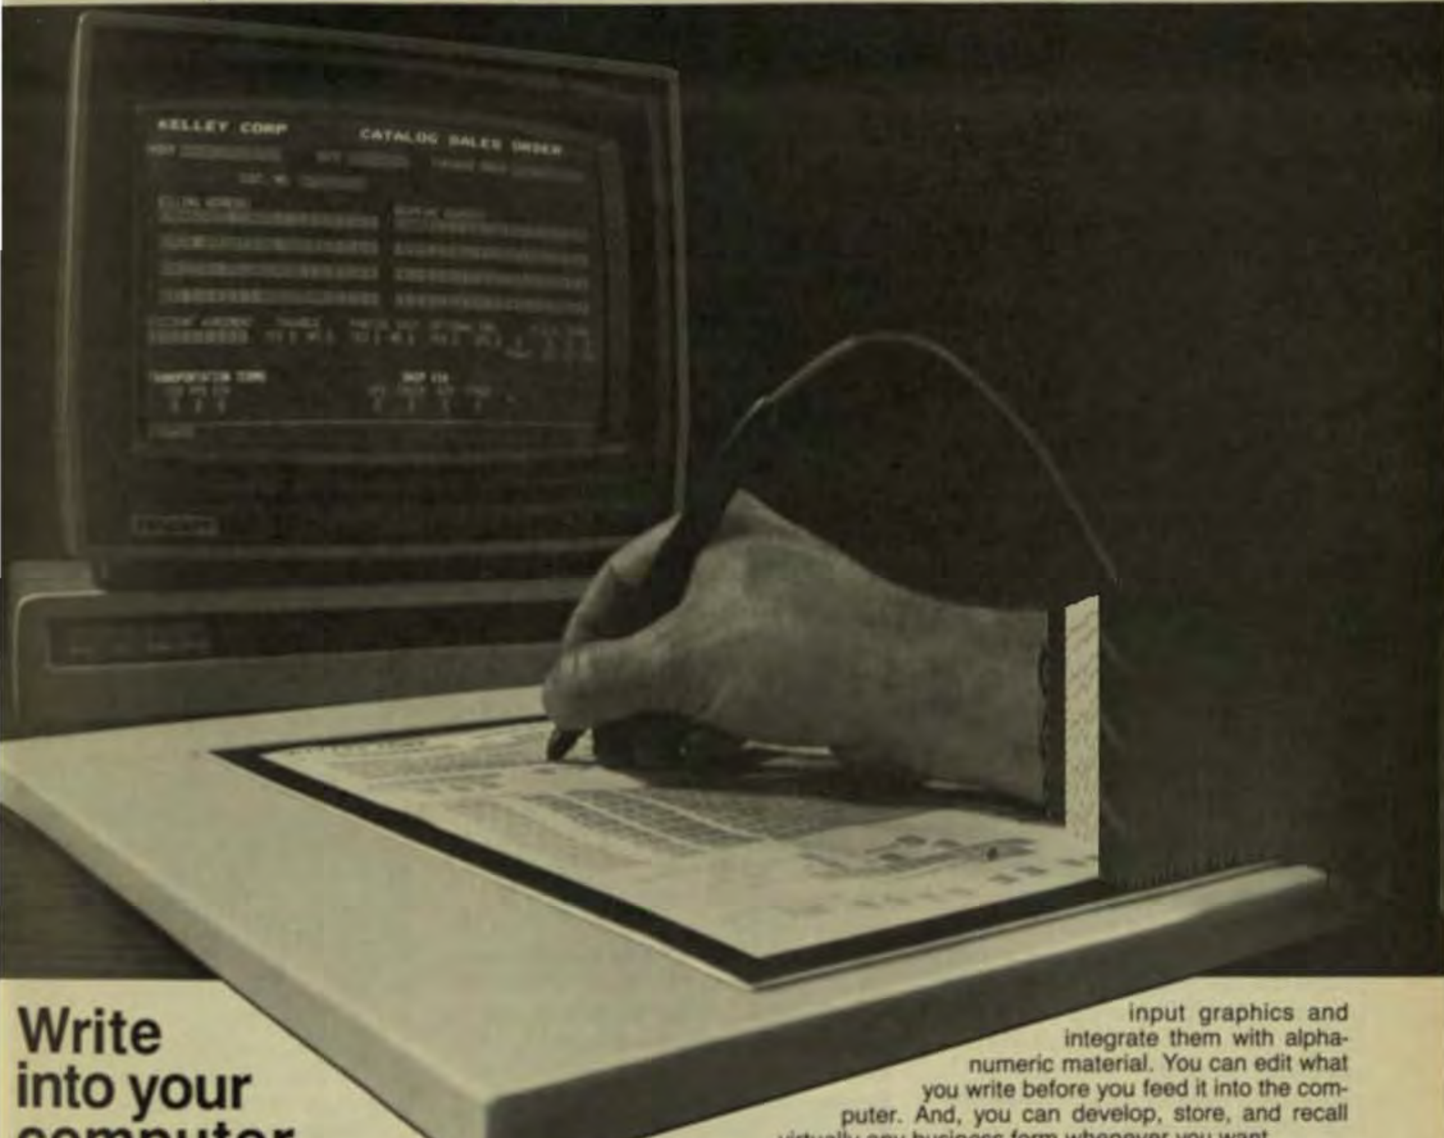

Write into your computer.

The most natural way to interact with a computer is now a practical reality. The PENPAD™ terminal offers an immediate alternative to keyboards or voice recognition equipment. As you write on a digitizing tablet, a high-speed microprocessor converts alphanumeric characters to ASCII code. And PENPAD™ accommodates a wide variety of printing styles as well as hand-drawn graphics.

Flexibility is foremost in its operation. You can program PENPAD™ to increase the efficiency of any business form by setting up a transaction area that reads handprinting and a command area that responds to the touch of a pen. You can access data bases or activate auxiliary devices. You can input graphics and integrate them with alphanumeric material. You can edit what you write before you feed it into the computer. And, you can develop, store, and recall virtually any business form whenever you want.

Unquestionably, Pencept's breakthrough in developing the technology of Dynamic Character Recognition (DCR) saves time and reduces the number of steps needed to create computer data. The person who fills out the form enters data into the computer at no additional cost. And, since the pen is the cursor, you can fill out the form in any order using only one hand. The form itself is a record of the computer transaction.

Any host computer with a standard, asynchronous RS-232C interface can use the PENPAD™ terminal. Personal PENPAD™ operates with popular microcomputers like the IBM PC.

Applications of this intelligent terminal that can read your writing are too many to list. Therefore, we invite OEM/System houses to explore this potential beyond the simple ABC's. . .

PENCEPT

Pencept, Inc. 39 Green St., Waltham, Massachusetts 02154 ☐ Telephone: 617/893-6390 ity, service, availability, delivery, price, and features.

As it turned out, few vendors were truly successful at emulating competitors' products. For while the emulation of a captive market terminal was fairly well-defined (having a limited number of identified applications and operating systems), the emulation of a general-purpose terminal wasn't.

By their very nature, general-purpose terminals are meant for applications and operating-system environments that aren't well defined or totally identified. Even the perennial dumb terminal has its quirks, usually undocumented; yet the user base learned to work around these anomalies, or even use them to advantage in their applications. Vendors who tried to emulate their older products with state-of-the-art technology ran into the same problems their imitators did.

The emulator terminal would usually work in a number of environments, but there were always applications that would show up an inconsistency. The number of permutations were enormous. The creation of a good emulator, let alone a multi-emulator terminal, thus required a bottom-up approach to terminal design.

But proper performance in an emulation environment isn't as simple as mere support for the correct code sequences. When the vendor adds functions that go beyond the capabilities of the emulated terminal, he must select code sequences carefully so that compatibility isn't affected.

Another current trend seemingly reverses the compatibility issue and places the burden on the application program. With the growth of microcomputers and standard application packages, many off-the-shelf programs have tried to support a multitude of different terminals via a menu selection. The goal is to make the applications hardware-independent, since they're being sold into a general mass market. In fact, this independence has been mandated by the same compatibility issues that the CRT vendor faces—only now it's the software vendor who's approaching the issue from the other side.

GRAPHICS MARKETS

The graphics market is growing rapidly but is still much smaller than the other markets discussed. In contrast to those, this area has had one predominant commercial success, one that encompasses both hardware and software. The Tektronix 4010 and 4014 terminals, combined with a package called Plot 10, have such a large share of their market segments that they can be thought of as de facto standards.

Other factors, however, also make this field ripe for emulation. Both those products are priced high enough that an emulator product could be produced at lower cost. Additionally, users of graphic terminals came to see the need for general-purpose computing and alphanumeric-terminal users simultaneously saw the need for graphics. The customer base ultimately wants to have both functions combined in a single unit.

Both these terminals are high-resolution monochrome graphics stations that are well suited for applications such as engineering design. They employ storage-tube technology, which doesn't permit selective erasure. Unlike the more common raster-scan technology used on CRT terminals and televisions, in which the screen is constantly refreshed (scanned), these terminals actually store the image once it has been written. Correction or modification of that image requires the entire screen to be erased and rewritten. This is a time-consuming procedure compared to a raster scan.

The two basic approaches that have been taken for the graphics market both provide a scaled-down version of the Tektronix products, but at a significantly lower cost that makes them ideal for use as preview terminals for initial design work that will be completed on a more expensive unit. These lower cost terminals can be used in a variety of general graphics environments.

In the first approach, which mainly targets the alphanumeric market segment that wants graphics, Digital Engineering and Selanar upgraded existing ASCII terminals with one or more printed-circuit cards customized to be integrated into the existing terminal hardware and display logic. Replacement of the CRT tube was also necessary to obtain the higher resolution available with some retrofit boards.

The advantage of this approach is that a user can upgrade existing equipment while maintaining his initial capital investment. Additionally, alphanumeric compatibility isn't in question, because the original terminal functions are maintained. The major questions left for the systems integrator involve retrofit issues (labor costs, reliability, multivendor support, warranties) and the quality of resolution required for the application.

The second approach targets both current Tektronix users who want a lower-cost solution and/or a general-purpose alphanumeric terminal, and alphanumeric users who want a lower cost solution and/or a general-purpose alphanumeric terminal emulator with 4010/4014 graphics capability.

The first such products, introduced by Visual Technology, provide scaled-down, medium-resolution Tektronix emulation as a standard feature. One version also provides selectable emulations of the DEC VT52, Lear Siegler ADM 3A, Hazeltine 1500, and Data General D200, and the other includes VT100 compatibility and ANSI X3.64.

The advantages of this approach are an integrated design, higher resolution than the retrofits, a single-vendor solution, and a lower total cost, depending on the alphanumeric terminal emulated. The major questions this approach raises for the systems integrator involve the issues of total cost. That is, will a new terminal be purchased if it's specifically targeted for retrofit? If there's no plan to add more terminals to the existing base, the retrofit approach could be the best answer. However, experience shows that often a new terminal is purchased specifically for retrofit, and in that case an integrated solution may be the best answer.

At any rate, the systems integrator who wants to incorporate graphics capability into his offering has a variety of choices available. Each must be weighed on several bases: resolution requirements, service requirements, cost, and packaging. ∎

Robert J. Pryor *is product marketing manager for Visual Technology Inc. and is responsible for the display terminal product line. He previously held positions with ADDS, including product manager and product planning manager.*

Introducing the Smart Terminal With a Dumb Price

$595.00
Factory Direct Offer

- Price/Performance Superiority
- Ergonomic Design Equals Convenience
- Z80-Based Intelligence
- Popular Industry Code Structure
- Keyboard Terminal Set-up
- Business Style Keypad
- More Features and Flexibility
- Unmatched 12-month Warranty

Now, there's an exciting new attraction in smart terminals. The "Challenger" family of ergonomically-designed Soroc terminals is the result of an innovative commitment to flexible performance and professional user comfort. These features combine to establish a new standard for price/performance comparisons and the ultimate in operator convenience and productivity.

Challenger's unique styling helps increase productivity; non-glare screens help reduce eyestrain—a problem affecting over 80%* of all visual display terminal operators; the detached low-profile keyboard with palm rest helps reduce typing fatigue.

A unique Softstart™ concept connects your Challenger almost effortlessly with its soft-key interface set-up capability. Absolutely no mechanical switch settings to adjust desired baud rate, word structure and display format—just type it in once, and your Challenger is ready to go; and remains in that condition unless you alter it. An industry code structure capability further assures straightforward interconnect.

Challenger comes standard with more features and more flexibility;
- programmable function keys
- intelligent flexibility with programmable operational characteristics
- a full range of built-in edit commands
- special graphics characters
- easy-to-use transmit and print functions
- tab and clear functions
- full feature video attributes
- cursor control flexibility
- display time out, self-test
- and much, much more.

Challenger is produced, tested and must pass stringent quality checks entirely in our Anaheim, California facilities, before it's packaged for shipment. This commitment accompanied by an unmatched 12-month limited warranty and support is backed by an expanding nationwide domestic service organization. These are sound reasons for you to check into all the facts before you make your selection. Once you do, we're confident you'll make the right decision.

SOROC TECHNOLOGY, INC.

*National Institute for Occupational Safety
Softstart TM is a trademark of Soroc.

161 Freedom Avenue, Anaheim, CA 92801 · (800) 854-0147 · (714) 992-2860

User-Defined Terminals Provide Flexibility for Changing Systems Needs

ost intelligent terminals incorporate the adaptibility users need to ons that were previously available only at the host computer.

ary M. Klinefelter, n Beland, and ucker Renshaw, ay

While terminals have been physically adapted to the user in many ways, functional flexibility to suit user preference and convenience has become available only recently. Some terminals now allow users to define a wide range of functions at the keyboard. Among the options that can be specified at the terminal are keyclick, keyboard auto repeat and rate, smooth-scroll rates, multiple-cursor character selection, character size, inverse video, margin columns, blinking or steady cursor, and screen saver. Most advanced terminals allow the user to define these items through a setup or selection-menu key rather than with hard-to-find switches.

COMMUNICATIONS AND DATA-FLOW CONTROL

The communications flexibility of adaptable ASCII asynchronous terminals is fairly extensive. For example, if two fully buffered, bidirectional, serial RS232C ports are available, as in Teleray Models 7 and 16, one port may be optionally equipped with either an RS422 or a current-loop interface. Both ports have independent keyboard selection, menu setup, or escape-sequence selection, and baud rates from 50 to 19,200. They also have a choice of even, odd, mark, space, or no parity checking, and full- or half-duplex modes.

Independent port selections are available for suspend-resume (XON-XOFF) and local echo (Fig 1). One port supports busy-ready with a selectable level of high or low. The availability of two fully configured ports allows the same terminal to interconnect to two networks, or a network and a peripheral, or a CPU and network, or a CPU and peripheral—thus giving the user much greater application flexibility.

APPLICATIONS ADAPTABILITY

As the applications of CRT terminals have mushroomed in recent years, the features designed into these terminals have been limited to the minimum required to support an application. While this trend has been justified on the grounds of low cost and ease of design, it has proved to be a false economy for users who find their applications changing and require new terminals for each application. Frequently, an application cannot be implemented or must be severely compromised because the features needed to support it don't exist in installed hardware.

Display-memory organization is one area of terminal design that has been extremely limited. Only in recent years has 132-column capability been provided—and frequently only in a rigid 132-column format. Little attention has been paid to other application-dependent aspects, such as variable line or page lengths.

This places the burden of maintaining file sense for these applications on the host computer or requires highly specialized terminals for each application. Frequently the task is impractical to perform in the host because the particular code that

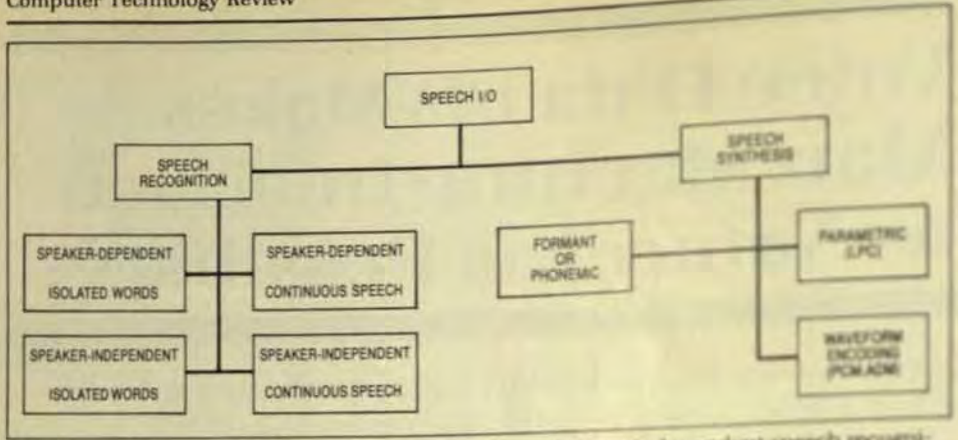

Fig 1 Speech input-output systems require speaker-dependent or -independent speech recognition, of either isolated words or continuous speech. The other part of the process is speech synthesis, which can often be very mechanical sounding (formant or phonemic systems). Linear-predictive coding (LPC) can be performed by LSI. With pulse-code modulation (PCM) and adaptive delta pulse-code modulation (ADM) the vocabulary is limited only by available memory

ments. There are three basic types of automatic speech-generation systems: speech synthesis, speech coding or voice response, and voice store-and-forward.

Speech-synthesis systems are based on the digital storage (magnetic RAM or ROM) of phonetic structures. The digital approach to speech synthesis permits creation of more than 300,000 words of working vocabulary—far more than the 50,000-word working vocabulary of everyday English. Speech-synthesis systems can be based on digital conversion of an actual human voice or on modeling of the vocal tract.

Speech-synthesis representations include text or formant synthesis (direct text to speech), parametric representations (analysis-synthesis techniques such as linear predictive coding), and waveform representations. Fig 3 shows the range of data rates for these various types of speech representations.

Formant or text synthesis is characterized by data rates of 600 to 800 bps and models the human voice by producing bands of resonant frequencies. Direct speech synthesis forms speech signals or responses directly from phonemes, the elements of speech. The physical structures that this type of synthesizer must accurately reflect include: an electrical analog of the human vocal tract, a program to specify the desired sound of the vocal-tract parameters, and the control interface of the vocal tract.

The major problem associated with synthesis by rule (constructive synthesis) is the proper encoding of phonemes. With the possible number of phonemes and emphasis commands, almost any phrase can be formed. The set of rules must consider factors like the garbling effects of word boundaries, stress variations, and pitch- and timing-contour problems.

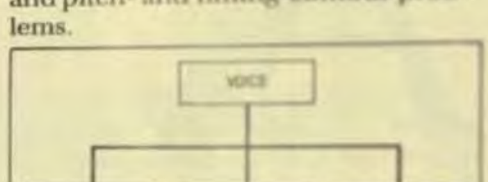

Fig 2 Voice technology can be divided generally into three areas

Text-to-phoneme conversion can compress speech data dramatically and hence enjoys advantages in unlimited vocabulary and low memory requirements. This process doesn't require speech input, but rather utilizes direct text-to-speech conversion or direct phoneme encoding. The major disadvantage is that the sound quality is poor and mechanical-sounding. Several alternative algorithms (allophones, diphones, demisyllables, and morphs) can be used to improve the speech quality, but extensive processing power and larger memory are needed to implement them. Moreover, the necessary algorithms are far from perfected.

Text-to-speech research at the Massachusetts Institute of Technology (M.I.T.) has been concentrating on the perceptual implications of waveform and spectral manipulations. Hence, detailed knowledge of the acoustic theory of speech production, phonetic theory, phonological theory, syntax, and semantics are required. Other research at M.I.T. includes the role of auditory modeling in systems for speech processing.

Parametric representations, including linear predictive coding (LPC) and partial auto-correlation
(PARCOR), have data rates of
10,000 bps and can trace the
tion to speech analysis. Spe
analysis techniques assume
parameters underlying the sp
model change slowly with ti
leads to a number of short-ti
sis techniques, such as Fou
sis and analysis by synthesis.

About half of the curren
synthesis-chip manufactur
the relatively simple LPC te
for estimating the paramete
speech signal (pitch and form
and to represent speech wit
bit-data-rate transmission a
age.

The basis for LPC is tha
sample can be approximat
ear combination of past spe
ples. Very accurate estimate
speech parameters can be pro
with relatively fast speeds o
tion. With LPC, fewer bits are
quired to produce each wor

Consequently, a compre
data on the order of 100 to 1
But the importance of linear p
tion lies in the accuracy with
the basic model applies to spe
The parameters of this mode
voiced/unvoiced classificati
pitch period for voiced spee
gain parameter, and the coeff
of the digital filter (all varyi
with time).

Unlike other parametric
ing techniques, LPC permits
plementation of an effective
synthesizer with available LS
gle chip. Rather than creat
from synthetic phonemes, LP
speech-synthesis systems ar
on digital conversion of spe
an actual human voice. Here
tical support structure is req
such as the speech data
development services offere
Texas Instruments and Hitac
offers good-quality speech b
cabulary is limited to stored
ters.

Waveform representation
ture data rates of 10,000 bps
(Fig 3). Here, speech is digit
sampling through various
encoding techniques and co
pressed to eliminate system
dundancy and silence interv
phase information is adjust
then the simplified wavefor
stored in a memory chip. To
ate a sound, it reverses the p
and recreates a sound that
the original.

With waveform encoding, vocabu-ary is limited only by the memory vailable to store digitized speech. peech quality is excellent; words nd phrases are digitized to retain he inflection of male/female/adult/ hild speech sources. However, the it rate is very high for high-quality ynthesis, so significant memory allo-ation is required.

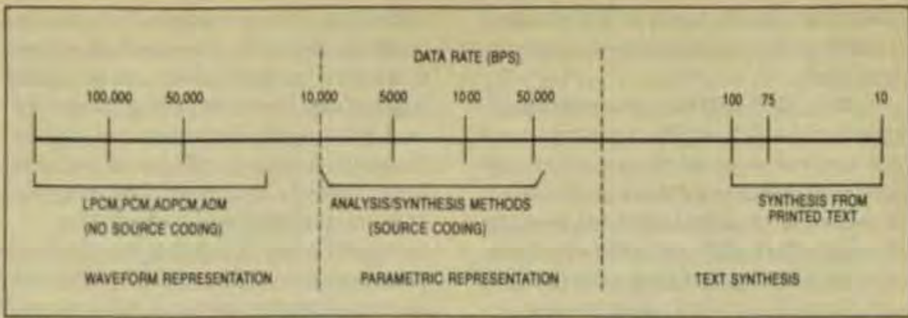

Fig 3 Speech synthesis by waveform generation features data rates of 10,000+ bps. Speech quality is excellent, but memory demands are high.

PEECH-SYNTHESIS CHIPS

lthough there are many different peech-synthesis chips available, the ain problem is to develop applica-ons software. There are many trade-ffs between speech quality and ata rate (memory efficiency), but for he most part, the specific appli-ation or market need is the determin-ng factor.

The traditional approach to peech synthesis involves several teps:

Define the size and contents of a vocabulary.

Choose a synthesis technique and vendor.

Send a vocabulary list to the vendor or hire a professionally trained speaker to record the vo-cabulary on tape and then send it to the vendor, who then digi-tizes, compresses, and fits the vo-cabulary into EPROM.

Listen to the resultant speech, edit it, and then place it into ROM.

The other way to develop speech-ynthesis applications is to use Signal echnology's ILS software package, hich must be run on a VAX minicom-uter and requires a lot of additional nowledge and acoustics/digitizing ardware. Fortunately, development stems such as Texas Instruments' ulti-AMPL and Centigram's Voice-are have become available over the ast year or so, so this route can be by-assed.

Most vendors are semiconductor ecialists, although TI does have an cellent staff of linguists at its re-onal technology centers. Other sup-ort methods include CP/M-based ftware from Speech Technology d National Semiconductor, or the in-ouse development services offered such companies as NEC Corp., gnetics/Philips, Speech Plus (former-Telesensory Systems), and Hitachi.

Another option is Computalker Consultants' software applications-development package that competes with the Signal Technology's ILS package. It will do waveform editing and linear-prediction analysis, and has a friendly command structure. As a general rule, manufacturers will have to provide the necessary support structure to penetrate and support speech-synthesis applications reliably. However, the following digital coding factors for speech synthesis must be considered:

- Information or bit rate required for speech quality.
- Complexity of coding and decoding schemes.
- Flexibility of the representation—potential modification of the vocabulary.

The speech-chip market is currently dominated by speech synthesizers, vocabulary read-only memories, encoding/demodulation chips, and digital signal processors.

Speech-chip manufacturers include National Semiconductor, AMI, Centigram, Epson, General Instrument, Hitachi, Mitsubishi, Motorola, OKI Semiconductor, NEC, Panasonic, Sharp Electronics, Texas Instruments, and Votrax. Most of these speech-chip makers are using CMOS technology.

Understandability and speech quality vary, depending upon the coding technique used. Understandability is the characteristic that allows a human listener to hear and differentiate specific words from other similar-sounding words. Speech quality is more subjective—it defines, in a rough way, the human-like quality of the synthetic voice.

Currently, LPC looks like the most popular synthesis technique and offers a good compromise between speech quality and memory requirements. At a rate of 1500 bps, the voice quality is good enough for listeners to distinguish the voices of various speakers.

On the low end of the spectrum, data rates of 600 to 1000 bps are typical of low-cost, self-contained speech synthesizers, which are ideally suited for games and toys (National's Digitalker II, for example). The lowest data rates belong to phoneme synthesizers (40 to 100 bps) like the Votrax SC-01.

Data rates above 2000 bps produce voice patterns identifiable as those of a particular speaker. Higher data rates, however, require increased memory capacity. Hitachi and Matsushita use the PARCOR technique with selectable data rates from 2400 to 9600 bps.

Memory capacity still dictates that lower bit rates be used where desirable—even though the availability of 64-Kbyte RAM, and 256-Kbyte RAM chips in the near future, will assist chip manufacturers who use higher digitizing rates, such as OKI Semiconductor's ADPCM line, National Semiconductor's Digitalker II, and Harris' CVSD chip. The SC-01 from Votrax has been tremendously popular, but the firm's new chip, the SC-02, can speak words from an unlimited vocabulary with 256 possible levels of inflection.

Toshiba and Sharp also have single-chip CMOS speech synthesizers. And NEC has an audio-synthesis chip that packs 64 Kbytes of ROM for up to 90 seconds of voice, music, or other sounds. In addition, General Instruments has developed an allophone-synthesis chip, the VSM2128, that uses LPC and enunciates 64 allophones.

The 32-bit TMS-320 signal-processing microcomputer chip from Texas Instruments may be the first VLSI device to allow designers to penetrate new applications that take advantage of digital signal-processing techniques to manipulate and interpret electrical signals such as those used in telecommunications and speech synthesis.

The TMS-320 has an operating speed of 5 MIPS (million instructions per second)—faster than many mainframe computers of the last decade. Based on a modified Harvard architecture, the TMS-320 provides separate spaces for program and data memory. Since it can fetch information from both memory spaces in parallel, it's able to fetch and execute instructions at the same time.

This approach, which allows crossovers between program and data memories, enhances the flexibility of the TMS-320 by letting the user perform program branches based on data computations. Chip throughput is increased because parallel operations can execute in a single 200-ns cycle.

The 32-bit TMS-320 (TMS-32010 and TMS-320M10) is fabricated with 3μ design rules by use of silicon-gate NMOS processing. Potential applications include:

- Translating speech into digital signals and compressing it for more efficient transmission over telephone lines.
- Digital filtering.
- Processing data for synthesized speech (data compression and LPC analysis).
- Analyzing data for speech recognition.
- Processing signals for a wide range of military equipment and sophisticated measurement systems.

SPEECH-SYNTHESIS BOARDS

Speech-synthesis evaluation boards are available from AMI, Epson, General Instrument, OKI Semiconductor, National Semiconductor, Texas Instruments, and Votrax. Speech-synthesizer boards available from several other vendors include Speech Technology's VR/S100, which is based on the SP-0250 chip originally designed by that firm and manufactured by General Instruments. The LPC-based VR/S100 has variable data rates—1200 to 1600 bps for low-quality speech and 2000 to 2200 bps for high-quality speech.

Speech Plus' (Telesensory Systems) Speech 1000/1100 speech-synthesis boards are Multibus boards with 50 to 200 seconds of speech capability, respectively. Computalker Consultants' CompuCorder can produce high-quality speech in any language and speak with any voice, male or female. It can also produce music and sound effects and is an IEEE-696/S-100-compatible circuit card. The CompuCorder is suitable for such applications as voice store-and-forward systems, paging systems, amateur radio repeaters, and computer-aided Instruction (CAI).

Centigram has a single-channel version of the Texas Instruments Module voice synthesizer, TIM II, which is form-factor- and connector-compatible with any Multibus single-board computer. It uses the TI VSP 5220 voice-synthesis chip at an average data rate of 1200 bps. The typical single-board computer can store 2.5 minutes of speech, while vocabulary storage on disk is virtually unlimited. LISA, Centigram's low-bit-rate synthesizer, is available for Digital Equipment Corp.'s Q-bus and Unibus. Talking boards are available from Analog Devices, Akerman Digital Systems, Applied Micro Technology, Data Voice, ICS Electronics, Hewlett-Packard, Master Specialities, Neutronics, Street Electronics, Votrax, and Votan.

Talking software is beginning to appear with packages like Peachtree Software's SPeachware, which is based on a proprietary digitizing technique (parametric waveform coding) developed by Centrigram. SPeachware is made possible through the SYBIL speech-synthesizer board, which will cost about $400 for an IBM PC. SPeachware is capable of storing about 8 hours of continuous speech on a 20-Mbyte disk.

Another speech-synthesis board is manufactured for the IBM PC—the Speech Master from Tecmar. It gives both a stored vocabulary (National Semiconductor's DT1050/Digitalker) which is limited to 143 words, and a phoneme synthesizer (Votrax's SC-01 speech synthesizer) with an unlimited vocabulary.

NEC's AR-100 Voice Output Terminal allows the computer to talk. It can remember up to 120 seconds of spoken words, messages, announcements, or instructions, and can repeat them over and over again as needed. Because messages are in digital form, they can be stored on floppy media at the host. The AR-100 utilizes adaptive differential pulse-code modulation for extremely high quality output.

SPEECH-TO-TEXT SYNTHESIS

Speech synthesizers are built into Kurzweil's Reading Machine, which has the ability to scan books and printed documents and read them aloud. Other vendors include the Maryland Computer Services and Perception Technology. The Prose 2000 system from Speech Plus is probably one of the best text-to-speech converters; it accurately and immediately converts keyboard inputs (using ASCII) to an intelligible voice speaking an unlimited high-quality spoken English vocabulary.

The incoming word is looked up in a lexicon of 1500 special words. If the search succeeds, the associated phoneme string is retrieved and passed directly to the allophonic process. If the lexicon search fails, the word is assigned a phoneme string and a stress pattern by a set of about 300 context-sensitive rules. By assigning allophones—variants of a specific phoneme based on surrounding speech sounds—the Prose 2000 is able to represent the required pronuncation accurately.

This process is sensitive to the syntactic structure and stress pattern of the sentence, as well as the speech rate and prosody mode selected. Prosodic rules determine the sound durations and fundamental frequency pattern of the words to be spoken. The resultant intonation contour gives the sentences the rhythm and melody of a human speaker. The rules are sensitive to the phonetic form, the part of speech of the word in the sentence, and the punctuation, as well as the prosody mode and speech rate selected by the user.

The parameter generator accepts the fully specified phonemes and produces an output that is a set of 19 time-varying speech parameters that use a constructive synthesis algorithm instead of pre-stored speech. The result is a good model of the human vocal tract and highly intelligible speech. These parameters control Speech Plus' proprietary speech synthesizer IC. The Prose 2000 is available in two forms—a single-board Multibus format and a standalone system.

Speech-to-text improvements are necessary if useful applications—such as the accessing of large databases that convey information by voice, and applications for the handicapped—are to become practical. The main problem with speech-to-text synthesis systems has been the lack of intelligible speech. Most of the current (first-generation) devices have poor quality, mechanical-sounding speech. That's passable for simple systems applications but higher quality speech is needed and will require additional linguistic principles and levels of detail.

The state-of-the-art techniques in speech-to-text synthesis developed by Dennis Klatt at M.I.T. include:

- Syntactic analysis of text—synthesis of good sentence rhythm and intonation.
- Morphemic decomposition of words so their pronunciation can be predicted better.
- Rules to guess the pronunciation of unexpected words.
- Rules to modify the pronunciation of words, depending on their sentence context.
- Rules to transform phonetic information into control parameters for a formant synthesizer.

The development work at M.I.T. was based on a PDP-11/60 minicomputer with limited memory and a high-speed array processor. This led to Digital Equipment Corp.'s development of a board-level product called DECtalk. This board uses a dual-processor architecture based on the Motorola 68000 microprocessor, which can address 16 Mbytes of memory and the Texas Instruments TM32010 DSP chip. The board is programmed in the high-level C language.

A 3-tiered software approach was used. The first level consists of preprocessing (numbers, end-of-sentence and nonspoken characters), lexicon (prefix and suffix stripping, exception words, phrase markers, lexical stress, common words, and user-loadable words), and rules library (numbers or amounts, letter-to-sound pronunciation, lexical stress, stress-modified pronunciation and sequence-modified pronunciation).

The first-level process searches for a text word in the lexicon and transmits the stored phonemic representation (if matched) directly to the second-level process. The lexicon—initially 1500 words (laboratory version) and being expanded to about 7000 words—would occupy 140 Kbytes of ROM. The letter-to-sound library is about 500 words. At the present level, DECtalk can pronounce about 20,000 words.

The second level of processing consists of phonemic recording, intonation, duration of word stress, target selection, phonemic smoothing, and parameter calculation. The third and last level is that of voice synthesis (through formant synthesis) and digital voice output to a D/A converter.

The Motorola 68000 handles first- and second-level processing, while the TI TM-32010 DSP chip handles the third-level processing functions. The system can be accessed through RS-232C ports from either a local terminal or host computer and has a variable speaking rate of 120 to 300 words per minute. Different voices and outputs—such as telephone and loudspeaker—can be chosen.

VOCODERS, VOICE RESPONSE, AND VOICE STORE-AND-FORWARD

Speech coding is distinguished from speech synthesis in that speech coding is derived from a human source (much like a tape recording playback, only from a digital storage/retrieval device). The most adaptable and life-like use of speech coding is voice response. Here, previous vocal inputs are combined for audio playback on command. Voice-response technology is similar to output methods that display output on a screen, only it outputs information in audible form.

In this case, prerecorded speech is strung together (concatenated) to form phrases. Technological advances now allow you to play back messages that have been digitized and stored in memory instead of motor-driven tape drives. The basic elements of a voice-response system include provisions for vocabulary storage, rules for forming messages from vocabulary elements, and a program to compose voice-response messages. These words or phrases are reconverted to an analog signal and replayed through a headset, speaker, or telephone.

Voice-response units (VRUs) may be configured as intelligent front-end processors or as standalone translation systems. These may be software-based disk storage or solid-state memory to achieve a voice response. Traditional access to voice-response systems has been accomplished via TouchTone telephones or with rotary-dialed telephones. Now, some systems are beginning to add speech recognition as an access mode. Votan's V5000, for example, integrates two major speech technologies—speaker-dependent recognition and speech generation (voice response, vocoding, and voice store-and-forward)—into one compact system. Typical applications for VRUs are order entry, credit authorization, inventory, and banking.

Voice store-and-forward involves real-time encoding, compression, and storage of speech messages for later retrieval. While voice reponse usually handles static messages, voice store-and-forward involves continuously changing messages that are only briefly stored. This technology is used primarily in the area of voice mail, in which voice messages are recorded, stored in disk memory, and recovered on demand. Voice mail is much like its counterpart, electronic mail; the vocal equivalent lets a speaker/caller leave a detailed message that is converted, stored, and available for subsequent recall, routing, and dispersal just like a piece of electronic mail.

Voice store-and-forward provides the technical basis for a very large, multi-user telephone-answering and message-distribution service. Busy signals or not, users are still able to "talk" to their associates, leaving a spoken message, of any length, that retains the speaker's original inflections, intonations, moods, and originality.

The typical voice-mail system will have "mail boxes" where voice messages from others are stored just like conventional paper mail. Voice-mail systems have the ability to deliver messages automatically to hundreds of people as simply as one. Further, messages can be created during the daytime and queued for subsequent unattended transmission at night.

The dictation of memos, letters, and other documents—even from remote locations like hotels—can be streamlined by such systems. For example, sales persons can call their home office and place a sales order directly with the computer, giving part numbers, quantity, and a customer number. The computer can respond with inventory figures for each of the

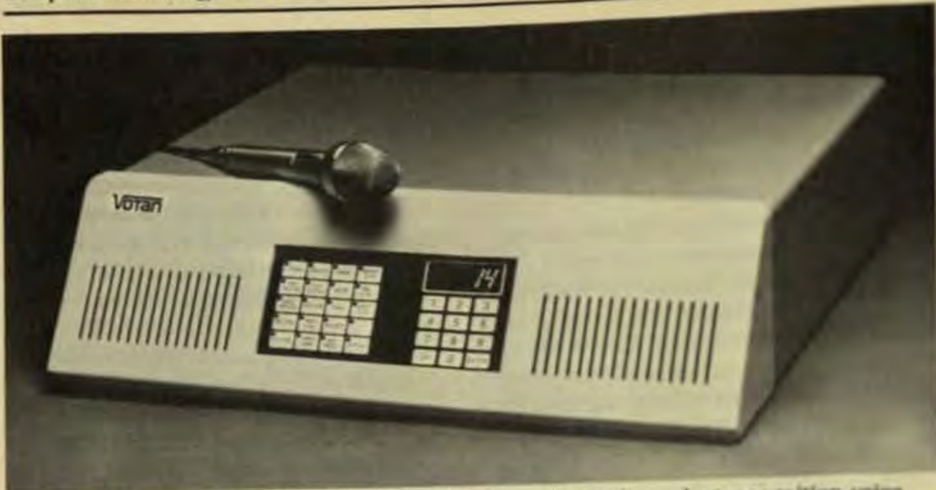

Fig 4 Votan's multifaceted V5000A includes speaker-dependent/independent recognition, voice response, and voice store-and-forward.

items and verify the shipping addresses and data. If the inventory is running low, the computer relays that information to the caller, who can then make a change in the order.

Vocoding is similar to voice store-and-forward. However, this application is primarily aimed at real-time telephone conversations in which speech is compressed to minimize the cost of transmission. Voice compression allows simultaneous transmission of several telephone signals and is the foundation that underlies cost-effective implementation of voice-generation technology.

The challenge is to compress more speech into the amount of memory space or data communications bandwidth. Thus, the goal is to keep removing redundant information, down to the barest minimum necessary for the ear to still recognize the speaker and understand the message.

The industry standard is 64 Kbps on telephone-line bandwidth. Pulse-code modulation and CVSD are the most popular techniques used by manufacturers of vocoders and voice store-and-forward systems. Typical data rates are 4.8, 9.6, and 16 Kbps.

Many companies—including ADP Computing Services, AT&T (American Bell), Datapoint, DEC, ECS Telecommunications, IBM, ITT, Northern Telecom, Rolm, Sydis, United Technologies Corp., Voicemail, VMX, and Wang—have entered the voice-mail market with products and services. The most commonly heard phrase is that voice-mail systems will attack real-time problems of business communications, or what is called "telephone tag."

VOICE-RESPONSE PRODUCTS

Other vendors of voice-response systems include AMNET, General Electric, Cognitronics, Digital Pathways, IBM, Infolink, ITT, Periphonics, Perception Technology, Rapidata, Sperry Univac, Speech Plus, and TSP.

Votan has two multifaceted development systems, the Votan Models V5000 and V4000. The V5000 (Fig 4) integrates speaker-dependent recognition and speech generation (voice response, vocoding, and voice store-and-forward). The V4000 combines voice response and real-time voice recording with voice store-and-forward. Voice-response capabilities allow the user to digitize, compress, store, and play back spoken words. Voice store-and-forward capabilities include digitization, compression, buffering, and transmission of spoken words to a mass-storage device or node in a communications network.

Both the V4000 and V5000 are user programmable, allowing the user to record or update his chosen phrases on-line without delay. The user is able to balance reproduction quality with storage requirements and select a data rate from 4.8 to 14.4 Kbps. At the top rate, voice reproduction is of high quality, equal to the sound of a human voice over a telephone line. At 9.6 Kbps, the reproduction is of communications quality, roughly equivalent to the sound of a human voice over a telephone line. At 4.8 Kbps, the speaker's identity, mood, and emphasis are clearly discernable.

Speech Plus' T/VIS voice-response telephone-interface system is a computer peripheral that brings together state-of-the-art electronic speech synthesis and telecommunications technologies to enable telephones everywhere to be used as simple computer terminals.

The voice management system used with the TI Professional Computer allows the user to create, listen to, and update voice-response messages. It can be used for voice prompting, dictation, or to record messages to be forwarded. Instead of using tones from a TouchTone telephone, this system combines speech processing, voice recognition, and telephone-management functions into a single system—based on the TMS-320 DSP chip—that fits into the desktop microcomputer.

One of the key features of this system is its ability to recognize and respond to voice commands. By using an innovative technique, it allows application programs designed for use with a keyboard to respond to spoken commands or inputs without modifications to the application. The internal option card includes the TMS-320 signal processor and a TMS-7000, as well as telephone, microphone, and headset. Speaker interfaces add voice store-and-forward capabilities that are equivalent to or better than good commercial telephones.

The telephone-management capabilities include software-driven dial-out of preselected numbers and automatic answering of incoming calls. TouchTone decoding is provided for incoming calls, and software selection of either tone or pulse dialing for outgoing calls.

Speech-recognition capabilities include: isolated and connected speech recognition of up to 50 (36 at 0.6 s nominal duration) user-trained utterances, fast download (0.1 s) from main computer memory of new word recognition vocabularies from menu-driven applications, and operational modes of enroll (train), update, and recognize. Thus, voice-response systems are beginning to be interactive—with the computer not only asking questions intelligibly, but responding to voice input.

The future of voice response and voice store-and-forward looks very promising, but the man-machine barrier will require multimedia communications in whatever mode is appropriate—voice, text, data, facsimile, graphics, or video. The key advances needed include:

- Voice annotations on existing text or documents.
- Facsimile images (optical representations of text, drawing, etc.).
- Hand-drawn sketches (drawn with pointing devices such as the optical mouse, digitizing bit pad, or light pen).
- Animation (frame-by-frame recording of speech, text, graphics, and video).

SINGLE-BOARD DECODERS

Some state-of-the-art voice-response systems use dual-tone multifrequency tone (DTMF) decoder techniques. Previous systems utilized discrete components—including separate filters for each tone. Current systems use hybrid LSI to achieve the entire decoder function on a single board consisting of microprocessor, ROM, DTMF decoder, modulation and channel controllers, and filters.

Sperry Univac's voice I/O module is an example of a step toward more efficient information transfer through voice processing. Utilizing the latest in analog- and digital-processing techniques, it allows a single, moderate-performance minicomputer to control and manage message-processing tasks on a large number of voice channels simultaneously. These tasks include voice store-and-forward, message editing, and dictation.

This module interfaces directly to the reconfigurable modular family (RMF) bus utilized in a number of computer systems. RAM provides storage for control information as well as temporary buffer space. Additionally, two independent processors reside on the module. The RMF processor handles the interface-protocol and data-transfer overhead, while the audio processor performs real-time audio data compression and manipulation.

The module provides A/D and D/A conversion; bit compression to 24 Kbps; placement and receipt of telephone calls while monitoring the line for dial tone, busy, ringing, speech, and silence; detection of DTMF characters; and pause removal, loudness control, and speech-rate control (speedup and slowdown).

Sperry Univac's Voice Information Processing (VIP) station is a computer-based voice message store-and-forward system for use with TouchTone-equipped telephones. It digitizes the analog voice message as it is entered and converts the message to an analog signal when it is played back. Station functions include listening to incoming messages, creating messages, obtaining the status of outgoing messages, and personalizing service.

Although currently implemented specifically for audio I/O in a record/-playback mode, ambitions for this module extend beyond voice-response applications to voice recognition. Here, the module could provide a powerful front-end processor for voice-recognition applications.

Interaction and communication with people will be improved with the new technological tools described, but the system must be designed with human factors in mind so that the interactions are natural. Automatic speech generation alone will not suffice—automatic speech recognition will also be necessary. ■

Speech recognition and related speech-input issues will be discussed in the Winter issue of COMPUTER TECHNOLOGY REVIEW.

One Board Says It All

© Computalker, 1983

IEEE-696 / S-100 BUS

TELEPHONE INTERFACE
- FCC Approved
- Initiate & Answer Phone Calls
- Touch-Tone® Generator & Decoder
- Detect Phone Line Status (busy, ringing, answered, etc.)

VOICE DIGITIZER
- Record Speech
- Speech Storage in RAM or Disk
- Speech Output
- Choice of 5 Bit Rates

CompuFone: COMPUTER INTERACTION BY TELEPHONE

Access and enter data to your computer from any telephone via the Touch-Tone® keys. Verify your transactions by recorded speech. So many exciting applications for the CompuFone are possible: voice mail, telephone answering machine, electronic funds transfer, telephone soliciting and data gathering, electronic Rolodex® data base query with voice output, order entry and inventory control, telephone banking—and that's just the beginning.

COMPUTALKER, 1730 21st St., Santa Monica, CA 90404
Call us at (213) 828-6546

Optical Mouse Provides Higher Quality Cursor Control

by **David Pruner**, and **Pamela Wilkey**, Summagraphics Corp.

The mouse is a hand-held pointing device that controls cursor movement to make computer control faster and easier. Named for its physical resemblance to the real thing, it reports relative changes in cursor location. The original mechanical design has been improved to increase the advantage of mice over other peripheral devices. Optical mice even have on-board intelligence that increases their reliability and ease of use for a growing range of applications.

The mouse operates by sending its horizontal (X) and vertical (Y) movements to the host to update the position of the CRT cursor; if the mouse moves down and to the right, the cursor on the screen will move in the same direction. When the cursor points to an on-screen menu selection or pixel, a button on the mouse can be pushed to alert the host software that you're at the desired screen location. Depending on the application, this signal will be decoded as a menu selection or as part of a geometric shape.

Before this new technology, computer users who wished to reduce their editing and menu selection time had a limited number of options. Light pens or touchscreens were quick, but tended to cause fatigue and soil the screen. Trackballs took up little desk space, but were slow and awkward in comparison to mice. And in graphics-intensive applications, an operator couldn't locate and hold the location of a single pixel and press a key on the keyboard at the same time.

Digitizers, although more precise than mice, were (until recently) relatively expensive and required a large amount of desk space. Joysticks, similar to trackballs, require a great deal of practice to manage any type of accuracy in positioning and can't provide the motion feedback provided by a mouse.

Mice out-perform all of these other devices. They are fast and accurate, and allow the user's hand to stay in the same position relative to the buttons at all times.

The Mechanical Mouse

The first mouse was mechanical, developed by Douglas Englebart at the Stanford Research Institute in the late 1960s and patented in 1970. It was a simple, round-edged box that rolled on two wheels placed at right angles to one another.

A variation of this original mouse was developed for Xerox in 1972 by Jack Hawley. This mouse was digital, used less complex software, had quieter wheels and eliminated the costly A/D converter.

A 1975 version traveled on a ¾-in. stainless-steel ball bearing that eliminated the drag of side-slipping wheels. Then Hawley came up with a mouse that worked on the same principle as a trackball, only upside down. As the ball rolls across a desk, it transfers its X and Y movements to two small cylindrical drums resting on the ball at right angles to one another. These cylinders rotate in proportion to the ball's forward and sideways motion.

Connected to the ends of these cylinders were small code wheels, alternately coated with conductive and nonconductive segments. The code wheels delivered electrical pulses that were decoded and sent to the computer in digital form.

Praised for its workmanship and tracking ability, this mouse gained instant popularity and Xerox began producing it internally. The only complaints were that this design tended to pick up surface dust, jamming precision parts that were difficult to repair.

Another type of mouse is electromechanical in nature, relying on an optical encoding system to transmit changes in position. This mouse, which features exceedingly high resolution, was developed by John Purbrick for Kenetronics.

Other companies that manufacture mechanical mice are Product Associates Inc. and 3G Company. PAI's mouse, fashioned after Englebart's original design, features potentiometer wheels, an A/D converter and an RS-232 interface. 3G's mouse is different in that its wheels are made from foam rubber instead of steel to eliminate the surface-gripping problems found in mice with traditional wheels. Also, foam wheels are less expensive, less fragile and less susceptible to dust.

The Optical Mouse

The solid-state optical mouse is a relatively new device, developed by Steve Kirsch at Mouse Systems Corp. This mouse travels on a felt pad and uses a light-emitting diode (LED) to track its position by decoding reflected light from a sheet of aluminum with colored lines (Fig.

USI International, another optical-mouse manufacturer, produces a microprocessor-controlled optical mouse that uses a black-and-white grid. This type uses optical technology to determine distance and a mechanical technique to determine direction.

The Summagraphics optical mouse has an RS-232 interface and uses an Intel 8051 microprocessor to provide the processing power and program memory needed for sophisticated data processing and bidirectional interfacing to the host.

Optical Technology

An optical system senses movement by bouncing light off a reflected surface. The reflected light is received by a four-quadrant photodetector and processed by the mouse's microprocessor, using sophisticated algorithms to determine distance and direction of movement. The mouse is moved on top of a special grid surface, known as a mouse pad. The mouse pad consists of a highly reflective material, printed with grid lines of nonreflecting ink.

As the mouse is moved across its pad, an LED inside the mouse focuses a beam of light on a particular spot on the pad. When the mouse passes over an absorbing grid line, the photodetector records the absence of light. Conversely, when the mouse is over a reflective area, the light will be reflected back to the photodetector. In this manner, the mouse senses motion in the form of flashes of light as it travels over reflective and nonreflective areas on the pad.

To determine the direction of movement and for higher resolution than the number of grid lines per inch, the mouse uses the four quadrants of its photodetector

The LED within the mouse shines light on the mouse pad; this light is either absorbed, or reflected back inside the mouse, focused through a lens, and imaged onto the four quadrants of the photodetector.

If the mouse is moved slightly, two quadrants report the presence of light and two will report no light, since the mouse is positioned half on a nonreflecting line and half on a reflecting line. If the mouse continues to travel in the same direction, it passes directly over a reflecting grid line and light will be imaged on all four quadrants of the photodetector. To determine the direction of movement, algorithms in the microprocessor look for specific sequences of patterns that identify the direction of movement relative to the grid.

As the mouse moves from one nonreflecting grid line to the next, it can sense four transitions in the patterns of light detected on the four quadrants of the photodetector. These four transitions, multiplied by the 25 grid lines in each inch of the mouse pad provide the mouse with its resolution of 100 lines per inch.

Optical mice are inherently more reliable than mechanical mice because they have no moving parts to break down or wear out. Also, unlike mechanical mice, optical mice don't require shaft encoders or potentiometers to detect movement. Such tiny mechanisms are prone to failure and can cause the mouse to act erratically when dust, dirt, and eraser shavings are picked up by the ball as it rolls across a desktop.

Another advantage of the optical mouse is its ability to work with an on-board microprocessor that enables it to interact with the host to change modes of operation and process data before it's sent to the host. In contrast, mechanical mice are considered a dumb peripheral because they continually send data to the host, where special circuits digitize the data for use.

Since the host has no way to communicate with the mouse, the host must constantly service interrupts from the mouse, convert the data to digital form, and then decide whether to use or ignore it. An optical mouse allows bi-directional communication, which eliminates hardware and host-processor time by sending the data to the host via a standard serial interface.

Sophisticated optical mice have several modes of operation controlled from the host via remote command. In incremental streaming, for example, the mouse sends data in a stream, at a rate selectable by the host, whenever the mouse is in motion. In this mode, the mouse's microprocessor outputs data only when it senses movement of at least one unit of resolution (0.01 in.). In this way, the mouse only interrupts the host when moved.

Prompt mode, an alternate method of data output, is a programmable I/O mode in which the mouse responds with a single report when polled by the host. Thus, the rate of data from the mouse can equal the rate at which the host can process it. In minicomputers where the host CPU is busy processing data, this mode allows responsive interaction without overloading the host.

A frequently encountered problem experienced by an operator when using a mouse is jitter, which causes the screen cursor to oscillate from one pixel to another on the display screen when the mouse is on the edge between two states. In a mechanical mouse, this happens when the mouse is positioned so that the shaft encoder's contacts are midway between the open and closed state. This phenomenon also occurs in optical mice when a photodetector quadrant can't decide whether it sees light.

In sophisticated optical mice, algorithms performed by the microprocessor can determine statistically whether the state changes of each photodetector quadrant are the random output of a confused photodetector or actual movement. This antiteasing algorithm is constantly acting on the data for the optical system of the mouse, so it essentially eliminates this annoying problem.

Another advantage provided by an optical mouse is that special firmware in the microprocessor can deter software piracy. For instance, the serial numbers of individual mice can be encoded into firmware and passed to the host during initialization, when the computer is configuring the mouse remotely, so that the system won't operate unless the serial number of the mouse matches the number encoded into the host. In this situation, any number of copies can be made of a software disk, but they will only work with one particular mouse.

At present, a relatively small number of mice are being used as cursor-control devices compared to the numbers forecast for the coming years. These pioneer mice are primarily of the mechanical type, but optical mice can be expected to grow in popularity as mice find increasing use in CRTs, graphics workstations, and personal computers, because of their reliability, sophistication, and lower cost.

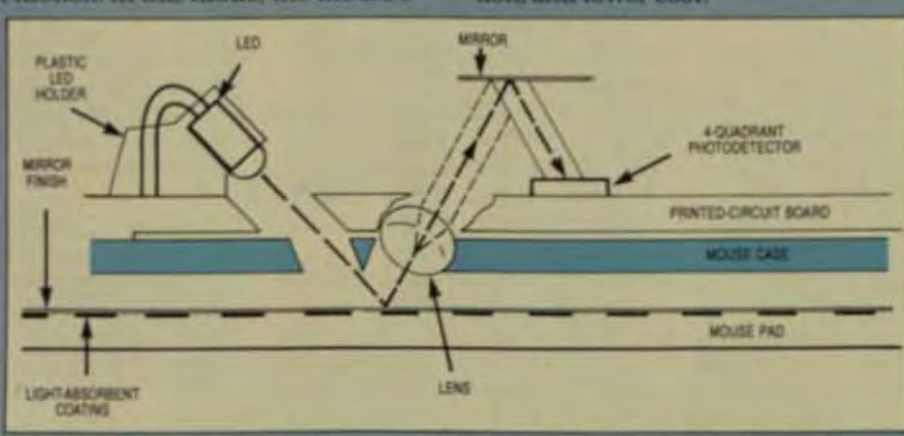

Fig The optical mouse travels over a mouse pad and is tracked by an internal microprocessor, which decodes reflected light from the pad. LEDs within the mouse shine light on the striped surface of the pad. Either this light is absorbed, or it's reflected into a 4-quadrant sensing device whose output is monitored and decoded for transmission to the host.

THE CENTRONICS PRINTSTATION 350.

THIS IS WHAT PRINTSTATION TECHNOLOGY IS ALL ABOUT.

Since its introduction in late 1981, the innovative Centronics technology behind the Printstation 350 Series has received OEM praise for its paper handling and reliability. With new Printstation family additions, we now offer new capabilities and higher speeds. Now, more than ever, the Printstation 350 family will provide OEMs with the flexibility to meet all their printing needs. Bar code printing. Large characters. Color. Graphics. More Multipass fonts. More speeds, from 50 cps (multipass) to over 400 cps (10 cpi). And more efficiency with an outstanding new breakthrough: a 1-, 2- or 3-bin automatic sheet and envelope feeder option.

Add these new capabilities to proven Printstation 350 innovations such as true multi-function paper-handling, and family design with 80% parts commonality—and you have the ideal OEM printer choice for all three information processing categories.

DATA PROCESSING.
Printstation 350 means exceptional throughput—approaching line printer speeds in DP applications such as: ▫ Program listings ▫ Business reports ▫ Data logging ▫ Spread sheets . . . using full 6-part, single sheet or fan-folded forms . . . and capable of operating at 100% duty cycle.

BUSINESS PROCESSING.
Whether in an office or on a loading dock, whatever a business needs, a Printstation 350 will print: ▫ Bar code tickets ▫ Mailing labels ▫ Insurance forms ▫ Purchase orders ▫ Sales charts & graphs ▫ Invoices . . . on business cut sheet, instant tear-off and sprocket-feed forms . . . with graphics . . . and without afterthought options.

WORD PROCESSING.
A Printstation 350 means complete job flexibility with a choice of fixed pitch or proportional fonts for: ▫ Business correspondence ▫ Office memos ▫ Proposals ▫ Personalized and form letters ▫ Envelope addressing.

And with our new automatic sheet/envelope feeder you can maximize operator productivity at an amazingly low cost.

Attractive and quiet enough for every office but right at home in a warehouse, teller station or shipping department. — That's Printstation 350. From Centronics—the first choice of OEMs worldwide. For a copy of our new Printstation 350 brochure, write Centronics Data Computer Corp., One Wall Street, Hudson, N.H. 03051. Tel. (603)883-0111

**CENTRONICS®
PRINTSTATIONS**

Analyze Specs Carefully To Find the Best Printer For Your Application

Careful attention to specifications for speed and print quality can help you tell the difference between products that look the same on paper.

by **Arthur K. Hyzer,**
General Electric Printer Products

Dot-matrix printers have achieved a level of sophistication, capability, and proliferation that can confound even a knowledgeable evaluator. Even after the field of choice is narrowed to two or three, you have to analyze the data sheets and functional specifications, which are written to put a product's best foot forward—to sell printers without actually misrepresenting them. Speed and throughput are two areas in which specsmanship may lead the unwary buyer astray.

When products are compared, it's useful to know what each one does or doesn't do to meet your needs. To make the best choice, you must be able to relate specified parameters that may not be provided clearly in the product literature. If you can "read between the lines," you may be able to identify the differences between products that look the same on paper.

SPEED CONSIDERATIONS

In the case of speed, the main factors that influence selection are the volume of printout required by an application and the waiting-time tolerance. Serial matrix print speed is specified in characters per second (cps). Typically this is the maximum print rate over a single line of print—the rate at which contiguous characters, including blank spaces, are printed. Yet columns per second would describe print speed better. This discrepancy is only the beginning of speed spec manipulation.

Some specs provide an additional piece of speed information called throughput. Throughput is measured in lines per minute (lpm) and is better than cps as an indication of a printer's capacity for output volume. Data sheets typically show throughput as a function of several fixed line lengths. Serial printers produce more lines per minute as the line length is shortened. Design parameters such as line-feed time, bidirectional printing, and logic-seeking intelligence to control printhead motion also affect throughput. The addition of any or all of these features is designed to reduce the percentage of nonprinting time.

For example, increasing performance of the paper drive system for a faster line feed yields an increase in lpm. Table 1 compares lpm throughput for three printers with different line-feed times. It shows the performance for three different line lengths. For a 100-cps printer, the comparison indicates some useful application sensitivity. Financial spreadsheets on 132-column computer paper gain only a 7% throughput advantage from the fastest line feed. However, business letters having a 40-column line length would be produced 18% quicker by printer C. Table 2 shows the throughput impact of faster line feed on printers with a 200-cps print rate.

The throughput increase for printer C is nearly 37% for the 40-column letter and over 11% for a 132-column spreadsheet. As print-speed increases, line-feed speed has

Selecting a Plug-Compatible Printer

by **Joe Bailey,** Printer Systems Corp.

The printer can be the most critical component in a successfully integrated system because the hard-copy output is often the end product on which the entire system is judged. If the printer creates a bottleneck that reduces throughput, or if the printout looks unprofessional, then the genius of the software and the cost-saving efficiency of the CPU will likely be overlooked.

Another reason that it's vital to keep up to date on plug-compatible printers is that competition is growing as such vendors gain acceptance in the micro-, minicomputer, and mainframe community. This acceptance is the result of several important factors:

- Plug-compatible vendors have been more flexible in introducing state-of-the-art equipment.
- Prices have been significantly reduced to compete with name-brand printers.
- Plug-compatible vendors have had much shorter lead times on deliveries.
- Nationwide service organizations have emerged, increasing end-user confidence in plug-compatible vendors.

These benefits can be realized by the end user who fully understands plug-compatibility as it relates to his specific printer application. For virtually every computer system on the market today, the end user can choose from a variety of sources when making a printer selection (Fig 1).

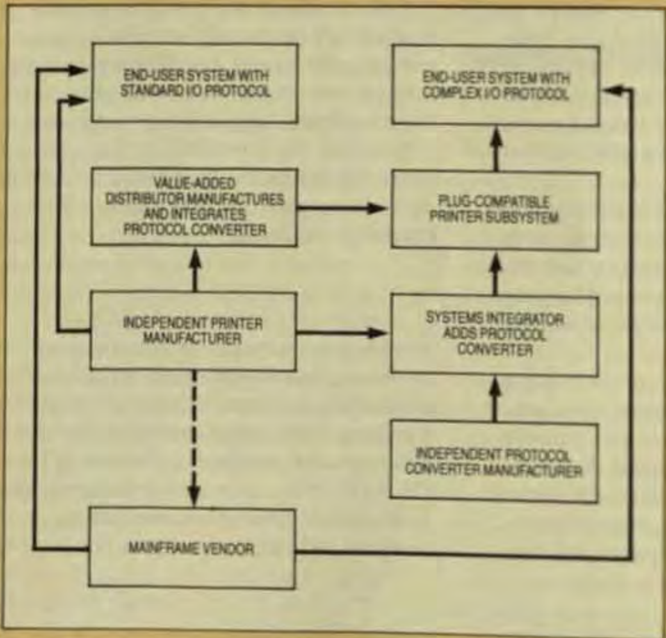

Fig 1 A variety of sources can supply the end user with the right printer for virtually any computer system or application.

In general applications, when performance isn't an important consideration, the user can simply shop around for the best price or the quickest delivery. In more specialized applications—such as word processing, graphics, or bar-code printing—if the system vendor doesn't offer a printer solution, the user is forced to look for alternate sources.

The first prerequisite, of course, is that the printer must be fully compatible with the user's system.

Defining Compatibility

The term compatibility has caused a great deal of confusion among end users, who tend to shy away from plug-compatible vendors and remain locked into a single source. Actually, to be plug-compatible with a given system, a printer must perform the required output function without system modification and be capable of replacing the device offered by the system manufacturer. In most cases the plug-compatible vendor must modify the interface of an independent manufacturer's printer to "look like" the printer it replaces.

For example, to the host system, the PSC M200-3270A looks like the IBM 3287. Either printer can be connected remotely to an IBM mainframe in a 3270 cluster network. By simply unplugging the 3287 printer and substituting the PSC, the user can more than double his print speed while lowering his costs. The M200-3270A is a version of the Dataproducts M200, modified by Printer Systems Corp. to look like the IBM 3287 and sold to IBM users.

Printer modifications must address both hardware- and software-interface requirements. The hardware interface defines the specifications of the connectors, cables and pinouts. The software interface can be thought of as the rule that governs the flow of data between devices.

Hardware Interface

Although a few I/O connectors have emerged as de facto industry standards, most computer manufacturers use their own exclusive versions. As a result, there are nearly as many connectors available as there are computer systems. Independent printer manufacturers usually use either an RS232C, a Dataproducts parallel, or a Centronics parallel connector.

The connector usually defines the cable requirements. For example, the twin-axial connector used on the IBM System 34 and 38 is used exclusively with twin-axial cable, while the 25-pin RS232C connector, used primarily in serial applications, can also be used in some parallel applications. Data can be transmitted in either serial or parallel mode. Parallel transmission uses multiple data lines, one for each bit of a data character and a control signal. Serial transmission uses a single data line, sending all data bits sequentially.

Parallel mode is faster because entire characters are sent or received at one time. In serial mode only one bit of each character is transmitted in the same time interval. Because of the multiple-wire cable requirement, parallel transmission is usually limited to relatively short distances.

Once the cables and connectors have been chosen, the pinout must be determined. Cables can be spliced to their connectors in a variety of ways, with each pin performing a different function. There may be as many as 50 possible locations for a given pin, but the computer and printer must agree on the pin locations for proper communications to take place.

If the printer is testing pin number 3 for a clear-to-send signal, then the computer obviously cannot send the signal through any other pin. If the pins on the printer end of the cable do not agree with the pins on the computer end of the cable, the pins must be relocated in the connector.

In some cases, simply modifying the hardware inter-
e of an independently manufactured printer enables
printer to be fully compatible with a computer system.
nnecting a printer to the IBM Personal Computer, for ex-
ple, simply requires the proper cable with a Centronics
rallel plug at one end and a 25-pin connector at the
er. In most cases, however, the software interface must
o be modified.

oftware Interfacing

e software interface defines the formatting sequence for
ta transmission, the manner in which electrical signals
translated into commands, and the error-recognition
d -recovery processes.

As with hardware interfaces, a few protocols have
erged as de facto industry standards. Some Honeywell
d Burroughs systems can communicate with remote
inters by using IBM 3780 bisync (binary synchronous), a
mmon batch-oriented protocol. A common interactive
otocol is Burroughs' poll/select, which involves the ad-
essing of each terminal on a line, one after the other.

In a typical application, the computer polls the first
minal, which responds NAK if it has nothing to trans-
t, or ACK followed by its message if it has one to trans-
t. The computer then polls the next printer in sequence.
ny terminal doesn't respond, the computer will time
t and proceed to poll the next terminal.

Polling takes place constantly in round-robin fashion.
tbound messages from the computer are transmitted to
h printer when it is due to be polled. Typically, this pro-
ol is used to verify the correct receipt of messages or to
quest retransmission. Polling is only possible if the print-
is smart enough to have an address, and able to respond
en it reads its address in a message received on-line.

e Controller's Role

ommunications controller or protocol adapter is a de-
e that performs handshaking functions between the
st and the printer. It provides adapters for different hard-
re interfaces and translates the protocols in the soft-
re interface. The controller can be located inside the
st computer, inside the printer, or in a separate enclo-
re external to both devices.

The controller is internal to the host only on systems
t provide for individual control cards for each I/O de-
e. Such systems include the Digital Equipment Corp.
P and VAX, the Data General Nova and Eclipse, and the
M Series I. A printer control card manufactured by the
g-compatible vendor simply replaces the system
dor's control card.

Many printers have unused internal card slots that
house controller boards. If the controller is mounted
this slot, the external I/O port on the printer may also
e to be modified to meet the host computer's cable
cifications. For instance, a twin-axial connector would
added for connection to an IBM System 34/38.

Controllers mounted in an independent enclosure
e some drawbacks but are generally the most flexible.
ese controllers provide easy access to switch settings
baud rate, self-test functions, memory dumping, and
n-on selection. External boxes can even provide for func-
ns not previously available on a system. For example,
the PSC 3780 controller gives IBM users the ability to send
job commands from remote locations, and eliminates the
need for a keyboard or CRT in most applications.

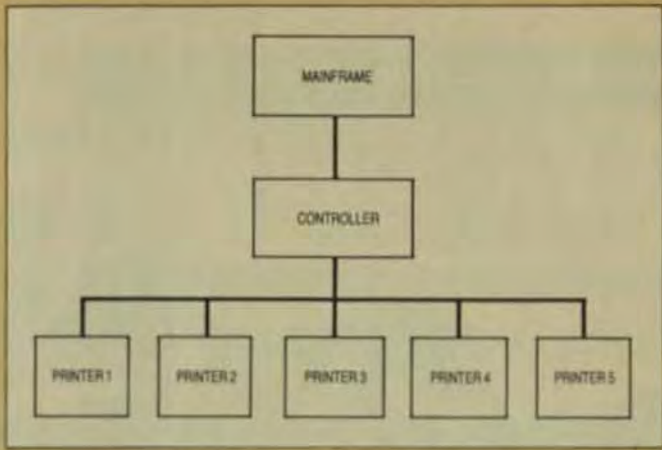

Fig 2 Some external controllers provide multiple printer ports, enabling the user to expand the system's I/O capability without sacrificing mainframe slots.

In some cases, external controllers provide multiple
printer ports, enabling the user to expand I/O capability
without sacrificing mainframe ports (Fig 2). Another ap-
proach is to daisychain, or cable-through, controllers and
printers, each controller providing one input line and two
output lines—one for the printer and one for the next con-
troller in the chain.

Choosing a Printer

As printer technology improves, the number of products offered by independent printer manufacturers increases at an accelerated rate. The plug-compatible vendor must constantly monitor new product introductions and modify his product offerings to fulfill his customers' printer application needs.

Printers are chosen with specific end-user applications in mind. The printer may offer faster throughput, quieter operation, better print quality, or lower cost. Reliability, serviceability, and availability are considered heavily. A printer may be chosen for specific vertical markets, such as hospitals, banks, or libraries.

Hospitals typically require printers at nursing stations, to print patient records or test results. These printers must be compact, very quiet, and reliable. As a result, most large hospitals use IBM cluster networks, which creates an ideal vertical market for plug-compatible vendors. Special characters, used only in library applications, create the need for a highly specialized printer in libraries. And banks may need point-of-sale printers with quick forms access and the ability to use special card-stock paper. Obviously, no single vendor can manufacture specialized printers for all vertical markets, but plug-compatible vendors can solve these application problems by integrating independently manufactured printers with the user's computer system.

Depending on the host system's hardware- and software-interface requirements, if the printer is to be plug-compatible, it will require varying degrees of modification. Vendors that specialize in performing these modifications are becoming increasingly competitive as awareness and acceptance grow among end users. ■

DESIGN EPSON PRINTER RELIABILITY INTO YOUR SYSTEM. AND GET SOMETHING EXTRA.

With our new serial dot matrix printers, for instance, you can choose features such as 160cps print speed, *correspondence quality*, pica, elite and over 60 character weights and sizes; 7 dot *addressable graphic modes* including 1 to 1 aspect ratio; user-defined character sets, and more.

But you get something extra whenever you design Epson OEM products into your system. You get confidence. Confidence because you're dealing with the #1 OEM manufacturer of serial dot matrix printers and mechanisms on the planet — we shipped over 7 million this year. Confidence because of the exceptional reliability designed into each unit — out of box failure rate is measured in tenths of a percent. And confidence because of our dedication to quality that has made that record possible.

You get value, too. Epson printers give you the lowest total cost of ownership in the industry. And they can be customized from the inside out to meet your unique system requirements.

We simply don't have room to tell you all the Epson extras. But if you call or write we'll fill you in on all the details. Do it today.

EPSON
EPSON AMERICA, INC
OEM Products Division
Printer Group
3415 Kashiwa Street, Torrance
CA 90505 (213)534-
TELEX 182412

SW Region (714) 751-1919 • NW Region (408) 985-8828 • SE Region (404) 456-9666 NE Region (617)
• CENTRAL Region (815) 455-2570

Band Printers Eclipse Drums—But Duty Cycle Determines Best Choice

Widely touted as a direct replacement for drum printers, band printers offer many advantages. However, duty-cycle limitations can affect reliability.

by **Edward J. Thomas Jr.** and **Robert A. Ross,** Digital Associates Corp.

A void was left in the line-printer marketplace with the recent phase-out of the heavy-duty drum printer, and the band printer, introduced in the mid-1970s, became the heir-apparent to fill this gap. Many microcomputer vendors are even selling them as direct replacements for drums. While this is the best low-cost alternative, the drum printer's duty cycle—approximately 8 to 10 hrs/day—far exceeds that of most band printers.

DUTY CYCLE

Before the band printer was introduced, users had two choices in line printers—a heavy-duty drum or a chain printer. Then, with the proliferation of the minicomputer, manufacturers began to offer band printers with duty cycles in the following categories:

- **Low**—3 to 4 hrs/day.
- **Medium**—5 to 8 hrs/day.
- **Heavy**—8 to 12 hrs/day.
- **Super**—Over 12 hrs/day.

Most band printers that are part of standard product lines fall into the low- and medium-duty-cycle categories. There's a distinct lack of heavy-duty band printers in the 300- to 1000-lpm speed range. Some alternatives will be discussed later in this article.

Manufacturers have made it difficult to evaluate duty-cycle ratings because they use mean time between failures (MTBF) calculated in hours and computed by a variety of means. Because MTBFs are generally rated in thousands of hours, users often find it difficult to translate these ratings into real-world (hours/day) requirements. Further, some band printers are conservatively rated in duty cycle while some are overrated.

When evaluating a band printer for your application, you should keep three points in mind:

- Try to obtain a copy of the product specification with reference to the recommended operational hours per day.
- Ask for customer references to determine user experience with the printer.
- Minimize bias on the part of any one system vendor by talking to an independent supplier who carries a wide variety of band printers.

Printer marketers like to refer to duty cycle in percentage figures (the ratio of recommended daily operating hours over the recommended daily power-on time). These figures are often meaningless. For instance, one band printer may have a recommended daily power-up time of 6 hrs and a recommended operating time of 3 hrs. It is therefore a 50% duty-cycle machine. Another may have a recommended daily power-up duration of 16 hrs and a recommended operating time of 8 hrs. It is also a 50% duty-cycle machine. Which would you choose for a medium- to heavy-duty-cycle print load?

ADVANTAGES OF BAND PRINTERS

In spite of possible duty-cycle limita-

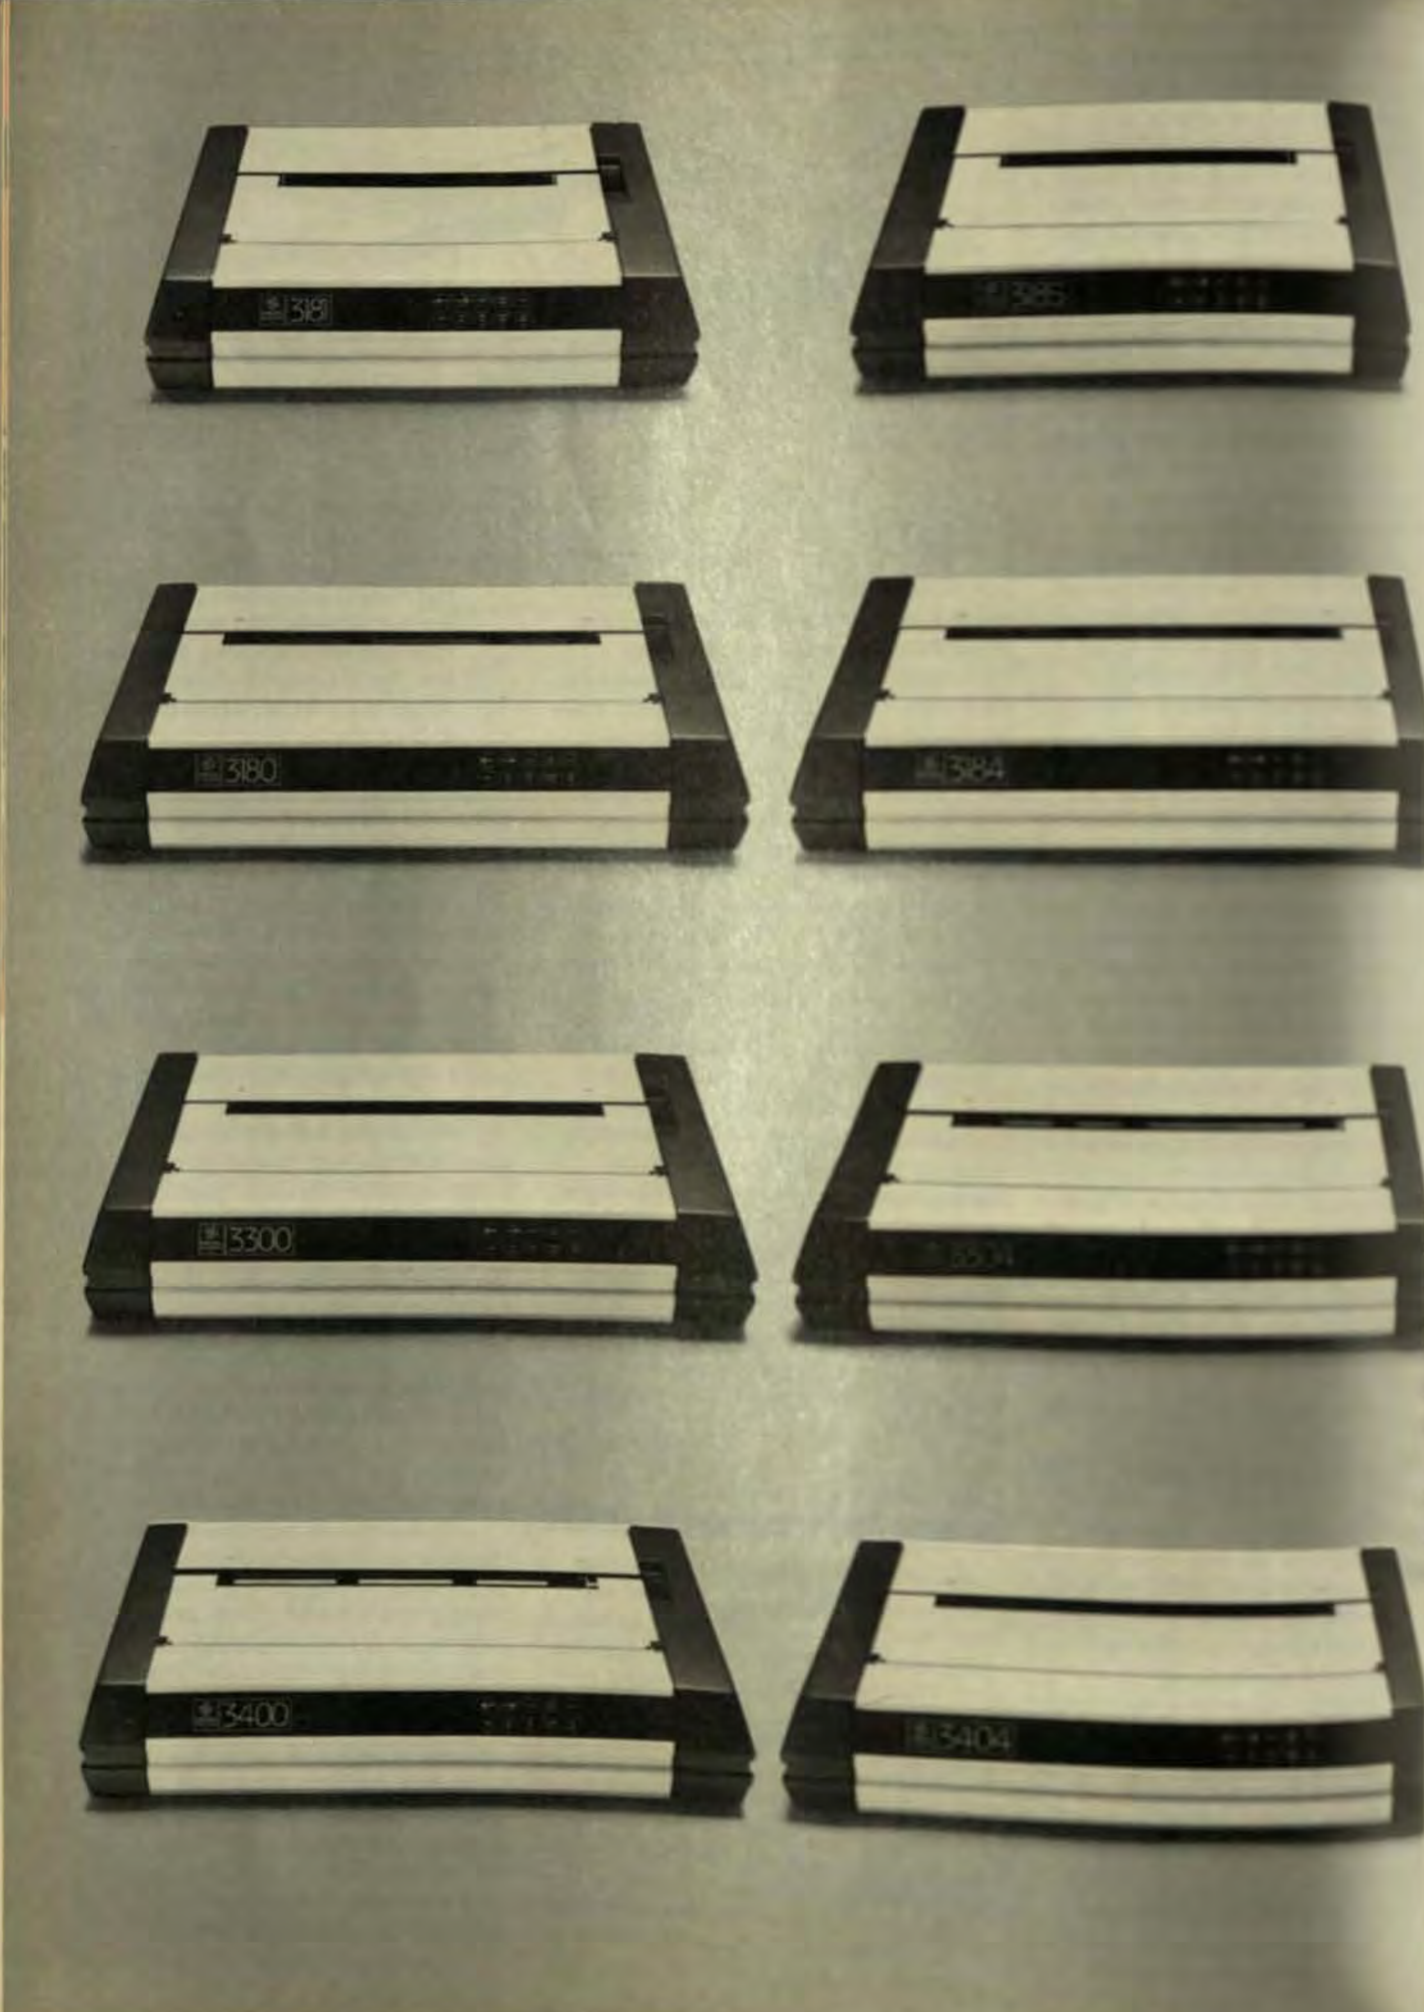

The GE 3000 printer family.

So many choices, it's the obvious choice.

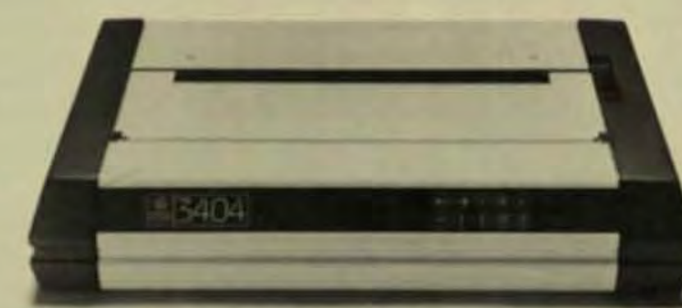

Speeds from 40 to over 400 cps. Accessible, easily programmable set up by either the operator or the system. Single or dual mode printing. Type quality from EDP to NLQ. Graphics. Multi-color printing. Selectable type fonts. 80 and 136 column models. Plus many more features and options to choose from.

All in a single printer series that's available in a wide selection of models to meet a diversity of needs...from main frames to mini's to micro-based personal computer systems.

Our basic concept is application driven price/performance matching. High parts commonality. Single source supply. Configuration and application flexibility.

At last, you can give your customers just what they need. With the GE 3000 series... printers that give you a choice.

General Electric. We introduced the first fully electronic printer with LSI circuitry in 1969. And our complete line today makes us the industry leader you should look to first.

First In Electronic Printing.

For the solution to your printing needs, call

TOLL FREE 1-800-GE PRINT

General Electric Company, Data Communication Products Department B312, Waynesboro, VA 22980. In Virginia, call 1-703-949-1170.

GENERAL ELECTRIC

helical pattern.

- **Swing-Open Band Gate.** Of the various designs used for easy loading of forms, the most convenient is that borrowed from the drum printer. The swing-open gate allows easy access for forms loading, ribbon replacement, and band changing, as well as a simple top-of-form adjustment.
- **Paper Tractors.** A band printer's paper tractors should have two important qualities: They should grip well and be able to handle difficult forms. Tractors should have as many pins as possible fully engaged with the forms' sprocket holes at any given time.

Some printers employ round or curved tractors in lieu of the conventional flat type. These tractors are dependable in most situations, but round tractors are less forgiving in situations where the forms supply is off center in the paper bay or in situations where increased drag is a problem. One common reason for increased drag is the use of multipart forms or those with many flaps or projections.

Most higher-speed band printers (1200 + lpm) have four paper tractors and can handle even difficult forms with ease since they maintain uniform paper tension in the print area. Some high-speed bands utilize two tractors and paper-tensioning solenoids in place of the two lower paper tractors for proper tensioning.

- **Construction Quality.** With printers, as with any device, perceived quality doesn't always indicate true quality. Ornate accessories and the excessive use of chromed plastic may at first catch the evaluator's eye, but to judge true quality you must look past the frills.

A quality machine performs its function as well and as simply as possible. Beware of elaborate mechanisms used for simple applications. For instance, one now defunct cartridge-ribbon system on a low-priced band printer used a cam-driven arm to tap the cartridge at regular intervals, thereby reducing the possibility of a ribbon jam. This mechanism wouldn't have been necessary if a better cartridge and ribbon drive mechanism had been incorporated into the initial design of the printer.

Inexpensive plastic parts are often used in place of metal ones, but a printer that uses many molded parts isn't necessarily of low quality. The lightweight and flexible properties of plastic make it an excellent material for many applications, such as cabinet enclosures. The evaluator should, however, beware of plastic parts that are subject to high stress or frequent use by the operator.

An examination of the printer's overall fit and finish is also recommended during the evaluation process. Pay specific attention to mating surfaces, like the cover and the body of the printer. Examine the tractor mechanism—when locked in position, tractors should be rigid with little or no side-to-side play. An inspection of the printer's paper path is also in order. The path from the paper supply to printer output should be straightforward, with as few bends as possible. A static eliminator should be included for proper paper stacking.

- **Serviceability.** A printer's serviceability or ease of repair is directly related to costly downtime and maintenance. Efficient band-printer designs incorporate features with the service technician in mind. For example, modular electronics allows for fast circuit-card exchange.

Accurate diagnostic LED status displays cut down troubleshooting time and simplify off-site preliminary service. An operator may inform a service technician of a specific diagnostic indication before he arrives on-site. This information often enables the technician to bring the proper spare parts to correct the problem.

Some manufacturers include LEDs on their circuit cards to show when a board is defective. Other machines use LEDs to simplify critical flight-timing adjustments that would otherwise require a time-consuming oscilloscope adjustment. One manufacturer also incorporates a double-hinged hood and swing-out card cage for easy service accessibility. In short, most state-of-the-art band-printer designs allow for simplified adjustments while greatly reducing a technician's need for sophisticated test equipment.

- **Power Paper Stackers.** Since the frequency of misfolded forms is directly proportional to a printer's speed, a power paper stacker is a worthwhile option for a high-speed (1000 + lpm) printer system. Power stackers are motorized devices, either freestanding or integrated into the printer by the manufactur[er]. Free-running rollers draw prin[t]ed forms into a stack in a movable tray that lowers until full (usually holding one 15-in. box of paper). Some stackers interface directly with the printer s[o] that the stacker operates only when power is applied to the printer. Such stackers also take the printer off-line when the stacker tray is full.

HEAVY-DUTY ALTERNATIVES

A gap exists in the 300- to 1000-lpm range for heavy-duty line printers. While no 300-lpm heavy-duty line printer is being manufactured currently, a chain printer can be substituted for a band at 600 and 1000 lpm. Chain printers offer the capability of printing 8 to 10 hrs/day and print quality equal to or better than the band printer. Their one drawback is cost, about 60 to 80% more than for a band printer of the same speed. However, for high-volume applications, price is no substitute for reliability.

An alternative solution is to buy two band printers instead of one and place a printer-transfer switch between them. That way, you won't exceed the duty cycle of either band printer and can almost guarantee 100% uptime because if one one printer fails, you can switch to the other.

If you are planning on an application requiring eight or more hours a day of print time in the 300 to 1000 lpm range, a band printer probably isn't the best choice. But if your application requires less in the way of daily print load, a band printer should fulfill your requirements nicely. At high speeds (1200 + lpm), there a sufficient selection of heavy-duty and super-heavy-duty bands to fulfill nearly every printing application. ■

Edward J. Thomas Jr. *is the manager of marketing communications for Digital Associates. Ed holds a BS in marketing from Sacred Heart Univ.*

Robert A. Ross *is a technical specialist with Digital Associates and will soon be promoted to production manager for the new Remote Line Printer System. Bob is currently working toward his BSEE at Bridgeport Engineering Institute.*

...and you'll unharness the pow nd performance of your computer

Florida Data OSP Printers increase computer productivity: fast, reliable and versatile.

Whether your computer costs $3,000 or $300,000, it's only as efficient and productive as your output devices. That's why from the beginning Florida Data has had one primary objective: to design and build the fastest and most versatile impact serial matrix printers for the wide variety of printing and paper handling requirements encountered in business, word processing and data processing applications.

Office Systems Printers unmatched price/performance

As a result of Florida Data's effort to date, the Office Systems Printers have been established as the performance leaders. In design. In speed. In versatility. In operator convenience. The OSP family simply outshines the competition in price/performance and cost of ownership.

Printhead technology leader

The power of the OSP is Florida Data's proprietary printhead technology, utilizing the "magnetic stored energy" principle and incorporating innovative materials that have advanced the state of the art. The eight-wire OSP printhead is the world's fastest: each wire generates 2700 dots per second. Florida Data's leading edge research and development has resulted in major patents on the printhead alone.

World's fastest matrix printer

Printing at 600 to 780 characters per second in o pass draft mode, 150 cps in two-pass corresponde and 100 cps in three-pass letter quality modes, t OSP breaks output bottlenecks for data and w processing applications. In fact, the OSP is 10 tim faster than daisy wheel printers in draft mode a twice as fast in letter quality mode.

In most applications the OSP will do the work of tw more printers. Florida Data printers really shine those end-of-the-month peak printing periods wh you need speed and performance to get the job do on time.

Florida Data's brightest highlight...versatility

Florida Data printers are in a class by themselves wh you consider the many different kinds of printing th can handle. As computers become more powerful a versatile the need to provide more printed output w continue to grow. Whether you are printing data, te equations, graphics or any other business applic tions, the OSP simply outshines the competiton.

Take font changes. You have virtually unlimited ele tronic font selection. The printer has enough memory store up to ten resident fonts. Or the computer can lo additional fonts, signatures and logos, etc. They a selectable by a flick of a switch or through compu commands. You can even change **type** *styles* a character SIZES in the middle of a sentence. The OS provides compressed print at 13, 15, 17 and characters per inch to allow up to 237 column print for spread sheet and planning data formats.

Take special applications. Florida Data has an au matic "video capture" font development system wh can provide you with custom fonts, signatures a logos. Bar codes, OCR-A and OCR-B, Greek-scien fic, and many other special fonts are already availab

Take media changes. The triple paper path mechanism easily and quickly handles many different kinds of paper for different jobs. You can be printing 8½ X 11 letters with OSP standard integrated automatic cut sheet feeder, and switch to tractor feed paper, multi part forms, labels, paychecks, legal size pages, or custom forms, in seconds, without any hardware changes. The OSP is the only printer with tractors and sheet feeder which can be installed at the same time!

Take word processing. The OSP 130 emulates the Diablo 1650/630 protocol supporting 37 word processing commands and 12 font commands for type style, print speed and character size selection. The OSP is the only printer which gives you draft copies at 600 cps using word processing features in the exact image of your final high resolution copy. Bold, shadow, underline, superscript, subscript, proportional print, horizontal/vertical motion indexing and page formatting are all part of the standard OSP firmware. You can even use these features in data processing applications for special effects on spread sheets, billing documents, program listings, etc.

The model OSP 125's interface is compatible with the parallel Diablo hytype I and Qume Sprint III. Both ASCII and spoke wheel addressing are supported along with resident multi-fonts which are operator switch selectable. The OSP 125 intelligently performs auto underline, bold and shadow completely transparent to existing software drivers, enhancing the printer throughput.

Take graphics. The raster graphics option provides bi-direction plotting, space compression, data replication and selectable resolution up to 360 X 384 dots per inch. The OSP graphics mode allows you to select from wide range of dot densities and speed for draft or final high resolution graphics output.

It's the kind of versatility that eliminates the bottleneck encountered with other printers and unharnesses the power and performance of your computer.

Rugged heavy-duty mechanism

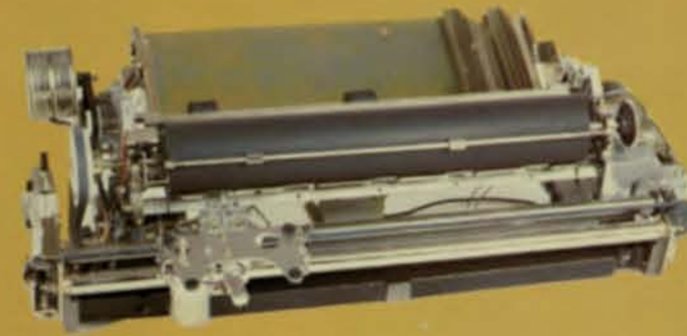

In 1981 Florida Data entered into an agreement with Brother Industries, Ltd. of Nagoya, Japan, one of the most respected electromechanical precision manufacturers in the world, to manufacture the FDC designed heavy-duty printer mechanism. This unique mechanism design has cast side frames, rigid cast rail support, a ball bearing carriage, a single-lever paper adjustment and a built-in automatic cut sheet feeder. The automatic bail control assures unmatched reliability in feeding cut sheet paper.

Displays printer status
- On-line
- Line number
- Fault codes

Reliable membrane switches
- Clearly labeled
- Adjustable bell
- Configuration switches under cover

FLORIDA DATA

OSP SPECIFICATIONS
All features standard unless otherwise noted

FONT SETS
1 each, 1, 2 and 3 pass pre-stored fonts: 10 or 12 cpi, plus one 18 cpi 1 pass font.
OPTIONAL: up to 10 fonts total, any combination of pre-stored and down loaded.
Available pre-stored sets include 10, 12, 13, 15, 17, 18 cpi plus proportional.

CHARACTER FORMATS
Horizontal — 45, 60, 90, 120 dpi density, resolution to 1/360 in. (.0706mm).
Vertical — 64, 128, 192 dpi (1, 2, 3 pass).
Two high or two wide or both.

PRINTHEAD
Patented stored energy principle, 8 wires, .0132 in. (.335 mm) diameter, 1/64 in. (.397 mm) spacing, long life ruby guide.
Pin cycle time — 370 microseconds.

PRINT SPEED
(10 cpi): 1 pass: 600 cps, 460 cps
2 pass: 150 cps, 115 cps
3 pass: 100 cps, 57.5 cps

PAPER HANDLING
Continuous forms — width 3 to 15 in. (7.62-38.1 cm), adjustable tractors, bottom feed, single lever paper tension adjustment.
Automatic cut sheet (cassette) — width 8 to 12 in. (20.3 to 30.5 cm) length 8 to 15 in. (20.3 to 38.1 cm).
Single sheet — width 3 to 15 in. (7.62 to 38.1 cm). Demand loaded from top.

HORIZONTAL FORMAT
Programmable: horizontal motion index, offsets, left margin, tab stops and absolute tabs.
Maximum print line 132 characters at 10 cpi, 237 characters at 18 cpi.

VERTICAL FORMAT
Programmable: vertical motion index, top and bottom margins, tab stops and absolute tabs.
Control panel: 6 or 8 lpi, single or double line feeds, top and bottom margins.
Line spacing: 48 lpi to 384 lpi including 6, 8, 12, 16 lpi.

COMMAND SET
Diablo, Qume, NEC compatible. Optional bidirectional graphics, font down load. Overstrike mode plus Diablo WP enhancements: auto bold, shadow, underline.

INTERFACES
RS-232 protocols:
Xon/Xoff, ETX/ACK, DTR/$\overline{DTR}$ (pin 20), RDY (pin 11)
7 or 8 bit data, even: odd or no parity.
DTR and RDY sync. on stop bit.
Baud rates: up to 19200.
OPTIONAL: Centronics, Data Products, IBM, NCR, Burroughs, Honeywell, UNIVAC, HASP.
OSP-125: Diablo (Hytype I, Hypro 6), Qume 13 bit parallel, ASCII or printwheel position addressing.

CONTROL PANEL
MEMBRANE SWITCHES & THUMBWHEEL SWITCHES:
FORMS LENGTH (1/6 to 16.5 inches in 1/6 inch increments)

MODE (Fonts: 10 positions)	SET TOP MARGIN
SET TOP OF FORM	SET BOTTOM MARGIN
BUFFER CLEAR	SELF TEST
FORM FEED	DOUBLE LINE FEED
MODE CHANGE	

INDICATORS:
Two 7-segment digits for printer status and line numbers

LINE FEED	8 LINES PER INCH
REVERSE LINE FEED	ERROR RESET
PAPER FORWARD	ON LINE
PAPER REVERSE	

Individual indicators: TEST MODE, DOUBLE LINE FEED, 8 LINES PER INCH, ON LINE.

BUFFER
512 characters, 2560 characters (opt)

RIBBON
Diablo Hytype II type cartridge, nylon ribbon.

DIAGNOSTICS
30 status conditions detected and reported.
Self test mode selectable from front panel.

GRAPHICS (OPTIONAL)
High speed, bidirectional.
RS-232, Centronics or Data Products parallel interface.
Horizontal dot density 45, 60, 90, 120 dpi, resolution to 1/360 inch.
Vertical dot density to 384 dpi.
Data stream features run length coding for compression.

POWER REQUIREMENTS
Standby 75 VA Maximum 500 VA
100 — 250 VAC (± 10%), 50 or 60 Hz.

ENVIRONMENTAL REQUIREMENTS
45°F (7°C) to 106°F (41°C) operating temperature
-20°F (-29°C) to 135°F (57°C) storage temperature

PHYSICAL
Height: 9.5 in. (24.1 cm) Width: 25 in. (63.5 cm) Depth: 28 in. (71.1 cm)
Weight: 85.5 pounds (38.8 kg)

FLORIDA DATA
FLORIDA DATA CORPORATION
600D John Rodes Blvd., Melbourne, FL 32935
Phone (305) 259-4700

Lack of Clear Standards Complicates Design of Low-Cost Color Printers

Because interface standards haven't been established, printer manufacturers must offer separate designs to accommodate each of the prevailing display technologies.

by **Art Cleary**
and **Dennis Buckley**,
Advanced Color Technology

Recent developments in the area of graphics, printers, and local-area networks all point to a proliferation of color hard copy in business applications. Since this printout will probably be generated by a system comprising products from diverse manufacturers, the need for well defined standards will be felt more acutely than ever before. These standards will have to encompass the graphics spectrum from intersystem communications of both vector and pictorial images, to the needs of the small-systems/color-printer interface.

It is understandable that the standards currently under consideration by ANSI are vector/line-oriented, since the generation of industry standards historically has lagged behind market trends. Vector standards are communications-efficient, but don't lend themselves to expanded features without a significant loss of image quality, nor can they accommodate the increased complexity that a pictorial scene requires.

While the need for vector-generated images will continue and the requirement for solid-color areas is satisfied by the inclusion of fill commands, similar standards for a raster-oriented image will require a significant amount of memory and computational power—which equals increased cost. The printer industry has historically developed diverse individual protocols rather than add any significant product cost.

RASTER STANDARDS

What is becoming evident is the need for an efficient, raster-oriented standard that would allow the manufacturers of graphics printers to adapt quickly to a diversity of systems. This goal could be achieved readily because a great degree of commonality already exists in the way raster images are generated for the color monitor. The system's raster generator scans the image in memory, a line at a time, then passes this information to a lookup table to generate the appropriate shade. The shade is converted to an analog signal and video-synchronization signals are added (Fig 1).

A printer interface could emulate this sequence with little modification, though a preface would be needed to stipulate the expected size of the image—information usually not needed for a video monitor because the lines and pixels in the field are fixed in hardware. Also, a second preface would be needed to load the lookup table with appropriate color values. While this takes considerably longer than the size command, it only needs to be transmitted once unless the color assignments are changed. The communication should be performed with RGB values to be consistent with existing video systems, which are based on additive primary colors. This can be accomplished without handicap because a direct relationship exists between the additive (red, green, blue) and subtractive (yellow, cyan, magenta) colors (see Table.)

Color lookup tables can be fixed or variable. The number of shades possible is equal to the product of the possible intensities of each primary

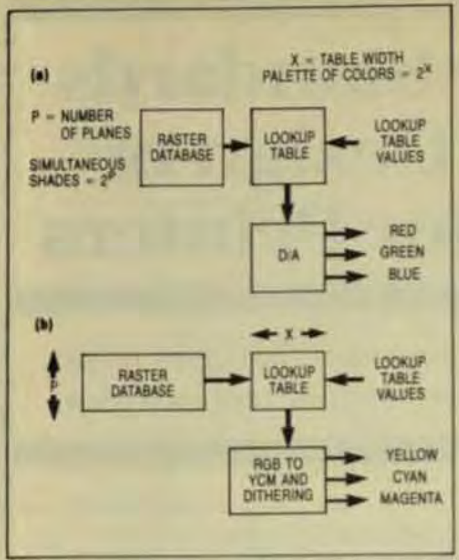

Fig 1 A degree of commonality already exists between typical video raster systems (**a**) and color printer raster systems (**b**).

color, but values beyond 8^3 are impractical since they exceed the human perception limit. Fixed lookup tables have the color values assigned to specific bits in the data word. For example, the lower three bits may equal red, the next three bits may equal green, and the two most significant bits may equal blue.

Variable lookup tables are the most common, with up to 256 intensities available for each of the primary colors. That would result in a possible 16.8×10^6 shades. Prior to database transmission, a final preface is required to identify whether the data will be in single, double, or triple precision.

Actual data communication requires three codes to ensure that the data is received in proper sequence. These would be: start of image, start of line, and end of image. While only the start-of-image command is absolutely necessary, the other two codes would allow the graphics printing sequence to be suspended for such other tasks as a status report or a change in the color lookup table.

The data-transmission portion of this sequence is the area in which vector-oriented data has an advantage since it's more efficient for the communication of simple images. Rasters

TABLE — RGB TO YCM COLOR CONVERSION	
ADDITIVE PRIMARY	SUBTRACTIVE PRIMARY
R = RED G = GREEN B = BLUE	Y = YELLOW C = CYAN M = MAGENTA
R = Y*M = C	
G = C*Y = M	
B = M*C = Y	
WHITE = R*G*B	BLACK = Y*C*M

can overcome their limitation in this area, for all practical purposes, by incorporating a run-length encoding scheme whose basic elements are commands that identify the start and end of fields of repetitive data so that entire portions of a picture can be duplicated. This compression technique depends on repetitive data and will lose its effectiveness as the complexity of the image increases. But if its use is selective, it will out perform vector data on all but the most simple images.

Lower cost printers will have trouble justifying the cost of the memory for the lookup table and the required line storage. A compromise standard for this class of applications is the transmission of raw RGB pixel data. A preface would still be required to establish the size of the image, and a start-of-image and start-of-line command, but data would then be transmitted directly without the use of a lookup table.

The sequence would be a byte of red, followed by a byte of green, followed by a byte of blue . . . until the line is complete. Each byte would be rotated so that each bit represented a pixel (LSB = pixel 1; MSB = pixel 8), and every three bytes would contain information for eight pixels. The limitations of this method are that any given pixel could only be one of eight colors, and if the number of pixels on a line were not evenly divisible by eight, then the unused pixels would have to be filled accordingly in the appropriate data bytes (Fig 2).

CURRENT DISPLAY TECHNOLOGY

The designer of today's color-graphics systems has more options than ever for hard-copy output. Impact, ink-jet, thermal, and camera technologies are competing for the buyer's attention. But this activity often overshadows the major deficiency of the color graphics industry, the lack of interfacing standards.

Interfacing is structured by the organization of the source database, the communications channel, and the output device. Of these three factors, only one is somewhat stable—the communications channel. The parallel and RS-232 standards have matured, and most local-area network (LAN) designers are striving to keep their designs data-independent. The database and output devices, however, are in a state of rapid change, complicating the establishment of standards needed to satisfy existing and future needs.

Graphics data can be visualized as having three levels of complexity: line drawings, solids generation, and pictorial. These depend on the number of graphic elements in a given image and their relationship. As the need for graphics data increases, operating conventions that were instituted to simplify the entry/storage process become counterproductive, and new functions need to be added to the operating systems.

For example, solid areas of color demand a fill algorithm for efficiency, even though the task can be accomplished with a series of vectors. Random pictorial scenes usually defy compression through the vector approach and favor a run-length encoding scheme that capitalizes on the raster orientation of the data.

Initial graphics data was, by necessity, vector-oriented, because memory was expensive and data was manually entered. In addition, the stored information could easily be mathematically manipulated and adapted to expanding system complexity. The analog display technology of the time satisfied these simple demands by producing a generation of stroke-oriented CRTs.

As system costs decreased, more functions were demanded, and graphics operating systems kept pace by offering the capability for arc generation, solid fills, and color manipulation. The limits of stroke CRTs were quickly reached, and raster technology became the mainstay of the industry.

Raster technology, while less efficient for simple displays, has the advantage of a constant display time independent of image complexity. In addition, the hardware necessary for the raster generator is amenable to LSI techniques, resulting in a series of CRT-generation chips that have placed this technique within the reach of every system designer. New advances in the writing speeds of the CRTs have ensured that raster will be the leading display technology in the immediate future.

Today's graphics systems use both techologies. Operating systems such as SISGRAF or PLOT 10 enable data to be stored as a sequence of commands that are then used to construct a raster-memory image for display on the CRT. This approach will continue, since manipulation of the graphic form is the goal of a system, as opposed to strict manipulation of color. Graphics systems need forms stored as mathematical entities, yet the display format needs the flexibility of a raster orientation.

Early graphics hard copy was produced by plotters, which took advantage of existing technology. This electromechanical approach has kept pace with operating systems but is rapidly approaching its natural limits as requirements for image complexity and operator ease increase. Plotters are still the choice for simple images since they take full advantage of traditional technology to produce smooth diagonals and fast line drawings at relatively low cost.

The simple tasks of solid fills and increasing the number of colors quickly demonstrate the shortcomings of this approach, since throughput decreases as image complexity increases, and the capacity and reliability of the ink supply comes under question for large or densely colored plots. Cost reductions also become more difficult because the plotter is primarily electromechanical, and unable to take full advantage of the semiconductor's price/performance improvements.

RASTER PRINTERS

With the introduction of the personal computer, the appeal of graphics fueled a development cycle whose first product used standard ASCII characters for a crude image representation. Since dot-matrix printers were the choice of this marketplace, it was only natural for the next generation of graphics hard-copy devices to utilize shaped characters to produce graphics images. This approach can be expanded by use of a writeable character generator.

As these developments were occurring with character-oriented systems, other microcomputer manufacturers were going directly to bit-addressable graphics. The dot-matrix printer manufacturers were quick to respond with electronic variations of their products that allowed the bit-addressable graphics image to be printed out. Until recently, however, these were rigidly oriented toward the mechanical configuration of the dot-matrix printer, and primarily restricted to black-and-white images.

The last few years have seen the emergence of color in the graphics-

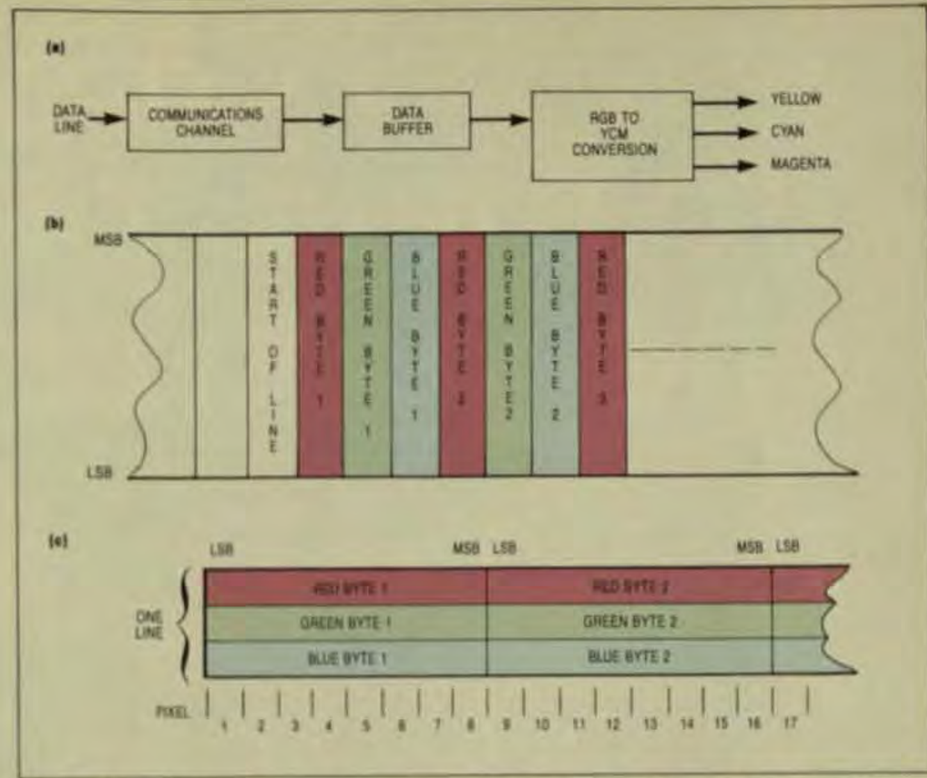

Fig 2 This sample block diagram shows a low-cost printer designed without a lookup table (**a**). In this design, a data stream must be performed by at least a start-of-line command (**b**). The sequence for a video raster line is a byte of red, followed by a byte of green and a byte of blue (**c**).

printer marketplace. The development of these devices has been preceded by a similar emergence of lower-cost color-graphics systems. At first, color was implemented as modifications to the ribbon systems of popular dot-matrix impact printers, but more recent entries are totally designed for color and utilize non-impact designs as well. Ink-jet color printers cover the application needs from the sophisticated Applicon plotter, through Advanced Color Technology's ACT series, to the low-cost Canon printer.

These products reinforce the trend toward hard-copy devices. Future products will continue to be raster-oriented and to generate an increasing number of shades to track those offered by the host system. While camera systems still dominate the upper end of high-resolution applications, paper-output printers are rapidly encroaching on their domain now that transparencies are available on an increasing number of products.

Since the demand for 8 ½ × 11-in. hard copy predominates in both the transparency and paper markets, raster devices will probably focus in this area, leaving the requirements for larger devices to existing plotter products. The price/speed ratio will continue to improve dramatically since a graphics printer has considerably more electronics than a simple character printer.

Driven by the personal computer, low-cost graphics systems, and the diversity of hard-copy printers, color graphics is rapidly entering the office environment. As with any area of rapid change, there's also a proliferation of different, and possibly conflicting, activities. The venerable ASCII code stabilized the text world many years ago and is still being used now. What is needed is a standard that's an order of magnitude more complex, to serve the emerging world of color graphics, both now and in the future. ■

Arthur L. Cleary *is electronics engineering manager and a founder of Advanced Color Technology. He previously worked with Applicon Inc. and Dataproducts. He holds a BSEE degree, and has been associated with ink-jet technology since 1976.*

Dennis J. Buckley *is vice president of product development for Advanced Color Technology. He holds a BS in electronics.*

to make points
when you're a pro.

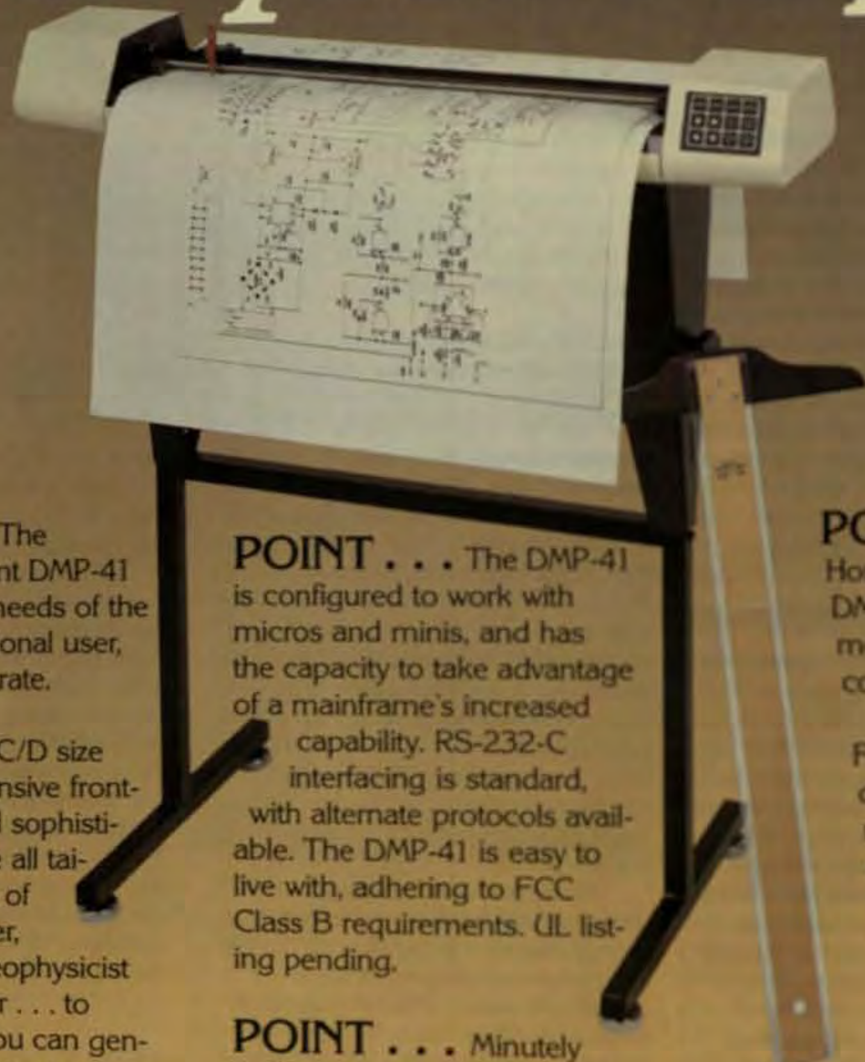

POINT . . . The Houston Instrument DMP-41 plotter meets the needs of the serious or professional user, yet it's easy to operate.

POINT . . . C/D size format, comprehensive front-panel controls and sophisticated firmware are all tailored to the needs of the surveyor, drafter, oceanographer, geophysicist and land developer . . . to name but a few. You can generate superior architectural elevations, contour maps, circuit-board layouts and assembly drawings quickly and accurately on bond, vellum or synthetic media.

POINT . . . The DMP-41 is configured to work with micros and minis, and has the capacity to take advantage of a mainframe's increased capability. RS-232-C interfacing is standard, with alternate protocols available. The DMP-41 is easy to live with, adhering to FCC Class B requirements. UL listing pending.

POINT . . . Minutely defined step size and high-resolution logic—combined with robust drives and optimized pen ballistics enable you to create plots of high precision and surpassing quality.

POINT . . . Th
Houston Instrumen
DMP-41 is one of y
most cost effective
considerations.*

For the name, ad-
dress and phone
number of your n
est distributor, w
Houston Instrum
8500 Cameron
Road, Austin, Te
78753. Phone
512-835-0900
800-531-5205 if outside
Texas. In Europe contact
Bausch & Lomb Belgium
Rochesterlaan 6, 8240 Gis
Belgium. Tel 059-27-74-4
tlx 846-81339.

BAUSCH & LOMB
houston instrument division

*suggested US retail $2,995

THE NEW AMPLOT-II
IS ALL BUSINESS.

PLOTTER APPLICATIONS . . .

Analytical data display • Mathematical functions • Geometric patterns • System/component characteristic display • Statistical charts • Workflow diagrams • Performance graphs • Organization charts • Creative graphics • Personal achievement charts • Instructional/reference diagrams.

The new AMPLOT-II, on computer command, can automatically produce hard-copy, 6-color business or engineering graphics in minutes. Designed for the end user and OEM applications, its high pen speed, automatic pen retrieval and high resolution assure fast accurate 10" x 14" plots. Compatible with most personal computers, the AMPLOT-II is economically priced at only $1299.00 retail. Just circle the reader service number for complete engineering details.

2201 Lively Blvd. • Elk Grove Village, IL 60007
(312) 364-1180 TLX: 25-4786
REGIONAL OFFICES: Calif. (714) 662-3949 • Texas (817) 498-

AMDEK

Amdek . . . your guide to innovative computing!

Interactive Digitizers Increase Flexibility of User/Computer Interface

Sonic and electromagnetic data-acquisition systems both allow 3-D input, but mechanical systems that use a hand-held pointer can simplify engineering applications.

by **Livingston Davies,**
Micro Control Systems

Data-acquisition systems can be divided roughly into two types—interactive and noninteractive. Available noninteractive systems include most coordinate-measurement systems, which are mechanical devices used primarily in such high-resolution applications as machine inspection. Interactive devices are generally less precise, but simplify the interface between the user and the computer by permitting a hand-held pointer to be moved more or less freely to a location of interest. This article will discuss both types of digitizer, with particular attention to the applications made possible by an interactive approach.

NONINTERACTIVE DATA-ACQUISITION SYSTEMS

A typical coordinate-measurement system looks a lot like a milling machine. It has a 2-dimensional bed that moves in X and Y coordinates (Fig 1) and a probe that moves up and down in the Z dimension. Such a system can easily obtain position information to an accuracy of 0.0005 in. or better for use in critical inspection applications. In some systems, the bed is movable in two dimensions and the probe moves in the third. In other systems, the probe itself moves along all three axes. To use this device, you fix whatever you are trying to measure (the workpiece) to the work table. Then you move the probe over so that it touches the item being tested, at a position of predetermined interest.

Usually, the probe has a little ball on the end of it. A computer monitors the three axes of motion and can therefore calculate the position of the probe. When the user moves the probe into contact with the workpiece, the computer calculates the center position of the ball and compensates for position offset due to the ball size. Since the computer is inferring that the workpiece is tangent to the probe ball, it must determine a unique tangency point. That is, it must figure out what part of the ball's circumference has been touched.

In a computer-driven coordinate-measurement system, a computer is programmed to move the reference tip (usually a touch-sensitive probe) to various locations on the part to be measured, and, when it gets there, to record its position or a deviation from a nominal position.

Using a coordinate measurement system, an engineer will typically look at the blueprints for an item to be inspected, determine interesting dimensions or inspection points, and then write a program that will move the probe to the appropriate locations and record the results. This is usually a noninteractive, very high-resolution task.

In facilities that support computer-aided design (CAD), the dimensions of interest are located in the parts database. However, the interface between CAD and computer-aided inspection is still in its infancy, partly because inspection typically is concerned with 3-D measurements.

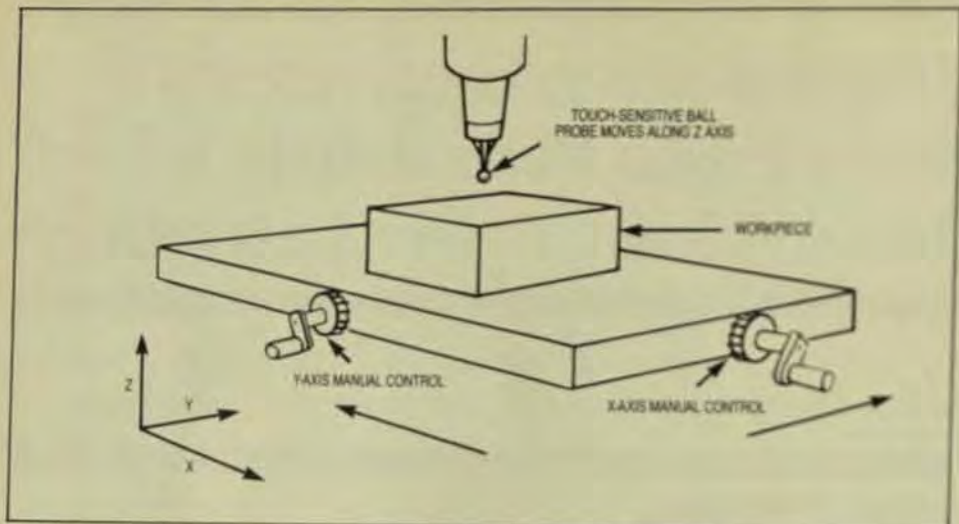

Fig 1 A coordinate-measurement system with mechanical controls allows measurements with tolerances of 0.0001 in.

and most CAD systems still use 2-D design practices. With increased use of solid-modeling and surface-modeling techniques that retain the true 3-D nature of an object in a computer's memory, this interface should become more direct.

INTERACTIVE DATA-ACQUISITION SYSTEMS

Interactive data-acquisition systems are intended to simplify the interface between the user and the computer by permitting the user to move a hand-held pointer more or less freely to a location of interest. A computer then calculates where the probe is located. Accuracies obtainable are typically an order of magnitude coarser than those obtainable with a coordinate measurement system.

There are three types of interactive systems:

- A mechanical system that uses a dead-reckoning approach, developed and marketed by Micro Control Systems.
- A sonic system made by Science Accessories Corp. (See COMPUTER TECHNOLOGY REVIEW, Winter 1982).
- An electromagnetic approach, announced of Polhemus Navigation Sciences, a division of McDonnell Douglas Corp., for release this year.

The sonic system uses a triangulation method of coordinate calculation. That is, a set of detectors situated around the work volume are used to sense signals transmitted from a hand-held probe. The source within the probe emits sonic radiation, and a computer calculates the source position by triangulation on all of the sensors (Fig 2). The magnetic system utilizes a low-frequency transducing technique to calculate the position and the orientation (X,Y,Z azimuth elevation and roll) of its single tri-axial sensor relative to its single tri-axial transmitter.

The accuracy of such systems depends on the positioning of the sensors and on the probe's position within the work volume. By using more sensors than needed to infer position, the host computer can use data averaging and other signal-processing techniques to reduce position-estimation error. Digitization accuracy being quoted for the sonic system is approximately 0.03 in. The projected accuracy of the magnetic system is 0.02 in.

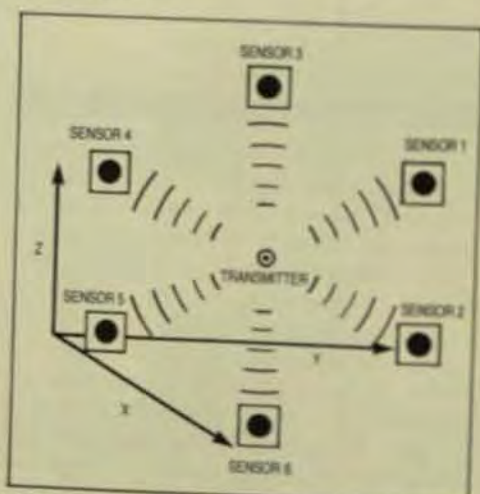

Fig 2 The phase-shift, which changes with displacement X as shown, provides an accurate measure of displacement within a complete grid-winding cycle of length d.

These systems are active, in the sense that they emit radiation. As a result, they constrain the objects that can be digitized and the environment in which digitization can occur. For example, with the magnetic system you can't digitize solid metal objects. With the sonic system, there must be an unobstructed line of sight between the digitized probe position and at least three sensors. Both magnetic and sonic systems can, at least in principle, be constructed to work in an arbitrarily large work volume. The ones currently marketed work in a volume of approximately $1m^3$.

An interactive position digitizer is a robot-like system consisting of a set of rigid links connected at rotational joints (Fig 3). It offers four degrees of rotational freedom and uses a dead-reckoning system of coordinate measurement. By monitoring each of the joint angles and knowing the lengths of each of the links connecting the joints, a computer can calculate the position of a pointer tip.

Starting at a coordinate-axis origin located at a rigidly mounted base point, a computer sequentially calibrates the position of the next most distant joint location. Each coordinate-transformation form can be thought of as a rotation about a certain axis—the rotation depends on the joint and its monitored angular position—and a translation along a rigid link by an amount equal to the length of the corresponding link member. These coordinate transformations can be implemented as a series of matrix multiplications.

Since the coordinate-transformation matrices turn out to be sparse, not all terms need to be retained. The transformations are actually implemented in a reduced nonmatrix form. To calculate a pointer-tip position, the computer needs to go through only as many calculations as there are degrees of freedom in the device—four, in this case.

With this design, the useful work volume is constrained by the size of the arm linkages. At maximum extension the arms reach approximately 15 in. from the base, so the digitizer functions in an approximate 15-in. radius. The useful work volume can be increased with longer arms; however, since the system uses a dead-reckoning system of calculating

We make IBM look good.

Our new ZETA 887 pen plotter runs direct from your IBM controller for the best color graphics you've ever seen.

You're looking at the only pen plotter made that can give you high resolution hard copy directly from your IBM 3274 or 3276 cluster controller. And we're operating system independent. No add-ons. No upgrades. Just plug it in and run.

We've been successfully interfacing our high performance pen plotters to IBM for years. They've got great machines and so do we.

Our ZETA 887 is no ordinary plotter. It's unique pen carriage holds up to eight different colored pens – virtually eliminating the annoying downtime for pen changes. No more plotter babysitting. We've designed it to go at high speed (20 ips) and with our "no stopping" feature you'll get unbeatable throughput.

You'll also find precision comparable to far costlier plotters. Pen resolution good to within one-one thousandths of an inch. Plot without operator intervention on continuous feed media over 11" by 120 feet.

Designed expressly for the SNA/SDLC environment.

The ZETA 887 connects to your IBM cluster controller the same way your IBM 3287 printer does. Just plug it in.

We're more than just hardware compatible. We take IBM's GDDM and convert it into a data stream all of our Zeta plotters have been designed to understand. In fact, we're a leader in the IBM interfacing world with a wide family of field-proven plotters running direct from IBM.

The same commands drive all Zeta plotters so you can be assured of running any application program using industry standard plotting routines. We're supported by the software leaders with such packages as ISSCO's DISSPLA and TELLAGRAF, SAS's SAS/GRAPH and Precision Visual's DI 3000.

Call (415) 671-0600. Or write: Nicolet Zeta Corporation, 2300 Stanwell Drive, Concord, CA 94520 TWX 910-481-5951

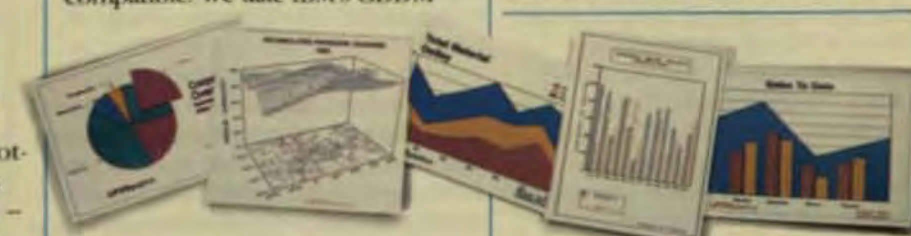

Nicolet

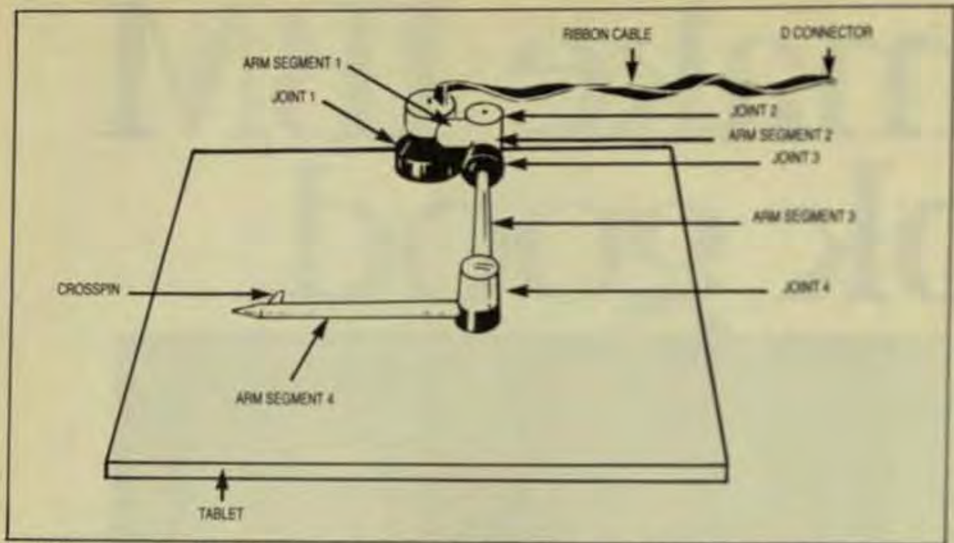

Fig 3 Mechanical 3-D digitizers employ an arm-linkage structure. By starting at a rigidly mounted base point, a computer sequentially calibrates the next most distant joint location.

pointer-tip positions, resolution decreases as the size of the work volume increases.

A popular design of the touch probe used in such systems was developed in England about 10 years ago. It consists of a spring-loaded probe that's supported in three places by electrical contacts (Fig 4). If the probe is moved in any direction, at least one of these electrical contacts is broken. The computer can detect the probe's deviation from its nominal position by monitoring electrical continuity through the probe. The resulting input data can then be viewed in a variety of ways (Fig 5).

Only three degrees of freedom are required to permit a pointer to move uniquely to a location in X,Y,Z, space. The redundant fourth degree of freedom substantially increases device flexibility and makes it much easier to reach certain points—especially if you're interested in reaching around an object. It also lets you hold the pointer stylus in the same way you might hold a pen.

The coordinate accuracy of a resulting digitized point is constrained by the accuracy of the computer's knowledge of the joint angles and the accuracy to which it knows the link lengths. Typical systems range in accuracy from 0.1 to 0.01 in. These systems are robot-like in that they consist of a set of rigid mechanical arms connected by rotational joints that can be monitored by a computer.

With the digitizers, the user must position a pointer tip manually and the computer simply tells him where the tip is located. In a robot, the task is not only to infer position but to have actuators force the robot to a certain position. Both systems use a classical geometric coordinate-transformation technique to infer the pointer-tip position.

ROBOT-TRAINING TECHNIQUES

An interesting application for interactive 3-D digitizers is that of robot training. Here, the problem is to have a robot move through a sequence of appropriate actions, such as: Go to point A, pick up an object, then move to point B. The most common method currently used for robot training is an interactive technique involving the robot itself. A trainer moves the robot through its paces and indicates, at appropriate points, the coordinates at which the machine is supposed to reside.

Another technique is a continuous-path training method in which the trainer moves the robot through the sequence of steps appropriate for a particular task. The robot records the sequence of positions through which it must move and plays it back to repeat the operation on subsequent workpieces.

Yet another technique is the point-to-point method. In this case, the trainer positions the robot at a certain place, indicates that as a point of interest, then moves the robot to a different position and indicates it as a position of interest. A controlling computer can then calculate a direct path from the first point to the second, which may or may not be the one through which the trainer moved the robot. Using this method, a trainer must be careful to include enough significant points for the robot to pass through without a collision between its arms and any workpiece.

A much more sophisticated method requires a computer to know the geometric characteristics of the robot itself and of any object that the robot might be carrying, and to have a model of the work volume in which the robot will be operating—including any workpieces that might be there. Given this information, a computer can calculate tool positions and tool orientations relative to the workpiece and can then go through a collision-avoidance calculation. The collision-avoidance calculation is made to ensure that the work tool does not run into anything.

Collision-avoidance calculations, which require heavy computational power, can largely be avoided by the use of an interactive 3-D digitizer. The trainer moves the pointer through the curve that the robot tool should move over, on a scaled-down model of the work volume. This can be done at an engineering workstation rather than a tooling workstation, thereby substantially cutting the costs of the hardware involved in the training. Since an interactive digitizer permits you to move a pointer tip through essentially arbitrary curves, it's an ideal aid for an engineer who wishes to trace approximate paths that the robot arm should move through.

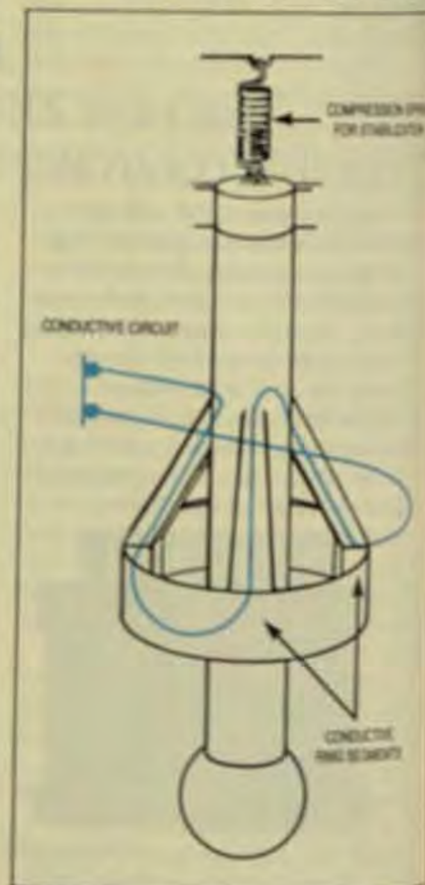

Fig 4 In a touch-sensitive probe, an electrical circuit (dotted line) is complete when the probe is undisturbed. If the probe's ball is moved in any direction, the circuit is broken.

Another useful application of interactive 3-D digitizers is that of field mapping. In this case, an engineer places a sensor at the tip of the digitizer's arm and moves the pointer tip to a variety of locations of interest. The sensor provides temperature, light intensity, sound intensity, or whatever that particular type of sensor is designed to provide. Together, the devices supply the engineer with the information necessary for field mapping—for example, temperature as a function of position in 3-space.

A RANGE OF APPLICATIONS

Interactive 3-D input is clearly an advantage in the modeling of irregular surfaces that can't be calculated through mathematical equations in software. For example, this technology is being used by a prosthodontist for cephalometric analysis, the measuring of the jaw and skull for orthodontics. Digitizing is also being used in the anthropological studies being conducted at the American Museum of Natural History in New York, where skulls of fossil monkeys are digitized for the study of evolutionary relationships between different species. Digitizing is also expected to help date the sites from which the fossils were taken.

One use for today's low-priced 3-D digitizers has been found by an acoustical engineer in Dallas, who's been using an interactive device in the design of public address systems. A cluster of speakers and horns is used to direct sound at a specific seating area and sound-pressure level. Once a long series of calculations is performed, the sound field is predicted and horns are chosen to issue the proper sound coverage. Software allows these horns to be rotated into position and assembled into a cluster. Once loaded, the horns can be scaled, moved, or rotated, and plan and section views printed out.

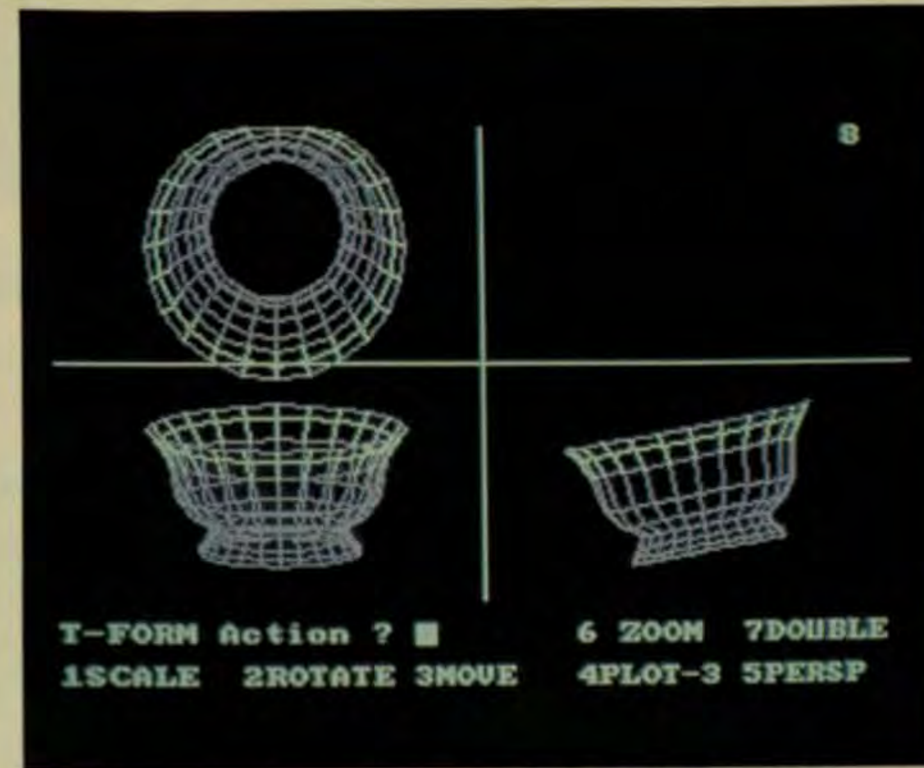

Fig 5 An image (a pewter bowl in this case) can be viewed from top, front, and side elevations simultaneously once the 3-D data has been acquired by a digitizer.

A related application is the acoustical shaping of rooms—specifically audio control rooms—where the shaping and treatment of room surfaces are critical for good audio reproduction through monitor speakers.

Another Dallas firm has added a 3-D digitizer to its microcomputer-based laser-cutting machinery. A plastics manufacturing shop, it had been using nonautomated techniques to cut, modify, and assemble plastic parts. Part completion used to take about a week, including the manual layout of patterns on sheets of flat plastic for cutting with hand tools. Now implementation with a personal-computer-based laser-cutting system has reduced cycle times to less than a day.

This system allows the product designer to interact directly withthe completed part, so that he can construct it interactively in 3-D, along with all of its associated detail parts. Using a color monitor, the designer creates a test-volume sample, and checks any dimension changes for possible interference. The cut files, after being saved on disk, are sent to the laser cutter where they can be cut immediately.

These applications are made possible by the low cost and flexibility of hand-held data-acquisition systems. These 3-D devices will continue to find new uses in engineering and other areas as the existing technology is adapted creatively to solve problems that have previously resisted computer technology. ■

Livingston Davies, *one of the founders of Micro Control Systems Inc., studied robotics and computer graphics at MIT while working on the development of automated imaging systems at C.S. Draper Laboratory. He is currently studying for his PhD in artificial intelligence at Yale Univ.*

The Best Just Got Better

Introducing the new Strobe Models 260 and 200 Graphics Plotters. You can pay more but you can't buy more.

How do you make the world's most cost-effective graphics system even more remarkable? Simple. Make it easier to use, more powerful, more versatile and more attractive. The hard part is doing all of the above while holding the line on a very affordable price. With the new Models 260 and 200 Graphics Systems, Strobe has done it all. And the benefits are all yours.

The Strobe 260 is the first eight-pen plotter and software package available for under $1,000. With just a few keystrokes on your computer, the system produces superb high-resolution charts, drawings and transparencies with up to eight automated color breaks. The even more economical Model 200 produces the same wide range of graphics with a single pen which may be changed manually.

The new Strobe plotters interface with virtually every popular microcomputer made today, including the Apple II, IIe and III, Victor, IBM PC, Kaypro and Osborne systems. What's more, the StrobeWare software programs provide every graphic format you will probably ever need, including multicolor pie charts, bar graphs, line graphs, text processing, and a broad selection of symbols and character styles. To further increase your capabilities, the new Strobe systems also interface with over 40 other popular graphics programs.

Every feature of the new Strobe Graphics System is designed to give you the highest quality graphics quickly, easily and inexpensively. The system also looks very nice on your desk. So visit your local dealer today and see how much better the best can be.

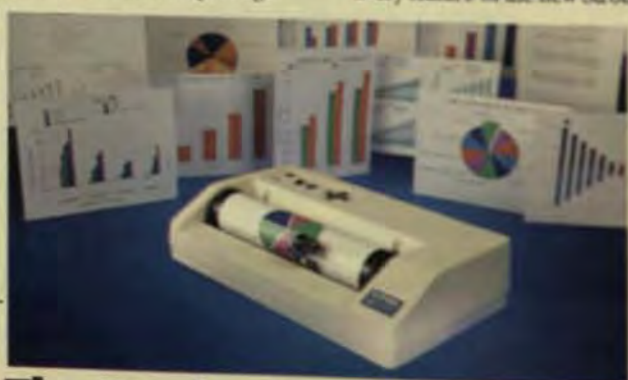

Apple II, IIe and III, Victor, IBM PC, Kaypro and Osborne are trademarks of Apple Computer, Inc., Victor United, Inc., International Business Machines Corporation, Kaypro Corporation and Osborne Computer Corporation, respectively.

STROBE

Strobe Inc.
897 5-A Independence Ave.
Mountain View, CA 94043
415-969-5130

The Strobe Graphics System
Seeing is believing

On-Board Intelligence Increases Accuracy of Plotters and Digitizers

Today's digitizers use a built-in microprocessor to process and format basic input data. Resident intelligence can be used to refine plotter output as well.

by **L. Robin Hulls**, Numonics

Decreases in computer prices have stimulated the design of a host of graphical I/O devices. The versatility of today's plotters and digitizers has been greatly enhanced by the presence of on-board microprocessors that allow special software routines to compensate for unwanted characteristics inherent in mechanical and electrical design.

In its most elementary form, the digitizer consists of two linear motion-transducing systems disposed at right angles. One is the X axis and the other the Y axis of an X,Y cartesian coordinate system. In the Numonics 1224 restrained-cursor digitizer, for example, cursor motion is transmitted to optical X,Y transducers by a mechanical arm and carriage assembly.

In the free-cursor design, the cursor is connected to the digitizer by a flexible cable and transmits information either electrostatically or electromagnetically to the underlying digitizing tablet. Apart from the obvious operator convenience, the advantages of the free cursor are improved accuracy and the optimal use of a pen. This system won't function in the presence of metallic material, which interferes with signal transmission. Similarly, the thickness of the document being digitized dictates the separation between cursor and tablet. There is a maximum thickness through which position information can be transmitted reliably (approximately ⅜ in.).

Some free-cursor systems utilize ultrasound to transmit position information through space. However, information from this type of system is often misinterpreted because of the presence of objects on the digitizing surface. Most of today's digitizers are smart devices that include an on-board microprocessor to process and format the basic X,Y position information.

Most digitizers use an absolute positioning system in which cursor position is defined by the X and Y coordinate distances from a reference point, not the X and Y distances from the previous position of the cursor. Thus, you can remove the cursor from the tablet and, in most cases, even turn the tablet off without losing the reference position for the cursor.

The important factors in digitizer specifications are size, accuracy, and digitizing rate. In any system that relies on a human operator to position the cursor, there's no advantage in being orders of magnitude faster or more precise than the operator can be. Digitizer accuracy varies significantly among designs, depending largely on the operating conditions. Disregarding operator error, the accuracy of the digitizer can vary with position and cursor height. In the case of pen-style cursors, accuracy may also be affected by the angle between the pen and the tablet surface.

Most manufacturers target their design to satisfy many applications, recognizing that in some instances accuracy will be greater than required. There are two primary factors that influence the accuracy of a design. Since all electrostatic and magnetic positioning systems use a conducting grid, accuracy is determined by the mechanical tolerances and position

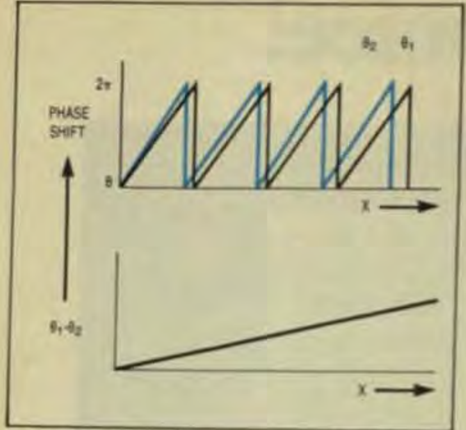

Fig 3 The phase difference $\theta_1 - \theta_2$ between the fine-position winding and the vernier winding provides a coarse-positioning signal.

ings, the phase-shift information is provided directly as a pulse count. An important advantage of phase measurement is that you can convert the information into digital form without using a conventional A/D converter.

Once the basic positioning information has been delivered to the onboard microprocessor, software procedures provide endless possibilities for processing of the position data. For example, you may scale X,Y axes, shift the origin of the coordinate system, or rotate the axes about the origin.

PLOTTER DESIGN

Plotting of digital data is the complementary process of digitizing graphical information. A plotter takes data and creates a graphical picture to represent it.

Several types of plotter are available—sprocket feed, flat bed, drum, and belt. Sprocket-feed plotters use paper that has holes along the edges and rides over driving sprockets. The plotting pen is moved in a single plane, typically the Y axis, and the paper is moved back and forth by the sprockets to create the graphic. These types of plotter are excellent for specialized applications that can tolerate the use of sprocket-feed paper, but not for applications where cut-sheet paper or Mylar must be used. For precise plotting, cut sheet is normally required.

In a flat-bed plotter, the plotting medium is paper or Mylar fastened to a flat surface or bed; the plotting device, usually a pen, is moved over the surface in the X and Y directions to create the graphic desired. In the drum plotter, the pen is moved in only one axis to create the plot. The belt plotter is similar, except that the medium is fastened to a large belt, instead of a drum, and moved back and forth.

These designs all suffer from the same limitation—the mass of the transport mechanism is large and therefore the mechanical resonant frequency is low. If it's so low that it falls into the normal plotting-velocity range, the mechanism's resonant vibrations will cause the pen to plot wavy lines instead of straight lines or cause the plotted lines to be wider than the thickness of the pen tip.

Grit-roller technology makes it possible to keep the system mass low and the resonant frequency high (above the range of the normal plotting velocity). In this design, the pen moves in a single plane, and uses a minimum-mass system to hold the pen and drive it back and forth across the plotting medium. To keep the mass to a minimum in the other axis, the medium itself is moved back and forth by two sets of pinch rollers, one on either side of the medium.

If ordinary rubber pinch rollers were used, the paper would tend to migrate from its initial position, so two types of roller are used instead. One is pliable and not unlike the type used in typewriters. The other is hard, of a precise diameter and covered by a single layer of very fine grit. The secret of this design is not that the grit provides a high frictional coefficient between the paper and itself, but that the grit embosses the surface of the medium on its first pass between the rollers.

As the paper or Mylar is passed back and forth between the rollers to create the desired graphic, the pattern embossed on the first pass registers with the pattern of the grit on the gritted rollers, forcing the medium to follow exactly the same path each time. This technique provides the kind of precise, repeatable plots required, and allows operation at resonant frequencies much above the normal velocity of the system.

The X and Y movement for this type of plotter is derived from two identical stepper motors, one for the pen mechanism and the other for the pinch rollers. Identical characteristics of motion are thus provided for the two axes. The stepping of these two motors is controlled when the pen is raised and lowered. Stepping motors are used because their motion is directly analagous to the digital data that drives them, and because they have a high rate of acceleration and deceleration. This is important because most plotters output short vectors and arcs and throughput is higher if the plotter can start and stop quickly.

In a plotter driven by a stepping motor, the X,Y movements are incremental rather than continuous, so proper selection of the gear ratios for pen and media drives is a key design factor. A low ratio produces high speed, but with large increments and poor resolution. A high ratio produces high resolution at the expense of speed. Plotting speed can be controlled by the user, but typical speeds are on the order of 5 ips.

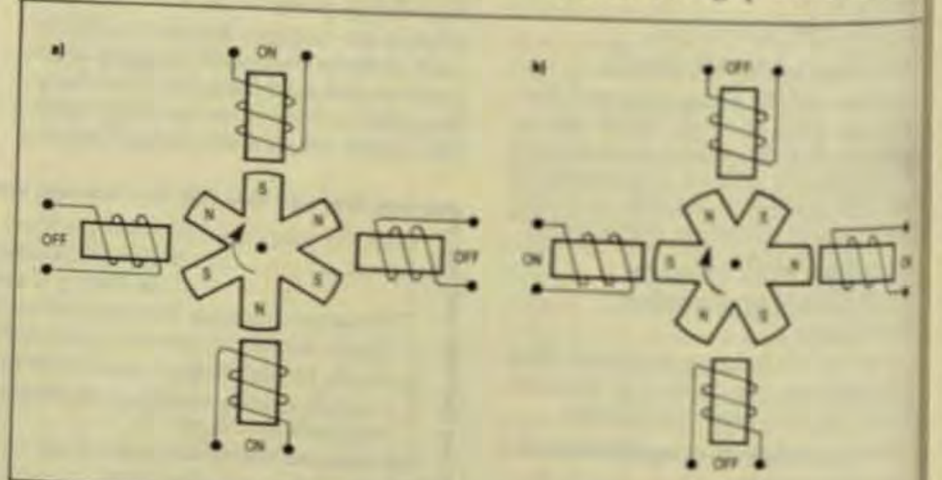

Fig 4 In a stepping motor, stepping is achieved by a sequenced field excitation that moves the rotor (a) one pole position at a time. On completion of a step, the rotor (b) remains in a position of equilibrium.

The pressure of the pinch rollers is also an important factor in this type of plotter. The rollers, mounted on a common shaft, drive the medium from both sides. The effective diameter of a roller is a function of its size and the pinch pressure. Increasing pressure causes the roller to bite into the medium, thus reducing its apparent diameter. Lack of symmetry between the two edges of the medium will result in skewing and jamming against the side of the machine.

The grit-roll diameter directly affects the accuracy of the plot. For example, if excessive pressure decreases the pinch roll's diameter, the scale of the plot on the Y axis will be reduced.

The plotter has a resolution of 0.005 in./step and an accuracy of ±0.5% for each 17 in. of axial motion.

Speed and accuracy are as important to the plotter as they are to the digitizer. Line quality is an additional characteristic that constitutes an important factor in plotter specifications. This quality is a function of the plotter pen's speed and its positioning precision.

Any spurious motion in the motor drives will generate artifacts and thickening of the line. Loss of registration in the drive will result in inaccuracies that appear as multiple lines when a plotter retraces the same path. These problems are common to all plotters, but only modern plotters like the Numonics Model 5400 that use resident intelligence to keep track of these factors compensate for them to prevent plotting quality from being degraded.

PLOTTER INTELLIGENCE

Like the digitizer, today's plotters have their own microprocessors, not only to control the plotter, but to print a complete 96-character ASCII set automatically in varying sizes.

One of the ways in which the microprocessor improves the plot quality in the Model 5400 is through the use of its resident arc routine. All plotters can plot straight-line vectors, and good plotters do this with uniform quality regardless of the angle of the vector. Arcs, however, are generally approximated by a series of short vector chords. These cause frequent plotter starting and stopping, slowing the plotting process and degrading the appearance of the plot. An intelligent plotter can compensate by use of a special command for clockwise and counterclockwise arcs that control the velocity of both motors for precise plotting of the arc desired. The result is a line plotted smoothly, from vector to arc and from arc to arc, with uniform line thickness.

Another useful feature the microprocessor performs is the buffering of data. Many newer plotters have buffers, but most use them only as a way to hold instructions in the plotter when they've been sent before it has finished plotting the last instruction. With a large buffer that can store up to 3800 instruction characters for replotting at any time, title blocks, logos, and other graphic features that are used repeatedly can be stored, and replotted by use of a single instruction.

A key factor in the design of a plotter driven by a stepping motor (Fig 4) is the use of special software to correct and compensate for the irregularities inherent in the motor characteristics. Without these software routines, stepping motors will not produce high quality plots. The stepping motor provides accurate incremental position control with less complexity than the conventional position servo. The motor makes a precise and repeatable increment in position in response to a step command.

The stepping-motor principle uses a salient, pole rotor, and stator construction. Stepping is achieved by sequencing of the field excitation so that the rotor moves one pole pitch. On completion of a step, the rotor remains in a position of equilibrium corresponding to minimum reluctance in the air gap between rotor and stator. When the rotor is displaced from this position, restoring forces are developed in proportion to the displacement. Inertia combines with this restoring force and results in an oscillating system.

If the stepping motor is suddenly stopped, it will tend to oscillate back and forth about the rotor's equilibrium position. If the motor's speed is sufficiently high, it may jump a whole pole pitch and stop one step later than commanded. Similarly, during acceleration, too rapid sequencing of the control windings won't allow time for the motor to complete its first step and the movement will consist of fewer steps than those commanded.

In a plotter, a failure to execute every step-command precisely results in unacceptable performance. A well conceived design minimizes the inertia coupled to the motors and uses current-source excitation of the field windings to reduce the inductive delays in flux buildup. Software routines can then be used to control acceleration and deceleration to avoid any miscounting of steps. ■

L. Robin Hulls, *a director of Numonics Corp., has previously worked with RCA, Sperry Rand, and Canadian Westinghouse. He holds a BS in physics and electrical engineering from Manchester Univ. in England, and an Intermediate BSc in chemistry from London Univ.*

THERE'S MORE SPACE ON JUPITER.

Jupiter 7 widened the horizons of colorgraphics, with its brighter and clearer colors, anti-aliased vectors, high resolution, and ergonomic design.

But now there is our newest model, Jupiter 7 PLUS with a larger screen, more display area (1024x780 pixels) and more standard features. Plus a price that's still very down to earth.

Eight 1Kx1K memory planes are standard on Jupiter 7 PLUS providing 256 simultaneously viewable colors out of a palette of 16.7 million. Further enhancing the display are five anti-aliased, proportionally spaced character sets.

The high resolution and large screen of Jupiter 7 PLUS make it ideal for applications like mechanical design and printed circuit board layout, where complex images must be seen clearly.

But Jupiter 7 PLUS is not just big, it's powerful. It has high speed graphics generation, zoom, pan and scroll, a fully programmable keyboard with dual joysticks, built-in sound synthesizer, and a hardware trilevel blue grid. A pixel block-moving system can shift any operator-defined shape at 1 million bits/second.

In addition, a bi-directional Centronics-compatible interface is available for driving a 125-color ink jet printer.

Jupiter 7 PLUS prices range from $15,200 for 768x585 pixels on a 13" display, to $19,000 for 1024x780 pixels on a 19" display. Besides the display, these prices include the processor with eight planes of memory and a detachable keyboard with dual joysticks. For information on these and other Jupiter products, call or write Jupiter Systems, 2126 Sixth Street, Berkeley, CA 94710. Call Toll Free 800-858-9309. (In California, 415/644-1024.) TELEX 176 248.

19" display, 1024 x 780 pixels

13" display, 768 x 585 pixels

JUPITER SYSTEMS

Computer Graphics, CAD/CAM & Image Processing

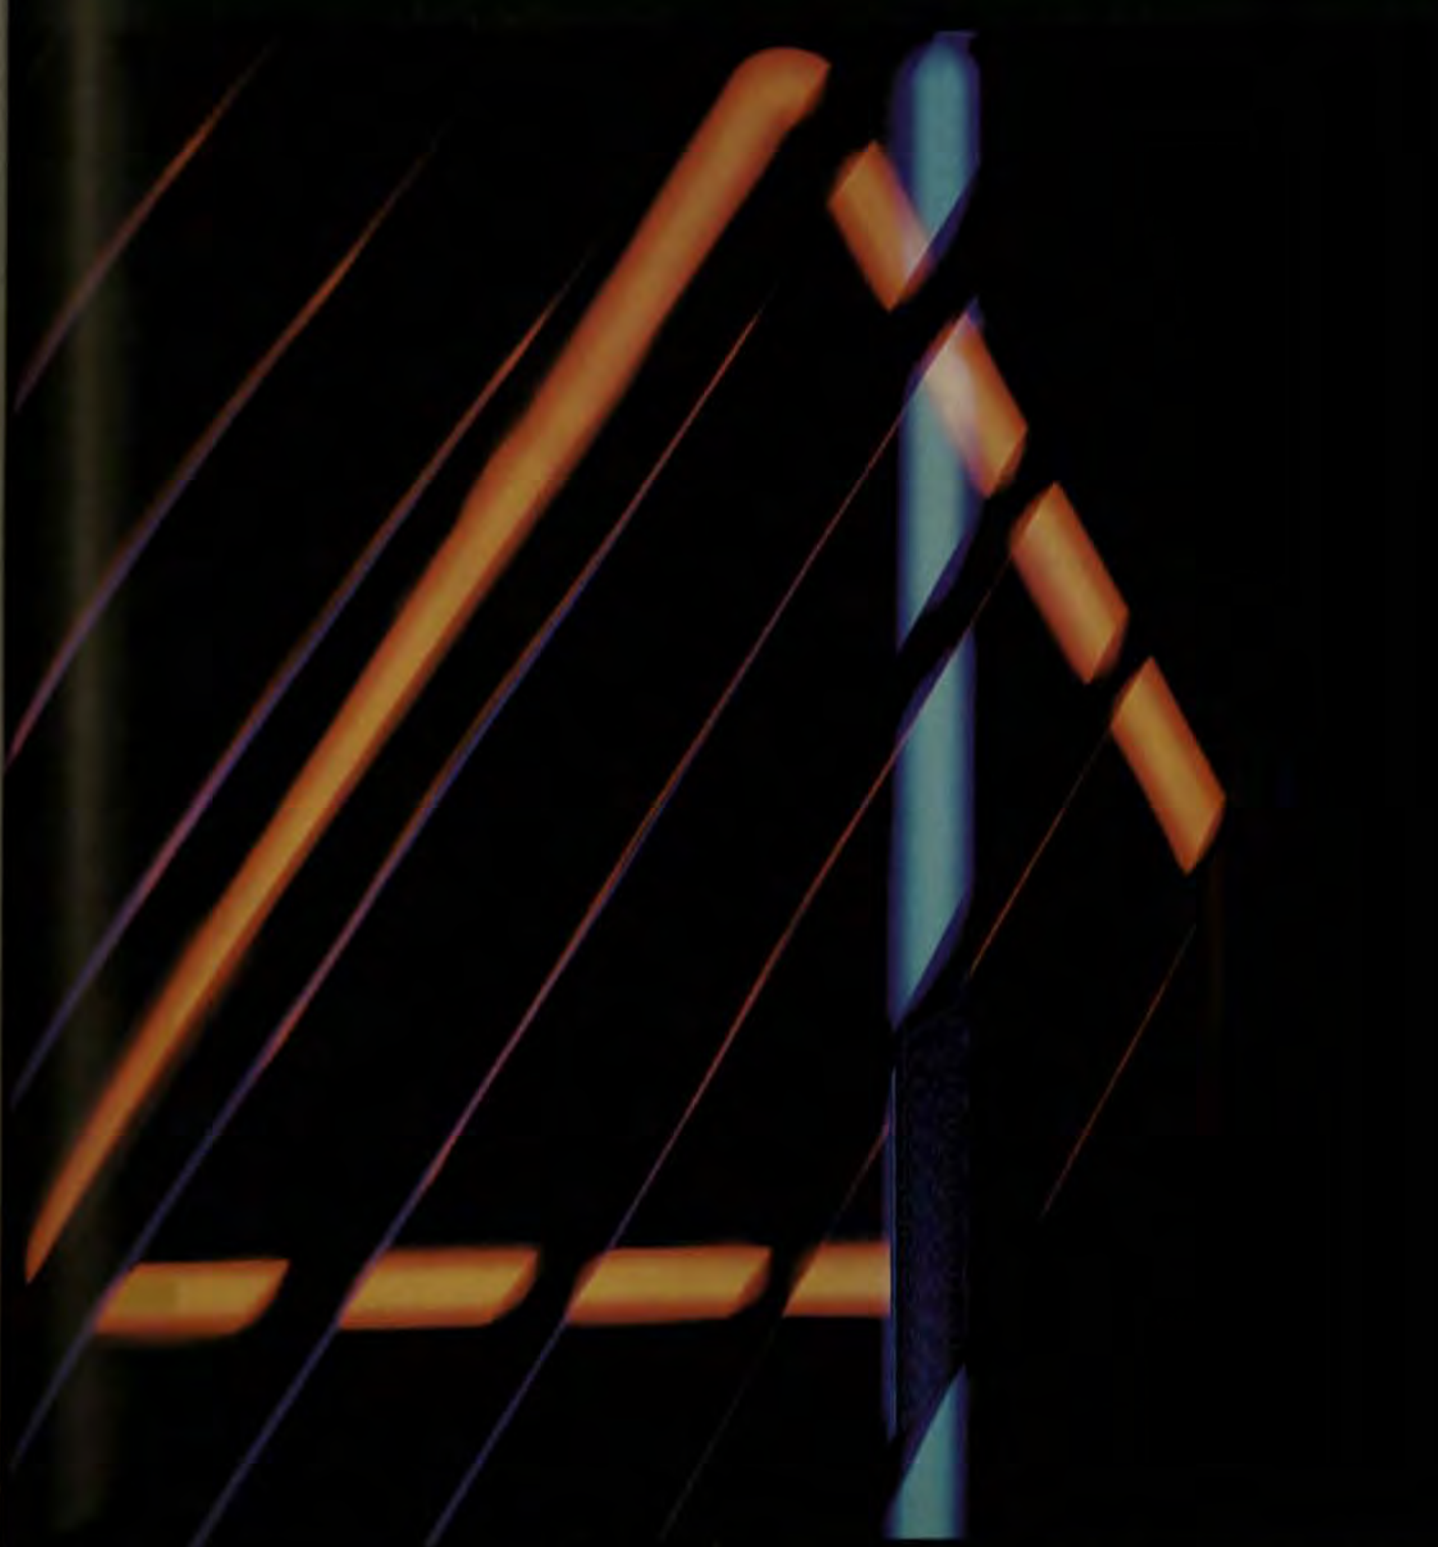

LEXIDATA · 3700

42 million pixels per second

The Lexidata 3700 has the speed and performance you've come to expect from the leader in high resolution displays.

In terms of speed, the 3700 writes in pixel blocks automatically (e.g. for fills and horizontal vectors). This is continuous speed, not burst speed, at 42 million pixels per second.

Its pixel-by-pixel rate is also continuous and fast. Fast when you're generating vectors. Fast for individual pixel update–crucial for imaging applications.

This speed is augmented by state-of-the-art 1280 x 1024 color at a flicker free 60hz refresh.

In terms of performance, the 3700 comes standard with Writable Control Store (WCS), enabling the user to implement his own high speed routines. Additionally, WCS also gives direct access to the main processor, so even complex functions run fast. There's also separate storage for over 1000 characters in user-definable fonts.

Remember, too, that the 3700 stems from the 3000 family, the most popular high-performance line ever developed. The 3700 is a *complete* performance package offering speed and features that live up to your expectations.

The Lexidata 3700. Fasten your seat belts.

For quick response call 1-800-472-4747 (in Massachusetts, call 617-663-8550) or write to us at 755 Middlesex Turnpike, Billerica, MA 01865. TWX 710-347-1574.

LEXIDATA
The clear choice in raster graphics

GKS Graphics Standard Replaces Core System For International Use

Germany's Graphics Kernel System has anticipated today's high-resolution rasters and broad low-resolution applications. SIGGRAPH's Core proposal did not.

by **Sidney Damron** and **Warren O'Buch,** Conrac Corp.

A clear winner can now be declared in the ongoing controversy over graphics standards. Germany's GKS (Graphics Kernel System) will soon be the central document for all other international graphics standards—in a decisive win over the widely implemented SIGGRAPH Core system developed in the U.S. in the late 1970s.

In retrospect, the triumph of GKS was inevitable and serves as still another example of the impact that lower memory costs can have on every aspect of computer technology.

The Core system evolved during a period when CAD/CAM applications dominated computer graphics and stroke-writing CRTs (storage-tube or refresh-vector) were the dominant display devices. Raster displays—limited in their resolution by the high cost of refresh memory—produced too many "jaggies" when diagonal CAD/CAM lines were superimposed on the horizontal lines traced by a raster-type CRT.

In comparison, GKS anticipated the day when lower memory costs would make it feasible to produce higher resolution raster systems such as those based on Conrac's new Series 7300 monitors.

High-resolution raster systems, with over a million stored pixel values on display, plus such added advantages as television-type animation and solid blocks of shaded color, are rapidly superseding stroke writers in all but the most demanding special-purpose applications, such as those involving a very high density of linear information, dynamic image manipulations, and 3-dimensional rotation.

The proposed new standard also anticipated the explosive growth of such new applications as business, educational, and home-entertainment graphics, all of which can live with the lower resolution of a conventional raster-type terminal or even a standard television monitor.

A LEAN AND CLEAN STANDARD

By emphasizing the special capabilities of raster-type displays and hardcopy devices through such raster-oriented graphic primitives as color-cell arrays for image-processing applications and areas filled with solid or patterned color (Table 1), GKS has proved to be the right standard at the right time.

A number of other features have contributed, however, to the success of GKS. To expedite its adoption as a standard, for example, authors of GKS chose to keep it "lean and clean." In its present form, it applies only to 2-dimensional graphics—which happens to coincide with the fastest growing computer-graphics applications (business graphics, for example).

In contrast, the Core system, reflecting its CAD/CAM origins, has complicated functions for viewing 3-D objects from any point in space. Two-dimensional X-Y implementations of the Core system are obtained by setting the Z or depth dimension to a constant value, typically zero.

Actually, in terms of today's technologies and applications, GKS occupies a middle ground. It's still too rich for many of the newer microprocessor-based systems—such as per-

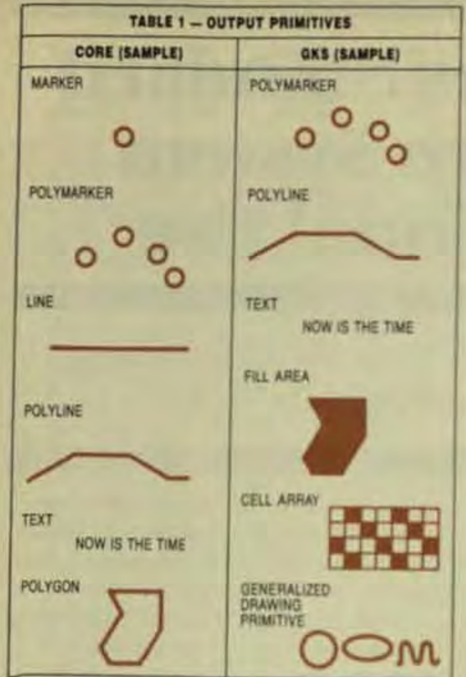

sonal computers, intelligent terminals, and desktop business systems—that have only an incidental need for graphics.

An ANSI committee that had been working on a simpler standard for such minimal applications has now integrated its work into a proposed new "m" level that is one step less demanding than the currently defined minimum GKS graphic-output standard.

Any new 3-dimensional standard that comes out of the international standards-writing organizations should also be compatible with GKS—producing a spectrum of common-denominator standards to meet the needs of every graphics application, from the simplest to the most sophisticated. The authors of 2-D GKS intended from the start to expand it to a 3-D system and an ISO (International Standards Organization) committee is now working on the project.

Core and GKS have common roots and share a number of basic concepts that are now firmly entrenched throughout the graphics community. Two-dimensional GKS is, in fact, an evolutionary product of 3-D Core—building on its successes, sidestepping its faults and weaknesses.

LOGICAL WORKSTATIONS

One of the pitfalls GKS has not entirely avoided is Core's relative inflexibility. Core implementation standards were set at overly rigid levels, and binding to a higher-level language, such as FORTRAN, was never defined—inviting noncompliance. Moreover, having never progressed beyond a committee proposal, Core has lacked the authority of a certified ANSI standard. Hence, the abundance of systems that have been "in the spirit of" or "compatible with" Core.

The American version of GKS, now in the process of adoption, has added a degree of flexibility to the original GKS structure. Work is also progressing on a standard binding (subroutine names and calling sequences) to FORTRAN. But even the original version of GKS took into account the fact that cost and performance considerations will always dictate a variety of graphics-systems implementations. Moreover, with the rapid development of new devices for interactive input and graphics output—many with unforeseen capabilities—ways have to be found to provide device independence without limiting devices to an arbitrary set of I/O functions and parameters.

The principal GKS vehicle for achieving this goal is the logical workstation—an elegant compromise between true device independence and a system in which each implementation defines its own I/O parameters.

The logical workstation is a software abstraction with a uniform, device-independent interface to the balance of the system. But whether or not a specific workstation will accept, reject, or substitute default values for parameters passed along by the system is another matter. Input/output capabilities and parameter limits are defined by entries in a workstation list derived in turn from a workstation description table that is 100% implementation-dependent (that is, arbitrarily set by the system implementer to reflect the characteristics of a particular device or set of devices).

Extensive inquiry functions allow the application program or operator to examine the entries in the workstation lists and adjust, if necessary,

TABLE 2 — LOGICAL-WORKSTATION CATEGORIES	
OUTPUT	OUTPUT ONLY
INPUT	INPUT ONLY
OUTIN	OUTPUT AND INPUT
WISS	WORKSTATION INDEPENDENT SEGMENT STORAGE
MO	GKS METAFILE OUTPUT
MI	GKS METAFILE INPUT

TABLE 3 — LOGICAL INPUT-DEVICE CLASSIFICATONS		
INPUT-DEVICE CLASS	FUNCTION	TYPICAL DEVICE
LOCATOR	PROVIDES A POSITION IN WORLD COORDINATES	DIGITIZER TABLET
STROKE	PROVIDES A SEQUENCE OF POINTS IN WORLD COORDINATES	MOUSE
VALUATOR	PROVIDES A REAL NUMBER	ROTARY SWITCH
CHOICE	PROVIDES A NONNEGATIVE INTEGER THAT REPRESENTS A SELECTION FROM A NUMBER OF CHOICES	FUNCTION KEYS
PICK	PROVIDES A PICK STATUS, A SEGMENT NAME, AND A PICK IDENTIFIER	LIGHT PEN
STRING	PROVIDES A CHARACTER STRING	KEYBOARD

to both the limitations and the capabilities of specific devices connected to the system.

GKS provides for six logical-workstation categories: output only, input only, output-input, and three specialized workstations for the storage and transfer of graphical information (Table 2). Within each category there can be any number of logical-workstation types, such as CRT displays or pen plotters in the output-only category, and any number of different, device-specific workstations within each type.

There are restrictions, however, on the number of physical devices a logical workstation can support. Output-only or output-input workstations are restricted to one display or recording surface. Input-only or output-input workstations can support any number of interactive-input devices in each of the classifications listed in Table 3, but an output-input workstation must have at least one device in each class. (Two or more device classes may be implemented by a single physical device, such as a keyboard.)

Thus, implementation of even a single-user graphics workstation may require two or more logical workstations to support, for example, a CRT/keyboard terminal (output-input) and a hard-copy device (output only). Systems with multiple users would have to support a proportionately larger number of logical workstations, although such devices as printers and cameras would typically be timeshared by periodically activating and deactivating logical workstations based on a single set of workstation description tables.

MATRIX OF IMPLEMENTATIONS

The logical workstation approach gives the system implementer a great deal of flexibility but could also lead

Color graphics from Intecolor makes sense to some very intelligent people.

For some very sensible reasons.
Color is the language of the human eye. Research has repeatedly confirmed that color displays convey more information, more quickly, and with greater comprehension than any other visual method. We learned that fact early on and at Intecolor, color is the only language we speak.

Competence. In 1974, the world's first color terminal using a microprocessor was introduced, the Intecolor 8000. Since that time, we have built quality and reliability into tens of thousands of color graphics computer products.

Now a world leader in computer graphics, Intecolor provides a complete line of color graphics desktop computers, CRT display terminals, and graphics work stations. And Datavue multi-user processors—superior price/performance "engines"—driving multi-user systems with user-friendly efficiency.

Value. Today, Intecolor retains its color graphics leadership with an aggressive, singular objective: To produce quality products substantially below competitors' prices, from basic color terminals to sophisticated, high-resolution, dot-addressable color graphics computer systems.

Commitment. Check the customer names in the graph above. These are just a few who decided that Intecolor makes sense. Join them. You'll be in pretty smart company.

Intecolor
AN INTELLIGENT SYSTEMS COMPANY
Intecolor Drive
225 Technology Park/Atlanta
Norcross, GA 30092
404/449-5961 • TWX 810-766-1581

to so much diversity that the main point of having a software standard—program portability between systems—would be lost.

GKS resolves this problem by creating, in effect, nine separate implementation standards within a single standard (Table 4). The nine upwardly compatible standards form a matrix based on three graphic-output levels—0, 1, and 2—and three interactive-input levels—a, b, and c. The addition of the proposed minimal "m" output level in the American version of GKS will increase the number of standards to 12.

Each of the standards represents a cohesive graphics system with predictable performance characteristics. Strict compliance is expected from any system that claims to meet the explicitly defined and required capabilities of a specified set of GKS output and input levels. As modified in the American version of GKS, however, these requirements are a floor, not a ceiling.

A system can meet the standard by providing all the required functions for a specified level but not all the functions for the next higher level. The result is a nearly continuous upward migration path for system performance improvements as new technologies or equipment become available. Yet users have the assurance that any applications written for a less capable installation will still be executable as graphics systems continue to become more powerful and sophisticated.

The lowest present GKS output level, 0, is designed for relatively simple systems with only a limited need for such display image transformations as scaling, rotation, and translation (from one location to another). GKS restricts within-the-picture image transformations to the "segment"—a group of graphic primitives that define an object or related elements within several objects. No segment functions are required at output level 0 (or the new level "m"). Users must therefore generate pictures based on just the GKS primitives shown in Table 1, without any opportunity to alter their relative positions except by starting the picture-building process over again.

Two other limitations apply to output level 0. The standard doesn't allow more than one active workstation at a time. Thus, if a physical workstation is equipped with both a CRT display and a hard-copy device (each requiring its own logical workstation), output to both devices will require the application program or user to build the picture twice, once on the display and again as hard copy.

Another level 0 limitation is that no functions are required for altering the primitive-attribute bundles in the workstation description tables or for adding new bundles to the workstation lists. The only way the appearance of a primitive can be altered is to start the picture over again and specify a different predefined bundle (or a different set of direct attributes).

Output level 1 removes these limitations. Any number of logical workstations can be active at the same time, allowing you to output a picture to several different devices without regenerating the picture for each output. Full segmentation facilities are also available, permitting a full range of picture transformations. And functions are provided for the changing or creating of attribute bundles.

But output level 1 is still limited in one important respect: Segment data is stored only in the logical workstations that were active at the time the segments were created. There are no facilities to let you transfer the picture descriptions from one workstation to another, or to store them for future use by the same workstation.

Output level 2 resolves these difficulties by providing a specialized logical workstation called a WISS (Workstation-Independent Segment Storage). Pictures can be built in WISS at the same time they are being constructed in any other workstation. Stored in WISS, individual picture segments can then be transferred to other workstations or returned to the original workstation at a later time.

TABLE 4 — GKS STANDARDS MATRIX

OUTPUT LEVELS	INPUT LEVELS a	b	c
0	NO INPUT. ONLY PREDEFINED BUNDLES, MULTIPLE NDC TRANSFORMATION FUNCTIONS. ALL OUTPUT FUNCTIONS. METAFILE FUNCTIONS OPTIONAL.	REQUEST INPUT. MODE SETTING AND INITIALIZE FUNCTIONS FOR LOGICAL INPUT DEVICES. NO PICK.	REQUEST, SAMPLE, AND EVENT INPUTS. NO PICK.
1	FULL OUTPUT, INCLUDING FULL BUNDLE, MULTIPLE WORKSTATIONS; BASIC SEGMENTATION EXCEPT WISS. METAFILE FUNCTIONS REQUIRED.	REQUEST INPUT, INCLUDING PICK.	REQUEST SAMPLE, AND EVENT INPUTS, INCLUDING PICK.
2	WORKSTATION-INDEPENDENT SEGMENT STORAGE (WISS).	REQUEST INPUT, INCLUDING PICK.	REQUEST, SAMPLE, AND EVENT INPUTS, INCLUDING PICK.

INTERACTIVE INPUTS

A similar progression applies to the three GKS interactive-input levels. The lowest level, a, doesn't require any operator interaction at all. The system is designed only for graphics output (at output level m, 0, 1, or 2).

Examples would be a process control or annunciator display application in which it would be inadvisable to allow a control-room operator to change the color coding that indicates whether a valve is open or shut. Only the application programmer should have access, in this case, to the form or appearance of the displayed images.

Interactive input starts at input level b (output level m, 0, 1, or 2). Again, however, the application is in control of the interaction. Input devices are limited to a request mode that permits the user to interact with the system only when the system has made a specific request for information.

Input level c (output level m, 0, 1, or 2) adds two new capabilities. The application program can sample an input device at any time, obtaining information without waiting for the user to respond. The user can also input information at any time, independent of the application program, by generating "events" that are placed in a queue and processed at intervals by the application program. In both cases, the input device must be in either sample or event mode for the application program or user to initiate the interactive input.

All three input modes—request, sample, and event—are applicable to all six of the input device classifications listed in Table 3. One of the input-device classifications is not applicable, however, to output levels m and 0. Only segments can be "picked" by such pick-type devices as light pens or cursor controls, limiting pick functions to output levels 1 and

REAL GRAPHICS POWER FOR THE Q-BUS!

ONE CARD DISPLAYS 256 COLORS. THE OTHER HAS HIGH RESOLUTION "WINDOWING" CAPABILITY. EACH IS AMAZINGLY COMPACT. AND YOU WON'T BELIEVE THE PRICE.

We made them small (just 5.2"x8.9").

We made them powerful.

We made them easy to use (with a complete library of graphics utilities).

And above all, we made them inexpensive, especially when you consider the alternatives.

Since 1979 Peritek has been a leading supplier of high-performance modular display components that plug directly into the Q-bus.* But we've never been more excited than at the introduction of our H-Series products.

Our **Model VCH-Q** dual high card can simultaneously display 256 colors from a palette of 16 million via Peritek's random access look-up table.

We packed lots of features into this board to make applications development easy. For example, under program control you can select alternate color map settings on a frame-by-frame basis so that special effects, animation and so on are possible. You can also double the number of apparent colors by displaying from alternate maps on adjacent pixels. Peritek's extensive software will provide all the support you need.

The **Model VRH-Q** graphics and character video display interface produces both a high resolution 1024 x 1024 monochrome bit-mapped graphic image and an overlaid 64 x 128 alphanumeric display.

Because the character sets are stored in random access memory, you can download an unlimited library of characters and symbols from a host computer.

The VRH-Q can simultaneously display several processes on different parts of the screen, making it particularly suitable for interactive applications. Typical uses include CAD/CAM, robotics and vision systems, and high resolution work stations.

EXTRAORDINARY COST EFFECTIVENESS

If you're presently building or using a Q-bus machine, you have a substantial software investment in your system architecture and instruction set. Doesn't it make more sense to simply add a compact plug-in board for $6,000 or less, rather than write more code or pay over $10,000 for a serial device?

Price of the VCH-Q is just $2995, with even lower prices in OEM quantities. The VRH-Q is also $2995. In addition, Peritek offers a variety of accessories and other Q-bus hardware to help you meet your design goals affordably.

For more information call or write Peritek at 5550 Redwood Road, Oakland, CA 94619. (415) 531-6500. TWX: 910-366-2029.

Peritek

*Q-bus is a trademark of Digital Equipment Corporation.

PrintaColor's new TC1040 is the perfect printer for your color graphics system. Hardcopy can be generated by virtually any computer or color display terminal—faster and easier than ever.

Color fast and picture perfect
The TC1040 was designed with full image-seeking and bidirectional capabilities that allow it to produce an 8½" × 11" image in one and a half minutes.

With resolution of up to 120 dots per inch, your images are printed crisply and clearly every time on a variety of media from 8½" to 14 7/8" wide.

4,913 choices of color
TC1040 ink-jet technology produces sharp text and precise graphics in just about any color you can imagine via PrintaColor's versatile method of ordered dithering.

Quiet and simple
Non-impact, whisper-quiet operation makes the TC1040 ideal in any office setting. Snap-in, disposable cartridges contain specially formulated inks that dry instantly on plain paper.

Built for long life
The TC1040 has a performance rating over 3,900 hours MTBF. The drop-on-demand printhead has a proven reliability of more than 10 billion operations. Its design simplifies routine maintenance checks and minor adjustments.

And the TC1040 has all applicable regulatory agency approvals.

The Logical Choice
For greater speed, more colors and high quality resolution in a color printer, it's PrintaColor's TC1040—the new color copy solution—for text, image and graphics.

THE NEW PRINTACOLOR

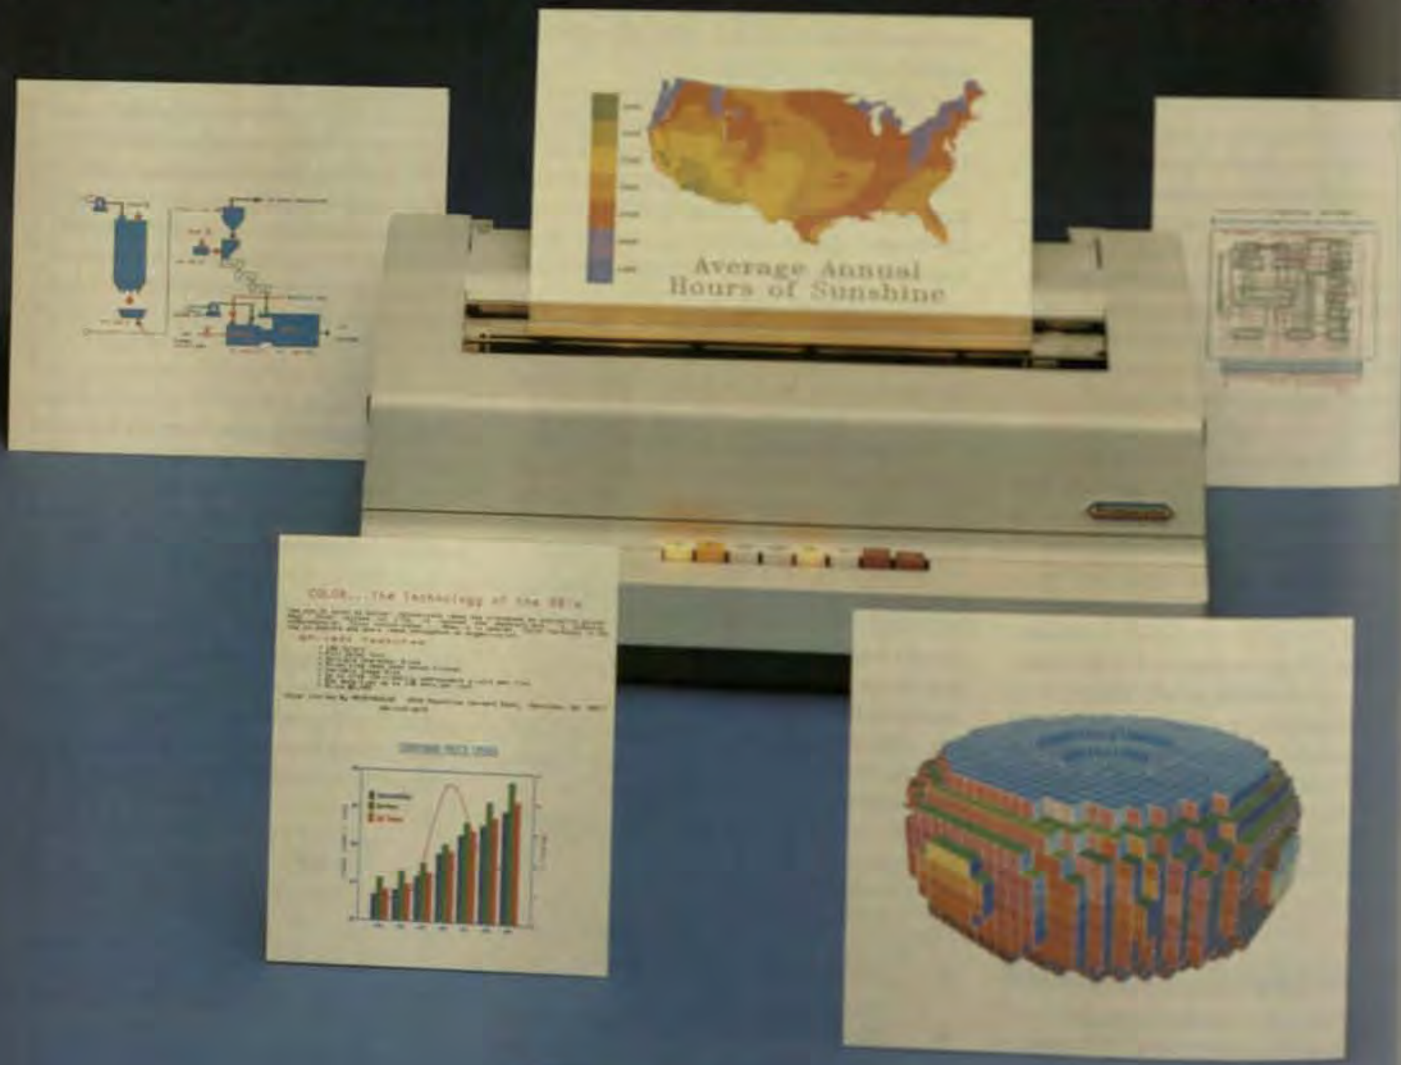

5965 Peachtree Corners East

PrintaColor CORPORATION

Norcross, GA 30071 (404) 448-

Both Vector and Raster Displays Show Strong Future as Costs Decline

These two display methods have emerged as the dominant display technologies in graphics. Each has attributes that dictate which is used for a given application.

by **Fred Zeiler,**
Vector General Inc.

Today, there are two principal ways to produce a picture on a cathode-ray tube. The vector display (also called the stroker display or calligraphic display) uses the electron beam directly to trace out the vectors that make up the drawing. In contrast, raster displays scan the entire screen repetitively to illuminate picture-element dots (pixels).

Where low resolution is adequate, as in business graphics, the raster display is dominant. For higher-resolution applications, such as computer-aided design and manufacturing (CAD/CAM) and command control, both technologies are widely used.

Color capabilities for both display types are comparable, and their costs are declining. Both technologies can be expected to hold their market share during the next five years, although rasters may dominate after that.

VECTOR DISPLAYS

In vector displays, the electronic beam is deflected continuously to trace out the image on the screen (Fig 1). After the beam has generated the complete image, it repeatedly refreshes the image before the phosphor fades. The displayed image is stored in memory and the refreshing is controlled by a special-purpose processor.

Since only the end points of each vector must be calculated, there is no need for high-speed processing, and memory requirements are modest. A reasonably complex drawing of 20,000 vectors, generated on a display of 4K-by-4K addressability, maintains high resolution with a memory of less than 500 Kbits.

Moreover, it permits a high degree of interactivity. A user can control the image in real time with interactive devices such as joysticks, light pens, and digitizing tablets. In addition, the user can erase portions of the display selectively, cause key features to be blinked, and magnify, rotate, or translate images on the screen.

Although most vector-display systems today are monochromatic, they can be designed to produce color images by modifications to a standard 3-gun color tube. A mask overlays the pattern of the color triplet of the three primary phosphor colors deposited on the CRT face. Phosphor masking ensures that the electrons striking the screen are striking precisely the desired locations. By adjustments to the intensity of three primary-color electronic guns, individual vectors can be drawn in an unlimited number of colors.

Traditionally, vector displays have suffered from flicker. Flicker occurs whenever unusually complex displays cause the refresh rate to drop below 40 Hz. However, modern high-speed analog circuits and deflection techniques have been developed to produce writing rates that make flicker a rare occurrence. Over 60,000 short vectors, for instance, can be painted within 0.024 s.

RASTER DISPLAYS

Raster-scan deflection circuits are designed to trace a fixed pattern of par-

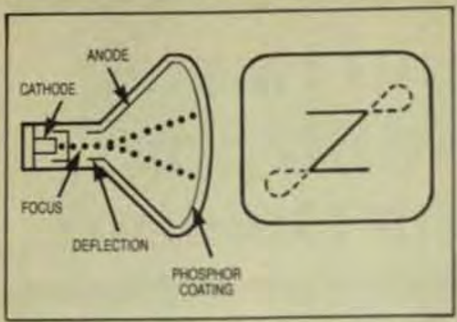

Fig 1 Deflection amplifiers in vector displays move the beam to trace out the image lines directly on the screen.

allel lines on the screen (Fig 2). A modulating signal changes the beam intensity at the proper point on the screen to illuminate each pixel selectively according to the pattern stored in memory. Each displayable point requires at lease one corresponding bit in memory.

The scanned lines are volatile, so the complete raster pattern must be redrawn at a repetition rate (typically 60 times a second) above the flicker level of the human eye. Unlike vector displays, however, the raster monitor is unaffected by the complexity of the drawing to be displayed. Since it addresses every pixel location on the CRT with every refresh cycle, it can portray solid or continuously varying areas of color or gray-scale intensity as well as lines and dots. Called fill, this capability is extremely useful in certain applications.

The principal drawbacks to raster displays have traditionally been low interactivity and resolution. To move a line can require recalculation of hundreds of pixel values. Practically speaking, raster displays have always been limited in the amount of displayed information the operator can alter in real time with interactive commands.

The low resolution that has marked raster displays to date is the result of the stair-stepping lines inherent to the pixel matrix. The image lines on the screen often appear jagged instead of smooth and continuous. Anti-aliasing algorithms somewhat improve the situation, but they incur additional hardware costs and produce thicker lines.

Resolution can be improved substantially if you increase the number of pixels on the screen (Fig 3). However, doubling the resolution quadruples the number of pixel values that must be calculated, stored in memory, and transferred to the monitor screen. A 1024 × 1024 array has over a million pixel locations and, hence, a 1-Mbyte memory requirement.

For this reason, 1024 × 1024 resolution has been the practical limitation for a raster display. To equal the 4096 × 4096 resolution of a vector display would require 16 Mbytes of faster display memory, a wider monitor-interface bandwidth, and a faster scanning monitor. For this reason, raster-scan users tolerate the jagged lines where their applications permit this compromise.

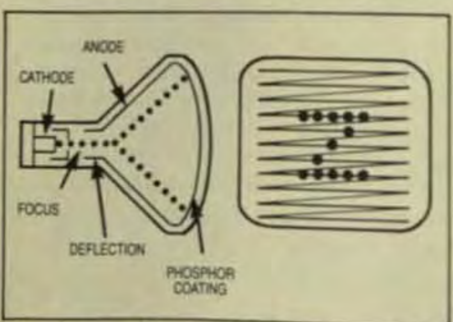

Fig 2 In raster displays the beam sweeps line by line across the entire screen. The beam intensity increases at each pixel that must be displayed.

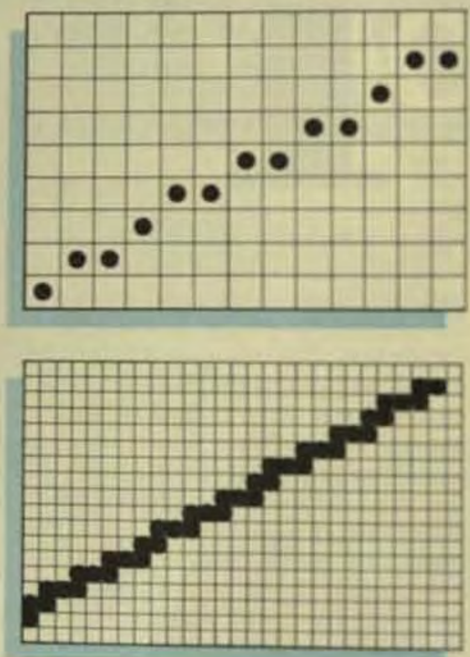

Fig 3 The staircase effect of raster displays can be minimized by an increase in the size of the pixel matrix.

SELECTION CONSIDERATIONS

With today's state of the art, there are significant differences that dictate the choice of one or the other display type for each given application. Primary selection issues are:

- Resolution
- Interactivity
- Area fill
- Color
- Cost

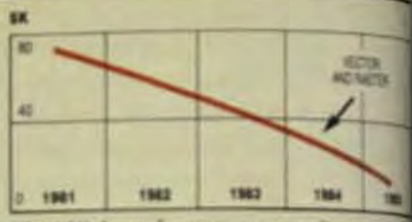

Fig 4 High-performance, monochromatic ras- ter and vector displays cost about the same to- day, and the costs of each are declining sharp-

If you require high resolution, ve- tor display is the clear-cut choice. The vector display's resolution of 4096 × 4096 elements on a typical 12-in. screen is four times the amount that's practical with today's raster technology. Vector displays also have significant advantages on interactivity. They give you far more speed and flexibility in manipulation of displays and communication via light pens or data tablets.

Conversely, only raster displays can provide area fill. Current vector displays are unable to display shading within the outline of an object or to display thicker lines.

From a technical standpoint, the color capabilities of the two display types are comparable. Both can display at least 16 different colors per frame from color palettes as large as 16 million colors. However, the cost of color has been far less with raster technology. Whenever color was an absolute requirement, raster displays were invariably selected because of cost.

Now, though, the cost picture is changing, and for black-and-white displays the change has already occurred (Fig 4). Since 1981, the costs of high-performance raster and vector black-and-white displays have been virtually identical. Equally significant, costs of both have been declining rapidly. Products now in the late design stage will sell for about $25,000 per workstation—one-third the cost of a comparable display just two years ago. Moreover, by 1985, black-and-white display systems will be selling at $10,000 per display station for either technology.

A similar decline has occurred in the costs of color display systems (Fig 5), with the most dramatic reduction occurring in vector-display systems. In 1981, a color display using vector technology not only was hard to procure, it often cost over $100,000 per workstation. The high-resolution raster, on the other hand, had a much

The Image Solution

The CAT 1600 Series lets you take full advantage of your color graphics potential.

Plug this powerful color video graphic system into your IEEE-696 bus and watch your computer open its eyes. Exercise your creativity developing new ways to study your world and discovering the flexibility of video imaging. Our real time frame grabber gives you instant availability of the image to be processed. The CAT 1600 is the creative link between machine, man and the world.

Resolution is the name of the game, and we've got it. Physically you're looking at 512 × 512 pixels up to 24 bits deep. And that's real color. Now, center on a pixel, any pixel, and roam the screen through an image space as large as 1K × 2K. Zoom in and explore a close-up of 32:1, not in the usual quantum leaps of integer zooms, but in smooth logarithmic steps of 1.1%. A smooth zoom... that's human engineering.

At the heart of the matter is a dedicated 8086 image processor. It blazes a 16 bit wide path through the various memories, lookup tables and image parameters as it executes high level commands from your host processor. Up to 48K of static RAM makes the image processor useful for downloading custom programs from the host.

When it came to adequate memory, we didn't forget. 768KB of dynamic memory gives you plenty of image. Our PROMs have a library of 64K organized into over 130 sophisticated graphics commands such as continuous live digitization, character and shape generation, global image manipulation and animation effects, to relieve the host computer from low level primitives.

At your disposal is a palette of 16.7 million colors and 256 shades of gray. Quantized lines or free-hand sketching completes the picture. Use a variety of pen widths, brush strokes, even airbrush! Now imagine what you can do with a superb quality image captured in real time from a color video camera. Contact us for an eye opening demonstration: 935 Industrial Avenue, Palo Alto, CA 94303. 415/856-2500.

DIGITAL GRAPHIC SYSTEMS, INC.

What You See Is What You Get.

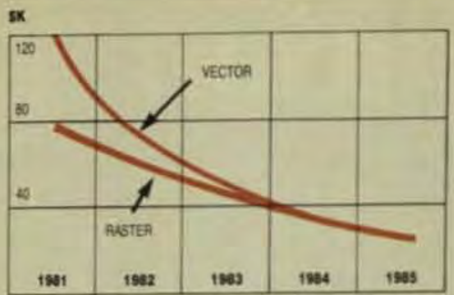

Fig 5 High-performance color vector displays have been quite expensive, but by the end of 1983 they should cost no more than raster displays.

more favorable cost—somewhere around $60,000 per workstation, declining to about $40,000 today.

At the same time, significant breakthroughs in the vector stroke-writing class of display, primarily in the processing elements, have produced a steep decline from the $100,000 model. This year, such a display system will be available for under $40,000. The price trend then will continue as for rasters, and by 1985 either type of display station will be priced somewhere around $25,000. Cost is no longer a significant issue in the choice of technology.

APPLICATIONS

With all these factors in mind, it's easier to see which type of display technology is likely to be used for the leading computer-graphics applications. In the case of command and control, the requirements can be equally well met by either technology except for high resolution and interactivity. Vector displays have significant advantages in each of these, so they are the logical choice.

The same situation exists for computer-aided engineering (CAE), typified by finite-element modeling and structures. With today's comparable price structure, the vector display is the preferred choice.

For simulation, the situation is quite different. Simulation requires area fill that can be performed today only by raster displays, and their resolution and interactivity capabilities are quite adequate for this application. An almost identical situation exists for imaging applications. Again, the choice today is raster technology.

Computer-aided design (CAD) is a more complicated issue. For the large semiconductor market, area fill is a necessity. Since the requirements for resolution and interactivity are not high, raster displays are today's most rational choice for this segment of the CAD market.

On the other hand, mechanical CAD/CAM used in the aerospace, automotive, and heavy-industry markets doesn't need area fill, but does require high resolution and interactivity. For this important market, vector displays are the predominant choice.

The projections shown in Fig 6 are used to determine the market share that each of these applications commands today and is likely to command in five years. The prediction is unusual: Each of these applications will maintain its market share.

This projection implies that vector will be the dominant technology five years from now. Analyses of each of the applications indicate that raster display is the clear-cut choice in business, research, simulation, and a

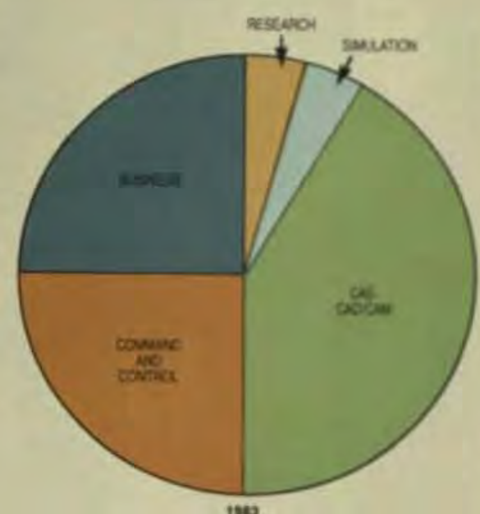

Fig 6 By 1987, the market for interactive computer-graphics systems will exceed $10 billion, yet the proportional usage will be the same as it is today.

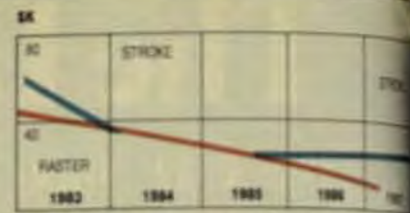

Fig 7 By 1986, memory and VLSI costs will have declined to such an extent that high-performance color raster displays will command a significant competitive edge.

segment of the CAE-CAD/CAM markets. Vector display, conversely, is the logical choice for the large market segments of command/control, CAE, and mechanical CAD/CAM.

But by 1987, raster, not vector, will be emerging as the dominant display system in the high end as well as the low end of the market, because a number of technological developments will alter the current situation significantly. These changes include the use of LSI and VLSI circuits, the constantly decreasing cost of memory in particular and of all electronics in general, and the higher volumes being experienced by the CRT industry.

In each of these cases, performance is being increased at a corresponding price reduction, and a high-performance raster-graphic system should result by 1985. Rasters with a resolution of 2048 × 2048 pixels will be produced economically, and a display that has a resolution of 2K-by-[illegible] pixels will meet the resolution requirements of almost any user.

These same development trends will also begin to give raster a new competitive edge in cost (Fig 7). By 19[illegible] the cost decline of vector displays will have about run its course, but raster displays will benefit from the continued decline of semiconductor prices as new, more dense VLSI and memory devices are developed. Th[illegible] by 1987 high-performance, full-color raster-display systems will sell at about $10,000 per workstation.

Clearly, though, both technologies will have a place in the 1980s. After all, even a minority share of a $10 billion market is significant, and both systems will have a strong market for years to come. ■

Fred Zeiler, at Vector General, is responsible for the program management of the VG 8250 and the VG [illegible] refresh stroke and raster graphics display systems predominantly used in the CADAM and CAD/CAM marketplace. He holds a BSEE from Case Western Reserve Univ.

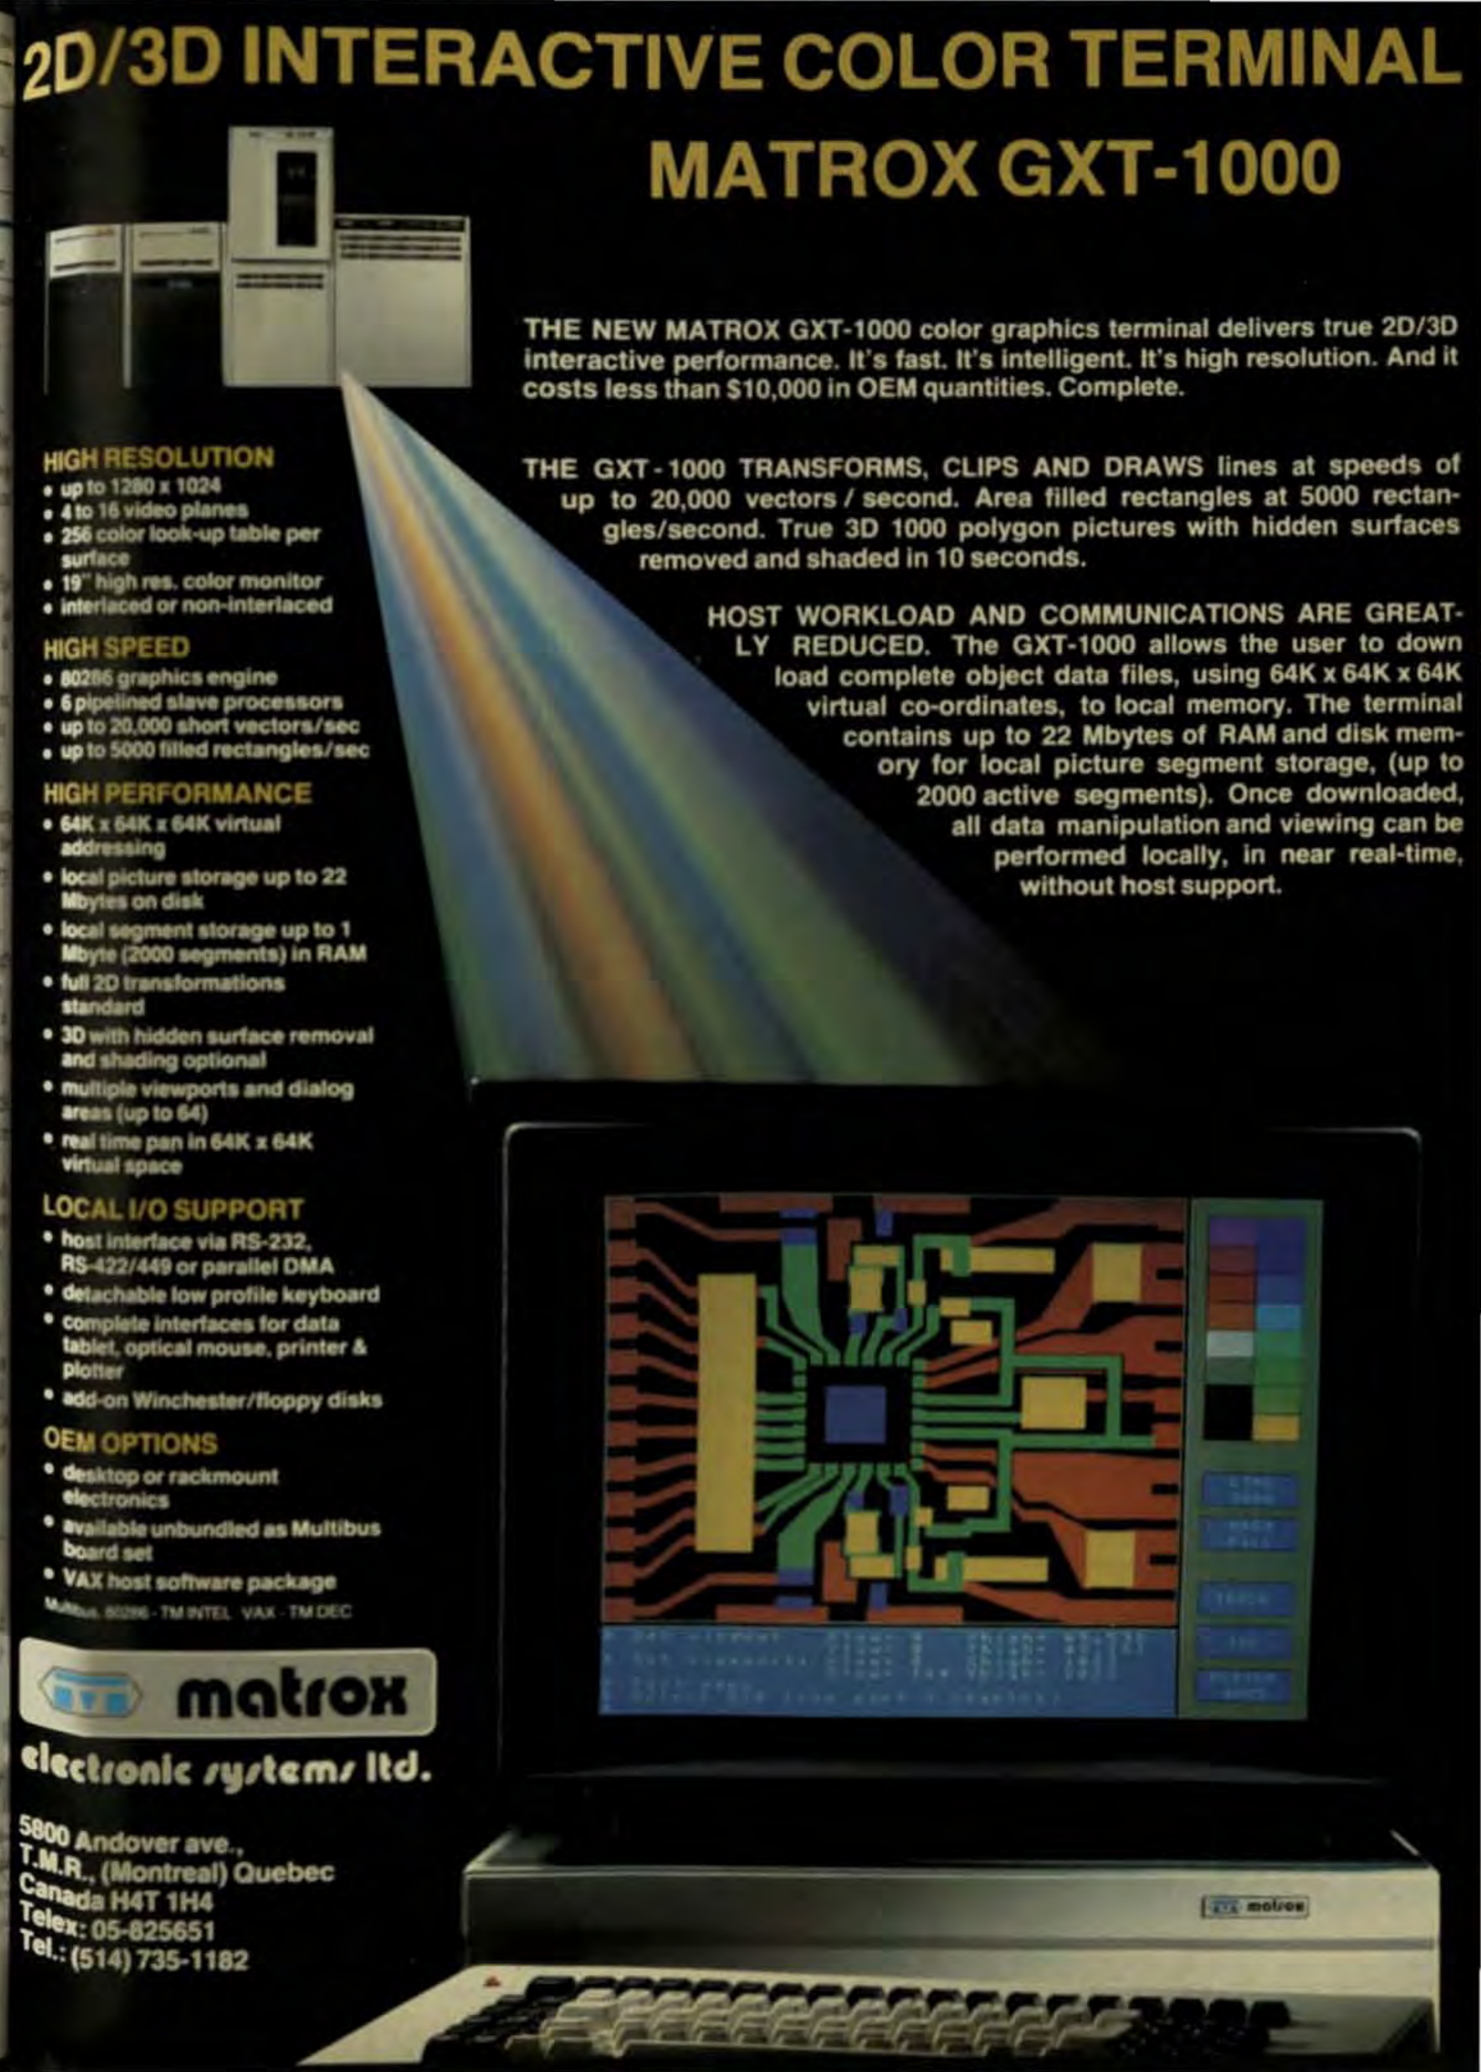

Antialiasing Provides Smoother Displays for Industrial Graphics

By smoothing out the jagged edges in raster displays, antialiasing techniques increase the value and viability of graphics in industrial applications.

by **Thomas H. Birchell, Stuart R. Dole,** and **Victor R. Gold Jr.,** Peritek Corp.

Antialiasing is a technique for "de-jagging" graphic images so that the boundaries of screen images look smooth, rather than stair-stepped or jagged. You find these jagged boundaries on raster graphics displays because the picture elements (pixels) on a CRT screen are organized in rows and columns that can only approximate diagonal and curved lines, giving them a jagged edge.

Storage-tube graphics can display smooth lines of any shape directly on the screen, but storage-tube technology hasn't dominated the market. Raster graphics holds two key advantages:

- It easily and naturally displays color or gray-scale images.
- It is indirectly supported by the television industry, which has developed low-cost monitors and equipment.

Antialiasing provides a bridge that combines the advantages of both technologies. With antialiasing software, moderately priced CRTs can display both a smooth line and color or gray-scale images (Fig 1).

SMOOTH CURVES AND TRUE DIAGONALS

In developing any raster-graphics system, the designer must take into account the human element. The eye integrates the pattern of pixels on the screen into a continuous-tone video image or a solid line. Either is an illusion in terms of the actual arrangement of pixels on the screen, but to the human viewer, the images look real.

For purposes of describing a dejagged or antialiased line, the ideal line will be considered the true image and any variation from the ideal will be called noise. Most raster-scan images include a certain amount of noise because the pattern of pixels on the screen can't directly duplicate the true image. However, when the noise level drops below the threshold that the human eye can distinguish, the image appears true and realistic. Below this threshold, the image is more pleasing, and the viewer can assimilate information quickly with less fatigue.

Antialiasing creates the illusion of a smooth, solid line by manipulating the intensity or hue of the pixels near the ideal line. A pixel that falls directly along the ideal line is displayed at the line's true intensity or color. A pixel that falls near the ideal line is displayed at an intensity intermediate between the background and that of the line (Fig 2).

The part of the pixel that falls outside the true image constitutes noise on the display. But, because of the lower intensity of the color and the eye-brain threshold, the image appears true to the brain of the viewer and the illusion of a smooth line is created.

Under normal viewing conditions, the human eye can't distinguish more than 64 levels or shades of gray on a video monitor, and just 16 levels of gray will produce acceptable images. As an example, the Ansel Adams zone technique for black-and-white photography uses only 13 or 14 levels of gray. A surpris-

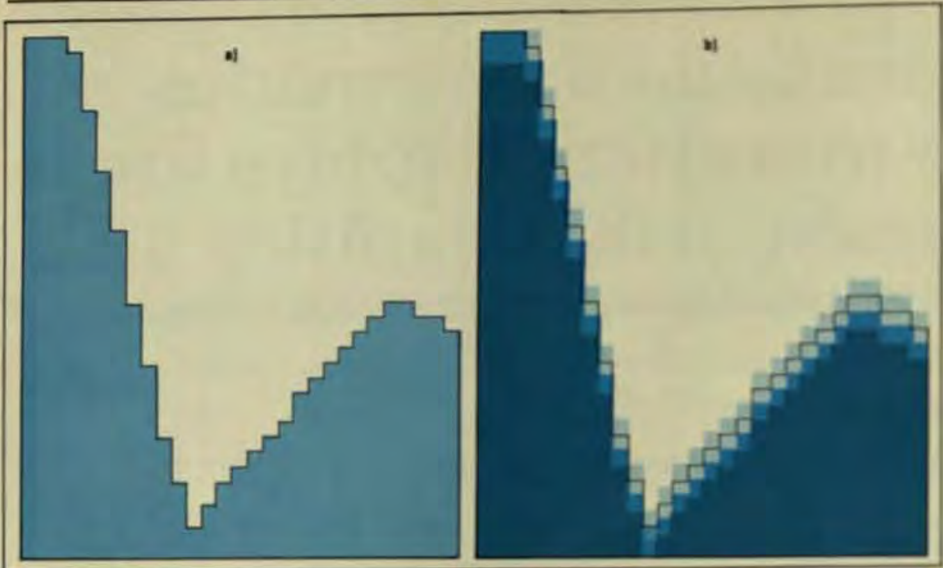

Fig 1 A typical business-presentation chart (**a**) looks better with antialiasing (**b**). In this example, four levels of intensities along the ideal line were used to smooth the jagged lines.

ingly effective antialiasing result can be generated with just eight levels of gray and, in some cases, only four.

Because antialiasing capabilities have been available only on large, expensive graphics systems, many users have assumed that it puts extensive demands on the processor and isn't appropriate for lower-cost systems. In fact, if you use only a modest gray-scale control, a low-cost system can do a very respectable job of de-jagging.

HARDWARE SELECTION

The hardware selected for an antialiasing scheme should be able to display enough simultaneous colors on the screen to accommodate the number of shades needed between each of the main colors used and the background color. The simplest case is one color against a background—for example, white on black. With eight shades from white to black, the brightest shade becomes the foreground color, the darkest becomes the background color, and the six intermediate shades create the antialiased effect.

A more common case uses eight distinct foreground colors, each with 16 intermediate levels fading into the background. Such a system would need to display 121 colors simultaneously. The equation to calculate this is:

$$8 \times (16 - 1) + 1 = 121$$

Eight-bit display systems can handle 256 colors, more than enough for this kind of application.

The intermediate shades must be carefully chosen so that the CRT will show balanced intensities. When the intermediate intensities are not well matched, the antialiased lines will have a candy-striped or beaded appearance. Therefore, the system must maintain control over the actual displayed shade and intensity of each individual color. This usually requires a color look-up table. An exception is the case in which the 256 pixel values are converted directly into gray shades with no color.

Because the intensity produced on a CRT screen isn't a linear function of the input voltage, the calculation of the values stored in the color look-up table is not straightforward. This departure from linearity is called gamma (γ), and must be compensated for by the values in the look-up table. The actual intensity correction depends on how the monitor is set up and may vary widely, but it usually approximates the 2.5 power of the input voltage (compared to the black level). The γ exponent varies between 2.2 and 2.8.

Color monitors with 512 × 512 pixel arrays are cost-effective and quite adequate for the majority of applications. State-of-the-art 4096 × 4096 black-and-white monitors and 2048 × 2048 color systems available can provide such a high degree of image resolution that antialiasing is less necessary, but the cost of these high-resolution monitors tends to preclude their use in most applications.

The addition of antialiasing capability to a graphics system with a 512 × 512 monitor is a modest expense, however. In effect, the process is a way to increase the apparent resolution of the monitor. Such a graphics system can offer the combined advantages of low cost and high performance. If the host processor is used to manipulate the image, the video output interface can in many cases be reduced to a single board, resulting in a truly cost-effective system with antialiasing.

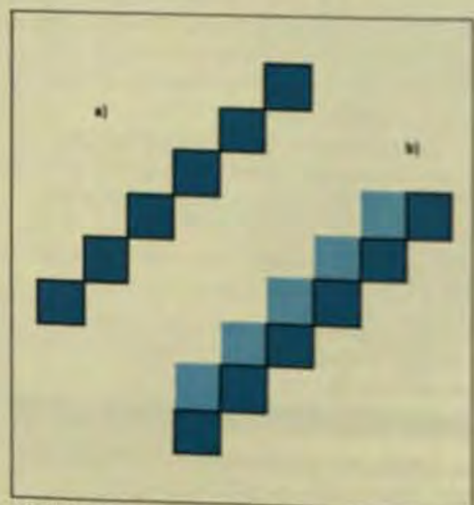

Fig 2 A typical angular line on a CRT display has a jagged appearance (**a**) that can be corrected by shading the area between the steps (**b**).

THE ALGORITHM

The fundamental idea of antialiasing is to adjust the intensities of pixels near the ideal line so that the center of gravity of the intensities corresponds to the true location of the line. If the total intensity is correct, the eye will interpret what it sees as the actual ideal line.

In theory, you could find the required pixel intensities by filtering and convoluting the ideal image, but this approach requires too much computation to be practical. The trick is to find short, easy techniques that approximate the ideal. In actual practice, two computational approaches have been used to de-jag vectors and lines.

In the first, the pixels are represented as a mosaic of little squares covering the view surface (Fig 3). A vector is represented as a band one pixel wide. The affected pixels are those that the band overlaps. On a scale of 0 to 1, the degree that the vector affects a pixel is a simple function of the amount that the pixel is overlapped by the vector.

SPECTRAGRAPHICS INTRODUCES THE SYSTEM 1500 MULTI-STATION.™

HURRY UP AND WAIT.
You're busy. Your company is busy. Your customers are busy.
How can you afford to waste productive hours waiting your turn at the workstation? You can't.
Patience may be a virtue. But it can cost you the marketplace.

STOP WAITING. START WORKING.
Thanks to advanced distributed graphics architecture, the System 1500 puts the power to create in your hands quickly, efficiently, colorfully.
Think of it. Multiple high-performance workstations with cost-performance unmatched anywhere else.
You'll not only find high-

resolution displays, but the fastest screen updates on the market, as well.
PCB designs and integrated circuits will spring to life. So will structural drawings, mapping, and the analysis of seismic data and sophisticated stress tests.
All quicker and more thorough than ever before.

HIGH-PERFORMANCE GRAPHICS. ANYWHERE.
The real beauty of the System 1500 lies in its flexibility of location. The result of its advanced architecture is staggering.
Workstations can be grouped, in the next room, the next building —or miles away.

THE MULTI-STATION CONTROLLER.
Key to the system is the controller. Providing the highest performance graphics processing available in a raster system. And supporting from one to four workstations.
Every single graphics-related function is performed by the controller.
Including high-speed interfaces to IBM, DEC and more.

ADDING THE THIRD DIMENSION.
The System 1500 controller ably performs all its graphics functions at high processing rates. Both 2D *and* 3D.
3D transformations, including rotate, translate and zoom—even of complex images—can be performed at real-time rates.
You'll find 1024 x 1024 resolution with flicker-free performance. And choose 16, 256, or 4096 colors from palettes up to 16 million.

STEP OUT OF LINE.
Now add to all this a full spectrum software package called Prism. And diagnostics as advanced as the system. And you'll get some idea of why we consider the new System 1500 MULTI-STATION concept the future of computer graphics.
So if you and your ideas have been patiently taking a backseat to someone else's ideas, we'd like you to do two things.
One, step out of line.
And two, give us a call for a demo of the 1500.
The hour or so you spend with us now can save you days of standing in line later.

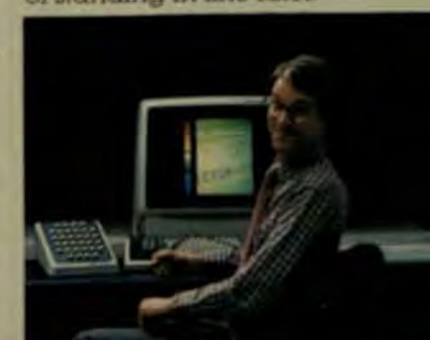

SPECTRAGRAPHICS
Spectragraphics Corporation/10260 Sorrento Valley Road/San Diego, Ca. 92121/(619) 450-0611.

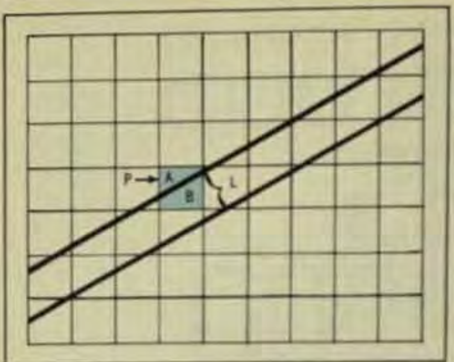

Fig 3 One algorithmic approach regards the pixels as square areas and the lines as bands that intersect them. A pixel P is affected by line L if L overlaps P. The amount of the effect is proportional to the ratio of the overlapped area B to the total pixel area A and B.

If the pixel is completely overlapped by the vector, it assumes the same intensity as the vector (effect = 1). If there's no overlap at all, the pixel keeps its original intensity (effect = 0). Otherwise, the effect of the vector on the pixel's intensity is an intermediate value represented as a fraction between 0 and 1, as calculated by the following formula:

$$\text{Intensity} = \text{Original Intensity} \times (1 - \text{Effect}) + \text{Vector Intensity} \times \text{Effect}$$

This is essentially a proportioning scheme between the two intensity levels by use of the standard color-mixing formula. The process determines which pixels are affected or overlapped by the line and by what percentage each is overlapped. Then, it calculates an intensity value by the above equation. This standard textbook approach works but is time-consuming.

In the second computational approach, the pixels are represented as points in a square grid on the view surface, and the vector is regarded as a line with zero thickness (Fig 4). The affected pixels in this case are those within some distance (d) from the line—usually one pixel.

The magnitude of the effect is calculated as a function of this distance, with linear and Gaussian functions being popular choices. The same color-mixing equation applies after the effect is calculated. Algorithms that calculate the distance can result in dejagging programs that execute very fast.

The two methods yield similar results, and in some cases are actually equivalent. The choice depends on computational convenience—the kind of hardware and algorithms that are available.

These methods can be extended to polygons, both with and without outlines. Special treatment is usually required for pixels near the end points of a vector and for pixels affected by two or more vectors. In addition, generalizations are required for nonsquare pixel-aspect ratios, but the same basic ideas apply.

These calculations will produce the intensity levels required for intermediate shades between the foreground colors and the background. At this point, the designer must decide how many discrete foreground and background colors are needed, how many of the foreground/background combinations will be dejagged, and how many levels are required for each combination.

For most presentation-quality graphs, not all possible combinations are required, and it's almost always possible to fit the requirements into the 256 colors of an 8-bit system. Presentation-quality graphs typically have only one or two background colors and few intersections between the different foreground colors. Such graphs are common in business applications or process-control displays. An 8-bit system with antialiasing can readily generate such displays with a dozen different colors active against a background.

OPERATOR BENEFITS

The combination of today's low-cost graphics systems, declining semiconductor memory prices, and technological advancements in high-resolution color graphics has driven a rising demand in the marketplace for user-friendly features like antialiasing. Complete hardware/software solutions with antialiasing capabilities are available to OEMs as board-level products to be integrated into existing host systems.

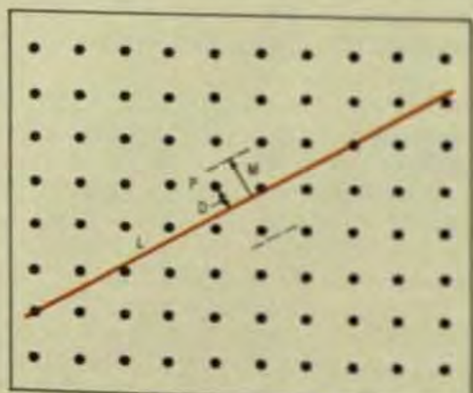

Fig 4 A second algorithmic approach treats the pixels as points and the lines as vectors with no thickness. A pixel P is affected by line L if it is within some function of the distance D. The effect is greater on pixels that fall closer to the line.

Products of this type can lower operator error rates, generate more accurate displays, and create more natural human interfaces in such applications as computer simulation, animation, education graphics, business-presentation graphics, process control, and even video games.

Certain computer-graphics applications have traditionally valued the naturalness of the image more highly than other uses. For example, cartooning, computer art, and graphic terrain simulation are applications where image quality and accuracy are of primary importance.

Solids modeling for CAD/CAM and computer-aided instruction (CAI) also shows a marked improvement when the subjects are rendered in antialiased form. Such applications might be termed passive uses of antialiased graphics, in which the graphics image, as in a painting, is valued for itself.

Another class of computer-graphics applications gives the display a more active role—the information presented on the screen is intended to elicit human action. In text processing, typesetting, medical imaging, and process control, a major benefit derived from a more realistic screen image is more accurate human perception of a given situation.

The result is operator decisions that are more timely and better informed, lower error rates, and sometimes even safer operation. Operator fatigue in these highly interactive uses can also be minimized by smoothed vectors and borders, avoiding the need for the eye-brain system to filter out the noise.

PROCESS CONTROL

The availability of lower cost graphics hardware devices and more efficient software algorithms for computing the boundary effects needed to dejag a given image are making antialiasing increasingly attractive for these active applications. In a process-control application, for example (Fig 5), a typical color-graphics system incorporates

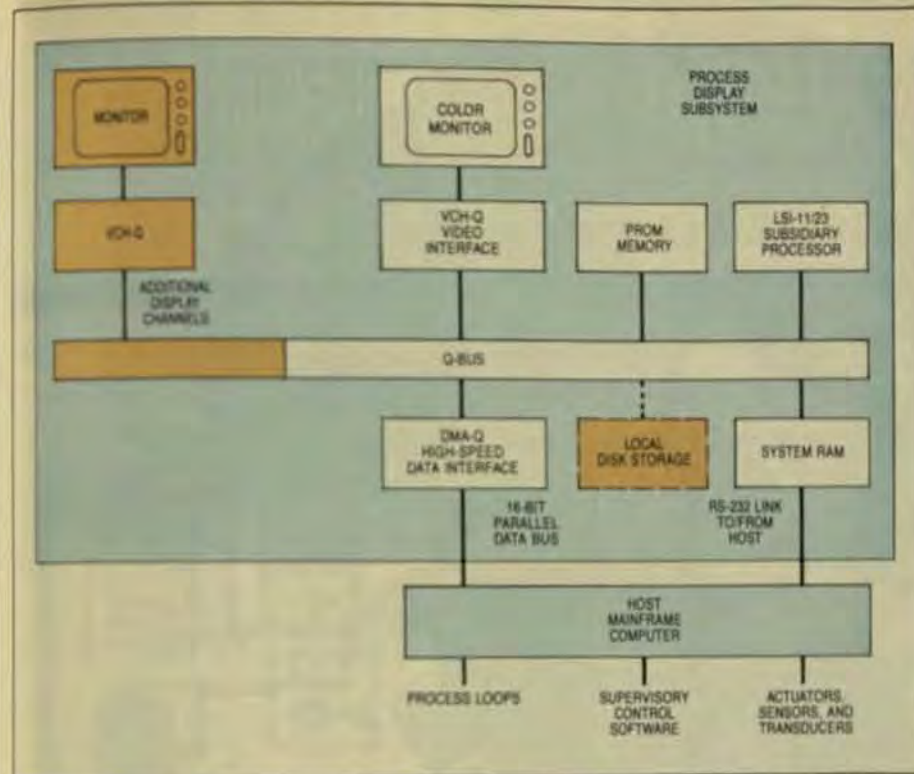

Fig 5 This monitor-control subsystem uses antialiased color graphics. High-speed parallel lines connect it to the host computer.

antialiasing software as an integral part of a monitor control subsystem connected to a remote host mainframe by one or more high-speed parallel links.

The subsidiary Q-bus processor is used as a multichannel data reduction and display system. Similar display systems are in use today in a number of process environments that use a variety of processors. In these applications, the host mainframe performs the supervisory control of one or more continuous processes (such as catalytic cracking), exercising direct and indirect control over the process. The display portion of this system keeps the operators informed of rapidly occurring process changes. Data displays are important in such an environment because the operators often have to respond quickly.

In this system, a subsidiary Q-bus processor receives high-speed data via a DMA channel from the host processor. Information may be acquired from the host in a condensed form to minimize transaction time and overhead on the host. A downstream Q-bus processor such as Digital Equipment Corp.'s PDP-11/23 performs the tasks of assembling, formatting, and presenting process variables and other information developed by the host processor as the process continues.

Local resources available to the auxiliary display processor may include:

- System RAM to store downloaded programs that allow the subsidiary display system to be flexible and fast without creating an inordinate processing load on the host.
- Nonvolatile PROM to store an operating system kernel and perhaps basic system constants such as character sets for alphanumeric presentation and display templates.
- Local disk to store and retrieve data structures not required by the host.
- Additional serial communication links for status verification and small-data-item exchange between the host and the display subsystem or other display processors.
- The display channel or channels that consist of color graphics hardware. The 16-bit parallel data bus rapidly transports display data and other exchanges between the host and the display system.

Software running in local system RAM or PROM, whether downloaded from the host or retrieved from local disk storage, generates the information displays in antialiased form along with accompanying alphanumeric annotation. The operator sees a simultaneous display consisting of quantitative information and a qualitative rendition of how the process is doing. In other words, the display includes precise numbers and, in very realistic form, a quick look at what's going on (Fig 5).

This type of system can be added to existing host-controlled processors with only a modest increase in overall system costs. And when requirements change, the data-display channels can be expanded modularly with the addition of a single board and CRT monitor for each complete display channel up to the practical limits of the processor. An auxiliary processor can handle anywhere from four to eight channels, depending on processor load, display contents, and display activity.

The main benefits of this approach to high-resolution antialiased color graphics are:

- Economical initial system costs.
- Easy future upgrades through modular additions of boards.
- A low-confrontation learning process for service people.
- The realistic, natural display images produced by antialiasing software. ■

Thomas H. Birchell, *president of Peritek Corp., is responsible for the overall definition, manufacture, and marketing of Peritek's line of small computer products.*

Stuart R. Dole, *president of SAMKHYA Software Research & Development, has been a computer consultant for biomedical and graphics applications since 1973.*

Victor R. Gold, Jr. *is vice president of engineering for Peritek and helped found the company in 1976. He has been responsible for hardware development of Peritek's color graphics video interfaces.*

At I²S, we know what it means to be an OEM—and what it means to be a responsible OEM supplier. So when we say I²S delivers total OEM image processing solutions, we speak with confidence—confidence based on proven capabilities and system experience. They're your assurance that—from design to delivery—we're the people to complete your image processing picture.

We'll deliver a standard system— or build one meeting your precise needs. Either way, you get top performance—affordably, the way we're known for delivering in all areas of image processing.

You see, from the beginning, we designed our hardware for easy integration—and our software for genuine transportability. We led the way in sensor interfacing, in high-speed high-resolution pipeline processing, in image storage and retrieval, and in multi-mode display formats. Today, we still set the image processing pace—in areas like image data compression and networking.

As an OEM involved with image processing, you've got important decisions to make. The first one is easy. Learn how I²S can deliver your total solution. Write or phone I²S, 1500 Buckeye Dr., Milpitas, CA 95035, (408) 262-4444. Do it now. Because your image is important to both of us.

OEM image processing solutions. That's our business.

International Imaging Systems

Hough Algorithm Plus Probability Calculation Aids in Image Analysis

When a number of densely packed circular objects need to be analyzed, the computer must first tell them apart. A unique program makes this possible.

by **Roger Button,** Omnicon Products

A problem in image analysis occurs when there's a need to separate objects that are packed relatively close together and which may be touching—for example, when parts in a bin need to be separated for an application involving an industrial robot. A unique feature of the solution to be described in this article is that probabilities are used by a computer to segregate the massed parts of objects by determining which portions of the whole image should be combined into each component.

This "circular de-aglommeration" technique was developed to enable a computer to recognize microscopic latex spheres mounted on a slide, under a glass cover. However "stuck together" the objects are, each is actually circular and located in a plane, so none obscures another from view. Because of an optical effect that occurs when an optical image is used, there may be a concave curve segment in the area where two circular objects touch. These concave segments are discarded. The goal is to enter the diameter of each object into the storage unit of a computer. Once these diameters are obtained, existing programs can perform statistical analyses on them.

Analysis begins with a picture of these objects—either a photograph or an image from a scanning electron microscope—with a black background and a white foreground. In the picture, the objects look like a large number of Ping-Pong balls. Some are separate, but most are clumped together in masses like frog's eggs.

After the entire picture is stored, it's processed to obtain diameters. The customary ways of measuring an object's size can't be applied until each sphere is separated from the background and from the other objects. However, in this case, the objects don't need to be separated because they can be represented mathematically. The algorithmic approach used is analogous to the Hough (pronounced "huff") method of curve detection, but has a probabilistic twist.

The points on the boundary of a circle are solutions to an equation with three coefficients, even when the boundary isn't contiguous because it touches other boundaries. These three coefficients are the locations on the X and Y axes, and the radius. For convenience, these can be labeled X, Y, and R. The picture shown in Fig 1a involves two circles. Since both have the same location on the Y axis, it can be omitted from the plot of a 3-dimensional histogram of X,Y,R versus a number of points (Fig 1b).

The program developed analyzes the picture's features to find segments of circular curves, then combines segments that have similar coefficients in the equation for a circle. The result hoped for is that all segments of circular curves actually belonging to a single circular object will be sufficiently similar that they can be combined, and that segments of circular curves that belong to separate objects will never be combined.

Due to a variety of conditions—including noise introduced by sensor and differences in the sizes of objects in the picture—there may be no sin-

YOU CAN FIND A LOT OF COMPANIES WHEN YOU WANT TO BUY A MONITOR...

...BUT HOW MANY CAN YOU FIND WHEN YOU WANT TO SERVICE IT?

There's one you won't have to search for: Conrac.

We're committed to giving OEMs the service and support you need—when you need it and where you need it. We're implementing an extensive network of regional and local technical centers throughout the country, and that means service and replacement parts along with sales support nearby, where you need it. And with our ongoing training program, the skills that our people can offer your people are constantly being polished and upgraded.

When you look at the total picture—our 36 years of experience, the broadest range of products in the industry, high quality, performance, reliability, serviceability, availability and responsive engineering and customer service—it's easy to see why Conrac is the first choice in OEM markets. You know that once you buy from us—you won't have to search for us.

For more details on Conrac's complete line of monitors, and our schedule of training programs, call or write:

Conrac Division, Conrac Corporation, 600 N. Rimsdale Ave., Covina, CA 91722.

Regional Offices:
Western (213) 966-3511,
Mid West (312) 833-3310,
Eastern (203) 348-2100,
New England (617) 272-6720.

CONRAC

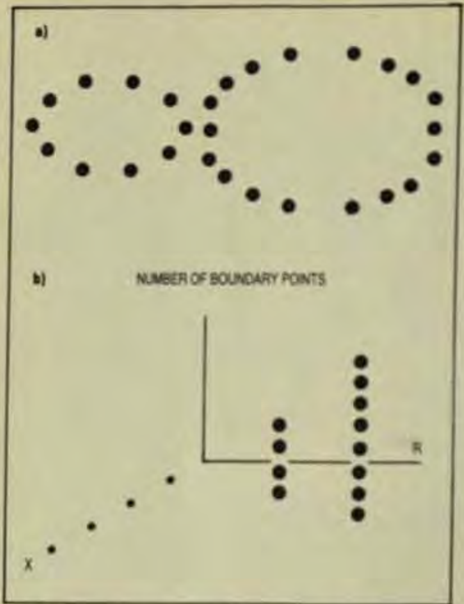

Fig 1 The points on the boundary of a circle are the solution to an equation with three coefficients that determine the points' locations on the X, Y, and R (radius) axes. The two circles shown (a) have the same Y-axis location, so that axis is omitted from the 3-D histogram of the circles (b).

gle threshold of similarity that will give the desired result. In pictures where the coefficients are similar enough to serve as a rule for whether to combine two curve segments, the coefficients for the whole curve could be estimated better if a way were found to combine two segments correctly.

Recursive combination, beginning with the two segments having the most similar coefficients, reduces the chance that two curves are combined incorrectly. In this program, when the two most likely curve segments are combined first, the correct results can always be obtained—even with a very simple computation for determining whether a combination should be performed.

You can easily see that two curves are represented in the plot shown in Fig 1b. "Seeing" the result in X,Y,R space is called partitioning the space—the mapping of every point on a boundary into a parameter space in a way that treats all points on the same mathematical form of a curve as if they were a single point in the parameter space. This mapping is the basis for the Hough curve-detection algorithm. Successful partitioning of the parameter space may require a great deal of "tradecraft," however, since the process may be complicated by a great deal of noise in the picture.

In looking at the program in detail, consider three algorithmic parts: boundary tracking, arc segmentation, and arc combination.

BOUNDARY TRACKING

In boundary tracking, only the points that lie on the boundaries are solutions to the curves, so there's no need to consider other points in the image. Also, a relationship is formed between a point on a boundary and its predecessors and successors. You can use this relationship to collect sets of contiguous points that have convex curvature. To do this arc segmentation, divide a boundary—a closed set of contiguous points—at its points of inflection. Then, combine these arcs to form a set of all the points that belong to the same circular object. When a set of arcs represents each ball in the picture, a list of the diameter coefficients for each set is a list of the diameters of all the balls shown.

For the purposes of this program, the picture can be viewed as a number of boundaries, each one a set of points in succession until the boundary closes on itself. The outer boundary outlines the ball or clump of balls, and the inner boundary outlines the holes in the clump of balls. Together, these form a set of boundaries that can be collected in sets that describe a feature, because only the arc segments that are members of the same set could possibly contain points located on the boundary of the same circular object.

This natural division of the data is used to reduce the number of possible combinations to increase computing efficiency. Each boundary can be segmented into convex arc segments that correspond to part of the boundary of a single object. Ideally the points on a boundary would all belong to a convex arc, but actually points may be contained in arcs that are either concave or convex (Fig 2a), so it's imperative to be able to tell the difference between the two types of curve segment.

This distinction is accomplished by controlling the order of succession for the points in a boundary—counterclockwise for the outer boundary and clockwise for a hole. Thus a boundary that curves to the left is convex, and the points that lie on the interior of a feature always lie to the left side of the boundary (Fig 2b).

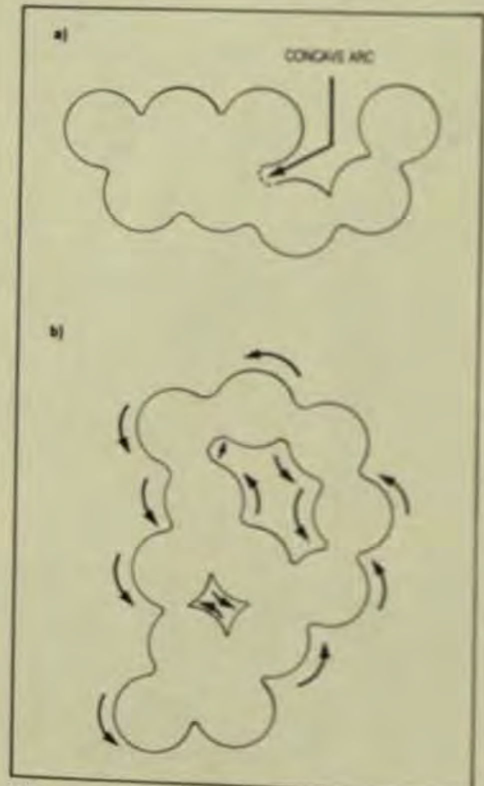

Fig 2 Boundary points may be contained in arcs that are either concave or convex (a). The boundaries are calculated so that points that lie on the interior of a figure always lie to its left (b).

REGRESSION OF BOUNDARY SEGMENTS

A textbook implementation of the Hough method will take each point that's found to lie on a boundary and map it into a parameter space whose axes are the coefficients of the equation describing the form of the curve. For a general conic section, that equation is

$$AX^2 + BY^2 + XC \times Y + DX + EY + F = 0$$

In the case of a circle, both A and B are 1, and C is zero. This means that you must solve for D, E, and F—the three coefficients equivalent to the center and radius of a circle. Where X, Y, and R are the center (in x and y) and the radius of a circle, then

$$X = -D$$
$$Y = -2E$$
$$R = \sqrt{(D^2 + E^2 + F)}$$

You can either solve for these coefficients for every point on a boundary, or, as in this case, you can recognize that if a contiguous part of a curve is convex, then the entire curve segment must be part of the boundary of a single circular object. The advantage of the second option is that the entire section of arc is analogous to a single point on the boundary and gives rise to only one point in the parameter space—thus reducing the number of points and the amount of data to process.

To accomplish this goal, solve a pseudo-inverse of a set of equations:

$$X^2 + Y^2 + D*X + E*Y + F = \text{error}$$

The solution allows you to obtain a best fit to the D, E, and F coefficients. To obtain a pseudo-inverse of the set of equations represented above, reorder, and rewrite them in matrix form:

$$\begin{vmatrix} X_1 Y_1 1 \\ X_2 Y_2 1 \\ \cdot\ \cdot\ 1 \\ \cdot\ \cdot\ 1 \\ X_N X_N 1 \end{vmatrix} * \begin{vmatrix} D \\ E \\ F \end{vmatrix} = \begin{vmatrix} X_{12}^2 Y_{22}^2 \\ X_2 + Y_2 \\ \cdot \\ \cdot \\ X_N^2 + Y_N^2 \end{vmatrix} + \begin{vmatrix} \text{error}_1 \\ \text{error}_2 \\ \cdot \\ \cdot \\ \text{error}_N \end{vmatrix}$$

In more concise notation, where M is the variable matrix, S is the variable-squared matrix, C is the coefficient-column matrix, and E is the error-column matrix:

$$E = M(X,Y)*C + S$$

To minimize the error, square E and then differentiate with respect to all the coefficients. Where E' is the transpose of E,

$$E'E = (M \times C + S)' \times (MC + S)$$
$$d/dC(E'E) = 2 \times (M'M \times C + M'S)$$

that gives the result:

$$C = (M'M)^{-1} M'S$$

WHEN CIRCLES SHOULD MERGE

So far, boundaries have been segmented into convex arc segments, and the points in these convex arc segments have been used to form a system of equations that result in three coefficients for each arc segment. A hypothesis is made that the collection of coefficients, taken together, represents all the circular objects in the picture. How can this hypothesis be tested?

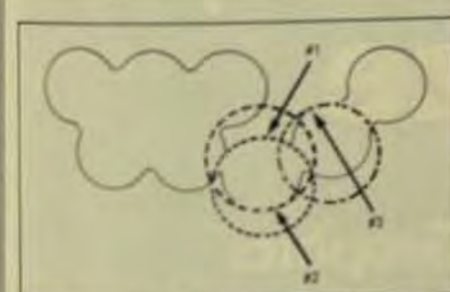

Fig 3 The probability that arcs 1 and 2 combine into a single circular object is 0.82, while the probability that arcs 1 and 3 combine is 0.44.

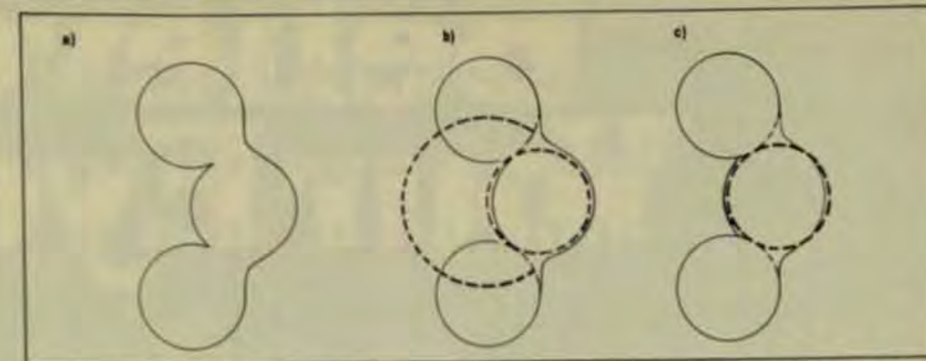

Fig 4 In the original image (a), the larger circular object has one side flattened. The coefficients calculated are all good candidates for combination into the larger circular object shown by the dotted lines (b). The situation is simplified, however, when the two most likely arcs are combined (c).

Since the eight circles result in 11 sets of coefficients, three of the arcs must be combined with other arcs in order for the hypothesis to be true. Say that the hypothesis is true if no two arcs have an area of intersection exceeding a certain threshold. This threshold is determined empirically to be some fraction of the area of the smaller of the two circular objects represented by the coefficients. When two sets of coefficients have an area of intersection greater than the threshold, compute the relative probability that the two arcs should be combined to solve for a single set of coefficients.

If two of the arc segments have exactly the same set of coefficients, the probability that both segments are part of the same circular object is said to be 1. More generally, the probability that two different sets of coefficients were obtained from the same object depends on how far apart the coefficients are.

Since the coefficients form an orthogonal space, their separation is the square root of the sum of the differences-squared in the parameter space. A more useful term is the ratio of this separation distance to the diameter of the smaller of the two circular objects represented by the arcs. This ratio forms the parameter of an exponential function that behaves well as a relative probability function. Thus, the probability that arcs 1 and 2 combine is 0.82, while the probability that arcs 1 and 3 combine is 0.44 (Fig 3).

CHOOSING WHICH CIRCLE TO MERGE

With such a means for establishing the probability of whether two arcs should be combined, you might simply establish an additional threshold in this function to determine whether a combination will be performed. That might prove adequate for many pictures, but one additional technique provides a great deal more robustness.

Fig 4 will help to illustrate the value of "most probable first" arc combination. In the original image (Fig 4a), the larger circular object has one side flattened. The boundaries have been segmented into convex arcs, and the points of these arcs have been submitted to a regression to obtain the coefficients for circular objects. These coefficients are plotted in Fig 4b.

Because of the flattened side of one of the objects in the original picture, all of the circular objects, as represented by their coefficients, are good candidates for arc combination. Fig 4c shows a plot of the same situation with the two most likely arcs combined. In the ensuing situation, none of the remaining arcs will be erroneously combined.

In actual practice, it's not necessary to recompute the relative probability of arc combination every time two arcs are combined. If the calculation that computes the eligibility for arc combination records the probabilities in a table, you can keep a second threshold that will allow the two arcs to be combined despite the fact that one of them was eligible for combination at a higher relative probability. ■

Roger Button *has been with Bausch & Lomb since 1964, and currently works in their Omincon Products subdivision, where he is responsible for supporting software development of new image processing and analysis applications on 32-bit microprocessors.*

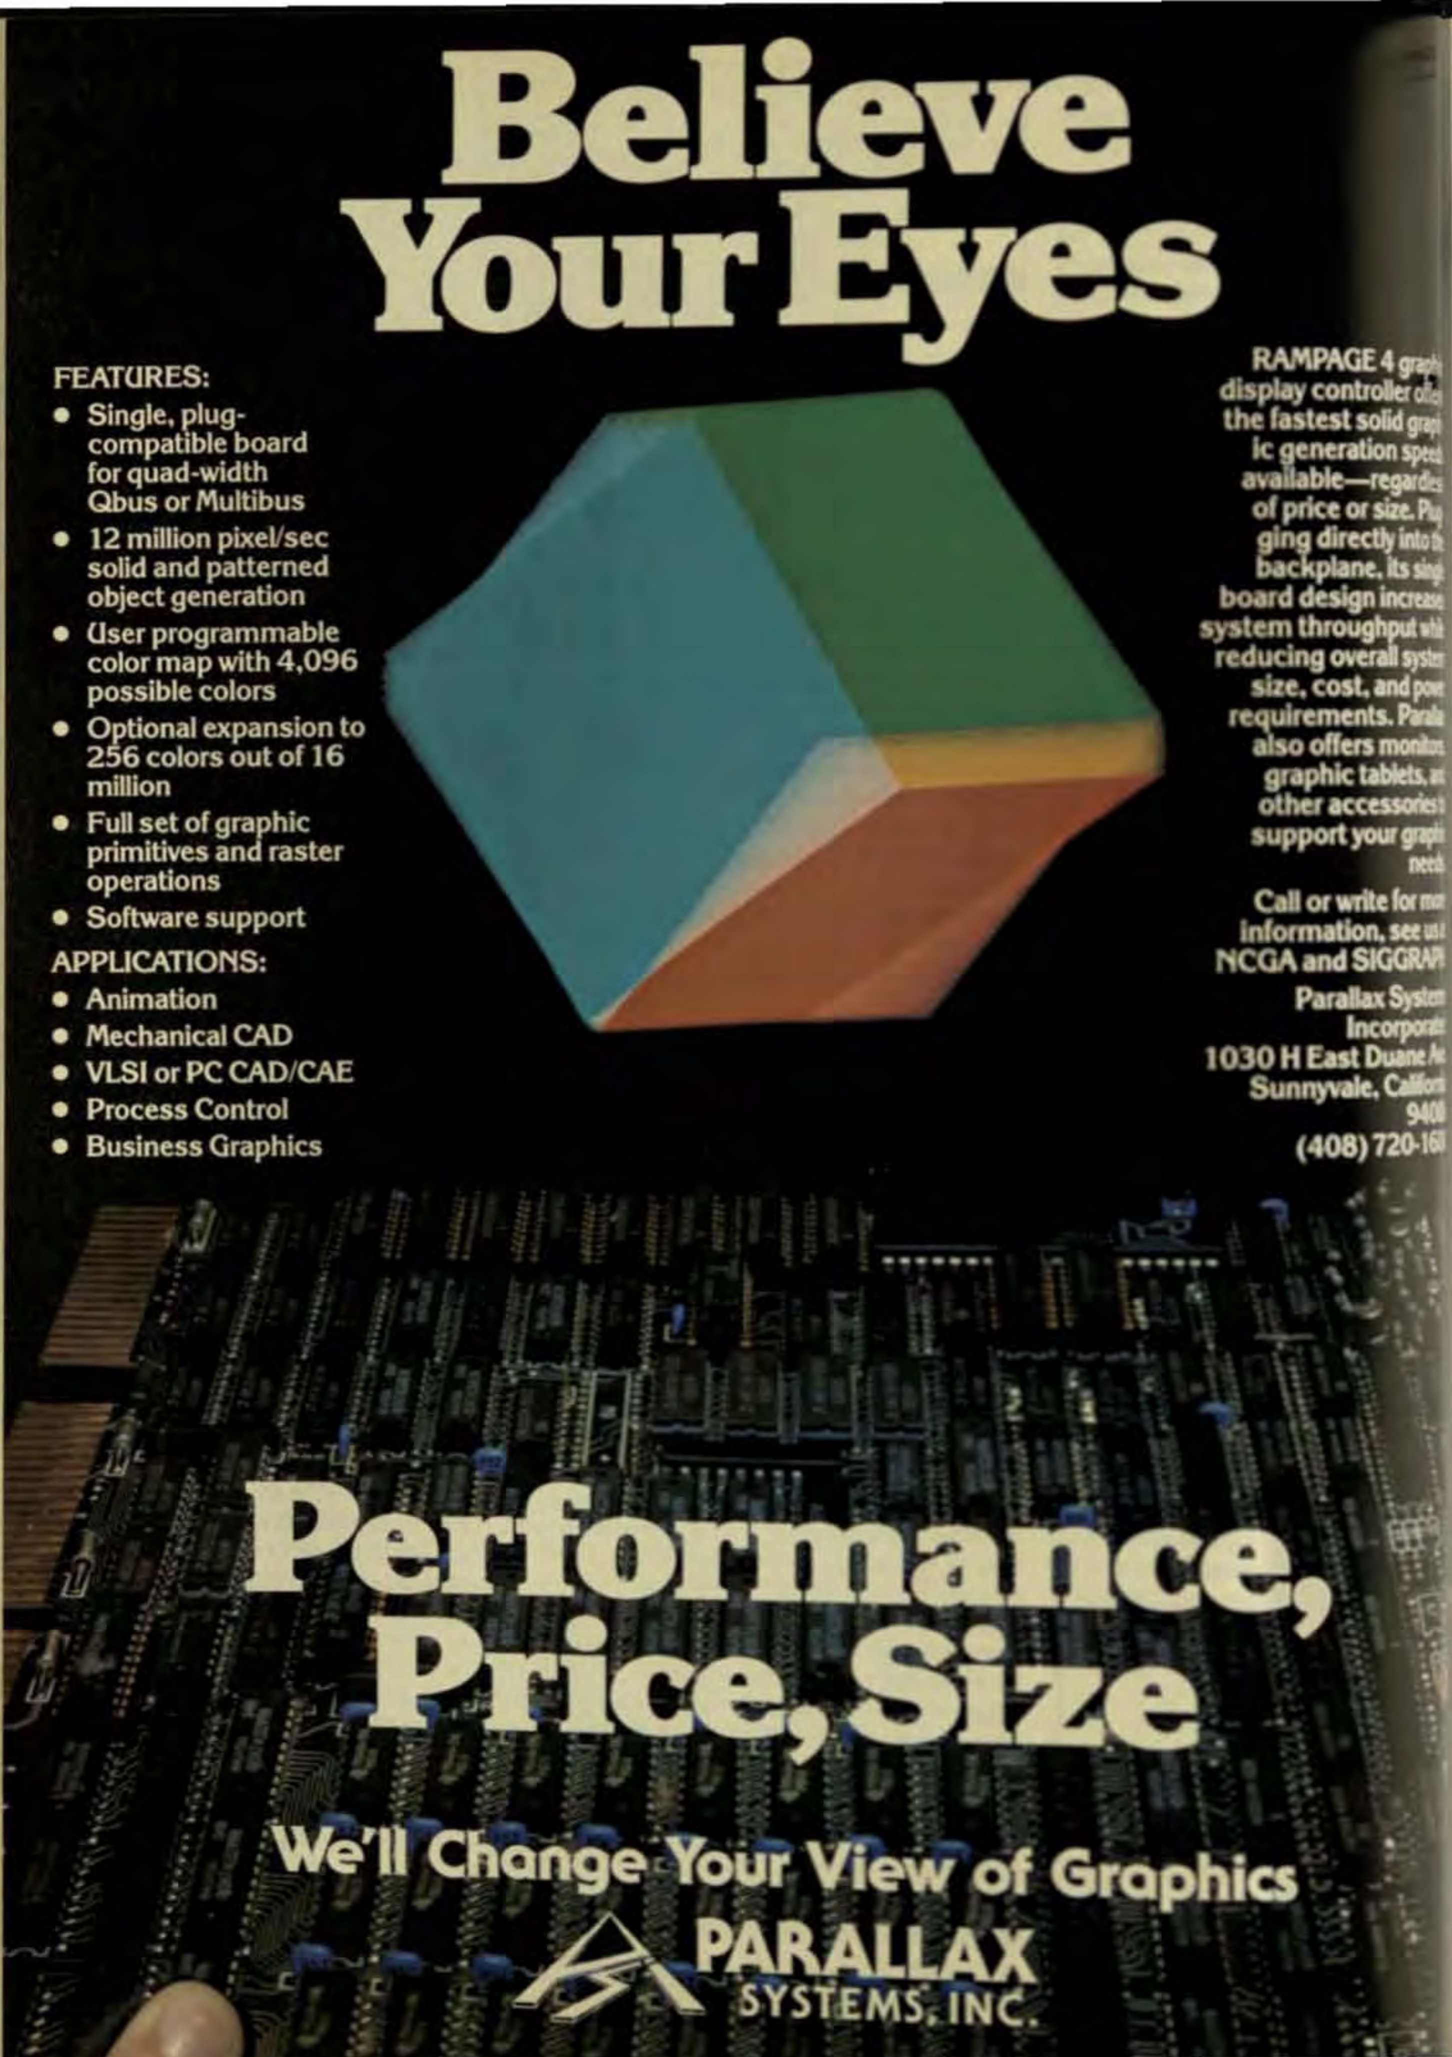

Array Processors Evolve For Modeling and Large Scale Simulation

The application of APs to modeling and simulation problems means that raw execution speed is no longer a sufficient criterion for evaluating these products.

by **Edward J. Kushner,** Floating Point Systems Inc.

Array Processors (APs) have been commercially available since the mid-1970s and are now used extensively for signal processing, image processing, and real-time simulation. Although these applications are quite dissimilar, they share the same computational characteristics—moderate requirements for precision (32 to 38 bits), modest demands for memory (64K words), and small computational kernels that typically consist of vector operations, matrix manipulations, or fast-Fourier transforms.

In these applications, the array processor runs in tandem with a host computer; the host handles I/O while the AP performs numerically intensive operations. Program code is usually written in assembly language to achieve the speed that this architecture makes possible. Because an application is run on a dual-processor system, the data-transfer rate is often crucial to the success of a particular application.

AP APPLICATIONS AND ARCHITECTURE

Since their first appearance, attempts have been made to apply array processors to the larger and more complex problems found in modeling and large-scale simulation. These applications require extended precision math (64 bits), make more severe demands for memory (1 Mbyte or more), and use much larger programs (many thousands of lines of FORTRAN).

In many instances, these applications are characterized by computational kernels that comprise a very small fraction of the code and consume a very large fraction of CPU time. Thus, attempts have been made to adapt the problems of modeling and large-scale simulation to a host/AP system. These attempts usually haven't been successful because of insufficient precision, lack of sufficient memory, or because of the time lost in data transfers between the processors.

One solution can be found in systems that allow the host and array processor to share memory. Such systems give the AP access to the much larger memory of the host without suffering the time penalty of data transfers. Another solution is to design a processor large enough to run the entire program. In this architecture, pipelined functional units are combined to operate in parallel (Fig 1).

Pipelining has been incorporated into a three-stage floating point multiplier, a 2-stage floating point adder, and a 3-stage memory pipeline. A single instruction can utilize these 3 units as well as include an integer operation, reads and writes to two separate data registers, a read or write to secondary memory and initiate a branch operation. Such an instruction can be initiated every machine cycle to give a hardware capability that allows 10 simultaneous operations.

To meet the precision requirements of most modeling and large-scale simulation applications, 64 bits are used to represent floating point numbers. The maximum memory

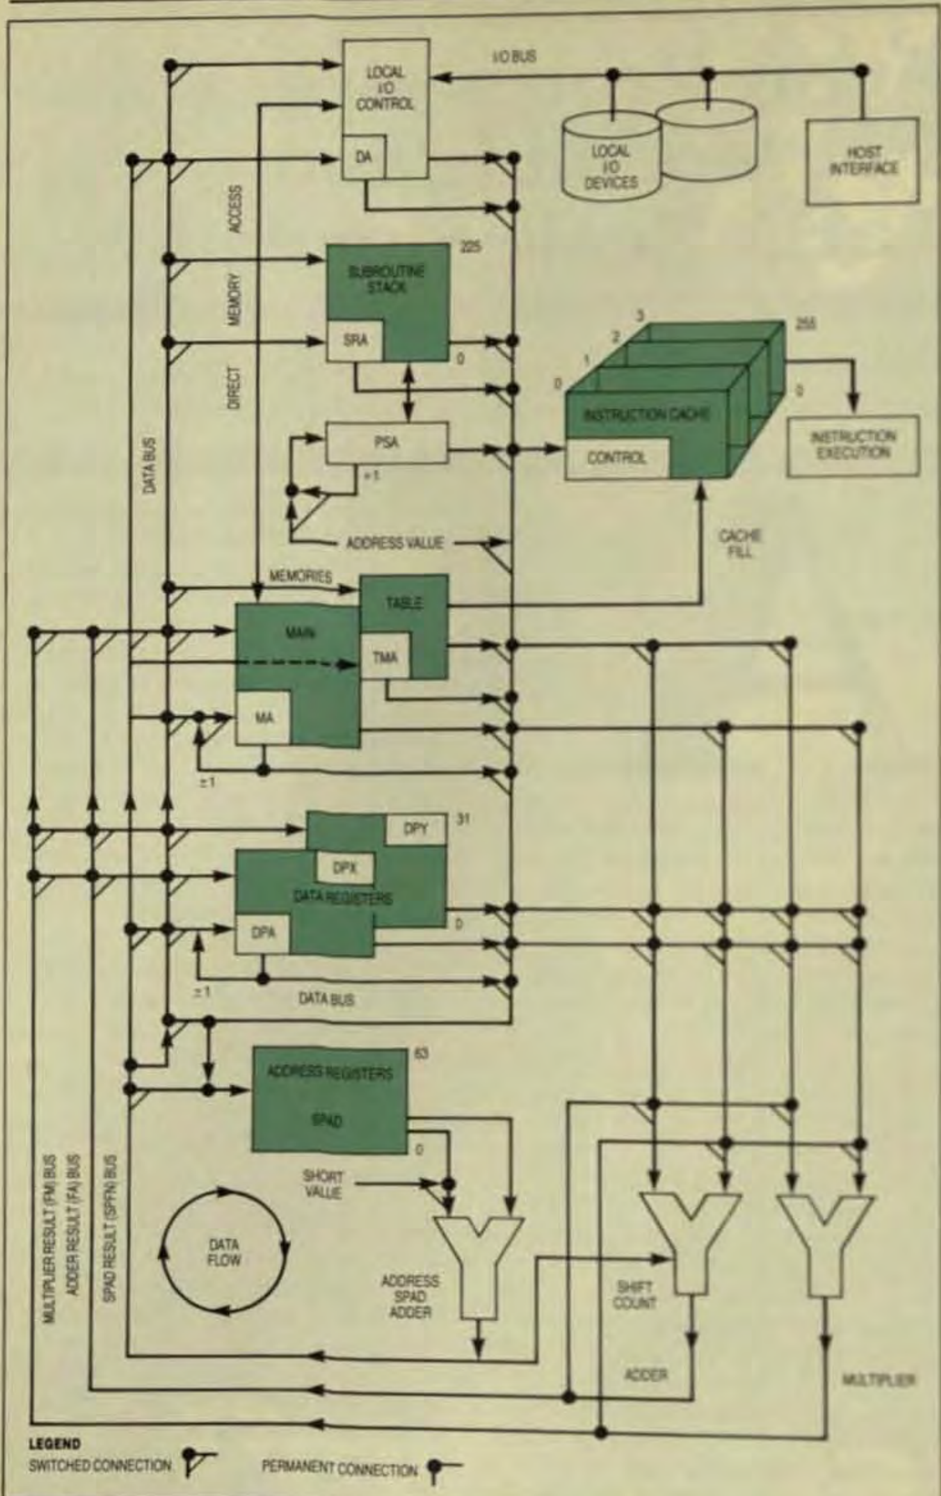

Fig 1 Array-processor architecture that combines functional units that are pipelined and operate in parallel can speed large numerical calculations.

available for such applications was recently increased to 58 Mbytes, and attached disks can be added when greater memory capacity is needed. Configurations are available from 135 to 3000 Mbytes.

SOFTWARE UTILITIES

Foremost among the software tools available for such systems is a FORTRAN compiler that adheres to the ANSI 77 standard and satisfies the U.S. Dept. of Energy extensions for asynchronous I/O. It runs on the host computer to produce object files. An object linker that also runs on the host combines several object files to produce a single file to be transferred to the AP.

At the lower levels of the compiler, instructions are generated that contain several of the operations allowed by the architecture's parallelism. Thus, code is squashed so that an instruction may include floating point arithmetic, integer arithmetic, memory reference, and a read or write to a data register.

A software pipelining feature allows the performance achieved by the compiler at lower levels to be significantly improved at the highest level. Thus, instructions that contain multiple operations are generated in inner loops so that the total number of instructions for a given loop can be significantly less than that achieved at lower levels. This feature can greatly improve performance on innermost DO loops that are determined to be pipelinable. The compiler considers a FORTRAN DO loop to be pipelinable so long as the loop doesn't contain any of the following:

- A logical operation (such as an IF statement).
- A call to a library routine (such as SIN, or SQRT).
- FORTRAN intrinsics that are generated in-line and contain a branch instruction (such as MAX or MIN).
- Exponentiation except to a power of the integer 2.

Pipelinable loops (Fig 2) may contain:

- Scalar operations.
- Recursion. For example: A(I) = C * A(I − 1)
- Indirect addressing. For example: A(I(J)) = X

For loops that qualify, the microcoded instructions contain operations that correspond to successive passes through the loop. Fig 2 shows that three types of instructions are generated: those that fill the pipeline and are executed once, those that comprise the basic loop and can be executed N times, and those that empty the pipeline and are executed only once.

The software pipelining feature is of little value when N = 1, but can have a significant impact on performance for values where N = 10 or more. Table 1 gives execution times with and without software pipelining for typical problems that arise in sci

COMPARED TO THE OTHER COMPUTER CAMERAS CELTIC'S VFR 2000 LACKS TWO FEATURES: GUESSWORK AND SOFTWARE.

Celtic's VFR 2000 Computer Camera does more than turn computer generated graphics into full color 35mm slides. It does so in seconds. For a few pennies each. And it doesn't miss.

You won't find better image quality anywhere. And thanks to Celtic's unique Image Preview Port, working along with a LED display of the film frame and a "ready" indicator, there's no guesswork to the system.

And you won't find any other computer camera that's so easy to use. Because unlike most, the VFR 2000 requires absolutely no software. Just take it out of the box and put it to work. Besides, your entire office will enjoy convenient advantages like a fully automatic film wind and rewind capability. And our Color-Enhanced Monochrome (CEM), a feature that lets you make color copies even when your display is in black and white! Use standard slide film, including new Polachrome® Instant 35mm film for immediate developing and mounting of your slides.

Use it to project crisp, clear images as big as wall size, the perfect way to dramatize business graphics for management meetings, training programs or just about anything. The VFR 2000 is a natural addition to your personal computer or computer graphics system capability. It hooks up universally to a large number of computer systems, from PCs to mainframe workstations like the IBM® 3279. It's also adaptable for use with a wide range of graphics terminals. And just like your color monitor, it accepts either RS 170 or TTL inputs.

We price our VFR 2000 model under $2,500. And we're ready to match its value to that of anybody else's computer camera— higher priced or lower. Because we're ready to work as is. Others aren't.

For a hands on demonstration call Celtic Technology at (213) 884-6767 or return the coupon below.

CELTIC TECHNOLOGY

Innovation in Computer Photography

CTR-S83

Celtic Technology, Inc., 6265 Variel Avenue
Woodland Hills, CA 91367 (213) 884-6767

☐ I would like a hands on demonstration.

☐ Please send me a product brochure.

Name ____________________

Title ____________________

Company ____________________

Telephone ____________________

Address ____________________

City ____________ State ______ Zip ______

© 1983 Celtic Technology, Inc.
Photography Examples Courtesy of: 1) Ramtek Corporation of Santa Clara, CA, 2) XIPHIAS, Santa Monica, CA, using their Software System, Videograph.

IBM is a registered trademark of IBM Corporation.

Vector General declares CAD/CAM independence for engineers

CAD/CAM independence means a complete turnkey system dedicated to the engineering department. It means a system powerful enough to handle the most sophisticated software, with the flexibility to grow as your needs grow. It means total compatibility to network and communicate with other systems, yet maintain security of information.

The System 8000 meets all those requirements, and more.

A complete turnkey system The System 8000 is a complete system, incorporating a powerful Motorola/Four-Phase CPU, 400 Mbytes disk storage, and a full range of support options, including plotters, digitizers and large screen displays.

An expandable system with local convenience The System 8000 is expandable from 1 to 12 stations at up to 3 sites. Displays can be located up to 400 feet from each site, and the sites can be up to 3 miles apart. All will serve the engineering department, even if it is in multiple locations. If you have more than one engineering department, each can have its own complete system.

Not minicomputer-bound The powerful Motorola/Four-Phase CPU has the processing power and data base integrating capabilities to handle the most sophisticated CAD/CAM software, including CADAM® N-CAD® or both simultaneously.

Full networking compatibility The System 8000 is compatible with synchronous network protocols and can communicate via modem and dial-up or dedicated phone lines with other systems wherever they may be.

Return on investment The System 8000 enjoys a cost advantage of at least 25% over any equivalent system. Typical payback time is one year or less. With this ROI, you can afford to give each engineering department its own system.

Productivity center The Vector General System 8000 turnkey CAD/CAM system. It makes every engineering department independent. And efficient. And profitable. Call or write for complete details today.

Vector General, Inc.
The Mind's Eye
21300 Oxnard Street • Woodland Hills, CA 91
(213) 346-3410 • TWX: 910-494-2764

N-CAD® is a registered trademark of Northrop Corporation.
CADAM® is a registered trademark of Lockheed Corporation.

Integrated CAE Improves Graphics/User Interface To Aid Concept Design

Improved CAE tools, combined with an integrated design environment that has rich geometric-modeling functions, eases evaluation of design alternatives.

by **Wayne A. McClelland,**
Structural Dynamics Research Corp.

The past 10 years have seen dramatic changes in the way mechanical systems are designed and manufactured—particularly in the application of computer-aided procedures for drafting, analysis, test, and numerical control. The productivity of individual tasks has been significantly increased over this period.

CAD systems can yield productivity gains of three to one or greater compared to the manual drafting process. Similar gains have been documented in such areas as finite-element stress analysis, laboratory vibration testing, tool-path generation, and many other islands of automation. Such productivity gains have rarely, however, correlated one to one with the end product's time to market and performance.

Two key ingredients show promise for closing this gap. One is improved computer-aided engineering (CAE) tools (Fig 1) to facilitate rapid evaluation of design alternatives and promote greater design innovation. The other ingredient is an integrated design environment that emphasizes data sharing among numerous applications and a consistent graphics/user interface.

Structural Dynamics Research Corp. has focused on developing an integrated environment for concept-design engineering. In this environment, systems and components are designed and analyzed in an iterative fashion, from the earliest phases of product planning and with successive refinements until performance objectives are met. The functions of design, analysis, and test interactively access a shared applications database and converse with the user via a consistent, menu-based, graphics/user interface.

SOFTWARE DEVELOPMENT

Software development for this CAE environment provides its own unique challenges. Certainly there's a great demand to increase software capability for individual engineering disciplines (application programs) and to provide a consistent and efficient level of integration. Additional effort must be expended to keep pace with rapid advances in graphics-display devices, attached processors, and powerful desktop computing.

The user interface must be consistent across applications, yet flexible enough to comprehend different user styles (such as hierarchical menus, command-driven modes, screen/tablet selection, native-language support, on-line documentation, and easy programmability). Because of such demands, the software-development process itself has become a critical element in efforts to achieve an effective CAE environment and, hence, the desired productivity gains.

It simply takes too long to develop software, and resulting systems all too often satisfy obsolete functional specs and are difficult to enhance and maintain. Also, software transportability between host operating systems is far from automatic.

The answer is to move from programmer-driven to user-driven software development. In the former, a formal process is followed—

IT'S SUPERCONTROLLER!

It's a 6250 BPI (GCR) controller that can handle dual and tri-density drives from STC and Telex.

It's a software compatible streamer controller for today's streaming 1/2" drives.

It's a TS11* emulator.

It's a 6250 BPI streamer controller for tomorrow's new technology GCR streaming drives.

It's a single board imbedded controller with a 64K byte on-board memory that acts like a large buffer in start-stop mode and as a multiblock staging buffer while streaming.

It's the TS-6251 Supercontro
from Western Peripherals, th
company that has put more
magnetic tape on DEC Unib
computers than any other
independent supplier. Call
write us for technical details

western peripherals
Division of WESPERCORP
14321 Myford Road, Tustin, CA 92680
(714) 730-6250 TWX: 910 595-1775 CABLE: WESPER

*Trademark Digital Equipment Corporation

Peripheral Controllers

Now you can make the most of your DEC® system's storage. Micro Technology's Winchester controllers achieve total RL01/RL02 emulation on drives never before possible. With our optimized recording technique, you can increase formatted data storage capacity up to 41.6 MBytes.

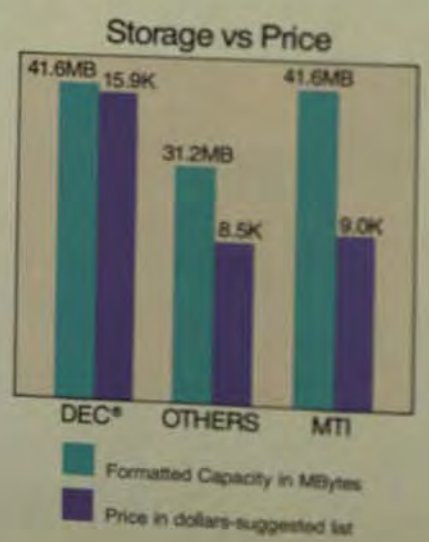

That's 10.4 MBytes more than competitive systems provide. Order the controller alone and do the simple plug-in integration yourself, or buy the complete storage system including disk drives and you're ready to go.

Write or call today for the key to 100% cost-effective data storage.

MTI
Micro Technology, Inc.

Proliferating Interface Standards Confuse Task of Controller Selection

Understanding the emerging interface standards for all sizes of magnetic media allows integrators to upgrade without costly hardware/software redesign.

by **Donald C. Peterson,** Intel Corp.

Technology advances have resulted in a dramatic increase in the capacity of magnetic mass-storage devices. Combined with a growing number of interface standards for the various media types, this has created a potentially confusing scenario for the OEM in his selection of mass-storage media.

An understanding of current and evolving interface standards should enable the OEM to choose an interface and peripheral controller that will allow him to remain competitive, and to upgrade an existing system without having to redesign controller hardware or write new software drivers.

There are two basic kinds of interfaces: device-level or dumb interfaces, and intelligent interfaces. An example of a device-level interface is the ST 506 5.25-in. Winchester interface. An example of an intelligent interface is the SASI (Shugart Associates System Interface). Device-level interfaces are most advantageous for large OEMs who maintain an engineering staff with controller expertise and believe they can gain an advantage over their competition by developing their own controller.

An intelligent interface is best suited for small- and medium-size OEMs who lack controller expertise, and OEMs who have multiple products and want to minimize the expense and resources needed for controller development. With an intelligent interface, an OEM can upgrade his product line by adding a different type of drive to the intelligent interface bus, without having to buy or develop a new controller for the new drive.

THE FLOPPY MARKET

Markets for the 5.25- and 8-in. floppy have standardized on the Shugart floppy interface. This has helped to expand the 5.25-in. market and allows relatively easy conversion to the smaller, lower-priced drives because the same controller can usually be used for either drive.

Four potentially different standards are developing in the sub-5.25-in. microfloppy market. Sony, the first company to have drives available, has chosen a 3.5-in. form factor and hard-shell diskette. Twenty-two other drive, media, and desktop computer manufacturers have chosen to align with Sony and promote the 3.5-in. microfloppy standard. The 3.5-in. drives use the same interface as the 5.25-in. minifloppy drives.

Three other form factors are also vying for microfloppy supremacy. IBM has a 4-in. drive that isn't minifloppy-compatible. Hitachi leads an industry group of about 17 manufacturers (primarily Japanese) supporting 3.0-in. drives that use a hard-shell diskette and are minifloppy-compatible. Dysan, Seagate, and Tabor support a 3.25-in. drive with a soft-shell diskette (similar to 8- and 5.25-in. diskettes) that's compatible with the minifloppy interface.

Because the 3.5-in drive is the first available, is interface-compatible with the 5.25-in. minifloppy drives, and has the largest number of backers (including Shugart), it likely will become an industry standard. More than one standard may exist in this large market, however, so other sizes may also be successful. When

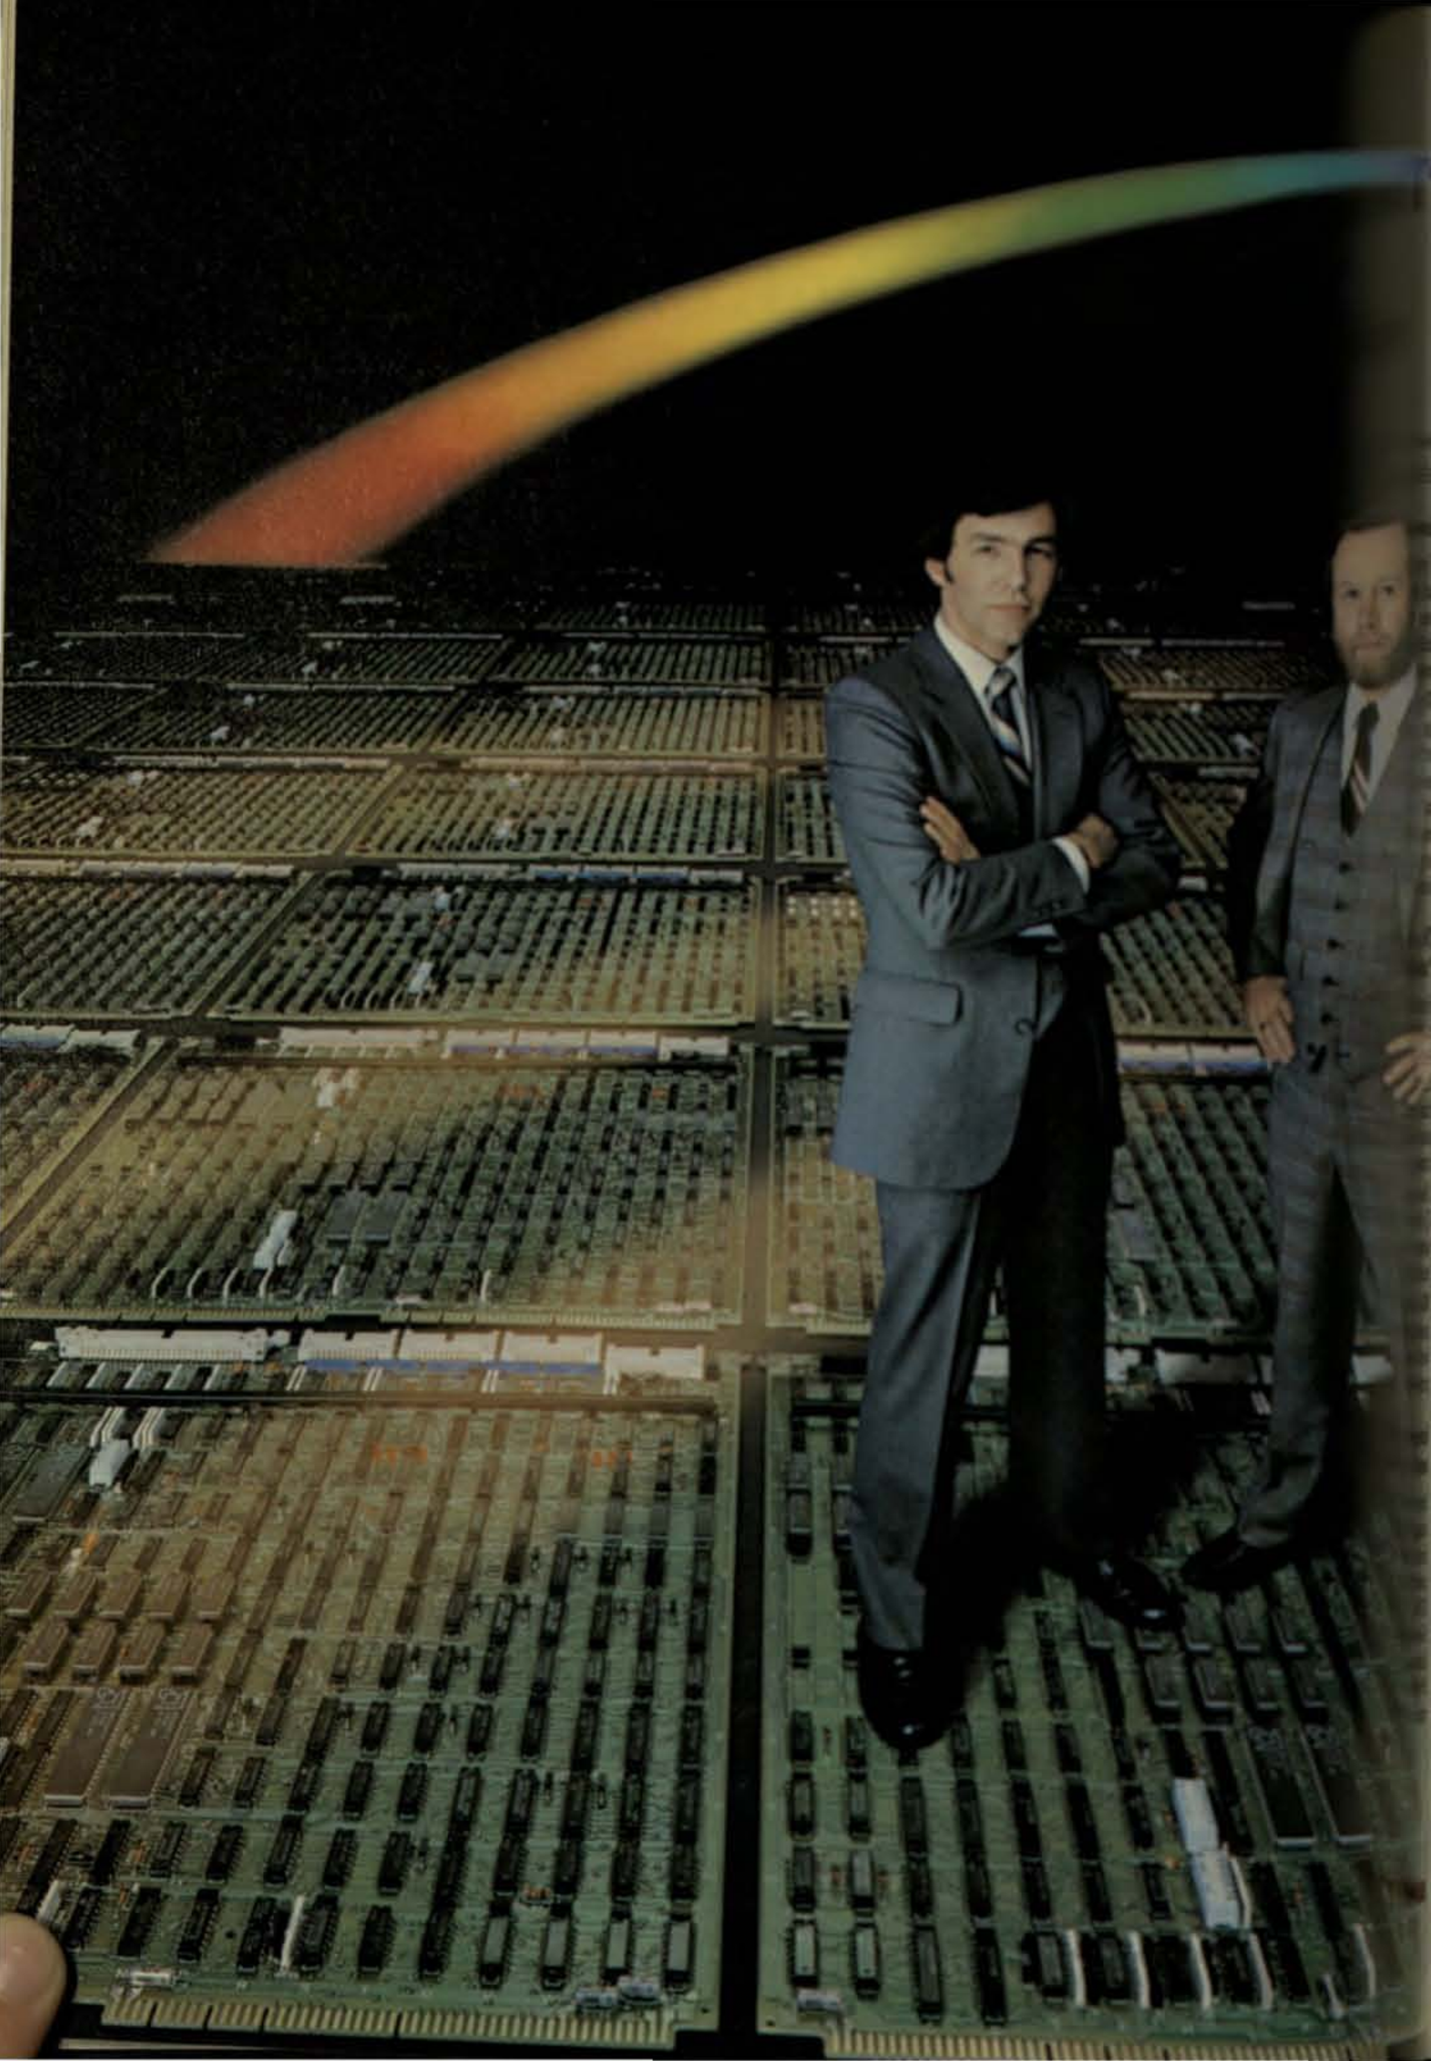

"Thousands of installations worldwide prove our concept: *multifunction controllers work.*"

Steven A. Roberts and Allan L. Abbott, co-founders

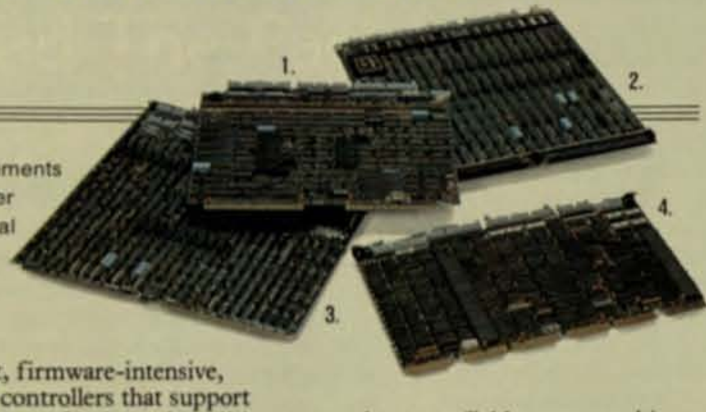

1. Texas Instruments
2. Perkin-Elmer
3. Data General
4. DEC

Tomorrow's solution today.

If you're an OEM or systems house and currently using disk and tape in your system, then Spectra Logic's field-proven multifunction disk/tape controllers are for you.

Why pay extra for separate disk and tape controllers when you can have both functions today on a single board? With improved performance and features. For example, unlike most single function controllers, all our products support the next generation of 2MB/second transfer rates.

The company behind "multifunction."

Spectra Logic first introduced the multifunction concept back in 1979. And we've been quietly revolutionizing the market ever since with controllers that provide the high-performance and added value you need to stay competitive.

Our new family of high-performance Q-bus disk/tape controllers is the latest example of Spectra Logic's technological leadership.

We also offer the industry's most comprehensive one year warranty, as well as responsive, nationwide professional service.

Improve system performance and reduce system cost.

Spectra Logic builds single and multifunction peripheral controllers for DEC, Data General, Texas Instruments and Perkin-Elmer minicomputers.

All are smart, firmware-intensive, single board controllers that support the widest range of independent disk and tape drives while emulating the CPU manufacturer's disk and tape subsystems.

Our multifunction controllers eliminate the need for separate disk and tape controllers, saving you money and CPU slots, increasing reliability, and reducing power and spares requirements. A unique dual microprocessor architecture and separate disk and tape buffering not only support simultaneous transfers, but appear to the system as two separate controllers.

To enhance system performance even more, Spectra Logic provides our own streaming software, Spectra Stream™, for DEC, DG and TI operating systems. It can back up an entire 80MB drive in only seven minutes, and is the ideal solution to the Winchester and ½-inch tape backup problem.

What are you waiting for?

No matter who you're doing business with today, Spectra Logic's high-performance multifunction controllers should be in your future. To find out more about our field-proven multifunction controllers and the company that's revolutionizing the controller industry, simply return the coupon or call our national sales manager Bob Carter at (408) 744-0930.

Mail to:
SPECTRA LOGIC — CTR-SUM
1227 Innsbruck Drive
Sunnyvale, California 94086

Tell me more about your disk/tape controllers for:
☐ DEC
☐ Data General
☐ Texas Instruments
☐ Perkin-Elmer

Anticipated quantity: ________

I'm interested in:
☐ Single function
☐ Multifunction
☐ Spectra Stream software

My need is:
☐ Immediate
☐ 6 months or longer
☐ Reference only

Name ________
Title ________
Company ________
Address ________
City ________ State ____ Zip ____
Phone (___) ________

1227 Innsbruck Drive, Sunnyvale, CA 94086
Telephone (408) 744-0930 TWX 910-339-9566
TELEX 172524 SPL SUVL

Western Regional Sales Office: (214) 934-9294
Eastern Regional Sales Office: (216) 826-3137
New England District Sales Office: (914) 623-0502

Spectra Stream is a trademark of Spectra Logic Corp.
©1983 Spectra Logic Corporation

The *Multifunction* Controller Company.

SPECTRA LOGIC

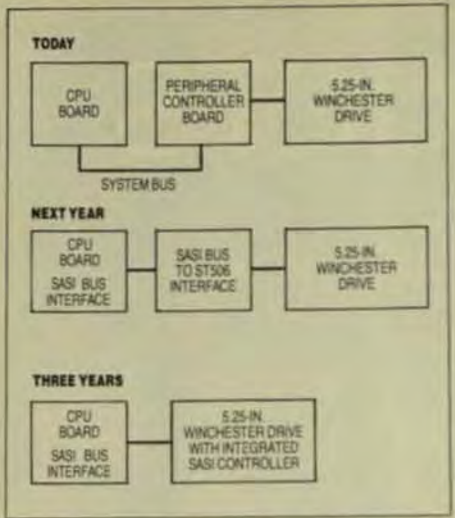

Fig The intelligent interface will evolve in the next three years, during which the intelligence will become part of the drive and the bus interface will be transferred to the host CPU board.

Two competing intelligent interfaces are now struggling for dominance of the high-performance drive and controller market. The ISI (intelligent standard interface) is supported by Control Data, which plans to use it on hard-disk drives 5.25 in. and larger, ½-in. streaming-tape drives, and optical disks. The ISI is a 16-bit interface that allows a transfer rate of 2 Mbyte/s today and up to 6 Mbyte/s in the future. CDC plans an enhancement to allow no-host involvement for drive-to-drive transfers.

The other proposed high-end intelligent interface is the IPI (intelligent peripheral interface) now being developed by an ANSI committee. The IPI appears to be trying to accomplish too many things within a single standard. Because of this, there are at least four transceiver types to handle varying distance requirements, two protocols to handle differing data transfer rates, and both an 8-bit and a 16-bit data bus.

With the large number of permutations possible, the IPI is no longer a single standard and care must be taken to assure drive and controller compatibility. At present, no IPI-based products are available. The Table provides an overall look at the present standards in various media and size segments of the market.

PERIPHERAL CONTROLLERS

The size and cost of controller boards have decreased over the past six years. Intel's first floppy-disk controllers were 2-board solutions that included a single-density controller priced at $995 and a double-density controller priced at $1290. Later, Intel introduced a single-board, single-density disk controller for $650, and then the VLSI floppy-disk controller chip made possible a single/double-density floppy controller that's only one-third the size of the original board. This board, however, doesn't contain the DMA controller included on previous boards.

In early 1981, Intel introduced its first Multibus Winchester-disk controller board for 8-in. drives. Because of the backup requirement for Winchester drives, other companies, such as Data Systems Design, Scientific Micro Systems, and Data Technology Corp., have introduced single-board products that integrate Winchester controllers with floppy and/or tape controllers. In addition, the data separator has been integrated onto the board (a requirement for the ST506 Winchester interface), which lowers costs and decreases card-slot requirements. The price for multiple controller boards is lower than for earlier boards that could control only the Winchester drive.

Controllers are also available for intelligent interfaces. Data Technology and Xebec offer controllers that provide an ST506 interface between the SASI bus and Winchester drives. This type of controller is available for interfacing to floppy and SMD drives and will be available for ¼-in. cartridge drives (using the QIC-02 interface) in the near future.

Intel has decided to include SASI as an option on future Multibus-based CPU boards. SASI will be incorporated into the parallel-port section of the CPU board and can be utilized by the insertion of two programmable array logic devices (PALs). With the increase in RAM and EPROM density, a single board can now contain the CPU, total system memory, and peripheral interface.

At present, several semiconductor manufacturers are developing VLSI SASI controllers and Adaptec is shipping samples of its SASI controller-chip set. It appears the chip set is still too expensive for drive vendors to incorporate into their drives but that should change as chip sets are driven down in size to a single component. As volumes increase and prices decrease, it will be feasible for drive vendors to make the SASI controller an integral part of the drive itself (Fig).

Also, Distributed Computer Systems is sampling an ISI controller with the Multibus interface, and CDC is developing an ISI-interface controller-chip set for use on its ISI-compatible drives.

Two or three years from now, when VLSI is available to allow cost-effective integration of the SASI controller into the drive, the peripheral controller module can be totally removed from the system. Then the SASI bus on the CPU board will connect directly to a wide variety of drives for the most cost-effective solution, one that doesn't require hardware or software changes when the system is upgraded. ■

Donald C. Peterson *is a marketing manager for MULTIBUS peripheral controller boards for Intel's OEM modules operation. He earned his BSEE and MS business degrees from the Univ. of Wisconsin.*

TABLE — CURRENT AND EMERGING INTERFACE STANDARDS

	EMERGING STANDARD →				ESTABLISHED STANDARD
FLOPPY		SASI			SHUGART
14-IN. HARD DISK	ISI			ANSI X3T9.3/128	SMD
8-IN. HARD DISK	ISI	SASI	SA1000	AND X3T9.3/128	SMD
5.25-IN. HARD DISK	ESDI	SASI			ST506
SUB-5.25-IN. HARD DISK		SASI			ST506
OPTICAL DISK	ISI	SASI			
½-IN. STREAMING TAPE					PERTEC
½-IN. START/STOP TAPE					PERTEC
¼-IN. STREAMING TAPE		SASI		QIC-02	
¼-IN. START/STOP TAPE	NO STANDARD				

CAD, CAM & CAN DO.

In the marketplaces of high technology, the company that has what it takes, takes nothing for granted. We think this healthy respect for the changing needs of our customers and the challenging advances of our competitors is the principal reason we are today's number one manufacturer of intelligent disk drive controllers.

It's a leadership position we intend to keep. And that intention is backed by a full spectrum of capabilities—in high performance and low pricing, in availability and reliability—that works synergistically to strengthen the value of each component.

In Our Business, Keeping Up Means Keeping Ahead.

It's done by maintaining technological and production edges. And at Xebec that means a couple of things. It means staffing one of the most elaborate computer-aided facilities in the business. Our CAD plant, featuring everything from digitizing systems to photoplotting systems, is used by our engineers—and often by our customers' engineers as well—to help define printed circuit board and LSI circuit design fabrication solutions, innovatively, accurately, quickly and cost effectively.

Computer-aided manufacturing helps translate this quality of design into quantities of the same order. And no competitor can lay claim to more sophisticated CAM facilities than ours. State-of-the-art robotics, including autoinsertion, autoloading and autotest equipment, are geared to handle volumes as high as 100,000 units a month to as low as a few thousand—keeping quality up and costs down, while keeping production here, on-shore, not just for patriotic but for very practical reasons that benefit our customers and ourselves.

The Performance Edge.

When we talk about "Can Do," we're talking about performance. We're talking about promises made and kept. Not low prices for products you can't get. Or that once you do get, don't work. We're talking about people. People who, from engineering to purchasing, from the manufacturing line to the management line to the bottom line, know what they are doing and where they are going. People who have an edge already. Because they draw from a larger bank of experience—and from an impressive and continuing investment in research and development. People whose worth is proven in the proven products they have taken to market. Like the S1410, today the best selling 5¼-inch hard disk controller in the world.

A Centrifugal Charter.

Being in a position of leadership has its prerogatives certainly. And of course it has its pressures as well. The pressures to keep what could well be called the centrifugal force of our business running at full speed—a force pushing out from the center, pushing prices down, pushing technologies ahead, pushing market share up and pushing competition back. We like it that way. And so do our customers.

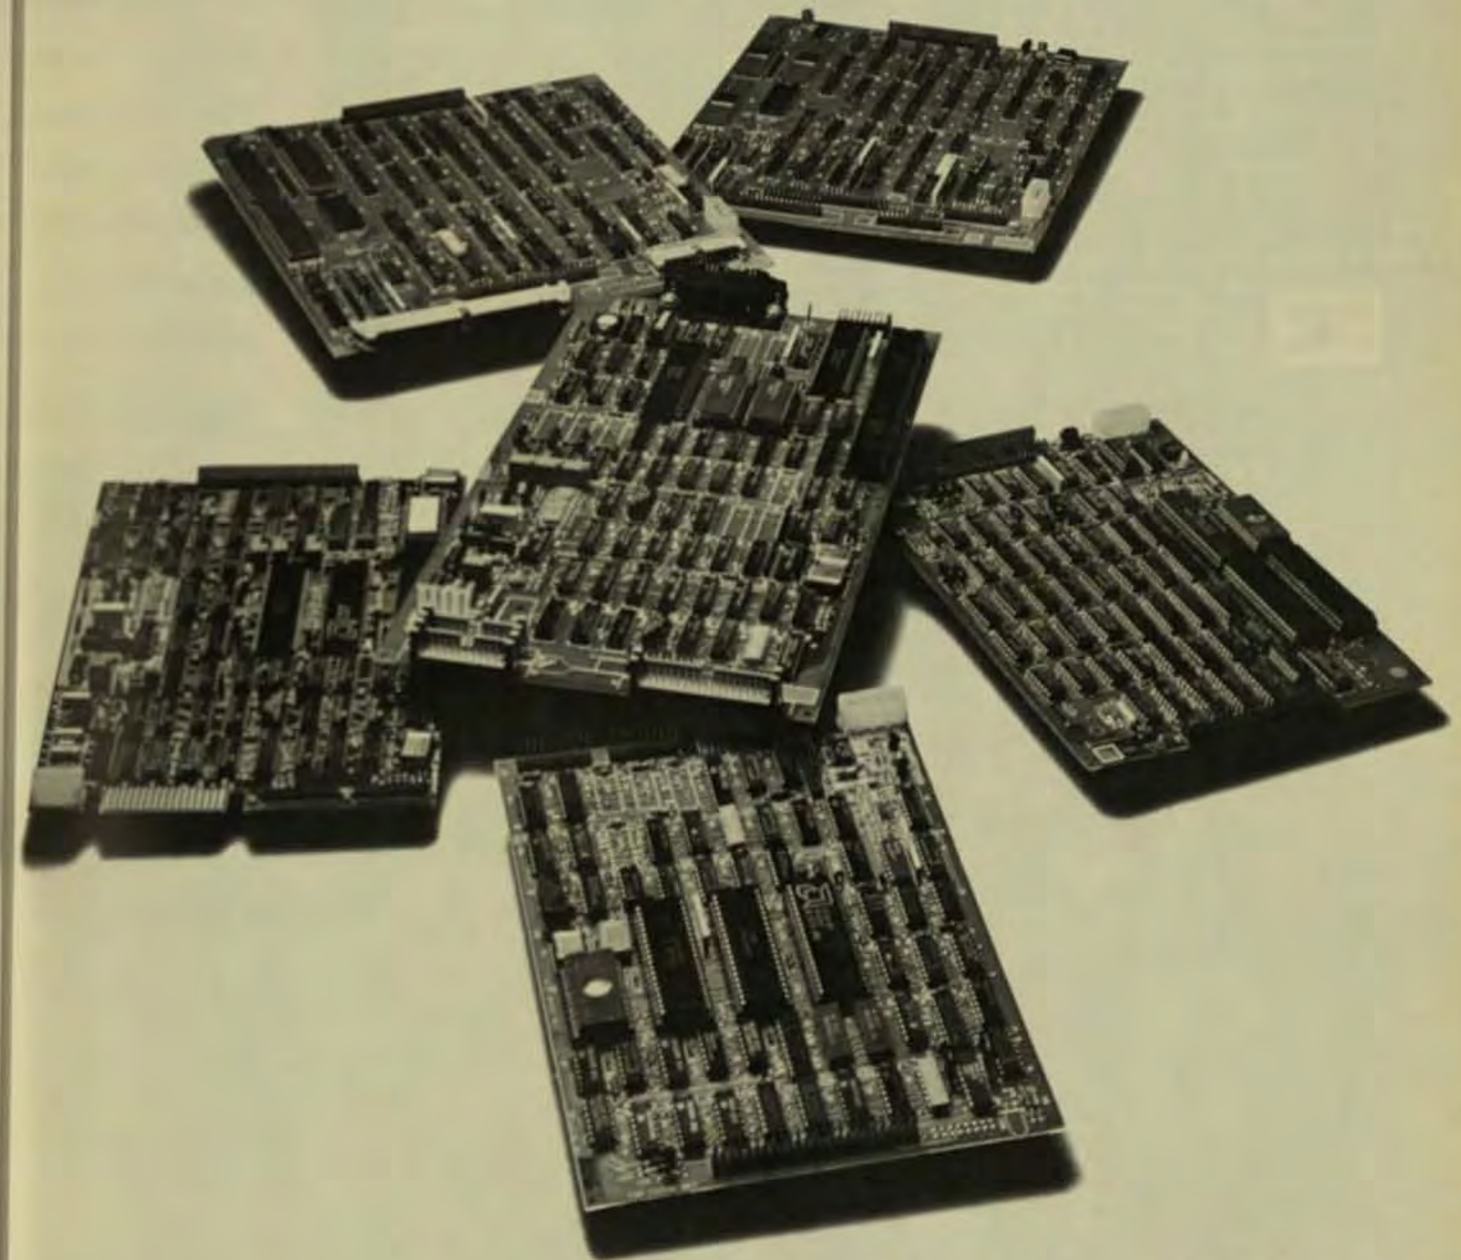

XEBEC
The Zero Defect Company
432 Lakeside Drive
Sunnyvale, CA 94086
(408) 733-4200

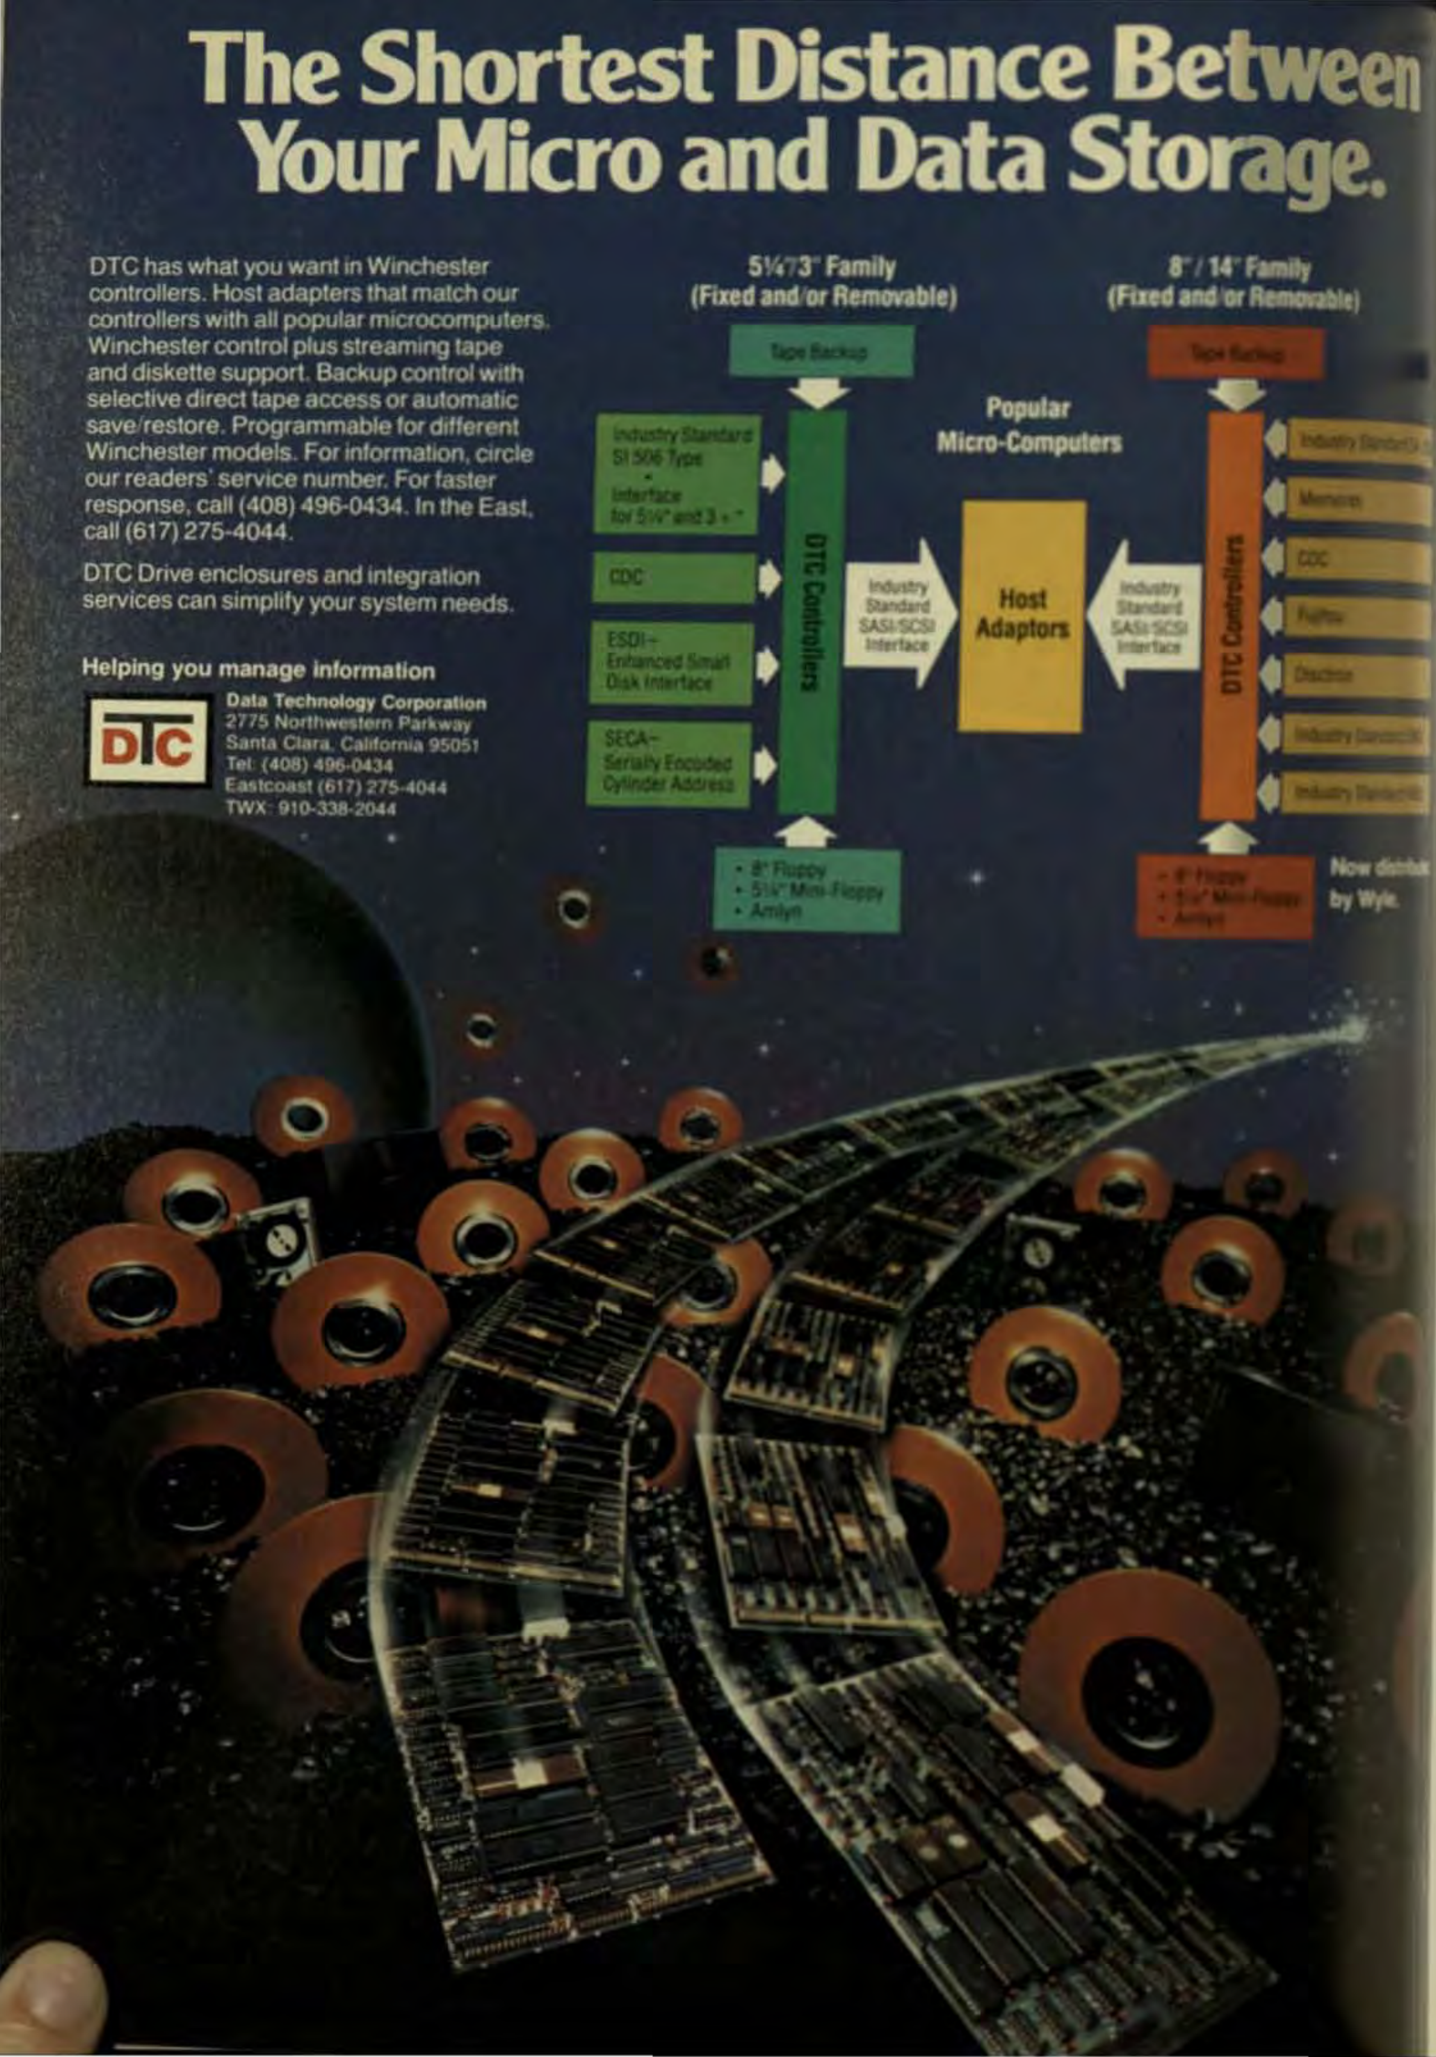

Graphics Standards Focus on Device and Programmer Interfaces

At the programmer level, GKS is the emerging world standard. VDI, a hardware/driver interface, defines a standard I/O protocol for device-independent software.

by **Mark Rawlins**
and **Mark James Brown,**
Graphic Software Systems Inc.

Two emerging standards are addressing the hardware/driver interface. One, based on videotex, allows consumers to access commercial databases by telephone line. For general computer-graphics applications, the virtual device interface (VDI) defines a standard I/O protocol and isolates the unique characteristics of the physical graphics device in the device-driver software module. The functions offered by the VDI are consistent with the conceptual model of the Graphical Kernel System (GKS) that is the emerging graphics standard at the programmer level.

Both the programmer and device interface standards now emerging are the result of over a decade of work by American and European organizations. The programmer interface refers to the conceptual model as well as the syntax the programmer uses when he incorporates graphics functions into an application program.

The device interface refers to the protocol between the device-independent and the device-dependent functions. The programmer-level interface standardizes the calling sequence and functions of a graphics procedure library, while the device interface defines a device-driver protocol that is consistent for all graphics devices (Fig 1).

THE GRAPHICAL KERNEL SYSTEM

GKS, the emerging standard at the programmer level, has felt the influence of many national organizations, including ANSI in the United States, and recently has been formally adopted as a draft standard by the International Standards Organization.

GKS achieves source-code portability. By providing a consistent interface in high-level languages such as FORTRAN it allows graphics-applications programs to be transported between different computer installations by providing a consistent interface in high-level languages such as FORTRAN. This ensures that graphics programmers will be able to work on different hardware without going through a costly learning curve.

GKS also provides a common graphics model and syntax to the programmer by standardizing the way in which graphics functions are accessed, and by providing graphics output on a virtual-device surface defined in normalized device coordinates. The application program may then control the way individual workstations interpret the normalized coordinates.

GKS supports a full set of drawing primitives, setting of primitive attributes, multiple workstations, segments, device-independent segment storage, and a comprehensive set of input primitives. It also supports raster graphics through area fill and cell (pixel) array primitives. While GKS provides device independence for standard functions, nonstandard operations are also made available through the generalized drawing primitive, a well-defined escape mechanism that allows access to the unique graphics capabilities of a particular device.

The basic drawing primitives in

GKS are the polyline, the polymarker, and text primitives. The polyline primitive draws vectors (straight lines) between a sequence of points specified as an array. A single line is merely a special case of the polyline that the operator defines by specifying both end points, rather than relying on a sometimes ambiguous and confusing

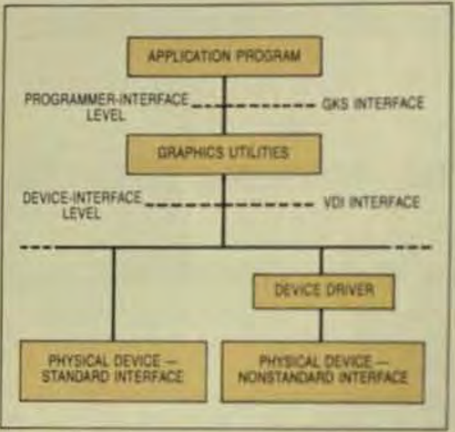

Fig 1 The two primary levels of graphics standards are the programmer and device-interface levels. The Graphical Kernel System (GKS) provides the standard interface between the application program and graphics utilities. The Virtual Device Interface (VDI) standardizes the interface between the graphics utilities and device drivers.

current-position model (proposed earlier by the Siggraph Core model). The polymarker primitive is similar to the polyline except that it draws a marker symbol at each specified point, rather than a vector. And the text primitive displays text strings at any position, with any orientation.

GKS also supports raster devices with fill and cell-array primitives. The fill operation paints the interior of a closed polyline (a polygon) with a specified color or a pattern such as a crosshatch. The cell-array primitive allows a 2-dimensional array of pixels of different colors to be defined. You can then replicate the cell over an arbitrary area by specifying the desired boundaries. These operations are critical in imaging applications such as video-frame displays, cartography, and other scientific areas.

Some graphics devices incorporate unusually powerful capabilities into their repertoire—for example, the ability to draw arcs, circles, and bars. GKS allows an application program to access these capabilities through a special escape mechanism called the generalized drawing primitive that allows you to invoke any unique feature of the device by passing a function number to the driver with the required parameters.

Associated with each output primitive are attributes that alter the object's appearance. For example, the polyline primitive has line type (such as solid or dashed), width, and color attributes. Polymarkers have attributes of type (. / o x), size, and color. Text primitives have attributes of size, color, and orientation. In addition, multiple fonts can be accessed if they are available in the graphics device.

You can define color indices by associating a desired color specified in RGB (red, green, and blue) intensities with a color-index number. The color values of the primitives are then given as the appropriate index.

A GKS workstation is a single display surface and one or more input devices. Multiple workstations may operate in a single graphics session that might include, for example, interaction with a raster display, a plotter, and a storage tube. GKS provides the logical interface through which the application program controls physical devices by redirecting graphics I/O at any time.

GKS maps the coordinate space of the application (called the world coordinate space) to device coordinates through two sets of transformations—normalization and workstation transformations. First, GKS transforms world coordinates into a normalized device coordinate (NDC) space by setting a world window. NDC space is actually an abstract viewing surface, or an intermediary space between applications and devices. Then it transforms the NDC space into the device coordinates of the workstation.

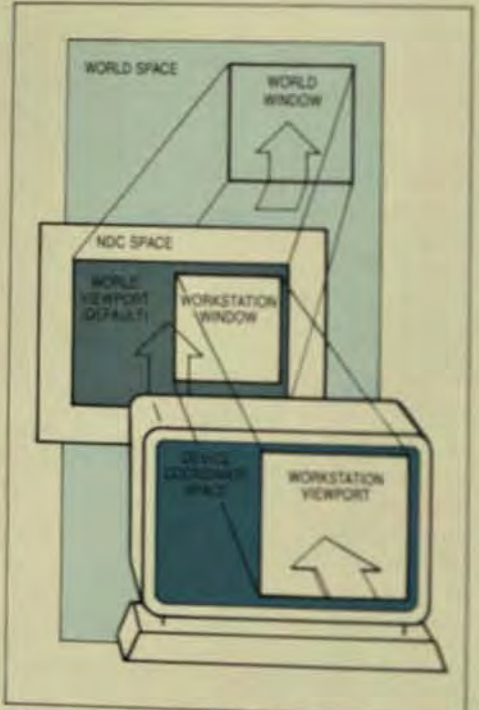

Fig 2 GKS performs two sets of transformations. First, it transforms world coordinates (the coordinate space of the application) into a normalized device-coordinate space by setting a world window that functions as an abstract viewing surface of an intermediary space between applications and devices. NDC space is then transformed into the workstation's device-coordinate space. Each workstation may, in turn, have its own distinct view of the application by setting a workstation window. Finally, each workstation can set a viewport that can be used to scale and translate the original picture.

In the case of multiple workstations, each may have a distinct view of the application by use of a workstation window. The last transformation lets the workstation set a viewport that can be used to scale and translate the original picture [...]

A full set of input operations allows an application program to receive input from a broad range of interactive input devices. The input operations are grouped into five classes: choice, locator, pick, string, and valuator. A vital capability, this input allows GKS to support the optimum input devices for a particular working environment. The result is improved interactivity through which the full potential of the graphics man-machine interface can be realized.

The request-locator function returns a position in world coordinates, while the request-valuator function indicates the current value of a continuous valuator device such as a potentiometer. The request-choice function returns an integer that represents one of a set of choices. The pick function returns the graphics segment number that corresponds to the objects being selected with graphics input. Finally, the request-string function reads character input from a keyboard device. These logical functions can be implemented in a variety of ways—joystick, mouse, function keys—depending on the workstation.

To aid the programmer, GKS also provides an inquire capability that allows the application program to determine the current operating state, primitive attributes, viewing operations and transformations, as well as device capabilities.

THE VIRTUAL DEVICE INTERFACE

Two emerging standards are addressing the hardware-driver interface. One, the North American Presentation Level Protocol Syntax (NAPLPS

How to choose the best peripheral controller for your DG or TI system.

1 Buy from a company with a complete in-house design and engineering staff.
Computer systems can change. Your controller should have built-in flexibility, and should get on-going support from the original design group. A "one-shot" consultant may not be around when you need him.

2 Make sure the software is compatible with your system.
A patched-in subsystem to modify your computer's operating system is a poor alternative when true emulation is available.

3 Don't get stuck with last year's technology.
Controller technology improves all the time. Independently manufactured peripheral interfaces should not only cost less than DG's or TI's, they should do more, too.

4 Buy from a company with a proven track record.
Promises are easy to make. It pays to find out just how happy other customers have been.

CSI is the leading independent manufacturer of peripheral interfaces for Data General and Texas Instruments minicomputers. All design and engineering work is done in-house. So you'll get the backup you need when you need it. And, our firmware is programmed to provide the most flexibility possible in your choice and upgrading of peripherals. All products are covered by our exclusive, two year warranty.

Most CSI controllers are true emulators, so there aren't any software patches to cause you problems down the road. When patches are required, a fully staffed software support team is available. Error detection/correction is standard. And our self-test diagnostics are among the most comprehensive in the industry.

If you're looking for a tape controller or streaming tape coupler, a winchester or large SMD disk controller, communications adapter, multiplexor, line printer controller, memory mapping controller, or remote diagnostic device, give us a call. We have a full line of products designed for your Data General or Texas Instrument's minicomputer. We'll be happy to send you our catalog, and give you the names of customers in your area that have been using CSI products for years.

CUSTOM SYSTEMS INC
6850 Shady Oak Road
Eden Prairie, Minnesota 55344, U.S.A.
Telephone: (612) 941-9480
Telex: 290975 CSI MPLS ENPE.

England: Tel: 0734/875464 Telex: 847395 • France: Tel: 736.87.00 Telex: 202878
West Germany: Tel: 0711/798047 Telex: 7255812 • Belgium: Tel: 02/345 9837 Telex: 64694
Sweden: Tel: 764/31580 Telex: 12986 • Denmark: Tel: 2/96 9822 Telex: 33514 • Switzerland: Tel: 1/520355 Telex: 822097

is based on the concept of videotex. Developed in Europe and modified in Canada, this idea was adopted by AT&T as a standard for transmission of text and graphics over telecommunications lines.

NAPLPS provides a way to connect television sets and other very low-cost raster devices to telephone lines so consumers can access commercial databases. Since it supports only one-way-at-a-time communication, it precludes interactive graphics applications. Although NAPLPS is important due to the number of devices that could potentially use this standard, it is not an optimal interface for general computer-graphics applications. It probably will sit below another more general device interface called the virtual device interface (VDI).

VDI is being developed by ANSI X3H33 technical committee as a standard interface between device-independent software and graphics devices. It makes all devices appear as identical virtual graphics devices by defining a standard I/O protocol and isolating the unique characteristics of the physical graphics device in the device-driver software module.

The VDI device driver performs several functions. It must communicate through a host hardware port or internal architecture path to a particular graphics device. (The port interface can be RS-232, Centronics, IEEE-488, or unique.) In addition, the drive must talk the language of the graphics device, using the appropriate control codes, escape sequences, or message formats to access its intrinsic graphics capabilities.

For example, a device driver translates between a VDI call to draw an output primitive such as a line, and the specific sequence required by a plotter (connected via an RS-232 line) to produce the requested line. The driver also controls graphics primitive attributes such as line width or color.

In some cases, the device driver must emulate functions specified at the VDI interface but not provided by the graphics device—for instance, a dashed line can be emulated by a series of short vectors (lines) with spaces between. On many displays, you generate marker symbols in the device driver by using special alpha characters or by using short vectors to draw markers. Information about the device characteristics can be returned to the application program to allow it to make some decisions about how to use the current graphics device.

The primitive functions offered by VDI are also consistent with the conceptual model of the GKS standard. The device drivers must perform the graphics functions supported: output primitives, primitive attributes, graphics and user input, control, and inquiry functions.

By defining functional capabilities, accessing methods, and parameter-passing conventions, VDI enables the development of device-independent software. A graphics application can be written to drive a generic VDI device with specific device drivers substituted as needed, either in a configuration process or at run time.

From an alternate perspective, VDI also allows equipment manufacturers to design software-independent hardware. In fact, the industry is more and more interested in building the VDI into the resident firmware of intelligent graphics devices. Ultimately, the VDI will be implemented in silicon as have been the IEEE-488 and Ethernet standards.

But how do you implement a standard in a technology that is anything but standard? Take, for example, the CMS operating system on an IBM 3081, CP/M on an Altos, Bell Labs' UNIX on a Digital Equipment Corp. VAX, and MSDOS on a Victor microcomputer. The virtual world of CMS lets VDIs be concatenated without regard to size.
to size.

On the other hand, CP/M requires "zaps" to the operating systems to load device drivers dynamically. MS-DOS, with its built-in alphanumeric-driver loader, can be fooled into loading graphics drivers also. And when UNIX and all of its derivatives are considered, the supply of standards that appear the same to the user borders on alchemy. But it will be done.

For the OEM, the adoption of this standard means that a VDI driver for a particular graphics device needs to be written only once. All graphics applications that conform to VDI will then be able the utilize the standard device driver. Long-range benefits will be more evident as equipment and semiconductor manufacturers begin to implement more of the software drivers's functions in hardware—in effect, moving the VDI interface down into the graphics device itself. This approach offers many benefits to the industry: less design effort expended re-inventing the wheel, numerous second sources, higher reliability with a proven design, reduced costs, and larger markets.

Several major vendors have already adopted the VDI standard as systems designers and applications programmers are beginning to experience the benefits of source-code portability. It's finally possible to develop a full set of graphics tools (plotting utilities, data analysis, drafting and IC layout) so that not only will programmers see a consistent interface to graphics functions at the language level, but compilers and graphics run-time libraries can be generic, with device dependencies residing in the operating system.

To the hardware OEM, adoption of this standard means that he only has to install the graphics communication to an operating system once. Compiler and other utilities that conform to the VDI standard can then access the virtual devices of a system without special adaptation. In time the hardware manufacturer, confident of a stable device interface, will begin to place higher-level functions into the device hardware or firmware. Eventually, graphics devices could incorporate a full VDI interface, entirely eliminating the need for device drivers.

The emerging graphics standards are a result of market demand. End users demand protection for their software investments, and they are asking for graphics solutions to the complex problem of how to transform data into information. What users don't need is a new set of problems involved with the selection and setting up of systems today and expansion and upgrades tomorrow. Manufacturers and OEMs are responding to the need. ■

Mark Rawlins *is director of marketing for Graphic Software Systems. He was previously director of marketing for ISSCO and has a BS in electrical engineering from the Univ. of Oregon.*

Mark James Brown *is a graphics software engineer at GSS, where he is a chief designer of graphics tools based on the emerging standards discussed in this article. He has a BS in math from Central Washington Univ.*

15% BACK TO VAX

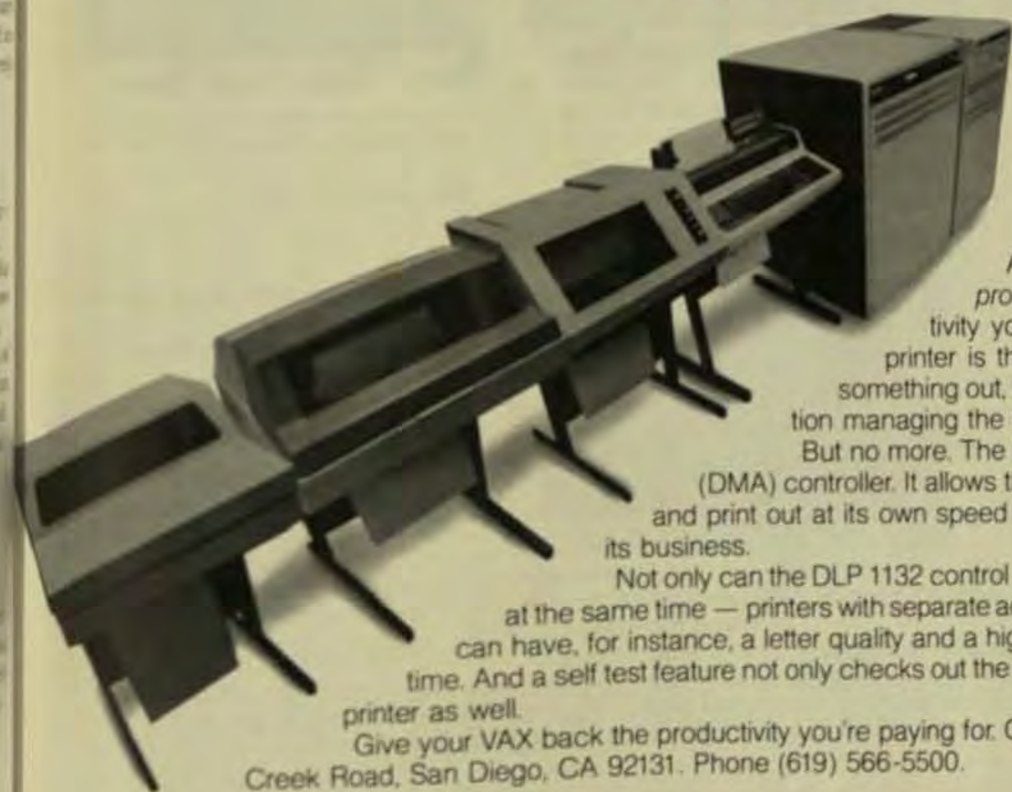

New printer controller increases VAX productivity 15%

Actually the DLP 1132 just gives *productivity* back to your VAX — productivity you didn't even know was lost. Your printer is the bottleneck. Everytime you print something out, the VAX has to spend time and attention managing the data output.

But no more. The DLP 1132 is a direct memory access (DMA) controller. It allows the printer to access memory directly and print out at its own speed while the computer goes on about its business.

Not only can the DLP 1132 control one printer, it can control two printers at the same time — printers with separate addresses and different speeds. So you can have, for instance, a letter quality and a high speed printer running at the same time. And a self test feature not only checks out the controller itself, but the cable and the printer as well.

Give your VAX back the productivity you're paying for. Call or write us today at 10072 Willow Creek Road, San Diego, CA 92131. Phone (619) 566-5500.

DATASYSTEMS™
a WESPERCORP subsidiary

The LSI-11 controllers with features comptrollers love.

Handles small 12-80MB disks, emulates RP02/03, RK06/07 and RL01/02.

Operates with new high-performance 80MB and larger disks on LSI-11/23.

Handles up to eight ANSI 8" Winchester emulates RK06/07 and RL01/02.

Handles 8-64 channels, 16-bit word transfers (DMA), 22-bit addressing. (16-channel price)

Dual density at up to 75 ips, handles up to four ½" transports.

Coupler works with all NRZI, PE, GCR formatters at speeds from 12.5 to 125 ips

Handles new cartridge Winchester disk drives, emulates RLV11/12.

*Quantity 100

Flinty-eyed finance types are push-overs for Emulex controllers.

And why not? Emulex controllers have great figures.

DEC LSI-11 users simply don't pay extra for higher Emulex quality, performance and reliability. In fact, Emulex controllers are priced to be highly competitive with the lower quality brands. And that's just the beginning.

There's no up-front investment because Emulex controllers are totally software transparent to LSI-11 CPUs. No special software drivers, handlers or tweaks required.

Smart comptrollers don't tie up cash in needless inventory. With Emulex's remarkably fast delivery, inventories can be kept economically small.

Quantity discounts are another opportunity for comptrollers with sharp pencils. With Emulex you can mix and match across a very broad product line of Q-Bus, Unibus & VAX controllers so it doesn't take long to bring prices down even further.

Knowledge that Emulex products work (at up to 72,000 hours MTBF) is important, too. But when support is necessary, comptrollers are pleased to know that nationwide applications assistance, training, and technical support is only a quick phone call away.

Great figures?

Write or call Emulex Corporation, 3545 Harbor Blvd., P.O. Box 6725, Costa Mesa, CA 92626. Telephone (800) 854-7112 toll-free. In California call (714) 662-5600.

EMULEX

The genuine alternative.

I/O Controllers Break Performance Bottleneck Between CPU and Disk

Advances in microprocessor performance and the increased capacity of mass-storage units require better controllers to optimize overall throughput.

by **Jeff Miller** and **Bill Breedon**, Adaptec

The phenomenal growth of the microprocessor market stems primarily from technological advances made in microprocessors and system memory. But these advances have occurred at the expense of the interface between the two, and I/O controllers were relatively neglected—the weakest link in a high-performance system. This article describes the advances made in microprocessor performance and the capacity improvements of mass-storage units, and compares them to developments in I/O control.

It wasn't until 1979 with Shugart Associates' development of the SASI I/O bus interface (later renamed small computer system interface (SCSI)) by the ANSI X3T9.2 committee) that a standard I/O interface capable of handling the data-rate requirements of Winchester disks even existed. The adoption of the SASI/SCSI interface standard by controller manufacturers was the first step toward solving an increasingly critical problem in system performance—the path between CPU and mass-storage units.

Breaking the I/O bottleneck between high performance CPU systems and Winchester disks required the elimination of several roadblocks, such as:

- Speed matching between the host and the disk.
- Defect handling on the disk.
- Error retries.
- Multi-user/multitasking support.
- Device independence.

HIGH-PERFORMANCE SYSTEMS

The microcomputer industry was spawned in 1974 with the introduction of Intel's 8080, the first truly low-cost industry-standard microprocessor. This introduction, coupled with the decreasing cost per bit of dynamic RAMs, established a base level for low-cost, high-performance microcomputer systems.

Shugart's 1976 introduction of the first commercially available 5.25-in. floppy disk expanded microcomputer memory beyond the capability of RAM. Similarly, Seagate Technology's 1980 introduction of the first commercially available 5.25-in. Winchester disk took microcomputer systems another step toward the high-performance capabilities previously available only in minis and mainframes.

Performance enhancements have occurred in two main areas. The first is in the direct-memory addressability of the microprocessor (Fig 1). The 8080 and its class of processors, followed by the 8085 and the Z80, could all directly address 64 Kbytes of memory. The 64K (8-bit) barrier remained until 1978 when Intel introduced the 8086, which directly addressed 1 Mbyte of memory.

This first 16-bit microprocessor was quickly followed by the Zilog Z8000 and Motorola's 68000, which pushed the addressability limit up to 16 Mbytes. To date, the addressability range of the popular 16-bit micropro-

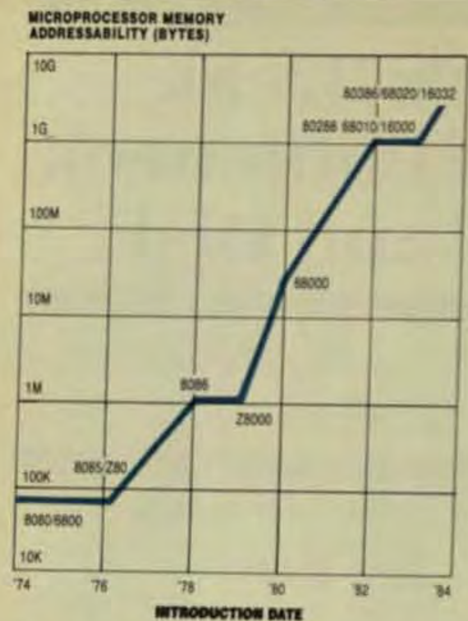

Fig 1 Direct-memory addressability in the microprocessor has increased by 64,000 from 1974, when the 8080/6800 was introduced, to the 32-bit chips that will be introduced soon.

cessors has been sufficient for the single-user microcomputer market.

Microprocessor capabilities have also moved into the multi-user arena, a marketplace previously reserved for superminis and mainframes. A microprocessor that supported virtual memory, the Intel 80286, was introduced in 1982. The 80286, coupled with the now available 68010 and the 16000, have pushed the addressability of microcompter systems into the gigabyte range. Soon, 32-bit microprocessors will extend this addressability to 4 Gbytes. Thus, across a 10-year period, the memory-addressing capability of microprocessor-based systems has increased 64,000 times.

Another way to measure the performance of a microprocessor is to look at the power and speed of its instruction set. With the 8080/6800 class of machines used as a benchmark and assigned an arbitrary performance level of 1.0, the performance

Fig 2 The 32-bit machines expected by 1984 will provide 100 times the performance of the 8080/6800 class of microprocessor that were released a decade earlier.

ranges associated with future processors is illustrated in Fig 2. Thus, over the same 10-year period, microprocessor performance has increased 100-fold from the 8080 to such 32-bit machines as the Intel 386 and the Motorola 68020, which are to be announced soon.

These improvements are all the more impressive when you consider that the basic pricing level has remained essentially the same. The 8080, the first 16-bit microprocessor, and the first virtual-memory 16-bit microprocessors were all in the $300 range when they were introduced.

Similarly, the capacity of mass-storage units has increased dramatically. Starting in 1976 with Shugart's introduction of the first 5.25-in. floppy disk, which had a capacity of approximately 100 Kbytes, the new floppy disks now contain 2 to 3 Mbytes. Likewise, 5.25-in. Winchester disks have grown from Seagate's original 5-Mbyte offering in 1980 to Maxtor's announcement of a 380-Mbyte disk.

SPEED MATCHING

Since high-speed disk drives transfer data at a much faster rate than an I/O bus can handle, throughput depends on a match between the processor and the mass-storage device. Buffers and interleaving are techniques used for speed matching. How these are used has a direct effect on system performance.

The interleaving technique involves the redistribution of sequential records on a disk track to decrease the effective data-transfer rate of the disk. An entire track of data, 32 blocks of 256 bytes, can be read in a single revolution of the disk if it is not interleaved. This is the highest performance possible. A typical interleave of three would require four revolutions of the disk to read an entire track of data.

System buffers are used to produce the same result and to prevent overrun conditions in the I/O device. These are caused when the I/O device is unable to start or stop instantaneously.

The system designer can use interleave factors and buffer space to fine tune the performance of the system to the application's needs. In most applications, it's inappropriate to use interleaving to reduce the performance of the disk to the lowest common denominator. Since the I/O controller should be capable of reading noninterleaved blocks, one bottleneck is eliminated. A dual-port buffer that functions as a circular FIFO can be used for this purpose.

Implementation of a dual-port buffer is very straightforward (Fig 3) using Adaptec's AIC-300 buffer controller. A 1-Kbyte buffer requires eight ICs, and if two more ICs are added, the buffer can be expanded to 64 Kbytes.

A controller can be tailored to simplify the buffering and increase the throughput of a block-oriented design. For example, a static RAM used as a dual-port circular FIFO can supervise simultaneous transfers on both ports and prevent overruns by stopping the transfer when appropriate. And logic to resolve simultaneous peripheral/host requests can grant priority to the peripheral and place a hold on requests to the host.

DEFECT HANDLING

Error-free media is difficult to obtain, so part of the system design must consider a scheme for the handling of media defects. The classic approach is to skip the media defect in question so that the system can never use the defective area. How the defect is skipped can have significant impact on I/O performance.

Most microcomputer disk controllers skip the entire track in which the defect is found. The replacement is a spare track on any innermost cylinder of the disk (Fig 4). Whenever a bad track is encountered, a seek must be issued to read/write data on each of the 32 blocks contained on that track. The seek operation is very time consuming (in excess of 400 ms on some drives) and has a drastic impact on the effective data-transfer rate. In addition, a single defect causes 8192 bytes (one track) of data to be retired from use by the system.

A more efficient method of defect handling is to skip only the block affected by the defect. With this approach, the replacement block is usually on the same track and almost always on the same cylinder. This ensures that a defect doesn't cause a data-transfer delay by performing an unnecessary seek.

This technique is used on some disk controllers that present a defect-free logical address space to

More Multibus.™ And less.

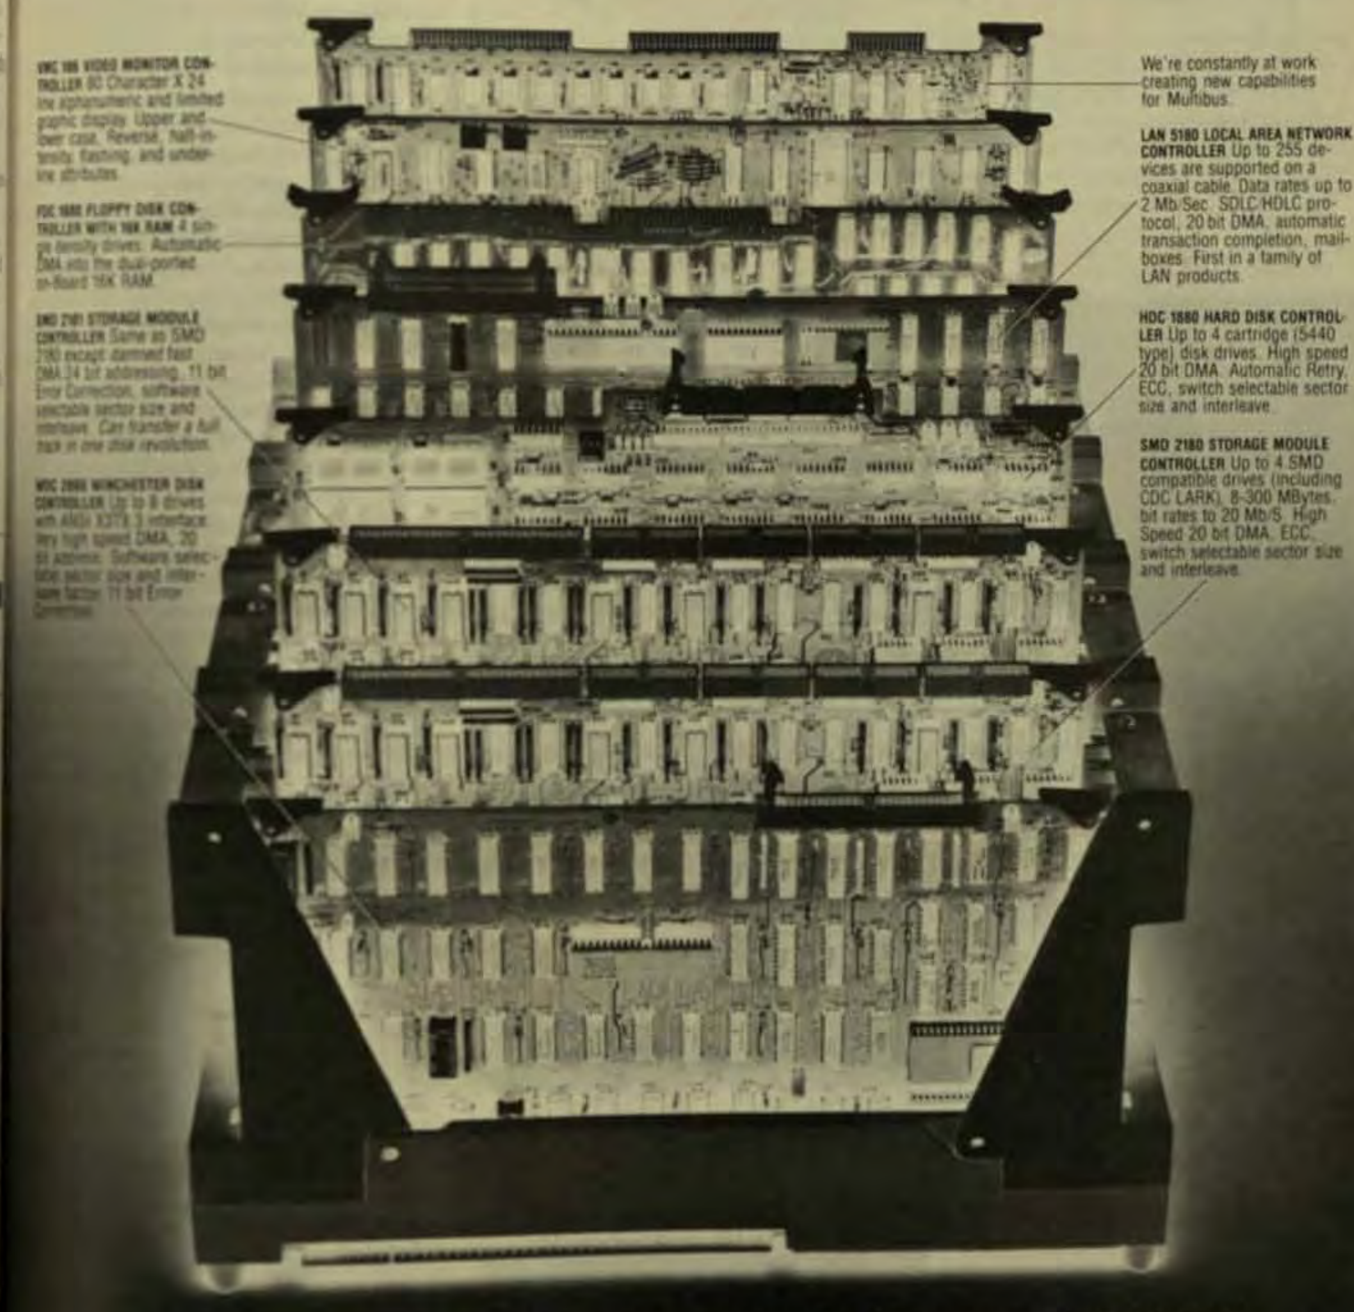

Now, the most talented Multibus specialists in the country bring you more Multibus options than ever before. And less.

More technology with less risk. More function with less cost. And more support with less hassle.

When you're ready to expand your Multibus capabilities, talk to us. We're the specialists who give you more. And less.

Entire rigid disk controller family is software interchangeable. Automatic completion of simple multisector commands, automatic recovery from errors, bad track mapping, full sector buffering, extensive diagnostic reporting. Variable burst DMA in 8 or 16 bit systems. Software drivers and disk subsystems also available.

INTERPHASE corporation

We stay ahead of our competition so you can stay ahead of yours.
2925 Merrell Road / Dallas, Texas 75229 / (214) 350/9000

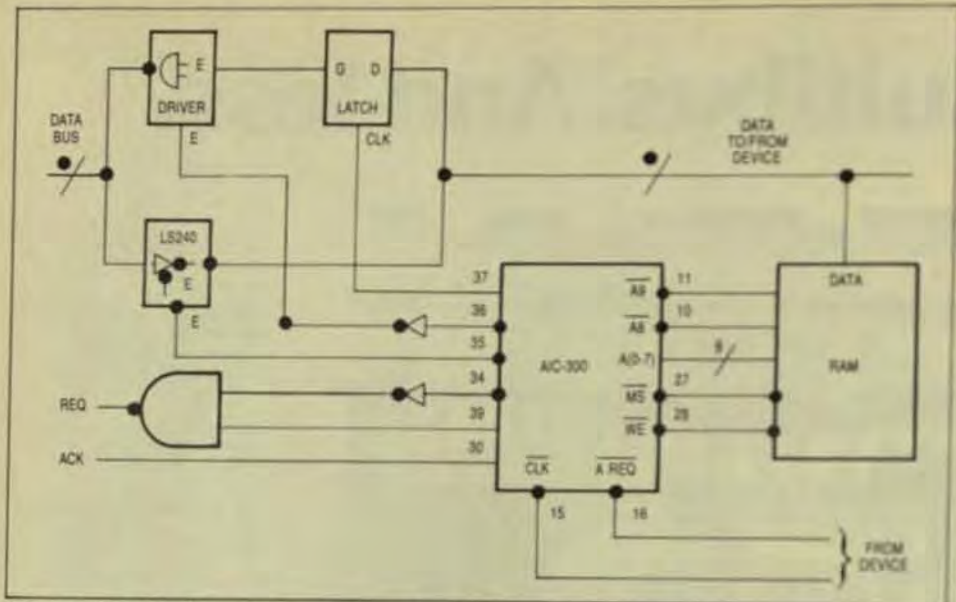

Fig 3 A 1-Kbyte dual-port buffer requires eight ICs, and two more if the buffer can be expanded to 64 Kbytes.

the host system. To preserve performance, the logical-to-physical address translation is corrected for missing defective areas by the controller so that seek time is minimized.

DEVICE INDEPENDENCE

A majority of I/O devices are electromechanical in nature and have unique control and error-recovery sequences as well as hardware interfaces. The unique handshaking requirements for each I/O device are passed through the controller, causing the creation of unique software drivers for each device. Relatively minor modification of an I/O device often necessitates software modification to many other system components.

New I/O devices can only be considered when a new system is in development because the cost to modify and support existing system software is prohibitive. This results in a limited range of I/O devices for the system and slow acceptance of new I/O devices and technologies.

A better approach is to incorporate the new devices into your product as soon as they are practical. But this requires device-independence in order to minimize the trauma of a new-product introduction. Device-independence can be achieved by moving device-control and error-recovery logic into the controller and having a standard system interface between the controller and the host processor.

The control differences between classes of I/O devices must be apparent to the host system because some I/O devices are output-only (printers), some are sequential-access (tapes), and others are random-access (disks). Within each class, the host system should treat all devices as functional equivalents. Standard functions should be implemented in a standard fashion, and only special functions should cause implementation methods to differ.

This approach means that any unique control sequences within a class must be absorbed by the device controller so that a standard interface will be preserved. In this manner, newer I/O devices may be added with little effect on other system components.

ERROR RECOVERY

Once a device-independent functional interface has been defined for each major class of I/O device, the only issue left is to handle exception or error conditions. If error detection and recovery isn't handled by the controller, the advantage of device-independent control has been reduced because these tasks have to be moved back into the system—and the recovery process can be very involved, consuming numerous system resources.

A defect-free logical address space is provided by some controllers, thus providing device-independence during normal read/write operation. However, to provide true device-independence, the controller must configure itself without host-system help, so the controller must know the physical parameters of the drive (number of cylinders, number of heads, step-pulse output rate, etc.). Some controllers require the host to pass these parameters to it every time the host is powered up or reinitialized. Since host software has to include this information, the principle of device-independence is violated.

Another common method for providing configuration information is to build it into the controller, but that doesn't provide enough flexibility given the number of drive types available. This problem is partially overcome by the use of some form of hardware switching, such as jumpers or dip switches, but the number of different drives on the market and the number of parameters required to specify a drive fully make this approach unsatisfactory. Even if it were economically feasible to provide jumpers for all the necessary parameters, initial and subsequent configurations would still be difficult because of manufacturing logistics and field-support requirements.

A recent approach to drive configuration is to write all necessary parameters on the drive at format time and recover them at each subsequent controller initialization. This approach requires these parameters to be supplied by a host system so that the drive can be formatted. The formatting process is drive-dependent anyway because the defect information provided by the drive manufacturer must be entered drive by drive. Once a standard is developed for manufacturers to record the defect locations and drive parameters directly on the drive media, this device-dependency can be eliminated.

Fig 4 Most disk controllers replace defective tracks by using a spare track on one of the disks in most cylinders as a replacement.

One way the industry can obtain device-independent operation is to adopt a standard system interface. The ANSI X3T9.2 committee currently is completing one, the small computer system interface (SCSI), that's supposed to present standard hardware and protocol interfaces to the system. In addition, SCSI allows the software device drivers to remain independent of drive characteristics.

The SCSI bus is a 50-pin flat cable with 18 active lines and 25 interleaved ground lines. There are nine control lines and nine data lines (eight data lines plus parity). The bus uses open-collector drivers and is terminated on both ends. The acceptance of this interface will allow both device and systems manufacturers to begin new-product production quickly and effectively.

DISCONNECTED OPERATION

In addition to device-independence, SCSI allows the implementation of a fully disconnected operation. This method allows the controller to receive a command and then disconnect, freeing the bus for further requests—a key requirement in a multitasking environment.

A multitasking system creates random-access requests to the disk system with an average of 50-ms access time spent for each request compared to data-transfer time of less than 1 ms per request. Disconnection allows the SCSI bus to be used by other controllers on the bus, or by the host when it issues a new task to a different unit on the same controller.

For example, on the Adaptec ACB-5000 controller, four operations could be active in the controller and the first unit to complete its access would be the first to transfer its data. This feature will allow systems to utilize higher performance Winchester drives more effectively, because drives with longer access times can continue to seek while data is being transferred from drives in which seeks are accomplished more quickly.

As multi-user systems proliferate, true multitasking can be accomplished by use of effective buffering and noninterleaved operation, efficient defect handling, and disconnected operation. Acceptance of a standard system interface provides additional benefits, including device-independent operation. Only through the implementation of these features will it be possible to eliminate the I/O bottleneck between high-performance microprocessors and I/O devices. ■

Jeff Miller, *vice president of marketing for Adaptec, was formerly employed as microprocessor operation marketing manager for Intel. He has also worked in a tactical marketing capacity at Fairchild Semiconductor.*

Phil Breedon, *Adaptec's vice president of engineering, has over 15 years experience in disk and tape controllers and I/O subsystems. He was formerly a partner in an engineering consulting firm specializing in intelligent disk drive design.*

MINI COMPUTER TECHNOLOGY

MULTIBUS*
SMD*
Controller

MULTIBUS*
ST506/SA1000
Controller

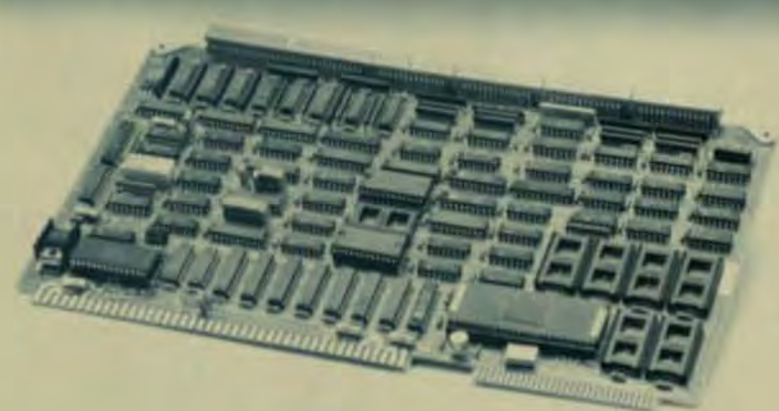

Mini Computer Technology offers microprocessor based intelligent disk controllers for Multibus applications. The Model MCT 4300 will control up to four disk drives with standard SMD interface - each drive can have different capacities. The Model MCT 4500 controls up to four disk drives of either 5¼" or 8" with an intermix of ST/506 or SA/1000 interface.

Call or write us for complete information

MINI COMPUTER TECHNOLOGY

Mini Computer Technology
125 Highgrove Street
Reading, Berks, RG1 5EN
United Kingdom
44 734 868033
Telex 847423 COCRG

Mini Computer Technology
696 East Trimble Road
San Jose, CA 95131
(408) 942-1616
Telex 910 338 2281

Mini Computer Technology
P.O. Box 1663
Apple Valley, CA 92307
(619) 247-4469
Telex 910 338 2281

*MULTIBUS is a registered trademark of Intel Corp. *SMD is a registered trademark of Control Data Corp.

UNCOMPROMISING DRIVE

"We produce Winchester disc drives of uncompromising quality. We build them in high-volume, and we deliver them to you on time."

—Alan Shugart, President, Seagate Technology.

The reputation of your system depends on the quality and reliability of key components. One bad drive can destroy customer satisfaction. That's why Seagate goes beyond industry-standard testing.

Our "zero-defect" program assures that every assembly, disc, electronic component, motor, and metal part meets our design and workmanship standards *before* assembly, then tested again in final tests where functional parameters are 100% verified and data integrity certified. Every drive is shipped with an individual printout of test results.

When you receive that drive, you can use Seagate's exclusive Receiving Test Station (RTS) that provides direct correlation with Seagate in-house test systems. When you assemble a Seagate drive in your system, you know it's right.

We back this promise with our industry-leading "105% Guarantee" program. If a Seagate drive doesn't pass your incoming inspection, we'll repair the unit and rebate 5% of your purchase price.

In 1983, Seagate will ship more than 300,000 Winchester disc drives from the most sophisticated Winchester production facility in the world. And every drive will be backed by the largest Winchester support team in the industry.

We also provide you with extensive technical and applications assistance, timely spares, and comprehensive documentation.

While we're proud that Seagate has been specified by the industry's largest and most quality-conscious OEMs, you don't have to be big to enjoy uncompromising quality, the lowest return rate in the industry, competitive pricing, and volume delivery. So call Seagate, the world's leading supplier of uncompromised 5¼-inch Winchester disc drives.

"Turning the tide in disc technology"

Seagate Technology

920 Disc Drive, Scotts Valley, CA 95066, (408) 438-6550, TELEX 176455SEAGATESCVL

Regional Sales Offices:
Hopkinton, Massachusetts, (617) 435-6961;
Newport Beach, California, (714) 851-9964;
Dallas, Texas, (214) 783-6711;
Schaumburg, Illinois, (312) 397-3727
Authorized U.S. Distributors: Arrow Electronics and Wyle Laboratories

European Sales Office:
Gassnerstrasse 5
8000 Munich 19
West Germany
49 89 177017

©1983 Seagat· ·ch ·c ·ogy

WE'VE GOT IT ALL

FROM 7MB TO 330MB IN 5¼", 8" AND 14" WINCHESTER FAMILIES.

Ampex's Winchester families deliver the performance you need in today's competitive marketplace. There's 5¼" Pyxis with up to 27MB; 8" Scorpio with up to 83MB; and 14" Capricorn with up to 330MB. All offer the features you want, like automatic disk and head locking; industry-standard interfaces and packaging; power-up self-test and diagnostics; and capacity expansion. None requires any preventive maintenance.

Our Winchester disk drives are backed by a continuing Ampex commitment to leading-edge technology research, automated processes, and vertical integration beginning with heads and media; by a highly professional worldwide service and support organization; and by manufacturing capacity in excess of a half million square feet located in four countries.

That's why you can count on Ampex to deliver whatever you need in Winchester disk drives, today and tomorrow.

Ampex Corporation, Memory Products Division, 200 North Nash Street, El Segundo, CA 90245. 213-640-0150. TWX: 910-343-6243.

AMPEX

Ampex Corporation • One of The Signal Companies

Head/Disk Interface Improvements Increase Data Storage Densities

Disk capacities and data accuracy can be increased by advances in disk-surface technology, data-bit-detection/track-location systems, and thin-film heads.

by **J.J. Newman** and **F.J. Sordello,** Memorex Corp.

As magnetic recording heads have improved from laminated metal (1960s) to ferrite (1970s) to the latest thin-film heads, particulate-oxide media has dominated and steadily improved. But four factors will continue to limit the use of particulate-media—signal-to-noise ratio (SNR), magnetic-coating thickness, coercivity, and defect density. Significant improvements in each of these areas can be expected to extend drive capacity until new technologies are available.

HEAD-DISK INTERFACE

Disk storage devices fall into two general categories (disk drives and disk files) based on whether the storage media is removable (disk drive) or whether the heads and disks are part of a nonremovable head/disk assembly (disk file).

The highest storage densities are obtained with files that employ a head/disk assembly (HDA) (Fig 1) in which the heads and disks are an integral unit. This packaging allows lower head flying heights and thinner media coatings since both head-loading forces and cleanliness problems are minimized.

Because the heads are designed to land on the disk surface as the disk rotation slows to a stop, the disks must be over-coated to minimize wear. The head-loading force at rest has been reduced to as low as 10G for the HDA compared to 350G for current removable disk-pack systems.

The technologies required for disk storage systems extend over many scientific and engineering disciplines, including physics, chemistry, mechanical engineering, electrical engineering, aeronautical engineering, material science, and manufacturing technology. The current status of disk recording includes the specific technologies as applied to storage media, magnetic recording heads, track positioning of the recording heads, and record/reproduce electronics.

Recording heads cover two key areas of magnetic transducer and air-bearing technology. Recording heads or read/write heads consist of an air-bearing slider carrying the magnetic recording transducer. This slider "flies" like a tiny kite lifted by a wind generated by the rotating disk at film thicknesses of micro-inches above the disk surface.

The head/disk interface is primarily concerned with magnetic transducer and media technology, as well as the initial electronics stages which directly affect the available system SNR.

RECORDING-MEDIA TECHNOLOGY

SNR is fundamental to any storage or information-transmission system because it directly affects the accurate recovery of data. The higher the data-recovery accuracy required, the higher the system SNR must be. If noise-free electronics were available, the fundamental limit would be the magnetic storage media SNR. Current systems designed for an intrinsic bit-error rate of 1 bit in error in 10^{10} bits transferred require a media SNR of 26 to 30 dB. This value is the ratio of base-to-peak volts of signal per RMS volts of noise (Fig 2).

Fig 1 High storage densities are obtained with files employing a head/disk assembly in which the heads and disks are an integral unit. This packaging configuration allows lower head flying heights and thinner media coatings.

The sources of media noise in currently used oxide (particulate) media are surface irregularities (scratches, roughness), particle dispersion within the binder matrix (clumps and agglomerations of particles), and the physical size of the magnetic particles used in the coating. A number of decibels of SNR improvement are still achievable. However, particulate technology is approaching its limit and new media technologies will be required. Contenders include continuous-oxide films (chemically grown, chemical vapor deposited, or sputtered) and continuous metal films (plated or sputtered). All of these technologies show viability with SNR improvements of up to 15 dB.

In many respects, the most consistent and dramatic technological

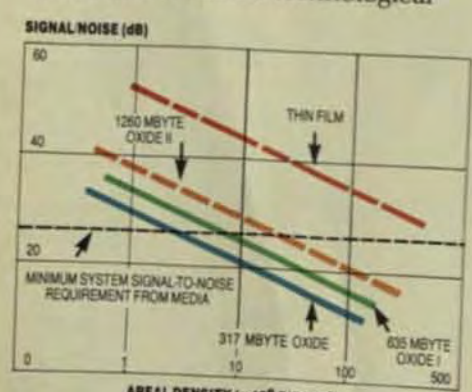

Fig 2 Current systems designed for a bit-error rate of 1 bit error in 10^{10} bits transferred require media SNR of 26 to 30 dB. As areal densities increase, the SNR varies, depending on the type of me ting.

change has taken place in the thickness of particulate oxide coating, which has decreased steadily over a quarter of a century from over 1000 µin. down to the present 25µin. used in the newest high-density storage systems (Fig 3). This reduction (40 times) has been accomplished while preserving the intrinsic SNR of the media and the relative defect density. At the same time, absolute defect density has improved over 5000 times.

Magnetic media coating thickness directly limits longitudinal recording density (Fig 4). High linear-recording densities require thin media coatings to minimize the effects of adjacent bit demagnetization. Current 12-Mbits/in.[2] disk-storage systems use media-coating thickness of 25 µin.

The next linear data increase for longitudinal recording appears to require thicknesses below 20 µin. The manufacturing process for coating thickness below 20 µin. will require significant improvements in binder, coating, and polishing technology. In this range, one of the new technologies (oxide film or metal film) may be required to maintain today's levels of data-recovery error rates at the higher data densities.

Media coercivity has increased only 25% over the past quarter century. This is largely because heads haven't developed higher magnetic-field intensities at the flying heights used. Present flying heights (10 to 15 µin.) are approaching contact recording. In addition, as the high effic of thin-film heads is exploited h head write-field intensities wil (Fig 5).

Increasing the media coerc has the same effect on adjacent magnetization as a decrease in th coating thickness, but without de creasing the volume of magnetic rial contributing to the output s When coercivity is increased to oersteds and above, particulate coatings sustain a significant in crease in areal density.

One of the primary aims of media development of the past years has been to maintain or de crease the relative defect density creased (recorded) track and lin (magnetic transition or bit) densi crease the minimum size of the fect that the system considers a error. Track densities have incr from 20 tracks per inch (tpi) to 1000 tpi while linear densities ha creased from 100 flux changes pe inch (fci) to over 10,000 fci. Conse quently, faults that cause hard e have decreased in size so that in stead of being visible to the naked eye, they can barely be seen with

Fig 3 A dramatic change has taken place thickness of particulate oxide coating last 25 years, from over 1000 µin. to the p 25 µin. This reduction has been accomp while preserving the SNR of the media a relative defect density.

high-quality microscope. Future s tems will require media with even tighter defect requirements; ulti ly defect density may even deter the relative viability of the new tech nologies.

Defects arise from a multitude sources during manufacture and each method of manufacturing pr duces a magnetic coating with de sources that are peculiar to it. The only sources that are relatively un sal across all coating technologies

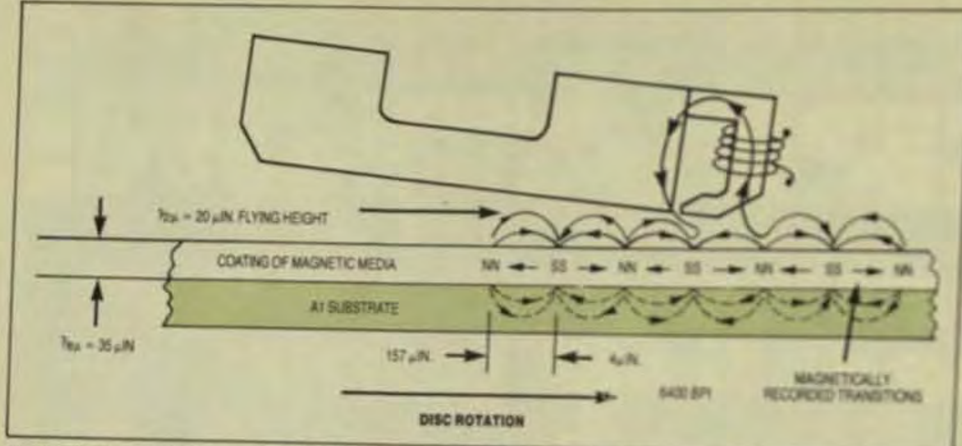

Fig 4 Magnetic-media coating thickness directly limits longitudinal recording density. High linear-recording densities require thin-media coatings. A Winchester technology ferrite-head/oxide-disk recording interface is shown with a coating thickness of 35 μin. and a flying height of 20 μin.

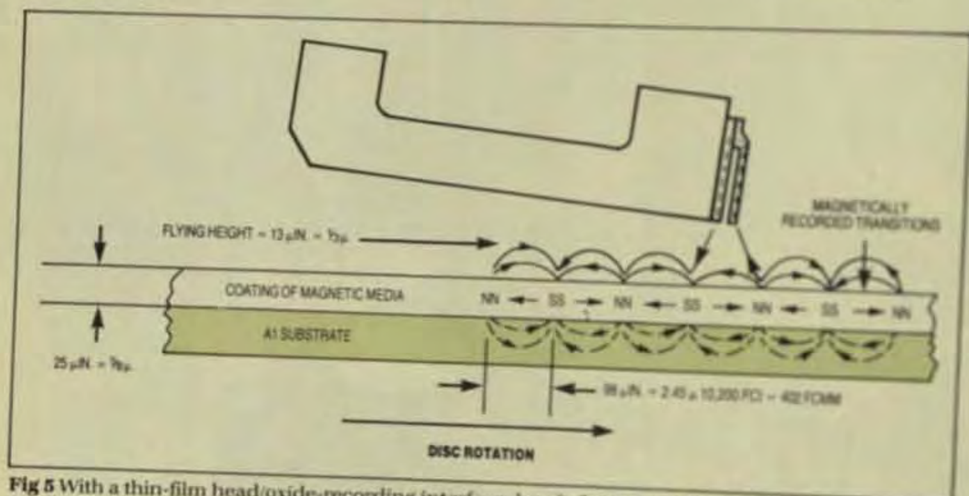

Fig 5 With a thin-film head/oxide-recording interface, heads fly at 10 to 15 μin. over a 25 μin. coating—almost contact recording.

disk-substrate finish, and cleanliness before and during coating. Thus, most disk manufacturing processes require a clean environment and significant effort is put into development of substrate-preparation techniques. The Table summarizes the most significant defect sources peculiar to several media-coating technologies.

RECORDING-HEAD TECHNOLOGY

Head-material technology in disk drives has come full circle. The first disk-drive systems (1956-1966) used heads made of permalloy laminations. These heads were replaced (1967-1980) by ferrite heads to obtain higher recording frequencies. Current ferrite-head manufacturing, which requires the transducer configuration to be machined by grinding and lapping, can generate enough stress to cause magnetically dead layers that appear as increased spacing losses.

Now, ferrite heads are being replaced (1980-onward) by thin-film heads made of deposited permalloy and other materials. The intrinsic high-frequency magnetic response of the thin-film permalloy, in addition to the absence of mechanically stressed materials, yields a head with wideband frequency characteristics, intense write fields, and outstanding trackwidth control.

The primary reason permalloy is once again usable at these frequencies is that advances in process-controlled semiconductor technology create thin layers in the magnetic

DEFECT SOURCES SPECIFIC TO MEDIA COATING TECHNOLOGIES	
TECHNOLOGY	DEFECT SOURCE
SPIN COATING (PARTICULATE OXIDE)	DISPERSION QUALITY FILM INTEGRITY SCRATCHES (POLISHING AND HANDLING)
PLATING	PLATING BATH CLEANLINESS COMPOSITION UNDER LAYER FINISH UNIFORMITY SCRATCHES (HANDLING) CORROSION (IF COMPOSITION OR DESIGN INCORRECT) PROTECTIVE LAYER (VOIDS OR PROTRUSIONS)
SPUTTERING	UNIFORMITY SCRATCHES (HANDLING) CORROSION (IF COMPOSITION OR DESIGN INCORRECT) PROTECTIVE LAYER (VOIDS OR PROTRUSIONS)

poles of the transducer, which minimizes eddy-current losses.

As areal densities (track density times maximum linear-recording density) have increased, there has been a steady decrease in trackwidth, gap length, gap depth (throat height), and flying height. Narrower trackwidths have increased the problems of adjacent-track pickup, so to minimize SNR degradation due to crosstalk, guard bands had to be provided between tracks and servo tracking had to be improved.

Shorter gap lengths require lower flying heights to keep write fields adequate and narrower gap depths to maintain the head's magnetic efficiencies. Consequently, a set of scaling rules came into existence regarding changes in trackwidth, gap length, gap depth, and flying height.

Thin-film heads have made it easier to satisfy several of these scaling rules. One significant change, for example, has occurred in the field gradients because the head can extend only a finite distance tangent to the track. And since the pole length of a thin-film head is typically less than five gap lengths, the vertical component of the head field is substantially increased in the region of the pole where writing takes place (Fig 6). Increased vertical field will be a necessary and welcome asset for future vertical-recording systems, which promise densities over 200 Mbits/in.[2]

The resulting modified magnetic transition and recontoured head-field response yield a slimmed output pulse on replay. Since better definition of the magnetic pole has decreased the requirements for a guard band for today's thin-film heads, a wider trackwidth for a given track density results. This improvement will be most significant at the higher track densities of future systems.

Another significant change is the ability to control the exact dimensions of each of the two magnetic pole (track) widths. The orthogonal fringing fields generated when one pole is made slightly smaller than the other minimize adjacent-track pickup. To further minimize adjacent-track crosstalk these poles can easily be made with 90° edges on the throat of the pole pieces. Photolithographic techniques will be used to manufacture the narrow-track thin-film heads required by future systems operating at 1000 to 2000 tpi.

One important advantage of thin-film heads is their potential for batch

Removable Storage With Hard Disk Reliability

A compact drive package which includes all drive electronics, power supply and SASI controller within 8.5" wide by 7" high by 19.4" deep. The VRC 8520 gives you 11 megabytes of removable storage on a standard ANSI 8" cartridge, and 11 megabytes on a fixed disk in a one-over-one configuration. Designed to fit practically anywhere, the VRC 8520 is ideal to produce backup files on a single peripheral. It can also accommodate an ANSI standard interface. Disk drive, controller and fully-integrated power supply are contained in a small, high-performance package designed for hostile environments. With the VRC 8520, you eliminate design costs for mounting enclosures, the costs of a separate controller and power supply and messy interconnecting wiring which radiates noise.

For more information about the VRC 8520 or the VRC 8010 — a single, removable disk drive with 11 megabytes of storage, contact Vermont Research Corporation, Precision Park, North Springfield, VT 05150, Tel: 802/886-2256, TWX: 710/363-6533. In Europe call or write: Vermont Research Ltd., Cleeve Road, Leatherhead, Surrey, England. Tel: Leatherhead 376221, TLX: 23280.

VRC
VERMONT RESEARCH CORP.
VERMONT RESEARCH LTD.

Memory Products for Systems That Can't Stand Failure

Unexplored space.

A few years ago Computer Memories launched an ambitious space program. We were determined to push the capacity limits of 5¼" Winchesters.

Without increasing the size of the drive. Without compromising performance and reliability. And without moving to unproven "the-jury's-still-out" media.

First, we broke the 12-megabyte barrier. Of course, others promised they could do better. But as they continued to promise we continued to deliver. And we went on to set new standards at 19, 26 and 40 megabytes.

In so doing, we created the industry of high capacity drives.

Today, with the lion's share of the market, Computer Memories is the leader in volume shipment of large-capacity 5¼" Winchester disk drives. And with our outstanding plug-and-play record, we also lead in reliability.

Of course, some things never change. Others still promise more and more megabytes. But when the industry needs a 60 or 80 or even 160 megabyte 5¼" Winchester drive, you can look to Computer Memories to deliver.

That's what leaders are for.

The capacity leader.

Computer Memories Incorporated

9216 Eton Avenue, Chatsworth, CA 91311 (213) 709-6445/ Ft. Worth, TX, (817) 332-7859/Costa Mesa, CA, (714) 241-9260/Springfield, MA, (413) 783-2386/In Europe: Bray-on-Thames, Berks, UK, 06283-4281 Telex 847591. For small quantities, call Schweber Electronics, (800) 645-3040.

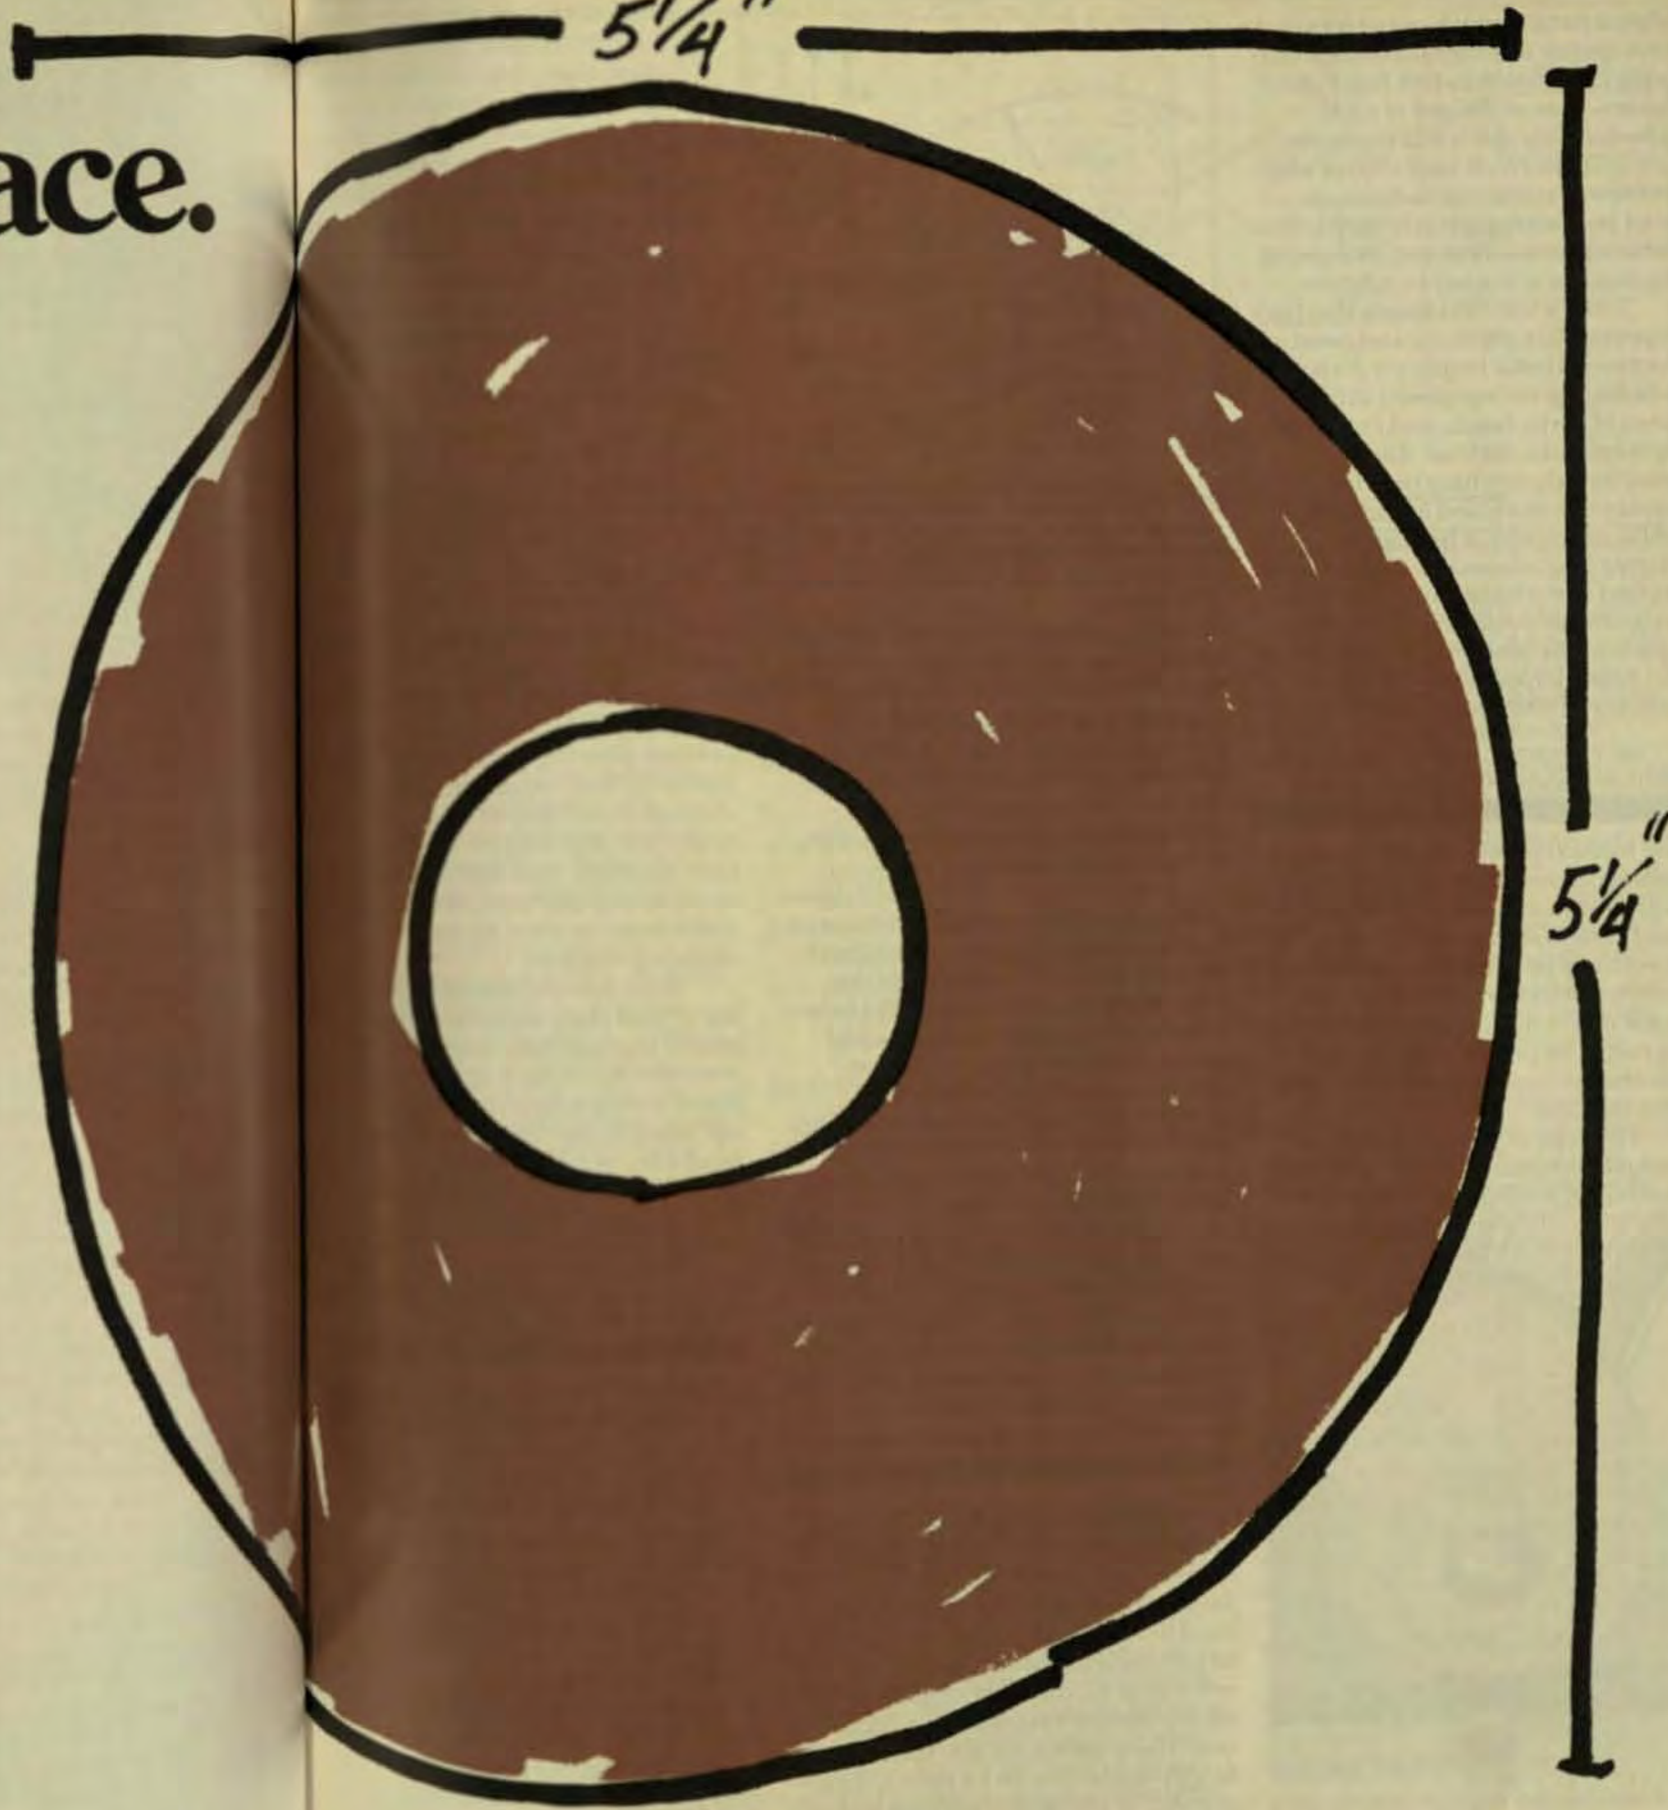

fabrication. Current head sliders that have two air-bearing rails (wings) normally carry two thin-film head transducers—one at the end of each rail—but only use one of them. Future systems could have sliders with multiple transducers to decrease head positioning actuator travel, thus reducing access time and increasing the number of tracks per cylinder.

Today's thin-film heads also have improved flux efficiency and head inductance. Useful frequency response is limited by the significant inductance of ferrite heads, and circuit capacitances. In contrast, thin-film head transducers have high frequency efficiency and low inductance, which allows high write-current frequencies at lower voltages applied to the head. This advantage is significant for systems with a high data-transfer rate (DTR). Current thin-film heads support data-transfer rates above 6 Mbytes/s.

DATA-BIT READBACK DETECTION

One method of storing data on disk is saturation recording. In this technique, the medium is fully magnetized at each storage cell. A change in direction of polarity or magnetization at each cell indicates the data. The heads detect this polarity change during read. The presence or absence of this change (readback pulse) constitutes the data.

This type of system requires a clocked detection window to identify whether an anticipated change (corresponding to a recorded data bit), has taken place. The information about a data bit's presence is contained in the peak location of the readback pulse. Since electronic circuits are adept at detecting polarity changes, the readback pulse is time-differentiated to convert peaks into zero crossings (Fig 7).

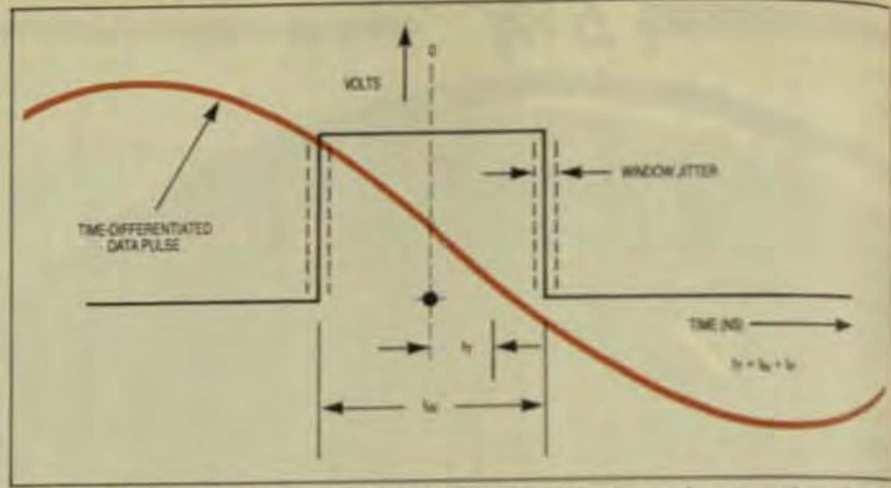

Fig 7 Data-bit presence is determined by electronic circuitry adept at detecting polarity changes in readback pulses. These pulses are time-differentiated to convert peaks into zero crossings. Since, ideally, the expected pulse occurs in the middle of the detection window, accuracy is maintained if the total peak shifts no more than half the window width available in either direction from center.

Over the years, the clocked detection window has evolved from a separate clock track that controlled both the writing and the reading of the data, to phase-locked-loop (PLL) electronics referenced to the instantaneous disk speed. Now, PLL electronics is combined with self-clocking data codes because high linear-recording densities require extremely high accuracy (nanoseconds) in data-detection timing.

These codes transform the input data into pulse patterns (sometimes referred to as clocked data) that include both the input data and the clocking information needed to synchronize the PLL-generated detection window. The pulses in these patterns each ideally occur in the center of a detection window, determined by the code.

Modified frequency modulation (MFM), in which a data bit is stored for each potential magnetic flux polarity transition, is a common data code used to store data on disk. Higher density systems are using data codes with up to 1.5 data bits stored for each potential magnetic-flux transition. These codes are also run-length-limited (no more than a finite number of transitions without polarity reversals), self-clocking, and use assigned code patterns to achieve increased density.


Fig 6 Since the pole length of a thin-film head is typically less than five gap lengths, the vertical component of the head field is increased in the region where writing takes place.

Several new detection technologies being investigated for future systems involve advanced detection techniques, optimal informational-content filtering, and channel linearization by bias recording. The overall strategy is an optimal trade-off between the pulse-width of the head-disk interface, and SNR capabilities to minimize the time shift of each pulse from its ideal location in the clocking window.

More advanced multilevel coding would allow more data bits to be stored per magnetic transition. Yet analysis of potential storage-density limits indicate that the ultimate storage capabilities (Shannon limit) of magnetic recording systems are far in excess of even the densities that can be achieved by the above improvements.

DATA-ERROR MECHANISMS

As with all data-storage systems, data accuracy and integrity are prime concerns. The intrinsic error rate is determined by the head/disk interface and its interaction with the data code. The most significant cause of error at the head/disk interface is noise- and pattern-induced peak shift.

Since the signals generated by magnetic heads are small, the electronic noise of the preamplifier must be included in any signal-to-noise considerations. Consequently, the total noise affecting the error rate includes components from media noise, electronic noise, system-generated noise, unwanted crosstalk from adjacent tracks, and inadequately erased signals from previous recordings.

When noise signals are added to the readback pulse, they produce a time shift from the ideal peak position of the observed pulse. The higher the noise amplitude, the larger the shift. Because noise amplitudes are statistical in nature, the distribution of noise-induced peak shifts is also statistical. A knowledge of signal and noise amplitudes then allows you to calculate the probability of peak shifts greater than the detection window.

Knowledge of the readback pulse shape, pulse width, and interpulse spacings is used to determine peak shifts generated by the pulse pattern that's written on the disk.

Several methods are used to make this calculation, including iterative computer modeling. However, simple pulse superimposition yields reasonably accurate peak-shift values as long as the minimum spacing is greater than two-thirds the width of the half-readback pulse, P_{w}.

In many cases, you can estimate noise- and pattern-induced peak shift by manual calculations using a graph of the time-differentiated isolated readback pulse obtained from the system.

You can obtain a worst-case value for the total peak shift, t_t, by adding the noise peak shift, t_n, and the pattern-induced peak shift, t_p. Then, since the expected pulse ideally occurs in the middle of the detection window, the system will function accurately if the total peak shift in either direction is less than the half detection-window width available in either direction from center, $|t_t| < \frac{1}{2} t_w$ (Fig 7).

In reality, an extra portion of the window must be set aside for other contingencies, such as jitter of the detection-window edges and manufacturing tolerances. This is referred to as window margin or phase margin and in good design practice is $\frac{1}{3} t_w$. Therefore, the total peak shift in either direction must be less than two-thirds of the half detection-window width: $|t_t| < t_w/3$.

As described, the intrinsic behavior of the head/disk interface is directly related to the isolated output pulse and the noise environment established by the media and basic electronic stages. The isolated pulse output is, however, a function of both media and head parameters. For example, consider a simple approximation to the output half-pulse width (in time units) from an isolated magnetic transition:

$$P_{w} = KS\sqrt{G^2 + 4(a + d)(a + d + T)}$$

where

G = (recording head) gap length (μin.)
d = (recording head) flying height (μin.)
T = media thickness (μin.)
a = media transition length (μin.) (length of magnetic flux change in media)
K = velocity (between recording head and media) (in./sec)/1000
S = slimming factor
{S 1 for ferrite heads
{S 1 thin film heads

for thin-film heads

Here gap length, flying height, and slimming factor are head parameters, and media thickness and media-transition length are media parameters. Both the pulse amplitude and other pulse-shape factors depend on the head and media parameters, so you cannot change either parameter without considering all of these factors.

FUTURE DIRECTIONS

The performance of the head/disk interface has improved many times over in the last 26 years. Areal storage densities (tpi × maximum bpi) have increased from 2 Kbits/in.^{2} to over 12 Mbits/in.^{2}. This is due to improvements in both recording heads and media. Better coating techniques have resulted in extremely low absolute error density for particulate media, and attempts to replace it in large-diameter disks by using other coating techniques (primarily plating) have failed.

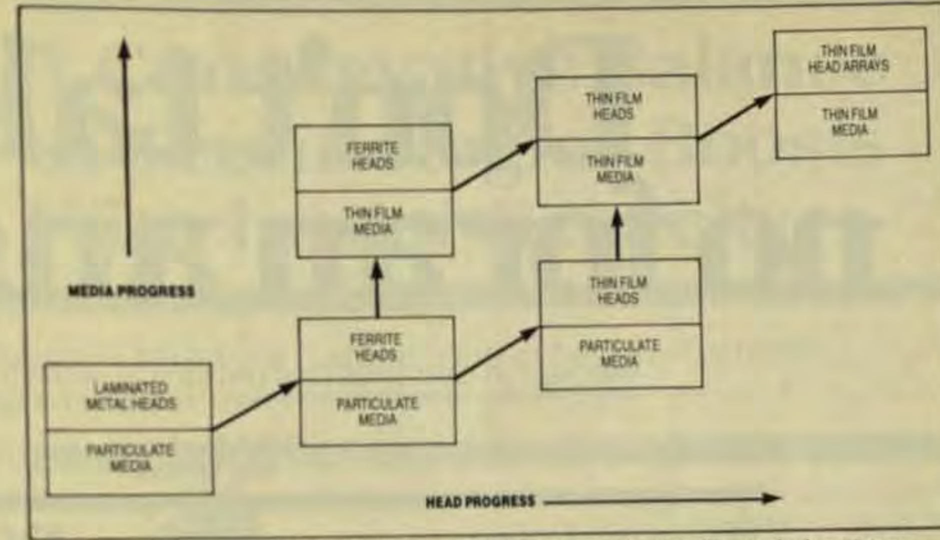

Fig 8 The performance of the head/disk interface has improved many times in the last 25 years. This is directly attributable to improvements in both recording heads and media.

Improvements in oxide particulate media, along with the perceived plated-disk problems of corrosion and durability, have allowed particulate media to maintain its dominance. However, the industry is approaching a transition point in media technology. Present densities are achieved either by thin-film heads on particulate media or ferrite heads on thin-film media. The next significant density increment will require the use of thin-film technology for both heads and media (Fig 8). This change is expected to occur at an areal density of 50 Mbits/in.^{2} sometime around 1984.

By 1986 the technology discussed in this article can be expected to produce an areal density of 50 Mbits/in.^{2} Extensions of these techniques, along with the application of vertical recording, promise to increase the areal density of magnetic recording to over 200 Mbits/in.^{2} by 1990. ■

John J. Newman *is the principal scientist and senior staff scientist to the recording technology center at Memorex. He is a member of the IEEE and the Magnetics Society, and holds BSEE, MSEE, and PhDEE degrees.*

Frank Sordello, *corporate vice president of technology for Memorex, holds 37 patents in the areas of disk drive and magnetic recording technology. He graduated from the Univ. of Santa Clara with a BSEE and an MSEE.*

Don't take no for an answer!

WINCHESTER CARTRIDGE DISK DRIVES FOR MICROCOMPUTERS

Feature	DMA Systems 5¼" Drives	3.9" Drives	Other 5¼" Drives
Now in production	yes	no	no
Data interchange	yes	no	no
Contamination control	yes	no	no
Retractable heads	yes	no	no
Standard cartridge	yes	no	yes
40 ms access time	yes	no	no
5 Mb capacity (formatted)	yes	yes	yes
Capacity growth capability	yes	no	no
Removable-only drive	yes	yes	yes
Fixed/Removable drive	yes	no	no

A comparison of Winchester cartridge disk drives shows there's really no comparison.

Only DMA Systems allows you to interchange data between drives.

Only DMA Systems has a unique retractable head that ensures data integrity by never touching the disk.

Only DMA Systems has a self-sealing clean air system that prevents contaminants from reaching the data —even after thousands of insertions.

And only DMA Systems gives you a choice of fixed/removable or removable-only drives. Both models match mini-floppy front panel dimensions and adapt to existing 5¼" Winchester drive controllers.

What's more, DMA Systems is the only manufacturer delivering microcomputer Winchester cartridge disk drives.

But DMA Systems drives don't just outperform other cartridge systems. They outperform all other types of microcomputer backup.

Consider the alternatives...

Floppies have low capacity, poor reliability and slow access time.

Streamers are unreliable and can't provide random access.

There's only one drive manufacturer who can answer "yes" to all of your needs: DMA Systems.

For more information write DMA Systems, 601 Pine Avenue, Goleta, CA 93117. Or call (805) 683-3811, Telex 658341.

DMA SYSTEMS

Just what the industry needs.

A Conservative Balance of Technologies Boosts Winchester Capacity

High-capacity 5.25-in. Winchesters require a combination of new and proven subsystem designs to produce reliable, manufacturable products.

by **Roger Stromsta**, Vertex Inc.

A new class of high-capacity 5.25-in. Winchester drives incorporates a number of sophisticated electrical and mechanical subsystems that differ substantially from those used by the first generation of low-performance, low-density Winchesters. The different design approaches that make possible these greater densities involve tradeoffs between proven and unproven alternatives. For the systems integrator, a careful analysis of the drive manufacturer's design choices can help in the selection of a reliable, high-performance product.

THE MARKET FOR HIGH-CAPACITY DRIVES

Advances in small computer systems, especially 16-bit microprocessor-based systems, have outstripped the advances in mass-storage devices (Fig 1). These increasingly powerful computers are being employed in multiuser systems that have multitasking operations, and in applications requiring large data bases. Previous 5.25-in. Winchesters were adequate for single-user business applications, but not for multistation applications or local-area networking.

To fill these needs, manufacturers are now releasing products with capacities of 30 to 100 Mbytes, or more, and access times of 30 ms. Drives in this performance class fit the same physical form factor as previous drives but incorporate major redesigns of their components and offer good price/performance characteristics as well. A rule of thumb is that a drive offering double capacity should command no more than a 25% price premium. These drives must be available in high volumes and at high levels of reliability, regardless of the technologies used.

A number of new and proven technologies in 30-, 50-, and 70-Mbyte drives strike a conservative balance between proven techniques and new technology. The major drive subsystems are:

- Read/write heads.
- Magnetic media.
- Spindle motor.
- Head-carriage actuator system.
- Servo system.
- Read-channel implementation.
- System controller.

READ/WRITE HEADS

Since the key to high capacity is increased track density, the read/write head is a major component in drive design. The IBM 3350 head and slider (Fig 2) has been field-proven for several years. When used with rotary actuators it permits a yaw angle of up to 13°. A wider yaw angle permits a shorter actuator arm—a vital consideration in the limited space available in the 5.25-in. Winchester package.

In contrast, some small-drive manufacturers are planning products based on the IBM 3380 Whitney technology. This type of head, slider, and flexure presents a lower profile, but

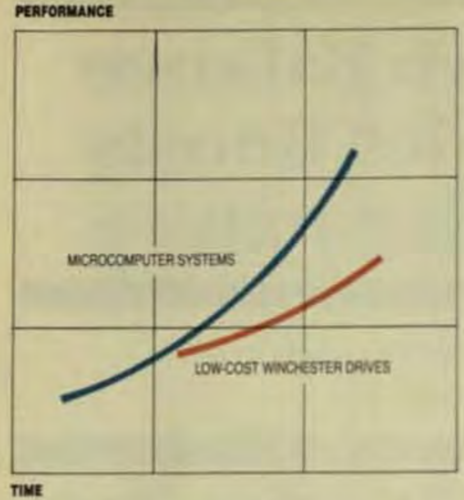

Fig 1 Technology-driven improvements in small computer systems have outpaced developments in mass-storage devices.

since it reduces the yaw-angle tolerance to about 8°, a longer actuator arm is required—and a longer arm is likely to be more sensitive to resonance, shock, and vibration.

The newer composite heads offer advantages in high track density applications. The core inserted into the ceramic slider can be machined to a 90° angle between the head's read/write core and the media. This angle concentrates the head's magnetic field and reduces track-fringing errors caused when adjacent tracks pick up the signals emanating from the head.

MAGNETIC MEDIA

Another major contributor to higher capacity in 5.25-in. Winchesters is the use of thin-film, plated media, which has a number of advantages over oxide-coated media. These include durability, better signal-to-noise recording characteristics, and increased areal recording densities of 50% or more.

In unusual circumstances the heads may contact the media and cause catastrophic head crashes that mar the recording surface. Plated media, while more expensive than oxide-coated media, is 10 times more durable. Thus, a reduction in overall costs may be obtained through decreased shipping and handling losses and increased yields in manufacturing. Capacity also increases, since plated media using standard ferrite read/write heads supports 14 to 15 thousand flux changes per inch (Kfci) compared to 9 to 10 Kfci with oxide-coated media.

Therefore, almost all manufacturers of high-capacity Winchester drives are planning to use plated media. But since early processes had oxidation problems and few drive manufacturers used plated media anyway, media manufacturers have been slow to supply it. Now that the manufacturing flaws have been eliminated and there's widespread agreement on the benefits of plated media, Ampex, Polydisk, and at least eight other firms will have it in production this year. Once availability problems are resolved, plated media may even become widely used in smaller-capacity drives, primarily because of its durability.

SPINDLE-MOTOR AND SPEED CONTROL

Newer high-capacity Winchester drives require a change from a 2-phase spindle motor to a 3-phase motor. These are more expensive, but provide more torque for faster start times and reduce the instantaneous speed variations required by the increased linear bit-packing densities. Vertex has chosen to stick with a proven external motor with a broad-base mounting because motor-in-the-hub assemblies aren't in volume production yet.

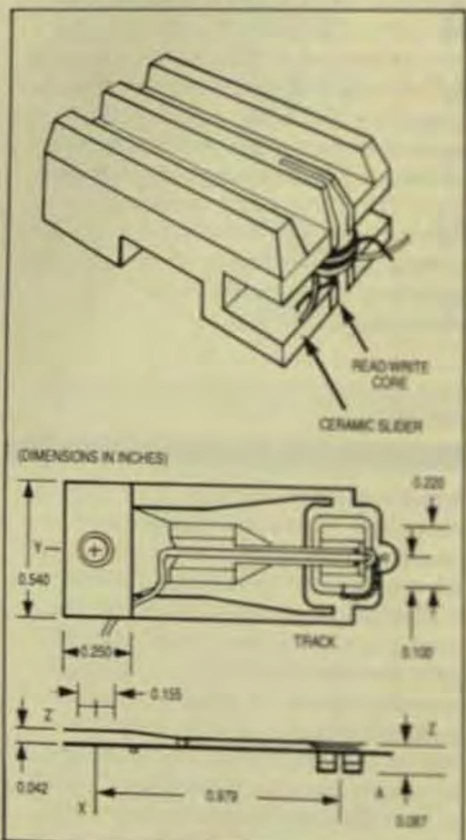

Fig 2 Composite read/write heads use ceramic sliders and manganese-zinc ferrite heads that reduce material costs. The 90° core sides minimize track fringing. Large winding windows allow more coil turns, and a lower inductance per turn is possible due to reduced flux leakage. There are fewer chips and voids along track edges.

Precise spindle-speed control is essential in disk formatting because the trend is to format in fewer but larger sectors. Two common types of circuits are used for spin-motor control—pulse-modulated and linear-drive. Pulse-modulated designs drive the control transistors completely on or off to produce a series of pulses. This switching induces noise into the system, affects data reliability, and tends to stress power supplies. The linear-drive circuitry drives the control transistor in the active region which, with a crystal-controlled phase-locked loop, provides less than a 0.5% tolerance error.

Systems that use a microprocessor for speed control have trouble maintaining tight tolerance since feedback loops with time-division errors are inevitable. Typically, these tolerances are no better than 2%.

ACTUATOR SUBSYSTEM

The three main types of actuator subsystems in use with high-performance drives are linear voice coil, rotary voice coil, and balanced rotary voice coil. Although some manufacturers are using linear actuators, these tend to be more complex and difficult to package in a 5.25-in. physical form factor.

While rotary voice-coil actuators cost less, take up less space, and require less energy to achieve fast access times, they are sensitive to shock and vibration and have natural resonance problems because of their cantilevered arm. These problems can be overcome by use of a balanced rotary actuator with a short actuator arm. This approach reduces the power consumption required to gain the desired access times.

The read/write heads and servo head mounted onto the actuator (Fig 3) are supported by precision ball bearings. A bobbin-type voice coil mounted between two permanent magnets provides the driving force to rotate the actuator arm. Track-stop limits are provided to protect the heads if a malfunction causes the actuator to lose control. When the drive is powered down, the actuator is automat

A Quality Family of Disk Drives. EXCLUSIVELY FROM FUJITSU.

Quality means more when it comes from Fujitsu. You get the latest Winchester disk drive technology. A total vertically integrated manufacturing operation. And a solid reputation built on more than 15 years experience.

Our broad base of customers have come to rely on this level of quality over the years. Quality that's exclusively Fujitsu.

For more information contact the Fujitsu America Sales Office nearest you. Northwest: (408) 988-8100, East Coast: (617) 229-6310, Southwest: (714) 558-8757, Europe: 44-1/493-1138.

	HIGH PERFORMANCE			MICRO-STORAGE	
	A 14-INCH	B 10½-INCH	C 8-INCH	D 8-INCH	E 5¼-INCH
CAPACITY (M Bytes)	84 / 168 / 336	474	48 / 84 / 168	24 / 48	7 / 13 / 20 / 27
AVG. POSITIONING TIME (ms)	27	18	20	70	83
TRANSFER RATE (K Bytes/s)	1,012	1,859	1,229	593 / 1,200*	625
INTERFACE	SMD	Modified SMD	SMD	SA4000	ST506/SA4000
POSITIONING METHOD	Rotary Voice-Coil	Rotary Voice-Coil	Rotary Voice-Coil	Buffered Stepper	Buffered Stepper

*48 M Bytes Configuration available only in 1200 K Bytes/s

FUJITSU

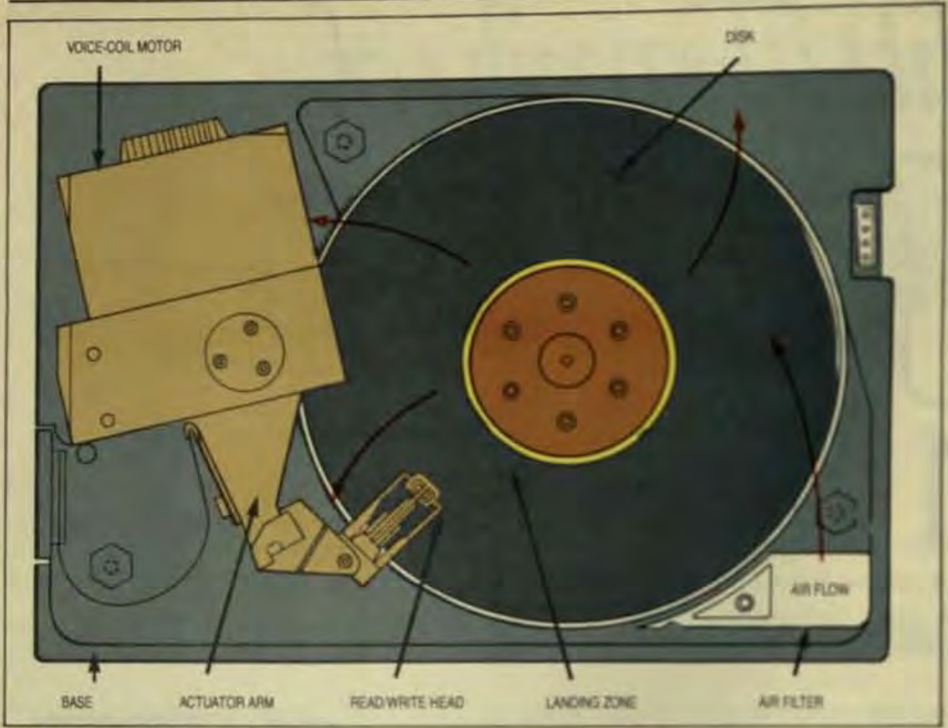

Fig 3 A balanced rotary voice-coil actuator may contain a number of mechanisms to protect the heads if a malfunction occurs. In a powered-down state, the heads are retracted to a safe landing zone and the actuator arm is locked into place to prevent any damage due to shipping or rough handling.

cally driven to a head-landing zone and locked in place on the inner diameter of the disk to prevent possible head or media damage.

SERVO SYSTEM

An important difference in some drives is the closed-loop servo, which uses a dedicated servo surface. The servo information is written on the bottom surface of the bottom disk. During seek operations, actuator movement is controlled by a dual-frequency technique based on amplitude differences between alternating servo tracks written at different frequencies (Fig 4).

On-track positioning is realized when the servo head is exactly between the two servo tracks. The servo head samples two frequencies, sums and averages the signals, and positions itself at the point where these signals are equal. This position will be exactly between the two tracks. Any resultant error signal is linearly proportional to the error (off-track) size, and indicates the off-track direction.

Dual-frequency servo positioning offers several advantages when compared to di-bit or tri-bit methods and is relatively insensitive to surface defects in the media. The error-free

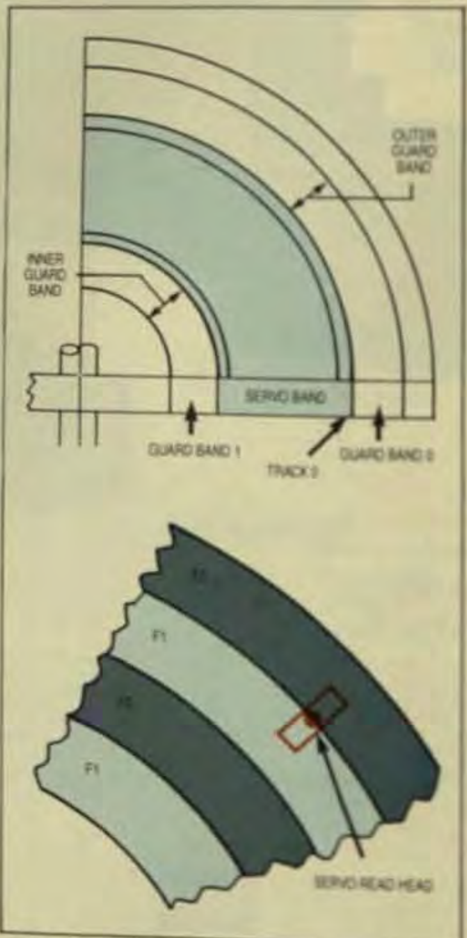

Fig 4 In a dual-frequency servo-tracking method that dedicates one surface to servo information, alternate tracks are written with different frequencies. The signal ratio is compared to determine precisely where the head is positioned.

servo surface generally required by di-bit/tri-bit methods reduces manufacturing yields—particularly for media with high track densities—and increases the cost of the servo disk. Since the dual-frequency servo continuously samples the two different frequencies, media defects are averaged out and not sensed by the servo.

In addition, servo-head azimuth alignment isn't critical because this method dosn't use pulse-position information, but in di-bit and tri-bit methods this alignment is crucial and further increases production costs and service requirements.

Because the two positioning frequencies are recorded well below the data frequency, write-data noise is easily filtered out. Also, track 0 is derived from the servo information so there's no need for a separate sensor as used in many drive designs.

Closed-loop servo systems offer not only higher performance than stepper-motor designs, but higher data reliability. The typical rule of thumb is that the maximum head-to-data track offset can be no greater than 10% of track pitch. Thus a high-performance 1000-tpi drive must maintain 100-μin. or less of off-track margin. Since stepper-actuator subsystems have no feedback loop to account for environmental temperature variations that can cause a shift in head-to-media registration, they're limited to an upper range of 400 to 500 tpi. Closed-loop servo systems can correct for head/media shifts and thus support densities up to 1000 tpi and higher.

READ-CHANNEL IMPLEMENTATION

One key to data reliability is the read-channel data path—which includes the data-separator circuits. These are currently located in the disk-drive controller, but because the performance of the data separator is critical in the prevention of data errors, future interface standards for high-capacity drives incorporate data separation directly in the disk-drive electronics. This change will help the drive manufacturer to match his data-separator design with the characteristics of his drives.

High-capacity drives use the same data-transfer rates as low-capacity drives—5.0 Mbps—and increased capacity is achieved through higher track densities. The limiting factors for track density are thermal

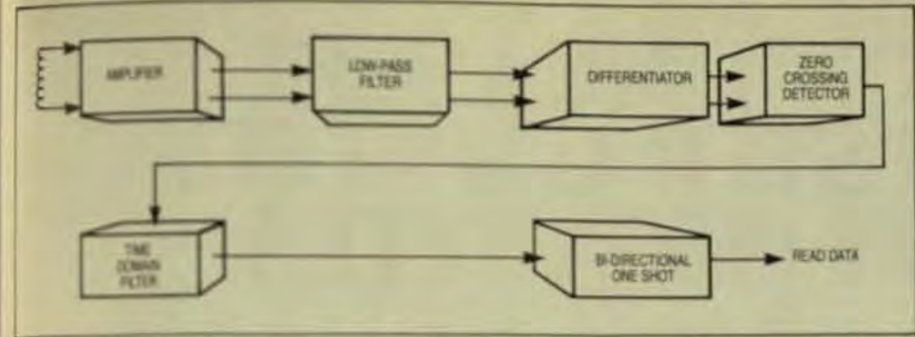

Fig 5 In the read-channel interface on a standard 5.25-in. drive, the data signal from the head is amplified, filtered, and differentiated. A zero-crossing detector denotes the peak of the readback signal.

expansion between servo and data surfaces, bearing noise, and servo tracking errors. In higher-capacity drives like the Vertex V100 series, the read-back signal from the head is 0.4 to 0.8 mV peak-to-peak, compared to 2 to 5 mV for a standard lower-capacity drive.

The read-channel interface on a standard 5.25-in. drive is comparatively simple. The data signal from the head is amplified, run through a low-pass filter, and differentiated. The output of the zero-crossing detector corresponds to the peaks of the read-back signal (Fig 5).

The methods used in the higher-capacity 5.25-in. drives are much more complex and incorporate a number of elements not found in lower-capacity units (Fig 6). In this process, signal pre-amplifiers are placed on the actuator arm very close to the read/write heads to improve the signal-to-noise ratio before the signal leaves the enclosed bubble.

Additionally, an automatic gain control compensates for variations in head-signal amplitude caused by variations in the distance between the flying heads and the media and by differences in the head amplifier's characteristics. The automatic gain control's amp output is differentiated, filtered, and digitized by the zero-crossing detector. The output, processed by the peak follower and the zero-crossing detector, becomes the read-data pulse. The time domain's signal-qualifier section eliminates data transitions caused by false zero crossings.

Such circuits are similar to those used in the large, IBM 3380-type drives and are required for reliable read-data operation.

MICROPROCESSOR CONTROL CIRCUIT

The internal microcontroller, an 8-bit single-chip processor, monitors all drive elements, and protects the read/write heads from damage in the event of failure. For example, the processor continually checks the spindle motor speed controller and deactivates the drive if the speed varies more than ±2%. It also monitors servo off-track and fault conditions on the read/write circuitry and keeps track of the head position to prevent the host system from seeking into guard-band zones. Further, it verifies and realigns track 0 position as necessary.

The most important function of the microprocessor is to monitor incoming step pulses, develop a velocity profile, and adjust the actuator to effect an overlapped seek function. This process permits the heads to begin their move to a track before all of the step pulses are sent, and thus dramatically improves system-access time.

The systems integrator who seeks to add the high capacity and faster access times of the next generation of 5.25-in. Winchester disks to his system should be aware of the design decisions made by disk-drive manufacturers. It isn't enough to promise the benefits of new technologies if their techniques are not proven. If risks must be taken to exploit a new technology, they should be tempered with other known methods to reduce the chance that units will fail in the field. ■

Roger Stromsta, *manager of engineering at Vertex, has 15 years of R&D and product development experience, with 11 of these years focused on rotating memory devices. He was formerly director of engineering at Shugart Associates.*

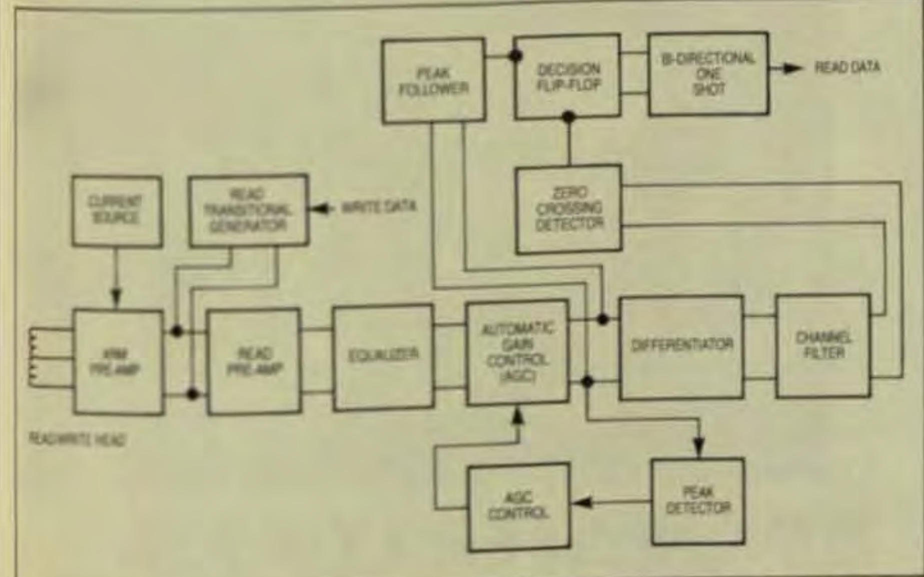

Fig 6 The read-channel interface on a high-capacity 5.25-in. drive is more complicated due to lower signal levels and the narrower noise margins that result from the denser packing of data bits. Signal pre-amplifiers on the actuator arm can strengthen the data signal before it leaves the enclosed bubble of the drive.

WE DON'T MEASURE WINCHESTER CAPACITY JUST IN MEGABYTES.

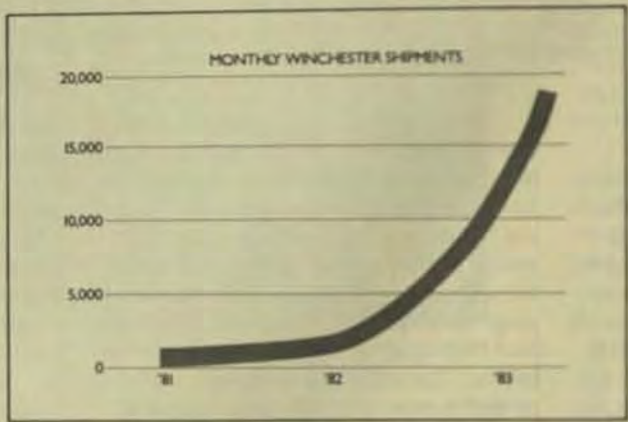

We measure it in terms of production capacity. Capacity for technological innovation. And capacity for building-in quality at the lowest possible cost.

We're the Tandon 5¼" Winchester Company. And we're at a disadvantage. Unlike our Tandon sister companies, we're not the world's production leader in our field. We're only number two. Though not by much. Like the Tandon 5¼" Floppy Company, we didn't start first, but we intend to finish first.

We build a wide range of drives. Open-loop and high performance closed-loop models. Right now, we have the capacity to build and ship 60,000 drives a month. That means that we have the capacity in place to be number one in 5¼" Winchesters. And we intend to be. By offering you what no other Winchester manufacturer can.

No other Winchester company makes so many of the parts that go into its drives.

No other Winchester company builds its own heads.

No other Winchester company manufactures more of its own plated media and its own PC board assemblies.

Our innovative products are backed with a 105% guarantee. Including our brand-new half-height 5¼" drives. With 6.4 and 12.8MB in half the space of traditional 5¼" Winchesters.

We're dedicated to building more drives with more capacity in less space for less cost.

Which means we're dedicated to being the world's leading producer of 5¼" Winchesters.

Tandon
THE MOST SUCCESSFUL DISK DRIVE COMPANIES YOU EVER HEARD OF.

TANDON 5¼" WINCHESTER COMPANY.

Tandon Corporation, 20320 Prairie, Chatsworth, CA 91311, (213) 993-6644, TWX: 910-494-1721, Telex: 194794. Regional Sales Offices: Boston (617) 938-1916 • New York (201) 449-7720 • Atlanta (404) 934-0620 • Chicago (312) 530-7401 • Dallas (214) 423-6260 • Irvine (714) 675-2928 • Sunnyvale (408) 745-6303 • Frankfurt, West Germany Telex 411547 • London, England (0734) 664-676 Telex: 848411. Distributors: Hall-Mark, Kierulff, Schweber.

Best assured for disk drive testing.

Whether you're a manufacturer of disk drives or a mainframe OEM, accuracy and performance are the keys to your success.

Precision test equipment that can measure heads, media, data channels and servo mechanisms in the design stage, or on the production line is essential.

In the past your test cycle may have consisted of a patch-work of equipment. Some dedicated, others adapted, usually about the only thing in common was inefficiency.

We've put together a team dedicated to disk drive test equipment. Over half are graduate engineers, and of those, fifty percent are disk drive specialists. That means expertise in the area where it's needed.

It also means you can have your choice of a unit capable of responding to your requirements from our existing line. Or an integrated custom system specially designed and built to your exact tolerances and accuracies.

That's probably why the world's leading disk drive related manufacturers are our best customers.

And, because we use a programmable design featuring a 50 MHz bandwidth, the system we build today will expand as you do.

Whatever your disk drive testing requirements, from media certification, to head testing, to window margin, speak first to the leader in disk drive testing technology.

And, rest assured.

Cambrian Systems, Inc.
The Leader in Disk Drive Testing Technology
31364 Via Colinas, Westlake Village, CA 91362, (213) 991-7026

Servo Design Extends Capacity and Increases Accuracy of Minifloppy

Widely accepted over the past decade, floppy disks will find further applications through the development of half-height, double-sided minifloppy drives.

by **Vladimir Langer**, Drivetec Inc.

System integrators face an increasingly difficult task in evaluating the many floppy-disk-drive offerings. The introduction of microfloppies (sub-5.25 in.) has only added to the confusion, while some low-end, relatively low-cost 5.25-in. Winchesters present new options for the replacement of floppy drives in some applications. Proper selection of these new devices requires an understanding of their limitations and the techniques that can be used to overcome them. This article discusses some of the advances in floppy-drive technology that will extend the applications of these products in the systems of the 80s.

Floppy disks, originally used for program loading in large computers, became the most widely used mass-storage device for small computers in the 1970s. Early 8-in. drives were followed by 5 ¼-in. minifloppies, and many design enhancements were made to boost capacities, so currently available drives can hold up to 1.6 Mbytes in double-sided 8-in. floppies, and 1 Mbyte in some 5 ¼-in. minifloppies. Today, sub-4-in. drives are further reducing the space requirements for convenient, removable mass storage.

Although floppies have been widely accepted, some limitations have caused problems, and some product introductions—especially the double-density, double-sided 8-in. drive—were plagued with technical problems. There are several areas of critical concern to system designers. These include:

- Diskette interchangeability.
- Data-storage capacity and reliability.
- Media longevity.
- System integration (compatibility).

While there are many ways of solving these problems, Drivetec has introduced a new minifloppy that presents some unique ways to solve these problems and manages to pack 3.3 Mbytes into a half-height, double-sided drive. Drives of this type of extended capacity will help floppy-disk technology retain its importance in small-system design for years to come.

DISKETTE COMPATIBILITY

One of the most important areas of concern is that of diskette interchangeability, the ability of one drive to read data from a diskette recorded on another drive. The mechanical tolerances involved in read/write head positioning, the media alignment, and hygroscopic effects are all small enough to allow accurate positioning of the read/write heads on the data track under varying environmental conditions.

The main problem is that the flexible substrate of the media itself expands and contracts as a result of heat and humidity variations (hygroscopic effects). A diskette recorded by a drive in one system environment and then moved to another undergoes slight changes over time. The concentricity of the data tracks will thus be somewhat altered. In conjunction with the different tolerances of the read/write head-positioning mechanism of the second drive, data track/head misalignment can result (Fig 1). The data track, which ideally

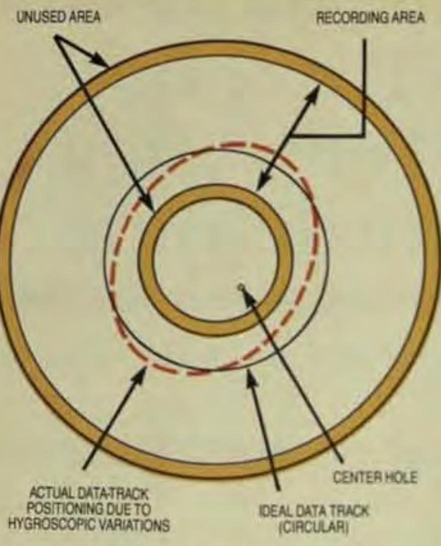

Fig 1 The ideal diskette has the data track in perfect concentric rings around the center hole. However, head and humidity effects on the polymer media and substrate, and mechanical stresses placed on the media during manufacture can cause the medium to expand and contract so that the actual data track varies.

should be at a constant radius from the center hole, weaves away from that ideal circle.

The second problem is misalignment of the media in the drive. This can occur after the media is inserted if the media-clamping device (a collette) doesn't accurately press the diskette into place. Any such off-centering results in off-track positioning of the read/write heads. Older designs brought the clamp down on the media at a slight angle. That procedure can cause some misclamping, so manufacturers often recommend a reinforced center hole. One way to overcome the problems of the past is to use a modified acme-threaded shaft to clamp the diskette vertically (Fig 2).

The shaft of the clamp has the same center point as the diskette spindle. But it's difficult to eliminate the effect of the accumulated tolerances of the several mechanical parts, because the spindle and the clamping mechanism are separately assembled modules that must be fitted together in the final stages of manufacture. Exact centering is accomplished by use of a wave washer that compensates for any off-centering and can be put into place at the final stage of manufacture. When the clamp is closed into place, the collette is precisely positioned and the wave washer locks it into place permanently.

The collette itself is tapered and supported internally by a flexible O ring. When the diskette is inserted and begins to rotate, the tapered clamp seizes the media and uses centrifugal force to help center it before clamping it into place. All of these improvements result in a drive in which no center-ring reinforcement is required in the media. The diskette is always precisely positioned.

The third factor in diskette interchangeability is the head-positioning or actuator technology. The two types in use now are band and lead-screw actuators. The band actuator uses a metal band and capstan connected to a stepper motor to move the head carriage. It's less expensive, but relatively inaccurate because the band can become stretched and because of tolerances in the capstan-to-motor connection. To support double-density drives requires precise and expensive machining of the band-actuator components.

The lead screw is slightly more expensive, but generally more accurate. These devices are usually found in double-density drives where precise positioning is more important. Mechanical backlash during positioning can be reduced if the traditional fixed guide rod that holds the head carriage in its linear travel path is complemented by a parallel spring-loaded rod (Fig 3). This assembly eliminates linear stress caused by a single rod bushing combination. The springs keep the carriage in place regardless of any mechanical vibration incurred in head positioning or caused by other system components.

Additionally, a third spring attached to the lead-screw nut can also

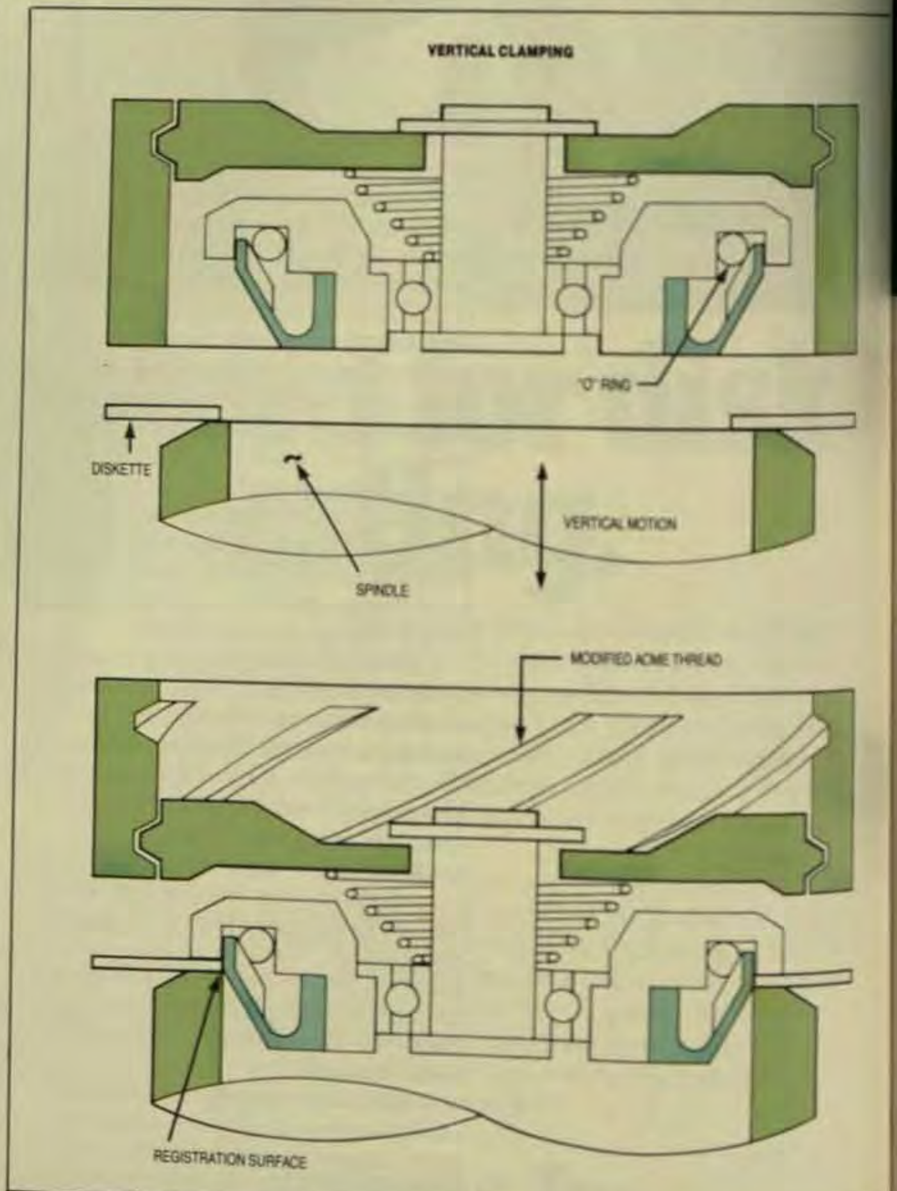

Fig 2 The media-clamping mechanism utilizes a precision-machined, power-threaded screw which is maintained at the exact centerline over the spindle shaft. As the clamp is lowered onto the media and the spindle (lower portion), the tapered collette grabs the media and centers it before clamping.

10 MBYTES IN A SLIM 1.63 INCHES

Keen engineering design packs 10 megabytes of formatted data into a slim 1.63-inch high 5¼-inch Winchester disk drive.

Plated media, closed-loop positioning, and guard bands all work together to put 12.7 megabytes of unformatted data into this slim-line package. The Microscience HH-612 combines proven technology with a unique design that utilizes dual microprocessors to maintain stepper motor closed-loop positioning, and precise spindle motor velocity.

Self diagnostics continuously monitor electronic performance, motor velocity, and head positioning accuracy.

The HH-612, designed to meet the needs of compact transportable computers and portable personal computers, is shock mounted and operates in any position. The drive also features sealed media with filtered air for reliable, continuous operation.

Storage compatibility is convenient and simplified with the industry standard ST506 interface.

Microscience second sources electronic and mechanical components to assure dependable delivery.

For more information on the HH-612, or to arrange immediate delivery of an OEM evaluation unit, write **Microscience International Corporation**, 575 East Middlefield Road, Mountain View, CA 94043 or call (415) 961-2212.

Microscience International Corporation

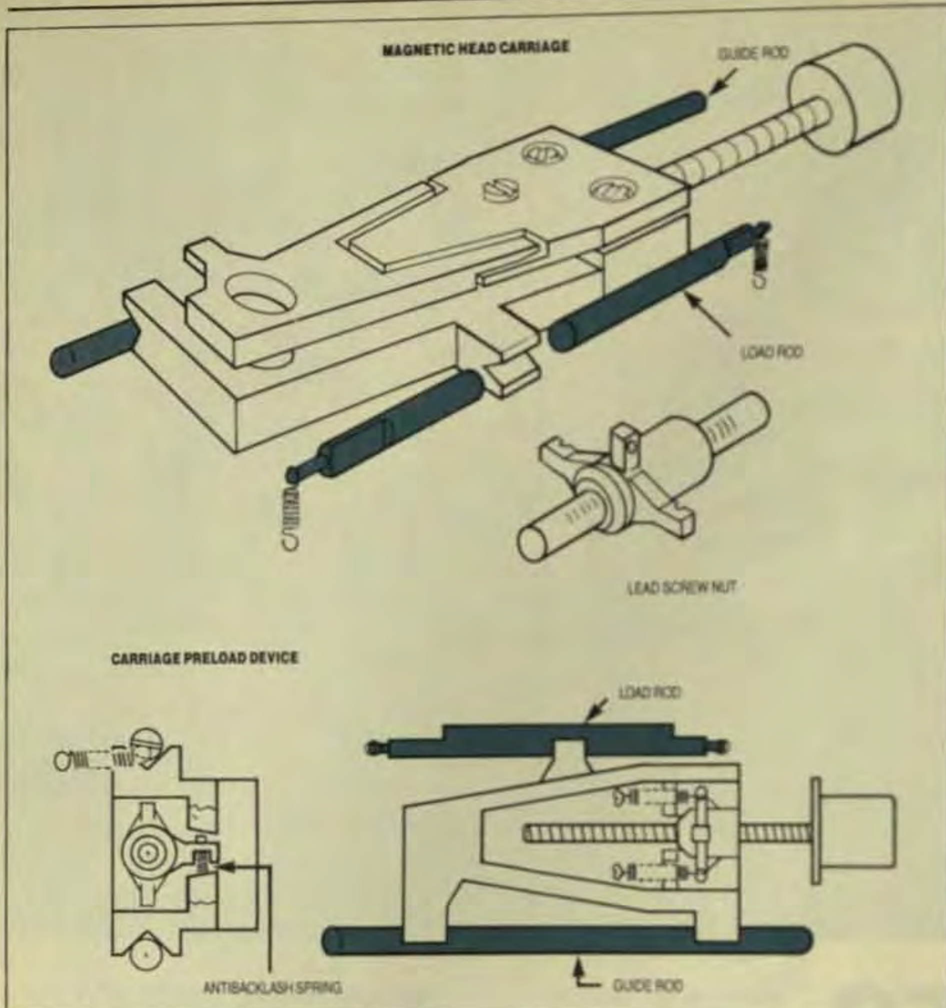

Fig 3 This design utilizes a fixed guide rod and a spring-loaded rod to self-align the head carriage. The springs hold the carriage in its correct position and absorb vibration caused by linear motion. The other springs eliminate backlash that results from linear and rotational motion.

sorb backlash created when the rotational direction of the screw changes. Together, these springs will correct for mechanical backlash caused by movement in any direction (rotational or linear) and self-align the head actuator. This design also eliminates any field-alignment requirements.

Another possible source of positioning errors is the small gap that inevitably occurs between the lead screw and nut. This problem can be eliminated by two springs that tilt the nut into the screw and hold it into place when it moves back and forth, so that no linear backlash occurs.

DATA-STORAGE CAPACITY AND RELIABILITY

Data interchangeability and storage capacity are directly linked because as track density is increased to achieve greater capacities, the component's error margin grows smaller. The critical breakdown in achievement of greater track densities is the open-loop head positioner. The solution is a closed-loop servo system that gives position information to the head-actuator subsystem.

This technique, widely used in high-capacity hard-disk systems, can be adapted for floppy drives as well. During manufacture, the servo information is written on the diskette by a device called a servo writer and is performed by the media manufacturer, or by large OEMs themselves. It's a straightforward process that requires no special media. Many minifloppy diskettes employ the 600-Oe 50 µin. magnetic-oxide coating that's become common among high-capacity-drive manufacturers. Eig[illegible] media manufacturers have comm[illegible]ted to supplying this media in vol[illegible]ume quantities. The price of the p[illegible]written servo-type diskette is ex[illegible]pected to be only slightly higher th[illegible] standard 96-tpi minifloppy diske[illegible]

A burst of servo information is placed in each sector (Fig 4). Micropro[illegible]cessor-based circuits in the control electronics continuously monitor [illegible] servo information and direct the placement of the read/write head. This closed-loop servo technology i[illegible]proves the accuracy of head-to-tra[illegible] positioning so that the drive can p[illegible]tion the heads to within 100 µin. of the exact centerline of the data tra[illegible]

With this accuracy the track in[illegible]sity can be increased from 96 to 1[illegible] tpi without a sacrifice of data relia[illegible]ity. In fact, data integrity is increa[illegible] because the heads find the track c[illegible]terline and follow it—regardless of the hygroscopically induced chan[illegible] in the media.

The drive employs two stepper motors, one coarse and one fine. D[illegible]ing a track-seek operation, the coa[illegible] stepper motor positions the head [illegible] the data track. Then the fine stepper motor responds to the servo-burst in[illegible]formation and steps the heads in 200-µin. increments until the exa[illegible] data-track centerline is reached. Th[illegible] head follows the track, remaining within 100 µin. of the centerline. [illegible]

Since the debacle of the first [illegible]ble-sided, double-density 8-in. flo[illegible] drives, one important concern in [illegible] new drive has been the head-to-media compliance and the resul[illegible] effect on media life. The head recording gaps must be in contact with the media to pick up a clear low-noise signal. But the heads m[illegible] not rub the media and harm the re[illegible]cording surface.

The most widely adopted tech[illegible]nique is to have one head rigidly placed and the second one hing[illegible] permit flexing and insertion of the diskette. Spherical "gumball" hea[illegible] (Fig 6) ease the passage of the med[illegible]

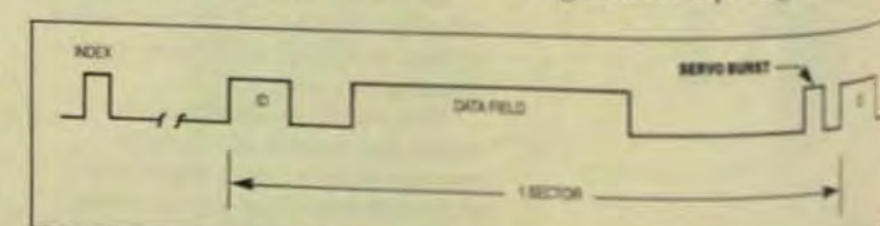

Fig 4 Each sector on the diskette contains embedded servo information, prewritten at the [illegible] each sector. This servo information includes a servo ID and a series of logical ONEs used by [illegible] servo-detect circuitry to obtain the track position.

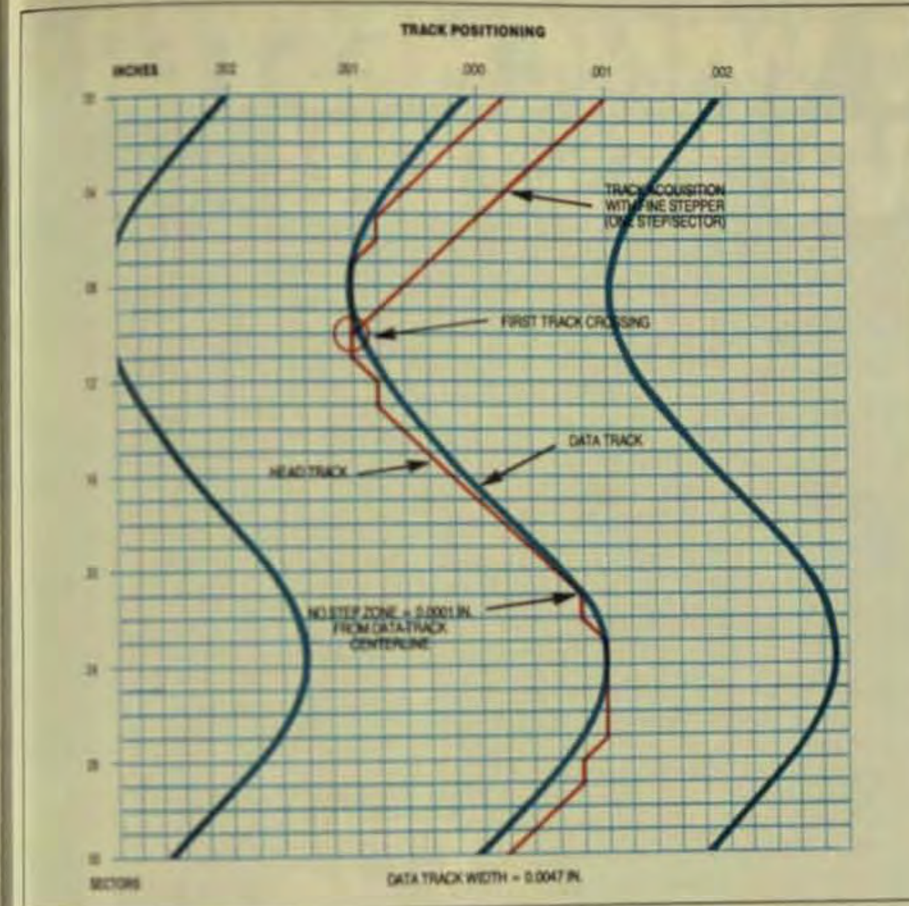

Fig 5 If the data track isn't at a constant radius from the centerline, it appears to wander or weave (shown greatly exaggerated). With servo-positioning feedback, the data-track centerline can be located and followed within 100 µin.

between them and have no sharp angles to scratch the media surface, so wear is reduced.

SYSTEMS INTEGRATION

Drive-integration problems involve three areas of concern:

- What controller electronics are needed to interface the drive?
- Will it be compatible with existing systems and existing databases already recorded on conventional diskettes?
- Will the physical package be suitable for use with systems already using the defined footprint of 5.25-in. drives?

For example, Drivetec designed its system to be Shugart-compatible, which means it's compatible with the standard SA4XX type of drive controller, and conforms to the same interface conventions as any other Shugart-compatible drive. Thus, current drive controllers can be used to integrate the drive into existing computer systems, and it can be used as a plug-in replacement for lower-capacity drives.

To aid systems integrators in updating their existing customer base, the drive's microelectronics allow it to read standard 48-and 96-tpi minifloppy diskettes, as well as its own servo-written 192-tpi diskettes, without system modifications. This responsiveness allows existing software and data to be transferred from standard systems to a higher capacity drive without interfacing problems or procedures.

Also, the drive is exactly half the height of a standard 5.25-in. minifloppy, 1.625 in. high, 5.75 in. wide, and 8.485 in. long. A single 1-Mbyte 5.25-in. full-height floppy can be replaced by two 6.6-Mbyte minifloppies without package redesign.

Floppy disks are the mass-storage device of choice for small-systems designers. The use of closed-loop servo techniques is the most important element in extending the data-capacity—and the market life—of these floppies. In all disk drives, the trend is toward smaller physical size, greater capacity, and removeable media. With its 3.3-Mbyte data storage and ultraprecise positioning accuracy, the half-height floppy represents a significant new development that will help secure the preeminence of minifloppies as the mass-storage device of choice for small systems. ■

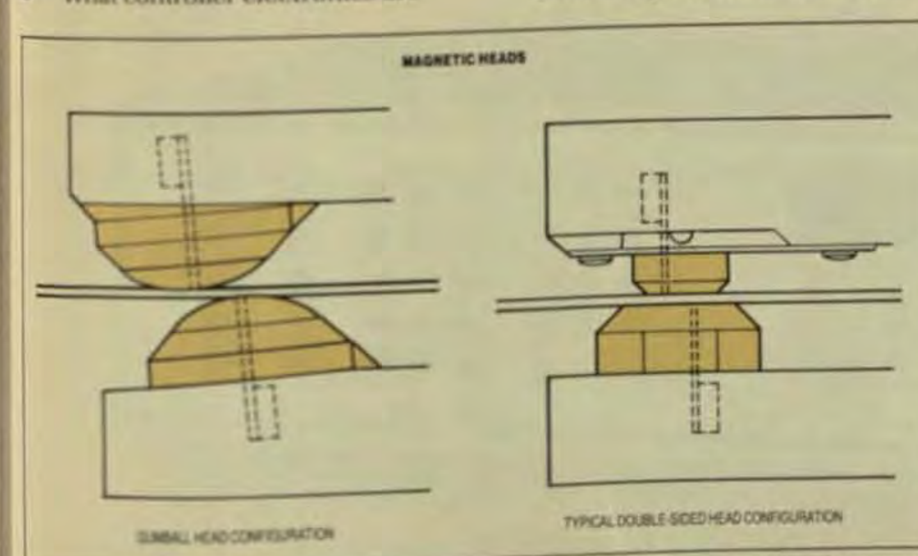

Fig 6 The rounded shape of the "gumball" head promotes easy passage of the media with no media wear.

Vladimir Langer, *product manager for Drivetec's 3.33-Mbyte SuperMinifloppy, previously worked for Shugart Associates as international marketing manager and development engineer.*

FASTEST WINCHESTER DRIVE TESTER!

DUAL MICROPROCESSOR BASED MWX-1000 PROVIDES FAST PRODUCT THROUGHPUT.

Compact portable and rack mounted versions.

Key up the drive configuration of any type Winchester drive having ST506 or SA1000 interface and the new MWX-1000 Winchester Analyzer is ready for work. Its dual microprocessors make it the fastest and most powerful QA, production or field service Winchester ever offered ...*the MWX-1000 is from Wilson, the 10-year leader in test exercisers.*

Imagine full interactive testing with operator prompting, and field programmable analysis modes to expose even stubborn intermittent errors. The MWX-1000 has them. And it all comes packaged in a carrying case that measures only 16½"x6"x20". The system version takes only 5¼" of rack space.

Big-system test functions include full basic drive test, format and verify test, drive performance test, and full drive margin test. These apply to 5¼" or 8" drives, multiple-head drives and drives with 4.35 MHz or 5 MHz data transfer rate. The basic unit supports up to four drives (16-drive system available).

The *go anywhere* MWX-1000 is your best partner to quickly analyze Winchester drives.

More Max Series Peripheral Analyzers

Wilson's giant step into the future with microprocessor logic now offers you a growing line of MAX Series Analyzers.

MFX-1000 Floppy Drive supports all standard floppy drives. It analyzes all basic functions including formatting and verification, drive performance margins, and head alignment without a

MPX-1000 Printer Analyzer and exercises all currently available printers. It checks or 96 character patterns, patterns, format printing, vertical and horizontal and set of characters-per- and lines-per-page.

SEND FOR COMPLETE INFORMATION

WILSON Laboratories, Inc.
2237 N. BATAVIA STREET, ORANGE, CALIF.
Telephone (714) 998-1980 · TELEX

Direct replacement for Cambrian, ACT, ADC, and Qubex test stations.

Analyzers Provide Range Of OEM Tests for Error Conditions in Floppies

As the need for floppy disks increases, sophisticated microprocessor-driven analyzers are finding wider OEM use to detect the many factors that can cause data errors.

by **Chuck Ouellette**, Brikon Inc.

Systems integrators planning to test floppy-disk drives before integrating them into desktop personal computers, word processors, or small-business systems now have available to them a wide range of sophisticated test equipment that can check the performance levels of several drives at a time. Some of this test equipment will verify only the performance of basic drive functions; other, more sophisticated equipment, however, will not only check these basic drive functions, but perform more critical tests as well.

This article will consider three types of equipment that integrators can use to test floppy-disk drives, and will discuss higher-level tests—window-margin analyses and asymmetry checks—that can be performed by the use of this equipment.

TYPES OF HARDWARE AVAILABLE

The three types of floppy-disk-drive test equipment available are:

- Exercisers.
- Testers.
- Analyzers.

Exercisers supply the logic to select a specific drive and to position the read/write heads. Some are equipped to write data patterns. These low-end devices are relatively inexpensive (about $750) and normally don't provide power to the drive, read data, or allow for any measurements. Typically, they require the assistance of an oscilloscope. All they really do is, as the name suggests, exercise the device.

Testers are more sophisticated and expensive (approximately $2000) than exercisers, and may or may not provide dc power for the drives under test. These systems can perform all the tests associated with an exerciser and, in addition, handle read/write functions and measure specific drive parameters such as the rotational speed of the spindle motor. Some testers can trap errors, but most can't identify where the error occurred.

Analyzers are microprocessor-driven test systems that typically are more expensive than either an exerciser or a tester but which can perform all the checks associated with each—and then some. These self-contained test systems have a significant amount of auto-test capability, and can easily measure the rotational speed of the drive, the positioning repeatability of the head actuator, and the read/write channels.

They generate worst-case data patterns to trap read errors, and can indicate where these errors occurred. Analyzers also can measure window margin accurately—an excellent indication of a drive's ability to read and write data—and perform asymmetry tests.

CAUSES OF DATA ERROR

Data generated from the magnetic impulses written on a floppy disk comes from the read/write heads in the form of an analog signal. In order for this analog data to be understood by the computer, it must pass through a phase-locked oscillator that references this data to time.

Each bit written or read to disk has a time slot (bit cell) allocated by the disk controller. In the case of a read, if there's no signal in the specific time slot (window) established by the controller, a logical ZERO is generated. If a signal is present in the time slot, a logical ONE is indicated.

In an ideal situation, the signal or lack of signal will always occur in the middle of the window (see Fig). But that is rarely the case. Instead, data appears to shift from one side of the window to the other; if it shifts too far, the controller may not be able to capture it. The result is a data error.

This bit shift, the analog signal's shifting from one side of a time window to another, can be caused by a number of factors, many of which can be detected by analyzers. In some cases, the window itself moves because of oscillator inaccuracies when the window is generated. The most common causes, however, relate to speed variations in the drive spindle motor, and to media centering.

Although the drive's rotational speed should remain constant, it can be changed by voltage fluctuations, rough spindle bearings, imperfections in the media jacket, friction of the media, or rough head loadings.

Depending on the combination of variations that come into play and the design characteristics of the drive itself, instantaneous speed changes can range from less than 0.5 to 2.0% of the theoretical rotational speed. This translates into instantaneous speed changes anywhere from 1.0 to 4.0 ms.

The impact of any of these conditions on drive latency (and hence, on window margins) is easy to see. The typical 5.25-in. floppy-disk drive has an average rotational delay of 100 ms. The average window for a bit in Shugart Associates' SA450 floppy-disk drive operating at a density of 5876 bpi is 0.017 ms.

Media centering is another condition that can affect window margins. If the media is not centered within acceptable limits, the disk will move in an elliptical path rather than a circular one, causing data to be alternately closer and further away and to shift in relation to the established window media used in floppy-disk drives. Since these diskettes don't expand and contract linearly in all directions with temperature and humidity, if the data is written at one temperature and read at another, certain areas of the track may move further away from the head than others. The result is bit shift.

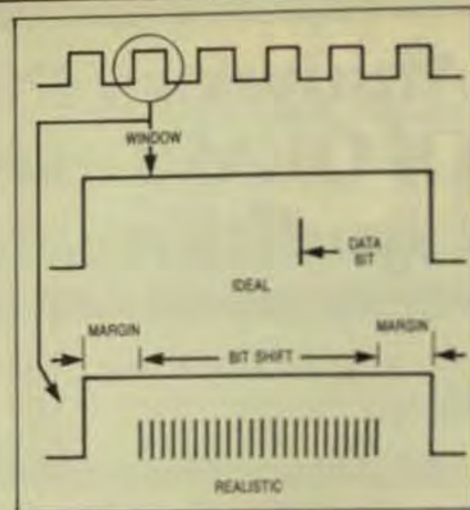

Fig While a signal should, ideally, always occur in the middle of the window (time slot) provided by the controller, many factors can cause the data to shift from one side to the other.

DATA-ERROR MEASUREMENT

An analyzer can be used in several ways to measure the speed shift and window margin. The greatest accuracy, resolution, and consistency is provided by the generation of a very accurate window. To do this, an extremely accurate crystal oscillator is used, plus linear ramp generator and instrument-grade components that generate a 1000 ns window that's accurate to within ± 2 ns. Measurements commence at the innermost track, where bit shifting is most prevalent, and a worst-case data pattern is used —such as $6DB6_{16}$ (0110 1101 1011 0110).

The analyzer writes this pattern on the diskette; then measurements can begin. Samples are taken for bits that shift to the left of center. Each shifted bit is registered, with the bit shifting the most stored for further analysis. The process is then performed again for bits shifting to the right of center. Once this process is complete, the negative and positive shifts are added together and their total subtracted from the window. The result: the window margin of the drive.

CAUSES OF ASYMMETRY

Bit shift can also result from asymmetry—imbalances in the drive electronics and read/write heads. An analyzer should be able to make this kind of measurement if it is to determine fully the performance level of a specific drive.

How an analyzer performs these tests is best seen by a closer look at the problem itself. The read electronics are supposed to operate as a balanced line to amplify equally both the positive and negative signals coming from the read/write heads. But if this line isn't balanced, one will be amplified more than the other and bit shift will result. Asymmetry also affects the write operation: If head cores aren't balanced, the magnetic flux they generate will be stronger in one direction than the other.

A less obvious contributor to asymmetry is the media itself. When data is written on a diskette, the particles in the media's magnetic layer will orient themselves in the direction of the flux transition. If the same data is written over and over again, these particles will ultimately end up in a perfect orientation.

If opposite data is then written, these particles will attempt to align themselves in the opposite direction but will not be able to do so completely. As a result, when data is read back, the partially reoriented particles can create an effective bit-shift condition.

To check for this situation and measure the quality of the drive's read channels, analyzers can be used to measure the time from the mean early bit to the mean late bit to produce an asymmetry value in nanoseconds. This test is normally performed on the outermost track by use of a low-frequency pattern so that jitter components other than asymmetry are reduced.

Self-contained microprocessor-driven analyzers are rapidly coming into widespread use as the market broadens for floppy-disk drives. The need for analyzers for preproduction and production test, and for incoming/outgoing quality-assurance applications is already well established. As their power is recognized, their use in other areas of testing is becoming more widespread as well—including field maintenance and depot maintenance, error and longevity testing, as sales demonstrators, and as required equipment in the development of new products. ■

Chuck Ouellette, *vice president of marketing at Brikon Inc., earned his MBA from Pepperdine Univ.*

LOOKING FOR 5¼" FLOPPY DISK DRIVES?

LOOK NO FURTHER.

THE REMEX® FAMILY OF COMPATIBLE INNOVATIONS

RFD480	RFD960	RFD485	RFD965	RFD1600	RFD3200
500 KB	1 MB	500 KB	1 MB	1.6 MB	3.2 MB
48 TPI	96 TPI	48 TPI	96 TPI	170 TPI	170 TPI
Double Sided	Double Sided	Double Sided	Double Sided	Double Sided	Double Sided
2/3, Full Size	2/3, Full Size	Half High	Half High	Full Size	Full Size

Regarded for over ten years as a leader in flexible disk drive technology, Remex offers a wide selection of 5¼" floppy drives to fit your needs. Our innovative designs and high performance standards give you the necessary compatibility for increased efficiency. And regardless of their size or capacity, all Remex drives utilize the same proven mechanics and quality levels — at an affordable price.

Call or write us today and see how there is a Remex innovation compatible to your needs.

REMEX®

Ex-Cell-O Corporation
REMEX DIVISION

2991 White Star
Anaheim, CA 92806

(714) 630-7020
TLX 4722112

To 40 megabytes.
Disctron 5 1/4-inch Winchesters.

Disctron offers you more capacity, higher performance, and easier integration in 5¼-inch Winchester. With capacities to 40 megabytes. With average access times as low as 35 ms. And with special features—electronically dampened actuator, buffered step mode, electrical spindle brake, integral microprocessor control and built-in AGC.

Easy integration. Slip these 5¼-inch drives into an existing minifloppy cavity. Use standard minifloppy voltages. Use standard ST506/ST412 interfacing. These drives are ready to go.

Plug in the Disctron difference in 5¼-inch Winchester.

To 160 megabytes.
Disctron 8-inch Winchesters.

Disctron has it all. Fixed and removable Winchesters with capacities to 160 megabytes. Thin film media. Internal diagnostics. Self-contained shock mounting. Internal carriage lock. SMD and ANSI-standard interfacing. DC power. ☐ Plug in 10.6 megabyte on an ANSI-standard, transportable cartridge. Linear voice coil, embedded servo and microprocessor control provide precise head positioning and reliable data access.

Plug in the Disctron difference in 8-inch Winchester.

Only Disctron offers a complete family of 5¼-inch and 8-inch Winchesters in fixed and removable, proprietary head technology and manufacturing, high volume manufacturing, and the solid financial foundations of CCT's more than $75 million in assets.

More than megabytes. Discover the Disctron difference. For more information, circle our readers' service number. Or call us direct at (408) 946-6692.

the disctron differen

Disctron, Inc.
A Subsidiary of CCT

For information regarding U.S. and world wide distributors, contact Disctron, Inc., 1701 McCarthy Boulevard, Milpitas, CA 9
Telephone: (408) 946-6692 TWX: 9103382068 Disctron MPTS

New Amdek 3" micro-floppydisk drive system!

AMDISK-III . . . OEM's choice:

- Works with standard 5¼" controllers
- Up to 1 megabyte of unformatted double density storage
- Ideal for industrial data collection/storage applications
- Microfloppy's hard plastic diskette offers increased durability vs. 5¼" floppies
- Automatic shutter protects media
- Small form factor for easy storage

Specifications

	Unit	Double Density
Capacity		
Unformatted Per Surface	Bytes	250K
Media		
Record Surfaces	2	
Tracks	80	
Recording		
Max Recording Density	Bpi	8946
Track Density	Tpi	100
Transfer Rate	bits/sec	250K
Access Time		
Average Access Time	msec	55
Track to Track	msec	3
Setting Time	msec	15
Average Latency Time	msec	100
Motor Start Time	sec	0.7 (min)
Disk Speed	rpm	300
Reliability		
Error Rates		
Soft Error		10-[illegible]
Hard Error		10-[illegible]
Seek Error		10-[illegible]
Media	3 inch Cartridge	
Drive Interface	Plug Compatible with 5.25 inch FDD	

External Interface

Connector: 34 pin (Shugart)

Pin No.	Signal	Pin No.	Signal
2	Unused	20	Step
4	In use (option)	22	Write data
6	Drive select 3	24	Write gate
8	Index	26	Track 00
10	Drive select 0	28	Write protect
12	Drive select 1	30	Read data
14	Drive select 2	32	Unused
16	Motor on	34	Ready
18	Direction	1-33	Ground

NOTE: Single head per drive

Evaluation samples $480

Includes two-drive Amdisk unit with built-in power supply, 4 diskettes and application literature. Call (312) 364-1180

The AMDISK-III Micro-floppydisk System is an engineering breakthrough in disk size, storage capacity, media protection and user convenience. Designed to serve many applications, the Amdek system is ruggedly constructed to provide years of trouble-free operation. Warranty is 90 days (parts & labor).

Put the new AMDISK-III to test . . . its recording format, data transfer rate and disk rotation speed are compatible with 5¼" floppydisk drives. Call, or write for evaluation samples at only $480.00 . . . or circle the reader service number for full technical details.

2201 Lively Blvd. • Elk Grove Village, IL 60007
(312) 364-1180 TLX: 25-4786
REGIONAL OFFICES: Calif. (714) 662-3949 • Texas (817) 498-2334

AMDEK CORP

Amdek . . . your guide to innovative computing!

3-In. CFD Format Vies For Acceptance as Microfloppy Standard

The 3-in. compact floppy diskette has a hard protective envelope and minifloppy compatibility. It offers reliability, high capacity, and low manufacturing costs.

by **Yasutaka Nakajima**, Matsushita Electric Industrial Co. Ltd.

Advances in magnetic recording technology, electronics, mechanics, and miniaturization have launched floppy-disk development into a new phase—microfloppies. Mass production has begun on the 3-in. compact floppy disk (CFD), and the 3-in. format has already won the support of 17 companies worldwide. Three other microfloppy formats—3.25-in., 3.5-in., and 4.0-in.—have recently been introduced and submitted to the ANSI X3B8 committee for standardization.

Since the introduction of the 8-in. floppy in 1972 (the first phase of floppy-disk formats), the trend has been toward decreasing size and increasing storage capacity. The development and subsequent standardization of the 5.25-in. format in 1976 was the second phase of disk development. Now, three Japanese firms—Matsushita, Hitachi Ltd., and Hitachi Maxwell Ltd.—are preparing to enter into the third phase with the 3-in. CFD.

Of the microfloppy formats competing for market acceptance, the 3 and 3.5 in. seem best suited for standardization. The media for these formats is protected with a hard case (envelope) and a shutter mechanism that protects the recording media from dust and fingerprints. If fingerprints on the media are moist they can be removed without damage, but once they have hardened or dried they can't be removed, and recorded information is lost.

This envelope also protects the media from rough handling, nipping, bending, and excessive pen pressure, and stabilizes the shutter mechanism. Since the 3.25- and 4-in. microfloppy formats don't provide a shutter, it's difficult to support their standardization.

In order to create a new floppy disk that would increase efficiency, maintain price-competitiveness, and win standardization approval, developers of the 3-in. CFD were guided by four major goals:

- Compatibility with existing 5.25-in. minifloppies.
- Handling ease.
- High-volume production for both diskette and drive.
- An expandable and nonredundant format.

Compatibility with the existing 5.25-in. format was essential in the development process. In order to gain wide acceptance, the CFD had to offer various applications, so the principal specifications of the new disk—such as rotation speed, data transfer rate, and recording capacity—were designed to be the same as those of the 5.25-in. minifloppy. Not only is the 3-in. CFD compatible with the minifloppy, existing measuring equipment and tools can be used to produce the new disk, thus cutting manufacturing costs.

Since the 3-in. CFD is a consumer product, it's crucial to be able to handle it easily and even carry it in your pocket. For this reason, it's provided with a rigid case (envelope) to contain the flexible magnetic diskette, which can be read or written on through windows on both sides of the envelope (Fig 1). This casing consists of an upper and lower plate 1.5 mm thick in those areas surrounding the diskette that the envelope protects. The case is easily manufac-

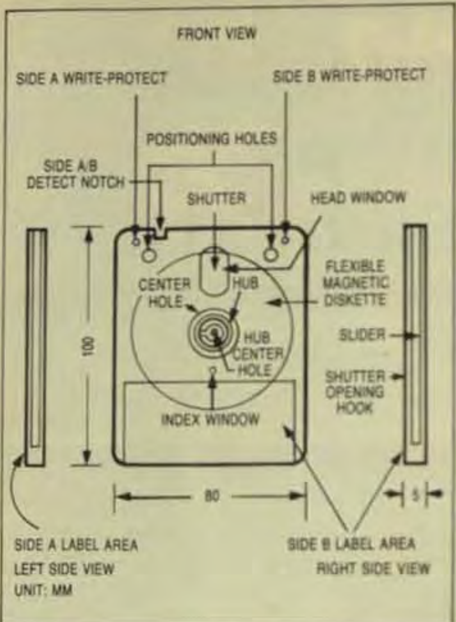

Fig 1 The 3-in. CFD diskette has a rigid envelope that contains windows (with shutters) on both sides to allow for reading/writing when the diskette is inserted in the drive.

tured by an ultrasonic bonding method and is inherently stronger than the thin case used in the 3.5-in. format.

The cartridge thickness is a key element of the disk, since thin cartridges are susceptible to the same defects as conventional jackets—particularly nipping and bending. As an additional labeling feature, the back face of the 5-mm cartridge can be labeled with a hard pen.

The 3-in. CFD also features a unique internal automatic shutter to protect the recording media. Since it's mounted within the cartridge, this shutter can't be opened easily when the diskette is removed from a drive. The shutter mechanism consists of a pair of shutter planes and a slider to open and close the windows.

The envelope is also equipped with a hardened plastic center hub that positions the diskette precisely and accurately, thus avoiding disk-chucking errors. The flange portions of the hub are located between the upper and lower plates so that the media never contacts the case.

Since these three components—shutter mechanism, hard shell (envelope), and recording media—are simple in shape, the disk cartridge can be produced inexpensively by high-volume manufacturing techniques.

An important design factor is that the case should be flat, not warped, and the liner within the case must make absolute contact with the media. Either you can use a plastic spring lever to press the liner against the media, or you can use a thicker compressive liner that is elastic enough to make contact. In either case, the liner is attached to the inner surface of the envelope, which must be very flat so that it can maintain the distance between the inner surface of the envelope and the media.

When inserted into the drive, the case is supported by three positioning pins (Fig 2) and the media sits on a spindle table. The spindle table has a drive pin to rotate the hub, which turns at 300 rpm. These pins and the spindle table determine the positioning relationship between the diskette and its envelope.

Note that relative to a conventional flexible jacket, the envelope results in increased cantilevering of the magnetic head from its suspension. However, this problem has been resolved with a new head-suspension design and production of a double-sided drive.

THE AUTOMATIC SHUTTER MECHANISM

The 3-in. microfloppy is provided with an internal automatic shutter mechanism that consists of:

- A pair of shutter plates, each 0.2 mm thick, mounted on the inner surface of the envelope. These revolve around the center hole.
- A moveable slider that's supported by the envelope and attached to the shutter plates. The slider has a flexible protrusion extending outside the case.
- A spring for forcing the slider to close the shutters over the head window.

The force of the spring causes the shutter plates to cover the head window, but when external force moves the protrusion of the slider along the channel in the direction against the force of the spring, the shutter plates are moved into the open position and the window is exposed.

One feature of this design is that the shutter opening mechanism is very simple. When the envelope is inserted, one fixed pin within the drive engages the slider protrusion, which forces the spring to open the shutter and expose the window to the head for reading/writing. The shutter plates and slider are installed within the envelope and the slider protrusion is positioned in the U-shaped channel, thus preventing damage to the automatic shutter mechanism when the case is not loaded in the drive. Erroneous insertion of the case into the drive can be prevented by a design in which one end of the U-shaped channel is open and the other end closed.

The 3.5-in. format originally had a manual shutter, but an automatic shutter mechanism was recently incorporated. However, the shutter mechanism is attached outside the envelope, which increases the potential for accidental breakage. Also, the direction of shutter movement is at right angles to the direction of insertion into the drive, which may make the drive more complex.

It might be thought that the shutter mechanism increases the cost of the disk, but since the shutter plates are easily manufactured by a punch press and the slider is made in high volume by molding, there's no significant cost increase.

THE HUB OF THE 3-IN. CFD

In the 3-in. CFD, a hardened plastic hub is attached at the center of the

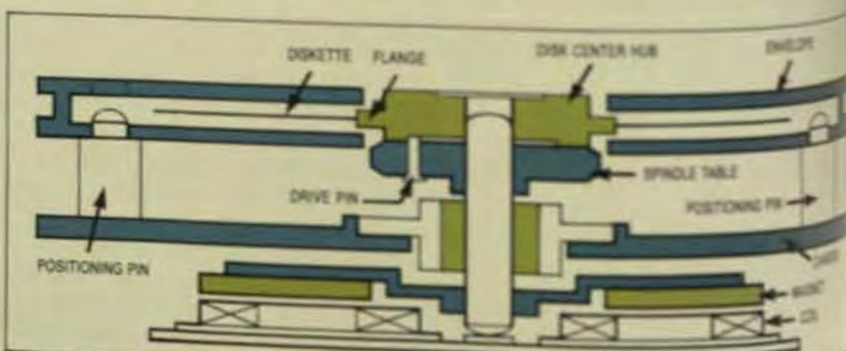

Fig 2 The envelope and diskette are supported within the drive as shown.

TECSTOR DELIVERS!

U.S. manufactured Winchester Series 3 fixed disk drives, offering, 83 to 332 MBYTES, CDC/SMD compatibility and rack mounting (25.2 inches deep). Tecstor delivers evaluation units in 15 days ARO.

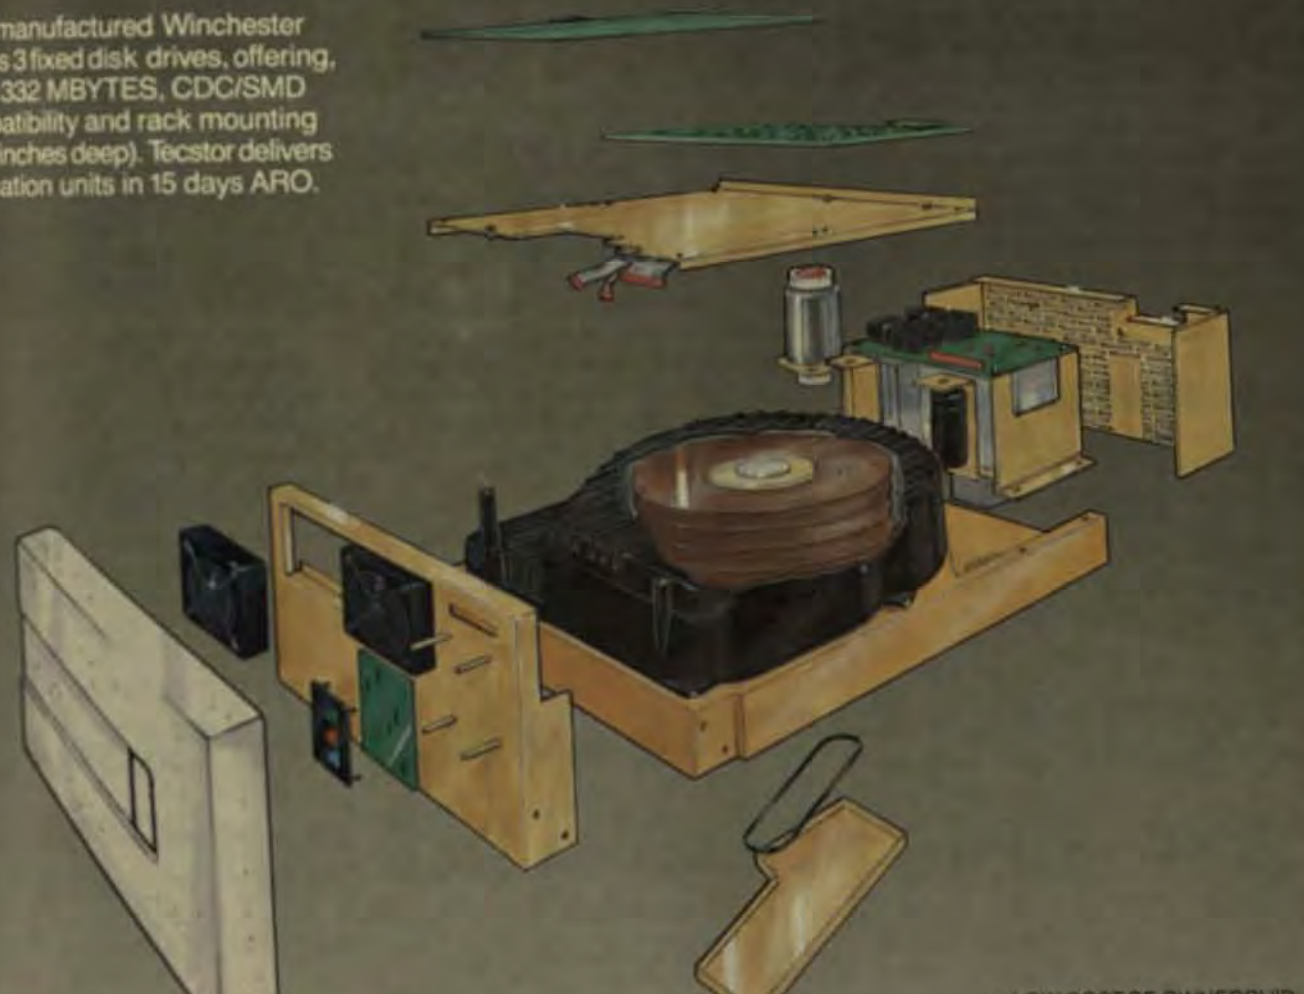

- CDC/SMD COMPATIBILITY
  - SMD Interface
  - 10 MHZ Data Transfer Rate
  - 20,160 Bytes Per Track
  - 823 Cylinders
  - 5, 6, 10, 12, 19, or 20 Heads
- PERFORMANCE
  - 29 MSEC AVG Seek
  - 8.3 MSEC Latency
- MODULARITY
  - 6 Series 3 Models
  - All Parts Common (except HDA)
- STANDARDS' CONFORMANCE
  - 19" RETMA Rack Mounting
    - 10.5" H x 16.5" W x 25.2" D
  - Safety
    - UL 437, IEC 435, CSA 22.2
  - Acoustic Noise
    - NC 60
- LOW COST-OF-OWNERSHIP
  - Common spares for all capacities.
  - 15 minute disassembly/assembly with one screwdriver.
  - No PM.
  - 10,000 hour MTBF.
  - Automatic spindle and head lock.
  - Power-on confidence check.
  - CE initiated diagnostics.
  - Automatic head retract to landing areas upon power loss.
- Less than $15/MBYTE (332 MBYTE, OEM quantities).

TECSTOR

Northern California (408) 732-2143

Southern California (714) 842-0077

Dallas (214) 644-1772

Boston Area (603) 889-5561

Metropolitan New York (201) 850-8400

velope. This hub consists of two parts, a substantial hub and a ring hub, that clamp the platter firmly between them. The centering accuracy of the diskette depends on the dimensional accuracy of the shaft it rides. Since the diameter of the shaft is 4 mm and its tolerance is 2 μm, the platter centering is quite accurate.

The hub is made of plastics that have a high tolerance to wear, and low creep. Centering accuracy is kept within 3 μm after 5000 times chucking at 40°C.

The diskette is centered by the insertion of the hub into the shaft. A recessed portion of the hub engages the drive pin mounted on the spindle table so that the hub and the diskette are rotated simultaneously by the spindle.

CHARACTERISTICS OF THE DISKETTE

The diskette is 72 mm in diameter and 77 μ thick. The recording area is defined by r_1 = 16 mm min., and r_2 = 35 mm max. on both sides.

The basic specifications of the 3-in. microfloppy system are shown in Table 1. The material of the diskette has a light transmittance of less than 2%. The torque required to rotate the diskette is 0.004 Newtons/meter or less. The recorded track width on the diskette surface is 0.150 ± 0.015 mm and the area between the tracks is erased. The nominal radius of the centerline of track 00 for single-sided recording is 32.5 mm for each side. For double-sided recording, the nominal radius is 32.5 mm for side A and 30.468 mm for side B.

TABLE 1 — SPECIFICATIONS OF 3-IN. COMPACT FLOPPY DISK

	COMPACT FLOPPY	CONVENTIONAL MINIFLOPPY
UNFORMATTED CAPACITY • SURFACE • DISK • TRACK	250 (MFM) 500 (MFM) 6.25 (MFM)	250 (MFM) 500 (MFM) 6.25 (MFM)
DATA TRANSFER RATE • KBITS/S	125 (FM) 250 (MFM)	125 (FM) 250 (MFM)
ROTATION SPEED (RPM)	300	300
RECORDING MODES	FM/MPM	FM/MPM
TRACKS/SURFACE	40	40
TRACK DENSITY TPS	100	48
RECORDING DENSITY BPI	8946	5536
DISK DIAMETER	72	130
DISK SIZE	80x100x5	133x133x1.6
DRIVE SIZE	90x40x150	150x86x228

The improvement of the media (diskette surface) is one of the most significant aspects to consider in achieving high density recording. A co-doped high-coercivity material has been selected for the 3-in. CFD to achieve a high linear-recording density. The magnetic characteristics of this material are listed in Table 2.

TABLE 2 — MAGNETIC CHARACTERISTICS OF HIGH-DENSITY MEDIA

	COMPACT FLOPPY	CONVENTIONAL MINIFLOPPY
COERCIVE FORCE — HC(OE)	634	270
REMANENCE — ΦR (MAXWELL)	0.080	0.14
RETENTIVITY — BR (GAUSS)	750	600-700
SQUARENESS — BR/BM	0.58	0.55
ORIENTATION RATIO	1.05	1.04
COATING THICKNESS (μM)	1.1	2.5

Since the coercive force of this media is stronger than conventional media, coating thickness can be reduced to 1.1 μ and the linear-recording density is still increased. The output characteristics of the new media are shown in Fig 3.

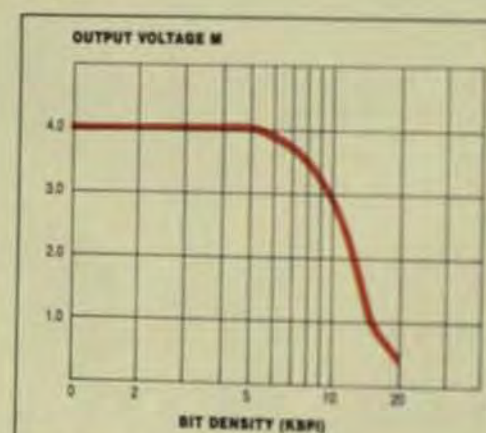

Fig 3 Output characteristics for the 3-in. CFD drive are shown in the above chart.

COMPATIBILITY AND DEVELOPMENT

The question of compatibility with existing systems has been of great concern from the start. In order to be widely accepted, the specifications of the new small-disk system (Table 3) must be flexible enough to adapt to many applications and be a good candidate for standardization. Therefore, the specifications for the 3-in. format (rotation speed, data transfer rate, recording capacity) are designed to be the same as those of the 5.25-in. minifloppy.

TABLE 3 — DRIVE SPECIFICATION

	SINGLE DENSITY	[illegible]
ENCODING METHOD	FM	[illegible]
UNFORMATTED CAPACITY • SINGLE SIDE • BOTH SIDES • PER TRACK	125 KBYTE 250 KBYTE 3125 BYTE	[illegible]
MAXIMUM RECORDING CAPACITY	4473 BPI	[illegible]
DATA-TRANSFER RATE	125K BPS	[illegible]
RECORDING SURFACE	2	
TRACK DENSITY	100 TPI	
TRACK NUMBER	40	
ROTATIONAL SPEED	300 RPM	
ACCESS TIME • TRACK TO TRACK • SETTLING TIME • AVERAGE • MOTOR START	3 MS 14 MS 95 MS 1 S	
MECHANICAL DIMENSIONS • WEIGHT	[illegible] 650 G	
POWER SOURCE	[illegible]	
MEDIA • DISK DIAMETER • DISK SIZE	COMPACT FLOPPY [illegible] 72MM [illegible]	

Signal interface for the 3-in. format is identical to the minifloppy with few exceptions. So, the 3-in. format can be used in place of the conventional minifloppy. Also, this compatibility makes existing test equipment and tools useful for production of the 3-in. microfloppy (Fig 4).

The 3-in. format is being developed according to the following plan:

- **First step.** 250 Kbyte/side, 100 tpi, 8946 bpi, 40 tracks/side.

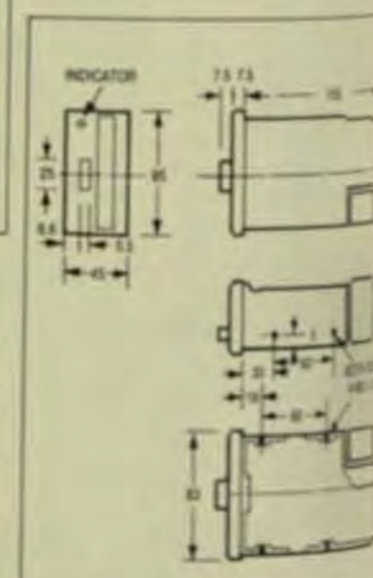

Fig 4 Dimensions for the 3-in. CFD drive are shown in millimeters.

- **Second step.** 500 Kbyte/side, 200 tpi, 8946 bpi, 80 tracks/side.
- **Third step.** 800 Kbyte/side, 200 tpi, 15,000 bpi, 80 tracks/side.

The diskette and the drive in the first step are already in mass production.

The second step, in which track density is increased from 100 to 200 tpi, will be achieved by improvements in the media.

The media expands and contracts due to the thermal and hygroscopic properties of the Mylar substrate. The magnitude of these expansions and contractions depends on the thermal and hygroscopic coefficients—conventionally, $(17 \pm 8) \times 10^{-6}$ cm/cm/°C (thermal), $(0 \text{ to } 15) \times 10^{-6}$ cm/cm/$^{-6}$ cm/cm/RH% (hygroscopic).

It also depends on diskette size—the smaller, the less expansion. This phenomenon restricts the track density to within 48 tpi for an 8-in. disk and to within 100 tpi for a 3-in. diskette. The total tolerance allowable for head-positioning error is basically decided by the track density of a drive. For a 100-tpi drive, it may be 45μ. However, in general, about half of the total tolerance is due to the platter size.

The improved media of which coefficients of expansion are $(17 \pm 4) \times 10^{-6}$ cm/cm/°C (thermal) and 5×10^{-6} cm/cm/RH% (hygroscopic) is under development and have been announced to be available by the end of this year. This new media will allow the second step in development to be realized. That drive will result in a total tolerance of about 22.5μ, but the required mechanical accuracy of the drive will be about 10μ.

Such accuracy will be made available for the drive by improved product technology and precise mechanical engineering. The 10μ accuracy is estimated to be the same as that required by 135-tpi format of 3.5-in. types when that format is based on a disk substrate with coefficients of expansion within a range from 0 to 2.48×10^{-5} cm/cm/°C (thermal) and to a maximum of 1.5×10^{-5} cm/cm/RH% (hygroscopic), as described in a January 26, 1983, ANSI proposal for a 3.5-in. cartridge.

The third version will be made available by increases in the linear density of the improved 200-tpi disk substrate.

The second and third versions can be achieved by use of improved media. In other words, the 3-in. CFD standard will be sufficient and not redundant as capacity is expanded up to 800 Kbyte/side and 1.6 Mbyte/diskette, making it compatible with an 8-in. disk. The development of the second and third steps is now being executed by the three sponsoring companies. The 200-tpi drive with 1-Mbyte capacity will be offered this year.

Since the 3-in. CFD has a 100-tpi track density, it exhibits almost no differences from the conventional minifloppy disk drive of 96-tpi track density. Accordingly, the technology for 96-tpi drives can be applied to the 100-tpi 3-in. drive and the CFD can be produced in high volume at low cost. ■

Yasutaka Nakajima *is chief engineer of the Mechanical Component Development Dept. at the Wireless Research Lab. of Matsushita Electric Industrial Co., Ltd., Japan. He has a BS in mechanical engineering from Hokaido Univ.*

Flexible Disk Drive Testing Should Be Flexible, But...

Uncompromising! At Brikon we won't budge a nanosecond when it comes to accuracy, resolution, and quality. Brikon testers are completely self-contained units with many extra built-in features.

Brikon flexible disk drive testers are accurate, easy-to-use, and less costly than most others. Major OEMs around the world recognize Brikon as the standard for flexible disk testing. Compare these features:

- Easy to use: names in English and numbers in decimal
- Self-contained DC power supply for drives
- Automatic testing option: tests and qualifies all major subsystems within the drive under test
- Multiple drive testing and radial alignment testing options

Contact Brikon today for in-depth specifications.

BRIKON INC.

Brikon, Inc., 22981 Alcalde Dr., Laguna Hills, CA 92653
Phone: (714) 768-3838, Telex: 181645

THE 8MBYTE FLOPPY SURPRISE

At last, a disk drive which accommodates both stand alone removable media data storage and Winchester backup applications — all with one device. The Amlyn drive is a superstar as the sole storage device on your system. It combines the benefits of high performance, high capacity data storage with removable, low cost user friendly media. All this for less than half the price of a Winchester drive and its backup device.

If you use Winchester drives, you need backup. Amlyn's drive is THE backup device which is interface, format and software compatible with the Winchester itself. So compatible, it can take over for your Winchester when it fails so that a Winchester failure doesn't shut your system down. Amlyn backs up your Winchester drive itself — not just the data on it. Now backup is more than something you hold in your hand when your Winchester fails.

Dynamic On-Line Operation

Amlyn drives provide 8 MBytes in the same space as a standard mini-floppy. Amlyn has combined the capacity of Winchester drives with the advantage of removable media.

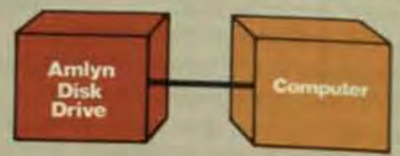

Available Now,
Contact Us At 408-946-8616

amlyn
2450 Autumnvale Drive
San Jose, CA 95131
(408) 946-8616
TWX 910-379-0029
TLX 171627 Amlyn

Random Access Winchester Back-Up

Now, you can have Winchester system backup, not just data backup with interface and software compatibility to your Winchester as well.

Removable Media

Amlyn's 8 MByte removable MiniPac cartridge holds and protects diskettes from damage and mishandling, and is easily removed and replaced. Each cartridge holds five mini-floppy diskettes.

Standard and Flush Mount Design

Amlyn drives allow OEM and systems house design flexibility with a variety of styles which assure form fit for compatibility with standard mini floppy drive dimensions.

5.25-In. Winchesters Must Handle Tradeoffs For Higher Capacities

Mechanical factors and lack of parts availability are among the problems that must be resolved before small Winchesters can achieve greater capacities.

by **Doug Mahon,**
Seagate Technology

As capacities increase, resolving the design problems of 5.25-in. Winchesters will require careful evaluation of which tradeoffs are practical. The universal acceptance of the ST506 interface is a good example of how a de facto industry standard can promote broad-based acceptance for a new class of product. At the same time, the definition of this controller and some fundamental design considerations have put limits on the capacities that designers of microcomputer-based small-business systems, word processors, workstations, and personal computers can expect from these drives over the near term.

In early 1980, Seagate Technology introduced the ST506, a two-platter, 6.38-Mbyte, 5.25in. Winchester disk drive. Since that time, nearly 35 other companies have announced products of this type, and this year, according to *Disk/Trend Report*, over 383,000 5.25-in. drives with capacities under 30 Mbytes will be shipped to builders of microcomputer-based small-business systems, workstations, word processors, and personal computers. By 1985, this figure is expected to reach 686,000.

Compatibility with controllers designed to handle the ST506 was important to the growth of this market for several reasons. First, it permitted different manufacturers of compatible drives to act as second sources—a relationship that's beneficial to both drive makers (who aren't barred from any part of the 5.25-in. Winchester market due to incompatibilities) and to OEMs (who can tap several sources for a drive with given capacities). Second, controller manufacturers have a benchmark by which they can make products available off the shelf to serve with any ST506-compatible drive.

This controller design has proven more than adequate for the vast majority of the drives projected to be shipped over the next few years. Meanwhile, however, drives with higher capacities continue to be announced. This raises the question of just how far the ST506 interface can go, given the need to produce large numbers of drives and given the new levels of quality and reliability required to compete in the market for microcomputer-based systems.

This article explores the issue of ST506 compatibility, and shows what kind of drive capacities can be expected if this interface continues to enjoy its current popularity. This article will also discuss some of the constraints imposed on high-capacity small drives, and possible solutions to the problems of building large volumes of highly reliable small Winchesters at prices the market will accept.

ST506 COMPATIBILITY

Two key specifications characterize an ST506-compatible controller. The first is a 5-Mbit transfer rate. The second is a specification that calls for 10,416 bytes per track (unformatted); 8192 bytes (formatted) (see Fig). These specs are important to both drive manufacturers and OEMs—not because they represent some optimum set of characteristics, but because they de-

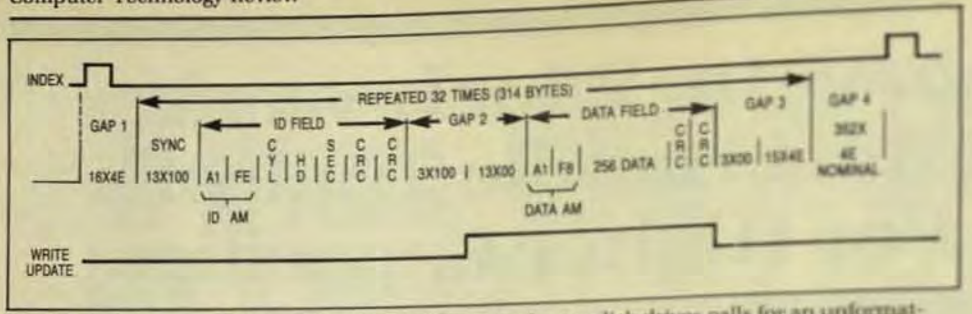

Fig The sector format of ST506-compatible 5.25-in. Winchester disk drives calls for an unformatted capacity of 10,416 bytes per track. Typically, the surfaces of these drives are divided into 32 sectors, each containing 256 data bytes yielding a formatted capacity of 8192 bytes per track.

fine the performance limits of a controller that has become widely accepted in the industry.

INCREASING THE NUMBER OF TRACKS

As a result of the need to maintain this compatibility, three methods have been developed for use singly or in combination to increase 5.25-in. Winchester capacities while permitting the use of widely available controllers. One method is to increase the number of tracks per surface by increasing the head stroke. This approach increases the area available for data storage without affecting the drive's ST506 compatibility. At the same time it minimizes the impact on actuator design because it doesn't increase the track density.

There's also a limit to this approach, however. To see why this occurs, consider how the optimum number of tracks is determined. Starting from a knowledge of the maximum flux density possible (established by limitations of the head/disk interface), calculate the maximum area of the disk you want to use. The outer radius (or the outer track) of the read/write area is set by the disk's size; the other dimension you're concerned with is the inner radius—the position of the innermost track on the disk.

It might seem best to make the inner track as close to the spindle as possible to maximize the read/write area, but the inner track is the smallest one on the disk and has the highest flux density—the maximum practical for that head/disk interface. If you make the inner track too much shorter than the outer one, the outer tracks aren't taking full advantage of the drive's maximum flux density; in effect this reduces capacity on the outer tracks. Any storage capacity gained by moving the inner track further in is lost on the outer tracks.

A few simple relationships help determine the optimum radius for the inner track:

$$\text{No. of tracks/surface} = \frac{\text{outer radius} - \text{inner radius}}{\text{track width}}$$

$$\text{Length of track} = 2\pi r$$

$$\text{Capacity} = (\text{max. bit density})(\text{length of inner track})(\text{no. of tracks})$$

By substituting as appropriate, and setting the first derivative of the third equation equal to zero, you find that the ratio between the outer radius and the inner radius should always optimally be two to one, given any maximum flux density.

As an example, the outer radius of a 5.25-in. Winchester disk is 2.625 in.; and say the outer flying radius for the head is 2 in. The best possible inner radius thus equals 1 in., giving a circumference of 6.28 in. This is equivalent to a linear velocity of 377 ips at 3600 rpm. Even though this is the optimum inner radius on a 5.25-in. Winchester and can lead to significant improvements in disk capacities in ST506-compatible drives, currently available read/write heads limit the extent to which the stroke can be usefully increased.

The problem is one of head stability. The standard 3350-type heads used on most Winchesters were designed to fly on a 14-in. disk where the inner radius is about 5 in. That radius gives a velocity of 942 ips at 3600 rpm—the slowest the head was originally intended to fly, yet much faster than the speed on the 5.25-in. disk. The result can be unstable operation on the inner tracks of the smaller disk as the head is moved closer to the spindle.

Just as a head normally pitches and yaws as it slows down and lands on the disk surface, it exhibits these same instabilities when it flies too slowly. Although the head is nearly perpendicular to the air flow from the disk when on the outer tra flow comes at an angle on the tracks, causing additional he instabilities.

The head also doesn't fly enough at very low speeds. Th flies at a nominal altitude of typically ±3 µin., at 470 ips. The surface is burnished to a toleran µin., so even under ideal cond there isn't much margin for er Slowing the flying speed to 37 lowers the head to approxima 15 µin.

Considering the toleranc disk burnishing, the adverse e temperature, and the drive's a (far above sea level the lower a sure doesn't support flight a the operating margin is gone. head simply won't work at sp slow as 377 ips, and thus can't used on the small radius that such low speeds.

Even though the 3350-type has been widened to improve formance at low speeds, the l speed at which it flies stably enough height is about 470 ip ing it further could create pro by extending over unburnish beyond the inner track. (The b nished area can't easily be ext because disk manufacturers ha same problem as drive makers burnishing head won't fly prop very short radii.)

What's needed is a head wi stiffer air bearing than the con al 3350 design. The minislide and 3370-class flexure combin provide this better air bearing as less sensitivity to the angle flow, but this alternative invol availability problems. Althoug heads should prove very bene in the future, they are only no coming available in high-volu production.

FLUX AND TRACK DENSITY

The recording media can als cate attempts to increase the ing area. Seagate's ST412, for stance, has a 0.9-in. stroke. B meet the two to one outer-to-

Today, people are solving their back-up problems with this fast, reliable, 10 MB disk cartridge drive.

IOMEGA's 10 Megabyte cartridge drive outperforms most winchesters.

So you can back-up 10 Megabytes from your fixed disk in less than 30 seconds.

The easy to use cartridge sports the industry's lowest price tag, only $30 each in OEM quantities.

IOMEGA's imbedded closed-loop servo guarantees interchangability of cartridges between drives. And the standard interface is SCSI compatible.

Solve your backup problems with the fast, reliable IOMEGA 10 Megabyte cartridge drive.

Call IOMEGA for a personal demonstration. **And ask about our OEM Special Evaluation Offer.**

IOMEGA Corporate Headquarters, 4646 South 1500 West, Ogden, Utah 84403. 801/392-7581. **San Jose**, CA 408/263-4476. **Coral Springs**, FL 305/755-1060. **Woburn**, MA 617/933-2000. **Dallas**, TX 214/458-2534. **Brookfield**, WI 414/782-5229. **Los Angeles**, CA 714/855-1211. In Europe, Sparrow Corp. **Slough**, UK (0753)76533. **Weisbaden**, (6121)700862. **Paris** (1)3621010. **Milano** (2)718531. **Brussels** (2)7626200. **Zurich** (1) 814-3131.

IOMEGA™

Copyright © IOMEGA 1983

track ratio because of considerations related to head-flying stability. But if a better head were available in quantity, it would take an increase in flux density to widen the stroke to 1.025 in. because the drive's maximum density is already employed at a larger recording radius. The smaller radius set by a 1.025-in. stroke would require a platter that allows the flux density to go from the current spec of about 10,000 flux changes per inch (fci) to about 12,000 fci.

The media to permit this flux level will cost more than the disks now used because the manufacturing yields diminish when the disks are certified against smaller flaws; a flaw that didn't use to count now makes a difference. Moreover, the tighter the certification spec, the more you run the risk of being caught by an avalanche curve that describes the yield on media.

This curve gives a measure of the number of errors per surface as a function of the size of flaws relative to the average read-signal amplitude. In other words, if your average amplitude equals 1V, and you count errors only at 0.2V, you won't find many flaws. At 0.6V, you'll find a few; at 0.7 V you'll find a lot; and by the time you reach 0.8V, there will be a tremendous number.

The result is that you have to accept more flaws per surface, or the disk manufacturer must yield—at a premium—better disks out of the fabrication process. The latter alternative won't serve high-volume needs. Accepting more flaws can be a valid course, but not for simple error-handling systems that spare out whole tracks instead of sectors to account for errors. At 10,000 tpi and 100 errors, for example, you would have to spare out a major part of your total capacity.

Once you've determined how much of the disk surface to use for recording, it's necessary to establish how many tracks will fit in that area. This value depends on the number of tracks per inch the drive will permit, and that parameter involves a determination of how narrow and close together the tracks can be.

The better the drive's positioning accuracy, the narrower you can make the tracks. The positioning accuracy, however, is limited by mechanical factors in the spindle motor. The main culprit here is nonrepetitive spindle-motor runout, which represents the play in the spindle motor's concentricity. The less precise the motor-bearing assembly is, the more runout there is. And because positioning accuracy depends on the disk being as concentric as possible, runout can defeat any design that demands too much accuracy.

The source of the runout problem lies in the spindle ball bearings and the race they roll in. The small size of 5.25-in. Winchesters means that this race is only 1 in. in circumference, yet to withstand the load they carry, the balls it contains must be the same size as those used in 14-in. drives.

The result is a greater sensitivity to bearing imperfections. Ball-bearing quality is measured by ABEC numbers, with higher numbers denoting higher quality. Standard motors for today's disk drives call for ABEC 9 balls, but even these aren't manufactured directly.

Thus, to improve on the current standard of about 40 or 50 µin. of runout in a spindle motor entails some improvement in bearing technology. Such mechanical advances come very slowly, so you can't expect an overnight decrease from today's average 40 µin. of runout. Nonetheless, you can buy a spindle motor that has only 5 µin. of nonrepetitive runout, and you can even buy 10 of them. But you can't buy the 1000 a day needed to support the broad-based market that has come to characterize the market for 5.25-in. Winchesters.

Having 40 µin. of spindle-motor runout means that you can never guarantee that any track will be any closer than 40 µin. to where you want it to be. This margin imposes a finite limit on track density in high-volume drives. And bear in mind that you can count on 40 µin. only at one temperature.

Because runout is nonrepetitive, only a servo with infinite bandwidth can correct for it. And even this assumes a perfectly accurate servo that's continuous rather than one of the more common designs that incorporate servo information embedded at each sector gap.

To see what specific limits runout imposes on increased track density, consider that at 960 tpi on a 5.25-in. Winchester disk drive, tracks are 800 µin. wide and 1041 µin. (Visual[illegible] between track centers. A spindle runout of 40 µin. represents [illegible] track width. Over the drive's temperature range, the spindle could [illegible] more, causing about another [illegible] of error and increasing the [illegible] tracking discrepancy to approximately 10%. Adding about another [illegible] for the inevitable hysteresis [illegible] brings the total tracking variance [illegible] 15%—where off-track errors [illegible] tronomical. This error level doesn't even account for inaccuracies in the head carriage.

Given the mechanical nature of these problems, track densities in excess of 600 tpi—much less 960 tpi—aren't going to become practical for 5.25-in. Winchesters any time soon, despite the appeal that this technique offers for significantly increasing drive capacities while maintaining ST506 controller compatibility. This level of performance is even difficult to achieve in 14-in. drives, where the bearing race is much larger and the units work in a more controlled environment.

Recording at 480 tpi is possible for high-volume drives, but going as high as 600 tpi is doubtful. After accounting for gross errors, all the error budget is used up at 600 tpi and costs will increase dramatically for better heads, media, motors, and servo electronics.

It might be possible to make these track-density levels easier to implement if you set restrictions on the drive's use. You could eliminate temperature problems, for example, by not allowing any reading or writing until the drive has warmed up for half an hour. Microcomputer users aren't likely to put up with such limitations, however.

INCREASING THE NUMBER OF DISKS PER SPINDLE

Another way to increase a drive's capacity without interfering with its ST506 compatibility is to use a greater number of disks. Most 5.25-in. Winchesters offer a maximum of three platters in one drive, so by increasing that number to four or more, you could easily boost capacity—if there were no restrictions. However, the standard form factor of 5.25-in. Winchesters limits the number of platters you can stack.

Another problem is that the temperature increases inside the head/disk assembly (HDA) as the head spins and, added to the maximum ambient temperature in which the drives must work, this heat can produce IC problems. For high track densities, ICs are placed inside the HDA to amplify signal levels before they become subject to noise. But if high temperatures occur inside the assembly, either the ICs fail outright, or their lives are shortened. And failed ICs inside the HDA can't be repaired in the field.

Further, if the amplifier chip that handles writing tasks is inside the package, its own self-heating can help destroy it. The chip can generate a 50°C rise at its junction while writing, so even with only three platters at 61°C, the part's junction temperature can reach 111°C.

Oxide-coated media can also suffer from the heat generated by the spinning disks—one factor that makes thin-film media more attractive to drive vendors from a reliability point of view. Thin-film media is more durable than oxide-coated platters, but here again, lack of availability prevents its use in high-volume production. Seagate, for example, recently stopped taking orders for its 6.38-Mbyte ST706 (removable) disk-cartridge Winchester because the 10 to 20 plated-media disks per drive per year needed to support these drives weren't available and didn't look as if they would be in the near future.

High temperatures can result in other difficulties, too. Spindle tilt, for instance, causes a difference in track position from one platter to another. And high temperatures aren't the whole problem, either. It's important to consider the entire range of temperatures—from cold starts at 4°C to extended operation at 65°C—that a drive might undergo. Because the drives aren't protected in a computer-room environment, their temperature span must be very wide. All these temperature factors probably limit a volume-production 5.25-in. Winchester to no more than four platters.

In the final analysis, for an ST506-compatible 5.25-in. Winchester, the highest capacity that can be expected in the near future from a reliable drive in high-volume production is somewhere between 47 and 51 Mbytes. This range assumes a 4-platter drive operating at 480 tpi.

Upcoming developments will enlarge this capacity. For example, if you increase the number of bits per inch, you increase the amount of data per track. Technologies such as thin-film media capable of supporting vertical recording, and thin-film read/write heads all offer the promise of higher capacity without resorting to longer stroke lengths, track densities, or disks.

But anything that raises the bit density also boosts the drive's transfer rate, requiring new controller designs. This is because the read/write head flies over more bits per second on a track and the drive must therefore transfer the data faster to the host computer. Any change in transfer rate makes a drive incompatible with the ST506 interface.

One solution is to employ an interface that specifies a faster transfer rate. Although at least two such 5.25-in. Winchester interfaces have been proposed (ESDI and the ANSI standard for 8-in. Winchesters), the industry is not in general agreement about what the alternative should be, and no broad-based support exists for either.

It isn't clear at this point what the market really needs in a high-performance interface, so for the time being it's important to consider what level of technology is practical, given the specifications of the controllers available. A drive's price will give some idea of its practicality, but evaluating small Winchesters generally will require an understanding of the performance factors involved and time for new technologies to prove themselves in the field. ■

Doug Mahon, *senior vice president of engineering at Seagate Technology, has been with the company since it was founded in late 1981. Doug directs all of Seagate's research and engineering activities.*

"Is there a reasonably-priced, reliable disc drive available now?"

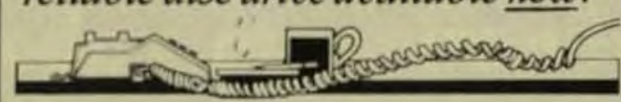

"YES! Contact Linda Smith at Disc Tech One."

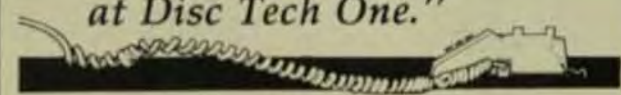

80 MB	160 MB	300 MB
$3,000*	$4,000*	$5,000*

Features Include:

- SMD interface
- Compact rackmount size (7" high)
- Integrated power supply
- Clean air package
- Quick performance
- DEC, DG, Perkin Elmer, TI compatible

*OEM Pricing

DISC TECH ONE, INC.
849 Ward Dr., Santa Barbara, CA 93111
(805) 964-3535

The 100MM Winchester
Removable. Half Size. Half Price. Full Performance

It's here. Winchester capacity and performance at half the size, half the price. And yes, available in removable or fixed disc drives.

The SyQuest 100mm (3.9") SQ306 packs five megabytes (formatted) in half the height of a 5¼" Winchester. And when the Q-Pak™ cartridge is full, just slip in another one. It's the best of both worlds—the reliability of Winchester with the transportability of removable cartridges.

A better drive.

SyQuest drives give you a better fit. Mount SyQuest drives almost anywhere. Under a keyboard. In your terminal. Fit two in one minifloppy space. SyQuest drives are only 1.625 inches high, 4.8 inches wide, and 8 inches deep.

Easy integration. The SQ306 has the same pin-outs, timing, data transfer rates, and track/sector formatting as industry-standard 5¼" Winchester drives. Use standard Winchester controllers and interfacing procedures, standard minifloppy DC power supplies.

Better price/performance. SyQuest delivers five megabytes with proven Winchester heads, positioning, brushless motors and air filtration. Buffered seek reduces average seek time to 75 msec. But the cost is half of comparable 5¼" Winchesters.

Q-Pak™—a better cartridge.

Better reliability. Closed-loop servo with imbedded digital servo (**DigiLok™**) ensures cartridge interchangeability while allowing variable sectoring. **Chromaflux™** graphite coated thin film metallic alloy discs protect against contamination. No long purge cycle required.

Available now.

SyQuest is shipping. Second sources are available. So order your evaluation units today. For more product information, circle our readers' service number. For delivery and pricing information, write or call Larry Sarisky, SyQuest Technology.

SyQuest
47923 Warm Springs Blvd. Fremont, California 94539
415/490-7511 TWX 910-381-7027

It Fits.

Open- Vs. Closed-Loop Positioning Systems for Winchester Disk Drives

Lower-cost open-loop positioning systems are ideal for 5- to 20-Mbyte single-user systems, but closed loops offer 20 to 100 Mbytes for multi-user applications.

by **Robert Abraham**, Tandon Corp.

In selecting Winchester disk drives, systems designers have a choice of drives with closed-loop or open-loop head-positioning systems. Both have advantages, so selection ultimately should be determined by the system and application requirements.

Typically, an open-loop system is lower in cost and offers a greater advantage in capacities from 5 to 20 Mbytes with access times of 90 to 100 ms, which makes it ideal for single-user applications. A closed-loop system, on the other hand, performs best at capacities ranging from 20 Mbytes to 100 Mbytes or more, with access times of 30 to 40 ms. As such, it's ideal for multi-user applications. An examination of the technologies, performance tradeoffs, and applications will make the systems integrator's choice obvious.

OPEN- AND CLOSED-LOOP SYSTEMS

Open-loop systems, regardless of the technology employed, share two common characteristics (Fig 1). First, stepper motors are used to drive and control the head-positioning mechanism. The electromechanical nature of the stepper motor limits the positioning accuracy of the system and dictates the access time. As such, it becomes both the heart of the open-loop positioning system and its limitation. However, its low cost and proven performance are more than adequate for lower-capacity, low-speed Winchesters and are warmly regarded by systems designers.

The second common characteristic of open-loop systems is the lack of feedback of positioning information. The positioning circuits send information to the positioning mechanism to move the read/write heads to the desired track location. Because of the lack of interaction and intelligence, there's no monitoring and feedback of this positioning information.

There are two common open-loop positioning systems used in Winchester drives today. One is the split-band positioner, which has two major variations of its own—rotary and linear—and the other is the rack-and-pinion positioner.

A split-band positioner consists of an anchored metal band wrapped around a capstan. A stepper motor rotates the capstan to move the positioner (Fig 2). The accuracy of this type of system is underscored by its wide use in Winchester disk drives. Positioning errors are minimized by accurately grinding the capstan's diameter and eccentricity. The split-band positioner is fast; track-to-track access time is typically less than 3 ms and long-term reliability is inherent.

In a rotary system, the head-positioning mechanism is placed on the end of the radius and swings in an arc. As the motor steps circularly, the arm moves across the disk (Fig 3). This method transfers the rotary motion of the stepper motor to rotary motion of the head carriage, thus achieving maximum positioning accuracy. The driving mechanism is mounted outside of the sealed storage chamber, thereby eliminating a major possible source of contamination.

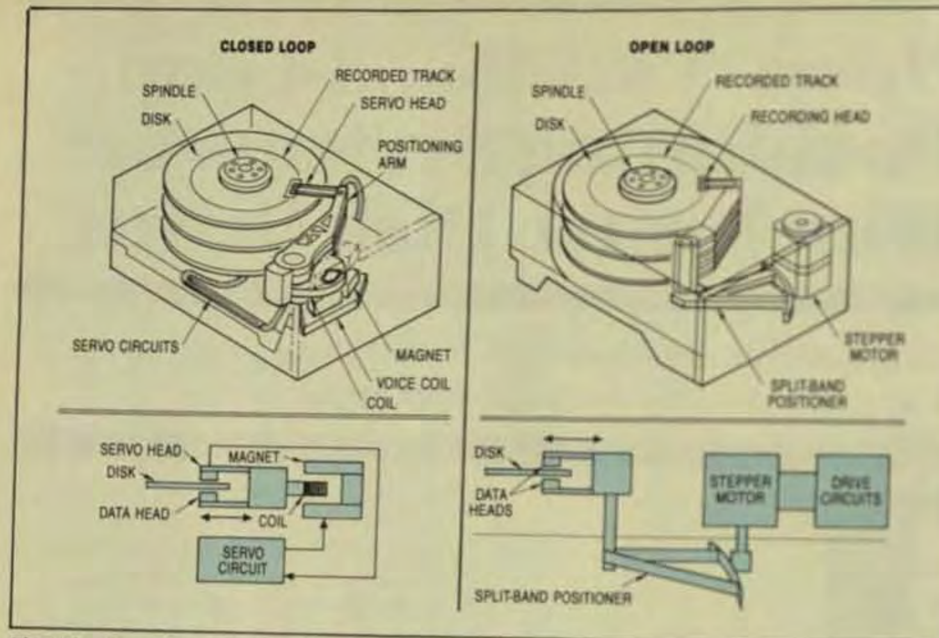

Fig 1 The closed-loop Winchester system (left) features sensors that monitor and feed head-positioning data to servo circuits that make in-course corrections. It also uses a voice-coil actuator to achieve higher densities and faster access times. Open-loop systems (right) employ a stepper motor to position the heads. This technique is cost-effective and accurate at lower capacities.

Another advantage to the rotary system is its shock resistance. Depending on the design, the mechanical advantage over a linear system can be 5 to 1, 10 to 1, or 20 to 1. If you place 10G of shock on the axis, for example, only a very small part of it will be transmitted to the head, so it's prevented from skating across the disk. Because of this mechanical advantage, rotary systems are less vulnerable to shock than linear systems.

In linear positioning systems, the rotary motion of the stepper motor is converted into linear motion of the head stack. The head-positioning mechanism moves straight in and out on a radius over the disk surface. A tighter packaging job is required to seat the elements in a 5.25-in. drive envelope. Also, since most of the mechanical parts must be mounted inside the sealed chamber, it's difficult to change the motor and other high-wear components. Further, a one-to-one mechanical transmission of force to the head makes the drive more vulnerable to shock.

The other type of open-loop system is a rack-and-pinion positioner. In this case, a stepper motor drives a toothed pinion (bar) on a toothed rack to convert the circular motion into linear motion. Read/write heads are mounted on a carriage and move across the radius of the disk surface (Fig 4).

There are several advantages to the rack-and-pinion positioner. It's fast—equaling the access times of the split-band positioner—accurate and cost-effective. The mechanism is easy to manufacture and construct with few parts. But there are drawbacks, too. Because it's a relative newcomer to the open-loop Winchester arena, it hasn't been field proven. And its longevity is questioned, primarily because the rack and pinion are mechanical parts subject to wear. Moreover, the mechanism must be mounted in the sealed memory chamber along with the lubricant required by the gear teeth. In addition, replacement or maintenance of the mechanical parts requires the memory to be opened in factory clean-room conditions—a not-so-trivial expense.

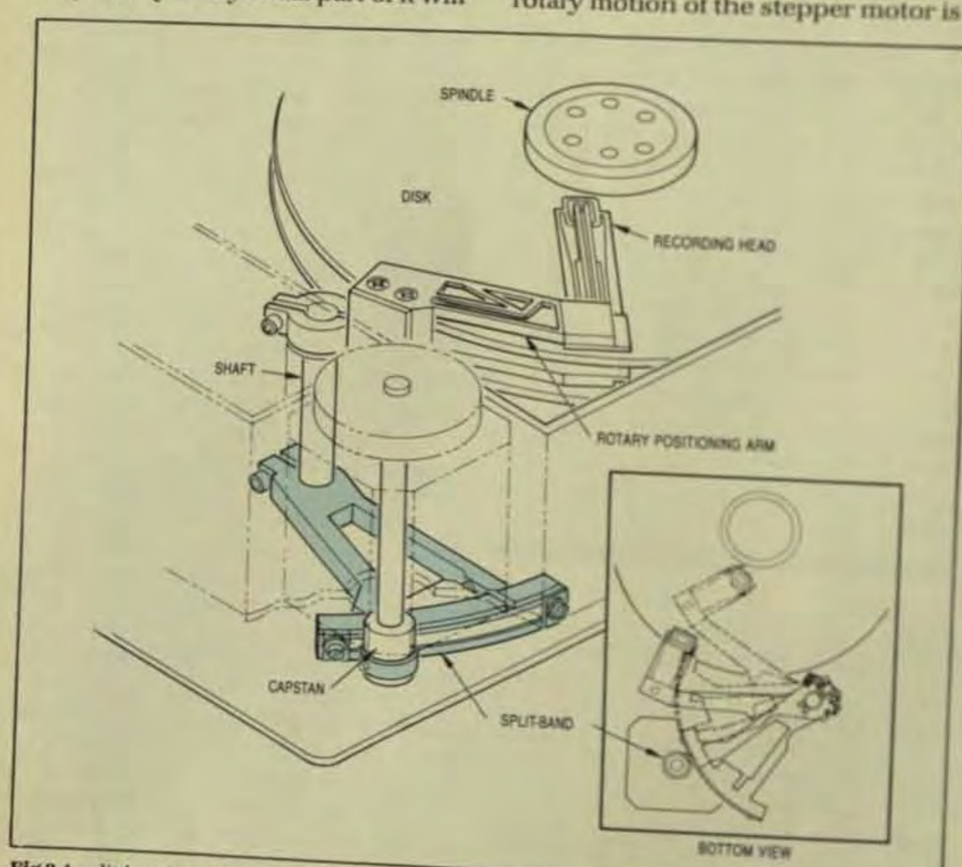

Fig 2 A split-band positioner provides fast access times and long-term reliability. The inset illustrates the split-band technique.

CLOSED-LOOP POSITIONING SYSTEMS

Two characteristics are common to closed-loop positioning systems. First, positioning information is continuously being fed back to the drive's servo circuits from sensors on the read/write heads to let the drive know the head's location at all times—whether it's on track and whether it's at the right location on the track. Thus, the drive corrects the head position as needed, homing in accurately.

Second, a voice-coil positioner is used. The voice coil is actuated by one or more magnets and moves the head carriage in and out over the su

Some Standards Are Untouched by Time

Standards of excellence and integrity are discovered and verified only from experience. These are qualities to be earned over time and, perhaps, that is why we most often look to the past to assign such values to a company or its products.

At MiniScribe Corporation, which, in less than two years, has become a major force in the manufacture of Winchester disk drives for microcomputer systems, today's business practices are motivated by the kind of unwritten standards that will be remembered tomorrow. This doesn't bypass the test of time, but it tends to compress past and future into a rewarding present experience for customers and employees alike.

At MiniScribe you will find innovations in technology, precision in craftsmanship, and value in product performance. But most important, you will find a contemporary, public standard of excellence befitting a private collection.

A classic jeweled Waltham timepiece, circa 1901, with a richly embellished solid gold case. From the Gary Stull collection.

The MII, 5¼" rigid disk drive with 6.4 or 12.8 megabytes of storage, very contemporary and noted for precision performance and high reliability. From the MiniScribe collection.

MiniScribe

MiniScribe Corporation
1871 Lefthand Circle
Longmont, Colorado 80501
(303) 651-6000
(TLX) 216-174 (MINI-UR)

See us at Systems '83 in Munich, October 17, and at COMDEX in Las Vegas, Nevada, November 28.

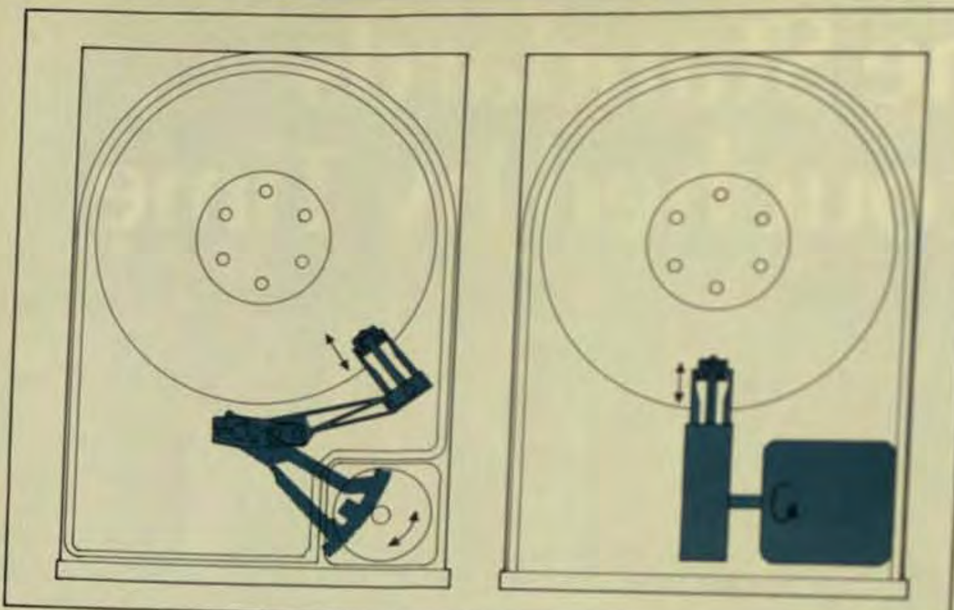

Fig 3 Split-band positioning systems come in two forms. A rotary positioner (left) uses the circular motion of the stepper motor to drive the head-positioning mechanism in an arc over the radius of the disk surface. A linear actuator (right) converts the rotary motion of the stepper motor into the linear motion of the head positioner.

face of the disk. The voice-coil positioner, essentially a linear motor, is extremely fast and accurate. In combination with servo positioning, it makes possible the higher densities and faster access times found in closed-loop disk drives.

There are some disadvantages to use of the voice coil as the motor. Magnets are heavy and rather bulky. They create magnetic fields that must not interfere with the read/write operation in the compact confines of the 5.25-in. Winchester world. They also consume a lot of power. Nonetheless, their accuracy and speed are well worth these tradeoffs.

There are three basic types of closed-loop positioning systems—dedicated servo, embedded servo, and optical servo.

Dedicated-servo system. In the dedicated closed-loop servo system, one of the disk surfaces is dedicated wholly to the storage and retrieval of servo-positioning data. For example, the Tandon TM705 disk drive has three platters (six surfaces). One entire surface is dedicated to servo data; the other five to operating data (Fig 5).

The servo-track information provides polar coordinates for the location of the data head(s) at any given instant. These polar coordinates are generally broken into one or more elements. For radial position, the servo tracks are located at defined radii and pitch (spacing) with respect to the center of rotation. These tracks are then grouped into bands—the outer guard band, data zone, and inner guard band.

The angular-position coordinates usually consist of an index, the start/end point of the data tracks. However, it's not uncommon to further divide the track into sectors, which are encoded by groups of pulses called servo frames. There are generally two kinds of servo frames on each servo track, labeled 0 and 1. The index and sector boundaries, as well as the outer guard band, data zone, and inner guard band, are defined by unique sequences of 0s and 1s in servo frames.

Since the pulse pattern for frames is different on each track, direction information can be transmitted to the head-positioning system. A common system is to have definitions, often referred to as A and B patterns, for odd and even tracks. Another system, a quadrature pattern, changes the servo-frame definition in groups of four tracks. Obviously, there are many other possibilities.

Once the servo head is positioned between the desired two tracks, it's moved in and out until amplitude of the position pulses of adjacent servo fields is equal. This technique defines the positions of the data tracks.

The dedicated-servo surface allows continuous closed-loop positioning. Track position is constantly being monitored, and head drifts—the drifting of heads away from the center of the track—are instantly corrected. As a result, tracks can be spaced very closely together, and capacity increases dramatically.

The ideal dedicated-servo system is a combination of a position servo that positions the head from the servo data it reads, and a velocity servo that positions the head by counting the tracks it crosses. A dual-mode servo system uses both techniques—velocity servoing to

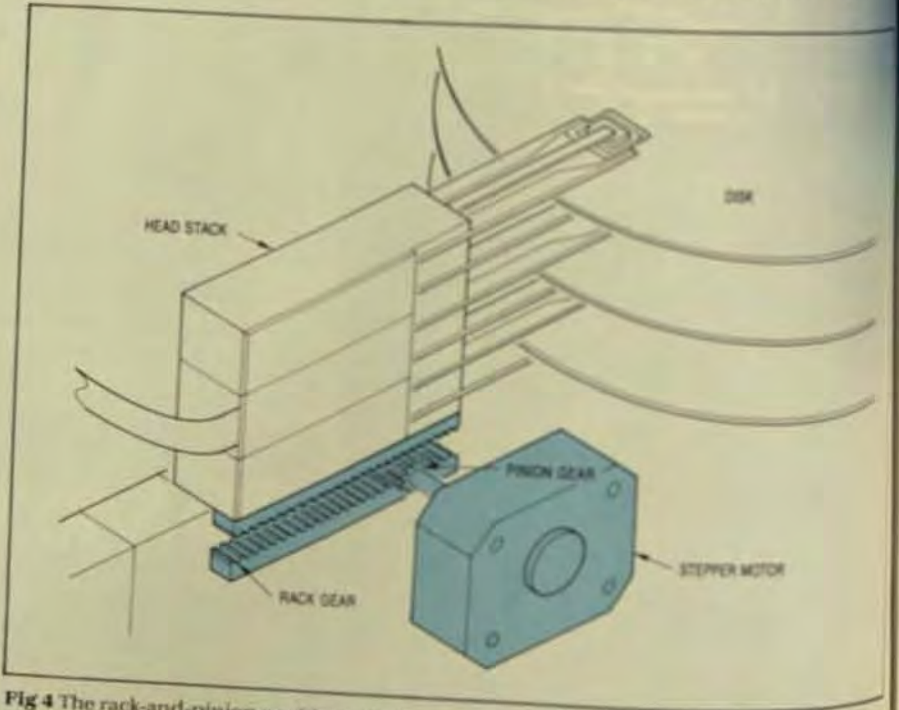

Fig 4 The rack-and-pinion positioner features a stepper motor that drives a toothed pinion on a toothed rack. Read/write heads mounted on the rack are moved across the radius of the disk surface.

NOW—SUPER-PERFORMING 5¼" WINCHESTERS WITH BUILT-IN RELIABILITY.

THE WREN™

The Model 9415 WREN disk drive gives you a host of value-added features. From microcomputer and LSI control logic to special-formula media, high technology has finally delivered a 5-1/4" winchester with more reliability.

Compare for yourself. Call your local Arrow or Kieruff distributor today. Or write: OEM Product Sales, HQN08H, Control Data Corporation, P.O. Box 0, Minneapolis, MN 55440.

CONTROL DATA

Information Hot Line 1-800-621-1776.
In Illinois, 1-800-572-6724.

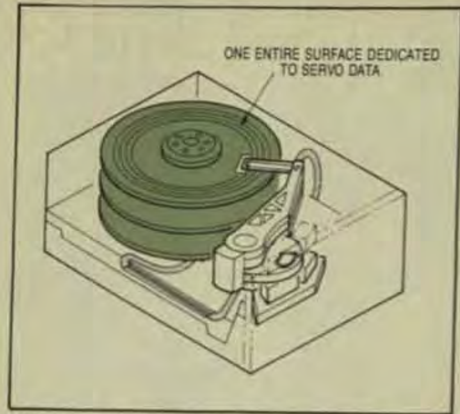

Fig 5 In a dedicated-servo system, one entire disk surface is devoted to the positioning information.

Fig 6 In an embedded-servo system, positioning data is recorded and retrieved in interrecord gaps on every disk surface.

move quickly to a designated track, then position servoing to home in on that track.

A major advantage of a dedicated closed-loop servo system is that a read/write head's position is updated and corrected continuously in real time. In every instant of time, the positioning system is working to correct itself and keep itself on track. It always knows where it is. In addition, positioning accuracy is much higher and heads can be positioned to a greater tolerance level. A track density of more than 1000 tpi is possible—approximately three times the track density of typical open-loop systems.

Embedded-servo system. In the second major closed-loop servo-positioning method, servo information is embedded between sectors on each track (Fig 6). The primary difference between the two types is that in embedded-servo systems the servo information is available intermittently, rather than continuously. Embedded systems have a lower track density than dedicated servos, and lower head-positioning speed and accuracy because of delays in head switching and settling-time. An embedded system is accurate immediately after it's updated on a track.

When servo information is only available intermittently, the positioning information is in a different format than with dedicated units and a different method of track identification is required.

Another disadvantage of the embedded servo is that interfacing requirements are different because the recording surface is formatted differently. The drive has to identify the servo data and separate it from other data recorded on the surface of the disk, so that the servo data isn't sent to the interface along with the data requested. A part of each disk rotation is dedicated to the servo-processing sequence, giving the controller an extra load to cope with.

But the embedded servo is not without its advantages. It doesn't require a dedicated head or disk surface, which is a cost-saving benefit over the dedicated servo.

Optical servo system. Though not as widely used as the dedicated and embedded closed-loop servo-positioning systems, optical servo systems present a clear alternative. Where dedicated systems employ an entire disk surface for reference-data positioning, and embedded servos use part of each disk surface, optical systems use none.

In an optical system, the head-positioner assembly is a separate modular unit mounted on the casting of the disk drive. Read/write heads are mounted on a carriage that's supported on a bearing structure on a shaft. A precision optical scale is attached to the head carriage. This scale, interposed between a lamp and photo transistor, consists of opaque and transparent areas in specific patterns. Density can range from [illegible] 2000 patterns (lines) per inch.

The position of the carriage is sensed with an optical position transducer, and carriage speed is controlled by a velocity signal derived electronically from the position signal. This servo system doesn't use any part of a disk drive's recording surface, yet achieves high accuracy (Fig 7).

TRADEOFFS

In selecting a disk drive with either an open-loop or a closed-loop positioning system, there are four main tradeoffs to consider: capacity, access time, cost, and reliability (see Table). These tradeoffs aren't all clear-cut; in a closed-loop system, for example, capacity and access time are often increased simultaneously, depending on the application.

One way to increase capacity is to space the recording tracks closer and closer together on a disk surface. Density of the most commonly used open-loop system is 345 tracks per inch; for closed-loop systems, 900-[illegible] tpi. This is a factor of nearly 3 to 1 between the two systems.

But users who specify the higher capacities also want a corresponding decrease in access time. The added capacity is invariably shared among multiple users of the same drive. Thus, if you double or triple a drive's capacity and its access time, it's not acceptable. To provide each user with the same, or faster, level of access requires a corresponding decrease in access time.

For example, a commercially available open-loop system with 15 Mbytes of capacity will have an access time of approximately 90 to 100 ms. In a multiuser environment, a 50-Mbyte drive—about three times the capacity of the open-loop drive—should have an average access time of 35 to 40 ms. This is exactly what a disk drive with a closed-loop positioning system offers. Of course, it will have a voice-coil magnet as a linear motor, will cost more, and will consume more power than an open-loop system. But it delivers higher capacity and faster access time.

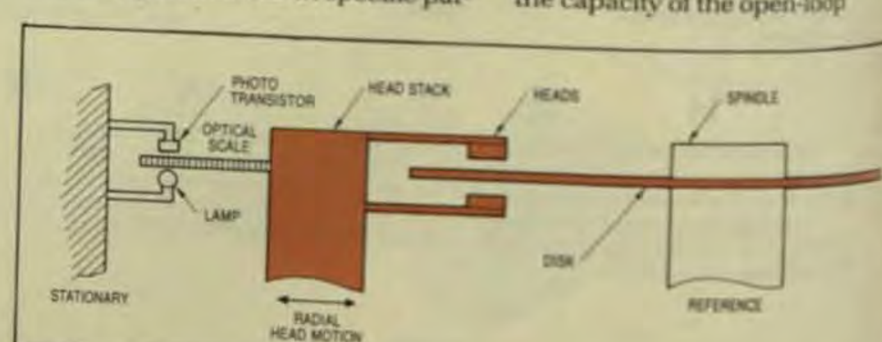

Fig 7 In an optical servo system, head positioning is accomplished with an optical scale mounted, typically, on the disk-drive chassis. No portion of the disk surface is used to store and retrieve the positioning data.

Only part of the closed-loop drive's higher cost is due to the positioning system, since other elements contributing to higher density can be more expensive. For example, higher track density means more expensive heads—whether a closed-loop or an open-loop drive is used.

Because the heads have higher resolution and smaller track widths, they're more costly to build. And since media must be certified to tighter specifications, it also commands a premium price. All of these costs are above and beyond the cost of the closed-loop positioning system.

In the borderline area, the two types of disk drives, capacities, and costs tend to overlap. For example, commercially available open-loop Winchester drives with a capacity of 20 Mbytes are within 10% of the cost of commercially available closed-loop 30-Mbyte drives.

Although capacities and costs may be about the same, access times are far apart. The closed-loop drive operates much faster—often three times faster. In larger capacities (50 Mbytes, for example) an open-loop drive would not be technically practical. On the other end of the spectrum, a closed-loop drive with 5, 10, or 15 Mbytes is feasible, but can't compete with the lower cost of the open-loop drive, even though it would provide much faster access time.

But users of 5-, 10-, and 15-Mbyte drives generally don't want faster access time. Typically, the application is for a single user, and 80 to 90 ms access time is quite adequate. When a system designer incorporates a higher-capacity drive, the application is for a multi-user environment. Each user in this type of system wants access in equal or better time than in a single-user system, so average access times of 30 to 40 ms are required.

Thus, the application dictates not only the capacity of a drive, but also its access time and cost. The positioning system then becomes a function of these three factors, plus a fourth—reliability (Fig 8).

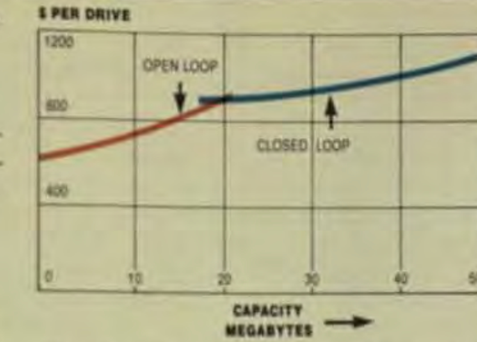

Fig 8 Open-loop drives are cost-effective at lower capacities and speeds, and closed-loop drives provide maximum value at higher speeds and densities.

Even though closed-loop systems are more complex mechanically and electronically than open-loop systems, they're designed as higher-performance drives and meet the same reliability requirements. The components used in high performance closed-loop systems have inherently higher reliability than their counterparts in an open-loop system.

A typical example is the stepper motor used in an open-loop system. It's not as reliable as the voice-coil linear motor used in a closed-loop system, but it meets the needs of the open-loop design, and the voice-coil magnet meets the higher requirements of the closed-loop drive. They both achieve the same level of reliability as measured in terms of mean time between failure (MTBF). Even though closed- loop systems have more complex electromechanical components, they're no less reliable than open-loop systems.

OTHER SELECTION ISSUES

Whichever positioning method is used, the choice of media can be critical. As tracks are spaced tighter, the media must be specified to tighter specifications. When tracks are bunched closer together, the track width and head-gap width become smaller and smaller. Signal levels therefore diminish. To compensate for this with standard oxide media, tighter parameters must be specified—but for a price. An alternative is selection of plated media, which has a much higher output-signal level.

The closer the tracks are packed, the more critical becomes the selection of heads to pick up the diminishing signal. The standard manganese zinc head performs well at densities in the 1000 tpi range, but beyond that range, new head technologies become necessary. The Whitney type of head, for instance, has the geometry of a thin-film head but has more conventional magnetic characteristics and is an attractive alternative when head geometry becomes critical.

A closed-loop drive generally consumes more power and dissipates more heat. Therefore, more care is needed when designing this type of drive into a system to eliminate hot spots. Also the large magnet in a closed-loop drive can affect other equipment around it. Therefore, the closed-loop drive must be packaged differently and better than an open-loop drive. The type of environment in which the disk drive will be used must also be considered. The closed-loop and open-loop systems operate equally well in all types of environmental conditions.

The majority of the Winchester disk drives commercially available today use a standard industry interface. In these drives, with capacities up to 50 Mbytes and beyond, data is transferred at a rate of 5.0 Mbps. But somewhere in the range below 100 Mbytes, a new interface will most likely be required in order to transfer data at a faster rate. The definition and parameters of this interface are now being worked out in the dynamics of the marketplace. Systems designers should expect to switch to this new interface in the near future for their very-high-capacity requirements. ■

Robert C. Abraham, *product marketing manager for Tandon Corp., was previously product manager and electrical engineer at Pertec Corp. He holds a BSEE from Univ. of California, Berkeley.*

Winchester Disk-Drive Market Trends

The overall driving force in rigid-disk drives today is the explosive growth of the office-automation and desktop-computer markets, and the resulting demand for the advantages of the Winchester drive in the same form factor as the 5.25-in. minifloppy already popular in these systems. This has made the 5.25-in. and sub-5.25-in. rigid-disk-drive field the fastest growing portion of the Winchester industry, and shock waves from its unprecedented growth are changing the entire market.

Intense competition among the manufacturers of small Winchester drives has caused marked price reduction, so these drives are now turning up in applications for which they were formerly too expensive. The result is a steady redefinition and broadening of the market as the technology matures. Three main trends can already be discerned.

Same Size/Higher Capacity

The first trend is the effort to pack more and more storage capacity into the popular full-size 5.25-in. form factor. This is the mainstream of development in the small Winchester market, and it's driven by the increasing popularity of small multi-user business systems. These systems demand both high storage capacity and fast access times, so that more users can be added without slowing down the access times. This combination of specifications can only be offered by Winchester drives that use closed-loop servo systems.

As a result of intensifying competition in the business-systems market, this part of the Winchester industry will be the major price battleground for the next couple of years, with most of the action taking place among Winchesters offering 30 to 100 + Mbytes of storage, with average access times around 30 ms. Storage capacities in this market may eventually go well beyond 300 Mbytes, as more of the advanced technologies originating in the larger formats are scaled down and mastered by manufacturers at this level.

It is this part of the 5.25-in. Winchester market that is having the most effect on the rest of the industry. As the capacities of these small Winchester drives edge upward, the makers of 8-in. drives are responding by increasing their products' capacities. This offers the upward compatibility needed to continue attracting high-capacity applications to the larger formats.

The capacities of 14-in. systems move upward for the same reasons. As a result, the storage capacities of 8-in. systems are edging into the 300 + Mbyte range, and those of 14-in. systems into the gigabyte range. Because of the competition from 5.25-in. Winchesters, *Disk Trends* sees production of 14-in. drives in the 30 to 100-Mbyte range peaking and beginning to decline this year, while 1984 is expected to be the peak year for 8-in. drives in this range.

Same Capacity/Smaller Size

The second main trend focuses more on the size of the package than on the storage capacity of the drive. This focus is the result of two factors:

- First, the open-loop technology used in the first 5.25-in. Winchesters has topped out at about 20 to 30 Mbytes of storage capacity. It's possible to build a larger-capacity open-loop drive, but it can't offer the faster access times demanded by users of higher-capacity drives, who are generally interested in multi-user applications.
- Second, the burgeoning portable and desktop computer market has created a demand for Winchester drives in a smaller package, as a step up from the slower, lower-capacity minifloppy. This application, which is almost entirely single-user oriented, demands neither astronomical storage capacities nor ultrafast access times.

These two factors have combined to direct open-loop technology into a search for smaller packaging and lower costs.

There are two areas of activity in this part of the Winchester market. The first involves the half-height Winchester, which has become very popular among manufacturers of small desktop and portable computers because it enables them to combine a floppy and a Winchester in the same space originally occupied by two floppies. Thus it offers up to 30 times the capacity and speedier response without the need to retool for a new cabinet.

Only one thing prevents half-height Winchesters from completely dominating the sub-30-Mbyte, 5.25-in. field, and that's their increasingly frequent use in portable systems. This usage exposes the Winchester mechanism—which is characterized by Swiss-watch-like accuracy and tolerances—to bumps and shocks, and changes in humidity and temperature.

The inclusion of safety features and shock mounts to overcome these problems raises the price, thus slowing down the market penetration of these drives. Still, it's certain that the half-height Winchester will be the largest part of the sub-30-Mbyte market within a few years.

The second area of activity in the low-capacity Winchester market is the sub-5.25-in. (micro-Winchester) market. The technology needed to compress a Winchester into a smaller package is well understood, and drives of this size are especially attractive to designers working with portable computers, electronic typewriters, and similar products. However, the susceptibility to shock and vibration of the sub-5.25-in. drives may be even more extreme than that of the half-height drives.

Whatever the eventual mix of full-size, half-height and micro-Winchesters turns out to be in the low-capacity market, the application of Winchester technology to this area has been so successful that the larger formats are being driven entirely out of the sub-30-Mbyte field. By 1985, *Disk Trends* estimates 5.25-in. and smaller Winchesters will completely prevail in this market.

Higher Capacity/Smaller Size

The third major trend is the search for ultrahigh capacity in a very small package. This is an area of intense technological activity, where all of the complex techniques originating with large-format Winchesters are brought to bear on the design. Closed-loop servo systems, plated media, vertical recording, and thin-film heads will all be used to develop micro-Winchesters with capacities that just a couple of years ago were found only in 14-in. drives.

Although portable computers and other space-sensitive systems are a factor in the growth of this field, a major impulse driving the development of ultrahigh-capacity micro-Winchesters is aesthetic. These products are attractive to designers because they make possible ever-smaller, sleeker, and more space-efficient designs.

Products in the 25- to 30-Mbtye area are already showing up, and may eventually go as high as 100 Mbytes. Since the micro-Winchesters are maintaining interface capability with their 5.25-in. brethren—just as microfloppy drives are doing with the 5.25-in. floppies—they'll mount a very credible threat to the larger format within a couple of years.

One other market trend of interest is the development of 5.25-in. and sub-5.25-in. removable-cartridge rigid-disk drives. At first glance, these drives seem to offer an attractive, random-access alternative to the use of tape drives for Winchester backup, at least in low-capacity applications. However, *Disk Trends* estimates that in 1985, only about 10% of rigid-disk drives will use a removable cartridge, though 77% of these will be in the 5.25-in. and smaller market.

There are a variety of reasons for this disappointing market performance. One is a lack of media standards, and the resulting problems with interchangeability. Related to this is the limit in capacity of any cartridge-type system. It's difficult to design a disk system that has both high information capacity and the ability to swap disks among machines. The eventual limit hasn't been reached yet, and may indeed be higher than many believe.

But the advent of vertical recording is rendering the standard 5.25-in. floppy disk capable of at least 10 Mbytes of storage, and probably more. Why deal with the expense of a rigid-disk system when a floppy drive will be able to offer nearly the same capacity, on a medium that will undoubtedly cost less than the rigid-disk cartridge? For higher capacity systems, the tape drive, with its low cost per megabyte, will continue to be the designer's choice. Finally, user uncertainty about the future of this product category, caused by awareness of all the above factors, is perhaps the strongest damper on this part of the market.

The rigid-disk cartridge will probably find its most enthusiastic acceptance in the high end of the home and personal computer field, where most users are unwilling to pay for both a disk drive and a tape drive. Until accepted by OEMs in general, the cartridge won't be a major factor in the rigid-disk-drive market.

All of these markets depend on technologies adapted from large-format Winchesters, but not all these technologies are immediately usable. The closed-loop servo system is well developed and widely used. Plated media has only recently reached the point where dependable mass-production techniques make it available at a reasonable price for 5.25-in. and sub-5.25-in. Winchesters.

The thin-film-head technology pioneered by IBM is still unable to offer a high enough manufacturing yield compared to the benefits it confers. Why pay five times the price of a standard manganese-zinc head when you only get twice the performance? Widespread use of this technology is still a couple of years off.

Oddly enough, considering its successful application to floppy disks, vertical recording is likely to be the most difficult technology to apply to Winchester drives, including the large-format drives. The higher bit density of vertical recording results in less output from the head. This is compensated for in floppy disks by the fact that the head rests directly on the disk.

In a Winchester drive, however, lower output means the head has to fly closer to the disk, and that requires both better air filtering to prevent particulate contamination, and closer tolerances in the manufacture of the media surface. This last-named reason, along with the difference in the thickness and composition of the media, makes the success of vertical recording in the Winchester field very dependent on the development of economical mass-production techniques for this exacting media.

Nonetheless, all of these technologies—and others not yet dreamed of—will eventually be successfully applied to the design problems of the small rigid-disk drive. The computer market's insatiable demand for more capabilities in smaller packages will ensure that this decade will see incredible growth and development in the small Winchester industry. ■

New Error-Management Techniques Increase Integrity of Disk Data

Enhanced error-detection and -correction techniques, combined with higher recording densities, improve data reliability in disk subsystems and reduce costs.

by **William A. Grace, Pete McLean, and C.M. Riggle,** Digital Equipment Corp.

Increases in data-storage capacities and the amount of I/O activity have put considerable pressure on the inherent data integrity of disk-storage subsystems. At the same time, and for the same reason, higher data integrity is required by the user. To cope with this increasing need, subsystem manufacturers have provided higher data integrity through better electromechanical design and error-detection, -correction, and -recovery techniques.

Error management has progressed from simple error detection via parity check to retry (reread with or without reseek), check sum and CRC, single-burst error correction, and such newer recovery techniques as automated track offset and detection-window offset. Interleaved Reed-Solomon error-correction codes (ECC) have provided improved correction of single burst errors and limited (nonindependent) correction of multiple errors.

Most recently, Digital has implemented a powerful version of noninterleaved Reed-Solomon ECC in its intelligent disk controller (UDA50) and mass-storage I/O server (MSC50). These devices are designed to work effectively with three high-capacity rack-mountable disk drives: the removable-media RA60 (205 Mbytes) and two Winchester fixed-media drives, the RA80 (121 Mbytes) and RA81 (456 Mbytes).

The data integrity of any mass-storage subsystem is the result of many hardware, software, and firmware functions in the controller and drive. Some of these functions check validity of control fields; others detect, correct, and recover from data errors. Many techniques applied in these controllers and disk drives are widely practiced in the mass-data-storage industry. Others have been selected to attain a higher level of integrity and performance for this architecture.

DIGITAL STORAGE ARCHITECTURE (DSA)

The main elements in any disk subsystem are a host-resident device driver, a controller, and one or more disk drives. A mass-storage architecture such as DSA describes the structure and functions of the subsystem that apply to any combination of hardware, software, and firmware. It divides disk-subsystem functions into layers and links the layers with standard interconnects (Fig 1). Commands and data are transmitted from layer to layer through the interconnects in the form of message packets defined by protocols.

Two protocols account for message transmission either way between the host and controller: mass-storage control protocol (MSCP) for generic I/O functions such as read and write, and diagnostics and utilities protocol (DUP). Messages between the controller and drive are described by the standard disk interface (SDI) protocol. Partitioning of the physical disk—an important element of integrity-oriented design—is described by Digital's standard disk format.

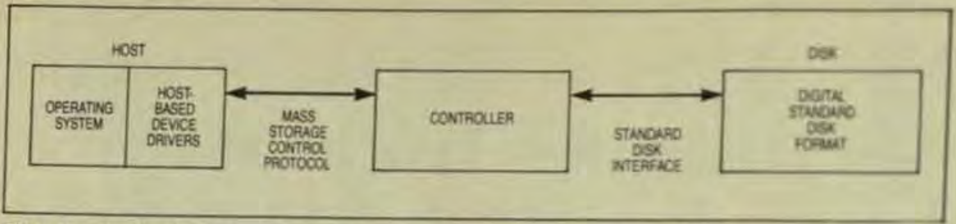

Fig 1 In a standard disk-based mass-storage subsystem, commands and data are transmitted in message packets defined by Digital Storage Architecture (DSA) protocols.

When the disk subsystem functions are divided into layers, the device driver, controller, and drive can each sense and correct error conditions with minimal interaction with the other elements. The device driver for disk storage can handle all message exchanges between an operating system and any DSA controller and disk drive. The device driver includes no error-handling software and no geometry data specific to a particular controller or drive.

A controller designed for this architecture can handle one or more DSA drives. The disk drive itself contains a parameter table of all its own characteristics, such as geometry and retry counts for error handling. This information is passed from the disk to the controller at startup.

To maintain data integrity both the disk drive and controller have embedded microprocessors appropriate to their positions in the subsystem hierarchy. Seek errors, for example, can be handled entirely by the drive, or recovered by the drive under direction of the controller.

Conversely, the controller itself is responsible for handling data errors. Data errors are the incorrect bits that can develop in data blocks between the time they are received from the host for writing on disk and the time they are read back from the disk for transmission to the host.

Together, the controller and disk are responsible for providing the host with "clean" data, for which all necessary error detection, correction, and recovery have already been done. Any hard errors—those unrecoverable due to major hardware failure in the controller or drive—are reported to the host in error logs.

SEEK ERRORS AND FAULTS

There are many integrity-oriented functions that are common to all Digital's DSA disk drives and to most removable and Winchester-type fixed drives on the market today.

For example, the microprocessor in a DSA disk drive monitors servo operation in order to detect misseeks on the fly. In addition, the header-compare process in the controller can indicate a misseek. If a misseek is identified, the respective micro in the controller initiates a specified number of retries (typically three), which are virtually always successful.

A quality-control check of the servo position is performed prior to a write operation. The magnetic head must be positioned within an allowable range of the track centerline, about 8% of center-to-center spacing, to avoid erasing adjacent tracks and to ensure a signal-to-noise ratio high enough to make subsequent read operation reliable. The drive's microprocessor monitors track position error and, if outside the allowable range, prevents initiation of the write operation and reports back to the controller. Again, the controller commands a specified number of retries.

Minimizing errors in read/write operations involves preventing the occurence of unacceptable electromechanical conditions in the drive or recognizing their occurence and counteracting them before they can cause any damage. When any of these faults is detected the controller is informed so that appropriate action can be taken. The system reponds to these conditions as follows:

- If more than one read/write head is selected because of failure in selector-bit logic, the microprocessor aborts the operation.
- If concurrent read and write signals are received, the drive is faulted.
- If ac power is lost, a circuit in the write amplifier clamps the write current to ground in order to avoid any writing in an uncontrolled state. However, the drive is able to finish writing the active sector prior to shutdown.
- If there is a transient power failure (no longer than 20 ms) enough energy is stored in the power supply to ride through this loss.
- If the head should be writing and isn't—or shouldn't be and is—the drive is faulted.
- If excessive heat is detected, the drive is faulted.
- In embedded-servo drives, fault circuits prevent accidental overwriting of the factory-written servo information and protect the drive and data.
- In dedicated-servo drives, patterns for generated index and guard-band data are selected to minimize the possibility of false seeking.

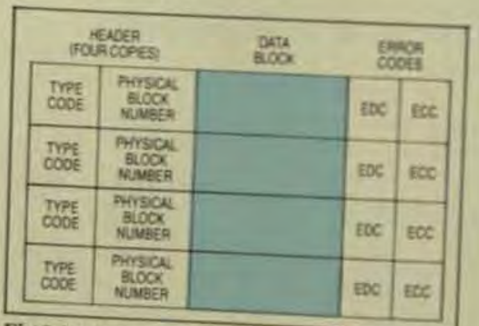

Fig 2 A sector consists of four copies of the header, a data block, and fields for error-correction and -detection code.

ERROR MANAGEMENT STRATEGIES

The recording medium in the user area of a DSA disk drive is divided into sectors (Fig 2). Each sector includes a data block (512 or 576 bytes), four copies of the block header (type code and physical block number) and fixed-bit fields for error-detection code (EDC) and error-correction code (ECC).

The disk controller uses error-correction and bad-block replacement techniques to detect, correct, and recover errors occurring in data blocks. Fig 3 highlights the effect of these techniques through the use of a Poisson curve that represents average distribution of errors (over many reads) among sectors in a recording medium.

Evidence to date suggests that the RA81 drive exhibits Poisson-like behavior. A very small number of sectors have a relatively large average number of errors, while most sectors have almost none. In Fig 3, the number of blocks appears in order of quality (not location in the recording medium) on the horizontal scale, left to right from worst to best.

Two types of errors may occur in a data block between a write to and read from the disk. Transient errors are induced by electrical, magnetic, or mechanical noise and occur more or less randomly. Errors due to media defects such as holes, inclusions, and thin spots may become permanent and recur in each read of a sector, or the sector may at least become more prone to error.

Physical deterioration causes the inevitable growth of physical defects throughout the life of the recording media. Since this phenomenon elevates the average number of errors per block (Fig 3), the controller must deal with a more intense error environment.

As the ECC circuitry in the disk controller detects and corrects errors in the user data blocks, it's possible to identify high-error blocks because of their repeated errors during production tests. All bad blocks—those having more than a threshold number of errors—are retired from service; the data is stored in replacement blocks; and all reads are revectored to replacement addresses. Later, as media deterioration takes place during normal field operation, blocks exceeding the specified error threshold continue to be retired.

The error-management strategy in the disk controller must be powerful enough to detect, correct, and recover all transient and fixed errors in user data blocks. And the bad-block replacement strategy must ensure that the user blocks continue to be available as a contiguous set of logical blocks of published capacity. The host's operating system will then see a fixed number of logical storage blocks throughout the useful life of the drive.

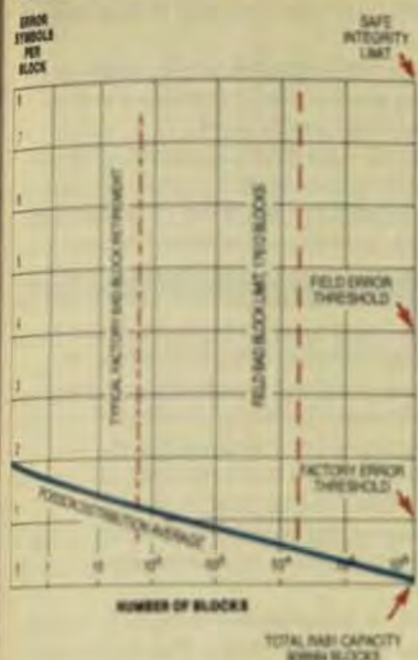

Fig 3 A Poisson curve represents the average distribution of errors among sectors in a recording medium. Evidence to date indicates that Digital's RA81 drive exhibits behavior of this type.

ECC/EDC

The ECC circuitry in DSA controllers generates a 170-bit ECC that is capable of correcting as many as eight independent 10-bit error bursts or one burst of up to 80 bits in any single data block. The bursts are aligned on fixed 10-bit boundaries so that worst-case single burst can be 71 bits long no matter where it happens to be positioned. In comparison, conventional processes typically generate a shorter ECC that ensures correction of only one error burst in each block, although some systems can correct multiple dependent errors.

Digital specifies that an unrecoverable error occur no more often than once in every 10^{13} bits read. However, analysis of DSA disk operation indicates even better performance. If electromechanical failure does result in one or more unrecoverable errors, they are reported to the host in an error log.

The EDC process—more limited than the ECC process—monitors hardware or firmware faults in portions of the controller. The CRC-type EDC circuitry guards against failure in the controller circuitry, but performs no data correction.

Both the 170-bit ECC and 16-bit EDC are calculated in the controller for each data block received from the host (Fig 4) and are added to each sector for writing to the disk. During a read operation ECC calculations are performed on the previously generated EDC and on the data block.

When a section is read from disk, error detection and correction proceed as follows:

- During read, a custom NMOS LSI chip (RSGEN) that handles the ECC process calculates a new ECC for the data block and stored EDC and adds the resultant modulo 2 to the ECC read from the disk for comparison. If the comparison results in all zeros, it's assumed that there are no errors in the data block or EDC.
- If the ECC comparison has a result other than zero and indicates the presence of error, computation is performed on the result to generate correction vectors, which are added modulo 2 to the data block and stored EDC. (The stored EDC may or may not have been corrected by use of the Reed-Solomon code.)
- A new EDC is calculated from the data block alone and compared with the stored EDC.
- If the comparison of new and stored EDCs shows that they're equal, the data block has been verified as correct. If the EDCs are not equal, the controller initiates the first retry for data recovery. A retry consists of another read operation, followed by the previous three steps. The disk specifies the number of retries at startup. Unless there is a hardware fault in the controller, it's

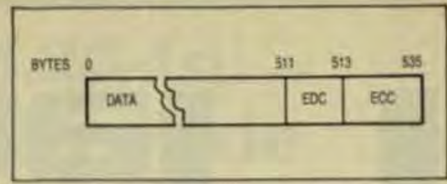

Fig 5 Each sector contains 512 bytes of data, 2 bytes of EDC and 22 bytes of ECC. During a read operation, both data and EDC bytes are used to generate the new ECC character, which is then compared to the ECC character read.

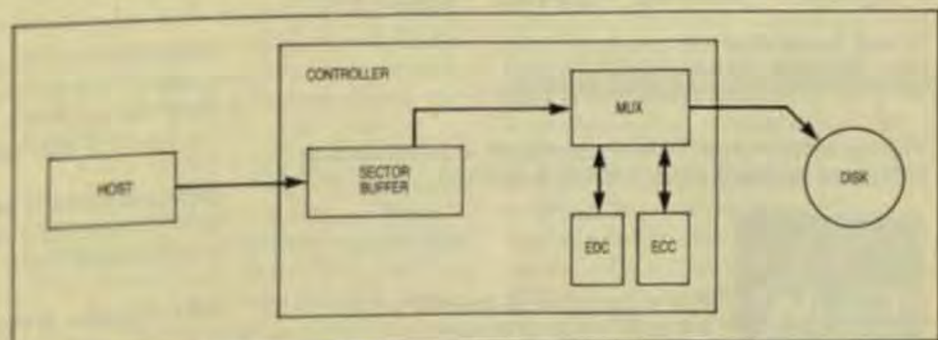

Fig 4 ECC and EDC are calculated in the controller for each data block received from the host, and added to each sector for writing to the disk.

HALF HIGH

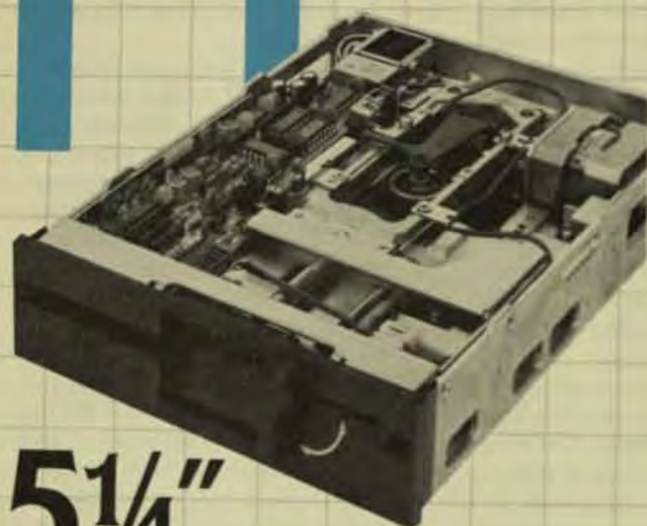

5¼″
Floppy Disk Drives
TEAC FD-55 Series

AT 1⅝″ HIGH, TEAC FD-55 SERIES 5¼″ FLOPPY DISK drives use half the space and run cool at half the power of conventional drives. High-reliability, low-noise brushless DC motors provide an MTBF of over 10,000 hours, backed by a one-year parts and labor warranty.

FD-55A	FD-55B	FD-55E	FD-55F
• 48tpi	• 48tpi	• 96tpi	• 96tpi
• 40 track	• 40 track	• 80 track	• 80 track
• 250KB	• 500KB	• 500KB	• 1MB
• single side	• double side	• single side	• double side

Power Requirements:
DC +12V ±5% 0.3A typical, 0.7A max.
DC + 5V ±5% 0.5A typical, 0.7A max.

Phone, write or wire TEAC Corporation of America for complete technical data, price and delivery.

TEAC Corporation of America
Industrial Products Division
7733 Telegraph Road
Montebello, CA 90640
213/726-0303

unlikely that EDC mismatches will induce these retries.

- If the retires are unsuccessful, the controller generates an error-log packet and transmits it over the communications bus to the host.

Normally, the corrected data block is passed to the user buffer in the host. The contents of the controller's sector buffers are parity checked for single-bit errors before the sector is passed along in either direction. If an error is detected, the controller requests retransmission from the data source, whether host or disk.

Digital's standard disk format includes logical data storage blocks and replacement blocks. The logical blocks are visible to application programs running on the host and are addressed by sequential block numbers from zero to the disk capacity. For example, the RA81 disk has 891,072 logical blocks (sectors) with 512 bytes of data per block, and 17,612 replacement blocks.

Replacement blocks are visible only to the controller and are addressed by means of revectoring to replacement block numbers.

As a logical block is determined to be bad, the data is transferred to a replacement block. The replacement address is written into a look-up table on the disk. To minimize revectoring time, the first replacement block is available on the same track as the bad block (primary revectoring). After that, replacement blocks from anywhere else on the disk can be assigned as needed (secondary revectoring).

Standard disk-interface protocol helps to protect data integrity in several other ways during transmission between the controller and drive. Check sums are calculated on the contents of an SDI command message and passed across the controller/drive interface in the message packet. In addition, the start of a message is identified by a preceding group of zeros and a synchronization character. An autocorrelation circuit provides for successful synchronization even if as many as three bits of the sync character are in error.

Although the operating system and host-resident driver aren't required to participate in maintainance of data integrity, the MSCP protocol supports several commands that the application programmer may use to introduce additional verification (perhaps when the stored data file is critical and there happens to be no backup for it). Two of these MSCP commands, read/compare and write/compare, can be written into user software to verify that what has been read or written on disk matches the copy stored in host memory ■

Pete McLean, *a consultant engineer with the Storage Systems organization, has been with Digital Equipment Corp. for 13 years. He holds a BSEE and an MSEE from the Univ. of Michigan.*

William Grace, *principal programmer for Digital for 5 years, has a BS from the Univ. of Wake Forest and a MSC from Lowell Univ.*

C.M. Riggle, *group manager, Storage Systems, has a BSEE from the Univ. of Minnesota. His technical specialties are systems and data integrity processes.*

Half-Height 5.25-In. Winchesters Provide Low-Cost Performance

Technology advances allow quantity manufacture of reliable half-height 5.25-in. Winchester drives to meet growing market demands.

by **Andrew Roman**, Cogito Systems Corp.

The driving force behind the dynamic 5.25-in. Winchester disk-drive market today is the incessant demand for smaller, more compact disk drives that offer high capacity, reliability, and performance at a lower cost. There is a continuous demand in the microcomputer industry for smaller, more compact physical chassis size. In 1981, half-height floppies appeared in the 8-in. form factor and in 1982 came the half-height geometric form factor for 5.25-in. floppies. This trend quickly evolved into a market demand for half-height 5.25-in. Winchesters.

Continuing new technology developments of disk-drive components have also resulted in the availability of pancake-shaped spindle motors; smaller more efficient, low-power stepper motors, and the new IBM 3380 high-resolution, low-profile read/write-head suspensions. In addition, advances in magnetic recording media have enabled higher track and bit densities and have resulted in higher storage capacity per surface at higher reliability than previously available.

Thus, today's half-height 5.25-in. Winchester disk drives offer the same or greater storage capacity as their full-height counterparts. A full-height drive measures 3.25-in. (H) × 5.75-in. (W) × 8.0-in. (D). The new half-height is only 1.625-in. high and models housing two platters can achieve 12.76 Mbytes of storage capacity per spindle.

One clear-cut advantage of using two platters intead of one to achieve 12.76 Mbytes is that this design provides complete plug-compatibility with the industry-standard ST412 interface. The areal density on Cogito's Model CG912, for example, is 345 tpi and 8783 bpi (about 3 Mbits/in.²) using coated ferrous oxide, which is today available from many independent sources. This density can easily be increased if you double the track pitch to 690 tpi to provide 25.5 Mbytes on the same two platters, and complete plug-compatibility is maintained in the same standard half-height 5.25-in. Winchester profile.

Other half-height 5.25-in. Winchester drives currently being produced offer 12.76 Mbytes on one platter, but close analysis reveals that these products have reached their capacity limit. In contrast, 2-platter half-height 5.25-in. models have taken only the first step with 12.76 Mbytes. And two platters offer a potential enhancement to higher storage capacities in the future and represent a family of 5.25-in. Winchester disk drives that maintain the same half-height profile and plug-compatible industry-standard ST412 interface.

THE PROMISE OF THIN-FILM MEDIA

Expected higher storage capacities packaged in a compact half-height 5.25-in. profile can only be achieved with higher areal densities. These can be achieved by use of thin-film media, which has the potential for 2000 tpi and 25,000 bpi. (Attaining these goals would yield an areal density of 50 Mbits/in.²).

There are two levels of sophistication in thin-film media. The first uses thin film the same way ferrous oxides are used now. The second exploits the ability of multiple deposits to create new crystalline magnetic-media structures to achieve ultra-high densities, as in vertical recording. Higher densities can be achieved with thin-film media because the metallic film is a continuous isotropic medium in which the molecules are packed closer and in a more orderly fashion, thus providing a thinner film without a loss of signal strength.

A reduction in film thickness, in turn, permits denser recording and faster flux reversals and results in higher signal strengths that can be switched at a higher data rate. But the fact remains that today there are only a few independent OEM suppliers of thin-film media, so it's unreasonable to expect several dozen drive manufacturers to depend on them.

Either sputtering or a wet chemical-plating process is used to produce thin-film media. Although control of such processes has been improved considerably, the stringent demands for high-quality plated media result in relatively low yields. Consequently, the cost of plated media is really higher than the cost of coated ferrous-oxide media, which is multisourced.

There are now more than six independent suppliers shipping standard 5.25-in. ferrous-oxide media in production volumes. This media has been in production for over four years and is a mature process; its quality and reliability have been improved and stabilized. The OEM sales price of 5.25-in. oxide-coated platters has, in turn, been reduced to approximately half of what they originally cost and their quality and reliability have increased many times over.

Two popular misconceptions today are that plated media needs to have an extremely hard surface to ensure resistance to head crashes, and that drives using such media are less sensitive to mechanical shock. In reality, the resiliency of ferrous-oxide media coated with a silicone lubricant is more likely not to cause a head crash when a particulate contaminant such as a piece of dust or dirt is lodged between the recording head and the disk.

During power-off periods a dedicated head-landing zone is provided for increased data integrity. This safety zone is located close to the hub, on the unused inner tracks of the platter. A mechanical constraint locks the actuator carriage during transport and secures the read/write heads to prevent them from moving across data tracks. A mechanical break on the spindle hub acts as a spindle-motor lock. A mechanical braking system is used to stop disk rotation within 3 s, minimizing wear between head and media during touchdown.

The future of perpendicular magnetic recording is based on switching magnetic fields in vertically oriented crystalline film structures, to achieve areal densities near 100 Mbits/in.[2]. Whether this density can be achieved depends heavily on advances in thin-film technology. Currently, at least two dozen independent media manufacturers are developing thin-film media processes. Half-height and full-height 5.25-in. Winchester disk-drive manufacturers are continually testing and monitoring their progress and assessing their availability to obtain production volumes.

ADVANCED HEAD SUSPENSIONS

IBM 3380 Whitney technology read/write heads are lightweight, low-mass, low-profile units. Available from several independent OEM head suppliers, they allow high-performance read/write head-and-suspension assemblies to be designed into half-height 5.25-in. Winchesters. In contrast, the first generation of 5.25-in. Winchesters, such as the Seagate ST412 introduced several years ago, still depends entirely on the older 3350-type read/write heads, which have a 3-rail, monolithic design.

Newer composite or monolithic manganese-zinc read/write heads are mounted on 3380-type assemblies and enable increased performance by permitting higher areal densities (on the order of 20,000 bpi and 2,000 tpi) and reduced head-setting times after track-seek operation. The low profile of the 3380 head-suspension assemblies permits a closer spacing between platters and thus allows the stacking of three platters in the 1.625-in. housing while maintaining the half-height 5.25-in. Winchesters.

Although the design and manufacture of 5.25-in. Winchester read/write heads are quite mature processes, one reason their costs have remained relatively high is the lack of universal standards. Thus, despite great similarities between the Winchester drives being offered, every manufacturer generates his own recording-head specifications to specify unique physical dimensions and characteristics such as number of turns in the coil, rail width, gap width, core width, and overall air-bearing surface design.

The development of thin-film heads was pioneered by IBM, which today uses them in all its high-performance Winchester disk drives including the 3370, 3375, and 3380. Primarily, IBM's objective was to enhance the efficiency of recording-head manufacture by eliminating much of the manual labor involved. Another goal was to create read/write heads capable of substantially higher performance in terms of enabling higher recording densities, higher bit-rate capability, and better signals at lower flying heights.

However, to exploit the full potential of thin-film media really calls for thin-film heads. Magnex Corp., San Jose, CA, was the first independent manufacturer to deliver thin-film heads specifically designed for the OEM disk-drive market. In 1981, Magnex began delivering evaluation units of thin-film heads based on its proprietary, internally developed, dry-sputtering process.

For the new generation of half-height 5.25-in. Winchester products, Magnex is the only head supplier currently delivering thin-film heads with the higher number of turns needed to produce the higher signal amplitude required to overcome the lower track velocities of 5.25-in. disks. The availability and implementation of thin-film heads will allow potentially higher disk-storage capacities.

BUFFERED SEEK STEPPER

Open-loop controlled actuators provide 50 to 100 ms access time (including settling time). In these devices a low-cost stepper motor operates in a buffered mode to drive a split-band actuator. Such steppers have been field-proven and are available from multiple sources, again permitting high-volume manufacturing.

The stepper-motor/split-band-actuator positioning mechanism will typically allow 85 ms average access

time, which is more than adequate for single-user applications. Today's 12.76-Mbyte half-height 5.25-in. Winchester is typically found in single-user personal computers used for word processing, electronic spreadsheet analysis, and general small-business system applications.

The control of the stepper motor was originally handled by the drive controller, which issued individual pulses to move the rotor one position at a time. Buffered seek operation implies that the seek argument is transmitted to the drive as a burst of pulses and stored. From that point on, stepper motor control is handled locally by a microprocessor in the disk drive. Rather than stop the rotor after each pulse, this intelligence produces a continuous and more efficient motion of the rotor. The result is that access speeds are significantly enhanced; step rates are 0.6 ms per step rather than 3 ms per step.

To increase the storage capacity by increasing the track density above 345 tpi might require a more precise positioning system, such as a closed-loop servo control, but today's multisourced stepper motors provide reliability while enabling the drive manufacturer to achieve high-volume manufacturing to meet the huge market demand.

ADVANTAGES OF HALF-HEIGHT WINCHESTERS

Market responses from several customer surveys have repeatedly shown that disk-drive size reduction by an integral multiple such as half-height design is particular desirable. Several OEM attempts to offer ⅔-height drives have been discontinued. But it may not be advantageous to reduce the platter diameter below 5.25-in. because these reductions make the outside dimensions of the target drive insignificant.

It's even risky to depart from the industry-standard multisourced 5.25-in. media diameter because there are as yet no established standards for sub-5.25-in. media. There are several reasons: multiple independent OEM suppliers for the media don't exist, product cost is high, and a manufacturer who depends on media with a nonstandard diameter may never be able to produce in high volumes.

Half-height 5.25-in. Winchesters allow two drives to be mounted in the same 3.25-in. panel height that previously held only one device. Ideally, the other half of the 3.25-in.-high slot could be used to mount a half-height 5.25-in. floppy that features the removable media needed for program or data loading, for portability, and for backup (Fig 1).

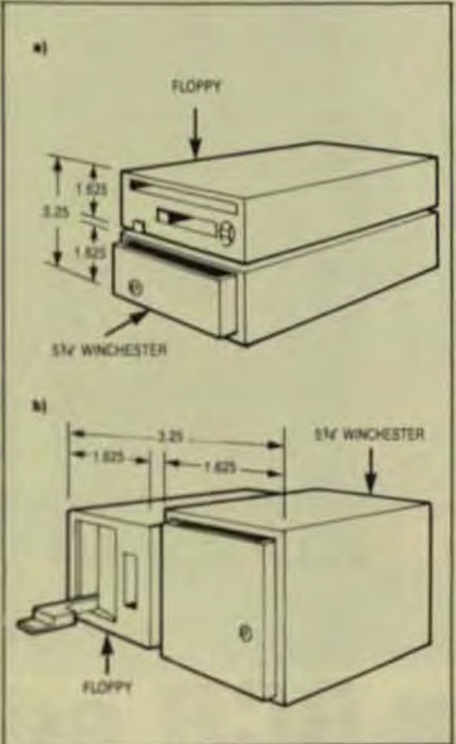

Fig 1 The half-height 5.25-in. Winchester can be mounted in either a horizontal (**a**) or vertical (**b**) orientation with an accompanying half-height 5.25-in. floppy-disk drive.

PRICING AND MARKET OUTLOOK

Half-height 5.25-in. Winchesters incorporate many new technologies and component advances. The parts count of the total number of mechanical and electrical components in the disk drive, however, hasn't been cut in half. Neither have the number of steps in assembly and test operation. In fact, the difficulty of assembly and test may even have increased because the parts are smaller.

Additionally, the smaller size results in reduced clearances, so the dimensional tolerances of the parts have to be tighter. But expenditures required for the capital equipment used for tooling, testing, burn-in, and in production-line automation still represent a sizable investment. A good case can obviously be made for a premium price on these compact

Fig 2 The market forecast of 5.25-in. Winchester disk-drive shipments graphically illustrates the significant penetration of half-height shipments for capacities up to 25 Mbytes.

drives, but actually, technology advances have lowered the unit sales price—though they haven't cut it in half.

So far, seven manufacturers have announced more than 10 different half-height 5.25-in. Winchester models, and several more announcements are expected, but last year, no production units were actually shipped (Fig 2).

Because 1983 is a formative year, slightly less than 5% of the OEM shipments this year will be half-height 5.25-in. Winchesters. However, this trend will increase dramatically, and by 1986, 66% of all 5.25-in. Winchesters shipped will be half-height drives. The average sales price per unit will be higher because of the projected increase in storage capacities and unit performance.

IBM's recent announcement of its PC XT personal computer, which features a 5.25-in. Winchester disk drive with 10 Mbytes (formatted) on 2 platters, virtually endorsed this capacity as the optimum disk drive for the personal computer. All the PC compatibles—and especially the new portable PC emulators being developed by other manufacturers—are offering half-height 5.25-in. floppies. Upgrades to 10-Mbyte half-height 5.25-in. Winchester are imminent. ■

Andrew Roman, *vice president of marketing and sales for Cogito Systems Corp., previously served in various marketing and engineering management positions with disk drive manufacturers such as Diablo Systems, Control Data, Pertec and NCR.*

Model 6253 Tape Subsystem

Telex GCR adds up to over 9,000 units shipped.

Today, more leading computer systems houses are specifying Telex GCR Tape Subsystems. The facts add up. With more than 9,000 units shipped (most for high speed 125-ips operation with tri-density capability), systems houses have discovered that Telex provides a design, manufacturing and quality maturity that is unmatched in the marketplace today.

Telex GCR Tape Subsystems are the most proven, high-performance rack-mountable GCR units available. Compatibility proven by five CPU manufacturers. Versatility proven in system house interfaces to sixteen different CPU's. Field proven in demanding seismic operations and office environments.

And their range is impressive. The Telex 6253 family is comprised of six models of tape drives, available in speeds of 125-, 75- and 50-ips, each with tri-density (6250/1600/800 bpi) or dual density (6250/1600 bpi) capabilities. The family also includes the model 6850 tri-density formatter. In addition, Telex engineers will provide experienced assistance in making Telex subsystems enhance your high-performance computer systems.

Start with proven capability. Specify Telex GCR. For more information, contact the nearest Telex OEM Sales Office listed or phone our OEM Marketing Department in Tulsa: (918) 627-1111. • Amherst, NH (603) 673-9272 • Southfield, MI (313) 352-2720 • Garden Grove, CA (714) 898-9833 • Houston, TX (713) 497-6770

The innovation continues... TELEX

Telex Computer Products, Inc.
Terminals/Peripherals/Systems/OEM Products
6422 East 41st/Tulsa, Oklahoma 74135

Magnetic Tape Systems

IBM Compatible 3200 bpi Tape...the new standard for high capacity disk backup.

CacheTape™ Improves System Performance—with 92 MB of Winchester Backup.

3200 bpi— A New Tape Standard

With the announcement by many computer manufacturers of the availability of 1600/3200 bpi tape drives, the primary requirement (a large user base) for standardization exists. Why 3200 bpi? The choices in tape density until now have been 1600 bpi tape (46 MB) or expensive 6250 bpi tape (180 MB). The new choice is 3200 bpi tape with 92 MB capacity at a low incremental price ($375) to the standard 1600 bpi tape drive. Substantially less expensive than 6250 bpi with better performance in most applications, Cipher's CacheTape also offers the added benefit of standardization.

Start/Stop Performance

CacheTape provides superior performance versus tension arm, vacuum column, 100/25 ips streaming, and 50 ips 6250 bpi (GCR) tape drives. As an example, the following benchmark comparison provides typical performance data for a file-oriented backup application:

	Benchmark Time*
Streaming Tape (variable speed, 1600 bpi)	23 min
Vacuum Column (125 ips, 1600 bpi)	7 min
CacheTape Model 891 (1600 bpi)	**7 min.**
50 ips 6250 bpi (GCR)	6.4 min. (calculated)
CacheTape Model 891 (3200 bpi)	**5.9 min.**

*Benchmark measured on a DEC PDP-11/34 under RSTS™ for available tape drives.

**OEM Quantities

RSTS is a registered trademark of Digital Equipment Corporation.

UNIX is a trademark of Bell Laboratories.

For higher capacity disk backup, 3200 bpi density offers 184 MB of disk backup with only one reel change.

CacheTape offers field-proven streaming mechanics, fully automatic loading and threading, and compact package size... and still performs disk backup and transactional applications as well. CacheTape is the total solution to your tape drive needs.

Call or write for a free benchmark brochure that explains the performance advantages and how to calculate in advance the benefits of CacheTape.

Tape Adapter Compatibility

CacheTape easily interfaces and operates with industry-standard tape adapters. CacheTape is completely interface compatible with existing couplers for products from DEC, DG, and TI, and couplers for Multibus, S-100, and other popular mini- and microcomputers. Take advantage of CacheTape's easy integration features and increase your system performance while eliminating extra time and expense from your budget. With CacheTape, you can use your current controller investment wisely and effectively... plug in CacheTape for immediate benefits.

Software Transparent

Cipher's CacheTape products are completely software transparent with current vacuum column or tension arm tape software. CacheTape provides start/stop tape performance for tape applications such as file-oriented disk backup, transactional journaling, tape sort/merge and data acquisition. Utilization of a cache memory in the tape drive means that CacheTape can provide higher performance than existing tension arm or vacuum column tape drives at much less cost. Just plug CacheTape into your system now... and benefit from total software compatibility.

Up to 40% Less Cost

CacheTape Model 891 (with 1600 bpi)	$3425**
CacheTape Model 891 (with 1600/3200 bpi)	$3795**
versus 125 ips vacuum column	$610[illegible]
versus 50 ips 6250 bpi (GCR)	$700[illegible]

UNIX Friendly

UNIX™—the emerging operating system of the 80s—now has with CacheTape an easily integrated, low cost, standard tape drive with superior performance. Do other alternatives make sense anymore... particularly if software development time and resources are scarce?

CacheTape with 3200 bpi... the new standard

• Up to 40% less cost • Tape adapter compatibility • Software transparent • Superior performance • UNIX friendly

Available Now

Cipher Goes Beyond

CacheTape

cipher data products, inc.

10225 Willow Creek Road, P.O. Box 85170, San Diego, California 92138
Telephone: (619) 578-9100, TWX: 910-335-1251

CIPHER DATA PRODUCTS (UK) LTD.
Camberley, Surrey, England
Phone: 0276-682912
858329

CIPHER DATA PRODUCTS S.A.R.L.
Paris, France
Telephone: (1) 668 87 87
Telex: 203935

CIPHER DATA PRODUCTS GmbH
Munich, West Germany
Telephone: (089) 807001/02
Telex: 521-4094

Streaming Tape Drives May Not Be Best Choice For Winchester Backup

Streaming tape has become synonymous with Winchester disk backup. But if file management is also needed, start/stop drives turn out to be a better choice.

by Kenji Eigner, Data Electronics Inc.

Start/stop and streaming tape drives, the two basic types of ¼-in. technology, have developed over the last several years for small-business computer systems, data-acquisition systems, telecommunications systems, industrial process-control systems, and word-processing systems. Start/stop drives are the classic full-function tape peripheral that can readily address records and manipulate stored data. These drives have been in wide use since the early 1970s. Streaming was introduced to ¼-in. digital cartridge tape drives so that it could be used to back up a small Winchester disk drive efficiently.

But streamers are not the answer for all applications. If mirror-image backup is all that's needed, streaming is the proper choice. But if file management is the desired application, start/stop drives can handle both Winchester disk backup and file management.

COMPARING THE TECHNOLOGIES

Both drives use the same ANSI/ECMA standard cartridge, which actually is the tape drive. All of the tape's tension control and guidance resides in the cartridge (Fig 1). An isoelastic band couples both reels of tape to a central drive wheel. The tape drive merely provides the motion control to the central drive wheel in the cartridge. However, the tape drive also contains the read/write recording head, motion-control electronics, position-sensing circuitry, motor, and sometimes the formatter/controller and codec.

In a start/stop drive, the motion-control electronics are very sophisticated. Its servo circuits must start or stop the drive in the interrecord gap prescribed by the ANSI standard for ¼-in. cartridge recording. A streaming drive, on the other hand, doesn't require interrecord gaps and therefore has more of the tape actually available for data storage. This increased area for data storage on the tape is one of the major advantages of the streaming format. The differences in the data format on the tape are shown in Fig 2.

The other major advantage of the streaming format is the speed at which an entire cartridge can be recorded. Whereas the start/stop drive typically records at less than 40 ips, streamers can record at up to 90 ips. Thus, a fast streamer can record the same data at more than twice the speed of a start/stop drive.

Unfortunately, all is not as simple as it appears. Data is commonly stored on a Winchester disk in predefined block lengths. As block lengths become shorter, the number of interruptions that the cartridge drive sees in the available data increase. When data to a streamer is interrupted for longer than it takes to fill the input buffer, the streamer stops. The same is true for a start/stop drive. When data flow resumes, a start/stop drive immediately ramps up and begins recording. The streamer, however, must go through a

THE 5¼" STREAMER OTHERS WILL HAVE TO MEASURE DOWN TO.

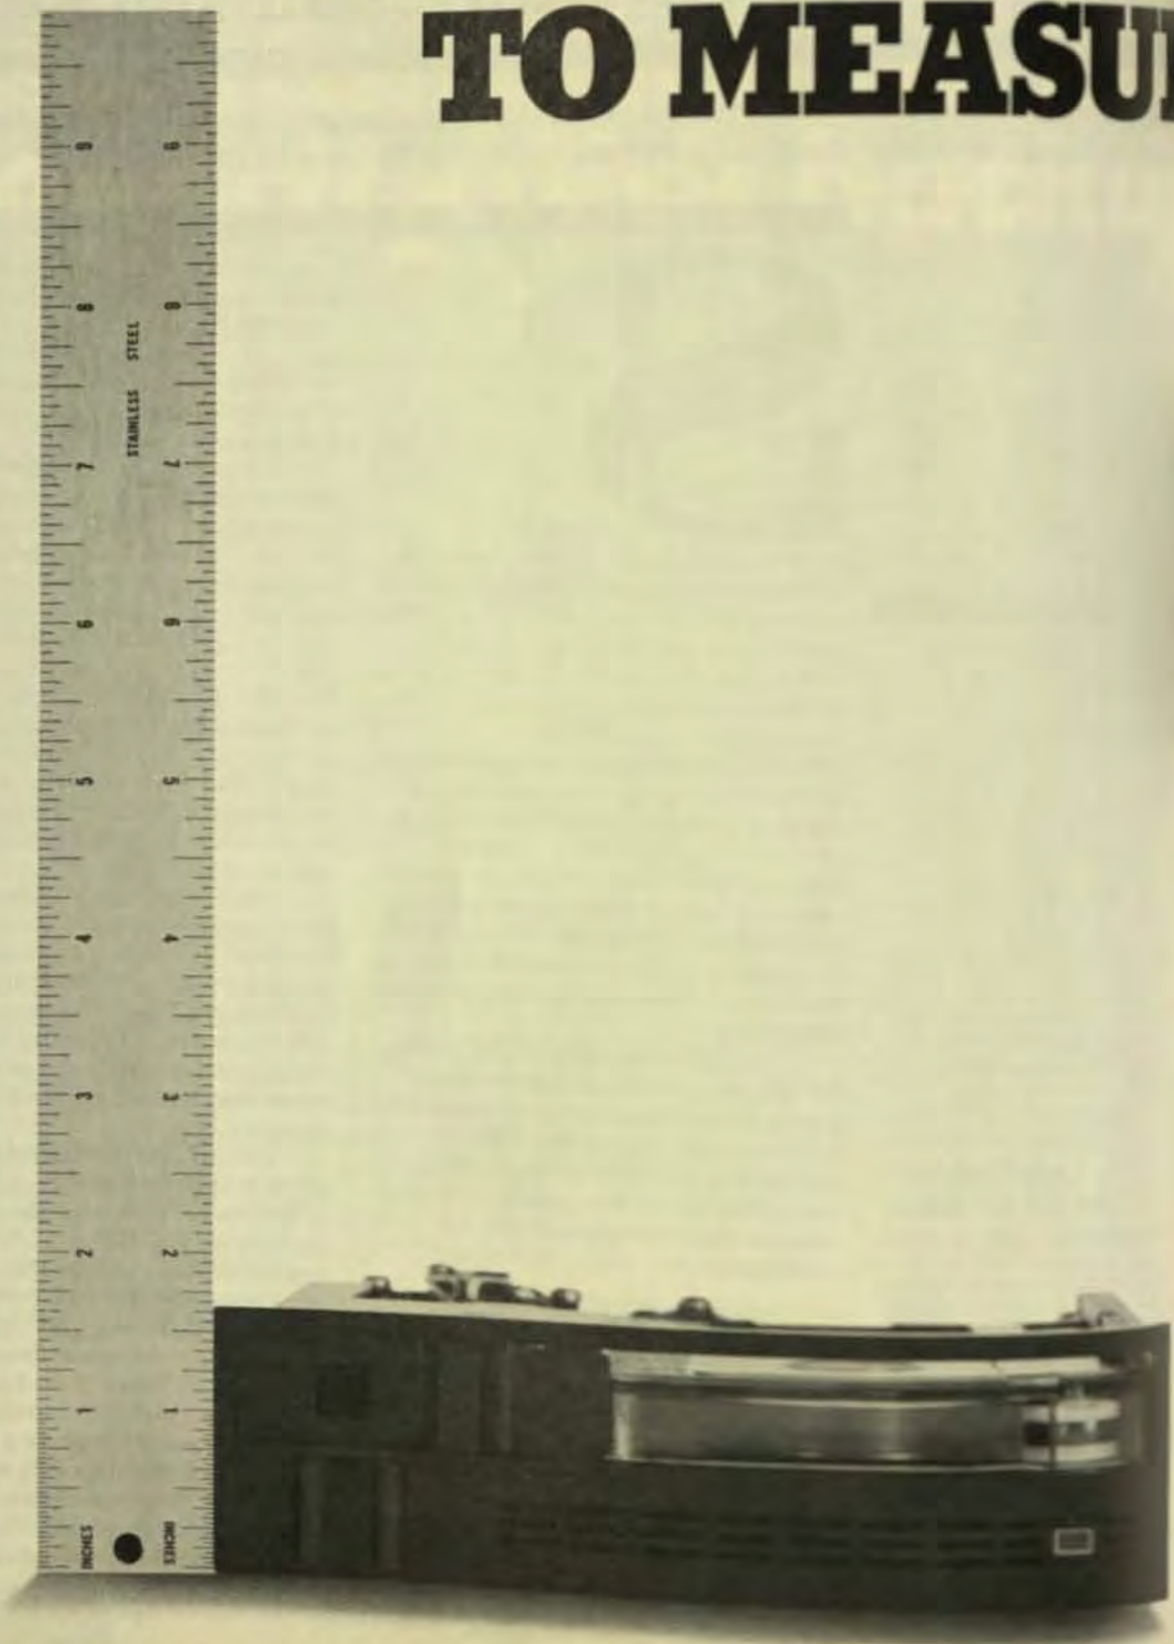

Introducing the Scorpion™ half-height 5¼" mini streamer.

Once again, Archive has the Winchester back-up others will have to measure up to. Only this time, they also have to measure down. Because the Scorpion is only 1.625" high.

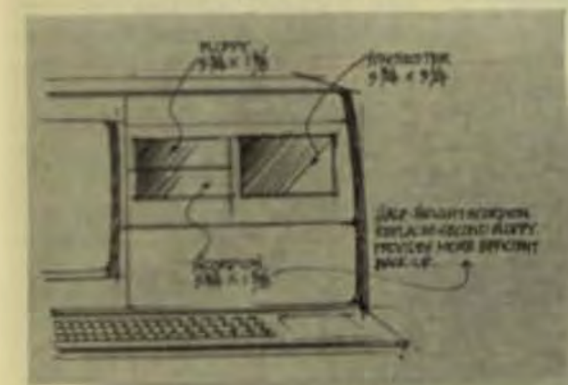

It fits the footprint of half-height floppy disk drives. So no change in existing cabinetry is necessary. Yet it still gives you up to 60MB of storage.

Now system designers can choose from a wider variety of storage combinations. For example, half-height Scorpion, half-height floppy and 5¼" Winchester.

And our standards meet your standards. The standard Archive (QIC-02) interface and QIC-24 format are both incorporated in the Scorpion.

But that's only part of the story. Thanks to Archive's advanced LSI technology, the half-height Scorpion gives you higher capacity and more functions than many 8" streaming tape drives. Yet it has fewer parts, uses less power and is competitively priced.

For complete details and a free copy of our 176-page streaming handbook, return the coupon. Or call Archive today.

*Scorpion is a trademark of Archive Corporation.

FREE!
STREAMING TAPE HANDBOOK.

To receive your copy, return this coupon to
Archive Corporation
3540 Cadillac Ave.
Costa Mesa, CA 92626
Attn: Marketing Dept. or call
(714) 641-0279

Name________
Title________
Company________
Address________
City________ State____ Zip____
Telephone (____)________

☐ Send me my free book.
☐ Send me additional information about Archive ¼" Streaming Tape Drives.
☐ Send an Archive sales representative.

Telex: 4722063 Domestic Telex: 183561

ARCHIVE
CORPORATION

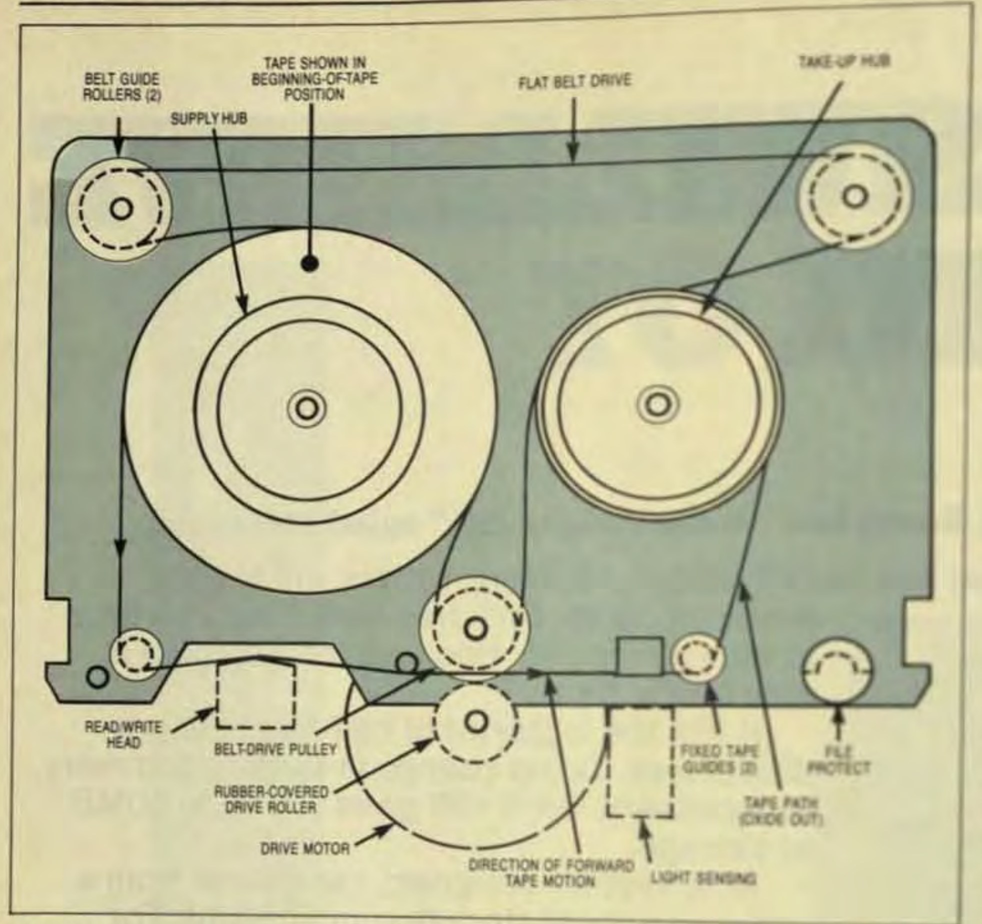

Fig 1 The ANSI-standard ¼-in. digital tape cartridge uses an isoelastic band as a drive belt to couple the reels of tape to the drive pulley, which is turned by the drive motor. The drive belt not only moves the tape but also maintains tape tension and minimizes instantaneous speed variation.

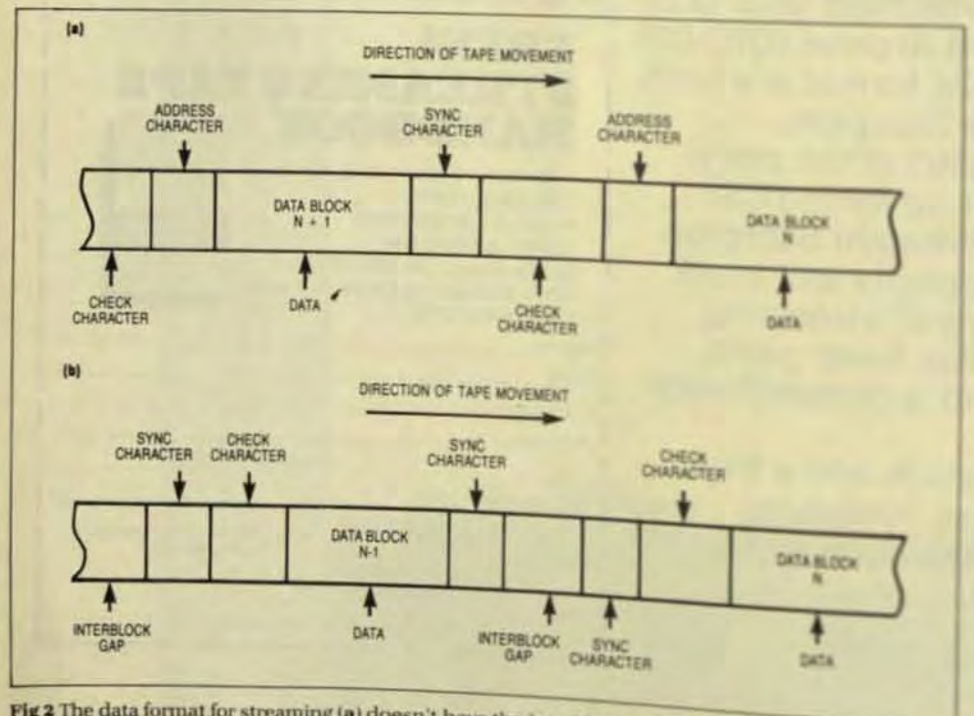

Fig 2 The data format for streaming **(a)** doesn't have the interblock gap that is seen in the data format for start/stop operation **(b)**. This design allows for more efficient tape utilization, and therefore higher capacity in streamers for a given length of tape.

more complex startup proced
fore it can begin recording the
data.

When data to a streaming
stops, the drive speed is ramped
down to zero. When data starts
the drive must reposition the
and be back up to speed at a
the tape adjacent to where the
block of data was recorded. To
reposition itself, the drive shifts
reverse and ramps its speed
and then down to zero at a point
prior to the end of the preceding
block. Then it ramps back up
speed in the forward direction
it's at full speed when it reaches
point adjacent to the end of the
er block.

This procedure assures
tape is wasted but takes a long
Each ramp, and there are four
about 300 ms for a 90-ips stream
100 ms for a 30-ips streamer. The
sulting intertransmission-block
is from 400 to 1200 ms for stream
Start/stop drives stop and start
ms, so their intertransmission
delay is only 50 ms. It's clear
given capacity, smaller transm
block sizes require more starts
stops. And it can be shown that
point is even reached where the
start/stop drive will record the
from the Winchester disk faster
the streamer (Fig 3).

LOOKING AT THE HEADS

The most expensive item in any
drive is the recording head. The
four basic head configurations
today and variations on each. M
commonly used is the 4-track
serpentine head used in indus
standard start/stop drives (Fig 4
configuration uses the middle
tape for tracks 1 and 2, and has
rate erase head for each track
over 100,000 high-density dr
operating in the field today
heads of this type and their
tion for high reliability is well

The 4-track serpentine he
fers from the nonserpentine
in that it records bidirectional
is, alternate tracks are written
read in opposite directions
ing the need for the tape to be
wound between tracks and
ducing the time required to
data onto the cartridge. This
head is useful in both start
streaming tape drives.

TRIPLE DEC'er.

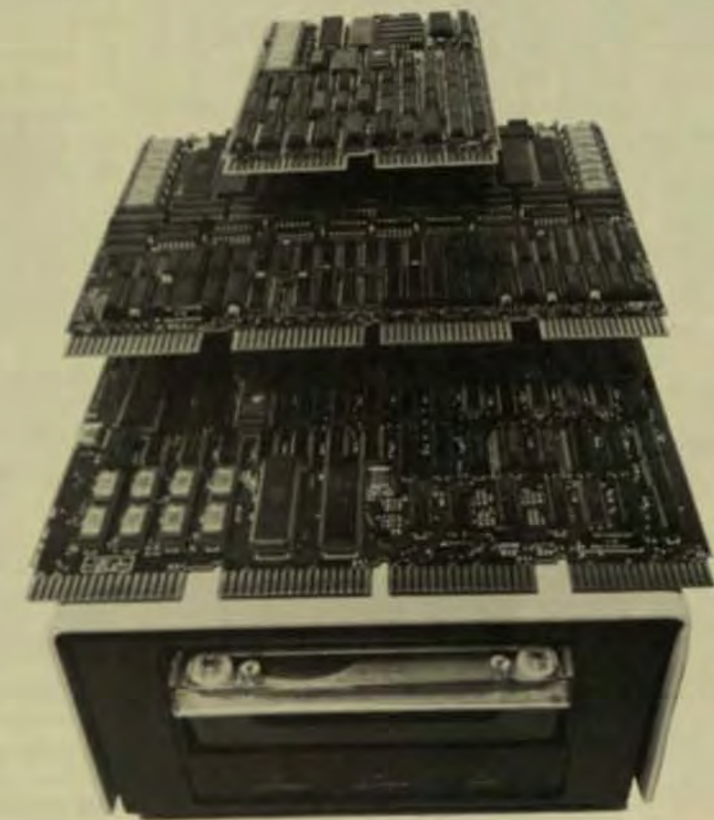

More Muscle, Less Money, For DEC Users
Now Alloy presents three new capabilities that greatly expand the usefulness of your LSI-11, PDP-11 or VAX computers.

CP/M For PDP-11, VAX
Our UCP-11 simply plugs into your PDP-11 or VAX Unibus to turn workstations into workhorses, performing 8 or 16-bit CP/M programs in parallel with programs running on your DEC processor which is called upon only to perform I/O activities. Result: You get to run any of thousands of CP/M programs with greater performance than is available with most CP/M machines.

CP/M For LSI-11's
Our QCP-11 is functionally similar to the UCP-11, but it plugs into the LSI-11's Q-Bus. Again, you get to select from a great variety of CP/M programs, plus high performance. Both the UCP-11 and the QCP-11 operate in parallel with the DEC processor, with no degradation in performance. What's more, you get direct access to data files, so you can work freely with either DEC or CP/M file formats.

Streaming Backup For LSI-11's
Alloy's new LSI-50 system combines our LSI-50 controller with CDC's remarkable Sentinel streaming tape cartridge drive to provide disk backup and data storage. This isn't your typical cartridge drive; the LSI-50 gives you 44 megabytes of storage, 2Mb/Min. transfer rate, file oriented backup/restore operation and error transparency. The LSI-50 is available both as a board and as a complete subsystem with the Sentinel.

All three of these units are exceptional in price as well as performance. For more information, call or write Information Services, Alloy Computer Products, 12 Mercer Rd., Natick, MA 01760. (617) 655-3900. TWX: 710-346-0394.

ALLOY
The Matchmakers

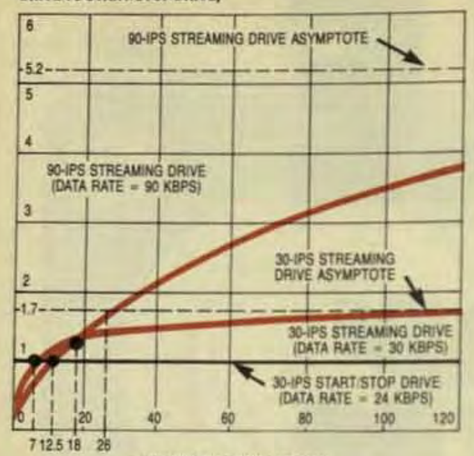

Fig 3 The throughputs of various classes of streaming and start/stop drives are given as a function of transmission-block size. A 30-ips start/stop drive is faster than a 90-ips streamer for block sizes of less than 12.5 Kbytes.

The 2-track serpentine head uses an erase bar and is often used for applications that require more than two tracks. Many of the new streaming drives use this head in a moving head configuration. The main disadvantage of this type of head is that the erase head isn't track-selective and doesn't always precede the write head.

The fourth type of head is of the serpentine variety, but contains more than four tracks. An example of this is the 7-track serpentine head that Data Electronics Inc. has used for several years on a military drive and is now using in the recently introduced 50-Mbyte start/stop drive. From a reliability standpoint, it's preferable to select a product with a fixed recording head rather than a moving head because each moving part is subject to potential failure. But some of the head movers are reliable enough that this isn't an issue.

SYSTEM INTEGRATION COSTS

It would appear that streaming is always a lower cost technology because of the lower base cost of the drive itself. Do not be deceived. The total cost of system integration—which adds the cost of the controller and host interface to the basic drive cost—is the figure that matters. An analysis of system integration costs for start/stop and streaming systems shows that they cost very nearly the same.

The basic start/stop drive costs about $200 more than a comparable streaming drive, but the controller and interface for the streamer cost about $200 more than for the start/stop drive. Also, controllers for start/stop drives are more widely available. One reason for this difference in controller cost relates to the additional buffer storage and DMA functions required to maintain streaming. Another is the difficulty of error correction in a streaming drive.

When an error is detected by the read-after-write head in a start/stop drive, an error-retry strategy is implemented. That is, the drive stops and rewrites the data on tape. If, upon rereading it, an error is still detected, the procedure is run through again, up to a certain number of retries. This usually results in the error being corrected.

With a streaming drive, an error-retry strategy can't be used because the drive keeps going continuously. Therefore, the controller must maintain enough memory space to remember any tape location at which an error has been detected. After noting the location, the data is rewritten at the end of the data stream. If an error is still noted, the procedure is repeated again. When it reads back the tape, the controller must remember which locations had erroneous readings, and must substitute the rewritten data from the end of the data stream. This is a far more complex controller function than start/stop drives require.

Until now, there has been a lack of standardization in streaming-drive interfaces, while start/stop drives enjoyed a common de facto standard—the DEI Funnel interface. However, there has been a significant move in the last year to standardize the streaming interface, and the industry's QIC (Quarter-Inch Cartridge) Working Group has developed a proposed standard interface for streaming tape drives. Known as the interface, this proposed standard been adopted by several and products that incorporate been introduced.

The next step for the QIC ing Group must be a proposed format standard that will interchangeability of cartridges corded on any manufacturer's A usable format standard is still a ways off, but the QIC mat is a step in that direction though it doesn't completely interchangeability.

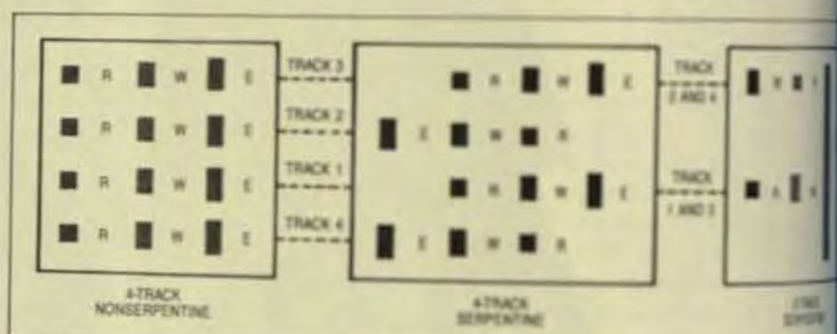

Fig 4 The three most common head configurations illustrate the difference between and nonserpentine recording heads. The 2-track head is often used as part of a tem with more than two tracks.

APPLICATIONS

While backup has emerged as jor application for ¼-in digital tridge tape drives, it is but their many possible uses. The of applications can be divided three basic categories: backup, I/O, and initialization.

Backup is a disk save function used for archival protection of invaluable data operator errors; to provide able media for data interchange other systems; to provide a data-storage medium that can

APPLICATION MODES	
APPLICATION	STREAMING
*DISK BINARY DUMP (MIRROR IMAGE)	++
*DISK RESTORE (MIRROR IMAGE)	++
SOFTWARE DISTRIBUTION/PROGRAM LOAD	+
*SELECTED FILE SAVE	+
*SELECTED FILE RESTORE	+
ARCHIVAL STORAGE	+
*TRANSACTION RECORDING/AUDIT TRAIL	
*DATA ACQUISITION	
TAPE-FILE RESTRUCTURING/EDITING	
ON-LINE DATA PROCESSING	

*BACKUP + GOOD CHOICE

NOW—GCR AND PE STREAMERS THAT REALLY DRIVE DOWN THE COST OF OWNERSHIP.

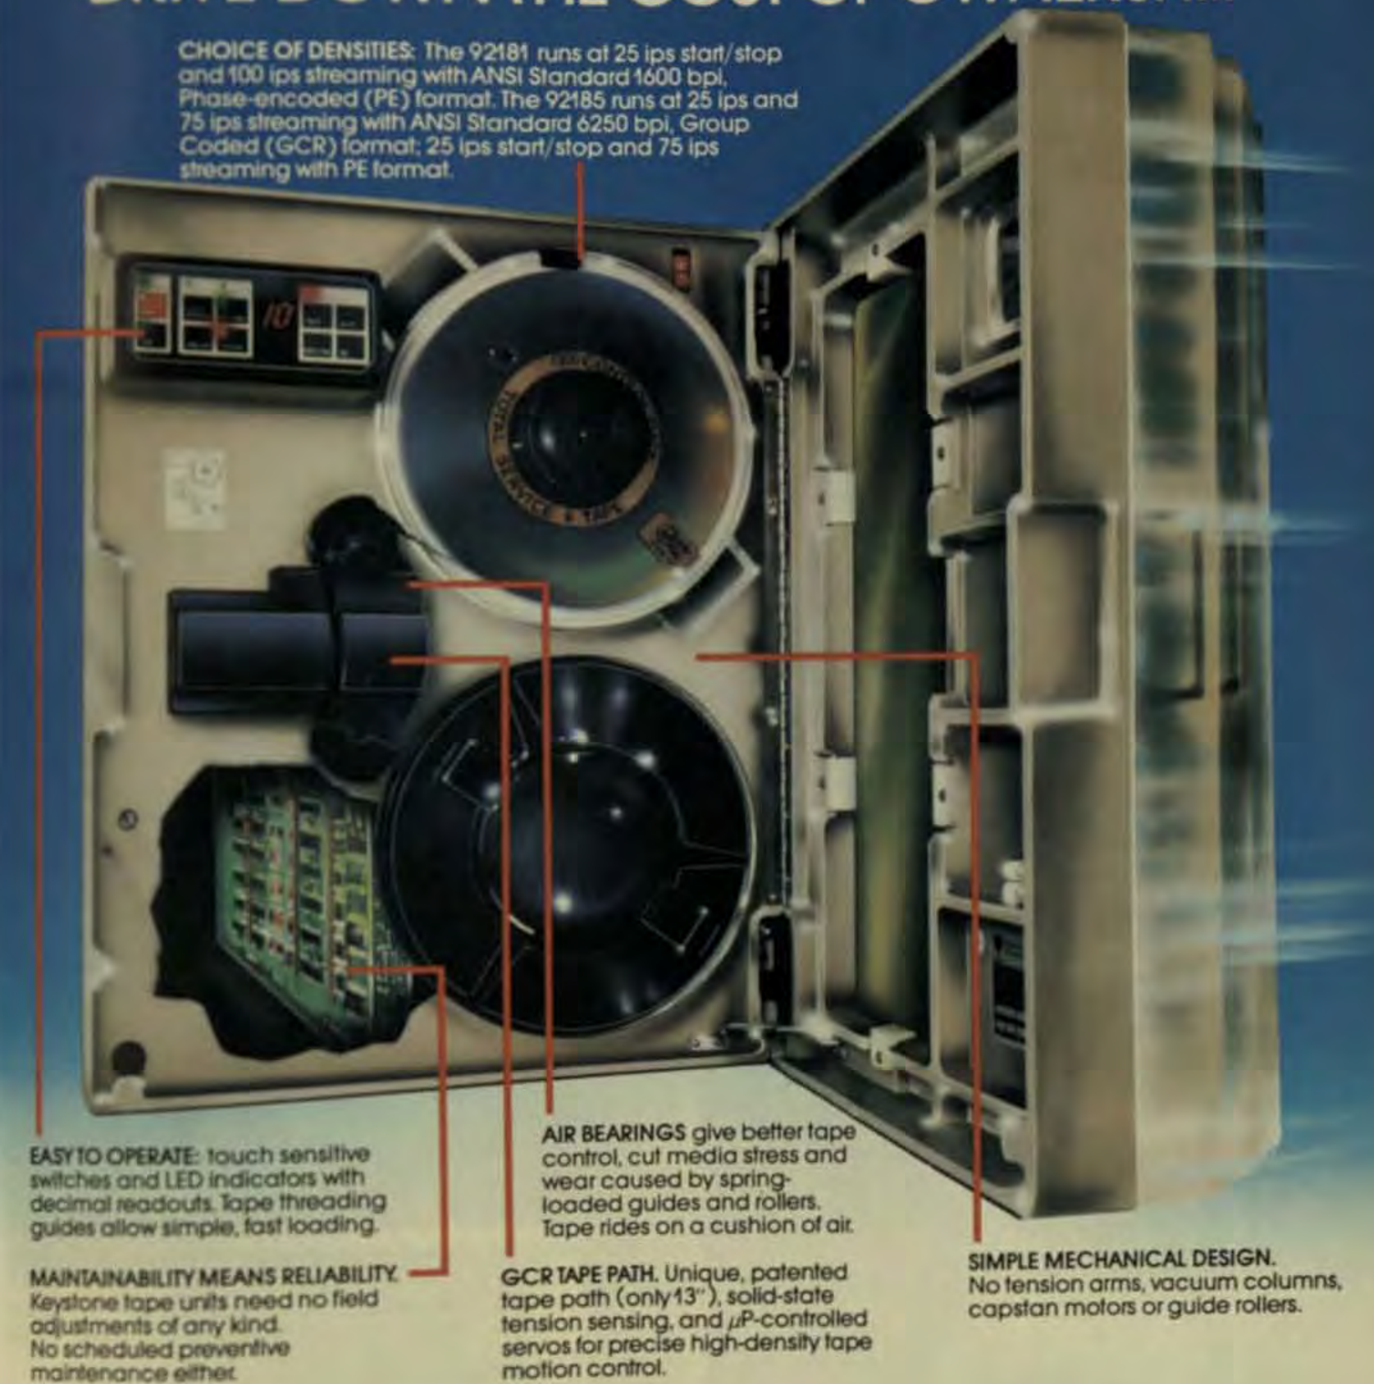

THE KEYSTONE™ SERIES

We designed the Keystone Series with built-in diagnostics and maintenance features that make it simple to operate, easy to service, low in life-cycle costs. For more information, call your local Control Data OEM Sales Representative or write: OEM Product Sales, HQN08H, Control Data Corporation, P.O. Box 0, Minneapolis, MN 55440.

CONTROL DATA

Information Hot Line 1-800-621-1776.
In Illinois, 1-800-572-6724.

stored for security purposes, such as fire, flood, or theft; and to protect against equipment failure such as head crashes on a disk.

Utility I/O is a real-time or monitoring function. Examples include: on-line data processing, where the cartridge drive is used as a real-time full-function peripheral; data collection, as in a data-acquisition system; word processing, as a convenient, high-capacity memory for the storage of long documents; and file management or restructuring.

Initialization is a software and data-file function. It includes: program loading; data interchange; software updates; and software distribution.

Applications of ¼-in. digital cartridge tape drives seem to be split into about 65% backup, 20% utility I/O, and 15% initialization.

BACKUP

Streaming cartridge tape drives are best for mirror-image backup. In this application, a straight binary dump from the Winchester disk to the tape is carried out with no regard for the location of data files on the tape. This procedure can also be carried out in the reverse direction as a mirror-image restoration of data from the tape to the disk. With mirror-image backup, there are a minimal number of tape stops, so concern about the length of time that it takes a streaming tape to restart after a stop is minimized. Start/stop drives can also do an efficient mirror-image backup, but for large record sizes they're slower than streamers.

Often selective backup is the desired mode of operation and the user wants to save only portions of the data on the disk. In that case, the disk is searched for the desired data, and only this data is sent to the tape. Since that requires a fair amount of time, selective backup often results in excessive stopping of a streaming tape drive and yields a bias toward the start/stop drive for this class of application.

Most often, mirror-image isn't enough and the user wants one or more of the utility I/O or initialization functions as well. Then the choice rapidly moves to the start/stop drive. The Table indicates the proper drive choice for various applications.

FILE RESTRUCTURING

File restructuring is one of the most common applications to accompany backup. For example, assume an organization desires to create a file on each customer, and that space is allocated within the disk file based on one activity per day. Occasionally certain customers will need more than one entry, which requires the disk controller to find additional space within the system for these files. Portions of the file may be allocated from the outer tracks of the disk, portions from the center tracks, and some portions from totally different disk surfaces.

The access times involved will degrade overall system performance when this file is being compiled. The operating system could retrieve this data wherever it may be, and a start/stop drive could be used to store it on tape in the most logical order. Then, when the data is transferred back to the disk, it will all be in one area of the disk, and the seek time will be minimized when the data is needed by the CPU. ■

Which Way Is the Market Going?

One of the hottest areas in the small-business-computer marketplace right now and for the foreseeable future is that of local-area networks and other forms of multi-user systems.

According to Frost and Sullivan, virtually every business with sales of more than $10 million already uses computers, and by 1985, 44% of the businesses in this country with sales of less than $10 million will have them as well. Since the trend is toward multi-user systems and local-area networks, it's clear that memory requirements for these systems will be significantly higher than for single-user systems.

With the increasing usage of Winchester disk drives in this class of system, there is a concurrent increase in the need to back up these devices. Thus, there's a dynamic and rapidly growing market for ¼-in. digital cartridge tape drives to satisfy these backup requirements.

The curious thing about this market for backup drives is that it's not following true to the form outlined by the market predictors. When DEI introduced ¼-in. cartridge streaming in 1979, there was practically universal agreement that streaming would quickly become the dominant form of Winchester disk backup. But while it has been a significant factor, it's never overtaken start/stop drives for backup purposes.

It appears that most system designers want to do more than just mirror-image backup. Once a Winchester disk drive is included in a system, it's no longer a bottom-of-the-line system; it is expected to offer great flexibility and multiple functions. This philosophy carries over to the backup tape; it also is expected to be available for file management and other full-function tape jobs. Thus, it must be a start/stop drive.

While newer and more capable streamers are expected to expand the market for streaming tape drives, they won't overtake the newer and more capable start/stop drives that are being introduced each year.

Kenji Eigner, *product manager at Data Electronics Inc., previously held various marketing and product management positions at companies such as Cipher Data Products. He has a BS in computer science from San Diego State Univ.*

SEEING IS BELIEVING

NOW WANGTEK'S ¼-INCH STREAMER PACKS 20 OR 45 MEGABYTES INTO A HALF-HIGH 5.25-INCH SPACE.

Wangtek now opens up new design and marketing options to the systems integrator and the OEM with its introduction of the Series 5000, the only available 20 or 45 Mbyte half-high ¼-inch streamer; an ideal backup for 5¼-inch Winchesters.

Check these advanced features:

Small size: The Series 5000 can be easily packaged with a half-high 5¼-inch floppy or half-high Winchester.

D.C. voltages: Uses the same voltages as a minifloppy and 5¼-inch Winchester (+5V, +12V).

QIC-compatibility: Packaged in an additional half-high module, the two-board intelligent formatter is capable of reading and writing the QIC-24 defined recording format facilitating data interchange ... is I/O compatible to the QIC-02 defined intelligent interface ...ard.

... **cartridge loading:** The Series 5000 incorporates a unique ...ction cartridge loading mechanism for easy operation.

Positive tape interchangeability: With its rigid and stable aluminum deck, stiffer head carriage assembly and track reference burst, the Series 5000 can consistently and precisely position the read-write heads over a desired track on a cartridge written on any ¼-inch streamer.

Host adapter: The half-high intelligent formatter module contains an additional slot for a custom PCB that could be used as an adapter to any host CPU bus.

The Series 5000 is easy to integrate into existing systems and a must for future designs. Call now for delivery, pricing, and facts on Wangtek's new Peripheral Integration Program.

Wangtek, 5845 Uplander Way, Culver City, California 90230. (213) 410-1444, Telex: 664425.

wangtek

Streaming Tape Drives Adapt Easily to Small, High-Capacity Systems

Connecting a ½-in. streaming tape drive to a small computer system via a standard interface provides high-capacity backup and other capabilities.

by **José C. Elaydo,** Rosscomp Corp.

The primary reason for streaming tape's popularity is the rapid increase in Winchester disk capacities. Eight-inch Winchesters with capacities exceeding 80 Mbytes are now becoming routinely available to designers of mini- and microcomputer-based systems, and by this time next year, 300-Mbyte OEM 8-in. drives will be readily available.

The growing disk-drive capacities have brought with them a demand for reliable, high-speed, high-capacity peripheral equipment capable of providing low-cost secondary storage. This secondary storage is required for on-line hierarchical memory-management schemes, off-line storage during large-scale file restructuring, and traditional file-oriented full-function backup of 8-in. Winchester disk drives. Low-cost streaming tape drives meet these requirements.

STREAMING DRIVE TECHNOLOGY

In a function such as the backup of a Winchester disk drive, a streaming tape drive ramps up to full operating speed, then records all data on tape at a continuous high rate. Depending on the way system software handles the task, the streamer can either record data in a file-by-file, sector-by-sector, and record-by-record manner, or as a "mirror" image.

The data is recorded on tape in a serpentine fashion: The tape drive makes one continuous pass from one end of the tape to the other, recording data serially on one track at a time. When it gets to the end of a track, it reverses the tape and repositions the read/write head on another parallel track. The drive now records on this track until it reaches the other end of the tape.

This process continues track by track until all the data has been recorded. The number of tracks available on the tape depends on the format used and the width of the tape. A ½-in. tape, for example, may employ 24 tracks for a total capacity of 160 Mbytes.

As data is transferred to or from the streamer, the tape controller usually performs error-detection and -correction as well as parity checks to ensure accuracy. If an error occurs during a read operation, the tape drive decelerates, repositions the tape, and accelerates back up to speed. If an error occurs when the streamer is writing the data, the drive keeps going and rewrites the data.

Like streaming tape drives, start/stop devices provide high transfer rates, but they handle data in blocks. At the end of each block, the drive decelerates, comes to a complete stop, and waits. When the next block comes, the drive accelerates up to speed and writes that block. Thus, the blocks on the tape are separated by interrecord gaps, and the drive must accelerate and decelerate quick-

HERE'S 330 MEGABYTES OF THE FASTEST, MOST VERSATILE, MOST INEXPENSIVE, HIGH CAPACITY WINCHESTER BACKUP YOU'LL FIND ANYWHERE.

With the growth of today's high-capacity Winchester disks, you simply can't ignore the backup issue any more.

Fortunately, MegaTape has solved the problem.

Using proven, reliable 9600 bpi tape technology, we've built the industry's most advanced file backup system. Once you see it in action, you'll agree: there's really no sensible alternative for backing up large disks.

MegaTape Features

High capacity in a convenient, inexpensive format. Your choice of 220 or 330 megabytes in the convenient cartridge shown here. Cartridges use standard 1/2" tape, and cost around $85.00 in OEM quantities.

Streaming and start/stop modes. In 200 ips streaming mode, MegaTape reads or writes an entire 330 MB cartridge in just 24 minutes! In start/stop mode, tape speed is 50 ips.

Random access capability. MegaTape's serpentine track layout and unique head stepping capability let it locate any segment of tape in an average 30 seconds, for fast, simple file restructuring.

Compact, low-cost drive. The ingeniously simple transport (pat. pending) occupies just 8¾" in a standard 19" EIA rack. Models start at $3,200 in OEM quantities—about 1/3 the price of standard 1/2" GCR drives.

Industry-standard interface. MegaTape plugs into a standard Pertec-type interface, and uses standard controllers and data rates.

If you have a backup problem for large disks, talk to MegaTape today.

There's really no alternative that makes so much sense!

Actual size: 9.8 x 6.4 x 1.1 inches

© 198[illegible] MegaTape Corporation

MegaTape, 1041 Hamilton Road,
P.O. Box 317, Duarte, CA 91010 (213) 357-9921

I'm intrigued by your claims, but I need backup. Send me more information.

Name ____________________ Title __________

Company ____________________

Address ____________________

City __________ State __________ Zip __________

Phone (____) __________ Ext. __________

MegaTape CORPORATION

ly (tens of milliseconds) to avoid wasting tape.

Getting the drive up to speed and stopping it within the fairly short gap—without breaking the tape—requires some complex and expensive features. For instance, multiple high-torque motors are needed to accelerate and decelerate quickly, and tension arms or vacuum columns must be used to avoid straining the tape.

The mechanisms required for start/stop operation make these units inherently more expensive than streaming drives. Because streamers work continuously, without stopping for gaps, they use less powerful motors that cost less. And the slower acceleration and deceleration puts less tension on the tape, so tension arms and vacuum columns aren't necessary in streaming drives.

THE SYSTEM ENVIRONMENT

Streaming tape drives are often relegated exclusively to the task of Winchester backup. In this activity, there are usually three levels of backup. Three tapes containing the three most recent days of Winchester disk contents are rotated, the oldest one being rerecorded with the current day's data. Thus, backup data is never kept more than three days.

While streaming tape drives do an excellent job of backup, they also serve in other ways, such as for secondary storage and software distribution. In archival storage, for instance, as opposed to backup, the data copied to tape is for long-term storage off-line. This data could be from a Winchester, or it might simply represent a large volume of information previously kept on paper.

Income-tax forms provide a good example of the latter. At year's end, an accountant could key his clients' tax data into his computer, then transfer it all to tape for compact storage in a safe place. The archive tape furnishes a history of past activities and, thus, an audit trail.

In their role as secondary storage, streaming tape drives can keep on line any data that's inconvenient or unnecessary to have on the primary storage devices (Winchester and floppy-disk drives). Some examples of candidates for secondary storage include very large data files, information gathered via a data-logging activity, and files for up- or downloading to or from other devices.

As Winchester capacities increase and greater numbers of small computer systems incorporate streaming tape drives to back up crucial data, the tape drives will also be used for software distribution. System manufacturers indicate a willingness to replace floppy disks as the medium for updating operating systems and distributing applications software because tape is cheaper and can more easily hold very large programs or operating systems (such as Bell Labs' UNIX). Floppies could therefore be eliminated entirely from many systems.

Streaming tape drives are also excellent for other applications, such as telemetry monitoring and seismic data acquisition, which call for capturing a fast stream of real-time data. And imaging applications such as radiography can make good use of the streamer's high capacity and transfer rate. Digitizing a single X-ray, for example, generates more than a million bits of data, and already about 25% of the X-rays taken in this country are never stored on film. This massive amount of X-ray data must be stored in readily available form; streaming tape provides this storage quickly and inexpensively.

STREAMER INTEGRATION

The integration of a streaming tape drive into a small computer system involves many of the same considerations as the incorporation of any other peripheral; the critical elements are the interface and controller. There are differences between streamers and other peripherals, but these differences are generally handled in the controller, so they need not pose a problem for systems integrators.

One factor that has simplified the use of off-the-shelf controllers for streaming tape drives is the standardization of host and peripheral interfaces. De facto standards for 8-in. Winchesters, such as the SMD and SA1000 interfaces, have been available for some time and make this side of the integration quite straightforward. As for the tape-drive interface, the proposed QIC (quarter-inch cartridge) standard originally designed for ¼-in. tape cartridge drives is perfectly adaptable for ½-in. tape units.

The newest standard available to systems integrators, the small computer systems interface (SCSI, formerly known as SASI), allows addressing in logical rather then physical terms, thus promoting instant compatibility between a host processor and a peripheral device.

Today's controllers can eliminate much of the burden of systems integration and provide valuable features that also eliminate a significant portion of CPU overhead. Some controllers handle the entire process of backing up and restoring Winchester data with a streaming tape drive; the CPU is not involved at all. The Rosscomp C160 SCSI-based controller, for example, can operate in an off-line mode that lets it transfer data to and from the four Winchesters it supports and a ½-in. streamer by file ID number. No streaming software is required.

Further, the controller handles direct file copy between the disk drives, automatic tape read/write error correction during backup, and auto tape retry during restore. In on-line mode, the C160 acts much like any other tape/disk controller. For the tape drive, it takes care of track selection and data transfers from the disk at approximately 3 Mbytes/min. The controller performs Winchester functions such as disk and track formatting, head seeks, recalibration, [illegible] sector buffering, and overlapped seeks.

To make use of the streaming tape drive more efficient, however, the controller selects the starting tape-track number if a streamer is used in split-tape mode. By dividing the driver format into logical drives of 2, 4, 8, or 12 tracks, the average access time for any file stored on tape is reduced.

An SMD interface version of the controller designated the C[illegible] offers capabilities similar to the C160 but with a choice of tape speeds. The higher of these speeds (130 ips, in contrast to the 90-ips rate used on other drives) lets the tape drive take advantage of the better throughput available from the SMD interface. Compared to the Q2000 or SA1000 disk standards, SMD provides a higher-performance Winchester interface with a higher data-transfer rate.

Note that this higher transfer rate is generally available only for transfers directly between the [illegible]

chester and the tape drive. If the data must travel through the CPU, only high-speed direct memory access (DMA) has any hope of keeping up with a 130-ips rate. If data isn't fed to the tape unit fast enough, the drive must stop, reposition itself, and ramp back up to speed—a time-consuming process that can wipe out the benefits of using the 130-ips speed.

But using a controller that handles 130-ips transfers off-line eliminates the problem of handling CPU-related bottlenecks as well as allowing the CPU to perform other tasks at the same time. Such a controller lets a system take full advantage of the SMD interface for high-speed backup.

A controller that uses the last interface discussed here—the proposed QIC standard—offers a different kind of advantage. Although the QIC interface was developed for ¼-in. tape cartridge drives, it's perfectly suited to ½-in. units, too. QIC-compatible processors can therefore gain greater backup capacity by using a ½-in. tape drive; all that's required is a controller, such as Rosscomp's C161, that connects to a host via the QIC's 8-bit parallel interface.

One way to simplify the often complex integration process is to use an evaluation system that lets you exercise a streaming tape drive to see how well it will operate with other equipment. These evaluation systems need not be complex. All that's needed is a tape drive; controller; product and interface documentation; cables and connectors; and an exerciser that allows you to manipulate the drive mechanically, write data patterns, step the head, and analyze signals with an oscilloscope.

Such a system can shorten the time needed to get a tape unit up and running with some specific mix of other peripherals and a processor. Then it lets you test the combination under actual working conditions to see how the tape subsystem performs. These steps are followed in any integration process, but a prepackaged evaluation system makes it a more cut and dried undertaking.

STREAMING TAPE ALTERNATIVES

When you integrate a tape drive into a computer system, you can choose among several alternatives, from cassettes to ½-in. reels. Because the performance differences from one end of this range to the other are tremendous, you can make the choice easier by comparing the applications involved.

At one end of the tape spectrum are 3M-type 0.15-in. and Philips-style cassette tape drives. Although these units are fine for small or inexpensive computers that need some form of mass storage, their low data capacity makes them ill-suited for Winchester backup.

Another alternative lies in ¼-in. 3M-type tape cartridge drives. Currently, these can store from 20 to 45 Mbytes of data, which is sufficient for backing up most available 8- and 5.25-in. Winchesters. However, attempts to provide greater capacities have required the use of specially certified and formatted media that will handle higher bit densities, and the use of other technologies that are not yet field-proven. Further, the cartridge's design makes it difficult to maintain constant tape tension at the read/write head.

Employing reels rather than cartridges and enlarging the tape width to ½-in. overcomes the limitations imposed by instantaneous speed variations in ¼-in. cartridges. As Winchester disk drives move into the 160-Mbyte range this year and to 300 Mbytes next year, ½-in. tape's higher storage capacity becomes increasingly important. And by making the tape self- loading, the ½-in. reel can be as convenient as a cartridge, but without the cartridge's performance limitations.

Another factor to consider is cost—both the cost per megabyte per box and the cost of media per megabyte. Half-inch streaming tape units generally offer a better cost per megabyte than other products. Of course, the original cost of a ¼-in. tape drive is less, so systems that have only a few megabytes of Winchester storage to backup might make good use of a low-capacity ¼-in. cartridge.

For higher capacity systems, though, the ½-in. alternative clearly costs less overall, especially when you consider the cost of media per megabyte. The least expensive tape package ever devised is the reel. A reel of ½-in. tape costs less than half as much as a ¼-in. cartridge with the same length of tape, and the ½-in. tape will store far more data. Cartridges offer convenience, but as noted earlier, self-loading reels are equally easy to use and also furnish equal, if not better, environmental protection.

OTHER BACKUP HARDWARE

In addition to the standard tape-drive alternatives for backup, you can also consider floppy disks, Winchester cartridges, and videotape recorders for the task. Here again, the value of these products depends on your application.

Floppy disks provide a readily available backup method. They can easily back up Winchesters with capacities under 20 Mbytes on a transaction basis. However, the relatively low capacities and/or low transfer rates of currently available 8- and 5.25-in. floppies make them impractical for applications where database security demands the backup of the entire file at the end of each day, or where a constant audit trail is necessary.

Typical floppy capacities range up to about 1 Mbyte, but this level would require five media changes to store the data on even a low-end Winchester for archival purposes. Another problem arises when large off-line storage capacities are required to restructure files and when Winchester capacity exceeds 20 Mbytes.

A number of vendors have introduced floppy-disk drives with capacities over 3 Mbytes, making these drives far more efficient in backup applications from the point of view of the number of diskettes required to handle an entire file. These drives have yet to be manufactured in large volumes, however, and it has yet to be determined to what extent their acceptance will be affected by the limited sources of the specialized media they require.

Winchester disk-cartridge drives offer a better alternative for some users. Since the introduction of the first fixed-disk Winchester, several vendors have announced products with the removable cartridges. These drives are well suited not only as systems devices, but as high-capacity, high-speed backup for fixed-disk drives.

At the same time, it's good to be aware of Winchester cartridge limitations, most of which stem from inherent properties of the media. First, the media is expensive—approximately $50 per cartridge in large quantities.

Remember TEAC for Memory Machines

- MT-2 Digital Cassette Transports with LSI-Controller
- Specifically designed and built by TEAC for computer-based systems
- Ideally suited for data logging and program loading applications
- Over 250,000 units in use
- Over 20 years of expertise in digital technologies — as well as quality stereo equipment
- Complete product support in all areas — marketing, technical and service

TEAC

For complete technical information, write:

TEAC Corporation of America
Industrial Products Division
7733 Telegraph Road
Montebello, CA 90640
(213) 726-0303
(213) 726-8417

Second, data interchangeability continues to be a problem, not only when cartridges are moved from drive to drive, but when the same cartridge is re-inserted into the same drive. Third, the use of proprietary servo patterns on many cartridges limits the sources for media.

Further, on many removable Winchester drives, the fixed systems disk and the removable cartridge are mounted in the same package and are driven from the same spindle. If for some reason the fixed systems disk goes down, the backup disk will go down as well.

Finally, the fast pace of fixed-disk Winchester technology tends to mean that a single tape cartridge can no longer store the entire contents of a fixed disk.

Currently, for example, 5.25-in. Winchesters with capacities in the 20-Mbyte range are readily available, with hardware at the 40- to 50-Mbyte level coming on fast. Disk cartridge drives with these capacities are not available, however. Similarly, the capacities of 8-in. fixed Winchesters routinely exceed 100 Mbytes, but 8-in. disk cartridges are so far limited to 25 Mbytes.

Some video-related technologies—most notably videotape recorders—have also been proposed for high-capacity Winchester backup. These drives offer the potential of enormous capacities, but they suffer from inherently long transfer times, high cost and large size, and the lack of suitable controllers. Video disks could someday furnish even higher capacities than videotape recorders. However, for the present they are even further from viability than the tape units and provide no ability to erase data once it's written.

THE FUTURE OF STREAMING TAPE

As computer systems get smaller in size, lower in cost, and incorporate increasing amounts of Winchester storage, the role of high-capacity tape drives will change significantly. Today, most microcomputer users have a relatively small amount of information stored on Winchesters. But as more data is acquired in the course of day-to-day business and new applications are added, the mass-storage capacity of small systems will balloon. Even if users continue to employ the same computers, they will be forced to add more mass storage through the years—a situation that will call for improved backup capabilities.

Although the small amount of data currently stored in most small computers doesn't warrant backup, the dramatic increase in data volume that will take place over the next few years will make backup mandatory. And as noted previously, the availability of tape drives in small computer systems will prompt their use for other purposes, such as software distribution.

Another trend in mass storage that's important to consider for backup implications is the increase of 5.25-in. Winchester storage capacities. Because it won't be long before 5.25-in. drives store as much data as 8-in. units do today, high-capacity backup systems will soon be needed in the 5.25-in. form factor. Half-inch tape on reels has an advantage here because it can be reduced to any size. Neither the media nor the technology limits future growth—even to a sub-5.25-in. Winchester chassis.

The characteristics of the tape are also crucial to consider. Some tape-drive manufacturers are now urging adoption of a media that could more than double the storage capacity available from a given length of tape. This improved tape would increase the standard coercivity from its current value of 295 oersteds to about 700 Oe, a level that can be handled reliably with today's technology. To keep the cost of this media reasonable, though, the industry (via an agency such as ANSI) must agree on a new standard.

The continually changing outlook for small computers requires that systems integrators account for tomorrow's needs as well as today's. For streaming tape drives, this means selection of a media format that can easily increase in capacity and decrease in size and cost—considerations that point to ½-in. tape as a solution for most applications. ■

Jose C. Elaydo, *vice president of marketing and sales at Rosscomp Corp., previously worked at Archive Corp. He has been involved in marketing, sales, and research within the computer industry for the past 20 years.*

The Perfect Complement

DEI® introduces a new dimension in magnetic tape cartridges. Total storage of data (at 6400 bpi/fci) is achieved with 555 feet of high grade ¼" tape.

This offers the end-user less expense for most applications by providing a lower cost per megabyte than other media types.

Model 401555
High Density ¼" Cartridge

For the All New Super Funnel

Based on the highly successful industry standard DEI Funnel, the Super Funnel is a high capacity start/stop, full function tape peripheral using the ANSI standard 450' and the new 555' magnetic tape cartridges. The Super Funnel has the added ability to read tapes generated at 6400 bpi with the faster through-put rate of 37.5 ips as well as read and write at 8533 bpi. with 50 megabytes of removable storage. Contact your local DEI office or Representative.

Model 310044
Super Funnel

® Registered Trademark of Data Electronics, Inc.

DATA ELECTRONICS, INC.

10150 SORRENTO VALLEY RD.
SAN DIEGO, CALIFORNIA 92121
TELEPHONE (619) 452-7840
TELEX 69-7118 • TWX 910-335-1150

Embedded Formatters Provide New Substance To GCR Tape Subsystems

by **Marold H. Lohrenz,** Telex Computer Products Inc.

Half-inch tape is the standard for information interchange and library storage. As new technology has been introduced, the costs have decreased and the systems have become more user-friendly. Since its introduction in the 1960s, ½-in. tape has enjoyed a position as the most cost-effective means of data storage and exchange. Fortunately, the standards and industry discipline have remained intact. You can transfer tapes from site to site knowing that the data can be read, the tape updated, and the new data redistributed as required.

Today, three primary ANSI standards co-exist—nonreturn-to-zero current (NRZI), phase-encoded (PE), and group-coded recording (GCR) (see Fig).

NRZI, introduced in the 1960s, provided a major step in capacity compared to the punched card. However, NRZI was limited at 800 bpi (bits per inch) due to mechanical skew in the tape drive. When PE was introduced a few years later, it doubled the capacity (to 1600 bpi) by incorporating a self-clocking technique that yielded a format that was more tolerant of mechanical skew. This standard also increased data reliability and allowed a single-track error to be corrected on the fly in a manner transparent to the user.

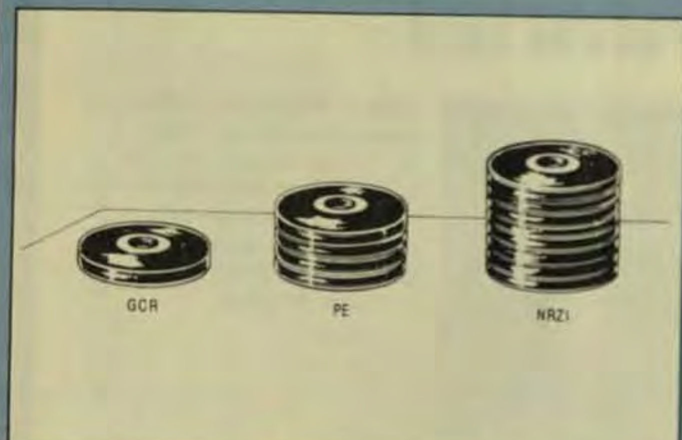

Fig The GCR format increases the storage capability of tape systems.

As disk systems continued to expand, greater tape capacity was required. The GCR format introduced in 1974 almost quadrupled the recording density (from 1600 to 6250 bpi) and cut the interblock gap (IBG) in half (from 0.6 to 0.3 in.). Data reliability was increased another order of magnitude as a result of the error-correction/detection facility embedded in the format, which allowed two track errors to be corrected on the fly and was completely transparent to the user.

The result was a capacity of up to 180 Mbytes per reel (compared to 46 Mbytes for PE), and a data reliability that was 10 times better. The GCR format could also be recorded on most tapes in existing libraries. The user [illegible]uld then upgrade his tape-storage capability with [illegible] as he introduced new, higher capacity disk [illegible]ms.

The GCR format is required when disk systems approach capacities of 100 Mbytes or more and when capacity and throughput are the critical factors. Capacity depends on the hardware capability and the size of the tape's data blocks. Typically, users will record 40 Mbytes per reel in PE and 150 Mbytes per reel in GCR. For a given tape speed, GCR will increase the effective transfer rate by 3 to 3.8 times.

Currently, the 10 ½-in. reel with GCR format is gaining in popularity for lower power computers, and the introduction of a 50 ips GCR start/stop subsystem is further evidence this format will find a continuing market. Now, following the lead of NRZI and PE tape subsystems, high-performance GCR products are using embedded formatters (controllers) that are compatible with existing interfaces. These will make GCR easier to use, reduce cost, and simplify systems integration.

EMBEDDED FORMATTERS

The formatter handles device-dependent details and presents an interface that transfers data in a conventional digital manner. It controls error detection and correction, encoding and decoding, detail parameters for individual devices, and tape-drive control.

To date, GCR formatters have been too large to be packaged inside the tape unit because they required 15 to 20 printed-circuit cards. However, today's semicustom techniques can deliver the contents of an entire PC card on a single chip.

For new CPU systems currently on the drawing board, embedded designs will be available with integral intelligent interfaces such as those being developed by the X3T9 ANSI committee. These designs will be developed anticipating the acceptance and implementation of interface standards. Then, with a hardware modification (card and cable) and a firmware change, the same GCR subsystem can serve existing interfaces, and the new intelligent-interface standard as well. In fact, these lower cost embedded subsystems will be supported by the same system software and job procedures used for current high-end GCR subsystems.

EMBEDDED ADVANTAGES

In the new low-price/high-performance GCR subsystems the quantity of components will be reduced along with the number of field-replaceable units. Therefore, diagnostic checks will be simplified and spares inventory will be reduced. Combined with the effective use of multiple microprocessors, the diagnostics will be more effective in fault isolation.

Embedded formatters more than double the mean time between failure (MTBF) of current GCR products. Adaptive designs implemented in firmware will be more tolerant of midlife variations inherent with an electromechanical device. Additionally, self-calibration routines included in the firmware compensate automatically for drift, component wear, and initial parametric variations.

These routines may be initiated at power-on and during specific operational modes, or controlled by the monitoring of certain operational characteristics. In other instances, adjustments will simply be eliminated.

Diagnostics resident in the subsystem can be used in a variety of ways. Traditionally customer engineers have used these routines to repair a subsystem. With embedded formatters, the operator has access to a limited subset of routines for a confidence check, and the customer engineer can use remote terminals to access a central location. RS-232 ports can be supported by special testers or controlled via modems at a central repair and product-support facility.

Maintenance is yet another area that benefits from embedded formatters. A major portion of product maintenance is devoted to compensating for the wear and physical changes of mechanical components. You can eliminate some problems by minimizing the number of mechanical components, of course, but sometimes a better solution is to increase their tolerance range. When more of the unit's real-time operational control is incorporated in firmware, microcode can track the inevitable physical wear and compensate for it.

SUBSYSTEM PACKAGING

With the formatter embedded in the tape unit there are more functions per module, and less space is required for new GCR tape subsystems—typically half that of current GCR subsystems.

Since power requirements are also reduced, convection cooling can be used, and that further reduces space, power, and the noise level. Integral brushless dc motors represent still other power, space, and maintenance savings. With smaller assemblies, packaging arrangements provide ready access for maintenance in a compact subsystem and reduce the time to repair.

Saving on space gives you more versatility in mounting arrangements. Tape subsystems capable of both horizontal and vertical mount can be packaged above disks in a meter-high peripheral cabinet or mounted at an angle to provide a low profile and still retain visual contact with the reels from across the room.

START/STOP VS. STREAMING

Tape units operate in one of two modes—start/stop or streaming. Typically these units operate with tapes using the ANSI interchange standards where the data blocks, at a predefined format, are separated by predefined interblock gaps (IBGs). For GCR, the standard specifies a 0.3-in. IBG and a storage format at 6250 bpi with an embedded error-correction facility. The PE standard specifies a 0.6-in. IBG at 1600 bpi. Thus, the decreased IBG leaves more tape available for data. Combined with the increased density, this typically yields 150 Mbytes recorded on a GCR tape and 40 Mbytes on a PE tape.

For start/stop units, when a data block is transferred, the tape may be stopped in the IBG and the next block transferred when the data is ready. The start/stop time will add 5 ms or less to the transfer operation. If the data is available immediately, the tape unit will not stop, thereby passing the IBG on the fly at the nominal tape speed. This sequence of transactional transfers continues down the tape, with the CPU transferring data as it becomes available and intermixing it among the various tape units as required by the application.

Streaming-tape units have been introduced to minimize product cost and optimize the unit for disk backup. This type of tape unit moves tape directly from one reel to the other and represents a cost/performance trade-off. To maintain the streaming mode, data must be continuously available at the tape unit, in turn maintaining the tape at a constant speed.

This requires special considerations at the system level and may require a dedicated CPU. If the data isn't available, or if an error condition exists on the tape, the unit will stop, back up, and then resume forward motion. This recovery typically requires 0.5 s or more as compared to 5 ms or less for a start/stop unit.

Throughput is controlled by the proportion of time used to transfer data, compared to the recovery time. If the application and the system parameters permit a continuous data flow, the throughput of a streamer and a start/stop unit will be equivalent. If interruptions occur or are unpredictable, then the recovery time—100 times longer for the streamer—becomes the controlling factor, and the start-stop unit will provide a significantly higher throughput.

A start/stop unit is also more immune to changes at a given installation. As users migrate and increase their capabilities in disks, CPUs, multiprocessing software, and other system enhancements, a start/stop tape unit will be relatively transparent to these changes.

Manufacturers today are concerned with reducing the cost of GCR subsystems without sacrificing performance so that their advantages can be obtained without major investments in system software, or changes to system-operational procedures and job mixes.

PRICE/PERFORMANCE

When GCR was introduced, the tape unit and formatter were each packaged in separate 30-in.-wide cabinets. The rack-mounted units introduced in 1978 were half the size, and units on the drawing boards will be only half as large as those. For this trend to continue, the price of a typical subsystem must also be reduced significantly while performance is maintained. Semi-custom technology and more power-efficient designs will accomplish that goal, and allow hardware designs to be implemented in a combination of firmware and more compact circuitry. These techniques will yield a subsystem that will be more reliable, smaller, and consume less power.

Half-inch tape has been the cost-effective means for storing data. The technology improvements currently on the drawing boards dominance through the 1980s and 1990s. The integrity and utility of the user library will be preserved. The flexibility afforded by the use of multiple low-cost microprocessors, and the advances in motor designs and power-efficient components will decrease product cost. And LSI technology will further decrease the cost of the control and data-handling circuits.

Also, servos will be implemented in firmware, not hardware. The user will have available tape units at one-half of current costs and, with the start/stop units, will have a substantial price performance improvement completely transparent to his current system operations. Manufacturers will introduce such tape units within the next two years. ■

You have an investment in power. The power that runs your computer, telecommunications, avionics, process-control or other sophisticated electronic equipment. You need a clean, consistent, ever-present line of power that adds to your bottom line.

In the computer, aerospace,

Your investment is on the line. Elgar can keep it there.

nuclear, medical and process-control industries and in many other applications, Elgar is shaping power for today's advanced technologies.

Elgar is known for quality. Since the 1960's we've been a leader in power conditioning, power modification and uninterruptible power supply. Elgar manufactures a full line of Precision and Ultra-Precision Line Conditioners (single- and three-phase versions to 15 kVA), High

Isolation Transformers (for 50 or 60 Hz lines, 1 kVA to three-phase 60 kVA), AC Power Sources (single-, two- and three-phase VA to 36 kVA), Oscillators and Digital / Analog Programmers Hz to 10 kHz), and Uninterruptible Power Systems (400 VA to 50 kVA).

Elgar power conditioning equipment is a vital part of many telecommunications satellites, off-shore oil rigs, large and small computers, hospitals and industrial facilities. Whether your power needs are simple or complex, Elgar can deliver the right product to keep you up and running—flawlessly.

At Elgar we know there's a lot on the line—your investment.

Powerful Solutions for Advanced Technology

Call today for complete information

1-800-854-2213

ELGAR

ELGAR CORPORATION
An Onan / McGraw-Edison Company
8225 Mercury Court
San Diego, California 92111
Telephone: (619) 565-1155
TWX: 910-335-1246 Telex: 6834028

Power System Technologies

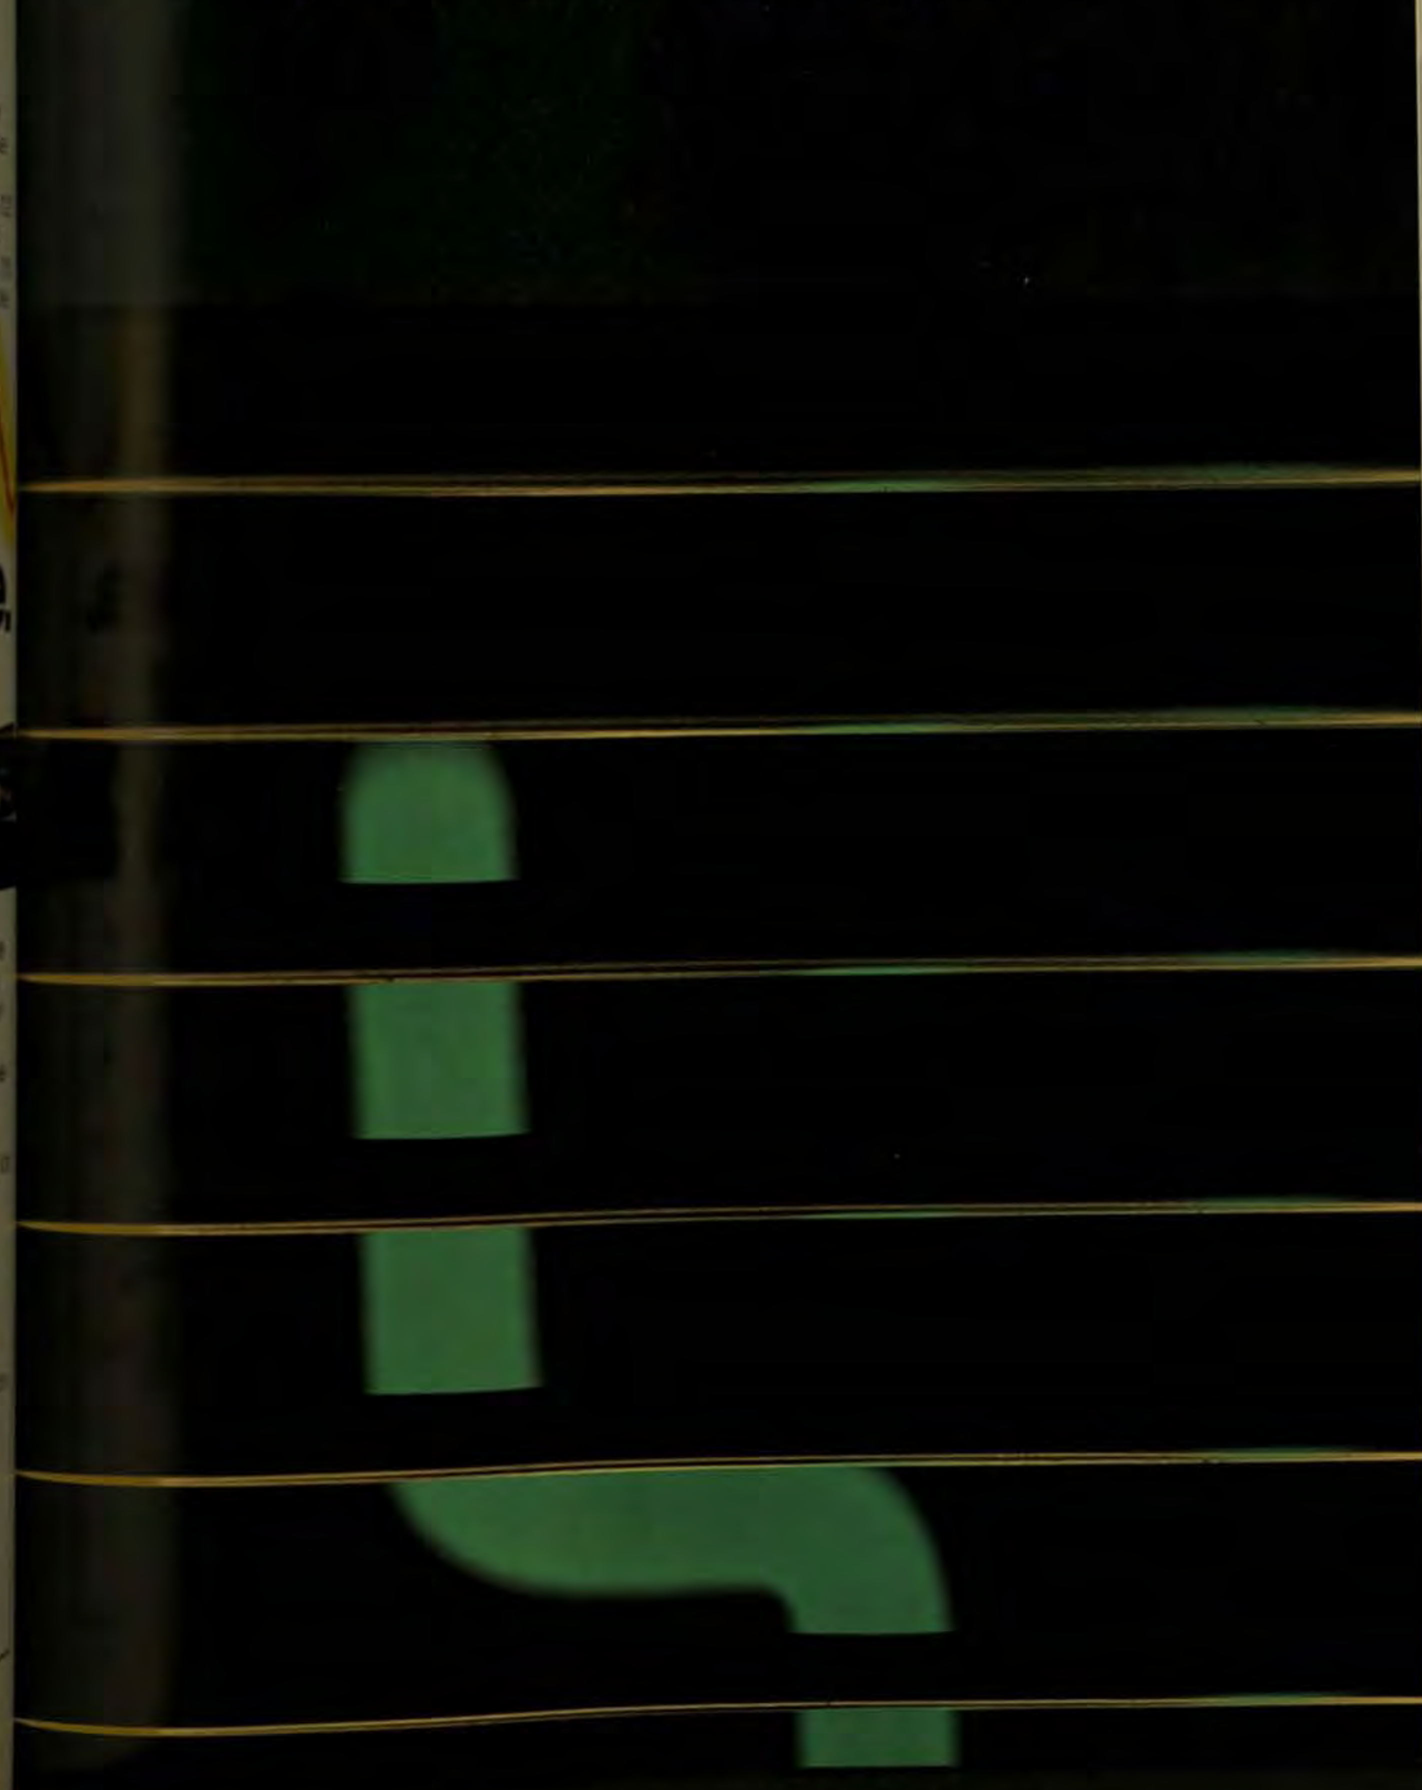

Compact, Fully Enclosed UPS Systems Move Into The Computer-Room

Manufacturers of UPS systems have developed products that are smaller, smarter, and more cost-efficient for today's smaller, more powerful computer systems.

by **William W. Campbell,** Exide Electronics

Recent advances in static power conversion and sealed lead-acid battery performance have resulted in affordable power protection for the small-system user. Uninterruptible power system (UPS) products are now available in packages designed for installation in a computer room. The size, weight, and appearance of these computer-room-installable systems and their associated sealed, maintenance-free battery packs offer UPS performance at more attractive prices.

These packages can be installed on standard, raised computer-room flooring and provide a 5- to 15-min. ac power-outage reserve. Their lower cost brings them more into line with non-UPS power-conditioning equipment while providing a complete solution to computer power problems. This new generation of UPS is especially attractive for retrofit applications to existing computer rooms and eliminates the need for specially constructed UPS equipment rooms.

QUALITY COMPUTER POWER

A recommended definition of computer-quality ac power appears in the *IEEE Orange Book* (IEEE Std. 446-1980), but the same text also acknowledges that not all computer-equipment manufacturers subscribe to the guidelines. In brief, quality ac power is provided to the computer room when:

- Electrical noise has been minimized.
- Transient or long-term overvoltage or undervoltage conditions are eliminated.
- Momentary or long-term power outages have no adverse affect on computer operations.

A variety of approaches are used to determine the presence of power-line disturbances within the computer facility. Existing electrical-noise levels, from low-frequency harmonics to high-frequency EMI, can be determined by measurement, but are often diagnosed and cured by on-site inspection and analysis of the electrical-distribution and grounding systems. Available products can continuously monitor the incoming ac service for transient and long-term over- and undervoltage conditions.

POWER-DISTURBANCE SOLUTIONS

Once you determine your needs, a wide array of products can provide the required degree of power conditioning. The minimum recommended solution for electrical-noise reduction is an isolation transformer that provides both a power-grounding system for the computer facility and a degree of electrical noise attenuation.

Line regulators are another common solution for ac service voltages that go outside of recommended operating ranges for extended periods of time. (Computer equipment operating ranges may vary from ±6 to +10, −15 typically; ANSI standard C84.1-1977 recommends +6 to −13%). Most line regulators fall into one of four broad categories of basic power-

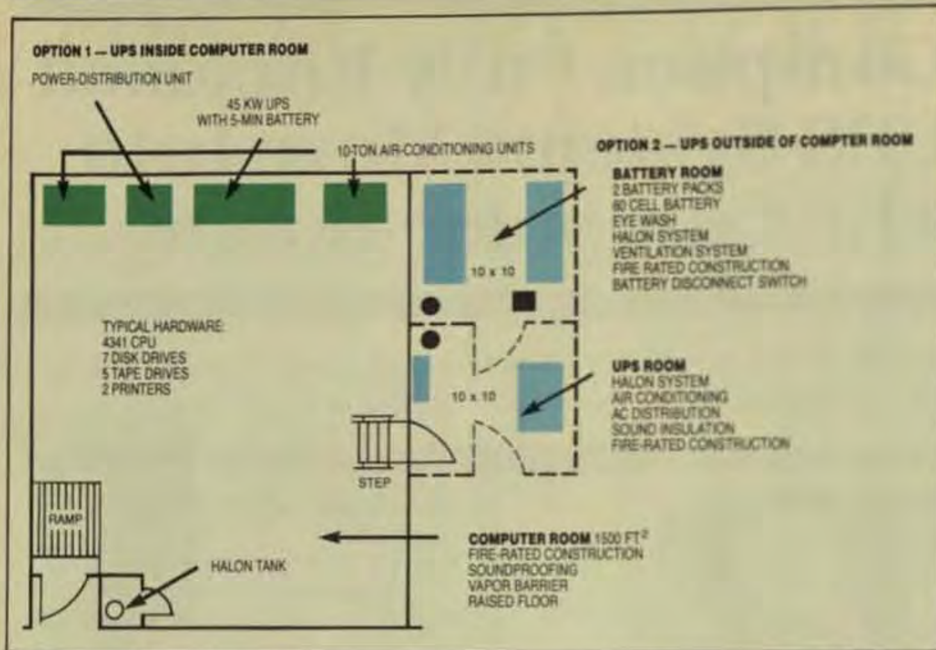

Fig Two uninterruptible power systems (UPS) options are considered—option one, in which the UPS is installed in the computer room using the new sealed, maintenance-free batteries, and option two, in which specially constructed rooms house the UPS and a wet cell battery.

circuit design: ferroresonant, electronic, and electromechanical or static tap changer.

The choice of line regulator is determined by the amount of correction needed, output-voltage regulation window, speed of response, and cost/benefit judgements. However, line regulators—with the exception of ferroresonant types—can't correct for ac service transients which are a subcycle to a few cycles in duration. The response speed isn't fast enough and the transient, perhaps somewhat attenuated, is passed through to the computer.

The next level of conditioned power is regeneration, either by rotating or static (solid-state) equipment. Regeneration eliminates subcycle line transients and includes grounding isolation and voltage regulation. Various degrees of ac-outage ride-through capability are also available with regeneration-type systems. In recent years, motor-generator sets packaged for a computer-room environment have become a popular, cost-effective choice and offer an ac ride-through capability of a subcycle to a few cycles.

To provide protection against longer-term ac outages (minutes to hours) a UPS, which uses the stored dc energy in batteries, is the most cost-effective solution available. For extended ac outages (hours to days), engine generators are normally incorporated to provide backup power to the UPS. Several rotary systems also offer charger/inverter/battery add-on packages that can be used to power the motor generator through longer-term utility losses of 30 min. or more.

POWER-DISTRIBUTION UNITS

An ac power-distribution unit provides an attractive, compact package for computer-equipment power distribution and eliminates the need to hard-wire your computer room in conduit. It's mobile and uses flexible plug-in cables that permit rapid low-cost installation or rearrangement of equipment. Since power-distribution units offer no industry-standard level of conditioning, you need to consider adding extra-cost options that may not be part of the basic unit. These include:

- **Isolation transformer**—Required to establish a separate computer-power ground system.
- **Surge suppression**—From a few cycles or less to 100% overvoltage.
- **Transient suppression**—For a subcycle event that may be 100 to 500% outside the regulation window.
- **EMI suppression**—Electrical noise in frequencies from low-order harmonics to megahertz.
- **Lightning arrestors**—Subcycle events that may reach thousands of volts.
- **Transformer electrostatic shielding**—For additional electrical noise (EMI) attenuation.
- **Line-voltage regulation**—Typical performance ranges accept a ±20% ac input to ±5% ac output.
- **Extra-wide input-voltage capability**—Typically +10% to −40% for sustained under- or overvoltage operation.

A power-distribution unit is also a convenient central location for the monitoring and recording of real-time conditions for incoming ac service and the ac power output, and for the collection and display of other computer-room and building alarms.

UPS FOR THE COMPUTER ROOM

Power-conditioning and power-distribution products can be readily installed in the computer room—with the notable exception, prior to 1982, of static UPS units. For UPS installation in a computer room, limiting factors—alone or in some combination—were the system's size, weight, appearance, audible noise, complexity of operation, and the need for wet-cell batteries.

Thus, UPS products didn't offer the ease of installation that had come to be expected with power-distribution units and rotary-power conditioners. Static-UPS products were introduced in 1982 with sealed maintenance-free batteries suitable for installation in new or existing computer facilities.

The versatility and cost-effectiveness of these new systems is illustrated by CompuDynamics Corp., a computer-facility planning firm in Springfield, VA. In this situation (which isn't unusual) a new computer facility was being built into an existing office/manufacturing operation located in a relatively new industrial park outside the city. The problems and solutions involved would also apply to a location in a suburban office center.

The computer-based activities of the example company had experienced performance problems as a result of ac-service disturbances, so a decision was made to install UPS protection for the firm's IBM 4341 system. Office areas are part of the plant and all ac power comes from one substation. The local utility is regarded

In the past, one of the hardest decisions was the choice of a power conditioning system . . .

Now the choice is very simple:
No. 1 EPE POWERBLOC
No. 2 EPE POWERBLOC/Distribution System
No. 3 EPE UPS Systems

EPE can simplify your power conditioning decisions by offering you the level of protection best suited to current needs and the flexibility to add incremental protection as your level of criticality increases. The effectiveness of EPE's "modular approach" to computer power conditioning is supported by the findings of a recently published national study by Bell Laboratories on AC power quality.

No. 1 EPE POWERBLOC
EPE's Powerbloc Conditioning System adds "outage protection" to EPE's Computer Power Centers or any other existing power distribution system. The Powerbloc Induction MG System was categorized as **95%** effective by the Bell Labs report.

No. 2 EPE POWERBLOC/Distribution System
Portable Powerbloc units which combine power protection and distribution in a single module can be installed directly on the computer floor. Categorized at **96%** effective, these modules provide the highest degree of protection possible for those applications where a UPS system cannot be justified.

No. 3 EPE UPS Systems
A complete line of UPS systems, coupled with EPE's Computer Power Center distribution units, can increase power conditioning effectiveness to **100% protection**.

EPE's Powerbloc II is the only system available to add battery backup to existing MG sets. This unique "upgrade" system is the only true off-line UPS available.

If you're faced with the decision of choosing a power protection system for your computer, we'd like to send you a copy of the Bell Laboratories Report "The Quality of U.S. Commercial AC Power." For your free copy, write or call us.

epe
EMERGENCY POWER ENGINEERING, INC.
3580 Cadillac Ave., Costa Mesa, CA 92626
Phone: 714/557-1636/ Telex: 183-535
The Leader in Computer Power Technology

Line power noise, interrupts & glitches got you down?

Noise from switching equipment, power interruptions and unsteady voltage are more than simple annoyances — they can cost you money. Errors and equipment damage are common results of faulty utility power. To avoid these problems rely on Computer Power, Inc. We offer a complete line of standby power and line conditioning systems. *And our price has the competition spiked.*

Computer Power has the system to handle virtually any line disturbance you encounter.

We're the UPS experts. Our UPS system protects against switching noise, lightning, brownouts and blackouts. At the instant of power failure, it automatically switches to standby power with absolutely no interruption. Models are available to 75 KVA. For isolation from common and transverse mode noise, our LCI series is your choice. Special separation between primary and secondary windings enable these units to deliver maximum isolation 10 times more effective than ordinary isolation transformers. The LCR series, our most cost-effective line conditioner, uses a ferro-resonant transformer to provide super high isolation, suppression of equipment switching noise and virtually foolproof regulation vs. input voltage changes. For line conditioning, isolation and utility voltage stabilization in a lower cost unit, our LCS series is ideal. Designed primarily for low cost mini and micro computer systems, the unit can be mounted on the floor under the computer.

Computer Power has the system you need and a price you can afford.

Whatever utility power problem you face, Computer Power has a system suited to it and your budget. Our brand new 32 page fact-filled catalog shows that we offer more for your money. Send for it today! COMPUTER POWER, INC., 124 West Main Street, High Bridge, New Jersey 08829 USA (201) 735-8000 Telex: 847481

COMPUTER POWER INC.

Analysis of a Typical UPS Installation

INSTALLED EQUIPMENT AT EXAMPLE INSTALLATION

Description	Use
IBM 4341, including remote access through modems	Material/manufacturing requirements planning
Dimension PABX	Telephones
Hewlett-Packard 9826 computer	Engineering
Apple IIe	Engineering
Texas Instruments micro-processor-development system	Engineering
Victor 9000	Marketing/Sales
Hewlett-Packard ATE	Manufacturing
Computer-controlled punch press	Manufacturing
Time clocks	Plant-wide
Supplier CRT terminals	Purchasing
Microprocessor-controlled test equipment	Manufacturing/ Engineering
High-speed collating copier	Office area
Timeshare computer terminal	Engineering

Computer-Room Features and Equipment

- Two 10-ton air-conditioning units (25% redundancy).
- Raised floor, 250 lbs/ft²-capacity, including extra floor pedestals under the UPS.
- Halon system.
- Power-distribution unit plus cables.
- AC power service.
- Room:
  - —Four walls, door, ceiling.
  - —Vapor barrier.
  - —Soundproofing.
  - —Fire-rated construction.
  - —Ramps, handrails, steps.

Total Installed Room Cost — $100,000

For the separate UPS/battery facility, the following requirements must be met:

- 10 × 10-ft UPS Room.
  - —Walls, ceiling, soundproofing.
  - —Fire-rated construction.
  - —Additional air conditioning.
  - —AC power service.
  - —Additional Halon.
- 10 × 10-ft Battery Room.
  - —Walls, ceiling.
  - —Fire-rated construction.
  - —Ventilation.
  - —Eyewash.
  - —DC service and battery-disconnect switch.

Total installed 200-ft² Room Cost — $31,000

INSTALLED UPS COST

OPTION ONE—In the computer room

• Exide Electronics Series 2000 45-KW UPS	
—5-min. sealed, maintenance-free battery for a critical load	$59,000
—Installation	$ 1,000

Total Installed Cost — $60,000

OPTION TWO—In a separate UPS/battery room

• Exide Electronics Series 2000 45-KW UPS	
—5-min. wet-cell battery for a 45-KW critical load	$63,800
—Installation	$3,000

Total Installed Cost—$66,800

COST SUMMARY

UPS installed in the computer room:	
• Computer-room cost	$100,000
• 45-KW UPS installed cost	$ 60,000

Total Cost — $160,000

UPS installed in a separate room:	
• Computer-room cost	$100,000
• UPS/battery-rooms cost	$ 31,000
• 45-KW UPS installed cost	$ 68,000

Total Cost*—$199,800

*The monthly expense of the extra 200-ft² needed for the UPS and battery rooms is an additional cost.

A Glitch Killer Is Not Enough!

Your Computer Needs Complete AC Power Line Protection

Sooner or later your computer will be subjected to power line problems that will result in costly loss of data or even damage to the equipment. You need protection against transients, noise and AC ground loops, also isolation and regulation. Exlin's Computer Power Conditioner will give you all that plus maximum reliability. It's cheap insurance and you pay for it only once.

Get all the straight facts about power problems and the solutions from Exlin.

Model EX 500P

Write for our FREE Computer Power Reference Guide

500 WATTS PROTECTION $375

250, 500, 1000 and 1500 Watt Models Available

Rackmount Models Too!

EXLIN POWER CONDITIONING EQUIPMENT

4888 Ronson Ct., P.O. Box 178471, San Diego, CA 92117 (619) 571-3134 CALL COLLECT

TABLE — COMPARISON OF POWER CONDITIONING PRODUCTS

MECHANICAL CHARACTERISTICS OF A VARIETY OF AVAILABLE POWER CONDITIONING PRODUCTS IN THE 50-KVA RANGE. THIS INFORMATION, INCLUDING LIST PRICES, WAS OBTAINED FROM CURRENT CATALOGS OF MAJOR SUPPLIERS DURING MAY 1983. THE INFORMATION DOESN'T REPRESENT ALL SUPPLIERS IN ALL CATEGORIES AND SHOULD BE REGARDED AS TYPICAL. LIST PRICES ARE SHOWN FOR REFERENCE COMPARISON. MIN/MAX VALUES ARE LISTED WHERE SIGNIFICANT VARIATION WAS FOUND. VERTICAL COLUMNS LIST MIN/MAX VALUES FOR THAT PARTICULAR CATEGORY. MIN OR MAX VALUES HORIZONTALLY ARE NOT NECESSARILY ALL THOSE OF ONE PARTICULAR POWER-CONDITIONING COMPONENT.

DESCRIPTION	RANGE	PRICE $K	RATED (KVA)	RATED (KW)	WEIGHT (LBS)	W (IN.)	D (IN.)	H (IN.)	FT.² (a)	EFF (%)	BTU	AMB. TEMP.	AUD. NOISE	NOTE
STATIC UPS WITH 5-MIN. BATTERY	MIN MAX	58.8 77.0	56.0 75.0	45.0 60.0	5000 9600	72 144	31 32	72 84	18.2 32.0	84 88	21000 35350	40 50	[illegible]	[illegible]
ROTARY POWER CONDITIONER WITHOUT DISTRIBUTION (f)	MIN MAX	22.4 27.6	50.0 50.0	50.0 50.0	2300 3000	48 83	32 36	60 60	11.3 13.3	84 88	23300 32500	40 40	[illegible]	
ROTARY POWER CONDITIONER WITH DISTRIBUTION	MIN MAX	26.0 34.1	50.0 50.0	50.0 50.0	3100 3500	48 72	32 34	60 60	15.1 16.1	84 88	23300 32500	40 40	[illegible]	
UPS CONVERSION PACKAGE FOR ROTARY UNIT C/O BATTERY-CHARGER-INVERTER (f)		36.3		50.0	3400	72	34	60	17.0	84.8	9800	40	[illegible]	[illegible]
FERROMAGNETIC POWER CONDITIONER		25.0	50.0	50.0	3600	74	35	72	18.2	82	14600	40	[illegible]	
POWER-DISTRIBUTION UNIT WITH ISOLATION TRANSFORMER (42 POLES)	MIN MAX	5.2 12.0	50.0 50.0	50.0 50.0	600 1100	25 36	24 36	42 60	4.7 7.8	96.1 97.5	3000 5100	32 40	[illegible]	[illegible]
LINE REGULATORS INCLUDING TAP CHANGER TYPES	MIN MAX	7.8 26.2	50.0 60.0	50.0 60.0	520 2000	28 48	21 27	42 60	7.5 14.2	84 88	3000 8000	32 40	[illegible]	[illegible]
HIGH-ISOLATION TRANSFORMER	MIN MAX	2.5 8.8	45.0 50.0	45.0 50.0	500 1200	22 53	17 31	17 37	2.9 11.4	97 98	3160 5260	45 55	[illegible]	[illegible]
50-KVA CABLE SET FOR POWER-DISTRIBUTION UNIT		5.5												

NOTE:
(a) SQ FT WILL NOT NECESSARILY BE THE PRODUCT WxD.
(b) FULL LOSSES ONLY OCCUR DURING POWER OUTAGE. PARTIAL LOSSES OCCUR DURING BATTERY RECHARGE.
(c) WIDE RANGE OF VALUES IS DUE TO OPTIONAL FEATURES.
(d) WIDE RANGE OF VALUES IS DUE TO OPTIONAL FEATURES SUCH AS HIGH-ISOLATION DESIGN AND POWER-DISTRIBUTION COMPONENTS.
(e) RANGE OF VALUES DUE TO DEGREE OF ATTENUATION AND OTHER FEATURES.
(f) THESE TWO COMPONENTS COMBINED FORM A ROTARY UPS.
(g) MIN VALUES ACROSS REPRESENT EXIDE ELECTRONICS' SERIES 2000 45-KW UPS.
(h) LISTED 75-KVA MODEL DOES NOT INSTALL IN THE COMPUTER ROOM.

as a reliable supplier, but plant operations create a poor electrical environment and brief, discernible power interruptions (a few cycles) occur about once or twice a week. Longer-term interruptions, perhaps 30 s or more, occur about once a month.

The accompanying box details the computer or computer-based equipment at the plant and features of the 1500-ft computer room. Additionally, the box details two UPS-based power-conditioning/power-protection options. In option one the UPS is located in the computer room and uses the new sealed, maintenance-free batteries. Option two uses specially constructed rooms for the UPS and wet-cell battery. The cost estimates are based on use of 45-KW UPS and includes the cost of the UPS with batteries, plus unpacking, moving, placement, and hookup.

The computer room layout in the figure is designed for a typical computer system and requires a 45-KW UPS for the following equipment: one IBM 4341 CPU, seven disk drives, five tape drives, two printers, and terminals, modems, and miscellaneous. A listing of various physical characteristics and list prices (based on May 1983 catalogues and vendor representatives) illustrates the product spread available from major power-conditioning-equipment suppliers. Using the Table, various other power-conditioning alternatives may be substituted for comparison.

SEALED LEAD-ACID BATTERIES

One of the key elements that make a computer-room-installed UPS possible is the development of the sealed high capacity lead-acid battery. Its primary features are minimal hydrogen gassing during recharge and a captured-electrolyte system that permits the sealed battery construction. The battery can be mounted in any position—even upside down if desired.

Before high capacity sealed cells were introduced (about two years ago), lead-calcium wet-cell batteries were the most cost-effective energy-storage systems that could provide ride-through capabilities for power requirements of 10 KW and above. To make proper use of these batteries a separate battery-room facility is needed for control of such items as hydrogen-gas evolution, accidental electrolyte (sulfuric acid) spillage, ambient temperature, and exposure to open high-voltage dc electrical conductors.

During recharge, wet-cell batteries give off hydrogen and oxygen, a mixture that becomes volatile when the hydrogen content reaches 4% of the enclosed volume. Safety codes require adequate ventilation to limit the hydrogen content of battery-room air to 1% or less. A number of variables are involved, but one to five air changes may be required per hour in a small, enclosed room.

By design, the new sealed lead-calcium cells produce minimal hydrogen during recharge, typically 0.01% of the total gas evolution. The rest is oxygen (99%) and other impurities (0.99%). The gas produced is normally recombined into the electrolyte system within the cell and is only vented (without loss of electrolyte material) when the internal cell pressure exceeds 7 to 10 psi. This condition may occur during a heavy recharge but results in gas concentrations well below the 1% limit, under normal usage.

Due to weight, volume, and cost considerations, sealed cells mounted in a compact enclosure are limited to UPS applications of about 50 KW which provide 5 to 15 min. of ride-through reserve. Distributing the weight of additional cells for extra reserve time is, of course, possible. Operating life for wet cells is typically [illegible] years, while a 5-year life is most common for the sealed cells. The sealed batteries are maintenance free, a major advantage over wet-cell types. ■

William W. Campbell, *product manager for the Series 2000 product line at Exide Electronics, has 15 years experience in engineering and marketing in the UPS industry. He holds a BSEE from Ohio Univ.*

Portable Power for Portable Computers

by David Dickey, Cuesta Systems

Truly portable microcomputers, small and light enough to be easily carried yet with powerful memories and displays, were initiated by the Osborne I in 1981. Since then, the Kaypro II, Otrona Attache, Grid Compass, and a dozen other portable computers have been introduced and have made it into full production.

However, few of these manfacturers were prepared to select appropriate battery power sources, to develop the required power inverter or converter, or to determine the size and demands of the market for their battery power options. Fortunately, independent firms have filled the void with products marketed by the OEM or directly to the end user.

Selection of the 12V rechargeable battery is between lead-acid and nickel-cadmium (NiCad). Sealed gel-cell types of lead-acid batteries are available from at least five manufacturers in shoulder-strap or belt-loop carrying cases with up to 12 Ahrs of storage capacity.

For portable microcomputers, which typically require from 40 to 80W of input power, an hour of battery operating time requires a 6- to 12-Ahr battery, weighing 4 to 8 lbs. NiCad batteries, while delivering about a third more power for the same amp-hour rating and weight, cost up to twice as much.

Battery power inverters (ac voltage output) or converters (dc volts) have performance requirements determined by the needs of the specific portable computer. The Osborne I was initially designed with a front-panel battery input connector that required well-regulated +12.5V dc and +5.5V dc for portable operation.

At the opposite end of the spectrum is the built-in battery-pack option for the Access portable computer system. This product includes a NiCad battery for 1 hr of operation, its own 2-step battery charger, a receptacle for an external plug-in 12V battery supply, a power-status red/green LED indicator and low-battery alarm buzzer, and a battery power converter that automatically steps in and provides uninterrupted prime power to the Access during line power loss or interruptions.

Some other portable computers being powered through their ac line cords require a 60-Hz battery inverter for their ac fans, or perhaps a 60- to 400-Hz inverter if a transformer type of power supply is used or a switching supply with an ac startup circuit. Tight frequency control, most likely from a crystal oscillator, may also be required to prevent annoying video jitter.

The market for portable power for portable computers is growing rapidly. These manufacturers estimate that from 5 to 20% of their customers intend to purchase a battery power option, and half of those include the carry-along battery. Manufacturers who consider their optional battery power needs early enough to affect the design have a head start, and they and the independent developers are busy pushing the state of the art to make portable computers truly portable. ■

UNINTERRUPTIBLE POWER... THE TOTAL SOLUTION

While some power conditioning devices are 65% effective against power problems and others 86% effective—Emerson's AP100 Uninterruptible Power Systems deliver 100% protection. Complete protection against line transients, voltage fluctuations and power outages.

The AP100, Emerson's most recent addition, continues a 20 year tradition of providing total solutions for power problems. Available in 5, 10, and 15 kva single phase ratings, the AP100 assures an uninterruptible source of power for micro/mini-computer systems, telecommunication networks, process control systems, security centers, and other applications requiring 100% power protection.

For complete information on the AP100 and the AP100-remote status panel, contact Emerson Industrial Controls Division, 3300 South Standard Street, Santa Ana, California 92702, telephone (714) 545-5581.

EMERSON Industrial Controls

The World's Largest Manufacturer of Uninterruptible Power Systems

© Emerson Industrial Controls 1983

Off-Line Micro UPS: a Cost-Effective Way to Keep Micros On-Line

High-speed transfer switches enable off-line UPS to provide essentially the same degree of protection as on-line systems, at a fraction of the cost.

by **Danniel H. White,**
Topaz Inc.

The high cost of conventional uninterruptible power systems (UPS) traditionally has limited blackout protection to mainframes and minicomputers. An off-line UPS that supplies power to the load only if the primary power fails or if voltage drops below a preset transfer point can save on both purchase and operating costs. However, this approach has been unsuitable because switching could not occur fast enough to prevent momentary lapses of power during transfer from the ac line to the UPS inverter. Now, the development of high-speed transfer switches has eliminated this problem.

All computers, from micros to mainframes, are vulnerable to ordinary ac power disturbances. Electrical noise, voltage fluctuations, and momentary power outages can cause a variety of costly computer problems, including data-entry errors, memory losses, hardware damage, and system shutdown.

Other than UPS, several types of power-conditioning devices are available to protect computers against these power disturbances. Isolation transformers and voltage regulators provide specific protection against electrical noise and voltage fluctuations, respectively. Power conditioners protect against both noise and voltage fluctuations. The type and degree of protection appropriate for a particular application depends on such factors as vulnerability of the application and the amount of money the user will spend.

The type and degree of protection appropriate for a particular computer application depend on a number of variables, including the criticalness of the application, the peculiarities of the ac power available at the user's location, and the amount of money the user is willing to spend for power conditioning.

Microcomputers, particularly those that use hard disks, are as vulnerable to power outages as larger systems, and microcomputer applications can be as critical as larger applications. In fact, blackout-related problems such as hardware damage and loss of valuable data can sometimes be more damaging to a small business than to a large one. Nevertheless, it simply isn't cost-effective for a business to invest $4000 or more in a conventional UPS to protect a $5000 to $10,000 microcomputer.

The rapidly increasing use of microcomputers has spurred efforts by the power-conditioning industry to solve this dilemma through creation of an effective, low-cost UPS. Success was recently achieved with the development of the fast-transfer off-line micro UPS.

Uninterruptible power systems used in computer applications traditionally have been on-line systems, which means that the UPS inverter is always on and the critical load always draws its power from the UPS, regardless of the condition of the ac power line. With off-line UPS, the load normally draws its power from the ac line, and the UPS is at rest. These systems are far less costly to produce and operate than on-line systems,

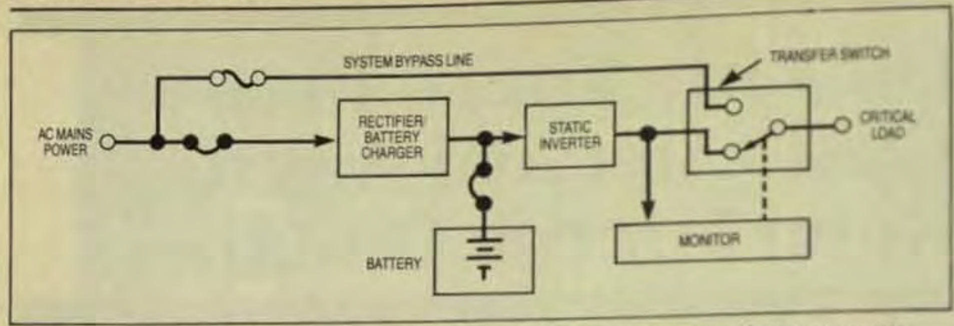

Fig 1 A typical on-line UPS consists of a rectifier/ battery charger, an inverter, a battery, a power monitor, and a bypass transfer switch.

but in the past, have suffered power lapses—though typically only a few cycles in duration—that were sufficient to cause serious computer problems because most computers can ride through only about one cycle of subtolerance voltage. The significance of the high-speed transfer switch in solving this problem can be more fully understood by comparing the operating principles of the two types of systems.

ON-LINE UPS

A typical on-line UPS consists of a rectifier/battery charger, an inverter, a battery, a power monitor, and a bypass transfer switch (Fig 1).

Under normal conditions, commercial ac power passes through the battery charger, which converts ac to dc. Potentially harmful noise transients are removed during this conversion. The dc power charges the UPS battery and then passes through the inverter, where it is converted back to ac to power the load.

If the primary voltage drops below the computer's tolerance level (typically, −10% of nominal), voltage from the battery is added to the primary voltage to boost the power back within tolerance. If primary power is lost completely, the load draws all of its power from the battery and continues to do so until the primary power returns or until the battery is depleted.

When commercial power is restored, the inverter automatically begins receiving dc power from the battery charger, and battery recharging begins. If the UPS malfunctions and the power monitor senses an ac voltage loss at the inverter output, the transfer switch automatically transfers the load to the commercial ac line, bypassing the UPS until the malfunction can be corrected.

The unit price of an UPS of this type typically is more than $4000. But the initial cost of the system is only part of its expense. Because the UPS is on line, it operates continually and therefore continually consumes power. This power consumption not only increases operating costs, but also produces heat inside the unit. This creates the need for a cooling system, thus adding to the cost of the unit. If adequate cooling isn't provided, the heat will eventually cause components to deteriorate, reducing system reliability.

Various attempts have been made to produce an inexpensive on-line UPS, but cost-cutting measures typically have resulted in system deficiencies. Some manufacturers, for example, have reduced product cost

Blackouts, Brownouts or Up-and-Downouts.

General Power Systems is the solution to costly power problems.

The unpredictable—a blackout has hit the city.

Thanks to General Power's uninterruptible power system (UPS), many businesses and vital emergency services continue to operate. The UPS provides conditioned power when changes occur in utility power caused by adverse weather conditions or any other unpredictable occurrences.

With a full range of systems, from 90 watts to 375 KVA, General Power can deliver clean power to any computer-based installation, all of the time.

Each General Power UPS is available in single or 3 phase versions and contains a battery, automatic battery charger, solid state power inverter, overload and short circuit protection, alarm annunciator, indicators and more.

Attractive to the eye as well as the pocketbook, it's easy to guard against down time and costly errors. Protect against blackouts and brownouts and insure that valuable data is never lost by including a General Power UPS in your system.

Put an end to your input power problems. Call General Power Systems today for the representative in your area, and ask for details on our full range of products.

GENERAL POWER SYSTEMS

1400 N. Baxter, Anaheim, CA 92806 (714) 956-9321 (800) 854-3469 TELEX 182283

Liebert
World leader in computer support systems

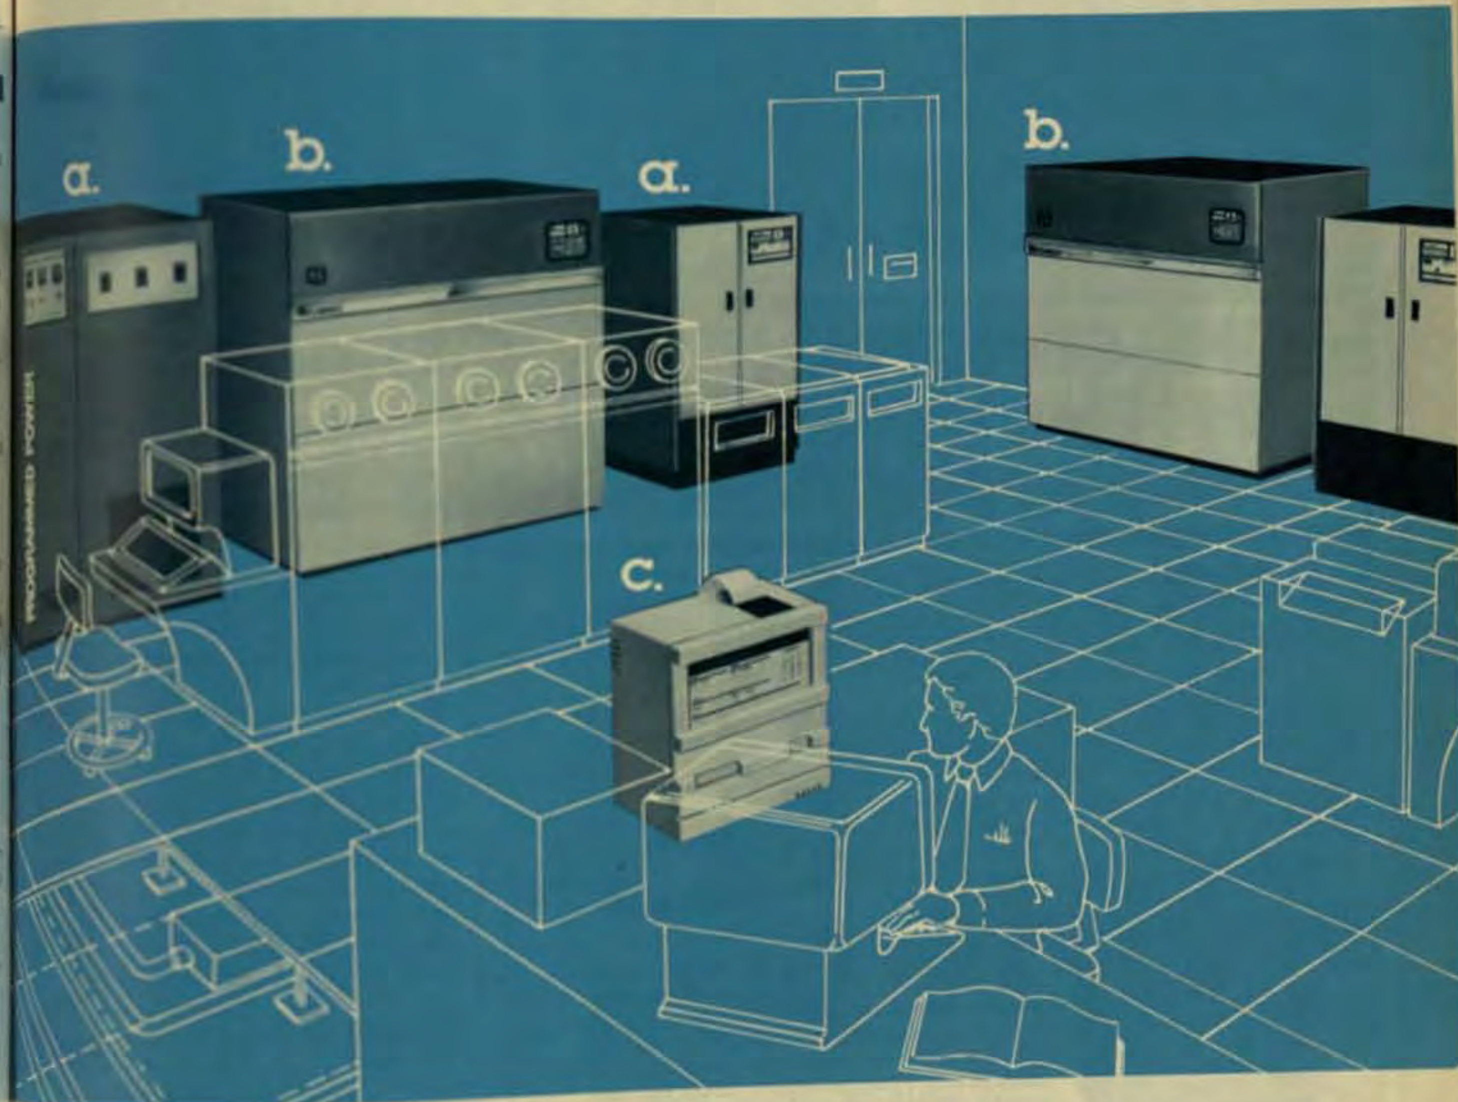

Is at your service with **products and people** that help make sure your computer investment pays off!

a. **Power Conditioning.**
Two outstanding Liebert products that provide 100% clean, computer-grade power.
Programmed Power Center™ for uninterruptible power and . . .
Power Management Center,™ a total electrical power system with unmatched cost-efficiency.

b. **Environmental Control.**
Liebert System 3, clearly the "industry standard" for effective control of computer room temperature, humidity and air cleanliness.

c. **All-system surveillance with Liebert Sitemaster.™**
A versatile, microprocessor-based, expandable monitoring system.

Contact Liebert today...because protecting your computer is as important as selecting it.

Liebert Corporation
1050 Dearborn Drive
P.O. Box 29186
Columbus, Ohio 43229
Phone: 614-888-0246
Telex: 246-655 LIEBERT WOGN

Liebert

Please send your free booklet on "Data Center Design"

Name
Title
Company
Address
City State Zip
☐ Have salesman call Phone CT83

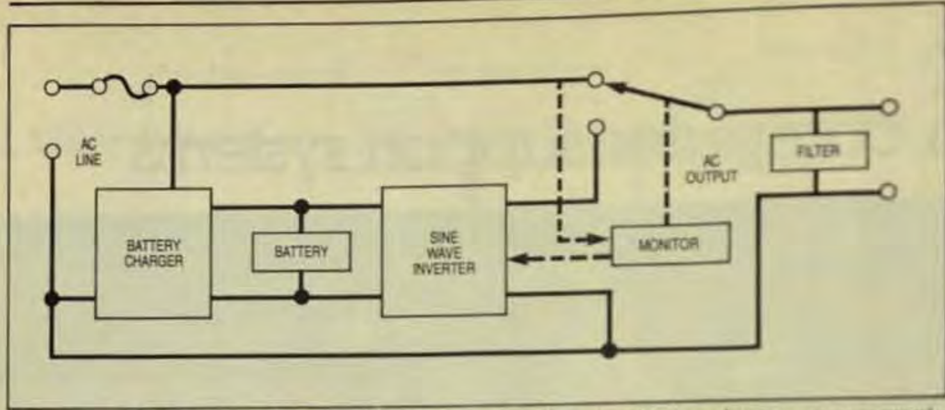

Fig 2 A typical off-line micro UPS consists of an inverter, a battery, a battery charger, a power monitor, and a high-speed transfer switch.

LINE CONDITIONING SYSTEMS BY HITRAN: 4 WAYS BETTER

1 More choices to fit the degree of protection you need. **UPS Series** uninterruptible power systems for the ultimate in protection against virtually all line problems, including prolonged blackouts. **SFR Series** sinusoidal ferroresonant regulation/isolation systems to effectively eliminate problems from brownouts, noise, lightning spikes, even short line breaks up to 1.5 cycles. And economical **UIS Series** ultra-isolation systems for isolation from common-mode and transverse-mode power line noise. All in a wide range of sizes from 500VA to 150KVA. Single and three phase. 60Hz and 50Hz.

UPS Series, 1000VA to 45KVA

UIS Series, 500VA to 150KVA

2 More reliability is built into each system. We design, build and, therefore, directly control the quality of all major components, such as transformers and circuit boards. We carefully select and conservatively rate other manufacturers' components to assure years of trouble-free operation.

3 More affordable. System for system, feature for feature, Hitran line conditioning systems cost less to buy, less to install, and less to maintain.

4 More available. Many frequently required sizes are stocked for rush shipment. Other systems are available with short lead times... because we make all the "impossible to get quickly" components in our own modern plant completely devoted to the manufacture of state-of-the-art, ultra-high reliability power systems.

So, if you are looking for the best system money can buy... as well as the best system for the money... call or write:

HITRAN
Division of AGEX

Route 31, RD 5-65
Flemington, New Jersey 08822
Tel: 201-782-5525
Telex: 833701/AGEX/FLEM

by eliminating the power monitor bypass section of the system. This does reduce the product cost, but without a bypass capability any UPS failure will result in shutdown of the critical load. This vulnerability, which is not much better than a vulnerability to blackouts, greatly diminishes the value of the UPS.

OFF-LINE UPS

One way to reduce the cost of uninterruptible power systems significantly without decreasing their effectiveness is to use off-line technology. A typical off-line micro UPS consists of an inverter, a battery, a battery charger, a power monitor, and a high-speed transfer switch (Fig 2).

Under normal conditions, the UPS inverter is at rest and the primary ac power passes straight through the UPS to the load. In most systems, a low-pass filter is placed across the UPS output to remove noise transients and voltage spikes. If the power monitor senses a drop in the primary voltage below a preset transfer point (typically, −10% of nominal), the inverter is automatically switched on and supplies steady ac power to the load from energy stored in the UPS battery.

New high-speed switches can complete this transfer in as little as 4 ms. When commercial power is restored, the load is automatically switched back to the ac power line and the battery charger begins recharging the battery.

By rapidly transferring the load from the ac line to the UPS inverter whenever voltage falls below computer tolerance, off-line systems provide protection against momentary power outages and voltage fluctuations. With filtering on the output to eliminate electrical noise, these systems can provide all the protection of on-line systems, at a fraction of the cost. An 800 VA unit, for example, can be purchased for less than $900.

Off-line systems cost much less to produce because the inverter is never powered by the ac line, so a large, high-powered rectifier/battery charger isn't needed. That's a major cost item in on-line systems. Also, since the off-line UPS operates only during power outages or voltage sags, its circuits receive very little stress and less expensive circuit components can be used. At the same time, the mean time between failures

Electric Utilities Do a Remarkable Job
... in delivering good, dependable power for millions of end uses, from light bulbs to ore smelters. But for computers power must be exceptionally clean and stable. Voltage variations that are harmless to other equipment easily cause faulty computations, data loss, violated programs and even damaged computer circuitry.

When Your Light Bulb Flickers
... you have a visible demonstration of a power line variation. It's harmless to your light bulb, but similar voltage fluctuations are responsible for untold computer problems often unjustly blamed on software or equipment.

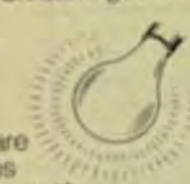

An Ounce of Prevention
... can save you a peck of power-related computer problems. ISOREG Computer Power Modules isolate computer equipment from raw utility power to keep out voltage spikes and electrical noise. At the same time, ISOREG Modules regulate voltages during sags and surges, long brownouts, and even ultra-brief interruptions. ISOREG Modules allow computers to operate without trouble even when voltages drop 30% below normal.

WHY ISOLATE?

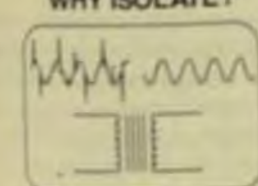

Voltage spikes are blocked when the computer is isolated from raw utility power.

WHY REGULATE?

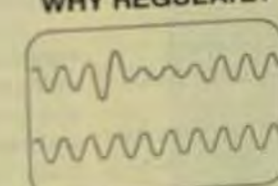

Sags, surges and brownouts are eliminated when the voltage is well regulated.

Full protection from raw utility power requires BOTH isolation AND regulation. You get it from ISOREG.

ISOREG and ISOSHIELD are Registered Trademarks of ISOREG Corporation.

Don't Be Misled Into Half-Solutions
Electronic surge suppressors only tackle spikes exceeding "clip voltage levels" that are considerably in excess of normal voltages. "Spike clippers" and "surge suppressors" do nothing to eliminate electrical noise and surges below the "clip level" and do nothing whatever to get rid of sags, glitches and brownouts.

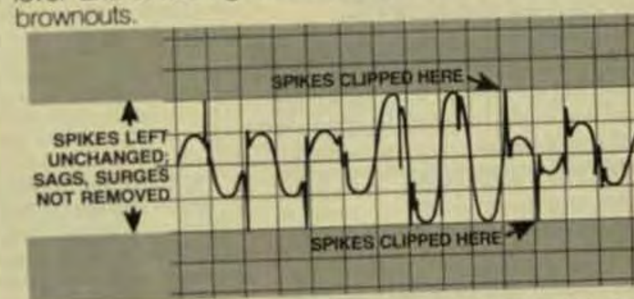

ISOREG Modules protect your computer from ALL these bothersome irregularities and from every type of spike—low or high level.

Do The Job Right
... when you set out to protect your computer from the ups and downs of "raw" utility power: Isolate and regulate—with equipment from ISOREG. We offer small modules for microcomputers, others for minis, for mainframes and entire computer rooms.

Get Full Information Free Call Toll Free: 1-800-225-5078.
In Massachusetts call 617/486-9483.
Or use the Reader Service Card.
ISOREG Corporation, 410 Great Rd.,
Littleton, MA 01460 USA.
TWX 710-347-6974 ISOREG LTTN

ISOREG CORPORATION

ISOREG Clean Stable Power for Computers

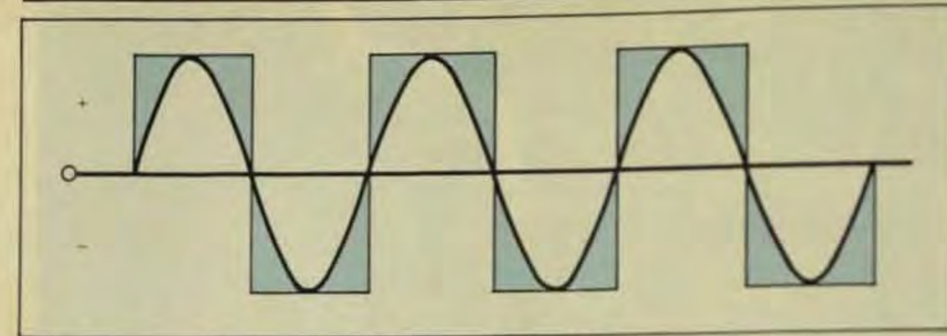

Fig 3 A square wave is a sine wave plus a large number of odd harmonics.

(MTBF) far exceeds that of on-line systems.

The cost advantage of off-line systems applies not only to purchase price, but to operating cost as well. An off-line system obviously uses less energy than an on-line UPS, since most of the time the inverter isn't operating. When the battery is fully charged, the only energy consumed by an off-line UPS is that required to power the monitor and control circuits. As a result, the typical off-line UPS is about 99% efficient, whereas a typical on-line system is only 65% efficient.

SELECTING AN OFF-LINE UPS

In selecting an off-line UPS, there are at least five basic features to consider:

- Transfer time
- Output waveshape
- Backup time
- Styling
- Unit cost

As mentioned earlier, most computers can ride through one cycle of subtolerance voltage. At 60 Hz, one cycle lasts 16 ms, so any UPS with a transfer time slower than 16 ms will leave a computer vulnerable to sudden voltage fluctuations. Typical transfer times for off-line systems equipped with high-speed transfer switches are a maximum of 10 ms—well within the range required for complete protection.

The waveshape of the ac power provided by a UPS is also very important, because the harmonic content of the ac signal can affect a computer's performance. Among the off-line systems currently available, some produce a square-wave output and others produce a sine wave.

A square wave is a sine wave plus a large number of odd harmonics (Fig 3). These extra harmonics can adversely affect certain types of printers and may distort visual displays. Thus, users should determine the harmonic tolerance of the system they wish to protect, before they choose a square-wave UPS.

Any computer system will operate satisfactorily from a sine wave. But the quality of the waveshape is critical. There are many different types of sine-wave inverters commonly used in UPS. Some produce a clean waveshape that nearly matches that of commercial ac power. Others, however, create harmonic distortion that can impair computer performance.

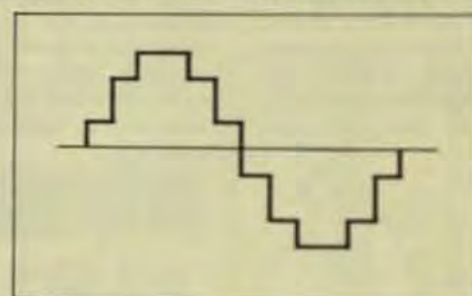

Fig 4 Step-wave inverters use square-wave generators that are controlled in phase and amplitude to produce a step wave. Usually the steps are evenly spaced across 360° and the height of each step is selected to eliminate low-order harmonics.

Among the most effective types of sine-wave inverters are those that employ ferroresonant technology, step-wave technology, or pulse-width modulation. Of these three, pulse-width modulation probably is the one best suited to off-line micro systems.

Filtering inherent in ferroresonant transformers enables ferroresonant-type inverters to produce a consistently clean sine wave by use of very simple circuitry. Ferroresonant inverters are, therefore, highly reliable. They are also bulky and inefficient, however, and because they are slow to respond to load changes, ferroresonant inverters aren't suitable for use in off-line systems.

Step-wave inverters employ several square-wave generators that are controlled in phase and amplitude to produce a step wave. Usually, the steps are evenly spaced across 360°, and the height of each step is selected to eliminate low-order harmonics (Fig 4). The remaining harmonics are removed by a relatively simple low-pass filter. Step-wave inverters produce a clean sine wave and are much more efficient than ferroresonant systems. Their high parts count, however, makes them economically unfeasible for use in low-power micro systems.

Inverters that employ pulse-width modulation produce a fundamental sine-wave output through the use of LC filters that eliminate low-frequency harmonics. The inverter also is designed to generate a reference sine wave, the precise shape of which is preset and unchanging. This reference signal is constantly compared with the inverter output, and if a distortion in the output waveform occurs, an error-signal feedback system electronically modulates the pulse width of the output to match the reference (Fig 5).

The error-feedback system maintains a clean, distortion-free sine wave despite load changes that might otherwise affect the output. And because this inverter technology is totally electronic, it is reliable, highly efficient, and economical.

Another feature to consider in the selection of an off-line micro UPS is backup time. Backup times vary considerably from product to product, but virtually all systems currently available meet minimum requirements. In fact, many systems actually offer more battery time than is normally needed.

Studies on the incidence rates and duration of voltage fluctuations and power outages indicate that about 90% of all power outages and severe voltage sags last less than 10 min., while 50% last less than 6 s. In that case, a micro UPS that provides 10 min. of backup power will allow equipment to ride through nearly all potentially harmful low-voltage conditions. This 10 min. also is sufficient to enable computer operators to conduct an orderly and controlled computer shutdown in the event of rare outages that last longer than that.

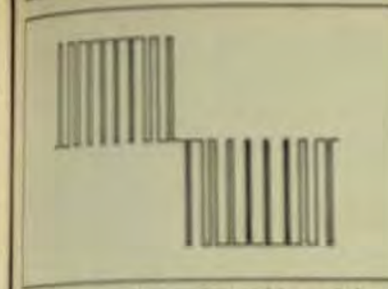

Fig 5 Inverters that use pulse-width modulation produce a fundamental sine-wave output through the use of LC filters that eliminate low-frequency harmonics. An error-signal feedback system electronically modulates the pulse width to match a desired reference.

The purpose of an UPS is to prevent system damage and loss of valuable data due to transient power disturbances—not to allow continued computer operation during prolonged losses of primary power. Highly critical applications that require this capability typically must employ a protective system that combines an uninterruptible power system with a motor-generator set.

A corollary to backup time is battery recovery time. This specification is not highly significant, however, because voltage sags or outages of more than a few seconds duration rarely if ever occur in close succession. Therefore, rapid battery recovery is not essential. Most UPS manufacturers use either nickel-cadmium (Ni-Cad) or lead-acid batteries, and both types typically can be recharged to 85% of full capacity in 2 to 3 hrs, which generally is more than adequate for complete low-voltage protection.

UPS styling, though certainly not a critical factor, is a feature worth noting. Many different micro UPS styles are available. Some systems are designed for industrial use, and many are styled for office environments. Size and weight are probably the most important aspects to consider, especially if the computer to be protected is a desktop model. For maximum convenience in such cases, the UPS selected should be small enough and light enough to be placed on or under the desktop, near the computer.

Perhaps the most important of all micro UPS features is price, simply because the specific purpose of these systems is to provide a cost-effective way to protect inexpensive microcomputers. Cost-effectiveness may be influenced to some extent by the nature of a user's application, but as a rule, the maximum cost-effective investment for an UPS is about 10% of the cost of the system to be protected.

The fast-transfer off-line micro UPS is an important technological development—nearly as important, perhaps, as the development of microcomputers themselves. Microcomputers have made affordable the many benefits of computer technology. The off-line micro UPS makes affordable the protection necessary to prevent power disturbances from taking these benefits away. ■

Danniel White *is engineering research analyst for Topaz Inc. He holds a BA degree from Eastern Washington Univ. and an MA from Portland State Univ.*

GARBAGE IN? PURIFIED POWER OUT!

FOR RELIABLE COMPUTER OPERATION

SERIES 800 POWER PURIFICATION SYSTEMS

GUARANTEED PROTECTION *from*
- Transients
- Noise
- Brown Outs
- Line Surges
- Other Line Disturbances

CONSTANT VOLTAGE TRANSFORMERS FOR LINE VOLTAGE REGULATION, CONDITIONING AND PURIFICATION *for*
- Mini Computers
- Large Computers
- Micro Processors
- Numerical Control Systems
- Communication Systems
- Programmable Controllers
- Other Sensitive Electronic Equipment

And other products to prevent costly malfunctions and downtime from input-power disturbances:
Series 500 - Line Voltage Regulators
Series 600 - Max-E-Isolation Transformers
Series 700 - Line Voltage Conditioners
Series 900 - Electronic Line Voltage Regulators

In The United States
Call Toll Free
1-800-521-4792
CONTROLLED POWER COMPANY
1955 Stephenson Highway
Troy, Michigan 48084
313-528-3700 TWX 810-232-3401

Replacing Bipolars With MOSFETs Improves Switching Power Units

Power MOSFETs have yet to achieve cost parity with bipolar transistors but they improve the reliability and performance of switching power supplies.

by Peter Wood,
International Rectifier
Semiconductor Division

The recent introduction of power MOSFETs has provided a significant improvement in the reliability of switching power supplies. Today, with more and more semiconductor manufacturers entering the power MOSFET arena, the initial objections of high cost and poor availability have largely disappeared. Despite their higher cost, the inherent performance advantages of these MOSFETs compared to bipolar devices make their use very attractive (see Table).

The majority of the switching power supplies manufactured are sold to OEMs who are integrating computer systems of varying sizes and types. These computers and peripherals all operate from standard dc voltages, which can be provided by the type of power supply that will be described in this article. Although the intent of this article is not to provide a cookbook recipe for an optimized power supply, the approach taken illustrates the use of small (low-cost) MOSFETs to process subtantial amounts of power. The current-driven configuration minimizes electrical stresses on the power switches and enhances reliability because of the benign operation of all active and passive components in the power train. This design yields an overall efficiency approaching 80% at full load and the control loop was stable for all line and load conditions.

A medium-sized computer requires 250W of processed power with the following specs:

- +5V dc, 10 to 20A with a total ripple and regulation of ±50 mV, overvoltage protected.
- +12V dc and −12V dc, 0 to 1 A with a total ripple and regulation envelope of ±100 mV and a common return.
- +26V dc, 1 to 3A with a total ripple and regulation envelope of ±1V and a common return.
- Input-voltage range 95 to 130V ac and 190-260V ac at 48 to 420 Hz.

Overcurrent protection is required on all outputs, and the power supply should have a minimum efficiency of 75% at full load. Other requirements such as VDE compliance and EMI attenuation are also required by most users, but are not considered in this article because they have been extensively covered elsewhere.

CIRCUIT APPROACH

To assure low cost and high efficiency and reliability, several design choices were made, including:

- A current-drive configuration was chosen to minimize the voltage stress on as much of the circuit as possible.
- The power switch was operated at a constant 50% duty cycle to maintain the best regulation of the unsensed output voltages.

- To reduce the number and cost of magnetic components, no output-filter chokes were used and the single primary inductor performed the dual functions of pulse-train integrator for the regulator and current source for the power switch.
- 50-KHz operation was chosen as the best compromise between conventional magnetic designs and minimum size (cost) magnetic components.
- 115/230V ac operation was obtained by use of a voltage doubler to supply a nominal 300V dc bus in the 115V input mode and a conventional bridge rectifier for the 230V input.
- Control-loop stability and response were enhanced by the use of a pulse-width modulator (PWM) with feed-forward capability so that line-voltage variations were regulated independently of the closed-loop control.
- 45V Schottky rectifiers were adequate to supply all outputs, because all input-line conditions are controlled by the front-end buck regulator.
- No power-wasting snubber circuits were used.
- Soft startup, under- and overvoltage protection, and current limiting were all provided by the PWM.

A block diagram representing the power flow of this approach is shown in Fig 1.

In the low input-voltage (115 ac) mode, ac power is applied between one of the input lines L_1 or L_2 and the neutral N. The reservoir capacitors C_1 and C_2 are each charged to the peak voltage of the ac waveform (approximately 150V dc with 110V ac input), and since they are in series, the total unregulated voltage is around 300V dc.

When a 230V ac input is used, L_1 and L_2 become the input terminals, and the rectifier functions as a conventional bridge circuit, yielding a 300V dc bus across capacitors C_1 and C_2.

In-rush surge currents are lowered somewhat by the use of surge-limiting thermistors, which have a negative temperature coefficient and minimize their dissipation under steady-state conditions.

50-KHz BUCK REGULATOR

The regulation of output voltages against line, load, and temperature effects is performed by the buck-regulator stage of the power supply (Fig 2b). PWM drive signals are impressed on transistors Q_1 and Q_2, which conduct on alternate half cycles of the drive waveforms. The switching function could be performed by a single transistor except that the duty cycle ratio needs to range from 0 to almost 100%.

Besides this wide duty-cycle capability, the use of two transistors affords two separate heat paths for power dissipation and results in smaller

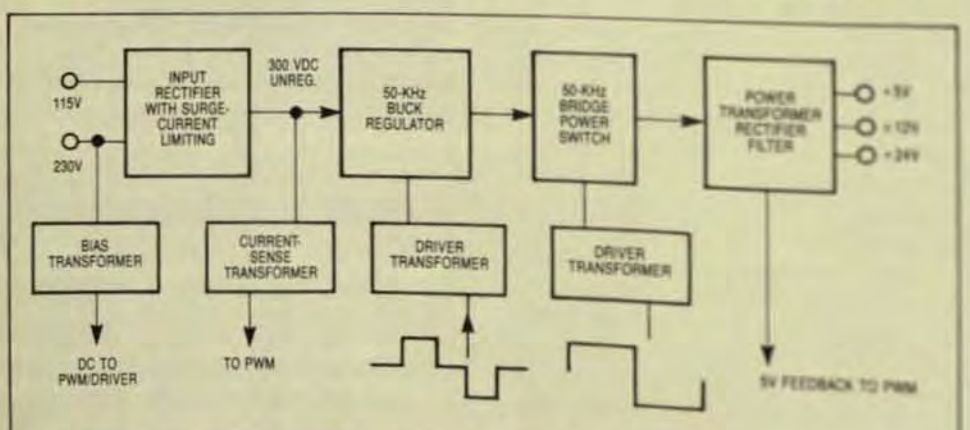

Fig 1 This switching power supply design incorporates MOSFETs for low cost and high efficiency

INPUT RECTIFIER

Dual input-voltage capability is achieved by the use of a diode bridge with split-reservoir capacitors (Fig 2a).

MOSFET switches and a much simpler drive circuit. (It can sometimes be less expensive to specify two smaller-chip MOSFETs than one chip with twice the active area.) The switches are placed in the negative bus so that capacitively coupled switching spikes don't appear in the drive circuits or current transformer.

The pulse train appearing at the inductor L_1 is at 100 KHz and is commutated by diode CR_1. No dc filter capacitor is used across the output of the regulator, since a high output impedance should be presented to the power-switching inverter circuit. In a current-fed converter, the output switch should have the following properties:

- DC bus current flows continuously.
- DC bus voltage is collapsible at constant current.
- DC bus voltage is transient-clamped to prevent bus overvoltages caused by high-impedance produced by leakage inductance from the power transformer.

The full bridge circuit was chosen (even though a push-pull circuit could have been used) because of its lower switch-voltage stress properties and its superior transformer utilization factor. Also, because MOSFETs are voltage-driven, the transformer driver is extremely simple. Unlike voltage-fed PWM converters, current-fed units don't need commutation diodes across each power switch because the power transformer always looks back into the closed switches that feed the primary (Fig 2c).

The bridge circuit delivers an ac output of the same magnitude as the dc input bus and is driven so that Q_1 and Q_3 conduct on one-half cycle and Q_2 and Q_4 conduct on the alternate half cycle. Thus, the power dissipation is shared by all four power devices, which can now be specified with half the active area that would be necessary with a 2-transistor design. For devices of equivalent total chip area, the 4-transistor design is more cost-effective.

With any practical power transformer, leakage inductance is always present, and in the case of a current-driven converter, the high source impedance of the power switch [illegible] ing the switching transitions [illegible] voltage spikes to appear on the [illegible] bus. A spike clipper comprising [illegible] C_3, and CR_1 connected directly [illegible] the dc bus effectively attenuates these spikes due to its low [illegible] impedance. The power dissipated [illegible] R_1 is directly proportional to [illegible] nitude of the leakage inductance [illegible] T_2, so it's important to [illegible]

IF YOUR DATA'S WORTH REMEMBERING, IT'S WORTH PROTECTING.

POWERMAKER® MICRO UPS

A split-second blackout or a sudden voltage sag can shut down your small business computer, completely wiping out critical data. Inventories, payrolls, receivables — whatever is in the memory may be lost instantly.

Although this type of data is just as important to a small business as it is to a large corporation, blackout protection has always been far too costly for small business applications. But now there is the Powermaker Micro UPS, an inexpensive standby power source specifically designed for small business computers.

This new rechargeable power system provides up to 35 minutes of steady sine-wave power, enabling even the most sensitive small computers to ride through blackouts and voltage sags completely unaffected.

Why sine-wave power? Because square-wave power impairs the performance of many printers, viewing screens and timing circuits. Powermaker produces a sine wave that exactly matches the wave shape of commercial power, ensuring compatibility with any computer system.

In addition to providing highly reliable blackout and brownout protection, Powermaker also protects against electrical noise, one of the major causes of computer errors and system malfunction. Powermaker is portable, completely automatic, maintenance free and plugs into any standard 120V outlet.

No matter how small your computer, your data is worth remembering. Protect it with an affordable Powermaker Micro UPS.

For complete information about the new Topaz Powermaker Micro UPS, please fill out this coupon or call us.

TOPAZ
SOLUTIONS TO POWER PROBLEMS™

TOPAZ; 9192 TOPAZ WAY
SAN DIEGO, CA 92123-1165
PHONE: (619) 279-0831
TWX: (910) 335-1526

☐ Please send me complete information about Powermaker®
☐ Please have a representative call me.

Name ______
Title ______
Company ______
Address ______
City ______ State ______
Zip ______ Phone ______

17

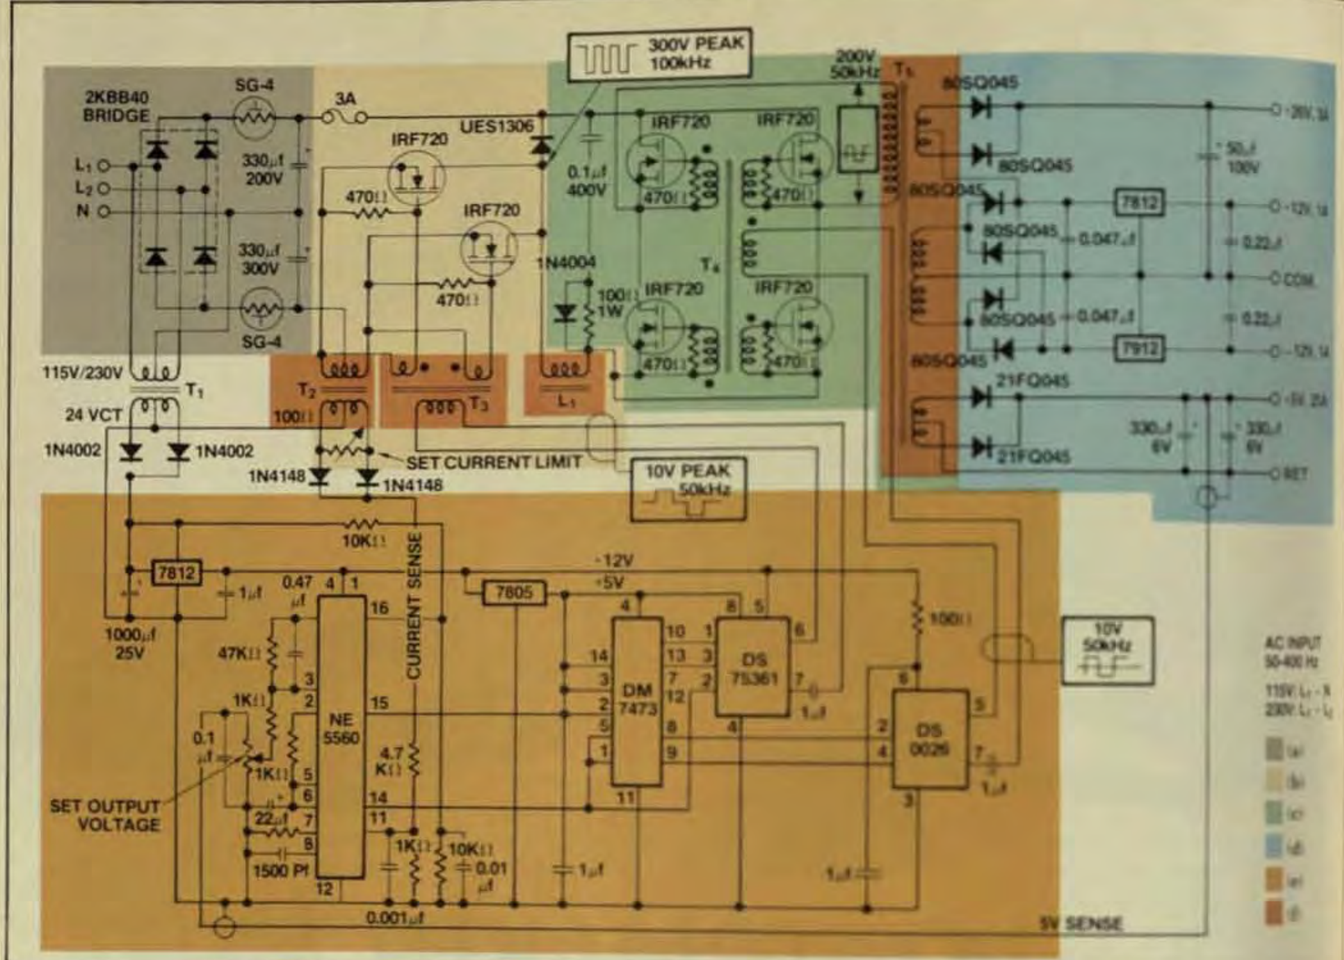

Fig 2 This schematic of a switching power supply using power MOSFETs shows the different components: input rectifier (**a**); buck regulator (**b**); power switch (**c**); output circuitry (**d**); pulse-width modulator and driver (**e**); and current, drive, and power transformers (**f**).

age inductance by careful magnetic design.

OUTPUT RECTIFIERS AND FILTERS

One of the main advantages of the current-driven square-wave converter is that all the secondary output voltages of the power transformer are accurate functions of turn ratios only, so the reverse ratings of the rectifier diodes don't need to accommodate any ac input-line considerations. As a result, 45V Schottky rectifiers can be used for all output voltages of this power supply (Fig 2d).

In the current-fed configuration, filter chokes at the dc outputs are replaced by a single inductor on the primary side of the power switch. Because of the very coarse turns resolution of the power transformer (approximately 5.6V per turn), it's necessary to regulate the 12V outputs with three terminal dc regulators. However, since the transformer output voltages are independent of ac input-line variations, the overhead voltages of the 12V regulators are small, and efficiency is not seriously compromised.

The +26V output is derived from the +12V regulator input voltage (approximately 16.2V dc) added to another 10.6V dc rectifier output, making a total nominal voltage of 26.8V dc. Schottky rectifiers are used here also. Since the bulk of the output power is in the 5V output, a feedback voltage for regulation is taken from this output.

PWM AND DRIVERS

In addition to the usual analog-to-pulse-width functions, the PWM control circuit (Fig 2e) has a feed-forward function that allows a constant volt-second output—a very desirable characteristic, since it reduces the gain requirements of the closed-loop control functions and hence simplifies control-loop stabilization. In this power supply, the feed-forward control voltage is derived from the dc bias supply obtained from a small line-frequency transformer.

The output pulse train from the PWM is at 100 KHz and is used to trigger a dual flip-flop in order to generate the required 50 KHz drive waveforms. Two different drive waveforms are required—for the buck regulator and the power switch, respectively. When power MOSFETs are driven through transformers, the accuracy of the drive waveforms must be preserved because MOSFETs can easily be damaged by excessive gate-voltage spikes that arise from poor transformer design or from drivers with uncontrolled high-impedance outputs.

The drivers in the power supply both have totem-pole output stages which always present a low impedance to the drive transformers. Also, the dc supply to these drivers is regulated at 12V, so that under high or low ac input-line conditions, the drive amplitudes are constant.

MAGNETIC-COMPONENT DESIGN

The choice of operating frequency depends largely on the types and complexities of magnetic components. Th

WITH A GOULD UPS, YOUR LAST RESORT WON'T BE THE FIRST THING TO BREAK DOWN.

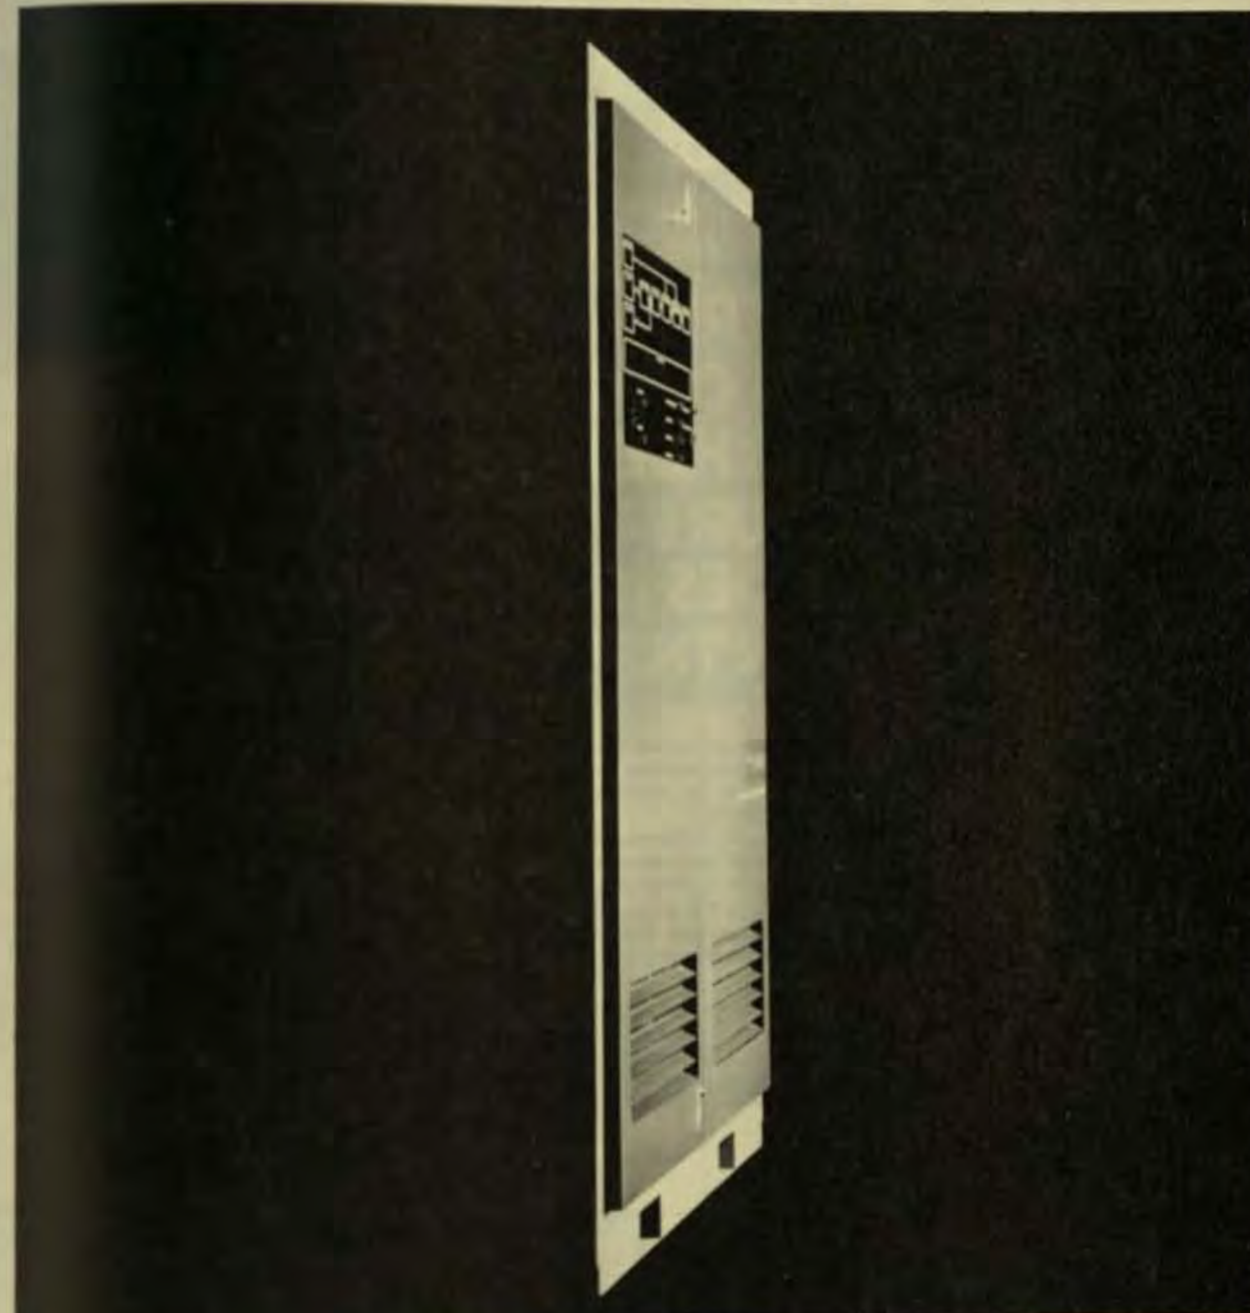

When your power fails, you need an UPS (uninterruptible power system) that won't. Because blackouts do more than leave you in the dark. They leave your computer without a memory.

Gould technology brings you an UPS that keeps your computer running smoothly throughout blackouts, brownouts, spikes, surges. In short, any disturbance on your power line.

Our UPS will always keep you in power because we keep our design simple. We use the minimum number of conservatively rated parts. That means fewer parts to break down. And fewer parts to keep up.

The result of this simple logic is an MTBF that exceeds 100,000 hours. And Gould uninterruptible power systems are available in ratings from 500VA to 45KVA.

So if you need an UPS you can count on under any conditions, put Gould to work. And discover the power of simplicity.

Gould Inc., Power Conversion Division, 2727 Kurtz Street, San Diego, CA 92110. Phone (800) 854-2658. In California (619) 291-4211. TWX (910) 335-1241.

GOULD
Electronics

like sine-wave transformers, switching power supplies demand wide bandwidth designs capable of supporting not only the fundamental switching frequencies but also the fast wavefronts associated with efficient power transfer.

The circuit isolation obtainable through transformers applies both for dc conditions, and for switching conditions where capacitive coupling between windings or even within a winding can cause unwanted ringing or common-mode spikes.

If switching frequencies are too high, leakage inductances cause inefficient operation because of dissipation in the snubber circuits that control voltage spikes. Conversely, if switching frequencies are too low, magnetic components become larger and the increased winding capacitances add to the common-mode problem, not to mention the additional cost of the magnetics. For the above reasons, it was decided that 50 KHz would be a satisfactory compromise between size, ease of winding, available cores, and cost.

CURRENT TRANSFORMER T2

Current from the unregulated dc bus is sensed by a current transformer (Fig 2f) in series with the buck-regulator switches. The primary winding for this toroidal transformer is made by passing the drain connection leads once through the core only, thus forming a bar primary. Errors in current sensing occur because of variations in magnetizing current at different flux levels, so two design considerations are important:

- The core hysteresis loop must be narrow so that magentizing current is small and low currents can be measured.
- The design's flux level must be small to minimize current-sensing errors.

The use of two switching transistors in the buck regulator allows an easy way to reset the core of the driver transformer because of the ac flux waveform produced (Fig 2b).

The secondary winding is bifilar wound with the winding occupying 360° of the core to distribute core flux uniformly and to form a center tap. The core is a small, tape-wound toroid of Square Permalloy 80, which provides the necessary high sensitivity and linearity (Fig 2f).

DRIVER TRANSFORMERS T3 AND T4

Because of the need to produce accurate switching waveforms for capacitive loading of the MOSFET gates, driver transformers mandate the following objectives:

- Maximum possible bandwidth by reducing leakage inductance.
- Minimum possible leakage inductance by use of:
  —Unity turns ratios.
  —Torroidal gapless cores.
  —Multifilar windings occupying 360°.
- Minimum number of turns to reduce winding capacitances and copper losses.
- Cores with small diameters and narrow hysteresis curves to reduce magnetizing currents to an absolute minimum.
- Select core materials for low eddy-current losses and high flux capacities at 50 KHz.

POWER TRANSFORMER T5

Because of the very coarse turn resolution, due to the high operating frequency, the design of this transformer (Fig 2f) is performed in the following sequence:

- Determine the secondary voltage for 5V dc output at full load (5V + one diode drop).
- Select core size based on a single turn winding at a flux density of around 2 KGauss.
- Calculate turns ratios for other secondaries.
- Calculate primary turns for 200V dc input to the bridge power stage. (This input corresponds to the maximum duty cycle of the PWM at the 95V ac input line.)

A PQ core yields low leakage inductance and simple winding techniques, and for these reasons was chosen for this application. The single-turn secondaries are made in the form of U-shaped strips of 0.031 copper insulated with adhesive Mylar tape.

International Rectifier's IRF720 HEXFETs were chosen for the original 230W design, but higher power levels can be achieved merely by suitable sizing of power-train components. The waveform generator and drive components would remain the same. These results aren't possible with bipolar transistors. ■

Peter Wood *is field applications engineer for International Rectifier. He gained a higher National Certificate in Electronic Engineering with a physics endorsement in England, and holds three patents in the field of power switching.*

EVERY MINUTE YOUR COMPUTER'S IN THE BLACK, YOUR BUSINESS GOES DEEPER IN THE RED.

That's why more and more companies that depend on their computers are turning to uninterruptible power systems for dependable emergency power. And for maximum on-line availability they're choosing Exide Electronics.

Exide Electronics UPS Systems combine state-of-the-art technology with an unexcelled reputation for solid, reliable construction. And they are backed by a commitment to service that includes 50 service centers around the globe, and a newly developed microprocessor diagnostic system that makes Exide Electronics UPS Systems the most reliable power source money can buy.

For more information on the world's foremost UPS supplier write: Exide Electronics, P.O. Box 58189, Raleigh, NC 27658. (919) 872-3020

EXIDE ELECTRONICS

POWER PROBLEMS ?

POWER CONDITIONERS FROM Power-Matic®

500W TO 15KW WITH 99% EFFICIENCY

HALF CYCLE RESPONSE TIME

HEAVY DUTY FILTERS—NO SPIKES

45–70Hz, ANY POWER FACTOR LOAD

Call or write:
Power-Matic, Inc.
(619) 292-4422

TEST EQUIPMENT FROM Power-Science®

SPIKE GENERATORS AND MONITORS

TRANSIENT GENERATORS AND MONITORS

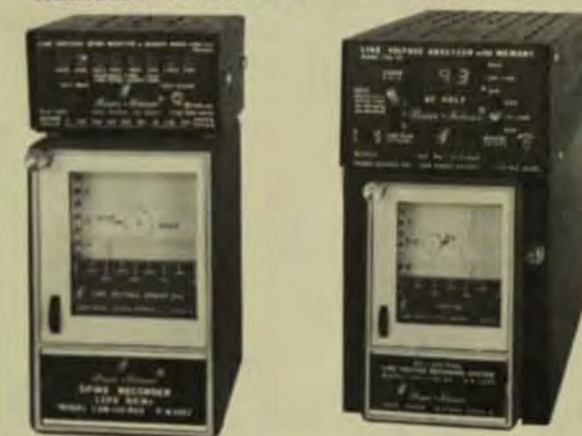

SPIKE AND TRANSIENT RECORDERS

POWER DISTURBANCE ANALYZERS

Call or write:
Power-Science, Inc.
(619) 292-4322

7667 Vickers St., San Diego, CA 92111
Ask for the booklet "POWER DISTURBANCE TECHNOLOGY" by E. Cooper

I THOUGHT I HAD TO CHOOSE BETWEEN TOO LITTLE POWER PROTECTION AND TOO MUCH COST.

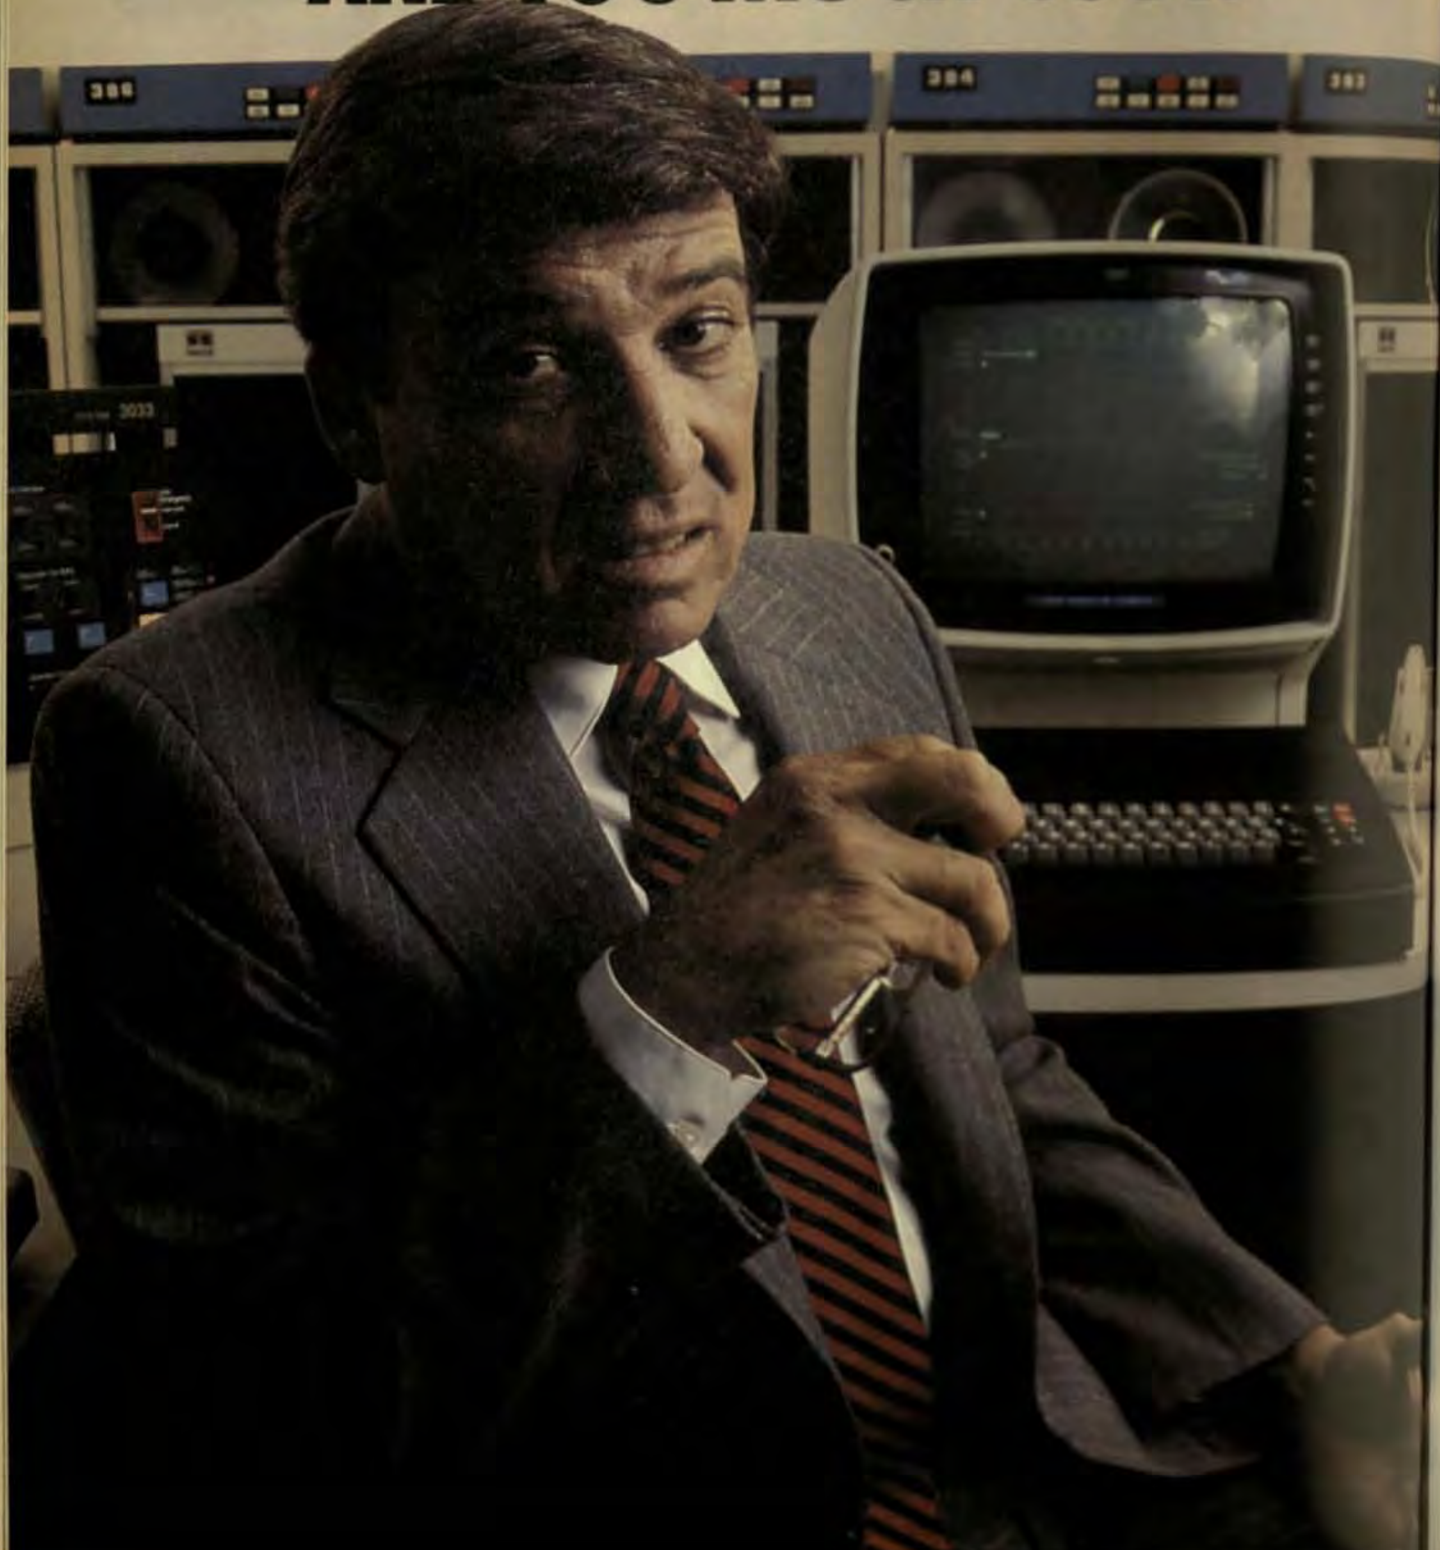

THEN I FOUND THE EXIDE ELECTRONICS SERIES 2000 UPS.

Traditionally, for moderate sized mainframe computer installations, it was a choice between two undesirables: too little power protection and too much cost. Voltage regulators or line conditioners left them vulnerable to flickers and power outages; uninterruptible power supplies provided all of the protection, but at too much cost.

Now there is a better alternative. The Exide Electronics Series 2000 UPS: it gives you complete protection at an installed cost as much as 30% below other UPS systems in its class. The Series 2000 provides complete protection against spikes, transients, line noise, brownouts and blackouts. Whatever happens, your equipment and your data are protected.

A UPS SYSTEM THAT BELONGS IN THE COMPUTER ROOM!

Just as the Series 2000 fits your budget, it also fits your computer room. The compact system, including battery pack, requires only 16 square feet, less than a third of the space needed for comparably rated systems. And it's light enough to be transported over computer room floors. Once in place, final installation is fast and simple.

UPS PROTECTION WAS NEVER THIS AFFORDABLE!

Exide Electronics' Series 2000 UPS is a new direction in uninterruptible power, designed specifically for the computer user who wants complete protection—without the cost and space requirements that have been associated with UPS. Look into the Series 2000 for yourself. It's the first UPS with big system performance at a small system price.

To: Exide Electronics, c/o Bill Campbell, PO Box 58189, Raleigh, NC 27658
☐ Please send me more information on the Series 2000 UPS
☐ Please have a salesman call

Name ______
Title ______
Company ______
Address ______
City ______ State ______ Zip ______
Telephone ______

EXIDE ELECTRONICS

P.O. Box 58189, Raleigh, NC 27658, 919/872-3020, TWX 510-928-0598

It's easy to interface your ½" drive to a DEC computer. When you have connections.

Dataram provides tape drive connections to your host LSI-11, PDP-11, or VAX computer, with a family of couplers/controllers that operate in NRZI, PE, or GCR modes. Dataram's couplers/controllers operate with ½" tape drives from all major manufacturers. As slow as 25 ips — or as fast as 125 ips. 200 BPI to 6250 BPI. With TM11 and TS11 emulations.

Start-stop or streaming. Efficient streaming is supported by a unique RSX-11M utility, FASTSAVE-11M, which provides optional backup and save capability for Dataram's streamer coupler. A full one-year warranty is standard.

For more information about ½" drive connections, call (609) 799-0071. We'll help you make the connection you need!

STANDARD AND STREAMER		GCR
AMPEX	KENNEDY	STC
CIPHER	PERTEC	TELEX
CDC	S. E. LABS	
DATUM	TANDBERG	
DIGI-DATA	TDX	

Dataram Corporation. Princeton Road, Cranbury, NJ 08512. (609) 799-0071.

LSI-11, PDP and VAX are registered trademarks of Digital Equipment Corporation.
FASTSAVE is a trademark of Computer Systems Advisors

DATARAM

Hostile Environments Require Rugged System Design/Configuration

Sensitive components can be used in harsh environments if the system is ruggedized to protect them from a wide range of adverse conditions.

by **Richard L. Strudwick, Kathy Bernstein, and Stanley R. Ulfeldt,**
Trilobyte Computer Corp.

Taking computers out of the clean, air-conditioned office and into factory-floor and field-mobile environments exposes them to the hazards of temperature extremes, humidity, vibration, shock, EMI and RFI, poor air quality, and power fluctuations. Each of these hostile conditions demands consideration in every element of system design and configuration.

Standard commercially available computer components can be ruggedized if they are integrated into a system that provides features to compensate for adverse environmental conditions. Temperature rise must be controlled by a high-volume filtered air flow. Electronic filters and shielding are needed to reduce electrical interference. Shock mounting must be installed to reduce physical damage due to vibration and shock. Air entering into the system also brings with it part of the surrounding environment (such as humidity, dirt, and heat) regardless of the quality of filtering.

To build an affordable ruggedized system mandates compromise between a broad range of options and requirement criteria. Ruggedization achieved by complete redesign or creation of components is beyond the scope of most projects. Rather than change the capabilities of internal components, it's more feasible to place them within a framework designed to expand the limits of their functions. Within this context, it's possible to discuss selected hostile environmental conditions and the available design solutions for them.

To date, the major user of ruggedized equipment has been the military. Mil-spec represents the state of the art, but it's economically unfeasible for commercial use and not readily available even at high cost. Because it's the only standard for the upper limits of system capability, however, it will be treated here as the optimum. Individual needs, naturally, will dictate individual combinations of solutions to ensure the most cost-effective design.

TEMPERATURE

Temperature is perhaps the most obvious adversity faced by transportable or process-control equipment. Few commercial computer systems will function outside a range of 10°C to 40°C. Temperature-range specifications for most commercial components are 0°C to +65°C. Some components are specified as narrowly as +10°C to +40°C. ICs used in commercial equipment have a specified temperature range of 0°C to 70°C; the mil-spec versions are −50°C to +125°C.

While absolute temperatures may range from −25°C to +125°C, it's unlikely that systems will be subject to these extremes because human beings can't function at such extremes either. Consequently, it's necessary to expand the operational temperature limits of computer systems somewhat beyond the still-air specs of their most sensitive components. This can be accomplished with proper cooling-system design.

Both Winchester and floppy-disk drives have mechanical problems with temperature change and extremes. In the case of Winchesters, the castings expand and contract with rapid changes in temperature, causing head and disk misalignment. When the temperature stabilizes, the servo tracks align properly with the data tracks, again providing stable operation.

The low-temperature limit for the drives in the Micropolis 1200 series is used by the stiffening of the grease used on the head arm pivot when the temperature drops below 0°C. The resistance introduced causes the arm to stop short of the desired track position. The drive senses this off-track condition and takes corrective action. During read operations, recovery techniques for offsetting and gain adjustments allow the data to be read. During write operations, the off-track sensor will trip and disable the write circuitry, causing the drive to return a write-error to the host computer.

With floppy drives that have a nonservo stepper-motor positioning system, the offset caused by differential expansion of the disk and media results in an off-track condition at extremes of temperature. If the system is normally used in hot or cold environments, the drives can be aligned to these temperatures for reliable operation.

The disk-drive media is most sensitive to high temperatures. The disk media used in most Winchesters is coated with a lubricant that softens at 65°C. Micropolis 1200 series drives have a 10°C rise between the disk platters and the casting. Thus, the maximum casting temperature is 55°C. Allowing 5°C for the heat transfer to the cooling air, the maximum operating limit is 50°C.

Quality floppy diskettes remain reliable somewhat above that limit. By forcing large amounts of cooling air around and through the disk drive, reliable operation can be achieved.

Commercial-quality integrated circuits used in the processor, memory, interfaces, and disk drives are rated from 0°C to 70°C. Most commercial modules will function reliably to 60°C with adequate air flow. Digital Equipment Corp. specifies the LSI-11/23 processor board at 60°C with a maximum rise of 5°C in the air flowing across the board.

To meet the cooling specs, air flow must be dispersed uniformly around all modules in the backplane. The boards are normally clustered together because of backplane priorities. If the unused slots are left empty, the air flow will follow the path of least resistance through the open space, leaving the boards with insufficient cooling air.

You can accomplish two things by using a printed-circuit card in all unused slots, with the priority grant lines jumpered. First, the active modules can be more uniformly distributed throughout the backplane, thereby reducing the heat concentration. Second, the air flows uniformly through the backplane.

Since all commercial modules are designed for direct air cooling, heat transfer from the components to the cooling air is the function to optimize for high-temperature operation. The major factor affecting the temperature of a component is the quantity of air passing over it. Of secondary importance but still significant is the turbulence of the air. High velocities result in increased turbulence.

If a large volume of high-velocity air is forced over heat-producing components, their temperature rise will be minimized and the maximum ambient operating limit of the system will be brought to within a few degrees of the maximum limit of the individual components.

You must carefully balance the air flow in the various paths to ensure that all components receive adequate cooling. The Trilobyte II illustrates the factors involved in air-flow design. A large, unobstructed air intake and outlet was needed to accommodate the large filter area required. The front and back surfaces were used because they remain unobstructed in virtually all installations. Filters are needed on both intake and outlet to prevent infiltration during shutdown.

Front access is required to insert floppy disks. A door with an adequate filter would cover most of the front of the unit and provide easy access to the floppy drive. The fans would be installed in the back, since mounting them in the door isn't feasible.

Air-flow direction was determined by several considerations:

- Fans are more efficient pushing air than pulling.
- Filters have better capture efficiencies at low air velocities and must be easy to change when clogged.
- Fans concentrate their output into high-velocity streams while their input is rather diffuse.
- Maintaining a positive pressure within the unit prevents infiltration of particles through cracks.
- Opening the front door during operation to change floppies should not allow dust to enter.

With the fans in the rear, back-to-front air flow allows the air to be pushed through the unit and maintain a positive pressure within, even when the front door is opened, so that dust is swept away by the outflowing air. Locating the rear intake filter before the fans in the air stream to allow easy servicing from the exterior of the unit might at first appear to be a compromise because the fans would have to pull the air through the filter.

But when the flow patterns into and out of the fans are taken into account, placing the filter before the fans in the airstream works out much better. If the filter were immediately behind the fans, the air would be blasted at high velocity at a small region of the filter and result in poor capture efficiencies, rapid clogging and higher back pressures. Placing the filter behind the fans would have required several inches of space and a specially designed diffuser to distribute the air uniformly through the filter.

With the filter in front of the fans the air flows uniformly through it at low velocity, resulting in good capture performance and low back pressure. This matches with the diffuse flow pattern at the intake of the fans and provides an even, efficient cooling-air system.

With the air flow established back-to-front, the backplane, disk drives, and power supplies have to be mounted along the front-to-back axis. Before the final package configuration can be determined, consider temperatures, filter type, humidity, shock, vibration, power supplies and EMI and RFI.

HUMIDITY

Powering up a commercial system that's stabilized at 0°C is generally successful. But even when the conditions in the room at startup are within the operating limits of the system, it can still be affected by previous chilling and dust accumulation.

In general, problems are not encountered until moisture condenses on components and causes changes

DISK EMULATION
THE COMPLETE SOLUTION FROM AMPEX

For Faster Throughput and Maximum Uptime

Now Ampex provides the complete solution to peripheral memory problems with a choice of Megastore* memories — either nonvolatile memory with up to 8 MB capacity or volatile memory with 64K DRAMs and up to 32 MB capacity. Both are nonrotating, have no latency, and offer access times in microseconds. They eliminate head crashing, have no moving parts and are insensitive to temperature, dust, shock and vibration. Popular emulations include DEC, DATA GENERAL and HONEYWELL.

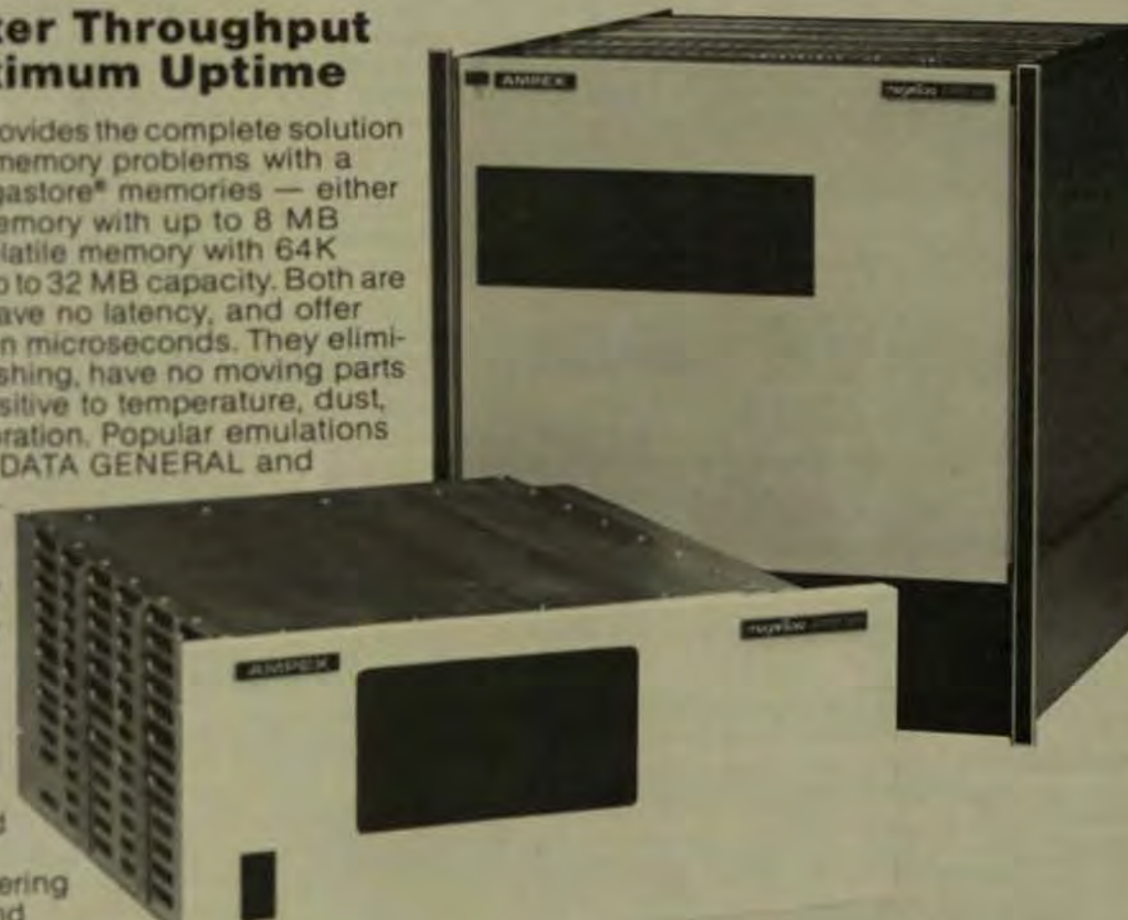

All Your Memory Needs

Ampex has answers to all your memory needs. To find out how our broad engineering experience and extensive production capability can help you, call Ampex now. Call toll-free 800-421-6863. In California call 213-416-1412. Ampex Corporation, Memory Products Division, 200 North Nash Street, El Segundo, California 90245.

AMPEX
Ampex Corporation • One of The Signal Companies

in electrical characteristics (including sparkdown), corrosion on some metal surfaces, and moisture problems on disk durfaces. This occurs when surface temperature is below the surrounding air's dew point—the temperature at which the relative humidity reaches 100%. The environment presents humidity ranges of 0 to 100%, while commercial electronic equipment performs best in noncondensing humidity of 0 to 80%.

Internal thermostatically controlled heating elements that keep the system above the highest dew point in the daily cycle will prevent both condensation and other cold startup problems. Adequate air filtration solves dust problems, and equilibrating the system within its operating limits at turn-on prevents thermal-shock failures of electronic components.

AIR PURITY

Particles as small as 0.001μ are abundant in the atmosphere in the form of oil and tobacco smoke, atmospheric dust, and carbon black. The air contains various sizes of dust, dirt, and (in factory environments) potentially corrosive and conductive materials.

High-velocity air flow circulates atmospheric contaminants and particulates as well as air, and results in higher particle velocity, which increases the likelihood of impurities permeating the system through pores in the filter media.

Filter media are generally specified for efficiency at 1μ particle size. For example, a 60% filter would filter approximately 60% of the 1μ particles impinging on it. Filtering efficiency also depends on the velocity of air flow and is significantly reduced at higher particle velocities.

If you increase the filter area, the velocity is reduced and filter efficiency is improved. At the same time, the backpressure for a given flow is also reduced. The Trilobyte II, with two 7-in. fans, a 4.5-ft² 60%-intake filter, and a 2.2-ft² exhaust filter achieves about 400 ft³/min. The air flow through this type of filter decreases rapidly with filter ratings above 60%. With an 85% intake filter, air flow is too low to cool the system at 20°C.

Individual applications may require you to specify particular filtering media to suit a specific environment. Filtering to remove corrosive chemicals is generally not effective and may require other solutions. In factories where metal dust is found, use extra filtering to keep the conductive particles out of the system.

SHOCK AND VIBRATION

Most commercial electronic equipment is not designed to withstand vibration and shock. In a process-control application such as a factory, the vibration and shock of the floor present problems. Equipment to be transported is subject to a great deal of shock and vibration regardless of the form of transit.

Shock and vibration protection needs to be considered in two contexts:

- Measures required for a system installed in an industrial or field-mobile situation.
- The additional measures necessary to protect portable systems from the hazards of transportation.

In both factory-installed and vehicle-mounted situations, the principal concern is vibration with moderate shock. Vehicle mounting assumes that the system is secured to the vehicle body, which is rubber-mounted in the conventional manner. Securely mounted electronic and mechanical components will survive and operate under these conditions.

We can divide the components of a system into two classes: those that will function well when securely mounted (including PC cards, mechanical fan assemblies, and power supplies) and those (such as disk drives) that malfunction or are damaged by the shock and vibrations present in industrial and mobile environments. For example, floppy-disk drives will survive but not operate reliably, while Winchester disks suffer damage from head vibration.

Components on PC cards may withstand high vibration and shock but the manner in which they are mounted may make them subject to additional vibration and subsequent lead failure. All components that can vibrate must somehow be fastened to the card by being foamed or glued into place. Board flexing due to vibration can cause the boards to work their way out of their connectors gradually. Similarly, the cable flexing due to vibration causes cable connectors to work free and even to fatigue and break.

Again, Trilobyte II can serve as an example. Here, the cards are mounted vertically to make it harder for them to work out of their sockets. Because DEC systems allow three widths of cards—dual, quad, and hex—the smaller cards may be left with only one edge, or none, supported by card guides.

All of the cards can be secured against lateral movement if you add intermediate card guides confined by the sides of the card cage and the adjacent cards. Semirigid foam compressed over the tops of the cards by the box lid retains the cards in the backplane, holds the I/O connectors

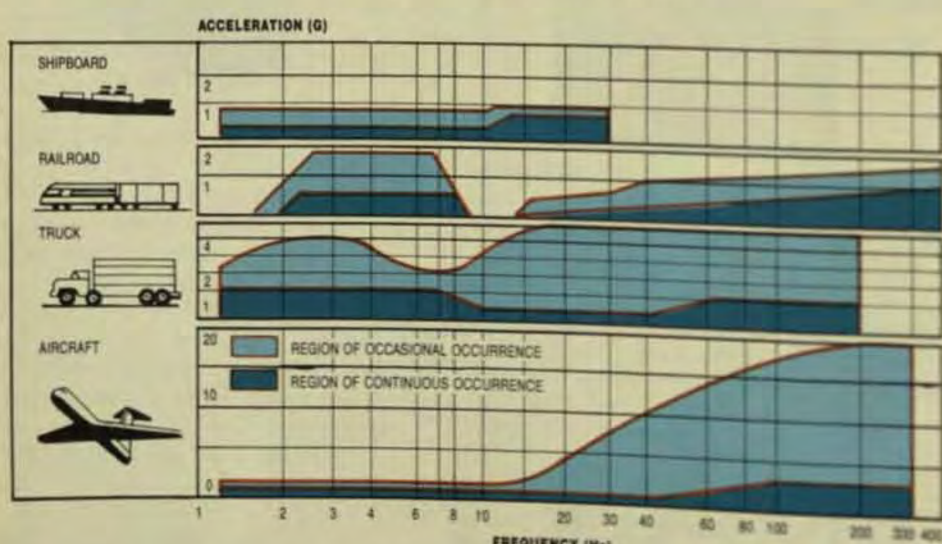

Fig 1 On a highway, vehicular-mounted systems will be subject to vibration. System amplification of 7 to 20 Hz lies in the low-amplitude region of the vibration spectra.

THE INTELLIGENT CHOICE

High Performance Disk Subsystems For Your DEC PDP-11 or VAX Computer From Data Management Labs.

DML's new intelligent disk controller with features like command stacking, seek overlap and lineal revectoring optimizes disk subsystem throughput. Your host computer is relieved of disk handling responsibilities freeing it to perform user applications. And, because it is DEC UDA50 compatible, our controller provides you an affordable way to implement DEC's Digital Storage Architecture. The 51 sector speed matching buffer and 1.2 MB data rate permit us to attach high speed disk drives and gate their transfers to the UNIBUS for maximum channel utilization. A powerful 170-bit ECC written with every sector virtually eliminates uncorrectable data errors.

DML's large capacity Winchester disk drives complete the package. Giving you over 450MB of data on a single drive. Up to four drives can be included in a subsystem for 1.8 GB of storage. And if that isn't enough, two subsystems can be attached to each UNIBUS. Our drives use the most advanced disk technologies, such as thin-film heads and linear actuator for better performance and data integrity. Add this to their outstanding reliability, self-contained diagnostics and affordable price for the best buy in mass storage available today.

Experience? We could teach DEC a thing or two about intelligent disk subsystems. DML has been successfully building them for years. We've shipped thousands of controllers and disk drives based on the same advanced architecture used in our new DEC compatibles.

For more information about DML's DEC compatible disk subsystems, call us at (408) 946-9424. Or write DML 2180 Bering Drive, San Jose, CA 95131.

DML
DATA MANAGEMENT LABS

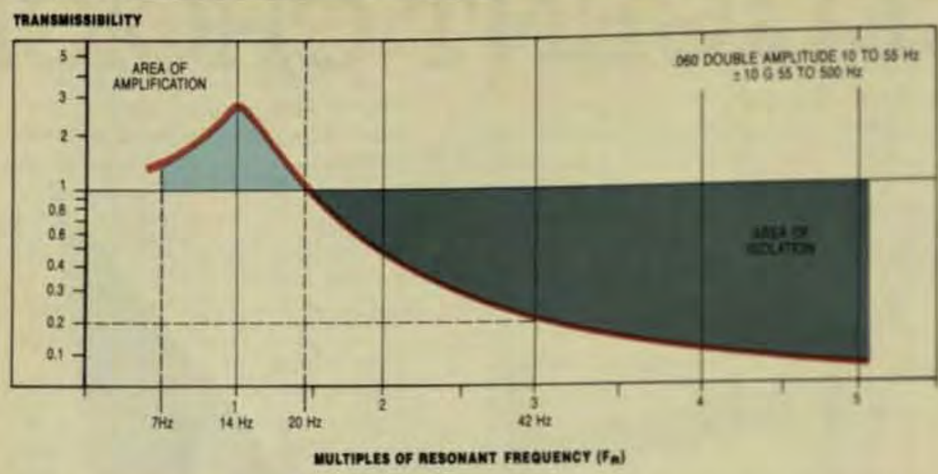

Fig 2 The vibration transmissibility curve shows that the 80% isolation point of the isolators at 42 Hz is reached before the rise in amplitude.

in their sockets, and prevents the cables themselves from vibrating.

The next problem is power supplies, which may contain heavy PC-mounted components such as transformers and large filter capacitors that must be attached securely. Several of the available switching power supplies, which use smaller components, have adequate vibration resistance.

The fans and other mechanical assemblies can easily handle the shock and vibration of installed-factory and mobile environments. Disk drives, on the other hand, malfunction when vibrations reach a few G's, and Winchesters are damaged at accelerations above about 20G. Both floppies and Winchesters require shock mounting to operate in installations experiencing only 1 to 2G. Winchesters require protection to prevent head vibrations (which occur above 20G) from damaging the disk media.

VIBRATION ISOLATORS

Vibration isolators are typically either bonded-rubber springs or coiled cable assemblies. When a disturbing force is applied to the carrier assembly, these isolators allow relative movement of the shock-mounted assembly to reduce the force by distributing it over a longer time. Isolators provide damping, and absorb energy by hysteresis loss in the rubber or friction between the wires.

The controlling property of all shock/vibration isolation systems is their natural frequency of oscillation, which depends on the mass of the supported object and the stiffness of the shock mounts. Since all systems with the same natural frequency, regardless of mass, will behave in the same manner, they can be characterized by their natural frequency.

Take, for example, a vehicular-mounted system for both on- and off-road use. The recommended natural frequency for smooth highway is 12 to 14 Hz. For rough roads, it's 15 to 20 Hz. The example we'll use is 14 Hz. The transportation vibration spectra (Fig 1) and vibration transmissibility (Fig 2) show that the system's region of amplification from 7 to 20 Hz lies principally in the low-amplitude region of the vibration spectra, and that the 0.2 transmission point (80% isolation) of the isolators at 42 Hz is reached before the rise in amplitude.

For occasional occurrence, Figs 2 and 3 show 7 Hz to be an obvious choice. However, while high-frequency attenuation is better, the continuous spectrum shows 1G at 7 Hz compared to 0.3G at 14 Hz—three times the driving force at resonance. Most of the time, the disk drive would be in resonance at 7 Hz with about 3G. A 14-Hz system would resonate with 1G average—a much better situation because disks run reliably at 1G but not at 3G.

For the 14-Hz system, the nearly 5G in the occasional spectrum at resonance will cause problems if they occur during operation. But since the 15G generated in resonance are well within the 20G fragility of Winchester disks, only recoverable disk errors are likely to occur.

Shock is the other side of the coin. The drop height vs. shock nomograph (Fig 3) relates natural frequency and drop height to deflection and maximum Gs transmitted. A 0.7-in. deflection gives a drop height of 7 in. for a maximum transmitted force of 16G.

The sample 14-Hz isolation system provides good protection for the disk drives when installed, and tolerates moderately rough handling. It also assumes that the outer chassis is not subject to severe shock.

In the portable shipping environment, the major concern is severe shock to the equipment due to dropping. The standard drop height is 30 in. for equipment weighing from 51 to 250 lbs (standard for most computer chasses).

A flat drop from 30 in., for example, subjects the main chassis and all of the rigidly mounted components to hundreds of G's over a broad spe-

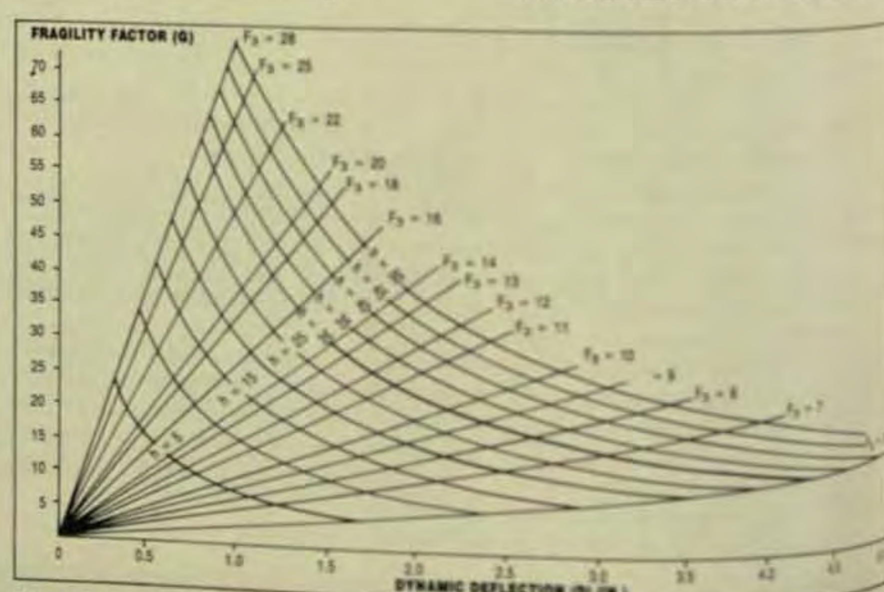

Fig 3 This nomograph plots deflection vs. shock for a variety of environments. By knowing the severity of the applied impact, the designer can determine the characteristics he needs in an isolation system.

70 megabyte CSS-800 storage system lets your micro work like a mainframe.

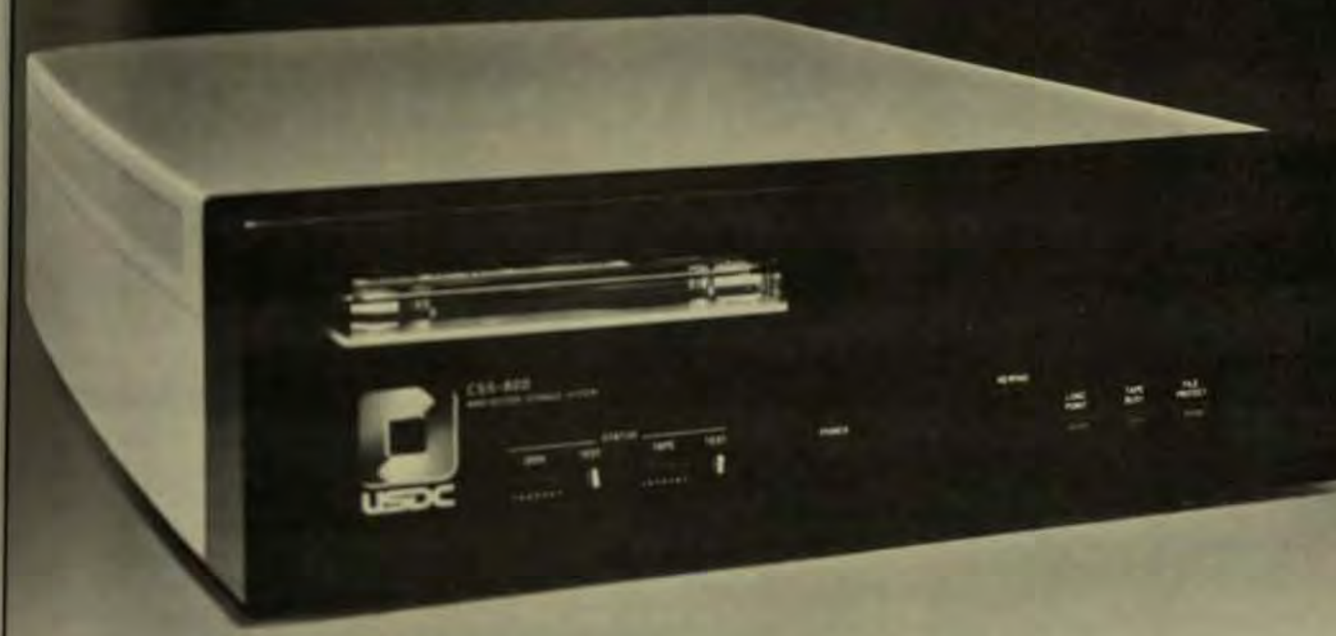

Advanced architecture

U.S. Design Corporation's CSS-800 storage system has proprietary disk and tape controllers which use advanced data management architecture. This architecture, usually found only on mainframes, will increase the speed of your mini or micro dramatically.

Multiported cache memory controller

The CSS-800 disk system improves data throughput in three ways. It keeps the most often used disk sectors in cache memory for quick access. It uses a look-ahead buffer to put sequential data in cache before you ask for it. And it reads tracks continuously without the delays of sector interleaving.

Self-contained diagnostics and ECC

The CSS-800's front panel switches initiate diagnostic routines that check disk and tape drives, controllers and interfaces. Status lights indicate faults, and the CSS-800 communicates any error to the host computer. Errors up to eight bits are automatically corrected with ECC. The CSS-800 also remaps defective sectors. This process is transparent to the host computer.

Emulates DEC disk and tape

The CSS-800 is plug compatible with LSI-11* Q-bus* and PDP-11* or VAX Unibus* host computers. It looks like an RK07* disk drive and a TU10* mag tape drive to your DEC host. And the CSS-800 is software compatible with DEC's RT11,* RSX* and RSTS-E,* as well as TSX** and UNIX.***

Software intensive controller technology

The CSS-800 microprocessor-based controllers are software intensive, which gives us the flexibility to adapt our hardware to advancements in host system technologies. The CSS-800 is a flexible system, designed to keep pace in a rapidly evolving industry.

*LSI-11, Q-bus, PDP-11, VAX Unibus, RK07, TU10, RT11, RSX and RSTS-E are registered trademarks of Digital Equipment Corporation.
**TSX is a registered trademark of S and H Computer.
***UNIX is a registered trademark of Bell Laboratories.

Let your micro think it's a mainframe. Call U.S. Design at our toll free number (800) 368-2811.

U.S. Design Corporation
5100 Philadelphia Way/Lanham, MD 20706
(301) 577-2880 TWX 710-826-0417

Ran·yan Com·put·er En·hance·ment Sys·tems

1. A company that was formed with the sole aim of becoming <u>the</u> specialist in providing systems and subsystems to enhance DEC* computer performance. **2.** A company that is currently manufacturing a family of unique products proven to enhance DEC computer performance. *e.g.* **Peripheral Processor Link Series** including (a) PPL-1 that provides a comprehensive interface between two DEC computer systems with a "memory window" technique newly developed by Ranyan. (b) PPL-2. Same as above but links Q-Bus or Unibus systems with VAX machines. *e.g.* **Bus Manipulator Series** including (a) BMA-2 Q-Bus To Unibus Adapter which permits DEC users to update older PDP-11 systems to Q-Bus processor performance. (b) BMA-1Q Q-Bus Repeater which doubles the Q-Bus processor load capacity by providing a physical and electrical extension of the standard Q-Bus. (c) BMA-1U Unibus Repeater which provides the same benefits listed above to the Unibus. *e.g.* **Intelligent Multi-Shared Memory,** a fully integrated subsystem specifically designed to enable memory sharing by up to four DEC computers, thus providing higher throughputs and expanded computing power while eliminating the need to invest in an entirely new system. **3.** A company that will continue developing and marketing unique products, such as those listed above, to further enhance the performance of DEC computer systems.

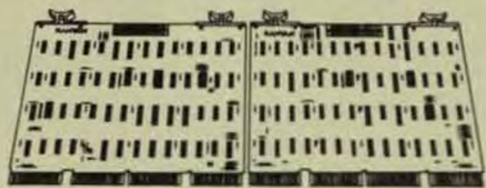
Peripheral Processor Link PPL-1


Q-Bus to Unibus Adapter BMA2

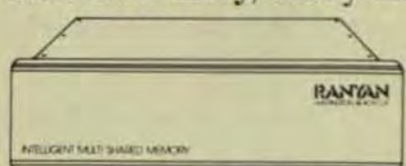

Intelligent Multi-Shared Memory MSM-1/5

*DEC, Q-Bus, Unibus, and VAX are trademarks of Digital Equipment Corp.

RANYAN GIVING NEW MEANING TO DEC SYSTEM PERFORMANCE.

RANYAN COMPUTER ENHANCEMENT SYSTEMS · 15239 SPRINGDALE ST. · HUNTINGTON BEACH, CA 92649 · 714-895-5504

Cou·pon *n.* **1.** A device at the end of an ad which allows you easy access to further information on the products mentioned in the ad.

PLEASE SEND ME MORE INFORMATION ON THE FOLLOWING: ☐ Peripheral Processor Link Series ☐ Bus Manipulator Series ☐ Intelligent Multi-Shared Memory

Name________ Title________

Company________ Phone________

Address________ City________ State________ Zip________

Mail coupon to: Ranyan Computer Enhancement Systems, 15239 Springdale St., Huntington Beach, CA 92649

tem. In that case, the disk assembly could bottom against the chassis, rendering a severe shock pulse to the drives. To prevent this occurrence, additional protection is required for the complete chassis.

A 30-in. drop with a 20G maximum shock limit requires a 9.4-Hz system with 2.3-in. deflection. Shipping containers with internal shock-mounted 19-in. racks are available with the desired isolation characteristics from several manufacturers. When a system chassis with internal shock-mounted components is externally shock-mounted, the natural frequencies of the two systems should be adjusted so that the frequency of the internal amplification is at least 1.4 times that of the external, so that their regions of amplification do not overlap.

POWER AND EMI, RFI

A system designed to be transportable will meet with power that varies considerably from place to place, including ranges of 12 to 120V dc and 115 to 230V ac, 50 to 400 Hz. Some of this power may come from small, gas-driven generators, which can be quite electrically noisy. The voltage requirements of military-type equipment may be more closely defined as 28 dc or 115 to 230V ac, 60 Hz or 400 Hz. Even so, both military and commercial power sources may have considerable noise to cope with as well as electrical power fluctuations.

Power supplies also should be selected to allow a wide variety of input voltages for the same housing dimensions, unless, of course, you plan to design your own. Power supplies (particularly switching power supplies) must have suitable shielding to prevent EMI and RFI from escaping and to prevent these fluctuations from passing through the power system. Ever more stringent FCC regulations are yet another element to be considered in design of shielding methods for electronic equipment.

Adequate filtering must be provided to protect systems from external electrical noise sources and magnetic-field-induced voltages. In a factory, you might find voltage spikes of as much as 10 KV on the power lines, caused by the operation of power relays and contactors. These voltage spikes, which decay in less than 1 ms, must be prevented from entering the system. The input power should be adequately filtered and some form of overvoltage protection provided, such as metal-oxide varistors. These limit excessive peak-to-peak voltages and bring them down to levels acceptable to the system design.

Another way in which EMI/RFI can enter or exit the system is through the air-inlet and -outlet vents. Honeycombed metals reduce this occurrence significantly. Honeycombed metals provide low resistance to air flow, high RF-shielding efficiency, reduced turbulence, and low weight. The most common honeycomb material is aluminum, which is also relatively low-priced.

Wire screens or meshes require as much or more thickness to provide the same attenuation as the equivalent honeycomb material. However, these cause considerably more turbulence and air-flow restriction. Aluminum suffers from the problem of decreased efficiency at low frequency ranges, starting at about 100 KHz. Three times as much thickness is required for the same attenuation at 10 KHz as at 100 KHz. Aluminum's lack of low-frequency attenuation is due to its low permeability. Iron and steel can provide relatively flat attenuation down to near dc, but steel honeycomb is five times more expensive.

PACKAGING

Packaging, as it relates to ruggedized computer systems, encompasses all of the elements already discussed. The package can either impede or enhance these protective measures and must include consideration of the type of mounting to be used, if any.

By taking account of the factors discussed, you can ruggedize available equipment to perform computer applications that would otherwise be untenable. ■

Richard Strudwick *heads the design and manufacturing teams for Trilobyte Computer Corp.*

Kathy Bernstein, *marketing manager for Trilobyte. She earned her BA at Univ. of California, Los Angeles.*

Stan Ulfeldt, *co-founder of Trilobyte, has been designing laboratory and field-mobile geophysical data acquisition and signal processing systems for over 10 years.*

A Typical Scenario for Startup Problems with a Nonruggedized System

It was −20°C the night before. The heat was turned on shortly before starting time and it's 10°C in the room now, but the system's still at −15°C internally. Some frost has formed on the Winchester disk's platters because the dew point was 10°C when the system was shut down last night.

The system is turned on. The fans rev up, blowing moist air, with a dew point of perhaps 5°C, over the PC cards, which are still at −15°C. The components already have a coating of alkali dust from the factory, and now frost starts to form on them.

The Winchester spins up, sweeping the frost from the landing zone; no damage yet, though. The disk is up to speed now and at auto-boot seeks to the first track. The disk is spinning fast now and the frost knocks the head out of alignment and the system fails to boot.

The operator tries again, but by now the temperature of the boards is rising above freezing and the frost is melting and mixing with the alkali dust to form a conductive solution, causing the processor to malfunction. The operator goes out for a long coffee break, leaving the system on—as he's done several times before on cold mornings when the system wouldn't boot.

The system warms up and dries out. The processor recovers and the Winchester's head is just a little bit further off track. But this time, the thermal stress has finally broken an internal lead in a memory chip.

The operator returns and tries to boot the system again. It boots—but slowly, because the disk has to retry nearly every read. He runs his program. There are some disk read-error messages and then a memory-parity error that kills the system. ■

Five reasons why DEC users should buy Emulex communications controllers.

Broad product line featuring our new DMF-32 emulation.
Nobody covers LSI-11, PDP-11, and VAX-11 users' needs like Emulex. More than 15 software-transparent controllers emulating DH11, DZ11, DV11 and DMF-32. All deliver improved line-handling capabilities, in a smaller package, at lower costs.

More channels.
Emulex's new DMF-32 emulation is typical. One controller board handles up to 64 lines, vs. only eight per DEC module. And Emulex offers *all* lines with modem control, not just two. For even more lines, Emulex's Statcon Series is the answer. We simply add a low-cost port concentrator, so that with one controller board you can connect up to 256 remote *and* local terminals.

Easy growth path.
As your system grows, upgrading is simple with Emulex controllers. Just change PROM sets. Example: DH to DMF for $350. In addition, Emulex's advanced microprocessor architecture is consistent throughout the product line. Think of the inventory savings.

Fewer backplane slots.
Emulex communications controllers pack so much capability onto each board that fewer boards are needed. Take a 64-line DH11 emulation. Emulex does on one board what it takes DEC to do on 36. Think of the savings in rack space, to say nothing of price.

Lower prices.
For instance, a DEC DH11 controller lists at $8,950 per 16 lines, with expansion chassis costing $3,000 or more. Compare that to Emulex's CS11/H at $4,500 for the first 16 lines and $3,000 for each additional 16 lines. At 64 lines, you suddenly have savings of about $23,000 and a lot of extra slots to boot.

Don't speculate with your communications controller dollars. Invest in Emulex. Phone toll free: (800) 854-7112. In California: (714) 662-5600. Or write: Emulex Corporation, 3545 Harbor Blvd., P.O. Box 6725, Costa Mesa, CA 92626.

EMULEX
The genuine alternative.

Data Bridges Can Move Stored Data and Records Between Environments

Data bridges—links between two separate computing environments—facilitate high-volume data exchanges that involve networks, protocols, and gateways.

by **Chris Shaw**, Xerox Computer Services

A central problem in computing today is the moving of stored data from one computing environment to another. In modern business, computer systems that can't communicate with each other proliferate, and the data produced by one system can't be used by another. With the rapidly growing use of microcomputers in the business office, this problem has become especially critical. To be really useful, in most cases these micros need to be linked into a broader information network so they can access the corporate data stored on central host computers.

When a transaction is processed by a computer that doesn't have the data it requires, there are two possibilities: either you move the data to the processing computer, or you move the transaction to the computer with the data.

Sometimes processing a transaction requires data from multiple locations, which complicates things considerably. A data bridge linking two separate computing environments provides a way to move stored data—records, files, even transactions—from one computing environment to the other. These environments may be similar or different. And they may reside in different, even remote, host computers, or they may reside within the same host.

A data bridge serves to exchange data between cooperating (but not necessarily co-located) processes, and preserves the meaning and usefulness of that data as it is exchanged between possibly dissimilar computer systems.

The goal of a data bridge is to provide independence from both data location and originating format. A data bridge thus entails both data transfer, which often requires remote communications, and data translation, which involves the re-coding, reformatting, and even restructuring of the data.

DATA TRANSFER

Data transfer involves:

- **Networks** to provide the transport mechanism to move data from one location to another. Networks provide geographic transfer and incorporate physical links, such as wires, cables, or physically transported media (floppy disks or tapes).
- **Protocols** to provide the synchronizing mechanism for data transfer. Protocols are rules specifying the format and relationships of messages exchanged between communication nodes. They provide for message control and timing.
- **Gateways** to provide a protocol interface between networks, workstations, and applications. Gateways provide protocol conversion, buffering, and coding.

The basic data-transfer process (Fig 1) requires networks, protocols, and gateways, and includes:

- File transfer.
- Record access.
- Message transfer.
- Document transfer.
- Job transfer.
- Transaction transfer.
- Remote procedure invocation.

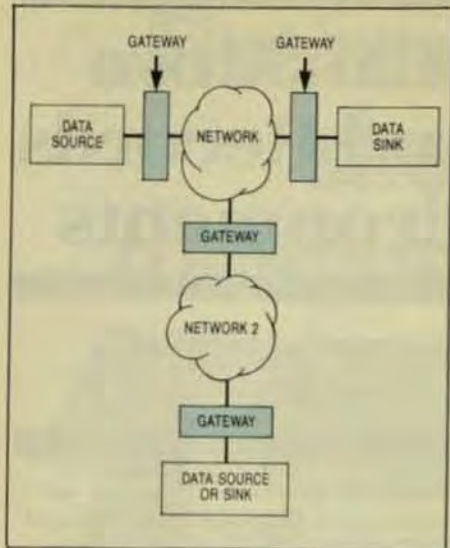

Fig 1 The goal of a data bridge is to provide data translation and transfer. These are implemented by means of protocols, gateways, and networks.

Communication networks for data transfer are almost always organized as a hierarchy of layers, where each layer performs a small set of closely related functions. The functions provided by one level support those of the next higher level and are implemented in terms of those at the next lower level.

The International Standards Organization (ISO) reference model provides a good framework for discussion of such multilevel networks (Fig 2).

Logically, each layer is communicating horizontally with the same level in the other machine. Of course, data is actually passed vertically down the levels of the sending machine and up the layers of the receiving machine. Only in level 1 does the physical communication occur (Fig 3).

When an application program running in host A wants to send data to the application in host B, it passes a message to the application-control level in its own machine. The application level adds a level 7 control header and passes the result down to the presentation level, which transforms the data and adds a level 6 control header and passes the result down to the session level, which adds its own header and passes the result down to the transport level, and so on. The complete path from one user process to the other is shown in the heavy flow line in Fig 3.

It should be understood that the ISO reference model presents a somewhat idealistic picture of things. In any real system, a user process typically has access to many of the lower-level interfaces (Fig 4). This access is necessary wherever higher levels don't provide needed services.

The ISO reference model is not, of course, the only layered architecture for data transfer. IBM's Systems Network Architecture (SNA) is layered, as is the Xerox Internet Architecture which splits layer 3, network control, into two sublayers: internet control and local-network control.

DATA TRANSLATION

In addition to data transfer, data bridges provide data translation to transform the data representation and structure of one environment to those required in another. Data bridges necessarily provide syntactic and semantic translation between applications and processes, as well as communication.

Data bridges imply both data transfer and data translation. Thus, a distributed database management system, which typically involves no data translation, is not in itself a data bridge, nor is an electronic mail system. But both these systems could be used to build a data bridge.

Data translation involves data representation—the coding of individual data items, and data structure—the grouping of data into larger elements. The problems of translating data representations and data structures are both so intractable as to be incapable of any useful, general solution. This simply means, however, that a data bridge must be fairly specific to the environments it links.

Unfortunately, when data is exchanged by systems that support different data formats and representations, the data-translation problems that occur prohibit one-to-one correspondence between the two data representations. These include loss of precision because of varying word sizes, format incompatibility, and data-type incompatibility (as when the source supplies age, and the target requires date of birth, or vice versa).

It's extremely difficult to move data files containing data mixtures (of integers, floating point numbers, and characters, for example) between machines. In theory, each data item could occupy one record in a canonical format, with the data type and value both explicitly stored. In practice, the idea doesn't seem to work very well, not only because there are problems of interfacing to existing software, but also because of the high overheads and the problems involved in converting floating point numbers.

Data structure translation is even more difficult. Restructuring operations like compress, merge, expand and partition are typical of those encountered in transferring data from one environment to another.

Data translation requires additional architectural layers, over and above those required for data transfer. These layers extract and insert the data, and format and reformat it.

The extract process serves to reduce the complex physical structure

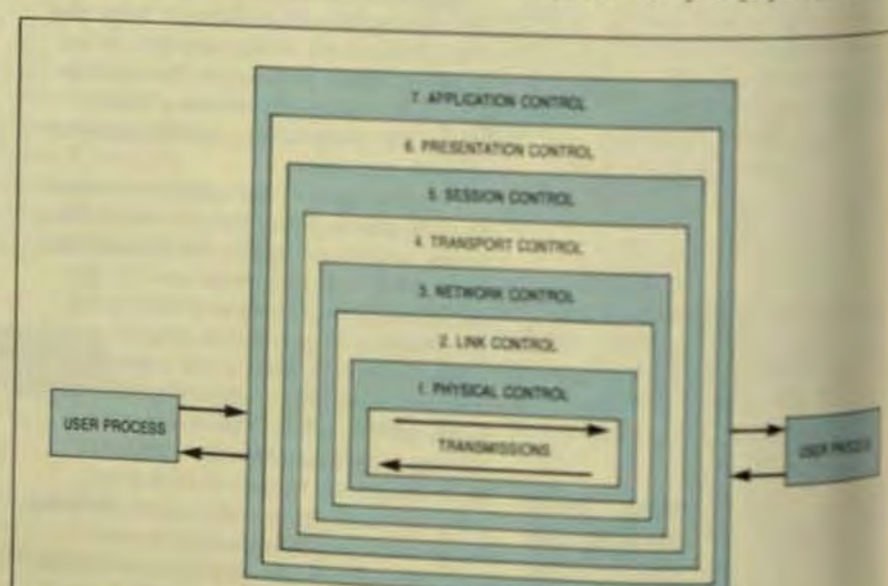

Fig 2 The ISO reference model is a 7-layer standard architecture for open-system interconnection.

Traditional Values

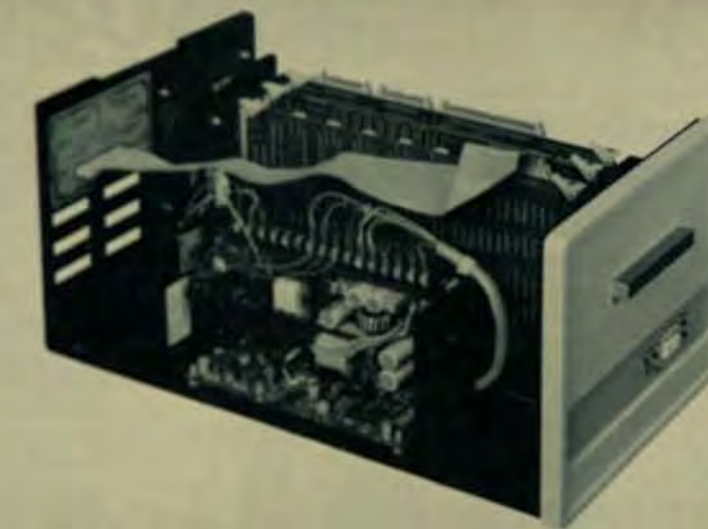

With some silver in the sideburns comes the conviction that you get pretty much what you pay for.

Check our MIKRON 600. Priced to proffer profit possibilities. It's a full 16-bit, 128KB DG-compatible minicomputer. Fast and powerful. Disc controllers for most kinds of storage: 10MB cartridge-types, Larks, CMD, SMD or Winchesters (5¼-ins). Even floppies. Up to ten multi-tasking users (or five on a single mux), with 512KB of auxiliary swapping, or cache memory, if response is less than instantaneous. Typical discounts bring you well under $5K for a basic system configuration.

Run IRIS, BITS, BLIS/COBOL operating systems. Sophisticated and powerful. With access to their loads of applications software.

You're the computer systems builder seeking a CPU to do your job with performance to spare. We're the manufacturer who knows you are his bread and butter. Like you, we've been there. That's why we strive for affordable quality, performance and versatility.

Traditional values in a word:

BYTRONIX

2701 E. Chapman Ave.
Fullerton, CA 92631
Tel: (714) 871-8763 Telex: 681-428

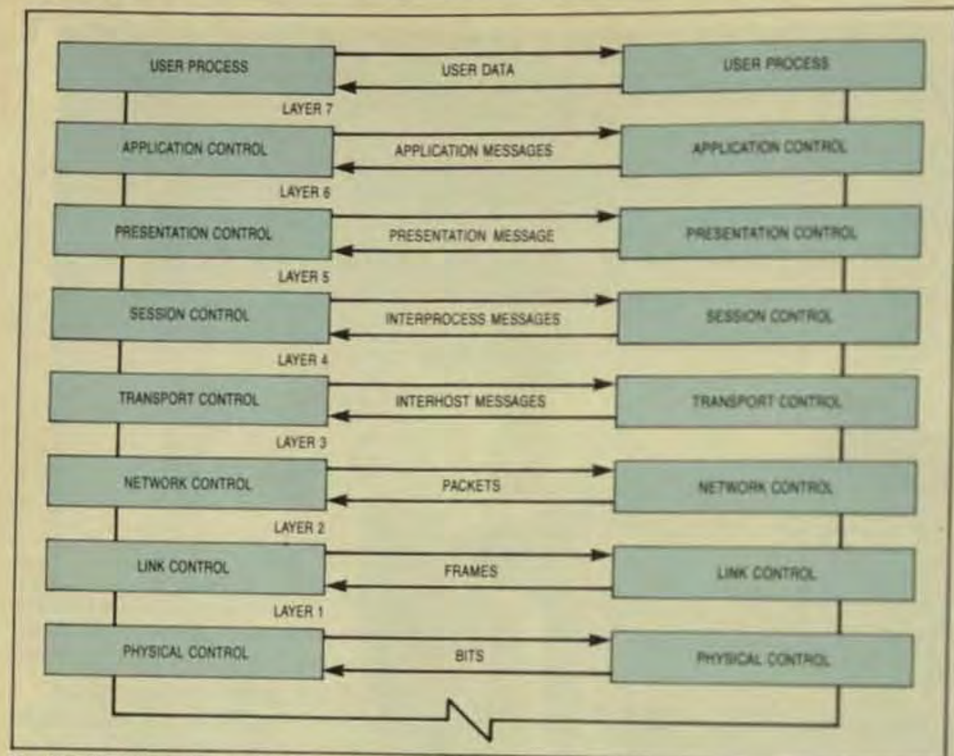

Fig 3 In the ISO model, data is passed vertically down the layers of the sending computer and back up through the layers of the receiving computer, as shown by the solid line. Logically, each layer communicates horizontally.

of the source data to a very simple physical structure, such as a table or flat file. The format process serves to create a common or standard representation of the source data, one that is acceptable to both the communication network and the target environment. The reformat and insertion processes reverse these translations.

As an example of the data translation approach, workers at the University of Michigan have developed an interpretive data-translation technique that uses a stored data definition language (SDDL) to describe the source and target data bases, and a translation-definition language (TDL) to define restructuring operations. Each host environment holds a description of its own stored data, and a description of the restructuring operations needed to translate that data into a form acceptable to each of the other host environments.

The use of a standard data representation in exchanging data among several environments is highly desirable as a practical matter. The number of translation routines can be reduced from N(N − 1) to 2N, with one algorithm in each environment defining the transformation to normal form and one defining the inverse.

Standard data formats thus facilitate data bridging. The continuing need and desire to exchange computer-readable information has given rise to numerous data representation standards.

Because of communication considerations, character-based standard formats are likely to provide the best general-purpose bridges for the exchange of structured data.

The exchange of self-describing data, in which data-descriptive tags accompany the data in its travels, is an obvious extension. A standard format for the exchange of structured data, which employs a data-description format based on a data-element tag is now being proposed by the American National Standards Institute.

Another example of a standard integrated approach to data translation is provided by XDEF, the XCS Data-Exchange Format used by Xerox Computer Services. XDEF serves as the foundation for the data bridges that are needed in order to integrate the various computing environments encountered by XCS and its clients, including a variety of microcomputer packages for business graphics, financial modeling, database management and even word processing.

An XDEF file is a self-defining flat, alphanumeric file. It has two main parts: a definition part, which contains the definitions and formats of the data in the file; and a data part which contains the actual data records. The data part of an XDEF file contains just one record type, with fixed-length records and fields and no repeating groups. Both data and definitions are in alphanumeric, textual form. XDEF files are thus well suited for the transfer of tabular data or data that can be converted to tabular form.

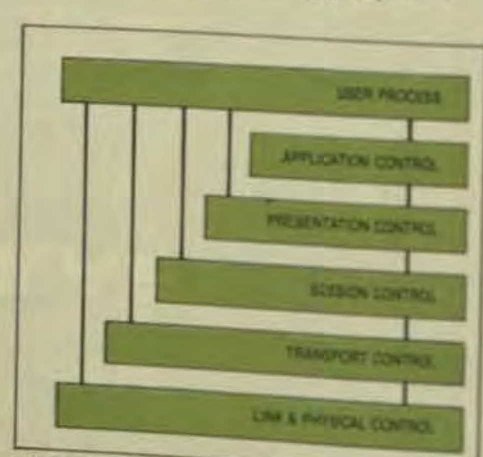

Fig 4 The ISO reference model presents an ideal situation. In reality, a user process typically can access many of the lower-level interfaces.

PROTOCOLS

To provide a data bridge capability, the desired data must be located, accessed, and transferred, and any data representation and structural incompatibilities must be resolved. These functions are incorporated in data-transfer protocols for interprocess communication.

Data-transfer protocols can be distinguished by three levels of difficulty, depending on whether the data consists of a given data type (such as characters), a pointer-free data structure (such as a COBOL record), or a data structure containing pointers.

A structured data-transfer protocol thus includes:

- A standard format for the exchange of structured data.
- The information required to describe the exchanged data.
- The control information (commands) needed to signal the establishment, maintenance, operation, and termination of a session between two processes.
- Data-flow and error-control responsibilities.

Stored data is transferred primarily for four reasons:

LSI-11 USERS:
SIGMA OFFERS
DEC* COMPATIBLE COMPUTER PRODUCTS

ENCLOSURES

MODULES

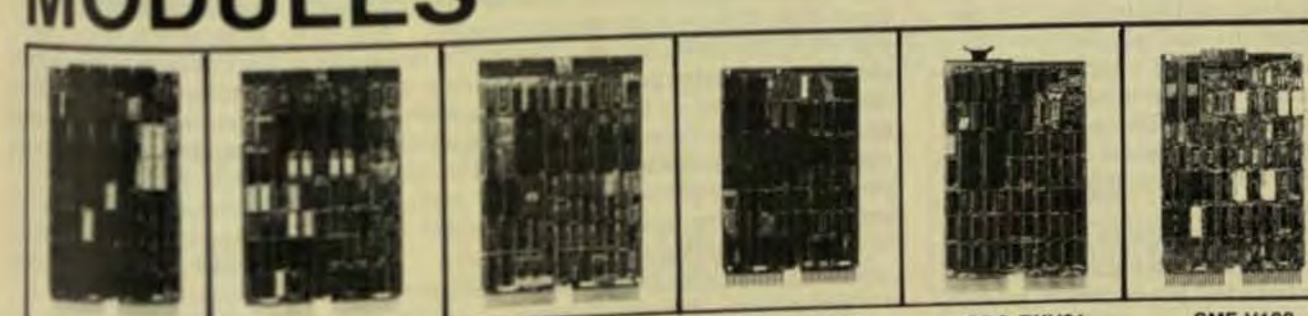

SCD-DLV11J Dual-wide, 4-channel asynchronous serial line interface.

SCD-DLV11J/8 Dual-wide, 8-channel asynchronous serial line interface.

SCD-DZV11 Dual-wide, 8-channel communications multiplexer.

SDC-RLV12 Dual-wide, 5¼'' Winchester controller.

SDC-RXV21 Dual-wide, floppy controller.

SMF-V100 Dual-wide module with two asynchronous serial line interfaces.

SIGMA'S PERSONNEL ARE EXPERIENCED in design and development of DEC* related products and have a comprehensive background in system compatibility. This expertise is available to develop single board, high performance interface and controller boards that are firmware intensive and DEC* software compatible; chassis with user-configured front panels, power supplies, backplanes and drive enclosures.

SIGMA OFFERS a complete line of DEC* compatible disk controllers, interface devices and system enclosures for the LSI-11.

CALL YOUR nearest sales office for information on all products.

Σ Sigma Information Systems

6505C Serrano Avenue, Anaheim CA 92807
Telex 298607 SGMA (714) 974-0166

Orange CA (714) 633-0652 / 771-7320
San Jose CA (408) 243-9112 • Valencia CA (805) 259-6559
Arvada CO (303) 431-0551 • Winter Park FL (305) 628-8896
Boston MA (617) 938-1310 • Minneapolis MN (612) 881-5015
Edison NJ (201) 572-3722 • Albuquerque, NM (505) 294-5423
Gresham OR (503) 241-2216 • Houston TX (713) 451-0466 X71
Vienna VA (703) 448-6800

* Registered trademark of Digital Equipment Corporation

- To store it for subsequent retrieval.
- To print it on the local printer.
- To submit it as a remote job.
- To use it as data input or output.

Another aspect of data transfer is data manipulation. Users often need to create, delete, copy, rename, and otherwise manage remote files. Most data-transfer protocols tend to concentrate on this aspect of the problem because it isn't as difficult as conversion.

Remote data-manipulation control functions include:

- Communication of data-manipulation requests via messages, including coding/decoding into/from messages, and sending/receiving those messages.
- Communication of status via messages.
- Communication of data objects via messages.
- Error control and exception handling, including aborts.
- Global flow control of the data objects generated by each data manipulation.

Session-control functions include:

- Initiation and termination of processes.
- Start/stop of process steps.
- Journalizing of updates.
- Commitment of updates and backup and recovery.
- Resolution of concurrency conflicts.

Three levels of interprocess communication can be identified: job level, call/return, and message-based.

At the job level, interprocess communication is supported only prior to or following a job step. Remote-terminal emulation is an example of a job-level approach to interprocess communication.

The procedure-call mechanism allows interprocess communication at any time, but poses certain problems that reflect uncertain delays, outages, and aborts, and can result in exceptionally long processing delays for the calling process.

Message-based interprocess communication allows a process to request some service and check later to see if the desired results have been returned.

A number of different architectures can support interprocess data communication within a computer network:

- A centralized architecture relies on a central database management system to supply data to each user process.
- A distributed, homogeneous architecture uses the same set of data-management services implemented throughout the network.
- A data-translation approach involves the restructuring of data as it is exchanged between processes.
- An integrated approach involves the use of standard interfaces and a common data language in conjunction with existing data-management systems.

An elaborate example of the data-translation approach is embodied in a data-presentation protocol being proposed by the National Bureau of Standards. A companion file-transfer protocol provides for data transfer.

The NBS data-presentation protocol provides the mechanism through which application processes exchange information between heterogeneous systems. The semantics of the information is preserved, although the syntax may differ. To accomplish this, the protocol negotiates a common transfer syntax from which the participating presentation-control layers translate the data.

In this negotiation, the data-syntax parameters serve to indicate the syntax of those data elements supported by the protocol user. Data-element constraint parameters indicate the constraints the protocol user imposed on the data transformations, such as minimum significance and precision for single- and double-precision floating point numbers. Character-encoding constraints indicate the sets of character codes that are acceptable, in order of preference, and, if the file exists locally, the character code it employs.

File-type constraints indicate the file structure and access method combinations that are acceptable, and which ones are employed in the local file, if it exists. The compression constraint indicates the protocol-user's preference for use of compression of the data to be transferred. Primary and secondary key-location parameters indicate the location of primary and secondary keys in indexed sequential files.

The negotiation is successful when all of the data can be presented in an acceptable syntax to the protocol user, when none of the user-imposed data constraints are violated, and a mutually compatible transfer syntax is selected.

APPLICATION INTEGRATION

The problems encountered in building data bridges to different environments are very similar to the problems of integrating separate applications within the same environment.

Consider the following situation: A new procurement application is to be installed in a manufacturing firm that already has an operational inventory application. Unfortunately, the procurement application interfaces with an entirely different inventory application. How is the new procurement application to be integrated with the existing inventory application?

Of course, the simplest approach would be to convert to the compatible inventory application, but assume there are good business reasons for not doing this. Assume also, for simplicity, that the new procurement application accesses inventory data but doesn't update it. What are the possibilities for accomplishing the necessary integration?

Automatic record access. The data-management system used by the new procurement application can automatically translate its requests for inventory data into requests to the old inventory application's data-management system, and it can translate the received data into the form needed by the procurement application.

One problem with this approach is that current data-management systems are not well equipped to ask foreign data-management systems for data. Such negotiations are technically possible, at least on a limited scale, but even if such an approach were feasible, the translation overhead entailed would almost certainly be unacceptable in any real system, since translation would be necessary between source and target formats as well as between source and target data-manipulation languages.

Programmed record access. The new procurement application can itself be modified to request inventory data either from the old data-management system or from the old inventory application, using an appropriate data bridge. It would then either use that data in the form received or translate it into a more convenient form for processing.

Aside from the not-insurmountable problems of modifying an application to interface with a foreign data-management system, the record-transfer overhead might be substantial enough to make this approach unacceptable. This is nevertheless a feasible approach and is, in general, worth considering.

File conversion and duplicate maintenance. The existing inventory data can initially be converted to the form required by the new procurement application, and the old inventory application can be modified to send updates (issue and receipt transactions) over an appropriate data bridge. The new inventory application (the one compatible with the new procurement application) would then process these updates to maintain a duplicate copy of the inventory data in the new environment.

This approach involves the implementation of the new inventory application in addition to the new procurement application. This implementation can be greatly abbreviated, since its only function is to maintain the duplicate inventory data in response to issue-and-receipt transactions received from the old inventory application. These transactions could, of course, be batched and transferred at any convenient and acceptable frequency, depending on the requirements of the procurement application for up-to-date inventory data.

It is highly likely that the duplicate inventory data, thus maintained in the new environment, will eventually diverge from the master inventory data maintained by the old inventory application. Therefore, it may be necessary to refresh the duplicate data periodically, using the conversion mechanism to retransfer it into the new environment. The frequency with which this refresh will be needed depends also on the requirements of the procurement application and is hard to predict.

Periodic file transfers. The required inventory data, maintained by the old inventory application, can be periodically transferred to the new environment and converted, either before or after transfer, into an acceptable form.

This approach is very similar to the previous one. The difference is that no maintenance of the inventory data is done in the new environment between file transfers. This is feasible only if the procurement application does not require access to the latest inventory data. Analysis of these requirements and of the various data volumes is needed for selection of an approach involving file transfers. ■

Chris Shaw *is manager of technology planning for Xerox Computer Services in Los Angeles. He is interested in the technologies of teaching computers how to know and do things and founded the ACM's Special Interest Group on Programming Languages.*

NET-488
The Laboratory Automation Network

There is only one networking scheme for the engineering development and test laboratory environment. . .NET-488.
NET-488 is a unique hardware/software package which allows the transfer of files between any combination of Unibus, Q-bus, Multibus, S-100 bus or STD bus computers. . .even IBM PCs.

Utilizing National Instruments field proven GATEWAY•488 family of IEEE-488 based interface products, NET-488 provides high data transfer rates, multiple station capacity, and a well defined interface protocol needed for local networking. Because it is IEEE-488 based, NET-488 puts your minicomputers, workstations and PCs together with over 2,000 varieties of test instrumentation.

There is only one networking scheme for the engineering test laboratory. . . NET-488. The Laboratory Automation Network.

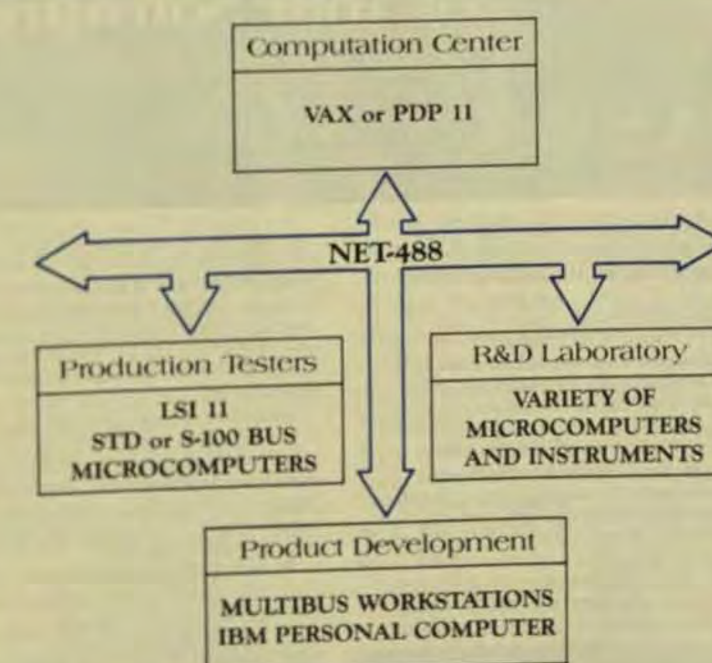

NATIONAL INSTRUMENTS
The IEEE-488 Leader
12109 Technology Blvd. Austin, Texas 78727
800/531-5066 512/250-9119

Selecting software is really a game of strategy: positioning the right product at the right time for the most effective results.

At Whitesmiths, Ltd., our full line of software products is designed to do just that and more. Our time-tested compilers are available for five architectures, complete with Pascal translators and cross-support for a wide range of environments. From the MC68000, to the IBM PC, to the DEC Professional, our compilers and cross compilers cut software development time and reduce dependence on operating systems and instruction sets.

Idris, our exclusive operating system, is powerful, small and extremely portable. Written almost entirely in C, it runs where other operating systems can't even fit – on an MC68000 with no memory management hardware, for example, or on a bank-switched 8080 or Z80. And it outperforms UNIX on comparable systems – an MC68000 with memory management hardware, for instance, or any PDP-11 or LSI-11. It even makes good use of a VAX, so you know you won't outgrow it!

In addition, our full line of Software A La Carte lets you select just the pieces you need to tailor our products for your unusual requirements. From portable libraries to special purpose tools, all A La Carte items are tested, easy to use, and very affordable.

To make things even simpler, Whitesmiths, Ltd. has introduced a new marketing tool: the Authorization Seal. Instead of hassling with complex license agreements and other legal stumbling blocks, you simply affix a serially numbered Authorization Seal, issued with each sale, to any machine that retains a copy of our software. End of hassle. All our customers – end-users, OEMs, and sublicensors – benefit from the elimination of multi-tier licensing agreements and detailed record keeping.

Compilers, Idris, Software A La Carte, and the Authorization Seal: together they can position you for today's competitive market – and tomorrow's. Before you make your next move, call or write to Whitesmiths, Ltd. today. We'll be glad to help you play to win.

Whitesmiths, Ltd.
97 Lowell Road, Concord,
MA 01742, (617) 369-8499
TLX 951708 SOFTWARE CNCM

Whitesmiths, Ltd.
Software Craftsmen

Structured Languages Such as Modula-2 Boost Programmer Output

Pascal's developer now offers a language with improved syntax and a modular approach that frees programmers from the limitations of standard Pascal.

by **David Ramsey**
and **Richard Gleaves,**
Volition Systems

Structured programming languages such as Pascal have shown their value in a number of applications but so far have proven to be far from optimum. As part of an effort to improve upon Pascal and provide programmers with a tool that makes their task easier, Pascal's developer, Niklaus Wirth, has developed Modula-2. To illustrate the benefits of this third-generation structured language, this article contrasts Modula-2's operation with that of other languages and provides some examples of typical Modula-2 code sequences.

Even a brief look at software development over the past 30 years reveals that programming language progress has not kept up with hardware advances. Although FORTRAN and COBOL were adequate for the 1950s, for example, they are outmoded today, even though still widely used for new systems development.

In the view of one computer scientist, "FORTRAN, 'the infantile disorder,' by now 20 years old, is hopelessly inadequate for whatever computer application you have in mind today: It is now too clumsy, too risky, and too expensive to use." Further, he says, "The use of COBOL cripples the mind; its teaching should, therefore, be regarded as a criminal offense."

These remarks come from Edsger W. Dijkstra, who is credited with much of the early thought on structured programming. They appear in his paper, "How Do We Tell Truths That Might Hurt?" (reprinted in *SIGPLAN Notices*, May 1982) and, while extreme, have found much agreement.

The continued use of FORTRAN and COBOL stems primarily from sheer inertia. An entire generation of programmers was raised on them and is reluctant to part with them despite the existence of available alternatives.

STRUCTURED LANGUAGES

The increase in programmer productivity that is possible with a methodical, structured approach has been realized for some time, at least in academic circles. And despite the nasty comments COBOL has received, it was actually the first computer language that had some semblance of structure. Its procedures allowed a program to be constructed in a nonlinear fashion, with separate sections of code handling separate functions. The first block structured language, ALGOL (ALGOrithmic Language), was developed in Europe in 1960. It introduced most of the concepts used in today's structured languages, including:

- Local identifiers, which allowed functional isolation of portions of a program.
- Recursion.
- Call-by-name references.

PASCAL

In 1969, Niklaus Wirth (pronounced "Veert") developed Pascal for teaching structured programming concepts. Implemented in 1970 and 1971, it has gained widespread popularity and has been adopted with particular fervor in the mini- and microcomputer communities. Texas Instruments Corp., among other companies, uses Pascal in the majority of its new system development.

The rationale behind structured languages is that a programmer can implement precise, maintainable code in a shorter time and with less effort than with nonstructured languages. Pascal was the first full realization of this philosophy. Once a programmer gains some experience with Pascal, he finds it natural to exercise the good programming practices that require conscious effort with other languages—to the extent that they can be used at all.

Pascal does have significant shortcomings, though, including:

- Array dimensions must be explicitly declared. Thus, it isn't possible to write a general array-handling routine. (This problem should be addressed in the ISO standard of the language, however.)
- Identifier scope is controlled implicitly by the program's structure. This factor can lead to difficulties.

For example, consider a procedure in an industrial situation that uses several other internal procedures to check the temperatures of reagents comprehensively. If a reagent overheats, the procedure sounds an alarm and returns a value indicating which reagent overheated:

```
PROCEDURE CheckReagents;
(variables, constants)
PROCEDURE CheckReagentA;
(code)
PROCEDURE CheckReagentB;
(code)
PROCEDURE CheckReagentC;
(et cetera)
```

Procedures CheckReagentA, CheckReagentB, and CheckReagentC are invisible to any code-calling procedure CheckReagents. Suppose in a later version of the program it becomes necessary to check some (but not all) of the reagents at times other than when the previous procedure checked them. The code has already been written, but because it resides in the old main procedure, it can't be accessed without involving the remainder of the old procedure.

Currently, the only solution is to code in the necessary procedure redundantly, or remove the inner procedures needed and make them global. The first solution is inefficient, and the second compromises the program's structure. There is currently no elegant way around such a dilemma in Pascal.

MODULA-2

One alternative is Modula-2. This language is an extension of Modula, a systems language originally designed for multiprogramming work and implemented in 1975. Modula-2 began in 1977 as part of a project at the Institut fur Informatik of ETH, Zurich, to design an integrated computer system including both hardware and software.

In his book, "Programming in Modula-2" (Springer-Verlag, 1982), Wirth explains: "Modula-2 emerged from careful design deliberations as a language that includes all aspects of Pascal and extends them with the important module concept and those of multiprogramming."

Modula extensions include:

- The Module concept.
- Control of scope.
- Low-level facilities.
- Processes.

MODULES, SCOPE CONTROL, AND SEPARATE COMPILATION

Any program in Modula-2 is a module. Modules may contain procedures and functions as well as constants, variables, and types. Objects and identifiers in module X may be used in module Y if they are made known to or IMPORTED into Y.

This Modula-2 feature supports separate compilation, a very useful attribute that allows often-used routines to be kept in object form and linked at run time. Many Pascal and FORTRAN implementations offer an independent compilation facility often called separate compilation. But the Modula-2 compiler has access to the objects and identifiers in the separately compiled modules and performs type checking of separately compiled objects, whereas in typical Pascal/FORTRAN independent compilation schemes, the routines are declared to be external to the main program, and their definitions (on a source-code call level) aren't supplied with type checking.

Modula's IMPORT and EXPORT keywords are the key to this function and also provide explicit control over scope. For example, in Pascal you might have:

```
Program Demo;
Procedure DrawLine
(X1,Y1,X2,Y2, : Integer); External;
Procedure Point(X,Y : Integer);
External;
Function PSet : Boolean;
External;
```

In Modula, on the other hand, you would say:

```
MODULE Demo;
FROM graphics IMPORT
Drawline,Point,PSet;
```

This Modula program fragment assumes the existence of another module called "graphics," which is stored in the module library. The graphics module may import items from other modules, but that isn't the programmer's concern at this level. In the reagent checking routine, the problem disappears with Modula-2 coding. CheckReagents is declared to be a MODULE, or functional unit of the program. It can then "export" objects if necessary:

```
MODULE CheckReagents;
EXPORT CheckReagentA, Check-
ReagentB, CheckReagentC;
PROCEDURE CheckReagentA;
(code)
PROCEDURE CheckReagentB;
(etc.)
```

Access to any of the inner procedure can be had with the IMPORT statement:

```
IMPORT CheckReagentB,
CheckReagentC;
```

Note that it isn't necessary to specify the entire heading when procedure and function procedures are imported and exported. The compiler provides type checking for object compatibility. An IMPORTed identifier might have a name that conflicts with another identifier, for example.

Some languages permit this conflict and call the bug "overloading." Modula-2 calls it an error. Fortunately, identifiers can be QUALIFIED in EXPORT lists. A QUALIFIED identifier must be referred to with a module name.identifier syntax, similar to the record syntax Modula shares with Pascal. For example:

```
EXPORT QUALIFIED Check-
ReagentB;
```


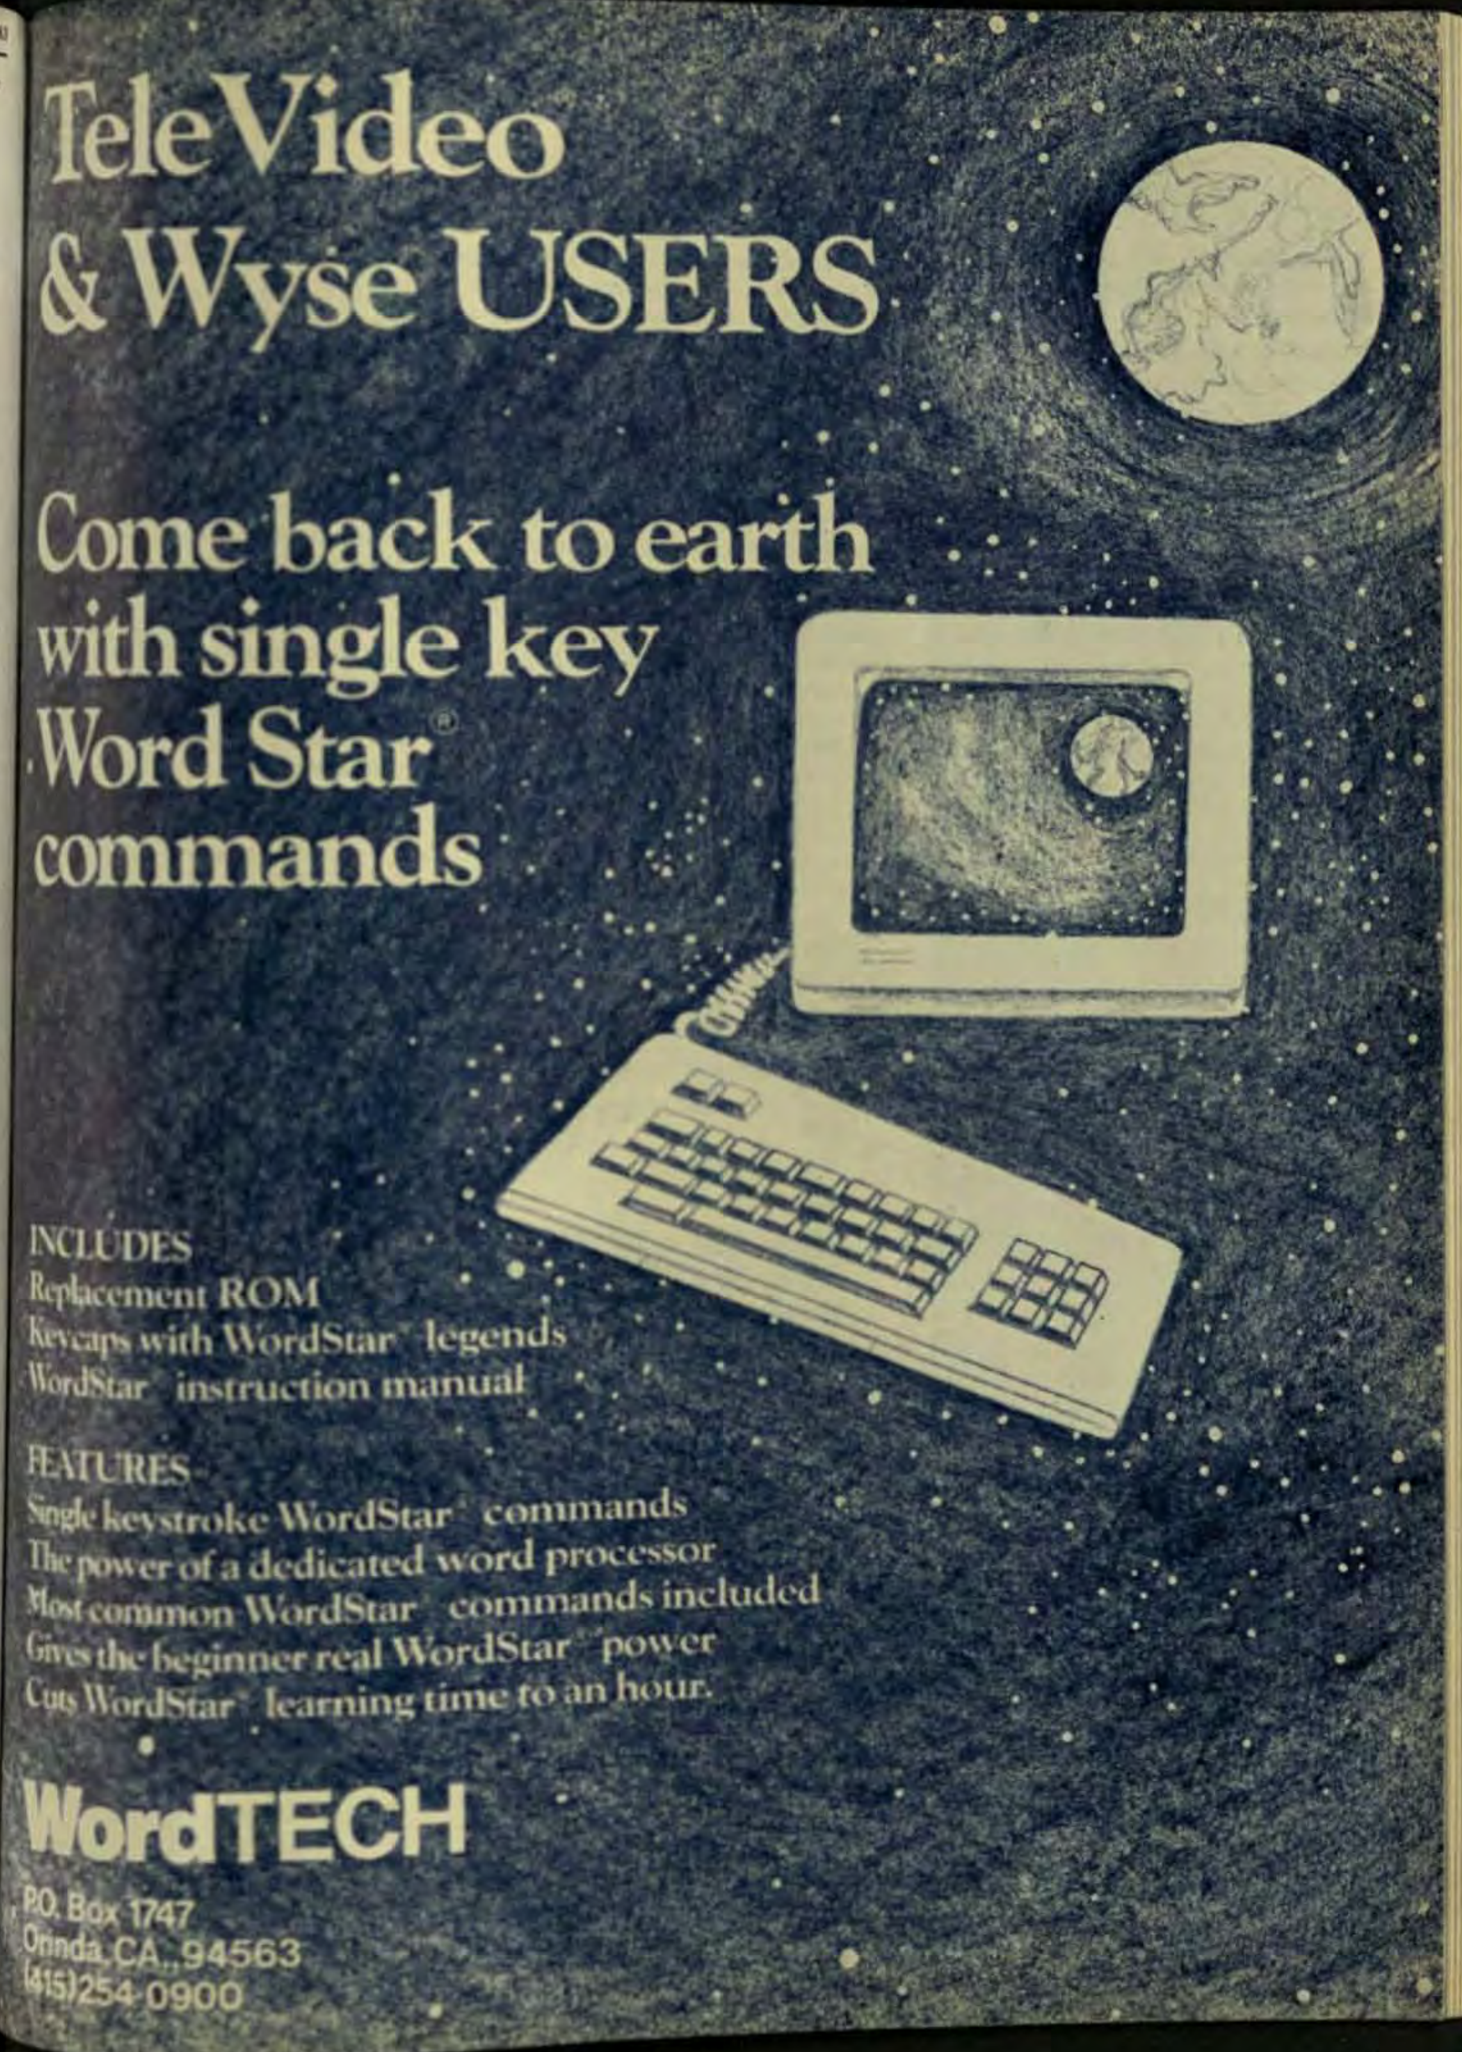

References to CheckReagentB in this case must be qualified by preceding them with

CheckReagents.CheckReagents.
CheckReagentB;

The IMPORT statement can also unqualify qualified identifiers. For example,

FROM CheckReagents IMPORT
CheckReagentB;

allows CheckReagentB to be referenced without the preceding module name:

CheckReagentB;

Modula uses modules very heavily in its implementation. No system I/O facility is specified in the formal definition of the language, but some standard procedures are assumed to be available in the standard module InOut. Several such standard modules are defined, including modules for I/O, math operations, processes, and graphics.

In "Programming in Modula-2," Wirth says, "The principle motivation behind the partitioning of a program into modules is . . . the establishment of a hierarchy of abstractions. Each module constitutes an abstraction if we regard it from the 'outside.' We even wish to go one step further: We wish not only to ignore the details of its innards, but to hide them."

The hiding of a module's inner functions has two important benefits. First, the programmer doesn't have to know about the workings, and he need not be concerned about possible side effects. Second, the module can be changed or improved without the need to recompile the programs that use it.

Modula thus provides for a functional decoupling of modules. This provision requires that the definition portion of a module be separate from the implementation part. UCSD P-system programmers who have written "units," with their "interface" and implementation sections, will be familiar with the concept.

A module's definition portion is defined as follows:

```
DEFINITION MODULE Demo;
⟨ export list ⟩
⟨ Constants, types, variables ⟩
⟨ Procedure declarations ⟩
END Demo;
```

This definition portion contains all the information that will be available to calling programs or modules, including all exported identifiers. Procedures are listed with their full headings.

The module's implementation portion (with the IMPLEMENTATION MODULE header) will contain the actual definitions and code for the module's operation. Separately compiled modules are separated into definition and implementation sections.

LOW-LEVEL FACILITIES

Modula-2 also provides low-level facilities that permit users to bypass the Modula-2 structure and type-checking operations. These facilities prove useful when it's necessary to manipulate data—such as machine control addresses or data from another system or program—that cannot be mapped into a Modula structure. The low-level facilities aren't part of the formal language definition, but are provided in the standard module SYSTEM included with the language. Use of these facilities will generally render a program nonportable, however.

One such facility is the ability to use type identifiers as type transfer functions. For instance, the following example uses the type INTEGER to convert a character variable to type INTEGER:

```
VAR
  Test : INTEGER;
  Q : CHAR;
PROCEDURE CheapOrd(C :
CHAR) : INTEGER;
BEGIN
  RETURN INTEGER(C)
END CheapOrd;
```

The WORD and ADDRESS types are predefined according to the machine they are implemented on. The standard function SIZE can be used to determine the size of variables. Type ADDRESS is defined as a pointer to WORD and can be used to perform pointer arithmetic. The standard function ADR returns the memory addresses of variables. Variables can also be declared to reside at fixed memory addresses.

PROCEDURE TYPES

Modula also allows the creation of procedure variables. Type assignments are made with dummy procedure headings. For example:

```
TYPE
  Proc1 = PROCEDURE
  (INTEGER);
  Proc2 = PROCEDURE(CHAR,
  INTEGER,CARDINAL);
```

Or, in the case of function procedures:

```
TYPE
  FProc = PROCEDURE(REAL):
  REAL;
```

Any procedure or function type may be assigned in this manner. Variables declared with this type can then take procedure names as values providing that the procedure whose name they take matches the declaration made for that variable type. A variable of type FProc, for example, could take on the value of a transcendental function such as sin(x), with the statement:

```
FProcVar := sin;
```

A call to FProcVar(x) would then be equivalent to a call to SIN(x).

CONCURRENCY AND PROCESSES

Modula-2's concurrent programming capabilities are limited to loosely coupled processes—processes that interact relatively infrequently and at well-defined points in the program. The standard module PROCESSES provides the following definitions:

```
DEFINITION MODULE
Processes;
  EXPORT QUALIFIED
    SIGNAL, StartProcess,
    SEND, WAIT, Awaited, Init;
  TYPE
    SIGNAL;
  PROCEDURE StartProcess
  (P:PROC; n: CARDINAL);
  (*Start a process P with a
  workspace of n words*)
  PROCEDURE SEND(VAR s:
  SIGNAL);
  (*One process waited for s is
  resumed*)
  PROCEDURE WAIT(VAR s:
  SIGNAL);
  (*Wait for some other process
  to send s*)
  PROCEDURE Awaited(s:
  SIGNAL): BOOLEAN;
  (*Awaited(s) = "at least one
  other process is waiting for
  s"*)
  PROCEDURE Init(VAR
  s:SIGNAL);
  (*Compulsory initialization*)
END Processes.
```

The exact manner in which Modula-2 processes execute depends on the system's hardware. Appropriate implementations on multiprocessor machines can provide true concurrent process execution. In most implementations, though, only one processor is available, so the processes are handled individually. Such processes are referred to as quasiconcurrent.

A process is started by a call to StartProcess(P,n). P is a procedure variable, and n describes the workspace size. The value of n is chosen by considering the number of local variables and number of internal calls made by the process.

Processes can communicate with each other via shared variables or signals. Signals carry no data, but they synchronize the operation of processes. A process waiting for a signal will start when a signal S is sent by any other process. A signal activates at most one process at a time. Note that signals aren't directed to a specific process. Shared variables are commonly variables that are global to the processes in question.

IMPROVED SYNTAX

When Modula-2 was designed some changes were made in the syntax (compared to Pascal's syntax) to improve its regularity. The two most notable alterations are the almost complete elimination of the BEGIN keyword and an integration of functions and procedures.

The BEGIN keyword is no longer used in compound statements. For example, in Pascal, usage is as follows:

```
IF (condition) THEN
  BEGIN
  (statement)
  (statement)
  (etc)
END;
```

In Modula-2, the same statement would be:

```
IF (condition) THEN
  (statement)
  (statement)
  (etc)
END;
```

In Pascal, procedure and functions are separate entities with separate declarations:

```
PROCEDURE Demo(X :
INTEGER; Name : String);
FUNCTION DiskAvail :
BOOLEAN;
```

In Modula-2, there are procedures and functions procedures:

```
PROCEDURE Demo(X :
INTEGER; Name : String);
PROCEDURE DiskAvail () :
BOOLEAN;
```

The only difference between Pascal and Modula-2 in the declarations of function procedures is Modula-2's assignment of a type in the heading. Note that both regular procedures and function procedures with no parameters carry parentheses. Function procedures return their value with the RETURN statement, because the assignment of a value to the procedure name conflicts syntactically with procedure variables. ■

David Ramsey *is a software engineer at Corvus Systems. He holds a BA in biology from Austin Collge and has been involved with Pascal since 1978 and Modula-2 since 1980.*

Richard Gleaves *is a software engineer at Volition Systems. He has a BA in computer science from the Univ. of California, San Diego, and was a member of the UCSD Pascal Project.*

WAS THERE A GOOF?

Are there some issues of Computer Technology Review that you either didn't get, or lent to friends and have never gotten back?

We understand that you want them all. So you'll be relieved to learn that a limited number of former issues are still available. But at the price of only $15 each they won't last for long. This time, don't let there be a goof. Stock up on CTR now.

Send us your check for the quantity you want at $15 each, to Computer Technology Review, 924 Westwood Blvd., Suite 650, Los Angeles, CA 90024. Or call us with your Visa/Mastercard/American Express charge number. (213) 208-1335.

PACXNET Welcomes IBM® with New PIN 3270 Multiplexers

Gandalf's PIN 3270 Multiplexers have just opened the door for easy access to IBM information display systems. Now both IBM and non-IBM computer resources can be readily accessed from ASCII devices in your PACXNET data network.

The new PIN 3270 multiplexers and attached devices appear to your IBM host computer as IBM 3271 (BSC) or 3274C (SNA) compatible cluster controllers with attached 3277/78 display stations and 3287 printers. The PIN mounts into PACX data switching node cabinets or is available as a stand-alone. Up to 15 asynchronous channels and one or two BSC/SNA synchronous channels can be supported, depending on the PIN model selected.

As part of Gandalf's PACXNET, PIN 3270 Multiplexers add system flexibility and insure system security for multi-vendor systems. Introduce your ASCII devices to your IBM resources. Welcome IBM to your PACXNET data network with PIN 3270 multiplexers! For complete details, contact your local Gandalf office.

gandalf™

Fully supported technology from concept to customer

USA: Gandalf Data, Inc.
Canada: Gandalf Data, Ltd.
Europe: Gandalf Services, S.A.
U.K. Gandalf Data Communications Ltd.

IBM® is a registered trademark of International Business Machines Corp.

Network Diagnostic and Control Play Vital Role In Large-Scale Networks

Computer disciplines can be applied to traditional data-communication management techniques for increased reliability through improved diagnostics and control.

by **Ronald R. Ritts,** Racal-Milgo

In most data-communications networks today, one minute of downtime can rapidly defeat the benefits of the best network design and equipment. And the failure of any element may disrupt service in the entire network. It may result in excessive mean time to repair, downtime, and error rates—any one of which might prove catastrophic.

Thus, when users plan data-communications networks of any magnitude today they emphasize network diagnostic and management systems that perform not only data-communications diagnostics, but also for communications management and control. Implemented entirely from a central site, these utilities permit operators to scan, identify, and record faults, failures, and unusual conditions in a network, as well as to take action to restore service and coordinate repair functions. Such systems offer managers an information overview of their networks and assist them in coping with the many problems large-scale networks can present.

MICROPROCESSOR-BASED DIAGNOSTIC CONTROLLERS

These systems evolved because network demands are changing from basically qualitative systems to quantitative systems. Qualitative systems monitor network performance for faults between preset points. Quantitative systems monitor network performance in a similar manner for faults, but also evaluate networks for specific undesirable conditions. Most quantitative systems today incorporate certain network-restoral functions.

One of the earliest and simplest qualitative devices, a microprocessor-based test set, is still used in limited applications. It functioned primarily as a diagnostic aid to be patched in as required to diagnose line faults. However, this device required human intervention to implement tests. In some cases, it became either impractical or too expensive for testing of specific remote locations, and it didn't permit the operator to perform preventive maintenance or run tests on one portion of the network while looking for problems in another.

This restriction led to the development of secondary-channel diagnostic functions that didn't interfere with main-channel data—a technique that became popular after a fundamental breakthrough was made in high-speed, narrow-band data communications.

The narrow-band technique concentrated high-speed, main-channel data transmission within a very narrow (800 Hz) portion of the available bandwidth (300 Hz to 3300 Hz) and centered it at 1700 Hz. As a result, channels beyond the 800-Hz spread could be used for asynchronous functions such as diagnostics.

The low-speed secondary channel, referred to in this discussion as T-7, transmits data at 75 bps and permits an operator at a central site to monitor selected remote modem/terminal interfaces manually. Another basic reason for the T-7 alarms is to pick up equipment quirks that could lead to faults, at which time central site can alert repair crews at the first sign of an erratic function. Often they can reroute the circuitry if a piece of equipment goes down so that while repairs to it are being made, the actual failure is transparent to the computer and data-terminal equipment.

This design, incorporated into two larger, microprocessor-based diagnostic controllers, still serves widely in small- and medium-sized networks. The controllers enable an operator to scan entire networks automatically for faults and degraded line conditions. In many cases, they enable the operator to alert service organizations even before faults and interruptions occur.

If the controllers detect a problem, they sound an alarm message and report the nature and location of the problem and the equipment or line involved. Using this data, an operator can enact a variety of automatic, noninterruptive, and interruptive tests to confirm a fault and obtain further information on its nature.

Such tests offer a number of advantages. They are automatic and routine—thus less time-consuming than the manual tests of the test sets described earlier—and they don't require highly trained operators to interpret the results, which are displayed in clear English. They forego the operator's need to change test equipment, and they enable tests to be conducted without the involvement of personnel at the affected sites. They also eliminate the human factor in determining whether particular test results are serious enough to warrant action.

Diagnostic controllers also have the ability to operate with up to 4064 modems. The addresses are arranged into 16 lines with 254 addresses per line. Optional digital-mixing modules can expand the number of lines and permit a more flexible distribution of the addresses (Fig 1). The controller can also operate through tandem or multiplexed lines with the use of digital-mixing modules, thereby extending control through all levels of a network.

Using automatic procedures, the diagnostic controller can spot line or equipment problems before or after working hours and on weekends. They also provide warnings of impending line or equipment failures during network operation. In the event a problem does occur, they permit automated test routines to simplify troubleshooting.

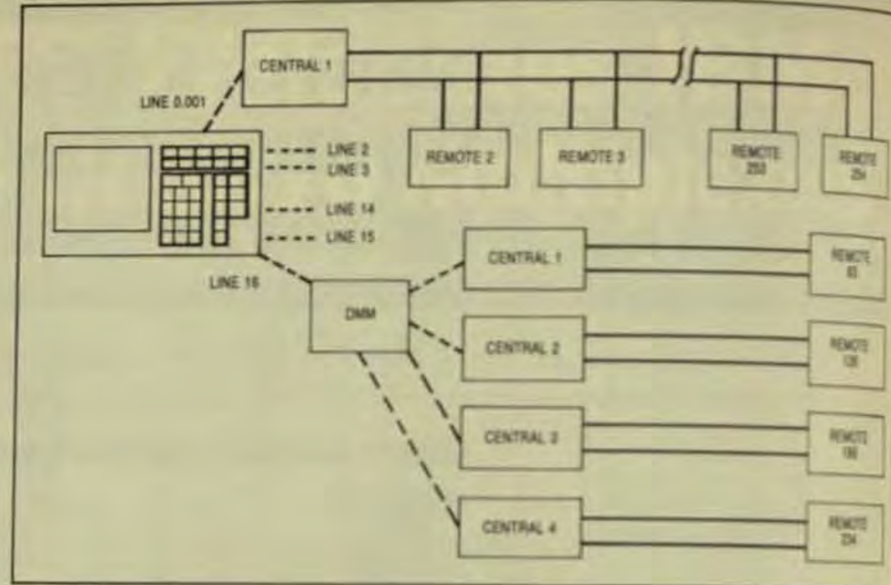

Fig 1 One diagnostic controller can handle up to 16 lines. Each line can individually address up to 256 modems for a total system capacity of 4064 modems.

Since the controller polls its channels and performs its sequences one at a time, rapid network expansion often presents complications. A network having 60 lines, for example, might expand to 160 lines or more, which would require much faster diagnostics than the controller's polling characteristics could provide. Obviously, much faster diagnostics are needed to maintain proportional operational and cost efficiencies.

A second microprocessor-based controller, which performs all the functions of the first but contains expanded capabilities, meets this need. It can parallel poll all channels at the same time and poll more types of modems (such as loop modems) as well as control remote transfer switches so that remote-site backup devices can be placed on-line even before a system component goes down.

Synchronous microprocessor-based modems used with this second controller contain a microprocessor-based T-7 test card. This feature includes a transmitter/receiver that provides a separate 75-bps secondary channel over which the remote modems and the diagnostic controller communicate. It also allows the addressed modem to decode 8-bit test-command words transmitted by the diagnostic controller to perform the tests requested, and to frame, format, and transmit the results in 8-bit words back to the controller to detect unusual modem conditions and automatically transmit appropriate alarm messages to the controller.

For identification purposes, each modem is given a unique electronic address to let the controller determine the location of a modem and distinguish it from all other modems in a system. The modem address also tags information being sent back to the controller from central or remote modems, which helps provide specific information on faulty lines or equipment.

Also, the controller is able to receive external alarm signals simultaneously from central-site and remote-site modems, even while tests are in progress. The T-7 circuitry in each modem stores enough power so that whenever it loses ac or dc power it can still inform the upstream modem, which reports the failure to the controller at a central site.

The T-7 can also send an alarm if a modem streams—that is, if it transmits an undesirable continous carrier in a polled network configuration. In the event of a streaming alarm, the diagnostic controller can be commanded by the operator to squelch the streaming modem and permit normal network operation until it's repaired. The feature proves particularly important in multidrop circuits where streaming can knock out an entire circuit.

In addition to its noninterruptive monitoring functions, the controller performs a number of troubleshooting routines automatically. Tests can be conducted promptly and without changes of equipment. One such test is the end-to-end test, in which a com-

Comtech allows the OEM to employ broadband techniques for interconnecting many types of hardware, using our M570B modem. This includes all types of terminals, computers, microprocessors, computer graphic stations, process automation devices, remote graphic stations, etc.

Transmission rates of 3.5 MBPS are possible with the M570B. For more information concerning the application areas of your particular hardware and our versatile broadband modem, write or call Comtech Data Corporation, 350 North Hayden Road, Scottsdale, Arizona 85257, call (602) 949-1155 or TWX 910-960-0085.

Comtech is a recognized leader in the field of high speed interactive data transmission using coaxial cable modems.

COMTECH Data Corporation
A SUBSIDIARY OF COMTECH TELECOMMUNICATIONS CORP.

Comtech Offers Broadband Capabilities at OEM Prices.

mand from the controller activates pattern generators in both an addressed central site and a remote site. The modems simultaneously transmit and receive this data, each comparing the received pattern with the one stored in its memory, taking an error count, and relaying the results back to the controller every 4 s for display on a screen.

These tests cover modem interfaces and the connecting lines between the modems. A high error count indicates a fault in one modem transmitter and a pair of transmit lines, or in a receiver and the receive lines. Other tests can then be conducted to determine exactly which component contained the fault.

SELF TEST

Another function of the controller is to enable an operator to put modems into a self-test mode to determine the integrity of central- and remote-modem transmitters and receivers. When a modem receives the self-test command, it sends itself a test pattern. The configuration loops the transmitter and receiver together exclusive of phone lines. From the test pattern, the modem takes an error count and automatically transmits the results back to the controller. If the error counts for both modems appear normal, a line test follows.

The line-loop test performed by the controller provides the operator with specific error counts for a line pair over a given period of time. The test can also be performed in conjunction with data scopes to provide a dot pattern for visual inspection of lines for such conditions as phase hits, and to determine error counts for the data-terminal equipment (DTE) loop path between the remote modem's receiver and transmitter.

In the event of line problems, the controller also permits operators to switch an affected remote modem from dedicated to dial operation by means of automatic line adapters. The modem and DTE can then operate over the dial-up line while the dedicated line is being restored. At that point, dial calls terminate at the central site by way of the controller, and normal operation resumes over dedicated lines. A multiline adapter ties the dial backup system into the network.

MINICOMPUTER SYSTEMS DIAGNOSTIC CONTROLLERS

Information-processing advances eventually called for reconfiguration of the more advanced microprocessor-based diagnostic controller for the same reason the earlier model had to be reconfigured—network expansion. While the more advanced microprocessor-based controller supports up to 16 channels, minicomputer-based systems support from 16 to 256 and permit multiple functions to be performed by up to four operators.

Essentially, a minicomputer-based diagnostic system functions as a tester, a controller, a database manager, a file organizer, and a report generator. As a tester, it monitors units for the assurance of communications, reports failures, and initiates a full complement of remote tests. As a controller, it controls and monitors special equipment, such as modems and switching devices. As a database manager, it keeps records on sites, units, and channels, and may maintain, modify, and display those records (in which case the database could be compared with the network automatically). As a file organizer, it keeps a significant-event file that records the occurrence of test and other system events. As a report generator, it can generate hard-copy reports on every major database item, including the significant-event file.

In many medium-size networks today, multiple controllers serve very effectively. However, the question the user of a medium- to large-scale network must ask is how long such controllers will suffice in view of future network expansions and the management utilities available in minicomputer-based systems.

Still another question must be asked by the manufacturer: Can the problems created by ever-increasing network expansions be overcome best by merely expanding the capabilities of a microprocessor controller, or by designing a computer system in itself that will handle all the controller functions previously discussed, keep the T-7 feature, and add more comprehensive diagnostics, reporting, recording, and management capabilities?

An interesting cost/price factor makes the second course preferable: To design and manufacture the new minicomputer-based diagnostic system would make the price of the network-management facility for large-scale networks low in terms of the cost per modem—lower, in fact, than a more complicated microprocessor-based controller that would require continual central and remote modifications.

Moreover, a more complex controller eventually would have presented additional expensive complications in most growing networks because the T-7 microprocessor operation contains context sensitivity—that is, the controller interprets data from T-7-equipped modems in terms of four basic contexts contained in the program stored in the microprocessor card.

For example, the second bit in a data stream returning from a T-7 unit may be interpreted as having one of four meanings, depending on whether the modem serves in multiport or point-to-point, central or remote configurations. This problem would clearly arise in tail circuiting, in which point-to-point lines radiate from multipoint lines, or when any other type of modem circuit hooks into different types of existing circuits. The T-7 function from modems in such instances might not be interpreted correctly by the controller, which treats all modems in a line with the same four-valued context.

Rather than change the stored program in the modems, which may number in the thousands in large-scale networks, the minicomputer-based diagnostic system can include a database with information on the location of any given modem, its type, whether it has a backup mode, which upstream modem it communicates with, and so forth.

Therefore, the T-7 feature can be retained—and for good reason. It typically serves in microprocessor-based modems and in the more recent of the two diagnostic controllers discussed. Yet it's totally compatible with minicomputer-based diagnostics. A new T-7 version for modems in minicomputer-based systems offers various enhancements for diagnostic and management functions but is also compatible with the previous version.

This upward-compatibility factor allows a network undergoing major expansion to be changed to a minicomputer-based system without replacement of all the modems (unless analog parameters, serial-number reporting, specific recordkeeping and

Meet the data communications experts.

Gandalf modems

They're experienced.
They're reliable.
And they're from Gandalf.

Gandalf modems are optimized for maximum communications link uptime and reliability. Check below for the Gandalf modems to fit your data comm requirements. Modems for the short haul, the long haul and the in-between. Modems for copper wire, phone company circuits, UHF/VHF, and satellite links. Coax and fiberoptic modems are also available. When it comes to modems, call GANDALF . . .

- ☐ Stand-alone or space-saving rackmounted units.
- ☐ Meet EIA RS232C and CCITT V.24/V.28 requirements.
- ☐ Excellent user ratings.

Model No.	Type[1] A	Type[1] S	Configuration[2] PP	Configuration[2] MP	Mounting[3] SA	Mounting[3] RM	Data Rates (bps)	Circuit Limitations[4]: Max. Distance (26 ga. in miles)	U	L	Act	Circuit Facilities: Private Wire	Leased Circuit 41404	Leased Circuit 3002
[illegible]	X		X		X		0-19.2K	10 @ 1200 to 1.75 @ 19.2K	DC			X		
[illegible]	X		X			1	0-19.2K	10 @ 1200 to 1.75 @ 19.2K	DC			X		
[illegible]	X		X	X	X	2	0-9600	5.5	X			X	X	
[illegible]	X		X	X		1	0-9600	5.5	X			X	X	
[illegible]	X		X		X	2	50-9600	6	X			X	X	
[illegible]	X		X		X	2	50-4800	12		X		X	X	
[illegible]	X		X	X	X	2	0-9600	12 @ 2400 - 4 @ 19.2K	X			X	X	
[illegible]	X		X	X	X	2	0-19.2K	12 @ 2400 - 4 @ 19.2K	X			X	X	
[illegible]	X		X		X	2.3	4800	200		X	X	X	X	X
[illegible]	X		X		X	2.3	Two channels @ 2400	200		X	X	X	X	X
[illegible]	X		X	X	X	2.4	4800	200		X	X	X	X	X
[illegible]	X	X	X	X	X	1	0-50K async 1200-38.4K sync	50 Feet (EIA/V.24 drive)	X DC			X		
[illegible]		X	X		X	2.3	4800	200		X	X	X	X	X
[illegible]		X	X		X	2.3	Two channels @ 2400	200		X	X	X	X	X
[illegible]		X	X		X	2	2400-9600	6 @ 2400 to 2.5 @ 9600	X			X	X	
[illegible]		X	X	X	X		2400-19.2K	12 @ 2400 to 4 @ 19.2K	X			X	X	
[illegible]		X	X	X		2	1800-9600	12 @ 2400 to 4 @ 19.2K	X			X	X	
[illegible]		X	X		X	5	9600	Any 3002 circuit		X	X	X	X	X
[illegible]		X	X	X	X		9.6-50K* (56K Optional)	.4 @ 50K*	X			X	X	
[illegible]		X	X	X	X	2.4	4800	200		X	X	X	X	X

[1] Type: A = Asynchronous S = Synchronous
[2] Configuration: PP = Point-to-point MP = Multipoint
[3] Mounting: SA = Stand-alone RM = Rack Mounted
1 = 100 series cabinet, 2 = 3000 series cabinet,
3 = takes two card positions, 4 = takes 3 card positions,
5 = 9000 rack)
[4] Limitations: Distance is in actual circuit miles.
U = Unloaded wire L = Loaded wire Act = Amplified Circuit
* 56 optional (3.4 mi @ 19.2K to 0.75 @ 256K)

gandalf™

Gandalf Data, Inc.
1019 S. Noel Avenue
Wheeling, Illinois 60090
TWX: 910-651-4951
TEL: (312) 541-6060

other features associated with the minicomputer-based systems' modems were immediately necessary.

All factors considered, the minicomputer-based systems were conceived as very flexible computers in themselves. They have their own CPUs, disks, I/O, terminals, and associated equipment (Fig 2). However, the system controller takes care of diagnostic and management controls only. It communicates with all remote devices—modems, switches, and terminals.

Microprocessor-driven backbone and remote modems have programming capabilities for such tasks as serial-numbers, transmission, reporting faults that have occurred within or around the modem, the nature of the fault, and in some cases analog parameters or line-level performance.

From such capabilities in general, a user could also obtain reports utilizing appropriate files from the system. For example, he could locate inventory, a decidedly advantageous feature in systems containing thousands of remotely located terminals, controllers, modems, cables, multiplexers, and other devices. The user could also maintain information on such items as service vendors, hours of operation, and contract terms.

Diagnostic systems don't interact directly with all components in a typical data-communications network. Instead, they monitor and control key checkpoints, which are defined in the database as sites, network-interface processors (NIPs), channels, and units:

- Sites are the geographic locations of a device or group of devices, such as modems, switches, and NIPs.
- The NIP is a semi-independent, intelligent processing node that links the host and downstream units.
- The channel is the communications line through which diagnostic data travels between the NIP and its associated units.
- Units represent the hardware devices capable of communicating over a secondary diagnostic channel.

The numbers of sites, NIPs, channels, and units defined in a database will depend on the size and structure of a given network. Loading of the database creates information files on all network components.

The database has been designed so that information on each network segment can be accessed in multiple ways.

If a large number of units must be added at a particular site, the communications manager may want to see the present organization of the site by displaying the database site function. Therefore, before loading the database, the user can identify and assign addresses to the devices that will be communicating with the minicomputer-based system. Accurate address assignment is important because the resolution of alarms and messages received during system operation depends on the addresses.

As far as future network and product support go, manufacturers are designing customized systems to enhance the capabilities of the minicomputer-based diagnostics. A wide range of supplementary dial-backup, switching, and diagnostic equipment can be coupled with a basic system to increase the flexibility and control capabilities of any given network. Future capabilities will provide greater command over the critical remote sites from the control site of the network.

Increased ability to restore and reconfigure network components from the central site is also under development, as are increased dial-backup, switching, diagnostic, and reporting capabilities. In addition, as networks continue to expand, other associated management situations are being anticipated. ■

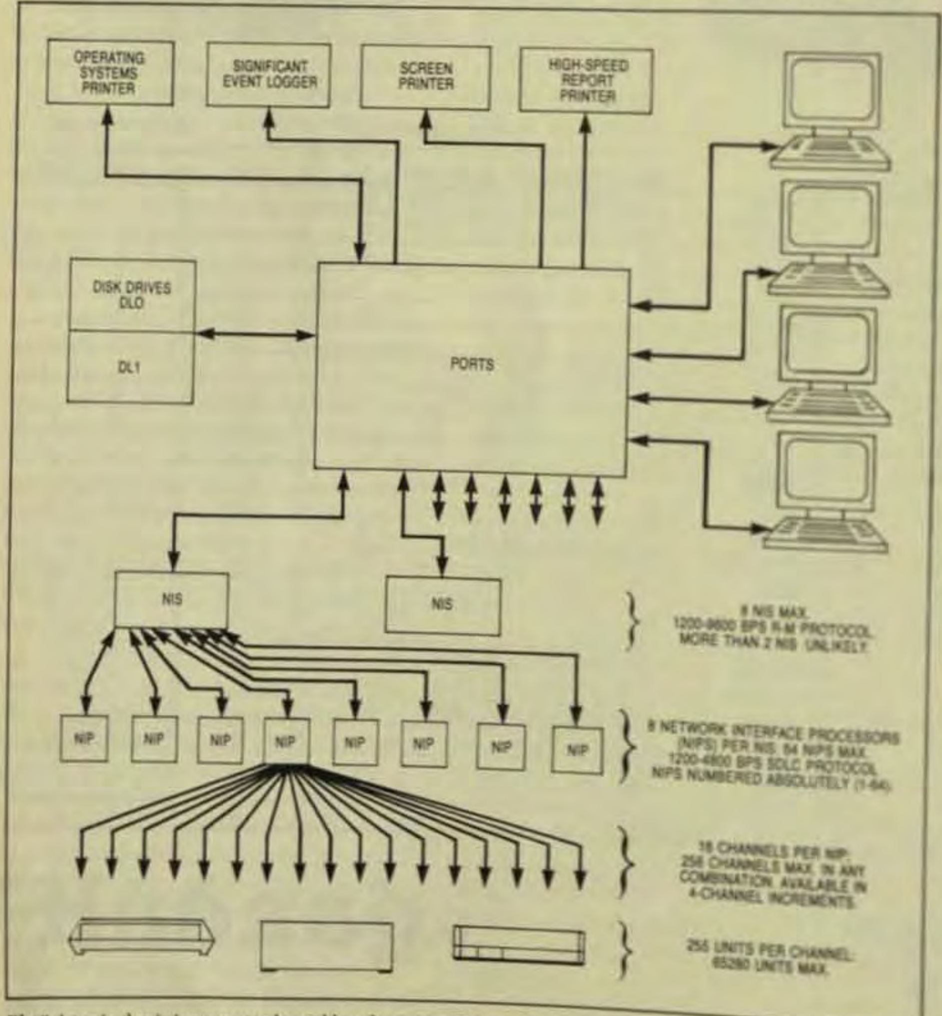

Fig 2 A typical minicomputer-based hardware configuration for network diagnostic control and monitoring can handle up to 10,000 modems.

Ronald R. Ritts, *data communications writer, has been with Racal-Milgo for the past several years. He earned a BA in general science from the Univ. of Pittsburgh.*

bo-sherrel co.

Nov. 29, 1982

Press Release

bo-sherrel co., today announced an improved version of their short-haul modem M-1. The new unit, called M-1A is fully compatible with the old units, but supports hardware handshake and control signals. This enables the use of Data Terminal Ready and Request to Send as control of data flow.

Lightning protection has also been incorporated as a standard feature. This was previously an optional item.

For ease of installation a switch is built in which exchanges the transmit and receive leads.

All of the improvements have been made available without any increase in cost. Availability is immediate.

6101 Jarvis Avenue • Newark, CA 94560 • 415 + 792-0354

Alspa

ZERO

ANNOUNCING THE HIGH SPEED, LOW COST NETWORK

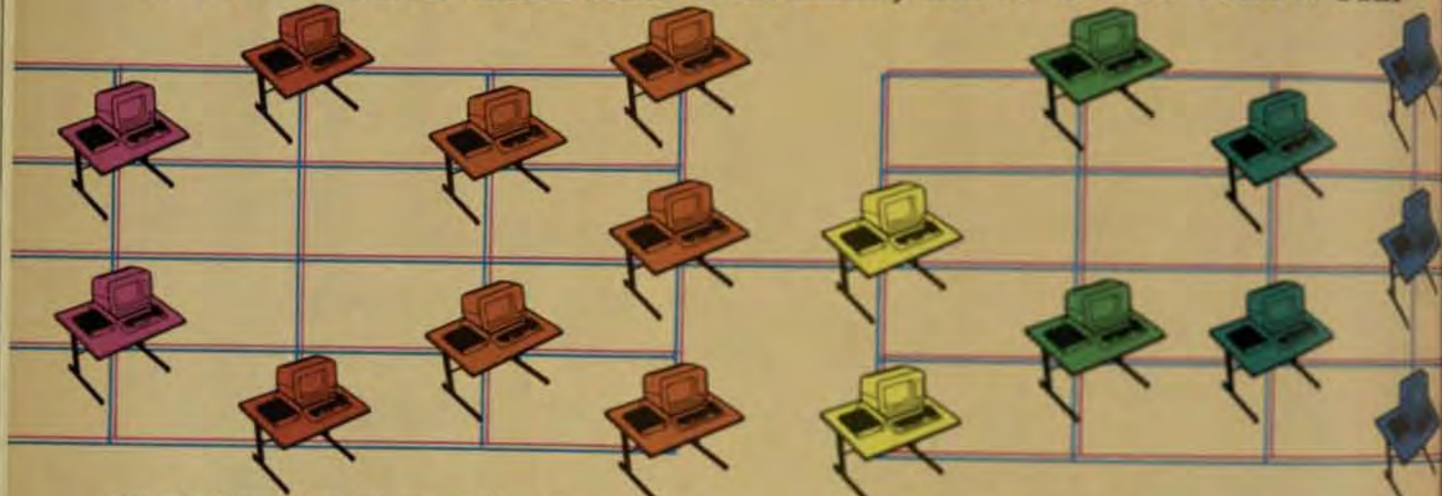

The **ZERO** is designed to bring high performance LOCAL AREA NETWORKING to users at budget prices.

The **ZERO** and **ZERO-NET** are unique. Any **ZERO** station can be a Network Master or Network Remote, permitting, for the first time, a low cost **non stop** network.

The **ZERO-NET** features a High Level Data Link Controller (message synchronous) at 400K bps carried over a simple twisted pair cable. To achieve maximum speed and reliability we use collision detection/avoidance circuitry and automatic CRC error detection/retransmission.

Each **ZERO** computer in the net can have Floppy Disk and/or Winchester Drives. The **ZERO** itself is a Z80 based Microcomputer with 64K Ram, 2K to 16K of EPROM, 2 Serial ports, 2 parallel ports, floppy disk controller, Z80CTC counter-timer and Z80DMA direct memory access.

The **ZERO** hardware design was optimized for TURBODOS*, (CP/M**, MP/M** compatible) including such enhancements as console type-ahead (buffering), 1.416 Mbytes per 8 inch double-sided floppy, multi-processing (background processing) such as print spooling, etc.

KEY PARAMETERS

- Local Area Networks up to 256 nodes per NET, with any mix of Master and Remote stations. Each station may support up to 16 logical drives, local or remote.
- Local Area Networks may be linked through gateways.
- Per Node — 0 to 2 floppies and 0 to 4 hard disks with appropriate Driver Modules.
- Per Node — parallel and/or serial printer.
- Each user can control print routing and/or spooling.
- Each node may reference a file system and/or printer on any other node.
- Each node may have an Autostart Log-on with security access protection.
- Each node may have a FIFO type Electronic Mailbox.

*TURBODOS is a trademark of Software 2000, Inc.
**CP/M and MP/M are trademarks of Digital Research, Inc.

- Each node may operate with MP/M compatible file/record interlocks, or with special TURBODOS relaxation rules.
- Maximum recommended buss length of 4,000 L.Ft.
- Full CP/M and MP/M compatibility.
- The TURBODOS operating system can support up to one GIGA Byte (1,000 Mega Bytes) per logical drive.

The **ZERO-NET** product family — the **ZERO**, the ZERO/FD, the **Z-DRIVE** and the **ZNT** terminal — all the components required to configure an entire system are available for immediate delivery. The **ZERO** provides a microcomputer network that OUTPERFORMS many multi-user minicomputer systems at substantially lower cost. While the competition is still studying it, **ALSPA** has done it!

ZERO-NET

Alspa Computer, Inc.

300 Harvey West Boulevard, Santa Cruz, CA 95060
(408) 429-6000 Telex 176279

Switched Data Systems Compete with Bus-Based Automated Office Links

Both baseband and switching technologies have reached a development level that allows informed analysis. The choice depends more on the application than on cost.

by **Larry Stolzenberg,**
Ericsson Inc

Digital baseband systems such as Ethernet may seem to be an obvious choice for office networks, but switched networks are growing in acceptance and must be considered as an alternative to bus-oriented technology. This remains a topic where there are more questions than answers—but the questions must be asked before a purchasing decision is made. In particular, such issues as voice capability, transmission speed, and cabling must be evaluated along with the system size and its expected growth and reconfiguration.

The private automated branch exchange (PABX) is an office telephone-switching system that has access to the public telecommunications network. The PABX allows calls to be made between local telephone extension and between the local phone and the public network, without intervention by the public operator.

Both local-area networks (LANs) and switching technologies encompass the requirements for voice, text, and data—perhaps even for video—in the workplace. Networking is the key to the computer revolution today because 80% of current data-communications tasks take place within the local environment. Particularly in the decentralized personal-computing environment, networking is valuable to prevent database division, which makes distributed computing inefficient and costly.

Until recently, a DP manager faced with the prospect of establishing a local-area network for his company might have migrated automatically to the system that seemed the most familiar—a baseband communication system. After all, baseband systems are digital, the DP man is a digital expert, and he appears to have a digital problem. Also, baseband systems are generally bus-oriented, a topology with which a person in the digital world is quite comfortable.

However, when you look at current events and future predictions in the office-network market, your decision may be different, as may his. The new factor to be considered is the growing importance of switched information systems as networks. Such networks have gained credibility from a number of vendor entries into this segment of the market. IBM's agreement with Mitel to exchange technology, and its recent partial purchase of Rolm Corp., endorsed the digital PABX as a principal communication method for future office networks. Digital Equipment Corp., one of Ethernet's creators, has a program to use PABXs to interconnect its own all-in-one integrated office-automation system. Also, Data General Corp. has a program that may lead to its equipment being interconnected via a computer-controlled switched system for business communications.

While there's no way to predict how these programs will evolve, the fact that they exist at all is an important sign that no major decisions about local-network installation should be made without serious consideration of switching systems.

Although the basis of PABX technology existed for many years, it received a major impetus for development from the Supreme Court's 1968 Carterfone decision, which allowed customer-owned equipment to be interconnected to the public telephone network. This prompted many manufacturers to enter the PABX market.

THE OPTIONS—BROADBAND

Broadband, frequently called cable from its use of coaxial cable as a transmission medium, employs frequency-division multiplexing (FDM) to divide a single physical channel into a number of frequency channels. Those channels can be assigned different bandwidths, which gives broadband one of its greatest potential benefits—the ability to transmit different forms of information (data, video, and eventually voice) on a single system.

The flexibility of band assignment and multiplexing of subchannels ultimately will make it possible for different types of devices—operating at different speeds, with various electrical channels, and in many modes—to function on the broadband network. But while these devices already may be operating on the broadband network at the same time, they have yet to be truly integrated because the equipment is still functioning on different physical networks within the single cable. Therefore, communication among the participating end points (branch junctions) of a network must still be accommodated. And that requires a complex system of costly logic, conversion, and switching capabilities.

The principal drawbacks of broadband LANs at present are cost, complexity, and the inability of today's systems to accommodate voice transmission.

Because broadband signals are limited to travel in one direction, 2-way exchange requires a single cable to be either divided or looped to create a dual-cable system.

Divided-cable systems require a central retransmission facility (CRF)—a costly front-end investment ($3000 to $4000 average) and must be centrally positioned to maintain optimum transmission and error rates. This positioning requirement affects the configuration flexibility of the LAN, and the addition of new connections to the system may require expensive and inconvenient repositioning of the CRF. Dual-cable systems require no CRF, but their cable costs are higher, and so are the costs of taps into the double cable for connections.

Broadband signal transmission is analog and its communications-interface unit can be a variety of radio-frequency (RF) modems. For low-to-medium speed (10- to 25-MHz) point-to-point connection on the LAN, inexpensive fixed-frequency modems can be used. The 55- to 75-MHz band, however, is made up of 128 switched channels and therefore requires variable-frequency (frequency-agile) modems. The initial cost of such modems, plus the cost and availability of switching controllers for them, will continue to limit broadband acceptance.

Another problem is that standardization of frequency allocation isn't widespread among broadband systems. For example, some systems use the IEEE-reserved television bands for data, thus limiting television transmission on those networks. Some groups of vendors are asking IEEE to standardize access techniques, but no decision is likely to be made in the near future. Also, while broadband has the potential for voice transmission through a dedicated channel, current technology doesn't make this capability available to the user.

All told then, the current limitations of broadband systems are sufficient to deny them serious consideration by today's network buyers.

THE OPTIONS—BASEBAND

Unlike broadband, where the physical channel is divided into many frequency channels, baseband LANs transmit unmodulated signals on a single channel. The transmission medium is usually coaxial cable, with data rates in the range of 10 Mbps. Since baseband systems were designed for data transmission, the signals and technology are digital.

Topologies used in baseband systems are most frequently ring and bus types (Figs 1 and 2). These provide distributed control rather than the centralized control of the star topology used in many PABXs (Fig 3). In ring networks, nodes are connected by point-to-point lines and each node accepts or retransmits messages on the network. With bus networks, on the other hand, nodes are connected to a single physical channel by cable taps or connectors.

Nodes must be able to recognize their own messages, but don't have to retransmit messages intended for other nodes. Therefore, node failure doesn't affect the status of the total bus. Bus networks are the most familiar baseband networks, with Ethernet being a popular example.

Baseband systems transmit messages by packet switching, in which whole messages are divided into discrete units of data. A packet of information of a specific length contains the address of its ultimate destination. This packet is placed on the single channel and is ultimately accepted by the correct node.

Due to this single-channel configuration, baseband networks require

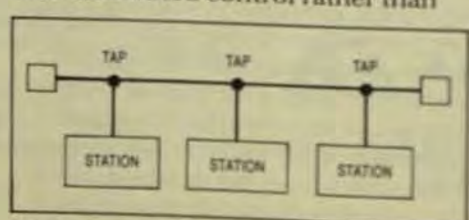

Fig 1 A ring network is a point-to-point link between nodes; messages are repeated by each node.

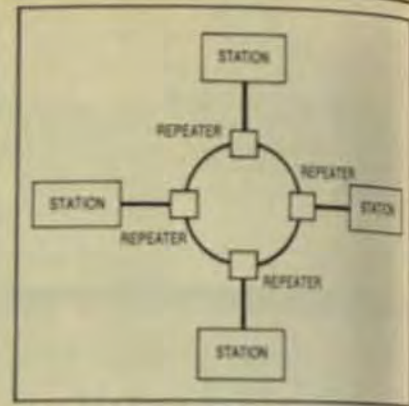

Fig 2 A bus network is a single line shared by a number of nodes that do not repeat messages passed on the bus.

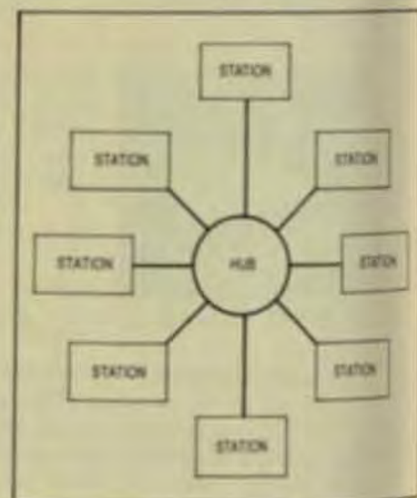

Fig 3 A traditional star topology has network nodes joined at a single point.

IBM

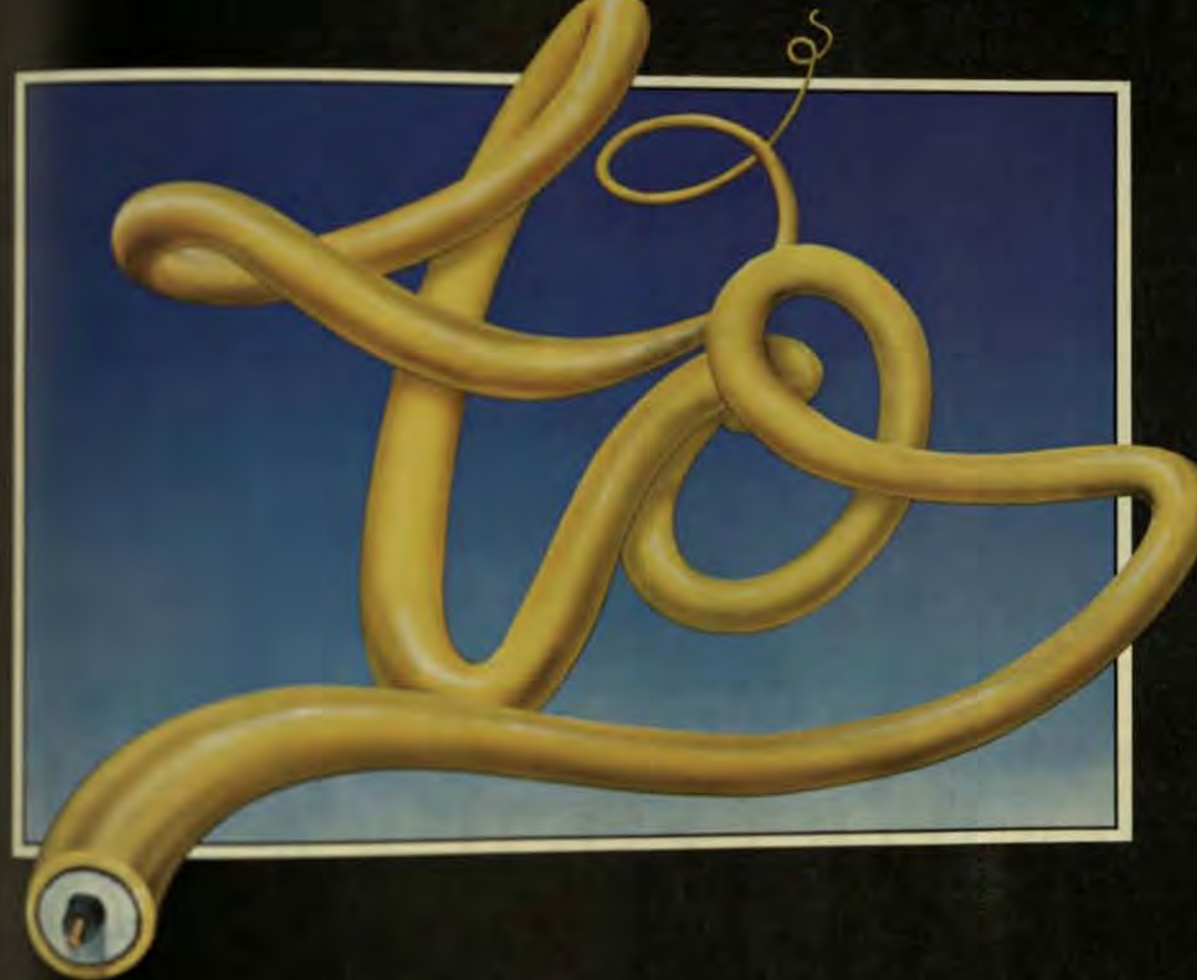

Ethernet

and other local area networks

AUSCOM introduces the proven solution for connecting your IBM mainframes to Local Area Networks. Using the 8911 Programmable IBM Channel Interface, AUSCOM can connect your network directly to the channel and, through software, provide the necessary drivers to incorporate the LAN protocol of your choice. You select the LAN; AUSCOM will make the connection to your IBM!

by AUSCOM, INC.

the "IBM to Anything" people

2007 Kramer Lane • Austin, Texas 78758 • 512/836-8080

access techniques that allow more than one message at a time to be passed by the system without collision. The two major types of access are polling and contention. Polling determines the order in which nodes can access the network, while avoiding the possibility of collision. The most popular polling method is token passing, used on ring systems. With this method, a bit pattern is circulated and messages can be dropped into it.

The most familiar type of access, however, is the contention technique—CSMA/CD (carrier-sense, multiple-access/collision-detection), which is used on Ethernet systems and has been proposed as an IEEE standard. CSMA/CD lets any node access the channel upon sensing that it's free; thus no waiting for a token is required. If, due to signal-propagation delay on the channel, two nodes have sensed a free channel and transmitted simultaneously, the collision is detected during transmission. In such cases, both transmissions stop and are retried at either a fixed or random interval.

THE OPTIONS—SWITCHING SYSTEMS

All of the switching systems in use in work environments, from very early PABX systems to the all-digital networks being developed today, use some variation on a star topology. Rather than try to embrace all variations, this article will outline a model for contemporary switching systems. The MD110 from Ericsson has been used as an example because it is both state-of-the-art in design, and has an installed base. It is, therefore, both a new system and a real one.

The MD110 is a digital voice-and-data switching system capable of supporting over 20,000 connections. Data is integrated with voice over a single pair of wires. Data and voice can be routed separately. A highly modular design permits easy geographical dispersion of network units. The system employs two active units, a line-interface module (LIM) and a group switch (GS).

The LIM is a nonblocking, microprocessor-controlled unit that can be equipped with any combination of line circuits, trunk circuits, and terminal devices. An autonomous switching system, it can function independently or as an integrated part of a larger system, and can handle as many as 200 devices.

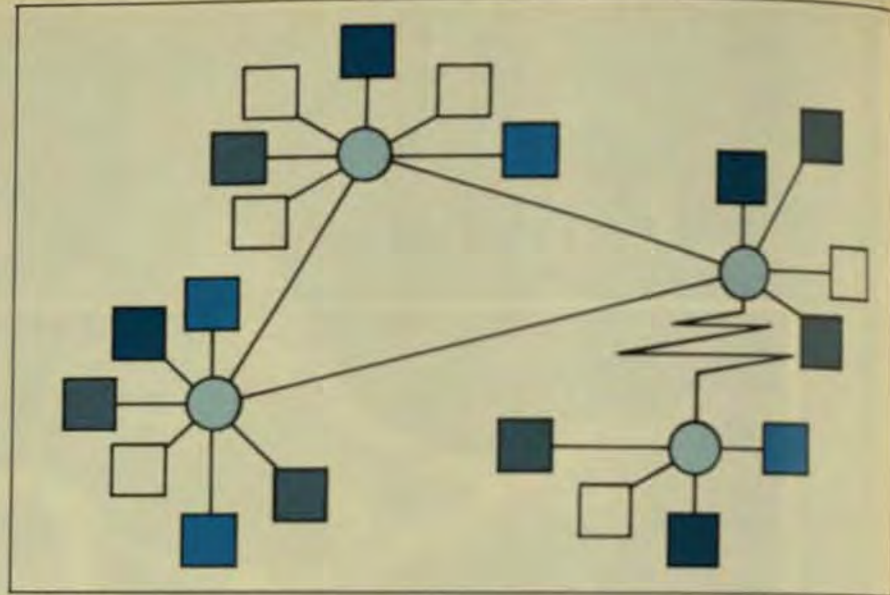

Fig 4 In a decentralized microprocessor-based switching system, stars are connected into "constellations."

Two LIMs can be connected directly via 32-channel PCM lines; larger systems employ the group switch—a modularly expandable, nonblocking digital switch that transmits PCM voice, data, and control signals among units. The GS is a passive device controlled by the connected LIMs.

The network topology resulting from this modular system can be characterized as a constellation—a system of stars (Fig 4). It offers the benefits of switch technology and the reliability usually associated with bus-oriented topology. When distances between the LIM and GS exceed 400 m, the 2-Mbps transmission system uses a transmission medium that may be normal twisted-pair wire, fiber-optic cable, or microwaves. Alternatively, 2-Mbps transmission may be converted for conventional 1.5-Mbps T1 lines when public lines are required.

All system facilities can be accessed by use of standard telephone receivers or digital telephones that can also serve as connection points for data terminals.

Many of the newest switching-system designs include advanced functions for system supervision, alarm sending, fault location, and traffic recording.

THE VOICE QUESTION

The need for voice communication in the local work environment (80% of local telecommunications remains voice) is probably the single greatest factor in the selection of switched information systems for networking. Regardless of the many features baseband systems have, single-channel multi-access techniques can't handle the lengthy transmissions and long delays of voice communication. Attempts to accommodate voice on baseband systems create unrealistic demands for bandwidth. In general terms, switched systems can effectively accommodate most data communication, and baseband can't efficiently transmit voice.

Therefore, if a baseband system is selected it must be in addition to or in combination with, a voice-optimized network. The principal deciding factor for many companies will, of course, be the cost of separate systems. For an environment requiring sophisticated voice-communication features as well as data transmission, the cost of a dedicated data network may well be impractical when the digital switched system can handle both requirements with little or no compromise in most environments.

A small company, housed on a single floor, with a frequent data-transmission requirement, may select a low-cost baseband LAN if the firm's current level of voice communications is adequate and no need for growth is anticipated. In such cases

The Digital DFM Intelligent Communications Processor*

Introducing computer status for the stat mux.

It's priced* like a stat mux, but that's where comparisons end.

Our new Intelligent Communications Processor is powerful enough to handle network control and management, data control and concentration, and channel switching and contention. All automatically.

Digital's unique Intelligent Communications Processor lets users upgrade from 4 to 16 channels by simply changing modular boards.

Multiple microprocessors make it possible and nobody else makes one like it.

Of course, it's all the stat mux you'll ever need, too. With synchronous and asynchronous support. And internal DMA.

Plus, its modular design makes it easy to upgrade. By you. In the field. Without recabling. Go from 4 to 16 channels—with and without 4.8K and 9.6K bps modems—by simply sliding in new boards.

Our powerful new Intelligent Communications Processor. Sold and serviced by Digital from one end of your network to the other.

It's the first data communications link that's just as smart as a computer, because it is a computer.

For more detailed information on the Intelligent Communications Processor and the complete Digital line of data communications products, call toll free: **1-800-DIGITAL**, Ext. 610, and request Digital's "Guide To Data Communications." When calling please reference code TI.

*Four-channel prices start at just $2,675. U.S. Dollars only.
© Digital Equipment Corporation 1983

digital

a baseband system may be more cost-effective because of its relatively low installation cost in small configurations.

TRANSMISSION SPEED

Baseband networks using coaxial cable are designed specifically for data transmission. The capacity of the cable is used to transmit one baseband signal at a high data rate, generally in the range of 10 Mbps. (This is in contrast to broadband, which divides the capacity of the coax into many lower-level TV channels or low-speed data channels.)

Switched systems are based primarily on twisted-pair wire with a data-transmission speed of 56 to 64 Kbps. Conditioned twisted-pair wires can be driven at rates up to 2 Mbps.

Therefore, if a user requires frequent transmissions over 64 Kbps, as is sometimes necessary for communication between two or more large computers, some technique other than a standard twisted-pair-based network will be required. The principal option will likely be a dedicated baseband cable that may be used either in a network context or independently.

Actually, such considerations are fairly infrequent. More than 90% of data communications in the local workplace occurs at data rates of less than 56 Kbps. The most frequent rates are in fact less than 4.8 Kbps. Therefore, speed considerations are not primary in most networking decisions.

CABLING

A frequent argument over whether to choose baseband or broadband modulation is centered on the presence of existing twisted-pair wiring (used in telephone systems) in virtually all work environments. While it's true that some existing wiring can sometimes be used in new building installations, it's not a major factor. A data-communications system employing twisted pair will usually require new wiring if only because existing wiring plans for a building may not exist. Therefore, the saving on installed wiring is not as great as it's sometimes represented to be.

The initial cost of twisted-pair wiring, however, is lower than for any other medium. Its installation costs are also low, since commonly available tools and less technical labor can be used.

The benefits of coaxial cable are well known. The widely used medium offers large bandwidth and the ability to support high data rates with low error rates and immunity to electrical interference. Cable is employed in telephone networks to multiplex many calls onto one cable and is best known for its use in CATV. This widespread use makes cable moderate in cost (though more expensive than twisted pair) and easily available. The installation techniques are substantially more complex than twisted pair's—requiring taps, controller, splitter, couplers, and repeaters—but the technology is well established.

Fiber optics, the medium most likely to provide ultrahigh bandwidth performance in the future, currently lacks a way to tap the medium efficiently and inexpensively. Therefore, any use of optical fiber in the near term will be limited to point-to-point lines within a larger network or in ring networks. Fiber is currently finding applications in hostile environments because of its immunity to radiation.

While the buyer's final network decision won't be based on the type of cabling used, the decision will be greatly affected by the cost, flexibility, and bandwidth considerations of the medium.

SYSTEM SIZE, GROWTH, AND RECONFIGURATION

Often the key issue in the manager's selection of a network will be based on the system size and configuration at the time of installation, and the user's growth requirements.

The baseband bus topology is inherently limited in its number of users. A typical system such as Ethernet can extend just under two miles (within a building) due to limitations in signal propagation. Within that distance, repeaters must be placed at distances of approximately 1000 ft to amplify the signal.

Taps into the bus for individual nodes can be made only at specified intervals, allowing a maximum of 1000 users. This is a sufficient number for many applications, but it pales beside the 20,000 to 30,000 possible users who can be supported by a digital switched system. Therefore, if the user anticipates any growth over 1000 users, he must look beyond a single baseband network.

Equally important is the configuration of the required system, not only at the time of installation but in the future. The expansion of baseband networks beyond a single floor of a building is very costly due to the need for an expensive bridge connection. Extension of the system across public streets is limited by public telephone regulations. For many companies, this may be a serious limitation on future growth. Switched local systems, of course, like telephones, have virtually unlimited flexibility for configuration between floors and buildings.

The ease with which individual nodes can be moved is also a factor. Baseband terminal connections are complex and as inflexible as the coaxial cable to which they are attached. The placement of users at the time of installation is somewhat limited in range, and the moving of a user is costly and difficult. A switched connection can be moved as easily as a phone can be unplugged.

While a system with a star topology was once considered much less flexible in configuration than a linear bus, today's modular stars are extremely easy to configure to user requirements. There's little argument that in this most important area of system growth potential, the switched systems have significant advantages.

RELIABILITY AND MAINTENANCE

A frequent criticism leveled against switching systems of all kinds is that the star topology creates a central point of failure (the control node) from which the entire system can be brought down. This vulnerability is often true for token-passing ring networks and in nodes that repeat messages to the next node in the ring. Failure of any node, or any other break in the ring, will generally cause the network to cease to function.

Bus networks, on the other hand, are inherently resistant to central-point failure. Nodes share a single physical channel by means of taps or connectors and, therefore, perform a message-repeat function. The failure of a single node has no impact on the operation of other nodes.

To overcome this traditional problem with star topology, the digital information-switching system eliminates single-point vulnerability with a highly modular and decentralized with processor control distributed to the board level, eliminating central computers. Line-interface modules function independently of one another. Within a LIM, redundancy of the control unit, time switch and other units eliminates the possibility of a major failure.

Today's switching systems are equipped with remote and self diagnostics and can generally be diagnosed, serviced, and maintained over phone lines from centralized diagnostic centers.

COMPATIBILITY

Among baseband systems, Ethernet, the combined product of Digital, Xerox and Intel, has become an unofficial industry standard. Ethernet conforms to the International Standards Organization's Open System Interconnect (ISO OSI) reference model for layered network architecture. That model calls for physical link, data link, network control, transport, session control, presentation control, and application/user layers.

This architecture also allows standard interfaces and protocols (such as RS-232 and X.25) to be incorporated into the layers with relative ease. However, if the much heralded IBM system does, in fact, employ a token-passing ring topology, this standard may quickly be supplanted.

While switched information systems also employ such communications standards as RS-232 and X.25, no single vendor has emerged as the industry leader and, therefore, no de facto standard has emerged. Different manufacturers have taken unique approaches to system architecture and signal processing. A strong European standard for switched systems—the Integrated Systems Digital Network of the CCITT, which applies to telephone central-office systems—has been proposed for adaptation to PBX-type systems. Such applications could provide the basis for a U.S. standard.

HUMAN FACTORS

The human element is a subtle but very important factor in many network decisions being made today. After all, networks are designed to increase the productivity and contribution of people. In today's office, networks should be usable by virtually anyone in the local environment—managers, engineers, clerical and secretarial personnel, sales staff, and scientists. To achieve maximum use with minimum training, the system must be considered friendly, and in this respect, few devices can surpass the telephone.

While CRT terminals, printers, and other computer-oriented devices may be every bit as easy to operate as some of today's sophisticated digital telephone handsets, the user may not think so. The phone is a totally familiar object. Any network, no matter how vast, that's even partially accessible by phone is apt to be seen as a normal part of the work environment and not as an exotic system for use only by computer scientists. Therefore, the installation of a switched network results in a minimum of disruption in the work environment.

COSTS

The cost-per-user quoted by vendors of both baseband and switched systems is approximately the same, between $1100 and $1500. A data-only switching system will cost somewhat less than a data-only baseband, while a voice-and-data switching system is priced slightly more than a data-only baseband.

The cost-per-user, however, doesn't vary enough to be the significant consideration. The real factors in cost comparison are user-dependent. Will the user have to duplicate systems and services to use a baseband system? How many dedicated baseband lines might be required? What is the anticipated growth of the system? (If today's system will be obsolete in one or two years, it's clearly no bargain.) How often will users change and the system be reconfigured? And, of course, which system has the short- and long-range performance features required by the user? It doesn't matter what a system costs if it won't do the job.

The intention of this article is to leave the manager with more questions than answers. Before the buyer can begin an analysis of vendors, however, he must do a thorough examination of his own company and application. What are the requirements of his business in terms of number and kinds of users, volume of data and voice traffic, types of services needed, etc.? He must examine the physical configuration of the office(s) or plant—how many buildings, floors, etc.? And, most important, he must look at where the business is going and how its needs will grow. If the manager can't reasonably project requirements for at least two years, he's not yet prepared to select a network.

Once this in-house analysis has been completed, the list of considerations and issues given can be used to evaluate prospective suppliers. The manager may seek a single source for all network requirements or prefer to separate voice and data systems between manufacturers.

No matter how challenging the user's system demands may be, he'll find 1983 an excellent year to make a decision. Both baseboard and switching-system technologies will continue to co-exist, often within the same office environment, and future developments will extend—not obsolete—current capabilities.

The era of local-area networking has arrived. The buyer has only to find the system and supplier that will carry him into the future. ■

Larry Stolzenberg *is product manager for MD110 and Supporting Systems at Anaconda-Ericsson Inc. He holds a BEEE from City College of New York and an MBA from Rensselaer Polytechnic Institute.*

Considerations in Selection of a Local-Area Network

The major considerations in selection of a local-area network for today's office environment are:

- Will the considered network accommodate the required types of communications in a single system, or will costly duplication of systems and services be required?
- Can the system provide the bandwidth required for the type of communications to be performed?
- Is the system affordable?
- What are the system's cabling requirements? Can any present wiring be used, or must totally new cabling be installed?
- Can the system accommodate the number of users anticipated in both the present and future?
- Is the system reliable, and does it minimize the effect of major failure?
- Is reconfiguration of the system efficient and affordable? Does it provide flexibility for growth and cost-effective maintenance and operation?
- What degree of compatiblity with other systems can be maintained?
- What are the operator factors? Can users be easily trained? Will they find the system easy to use? What level of acceptance will the system find among users?

Network Vocabulary

Amplitude Modulation. The amplitude of the carrier signal is altered with time to conform to the message signal.

Analog. Wavelike signals (cycles) proportionate to an original form of communication. The number of signals to occur in one second is called frequency (measured in hertz).

Baseband Signal. An analog signal sent at its original frequency.

Bus Network. A single line shared by a number of nodes.

Capacity (Bandwidth). The range of analog signals that can be carried by a physical medium. Bandwidth also describes the difference between the highest and lowest frequencies of a signal.

Carrier Signal. A continuous signal transmitted over a medium onto which an original frequency signal can be modulated to be sent at a different frequency.

Centralized Control. Access to the network and channel allocation is controlled by one node.

Centralized Polling. A method of access to the network channel wherein a master node queries each node in turn.

Coaxial Cable. A central carrier wire surrounded by fine copper-wire mesh and/or an extruded aluminum sleeve. Offers large bandwidth and supports high data rates.

CSMA/CD. Carrier sense, multiple access with collision detection. A channel-access technique that allows any node to send a message when it senses the channel is free. If two nodes transmit at the same time (possibly due to signal delay), the resulting collision is detected during transmission, and the transmission is abandoned. After a brief wait for a fixed or random interval, a retry is made.

Digital. Discrete signals (bits) represent two possible states (ON or OFF).

Distributed control. More than one node (or all) can establish connection and access the channel independently according to accepted rules.

Distributed Polling. Access control to the network is located within each node (see **Token Passing, Slotted Rings**).

Fiber Optics. Very high performance transmission medium. Fibers made of plastic or glass carry electrical signals that are translated into light pulses.

Frequency-Division Multiplexing. Divides the available bandwidth of a single physical channel into smaller, separate frequency channels. Analog only.

Frequency Modulation. The frequency of a carrier signal is altered to conform to the message signal.

Link. Communications path between two nodes.

Multiplexing. The method whereby a single channel is divided into multichannels for transmission of a number of independent signals.

Multipoint Line. A circuit shared by more than two nodes.

Node. An end point of any network branch, or a junction of two or more branches.

Point-to-Point Line. A circuit that connects two nodes.

Ring Networks. Point-to-point link nodes form an unbroken circle.

Route. The ability of a node to pass messages to an adjacent node.

Slotted Rings. Noncontention access control. A number of frames of fixed size circulate on the ring network, each containing bits for addresses, control, parity, and data. Nodes insert data into unused fields.

Space-Division Multiplexing. Creation of a single channel by the grouping of many physical channels.

Star Networks. Network nodes are joined at a single point (frequently seen in switched information systems).

Time-Division Multiplexing. Allows each node of a network a small time interval in which to transmit information. Digital.

Token Passing. Noncontention access control: a bit pattern circulated on the network is "grabbed" by a node for exclusive access to the channel.

Twisted-Pair Wire. Generally, copper wire twisted to avoid interference in multipass cables. The least expensive and most widely used medium for voice and data transmission.

Unconstrained Topology. A network that takes the shape of the connections as dictated by the application. ■

YOUR LOCAL COMMUNICATION NETWORK HAS TO WORK 24 HOURS A DAY.

SHOULDN'T IT HAVE A BACKUP OPTION?

Your local communication network has to work 24 hours a day, so we designed the IDX-3000 digital data switch to be fault tolerant: if a computer resource fails, the inherently reliable IDX-3000 keeps the network running. We also offer optional redundancy features throughout its distributed architecture: at the control program, Network Exchange, and multiplexers.

For backup and load sharing, you can run as many as 16 system control programs simultaneously, and you can install up to four control modules in your Network Exchange to handle network control functions. Extra switch modules can be installed, with automatic switch-over when a failure is detected. Continuous fault monitoring detects potential problems early, and continuous self-testing isolates problems to the board level. With the IDX-3000, you can have as much redundancy as you need—even out to the distributed multiplexers.

For the office of now

Simple installation, distributed architecture, and practical, fault-tolerant design keep costs down. Expandability to more than 3000 lines means the system can grow with your needs. The IDX-3000 isn't one of those products of the future. It's here now—proven, in use, and backed by M/A-COM, one of the nation's largest communication companies.

If you have or need more than 200 terminals, contact us. We'll give you the full story. And we can back it up.

THE NO-RISK NETWORK

M/A-COM

M/A-COM LINKABIT, INC.
Attn: Ruth Stoffel
3033 Science Park Road
San Diego, CA 92121
Toll Free (800) 626-6640
(619) 457-2340

© 1982 M/A-COM Linkabit, Inc.

"What good is 9600 baud if your stat mux slows you down?"

Speed can be addictive, that's a fact. Once users get accustomed to fast feedback they just don't want to hear that your lines are overloaded or your mux just can't keep pace. They want their data and they want it now.

To help you avoid these baud rate bottlenecks, Teltone is happy to introduce our new M-860 Statistical Multiplexer.

You want speed? It's a thoroughbred at up to 78,600 bps, so you can take full advantage of 56K bps DDS circuits. You need flexibility? It's a chameleon that offers six different data link interfaces—two with integral modems.

Efficiency? It's an unparalleled performer that delivers multiplexing efficiencies of 800% and more. Friendliness? This is your mux—fully programmable from the local or remote end with an interactive menu system that also provides access to diagnostics and statistics. You'll never deal with DIP switches again. The M-860 is compatible with our own DCS-2 Data Carrier system and voice-grade lines, but it works equally well over microwave, coaxial cable and fiber-optic hookups.

Want more? How about local or remote echoplex, full field expansion, CRC error detection and ARQ retransmission, speed conversion—plus the same dependable service we've been giving the telecommunications industry for more than 15 years.

It's a combination that's really hard to beat, in a mux for people who'd rather be working than waiting. Call 1-800-227-3800 Ext. 2301 or write Teltone Corp., PO Box 657, Kirkland, WA 98033. In Canada call (416) 475-0837 or write 183 Amber Street, Markham, Ontario L3R 3BR.

TELTONE

Intelligent Hardware Meets the Demand for Network Performance

Programmable front-end processors and modular statistical multiplexers offer alternatives for easy, cost-effective upgrades of communications systems.

by **Dave Korf**
and **Maureen LaGro**,
Digital Equipment Corp.

Today's intelligent communications processors (ICPs) and multiplexers offer a variety of options to enhance existing networks and realize a fast return on investment. ICPs can handle the computer processing needed for communications between CPUs and public packet-switched networks (PPSNs) and between CPUs and other network architectures. Not only can ICPs ensure the sequencing and integrity of data transmission between networks via the X.25 protocol, they offer extensive user-programmability so that unique non-X.25 features can be incorporated. Statistical multiplexers also help provide cost-effective communications and system upgrades.

PUBLIC PACKET-SWITCHED NETWORKS

Computers may be linked through a PPSN based on the International Telegraph and Telephone Consultative Committee's (CCITT) recommended X.25 protocol. This protocol is a set of rules that ensures the correct sequencing and integrity of transmitted data. X.25 networks such as Telenet in the U.S. and Datapac in Canada vary in detail from country to country; each requires certification of X.25 interfaces that link to the network.

Computer manufacturers customarily offer software packages that enable their computers to implement the X.25 network protocol. Any computer operating with X.25 capability may be linked directly to a PPSN and communicate with other linked X.25-compatible computers (Fig 1). Public or private communications links may be telephone lines, satellites, on-ground microwave links, or any combination of these.

In addition to X.25 PPSNs, there are other network architectures—such as Digital Equipment Corp.'s Digital Network Architecture (DNA) and IBM's Systems Network Architecture (SNA)—that have been implemented for the manufacturer's own operating systems and hardware, which define specific message protocols for data transmission between nodes in their networks.

New hardware and software allow computers in one private network to communicate with computers in another private network, either via direct connection between networks or by linking the networks through a PPSN.

X.25 COMMUNICATION

A communications processor is a front-end device that performs communications functions within a single computer system or in a computer network. What these functions are depends on the configuration or protocol requirements of the particular system or network. X.25 defines three levels of protocol: Physical, Link, and Packet. The International Standards Organization (ISO) Open Systems Interconnect (OSI) model describes seven communications levels for network architecture.

The Physical Level (Level 1) defines both the mechanical and electrical functional and procedural characteristics of the physical interface between the Data Terminal Equipment (DTE, an intelligent system receiving or transmitting data) and the Data Communications Equipment (DCE, an intelligent system that serves as the DTE's access point to a PPSN). In Fig 1, the DTE is the CPU at the interfacing node and the DCE is the packet-mode facility at the local-network exchange.

The Frame Level (Level 2) defines the link-access procedure for error-free data exchange over the physical link between a DTE and DCE (Fig 2). In this format, a header precedes the data field to specify the channel being used and the sequencing of the packets being sent. Each packet is framed with the information needed for error-free transmission:

- **Flag.** All frames start and end with a flag to indicate the beginning and end of a frame.
- **Address.** This indicates whether the frames are created by the user (DTE) or by the PPSN (DCE).
- **Control.** Control functions ensure that data packets are sent, received, and re-assembled in the correct sequence. They also acknowledge receipt of data or request retransmission, and connect and disconnect the link.
- **Frame-checking sequence.** A check is performed on all the bits (other than flags) to ensure error-free transmission.

The amount of data contained in a packet varies among networks; a maximum of 1024 bits is typical.

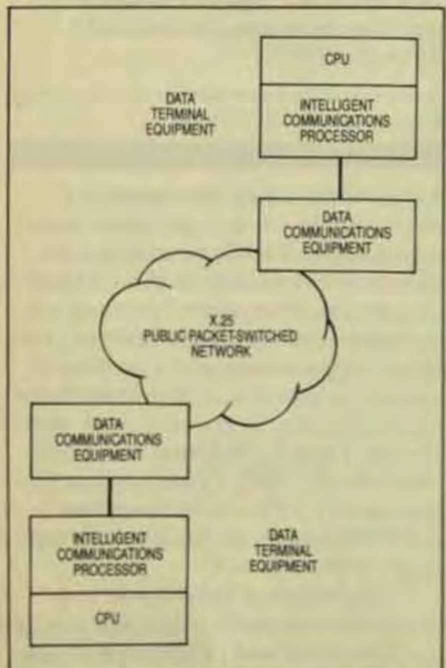

Fig 1 Computers connected through an X.25 need not be from the same manufacturer because their communications are predefined by their X.25 capability.

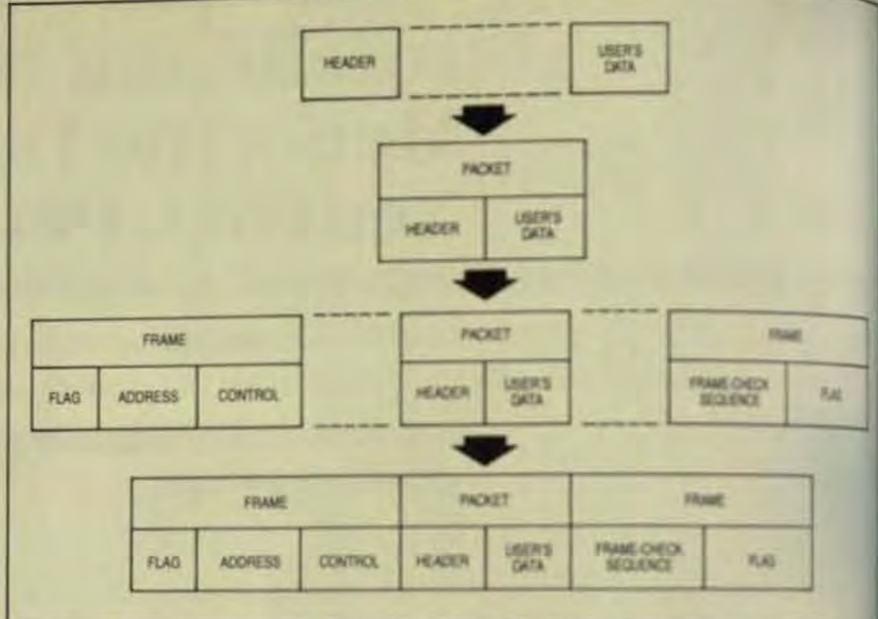

Fig 2 A packet as defined by X.25 consists of a header and data framed with a variety of information to ensure error-free transmission.

The Packet Level (Level 3) defines packet-format and -control procedures for exchange of packets across a network. Control packets precede data packets across the network to set up the virtual circuit. After all data packets in a message have been sent, control packets are again transmitted to conclude the circuit.

The X.25 communications functions reside in the four upper layers of the ISO network protocol and are executed by the CPU at the DTE node. These functions are implemented quite differently in custom and proprietary networks. Software may be rewritten in-house by the network owner, computer manufacturer, or PPSN developers. Communication is transparent to application programs running in a network whose architecture has been fully defined and implemented in software at all protocol levels by the computer manufacturer, but a user who writes his own software must take full responsibility for proper implementation.

INTELLIGENT COMMUNICATIONS PROCESSORS

The most common processors for communications interfaces are nonintelligent single- or multiline synchronous devices that perform only the Level 1 functions and operate at maximum line speeds of 19.2 Kbytes. If the local CPU has excess processing capacity available to perform the Level 2 and 3 functions, such synchronous interfaces are not only more than adequate, but the least costly design.

In new and expanded networks these interfaces are being replaced with intelligent communications processors (ICPs) (Fig 3) that have their own 16-bit microprocessor, ROM and RAM.

Today's ICPs vary in the number of levels/functions they can perform, the line speeds at which they can operate, and the amount of on-board ROM or RAM. Most ICPs designed for X.25 compatibility perform Level 1 and 2 functions, and the most advanced models even handle Level 3. These same ICPs are typically designed for line speeds of up to 56 Kbaud (64 Kbaud in Europe), with the more advanced models handling up to 256 Kbaud or even 1000 Kbaud (1 Mbaud).

Based on the desired performance level and degree of program-

Fig 3 An intelligent communications processor includes a built-in microprocessor, ROM and RAM. It performs X.25 Level 2 and Level 3 functions (link and packet), off-loading communications overhead and freeing the CPU for other processes.

GET MORE MUX FOR LESS BUCKS.

The way we figure it, you should get a lot more out of a low-cost stat mux than just low cost. That's why we put so much more *into* our Model 5300 statistical time division multiplexer.

In fact, you're getting a lot of high-end features at low-end prices. Features like 4 to 16 channel capability in a single enclosure. Selectable local/remote diagnostics. ARQ error correction. Asynchronous port speeds of up to 9600 bps. And autobaud.

And the price? Just $1,630 for the 8 channel version. Only $2,353 for full 16 channel capability. Best of all, if you order now, you can take another $50.00 off the purchase price. (This offer extended til December 31, 1983). When you consider the price/performance advantages, that's a deal you really can't afford to pass up.

For more information just call us at (408) 743-4300 and ask for Communications Products. Or send in one of the coupons below.

SEND ME A MUX.

I'm sold. Please . . .
☐ Send me ____ Model 5300 stat mux(s) with ____ channels right now. Enclosed is a check for $______.*
☐ Put some on the truck and bill me at the following address.

NAME ______ TITLE ______
COMPANY ______
ADDRESS ______
CITY ______ STATE ______ ZIP ______
PHONE ______

*F.O.B. Sunnyvale, Calif.

SAVE ME SOME BUCKS.

I'm interested. Please . . .
☐ Send me complete product and pricing information.
☐ Have a sales rep call me about your introductory offer.
☐ Send me at least two.* I can't resist the offer.

NAME ______ TITLE ______
COMPANY ______
ADDRESS ______
CITY ______ STATE ______ ZIP ______
PHONE ______

*F.O.B. Sunnyvale, Calif.

OFFER GOOD FOR A LIMITED TIME ONLY.

Intersil Systems
1275 Hammerwood Avenue
Sunnyvale, CA 94089
(408) 743-4300

mability, the user can select various amounts of ROM or RAM. For example, Digital's KMS11-P performs Level 1 and 2 tasks for one synchronous line and has 4K words of data and instruction RAM. Higher performance devices may have 32K words of ROM to store X.25 Level 2 and 3 programs, and substantially more RAM for high-speed line buffering.

Firms that base their networks on X.25 PPSNs generally prefer to purchase Level 2/Level 3 software with the ICP hardware, as part of a specific PPSN-oriented package. Since the X.25 network link must be certified, this option is simple, fast, and more cost-effective. On the other hand, privately owned and operated networks may have communications protocols that, for the operator's own reasons, have unique non-X.25 features. In those cases, custom Level 3 software must be written internally or contracted for by the network operator. The extensive user programmability in the newer ICPs provides this option. Additional software modifications can easily be made as needs change after implementation.

Overall network efficiency isn't affected by X.25 software being executed in the CPU or ICP, but the execution of communication processing within the CPU wastes valuable CPU time and requires costly line hookups. The servicing of a 9.6-Kbaud synchronous link through a basic Level 1 communications interface can consume up to 10% of available CPU time, although this time can be reduced to 2% if the service is linked through a Level 2 ICP. Because most ICPs are direct-memory-access devices, they reduce CPU overhead by accessing the host's main memory without CPU intervention except at the beginning and end of the message.

Selective upgrading of ICPs as new models become available may well be the most economical way to increase data-processing power within a network. Before adding a powerful computer, you should consider the ICP. After all, return on the investment can be determined best in terms of the additional CPU time cycle made available for application programs. Besides lowering the cost of new processing capacity, ICPs require less space in the computer room and, with the newer models, less space inside the mainframe, too. Many ICPs are single-board modules that insert into the CPU's backplane slots.

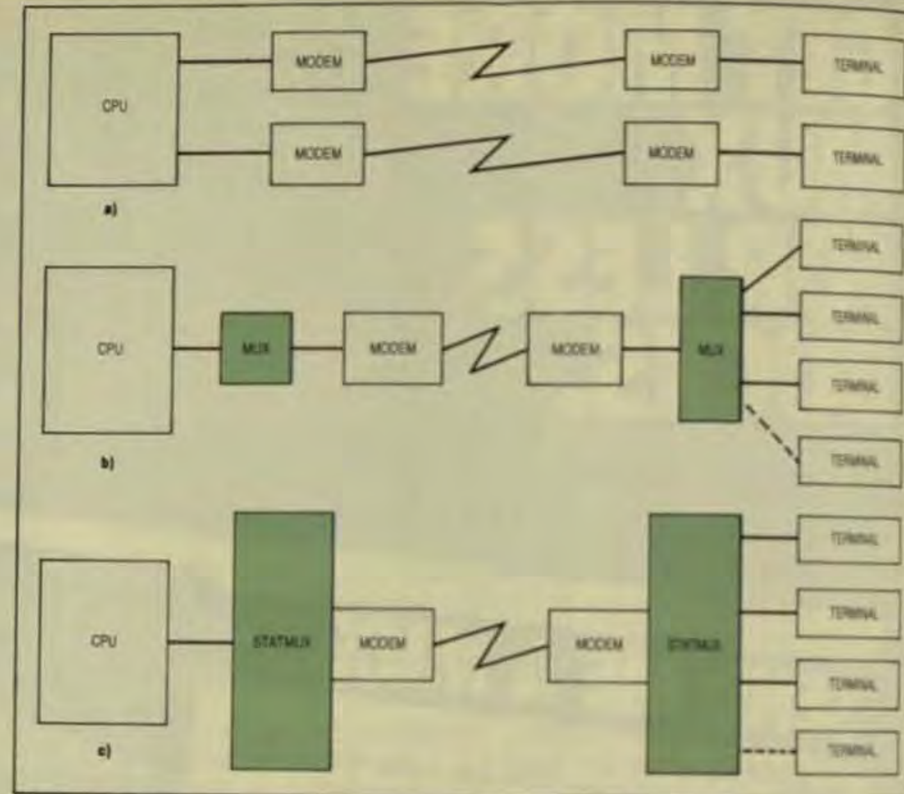

Fig 4 Computer-to-terminal communications has gone through several stages since the late 1970s. Terminals may be linked to computers over individual leased lines with modems on each end for distances over 50 ft (**a**). Data transmissions between a computer and two or more terminals may be multiplexed over a single leased line (**b**), and statistical multiplexers with integral modems can provide economical long-haul communications (**c**).

Significant changes have occurred in communications between computers and I/O terminals. At one time, terminals were grouped in tight clusters and individually hardwired to a computer in the same room. This is still common practice wherever on-site clusters are needed (Fig 4a).

In the past, leased lines were used to link remote terminals to the CPU. But leased lines are expensive. An often-cited rule of thumb suggests a monthly lease charge of $1 per mile per line. Thus, linking eight terminals to a remote CPU involves buying 16 modems and paying operating charges for eight lines. Dial-up lines used in place of leased lines saves money if the terminal is transmitting or receiving approximately less than 2 hrs a day.

A further cost savings can be made if you multiplex transmission between computers and two or more terminals over a single leased line (Fig 4b). This is accomplished by installing one of two types of multiplexers—either time-division multiplexing (TDM) or frequency-division multiplexing (FDM). Time division provides each terminal with its own slice of time on the shared line; frequency division provides each terminal its own slice of the line's frequency band. For example, 10 terminals would be assigned 300-baud portions of a 3000-baud line.

Because of inherently low line utilization in both types of transmission, virtually all new multiterminal communications installations use statistical multiplexers. A STATMUX strips out the idle time so that a line is actively transferring data all the time—as long as data is available—and at full line speed rather than at a proportionate fraction.

Fig 5 shows the major elements in the host and remote STATMUXs that can support a cluster of video display terminals. On the host side, the STATMUX includes a built-in microprocessor connected to a PDP-11 or VAX via a Unibus. On the remote side, the STATMUX is linked to as many as eight asynchronous terminals.

The composite data-transfer rates for full-duplex synchronous communications are switch-selectable at up to 19.2 Kbaud. Maximum data transfer between the individual asynchronous terminals and the remote STATMUX is 9.6 Kbaud. Communications parameters for each terminal/channel—speed, number of bits per character, buffer priorities—reside in the host STATMUX, rather than in the host CPU.

Most STATMUXs in use today handle from 4 to 16 channels, although 32- and 64-channel devices are specified for systems with large numbers of terminals. The main reason is space savings rather than cost savings, because one 32-channel STATMUX is not necessarily less expensive than two 16-channel devices.

Many of the newer devices are modular and allow network managers to upgrade the number of channels easily on-site, without having to return the unit to the manufacturer. For example, Digital's DFM series provides 4-line and 8-line channel-expansion cards that are simply inserted in available backplane slots.

Most intelligent STATMUXs also provide error-control and network statistics. CRC error checking is inherent in the communications protocol used in synchronous transmission. For example, X.25's bit-oriented protocol, based on High-level Data Link Control (HDLC), provides for retransmission until error-free data is received. Although common practice is to use conditioned lines for bandwidths of 9.6 Kbaud and over, the STATMUX error-control function allows networks to tolerate higher noise levels on multiplexed communications lines.

Statistics on buffers, data links, channel status, channel parameters, and EIA functions are useful in identifying potential problems. Overall system and channel-oriented statistics are requested by the network manager under password control and aren't available at user ports.

The network manager can run modem and channel loopback tests to any terminal in a system through a dedicated supervisory port. In some multiplexer designs, that port can be linked only to the remote STATMUX. In others, it may be connected either to the host or to a remote STATMUX. In Digital's DFM series of STATMUXs, users can perform loopback tests on their own channels from VT100 terminals.

Depending on line length, modems aren't always needed in conjunction with STATMUXs. The allowable length depends on the particular model used. With DECmux devices, the remote terminal cluster may be located as far as 3300 ft from the CPU.

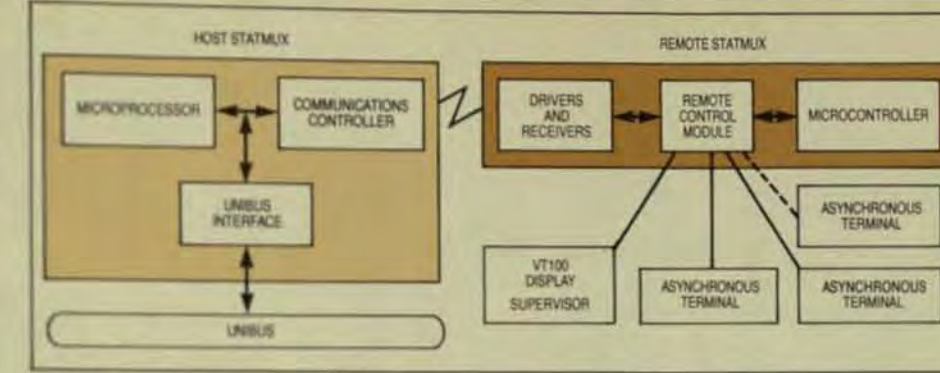

Fig 5 Digital's DECmux statistical multiplexer system configuration can include a microprocessor-based host STATMUX and a remote STATMUX that can handle up to eight asynchronous terminals.

More recent STATMUXs contain integral modems and are used mainly in long-haul communications. The main advantage of integral modems is reduced hardware cost, although prices can vary widely (Fig 4c). The hardware cost for multiplexed lines is significantly lower than for a group of dedicated leased lines. Field experience has shown that the investment required to convert existing dedicated lines to multiplexed lines can be returned in less than a year through reduction in monthly line costs alone.

STATMUXs can be purchased either as built-in system devices or free-standing cabinets. When the STATMUX is built in, a system only requires a single board inserted in the CPU cabinet on the host side, and two smaller boards mounted in the remote terminal.

Free-standing STATMUXs are more flexible, don't obsolete hardware already in operation, and aren't specialized to a particular host bus structure. They can be added to an existing network without further cost since new installations require a communications interface card in the host CPU. This function is integral in system STATMUXs.

ADDITIONAL STATMUX FEATURES

Several features unavailable in the past are now offered on STATMUXs:

- **Synchronous data handling.** A computer can be connected to a normally asynchronous channel at the remote location and exchange data synchronously with the host computer, allowing increased use of a single leased line. The STATMUX manufacturer may provide for conversion to computer-to-computer communications on any number of the available channels. The computers must utilize the transmission protocol the STATMUXs support.
- **Channel switching.** A user at one channel of the remote STATMUX can specify at the terminal keyboard that transmission from the host be switched temporarily to another channel. This feature is useful when a file maintained at the host CPU is to be generated on a high-speed printer on another channel.
- **Channel contention.** A network is likely to involve communications between nodes that have different numbers of user channels. In the event of contention—when 16 users at one node are exchanging data with eight users at a second node, for example—the STATMUX provides for the necessary waits.
- **Speed conversion.** The STATMUX automatically handles speed differences between the host CPU and the terminal channels. For example, without any interaction from users, remote terminals can receive data at 1.2 Kbaud that has been transferred at 9.6 Kbaud from a CPU to a host. ■

Dave Korf *is a product manager for Digital Equipment Corp.'s Special Systems Group. He received a Bachelor's degree in quantitative methods from the College of Saint Thomas, MN.*

Maureen LaGro *is a marketing manager for Digital Equipment Corp. in the data communications area. She earned a BS in engineering physics from the Univ. of Toledo and an MSEE from Lowell Univ.*

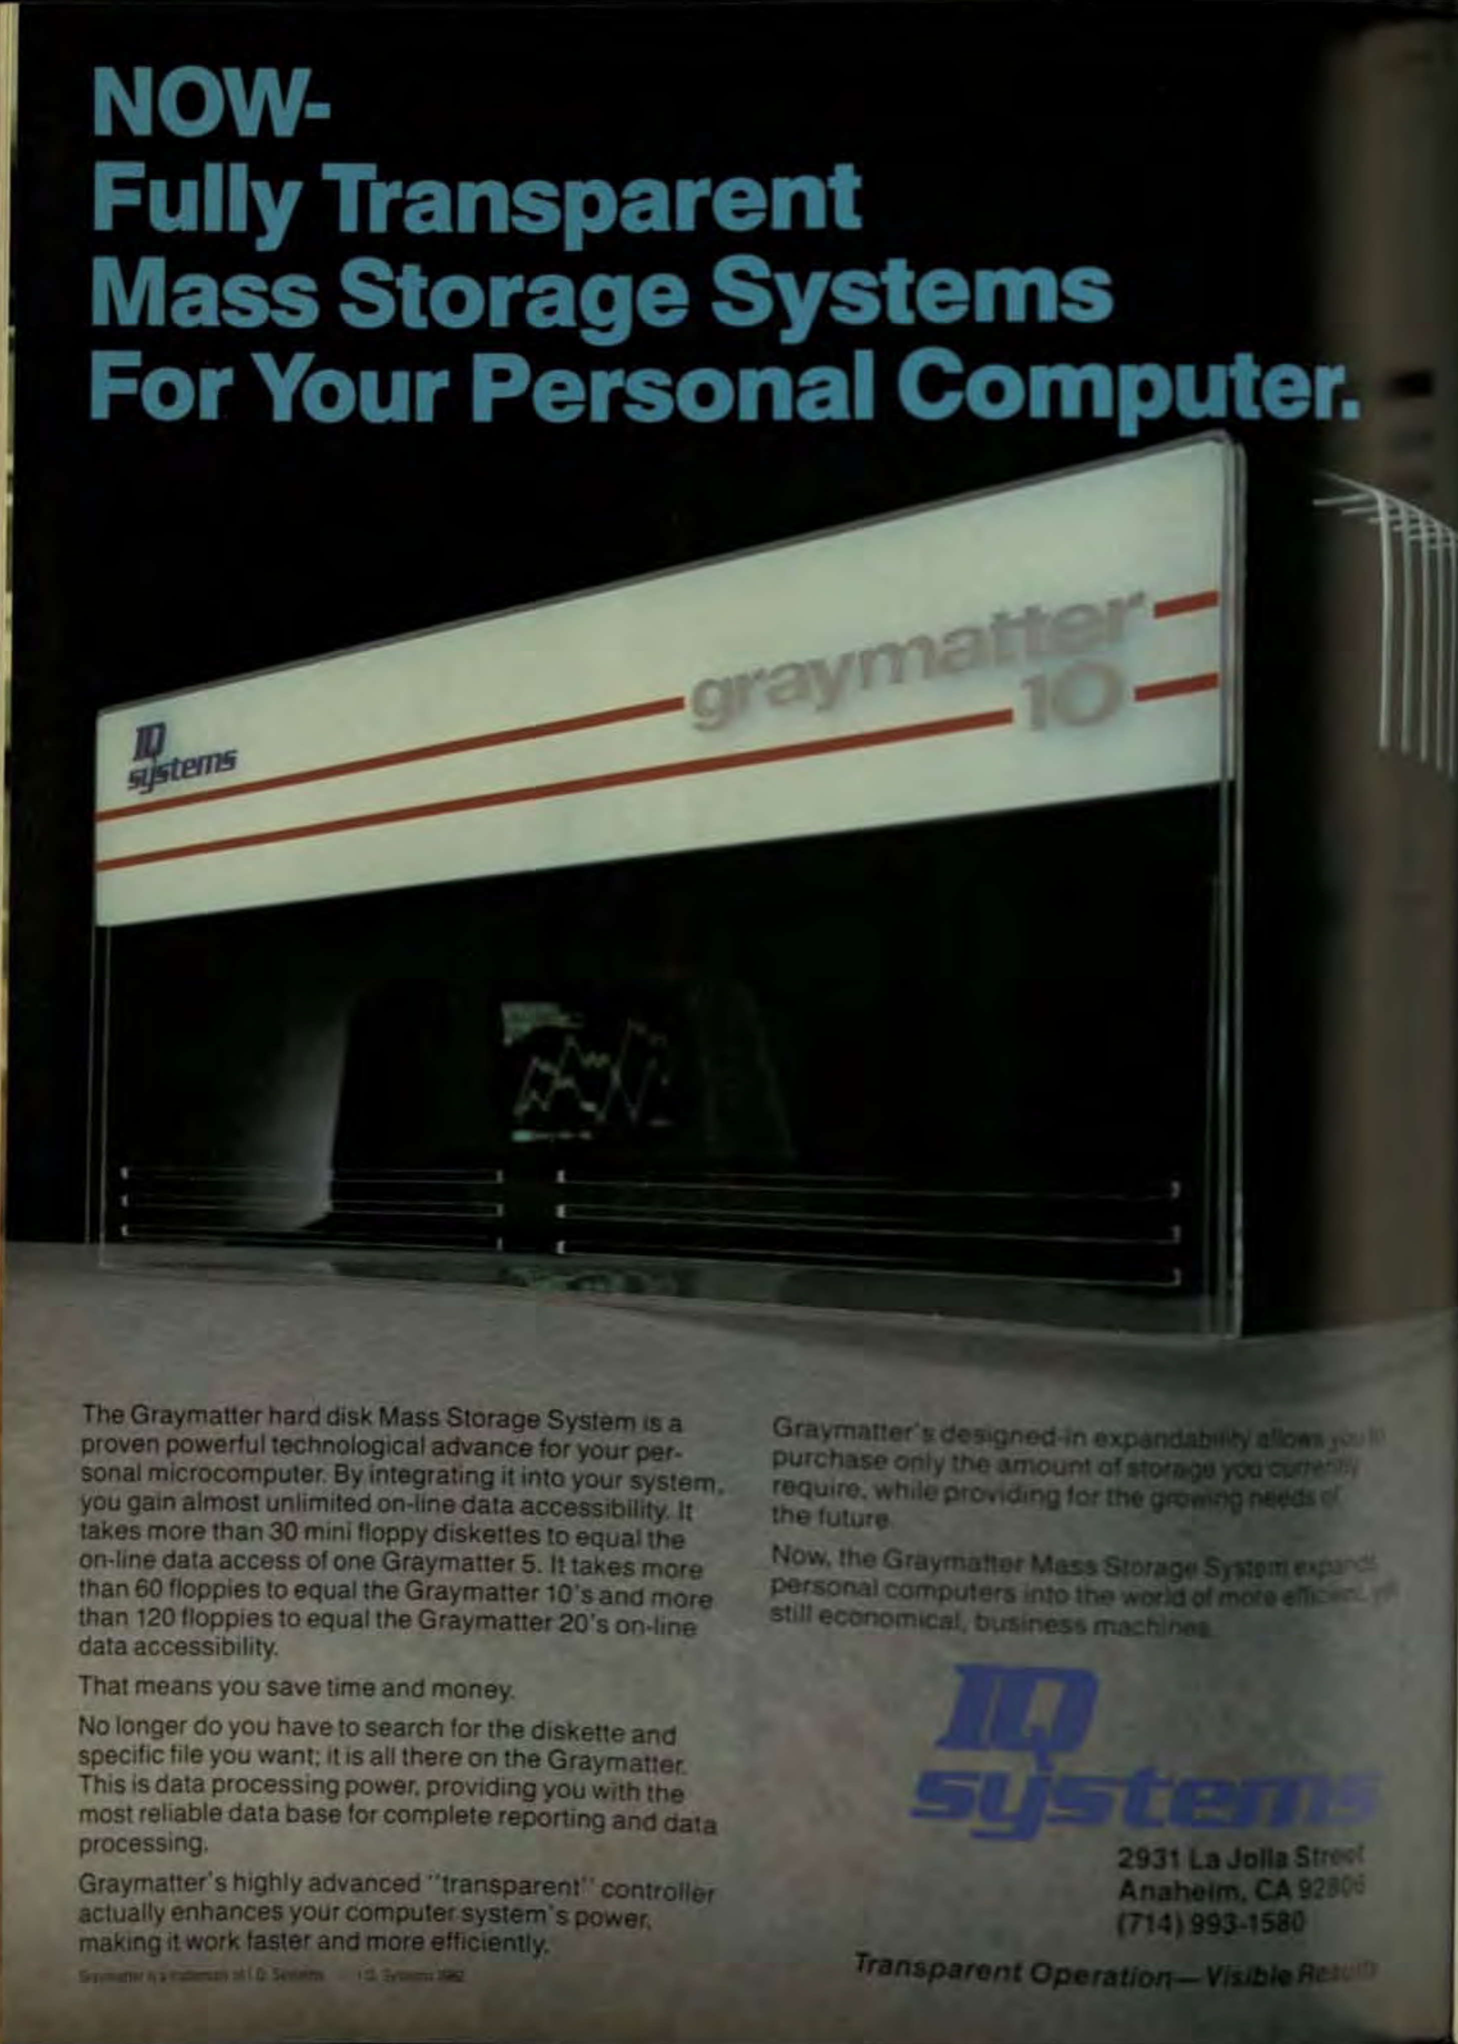

Electronic Spreadsheets Move From the PC Arena to Minis and Mainframes

When electronic spreadsheet programs gain access to company databases, the applications and cost savings are limited only by the user's imagination.

by **Brent Vanderwood**, Datatrend Inc.

Since the introduction of VisiCalc in 1979, electronic spreadsheets have been a major driving force behind the explosive growth of personal business computers. Spreadsheets have given nontechnical people the ability to control computers to solve their personal DP problems. These programs have now moved onto the larger systems that handle an organization's entire DP operations. Although implementation on timeshared systems doesn't provide the responsiveness of a dedicated system, a number of benefits are gained: the ability to access company databases; savings on software, hardware, and maintenance costs; multi-user access to the same spreadsheet; and the consistency of operation that results because all applications can be run on a single system.

SPREADSHEET ANALYSIS

Almost everyone who has anything to do with computers has heard of spreadsheet technology, but even among those who can name some of the more familiar versions, such as VisiCalc and SuperCalc, few are familiar enough with electronic spreadsheets to appreciate their versatility and the breadth of their applications.

Electronic spreadsheets are designed to facilitate the computerization of a basic accounting function: the use of paper with rows and columns to itemize entities and associated numerical data. Most financial people and managers who prepare budgets, forecasts, amortization schedules, and price lists frequently use spreadsheets. In essence, an electronic spreadsheet replaces paper, pencil, eraser, and calculator. A typical spreadsheet might be a price list (Fig 1).

A typical electronic spreadsheet is made up of 63 columns by 254 rows, although the size varies from package to package. The column width is usually variable from 3 to 77 characters. A typical CRT terminal is only 80 characters wide by 24 vertical rows, so it obviously cannot view the entire spreadsheet at once. For a column width of 9 characters, with 4 rows allocated for control information, a typical CRT would be able to display an 8-column by 20-row window into the spreadsheet (Fig 2).

The window on the CRT screen can be scrolled to view different sections of the spreadsheet. In addition, the CRT can be split into two or more windows, each of which can be moved separately. For example, two windows may each show 20 rows by 4 columns, or 10 rows by 8 columns. Again, the width of the columns is variable, so the number of columns the CRT can show at one time depends on the width of the columns.

The intersection of a row and column is called a cell, a box, or an entry position. If the 63 columns are labeled A through BK, and the rows are 1 through 254, then A1 would be the first cell in the upper left corner, A2 the next cell down, B1 the next cell across, B2 the next down, and so on. BK254 would be the last cell in the lower right corner (Fig 3).

The Ever Expanding IMS Product Line

The IMS Family

IMS Computer products not only fulfill the requirements of stand alone applications, they are designed to be cost effective, intelligent nodes in a total network environment! Each product fulfills a particular requirement of the network with a conservative functional overlap of the system above and below it in the Family Tree.
The IMS family is growing rapidly—keeping pace with technology and the ever increasing needs of industry.

Ⓐ Ⓑ 5000SX/8000SX

The SX table top computer systems are the ideal choice for companies with expansion in mind. These systems are easily expandable from one to eight users, each having his own Microprocessor, 64K of memory, and local peripheral control.
The typical SX system is configured as follows:

- Bus monitor processor
- 64K of main memory with byte parity
- Floppy disk controller
- Winchester disk controller with error correction
- Shared peripheral controllers
- Up to eight user processor modules (Each with a Microprocessor and 64K of memory)

The bus monitor with its operating system controls all of the data transfers between the user processors and the shared data base devices. File protects, data lockouts, and print spooling along with the standard operation system housekeeping functions are handled in this logical area. The operating system required by each user processor is relatively simple requiring only a small amount of memory, thereby allowing large application program and buffer area for each user.

Ⓐ The 5000SX chassis has three full width 5¼ inch Floppy/Winchester slots available. Winchesters from 6.3 MByte to 19.2 MByte require one full width slot each. Floppies from 500 KBytes to 1.0 MBytes of storage require ½ slot per drive. The high performance switching power supply in the 5000SX will support any combination of Winchesters, Floppies and User Processors.

Ⓑ The 8000SX chassis has two full width 8 inch Floppy/Winchester slots available with the added capacity to house a magnetic tape bulk memory subsystem. Winchesters from 6.3 to 85 MByte storage require one full width slot each and 1.6 MByte Floppies require ½ slot per drive. The power supply in the 8000SX will support any combination of disk drives, tape subsystem and user processors.
These compact systems fully configured give you performance comparable to that of large Mainframes at Micro Computer prices.

Ⓒ 8000S "MAXIMA"

The IMS 8000S "MAXIMA" Computer system is designed for the company where many people must have access to a large common pool of information. Basically the system configuration of the 8000S is similar to that of the SX Table Top system with the added capability to support up to 16 users each with his own Microprocessor and 64K of memory. The 8000S has five full width 8" Floppy/Winchester slots available supporting any combination of full width Winchesters and ½ width Floppies plus a magnetic tape bulk memory subsystem. This system is the highest performance lowest cost system in its class.

Ⓓ 5000IS - "The Desktop Mainframe"

The 5000IS is the most versatile integrated system available—without compromise. From the crystal clear monitor with a true typist keyboard to the high performance switching power supply—an engineered solution—the 5000IS system is designed with IEEE standard S100 bus architecture, and can be configured as either a 16 bit or an 8 bit system.
You can have 1 MByte Floppies, 19.2 MByte Winchesters, extended RAM memory (beyond the basic 64K), special purpose peripheral controllers, and best of all, the 5000IS can serve as the host processor of a multi-user, multiprocessing system. Up to four I/O processors may be resident in the 5000IS, each with its own Microprocessor, RAM memory and local I/O Channels. With this flexibility you can configure the highest performance, lowest cost multi-processing system available.

Ⓔ IMS "ULTIMA" Terminal

The "ULTIMA" CRT Terminal is the finest stand-alone terminal available. It has a separate microprocessor for each of its major functions: Local Intelligence, Screen Handling, and Keyboard Control. It can be customized to emulate any of the conventional CRT's, with added capability to perform functions the others don't even talk about.
The contoured keyboard has a touch the best typist will really appreciate. We know that differences in keyboards may affect typing speed by as much as 50%. This difference alone can pay for the "ULTIMA" Terminal in a matter of months.

- Keyboard - electrically non-contacting for long life and reliability.
- Monitor - High resolution green phosphor, eliminating eye fatigue associated with lower quality monitors.
- Characters - 7x11 characters in a 9x12 field.
- Character Sets - English, French, German and Spanish.

Ⓕ Portable Cartridge Tape Back Up Unit

The new IMS Stand Alone Tape Subsystem may be added to any IMS system. This high reliability drive uses the industry standard 3M cartridge and features start/stop capability. This permits storage or retrieval of files on an individual basis.
Maximum capacity of each cartridge is 17.5 megabytes. The Stand Alone enclosure includes the drive and power supply with the controller requiring a single slot in the S100 bus mother board.

SOFTWARE:

A wide variety of fully tested, industry standard single and multi-user operating systems along with the finest applications packages available. All are functionally proven for completeness, compatibility and performance to complement the IMS system.

WARRANTY:

Our 2-year warranty on all products manufactered by IMS is a dead giveaway that we produce rugged, top quality professional equipment engineered for reliability.
In addition, all peripheral products supplied by IMS with IMS equipment are covered with a 90 day warranty.

We Build Computers As If Your Business Depended On Them.

For complete information and specifications plus the location of your nearby IMS International dealer, call or write today!
(702) 883-7611

IMS INTERNATIONAL

2800 Lockheed Way, Carson City, Nevada 89701 • Telex 910-395-6051

7 (V) +D7-E7

	C	D	E	F	G	H
1				BUDGET FOR 1983		
5		SALES	COST	PROFIT	EXPENSES	NET/B TAX
7	JANUARY	1500.00	900.00	600.00	500.00	100.00
8	FEBRUARY	1515.00	909.00	606.00	500.00	106.00
9	MARCH	1530.15	918.09	612.06	500.00	112.06
10	APRIL	1545.45	927.27	618.18	500.00	118.18
11	MAY	1560.90	936.54	624.36	500.00	124.36
12	JUNE	1576.51	945.90	630.60	500.00	130.60
13	JULY	1592.28	955.36	636.91	500.00	136.91
14	AUGUST	1608.20	964.92	643.28	500.00	143.28
15	SEPTEMBER	1624.28	974.57	649.71	500.00	149.71
16	OCTOBER	1640.52	984.31	656.21	500.00	156.21
17	NOVEMBER	1656.93	994.15	662.77	500.00	162.77
18	DECEMBER	1673.50	1004.10	669.40	500.00	169.40
19						
20		19023.75	11414.25	7609.50	6000.00	1609.50

Fig 5 A sales forecast in the DataCalc format displays both input data and figures calculated by the program.

growth rate of 1% per month. The contents of cell D8 is 1.01 × D7 (the contents of cell D7 multiplied by 1.01), making February sales 1% greater than January sales. The Replicate command was used to copy the contents of cell D8 to cells D9 through D18, and the cell reference (D7) was automatically changed by the computer to reflect the new location of the data. In other words, .01 × D7 copied from cell D8 into cell D9 becomes .01 × D8.

Similarily, cell E8 contains .6 × D8 to indicate that the cost of goods sold is 60% of the sales amount. F8 contains D8 − E8 (sales minus cost of sales) to give the gross profit. Cell G8 is a fixed expense of $500, and cell H8 contains F8 − G8, which gives the before-tax profit. Once row 8 is created, one command can replicate the entire row (from column D to column H) into rows 9 through 18, thus creating the rest of the forecast for months March through December. Fig 6 shows the actual contents of the cells for the spreadsheet in Fig 5.

Replicate, which is sometimes called Copy or Reproduce, is a very powerful command. Other commands allow users to Insert, Delete, or Move rows and columns, with affected cell references automatically adjusted to conform to their new locations. In addition, cells may be edited as in a word processor, and the results of calculations can be displayed in many different formats—including floating point, dollars, and integers. Spreadsheets can also be saved on disk for later use and printed out in a format that is of boardroom quality.

SPREADSHEETS ON PCs VS. LARGER SYSTEMS

Electronic spreadsheets originated on the personal computer and, at first glance, seem best suited to this type of system.

Spreadsheets are highly interactive programs, similar to word-processing software, and a dedicated PC system will provide maximum responsiveness. Spreadsheets on personal computers typically have such features as vertical scrolling, which allows users to move sideways on a spreadsheet without the entire screen having to be rewritten, and memory-mapped video, which enables a screen to be refreshed almost instantaneously whenever the window moves or data changes.

Larger systems typically utilize terminals that are less versatile. Few, if any, will do vertical scrolling, and they commonly run at baud rates less than 9600, such as 4800, 2400, or even 1200. Also, timesharing can cause delays that are not experienced on dedicated PC systems. There are, however, several significant advantages provided by large system implementations that compensate for the loss of the better responsiveness of a dedicated system. These include:

- **Dual usage.** Terminals tied into a data-processing system can be used for standard DP applications as well as to solve spreadsheet problems, and usually the system printer can be used. In addition, there is a one-time cost for the spreadsheet software, and it can run on every terminal attached to the system. The result is a savings on software, hardware, and maintenance expenses.
- **Consistency of use.** Users do not have to become familiar with two different systems to do spreadsheets as well as other applications. Also, speadsheets for large systems are typically done by software engineers who are familiar with the de facto standards for those systems, such as the use of the "gold" key to initiate software commands on Digital's systems. In addition, the filing and retrieving of spreadsheets is handled by the host system, which is usually faster—and certainly easier—than the juggling of floppies.
- **Consolidation and standard templates.** A financial manager (or any manager) may want to develop standard outlines for spreadsheet reports to be prepared by different departments. On a multi-user system, the format can be prepared and saved in an account where the various departments can use their own terminals to access it and provide the required data. Then selective information, such as totals, can be extracted from the individual reports and consolidated into the main report. The result is less work and more management control.
- **Access to company databases.** Integration of spreadsheets into a total DP system is much more effective and easily accomplished when spreadsheets are on the DP system itself than with personal computers that have a communications capability.

	C	D	E	F	G	H
4						
5		SALES	COST	PROFIT	EXPENSES	NET/B TAX
6						
7	JANUARY	1500	.6*D7	D7-E7	500	F7-G7
8	FEBRUARY	1.01*D7	.6*D8	D8-E8	500	F8-G8
9	MARCH	1.01*D8	.6*D9	D9-E9	500	F9-G9
10	APRIL	1.01*D9	.6*D10	D10-E10	500	F10-G10
11	MAY	1.01*D10	.6*D11	D11-E11	500	F11-G11
12	JUNE	1.01*D11	.6*D12	D12-E12	500	F12-G12
13	JULY	1.01*D12	.6*D13	D13-E13	500	F13-G13
14	AUGUST	1.01*D13	.6*D14	D14-E14	500	F14-G14
15	SEPTEMBER	1.01*D14	.6*D15	D15-E15	500	F15-G15
16	OCTOBER	1.01*D15	.6*D16	D16-E16	500	F16-G16
17	NOVEMBER	1.01*D16	.6*D17	D17-E17	500	F17-G17
18	DECEMBER	1.01*D17	.6*D18	D18-E18	500	F18-G18
19						
20		SUM(D7...D18)	SUM(E7...E18)	SUM(F7...F18)	SUM(G7...G18)	SUM(H7...H18)

Fig 6 The numbers in Fig 5 are based on different types of calculations made by the program

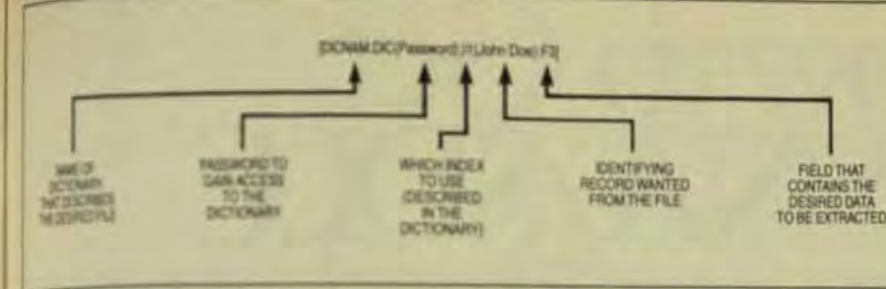

Fig 7 The format used by DataCalc to allow you to access external data is simple and interactive.

INTEGRATED SPREADSHEETS

An integrated spreadsheet is desirable because it helps to eliminate repetitive data entry and provides access to up-to-date information. Managers have ready access to information they need for reports they design and build themselves. For example, in the preparation of a budget or forecast, the projections are based on real figures, such as the last month's or last year's sales and expenses. This information is available in the general-ledger databases. With an integrated spreadsheet/DP system, the information can be extracted from the general ledger automatically and placed into the spreadsheet.

Another example is the data managed by inventory-control software. With access to information in the inventory databases, all kinds of relevant reports can be generated. Perhaps a parts list of a special-order product is needed, or maybe management needs the figures to develop a purchasing plan. With the information readily available, use of an electronic spreadsheet is limited only by the user's imagination.

Electronic spreadsheets have become popular because they are easy to learn and use by people who are not computer experts. The method used to access data in external databases must follow this example to be truly usable by nontechnical people. In addition, it's important to restrict a spreadsheet's ability to access data in the company databases, so that only authorized personnel may do so. And finally, for a spreadsheet to reach the largest audience and gain popularity among systems integrators, it must be flexible enough to be interfaced to a variety of software products.

With DataCalc, when a spreadsheet user wants to extract data from an external database, the format described in Fig 7 is entered into a cell. When the entry is completed, the program will pull the indicated data from the database and enter it as a value that can then be accessed by other cells in the spreadsheet.

Data security is maintained through the operating system's security mechanism and through use of an encoded password within the dictionary. Users can be restricted from gaining access to the software tools that are used to build dictionaries, and the data in external files can only be read, not written.

The dictionary approach provides flexibility for the interfacing of DataCalc to a variety of data file types, such as those using the Mini-Computer Business Applications (MCBA) indexing method, Digital's Index Sequential Accessing Method (ISAM), or even straight sequential files. File characteristics are described in the dictionary, as well as information such as the primary and secondary keys, location of the file, field layout, and the encoded password.

Electronic spreadsheets have brought about a revolutionary change to the world of data processing. As a result, many applications that previously required a programmer can now be handled by managers themselves. ■

Brent Vanderwood, *president of Datatrend Inc., was previously product marketing manager for Plessey Peripheral Systems. He has a BA in political science from California State Univ., Fullerton.*

DON'T SAY IT. WRITE IT.

If you've got something important to say about the technology in your field of interest, write it down—in a feature article for Computer Technology Review. Speak your mind in the one publication where working technology speaks *its* mind.

Call us at (213) 208-1335 to discuss your ideas and obtain your author's guide, or write:

Computer Technology Review, 924 Westwood Blvd., Suite 650, Los Angeles, CA 90024.

SNA Is Key to IBM Strategy for Office Use Of Personal Computers

For personal computers to be an integral part of the automated office, software based on Systems Network Architecture will be needed to link them to corporate mainframes.

by **George Haskell,**
Communications Solutions Inc.

Systems Network Architecture (SNA), IBM's data communications architecture, once epitomized the bigness of IBM's approach to data communications as well as to data processing in general. Major host-based SNA access products such as ACF/VTAM and database/data communications applications like CICS/VS require huge amounts of memory and external storage. Front-end processor software such as the Network Control Program (NCP)—and even distributed operating systems such as DPPX, which runs in the 8100—are generally characterized as large and cumbersome.

It seems strange, therefore, that the personal computer seems destined to form a major component in IBM's networking strategy. This article will examine where the personal computer fits in the SNA arena, with particular emphasis on the capabilities of IBM's Personal Computer (PC), which will be the target system for other personal computer manufacturers to shoot at when they attempt to offer SNA-compatibility on their own systems.

In addition to looking at current SNA capabilities, this article will take a close look at the additional capabilities that are required if the personal computer is to achieve its full potential as an SNA network component.

While the personal computer may originally have been thought of as a home computer, it's become obvious that the main marketplace for these systems is in the corporate world. Large numbers of personal computers are being sold to individuals for home and small business use, but an estimated 30% of all personal computers are being sold to large companies for use as a management productivity tool.

As the pricing battles continue to shape up over the low-end home market, it's fairly obvious that the personal computer vendors destined for success are those that make the PC into an effective corporate tool. That means they must go beyond strictly local processing capability, as exemplified by word processors and spreadsheet managers, and provide access to the power of firms' central-site mainframe computers. And given the dominance of IBM in this mainframe marketplace, additional power means SNA compatibility for personal computers.

PERSONAL COMPUTERS VS. DUMB TERMINALS

Before getting into a discussion of the SNA capabilities of existing PC products or requirements for future products, it will be useful to differentiate between the communicating personal computer and the type of dumb terminal it will be replacing in some applications.

The obvious difference between personal computers and dumb terminals is local intelligence. With a personal computer, a user has the ability to process data locally, rather than merely look at it, which is all he can do with a dumb terminal connected to a host application. Inquiry terminals depend on host-based applications to process data, and the only flexibility the terminal user has in

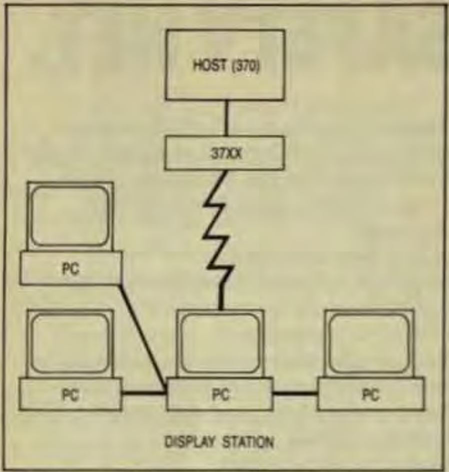

Fig 2 Non-IBM vendors offer packages to allow up to three PCs to be connected to a controller in one PC. The PC containing the controller board appears to the host as a cluster controller supporting four display stations.

rate 3274 cluster controller with a single display station attached. A normal 3274 cluster controller supports the connection of up to 32 display stations and/or printers, and a 3276 supports the connection of up to eight display stations and/or printers.

This shortcoming can have a major impact on a user attempting to utilize PCs in 3270 emulation mode when he's configuring an SNA network. For example, if used in dialup mode, each PC would require its own individual data link and its own port on the 37XX front-end processor. A real 3270, on the other hand, could support up to 32 users on a single dialup link (Fig 1).

If leased lines were utilized instead of dialup, it would be possible to multidrop multiple PCs on a single data link attached to a single 37XX port. However, it seems likely that dialup access will be appropriate in many cases, given the nature of the PC user. Again, 3270 emulation packages from other vendors do allow PCs to be clustered on a single link. The Access/SNA package mentioned above, for example, allows up to three additional PCs to be connected to the synchronous communications controller board in one PC.

In this configuration, the PC containing the synchronous communications board appears to the host as a cluster controller supporting four display stations, and to the three PCs connected to it as a protocol converter, converting their asynchronous data streams to SNA 3270 data streams (Fig 2). This type of configuration greatly enhances the configurability as well as the price/performance of the PC and the 3270.

At this point, it's worth speculating on IBM's reasons for keeping the level of SNA support on the PC relatively low. Even with the existing level of 3270 support, the PC is beginning to put price pressure on the 3270 itself. For example, in small configurations (four displays, for example) it's currently cheaper to utilize four standalone PCs running IBM's 3270 emulation package (which costs approximately $1000 for the hardware/software combination) than it is to utilize a 3276 with four display stations (see Table). If a PC/3270 package that allows clustering is utilized, the price comparison becomes even more favorable.

IBM has an extremely large installed base of 3270s; estimates approach 1.8 million display stations/printers. The firm is quite probably not ready to see that base affected in any major way, even by its own product.

One way IBM is attempting to protect its 3270 installed base while satisfying user desires for the additional capabilities of the PC is by introducing a personal-computer attachment option for the 3278 display station, which allows the user to operate in two modes—one that utilizes the 3278 display as a monitor for the PC, and one that uses the 3278 as a display station connected to a mainframe-resident application program via a 3274 controller.

TABLE — PC/3270 EMULATION PRICING

	1 DISPLAY	2 DISPLAYS	3 DISPLAYS	4 DISPLAYS
3270	7,755	9,475	11,195	12,915
ALL PCs	3,794	7,558	10,827	13,696

The user is able to switch from one mode to another without affecting the host or PC application. This package also includes sample programs that enable users to transfer screens of data from the host-connected 3278 to the PC and vice versa. In addition, as a monitor for the PC the 3278 offers such features as 16 levels of intensity (when used with the PC color-graphics adapter), user-defined character sets, character attributes, and business graphics.

The 3278 is an extremely expensive monitor for a PC. IBM's 3270 Personal Computer attachment option is clearly intended for users who already have 3278s. The option allows them to buy PCs without monitors and utilize their existing 3278s instead, gaining a PC/3270 interface in the process.

One other type of SNA/3270-oriented product that is currently available (IRMA from Technical Analysis Corp.) enables the PC to emulate a 3278 by being connected directly to a coax cable that's connected to a 3274. In this way, the user can have alternate access to host-based applications or local PC-based applications.

IBM also offers an adapter that allows PCs to be connected to the 5520 Administrative System, a major IBM entrant in the office-automation marketplace. This option lets the user access the 5520 (and through it, host-based applications) by emulating a 5250 terminal (the terminal that usually connects to the 5520) while alternately functioning as a standalone PC.

WHAT'S REQUIRED?

From a communications standpoint, for personal computers to mature as management tools, the types of features that are required involve expanded connectivity and applications that provide easier and more flexible access to mainframe databases.

As far as applications are concerned, the market is already seeing the beginning of the type of applications described above. For example, VisiLink from VisiCorp gives PC users access to the Data Resources database. This is one of the largest economic databases in the world. Based on a catalog of available data, the user can ask for a package of information (for example, economic performance of the automobile industry over the last 10 years) that will be sent from the mainframe-database to the personal computer.

What makes this package especially useful is that the data is put into VisiCalc spreadsheet template format when it is sent to the PC. Then, the user not only has access to a large database, the data is presented in a way that the PC user can perform extensive local (pc-based) manipulation of the data to suit his particular needs.

The VisiLink package, like most other PC-based communications packages, utilizes asynchronous communications not SNA. This protocol is not nearly as accurate as the SDLC link-level protocols utilized by SNA and generally requires some type of error checking to be built into the applications to ensure the integrity of

Don't miss America's #1 computer trade show. Whether you're an old pro or new-to-market ISO, COMDEX is for professionals who want to turn ideas into profit.

It's not easy for today's ISO to keep up in the ever-changing computer market. Hot today, gone tomorrow.

That's why COMDEX is America's #1 computer show. It's the place for competitive ISOs who want to stay up front.

And COMDEX/Fall '83 is where you'll make contacts, pick up new lines and learn ways to improve the profitability of your business in 1984.

Where you'll find the world's leading manufacturers and suppliers as well as the up-and-coming stars. More than 1200 exhibitors in 5000 booths.

Where you'll find the newest products and services that will make headlines in the trade press after they debut here.

And the COMDEX Conference sessions are where you'll keep pace with the latest business strategies. Where internationally-prominent industry experts keep you abreast on everything from computer retailing to marketing strategy and management methods.

Profit from America's #1 Computer Show

COMDEX™/

November 28-Decemb

LAS VEGAS, NEV

• Las Vegas Convention Center •
• Las Vegas Hilton Hotel • Sahara

I WANT TO TURN IDEAS

☐ Send me information on attending COMDEX/Fall '83

Name ______
Title ______
Company ______
Address ______
City/State/Zip ______

World
expo
COM
and

the data transmitted over the communications links.

Asynchronous communications is often used to access dialup databases because of the wide number of asynchronous terminals that exist in the user marketplace. (Most personal computers as well as dumb ASCII terminals offer asynchronous communications.) However, most mainframe sites that maintain these large databases are migrating toward SNA communications. As SNA communications capability becomes more prevalent on personal computers, it will be advantageous—to both the provider of the database service and the user—to access the data via SNA communications.

Another example of the type of software package that is making the PC a much more valuable management tool is another offering from VisiCorp—VisiAnswer. This product (which is not yet released) was produced as a joint venture with Informatics, a company that has long been providing mainframe software packages.

With VisiAnswer, the personal computer will have the ability to build high-level queries against IBM mainframe databases (primarily IMS databases). Given a basic knowledge of the structure of the data records in the database, the user will be able to select the specific data he requires and have it transmitted to the personal computer.

While this application is also being implemented by use of asynchronous communications, it is exactly the type of product that will free the personal-computer user from dependence on the application programming staff in large organizations. Packages like this can give the user almost unlimited access to the specific information he needs out of the large databases normally maintained by most large companies.

This type of product is designed for use within private companies (unlike VisiLink, which provides access to a public database). Migration to SNA capability is even more critical because most large companies are migrating their communications networks to SNA. It's estimated that 90% of IBM's Fortune 1000 customers will be using SNA as their networking standard by 1985. Under SNA, support for asynchronous terminals is awkward, to say the least, requiring the use of protocol converters or special software (the Network Terminal Option) running in the 37XX front-end processor.

The evolution of these database-access products will follow the evolution of mainframe database-management products, which are attempting to give users easier access to databases that often are quite cumbersome to access even from mainframe-resident applications.

VisiAnswer may exemplify the direction from which the new pc/mainframe products come. Traditionally, vendors of distributed systems (chiefly minicomputer vendors) have avoided the production and marketing of mainframe-based software. Such vendors have restricted their distributed products to interfacing at the highest level possible with mainframe-resident subsystems such as CICS/VS, IMS/VS, and JES (teleprocessing monitors that are utilized in most IBM mainframe environments).

While these vendors may have avoided dealing with mainframe software for good reasons (keeping to their own area of expertise), such software is going to be critical in order to give the remote-personal-computer user flexible access to mainframe data. Joint ventures between personal-computer software vendors and mainframe software vendors seem to be a reasonable approach to providing a total pc/mainframe-based package that will provide access to information and the ability to manipulate it in any way required.

ENHANCED CONFIGURABILITY

The second area in which the personal computer will likely evolve is that of configurability within the SNA network. As discussed earlier, current products simply emulate SNA cluster-controller nodes and are restricted to a direct connection to a host through a front-end (37XX) processor.

This type of connectivity is common in SNA and is exactly the way most terminal products (3270, 3770) are still configured in current SNA networks. Their only access is to the host, not to each other. However, as SNA moves into more truly distributed application environments like office automation, a strong demand is being felt for peer-to-peer connectivity between terminals as well as connectivity to the host (mainframe).

This type of connectivity has been introduced over the past couple of years in products like the 5520 Administrative System, the 6670 Information Distributor (a laser printer), and the 8100 distributed computing system. These systems use SNA host-based communications for document archival and retrieval, and SNA peer-to-peer communications for electronic document distribution (electronic mail, remote printing and display, etc.).

IBM's original implementations of SNA peer-to-peer communications between remote devices were limited to communications between similar devices (5520 to 5520, 8100 to 8100). However, peer-to-peer communications capability under SNA is rapidly evolving to allow this type of connectivity between dissimilar systems.

IBM has recently introduced a new subset of SNA capabilities referred to as Logical Unit Type 6.2 or Advanced Program-to-Program Communications (APPC). This capability has been announced for products like the Scanmaster I facsimile terminal, the Displaywriter, and the System/38, dissimilar systems. LU 6.2 also allows for distributed transaction processing, in which a transaction may be started on one set of systems and continued or completed on a successive series of connected systems.

IBM has indicated that the capabilities provided for in LU 6.2 are going to be the standard for SNA communications products, with full connectivity between a very broad range of LU 6.2-compatible products. Given the local-processing capabilities of personal computers, these systems seem to be naturals for this level of SNA capability.

Given the combination of continually falling prices, full peer-to-peer as well as host-SNA connectivity and a series of host/pc-based application software packages, the personal computer seems destined to take its place as an important tool for management productivity. ■

George Haskell *is director of publications and training at Communications Solutions Inc., which provides consulting, technical publications and public and in-house seminars on SNA. CSI also markets microprocessor-based SNA-compatible software products to provide non-IBM vendors with SNA compatibility.*

SYSGEN ANNOUNCES TAPE BACKUP FOR THE IBM PC.™ WITH OR WITHOUT HARD DISK.

Sysgen II-10 Fast streaming cassette backup with 10 Mbytes of disk: **$2995.***

Sysgen II-20 Fast streaming cassette backup with 20 Mbytes of disk: **$3795.**

Sysgen Image™ Fast streaming backup for IBM's XT: **$995.** (Right. Nine hundred ninety-five dollars!)

Sysgen is delivering complete low-cost Winchester systems with high-speed tape cassette backup.

Low cost means that everyone can afford to keep their valuable data safe and movable.

High-speed means you can back up an entire Sysgen II 20 Mbyte hard disk system in *four minutes max.*

Our hard disk systems operate under PC-DOS™ or CP/M-86.™

Under PC-DOS, they look like disk drives C: through L:

Under CP/M-86, you get fourteen new "disk drives."

And if you own an IBM XT, which includes a hard disk, you can get safe, fast backup with our new Image streaming tape cassette backup system for only $995.

Look to Sysgen for fast, inexpensive backup. With or without hard disk.

For complete information, see the smartest computer store in your neck of the woods. Sysgen, Inc., 47853 Warm Springs Road, Fremont, CA 94539. 800-538-8157 ext. 970; in Calif., 800-672-3470 ext. 970.

SYSGEN
·INCORPORATED·

Machines. CP/M-86 is a trademark of Digital Research, Inc.

Industry faces an unprecedented challenge. To compete successfully in the world marketplace, we must not only look to new plants and equipment, but must be more productive through automation of existing machinery.

Gould is in a unique position to help you automate through a skillful blend of electronic hardware, software and engineering expertise. Our broad range of products enables you to create the precise system you need today. You can start with a single process and upgrade to full automation as your requirements dictate.

Complete automation often requires powerful minicomputers, graphics terminals and displays, programmable controllers, servomotors and controls, solid-state switching devices and measurement sensors, all interconnected by a flexible communications and control network.

Gould offers all of the integrated electronics products required to do the job. All standard and all compatible, right from the start.

Gould expertise in automation and control can improve productivity for you...now. Call or write:

Gould Inc., S.E.L. Computer Systems Division, 6901 West Sunrise Boulevard, Fort Lauderdale, Florida 33313, 1-800-327-9716.

GOULD
Electronics & Electrical Products

Computerized intelligence for factory automation.

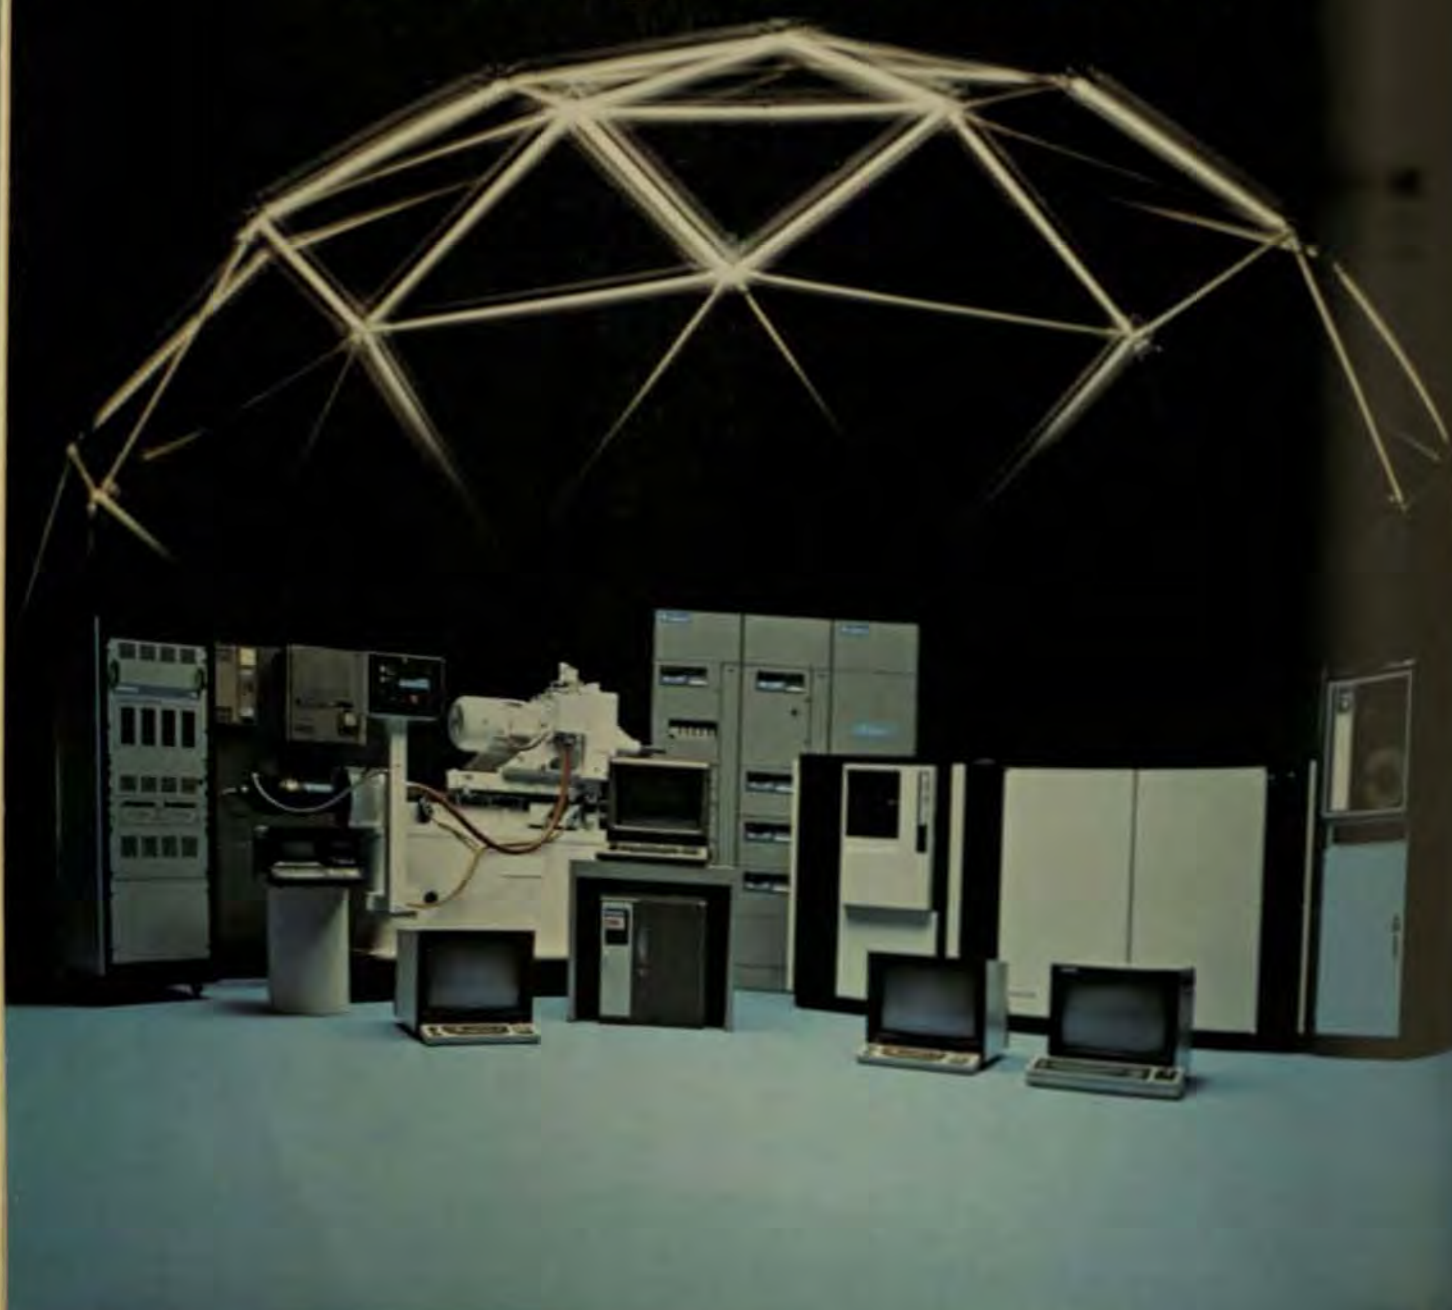

Nuclear Fusion Reactor Uses Highly Integrated Monitoring and Control

Controlled nuclear fusion requires a system designed for complex experimentation that can be reconfigured or upgraded weekly as needs or new technology dictate.

by Robert Daniels,
Princeton Plasma Physics Laboratory

At Princeton University's Plasma Physics Laboratory (PPL), computers are being applied in what may prove to be one of the most important research and development projects ever undertaken—unlocking the knowledge of how to control the nuclear-fusion process that's responsible for producing the energy of the sun and other stars. To acquire this technology may mean the definitive solution of the energy crisis and the availability of inexhaustible energy on earth for as long as civilization exists.

Slated to play a critical role in realizing these hopes is a large multiprocessor system with shared memory and high-speed data links, called CICADA (Central Instrumentation, Control, and Data Acquisition). This powerful system is used to control and monitor, in real-time and on-line modes, the extraordinarily complex, intricate, and formidable experimental reactor. Among other things, the Tokamak Fusion Test Reactor (TFTR) must create temperatures more than six times higher than those in the center of the sun.

In this race to open the source of endless energy by fusing the atoms of omnipresent hydrogen, the United States appears to be in the lead with the $315,000,000 TFTR nearing completion at the PPL under funding from the U.S. Department of Energy.

The first successful tests with the new reactor were made last Dec. 24, and the historic test of scientific feasibility—where enough energy is obtained from fusion so that the reaction could be made self-sustaining—is scheduled to be achieved in 1986. In this historic test, the TFTR will produce 30 MW of pulsed fusion power, paving the way for actual engineering test reactors.

In scope and complexity, the reactor and the computer system are comparable to the systems used by NASA to send men to the moon. And like the conquest of space, the conquest of controlled nuclear fusion would be equally impossible without the aid of computers.

THE BASIC CONFIGURATION

CICADA was designed to perform two different but related types of tasks: to control and monitor the operation of the reactor, and to analyze the results of the fusion tests through a diagnostic system that incorporates data acquisition, reduction, processing, display, and archival (Fig 1). The system also provides the facility's timing and synchronization controls as well as the hard-wired personnel-safety and equipment-protection systems.

The system must handle a total of 20,000 different control and monitoring points for the reactor and must acquire up to 20 million 16-bit words of data generated during a fusion pulse that lasts one second.

In line with the scope of machine-interface points, the number of user interfaces is exceptional. Not counting continuing program preparation on a dedicated computer with 32 ports, the system will be able to interact with 32 terminal operating stations and 26 operational consoles. A planned off-line-analysis computer will handle 50 terminals.

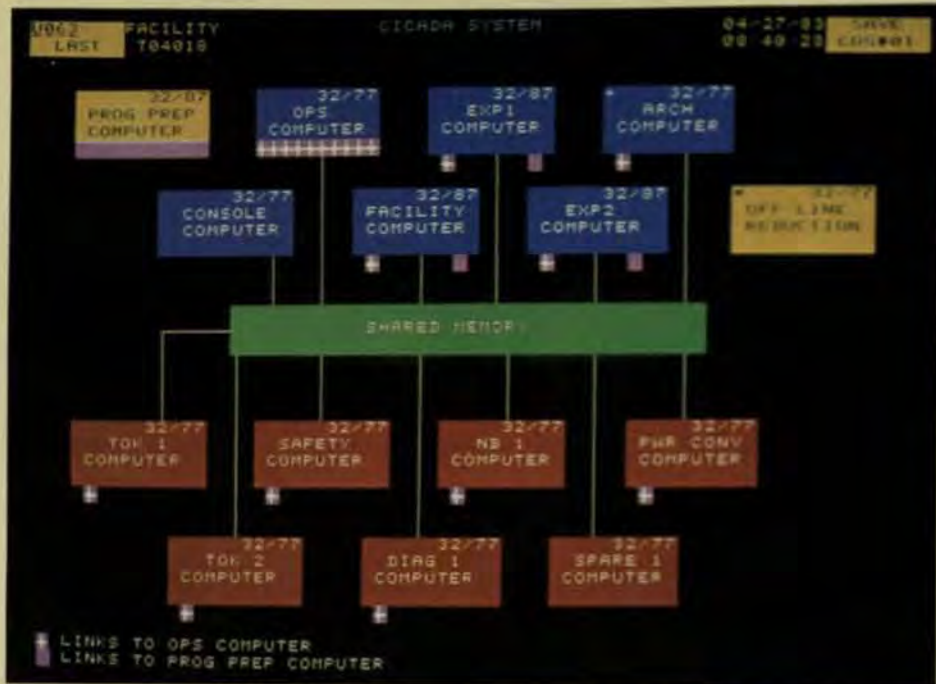

Fig 1 The configuration of the CICADA system, as it will look in the latter part of 1983, is displayed on the console operating system. The blue computers are central system processors that execute application programs and other overall system tasks. The red computers control and monitor Tokamak operations and are interfaced to the CAMAC data-generating system.

The 26 console operating stations handle the applications that demand interactive real-time control and updating of the reactor. The 32 terminal operating stations are for data-diagnostic functions that demand the ability to interact with tables of set points, sequence and data-reduction parameters and limits, as well as to issue commands directly.

In addition to the 50 off-line terminals, an extensive system of computer-generated diagnostic status displays and waveform graphics is distributed throughout the facility over a 70-channel carrier-television network. This has been dubbed the comfort display system, after the name for a similar arrangement at the Fermi National Laboratory.

The basic computer configuration is a tier of central-system computers, linked by shared memory and tied into a communications computer (called Operations) that acts as the central node of a star network. CICADA user-application programs, written in FORTRAN 77, can run in any of the central-system computers and can access the resources of any other computer in the system.

A secondary tier of subsystem computers is linked directly to the operations computer and, through it, linked indirectly to shared memory.

The subsystem computers provide the current and real-time control requirements for the system through interfaces to a Computer Automatic Measurement and Control (CAMAC) system, which is the international standard for computer interfacing.

There are important real-time demands on this system, beginning with the requirement that CICADA take a snapshot of all 20,000 control and monitoring points in the reactor and deposit it in shared memory once every 2 s and immediately before firing the reactor. Every 200 ms, a central computer dedicated to graphics for the entire system must poll user demands. And every millisecond, a computer for the power-conversion system must control the power supply in a process-control mode.

In addition, other real-time and process-control tasks are handled by 32 microprocessors located in the CAMAC system. Approximately half of the CICADA hardware budget was spent on the computer system. Included in this are 14 Gould S.E.L. minicomputers using the MPX-32 operating system and four 600-Mbyte disks, five 300-Mbyte disks, fifteen 80-Mbyte disks, several 6250-bpi tape drives, four sets of background peripherals, a multiport shared memory, and the necessary networking to tie the system together.

These computers will soon be supplemented with three more superminis, and by the time full operation begins in 1986 there will be three additional units.

CONTROL AND DIAGNOSIS

The reactor is essentially a doughnut-shaped magnetic container (Fig 2) called a Tokamak (see Box). Here a plasma of hydrogen atoms can be contained long enough—up to one second—so that an ohmic current and particle-beam injections can raise the temperature of the plasma sufficiently to fuse the atoms.

The plasma diagnostic system contains an array of highly sophisticated measuring devices to provide the critical monitoring of the state of the plasma and its density, the plasma ohmic current, the temperatures attained, the number of fusion reactions achieved, and other critical parameters.

These measuring devices range from microwave interferometers and doppler expanders of laser light to spectrometers and residual-gas analyzers. Still another varied set of instrumentation monitors the state of the reactor itself, including its powerful magnetic fields and massive electrical system.

Both sets of devices are interfaced to the CICADA computer system through CAMAC. The CAMAC system contains the varied data-generating devices and electrical interfaces for the computers and is linked by fiber-optic cable from the reactor cell about 900 ft away. The CAMAC system must collect as many as 20 million words of data relating to the plasma diagnostics as well as the data for the 20,000 control and monitoring points.

There are more than 50 different types of CAMAC modules, from 32-channel transient digitizers and scalers to strobed and unstrobed digital inputs. Each module holds up to 16 digital-control points and is capable of interfacing with a variety of data devices.

Thousands of analog channels are digitized and put into memory in the modules, which currently fit into 100 CAMAC crates, each of which holds up to 25 modules. Eventually, there will be 450 CAMAC crates holding 6400 modules.

The interface between the CAMAC modules themselves and the application programs is provided by the device control system, which provides the means to define the 20,000 monitoring and control points and their related devices, using 15-character names. These devices can be individual digital or analog components, including microprocessors resident in CAMAC as well as groups of components. Such devices can be added, modified, or deleted, and can be controlled automatically by up to three digital-control commands.

Every 2 s, the device control system automatically scans all devices defined for Tokamak and transfers the status parameters to shared memory where they are available for access by application programs and for periodic display. The monitoring information is available to any program, but control authorization must be assigned.

A subsystem table contains all the information on the location, components, set points, alarm and trip limits, and parameters for conversion to engineering units for all the control and monitoring points. At each 2 s scan, all digital and analog points are compared against alarm, trip, and transducer limits. Violations can change the absolute status flags and enact subsystem actions. Digital points, in particular, are also checked for a transition from normal states or to normal states, and the results are logged on a terminal if appropriate.

The operating physicists and engineers who control the operations of the Tokamak do so through the console-operating stations and the console-operating software (Fig 3).

Among other things, the software includes the device-entry utility that permits the operations chief to add, modify, or delete any of the Tokamak control devices. Besides providing the services to update the console displays and read console devices, the software also allows a user to log-in at a console and select the desired application program (under assignment of control by the operations chief).

Since all application programs executing under console control reside in memory, there is immediate response to user demands. This speed allows such features as software knobs that have real-time feedback on adjustments.

Each console operating station contains two 512 × 512 color-graphics displays that are connected to a gray-scale hard-copy unit. One display is for color graphics and the other is dedicated to the alphanumeric/process-control character modes, with cursor control via a trackball, an attention button, a knob, and a keyboard. There are also four video monitors that carry comfort displays of summary waveform data generated immediately after each pulse.

The data-management system provides utilities and services to schedule the arming of CAMAC data-acquisition hardware, to read data either directly from logically specified CAMAC devices or from on-line or archival files, to write these files, and to write results and status data for intersystem access.

The diagnostic data is available for on-line interaction by physicists through the 32 terminal operating stations. Physicists may also use the terminals to set up parameters for measurement of the plasma during a pulse.

When diagnostic data is requested by an application program, it's transferred to the program with a header containing the parameters necessary to analyze the data. Application programs running on these stations can also be executed on the group of approximately 50 off-line ter-

How to become a successful

COMPUTER CONSULTANT

Have you ever wished you could quit your job and start working for yourself?

Well maybe you can! Many people are amazed when they discover the tremendous amount of professional experience and specialized knowledge they've accumulated — experience and knowledge that others will gladly pay for. Literally thousands of people who made that discovery are now prospering as **independent consultants**.

The way to begin is by reading *How to Become a Successful Consultant in Your Own Field*, by Hubert Bermont.

Clear, straightforward, packed with solid information and advice, this authoritative manual tells you everything you need to know to establish your own independent consulting practice. Here's a sampling of the contents:

- What does it take to be a successful consultant? (See Chapter 1.)
- How to get started. (See Chapter 3.)
- How to operate your business — a collection of "tricks of the trade." (See Chapter 5.)
- What to charge your clients — plus five helpful rules on fees. (See Chapter 6.)
- Why you should **never** work on a contingency (speculative) basis. (See Chapter 7.)
- Ingenious ways to promote yourself — and make people want your services. (See Chapter 9.)
- Contracts: why you should **avoid** them at all costs. (See Chapter 10.)
- Just what do consultants do all day? (See Chapter 11.)
- How to market your ideas. (See Chapter 11.)
- Why you'll never have to worry about competition. (See Chapter 13.)
- And much more!

Perhaps no one is better qualified to have written this book than Hubert Bermont. He has served as consultant to more than 70 major corporations and trade associations, including the U.S. Chamber of Commerce, McGraw-Hill, the Electronic Industries Association, Evelyn Wood Reading Dynamics and the Smithsonian Institution. Yet he made the decision to become a consultant only after being fired from an executive position at the age of 43. You'll learn first-hand how he did it — and how **you** can do it, too!

How to Become a Successful Consultant in Your Own Field is just $20 (tax-deductible if you use it for business purposes), and you're fully protected by this **unconditional money-back guarantee**: Keep the book for three weeks. If you're dissatisfied with it for any reason whatever, simply return it and **every penny of your $20 will be promptly refunded** — no questions asked!

How many times have you told yourself that you're not getting anywhere — that it's time to think seriously about a major change in your career? **Don't put it off another day!** Clip and mail the coupon now!

Enclosed is my check or money order for $20. Rush me, postpaid, *How to Become a Successful Consultant in Your Own Field*, by Hubert Bermont. I understand that I have the right to return the book within three weeks for a complete refund if I'm in any way unhappy with it.

Name ____

Address ____

City ____ State ____ Zip ____

The Consultant's Library

Dept. **CT**, 815 Fifteenth St. N.W., Washington, D.C. 20005

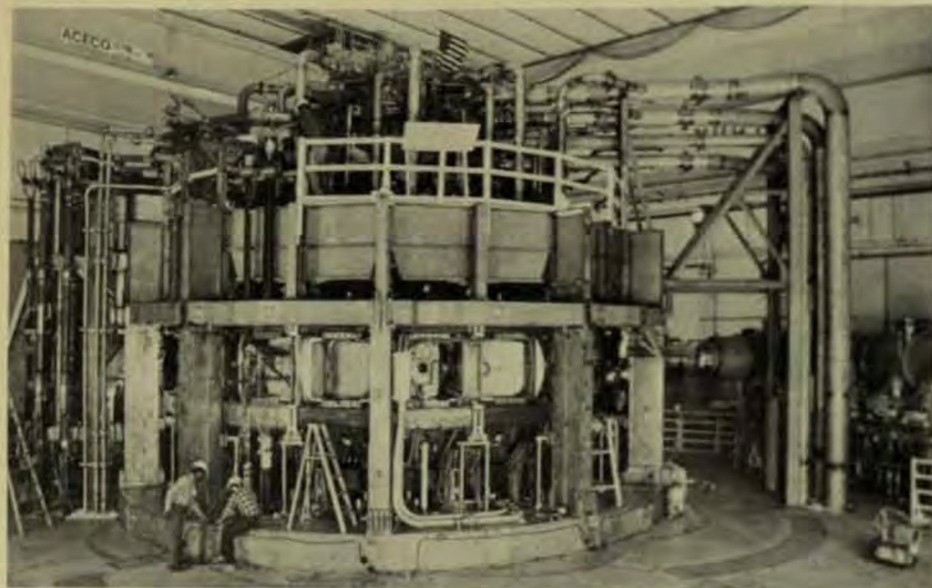

Fig 2 The Tokamak is 25.3 ft across, 7.9 ft high, and weighs 80 tons.

minals that are planned.

The on-line terminal operating stations consist of an intelligent alphanumeric terminal, a CAMAC-interfaced high-resolution graphics display, and three video monitors driven by raster-graphics CAMAC modules—two for comfort displays and one status monitor that provides page displays of monitor values and device statuses.

A number of data structures are generated under the data management system for use by the application programs. These include:

- The machine-state database, which contains high-level machine parameters for Tokamak operations and which are accessed by Tokamak's setting of tasks prior to each pulse.
- The diagnostic-settings database, which provides information needed to set up parameters in the measuring devices.
- The event-system database, which provides a record of the evolution of tasks, such as the arming of Tokamak systems or generation of summary waveforms.
- The machine-pulse database, which provides a store for specific prepulse and pulse snapshots, characteristic scalar data, and summary wave forms.

Still another major piece in the system is a program-control supervisory program to coordinate the application programs operating in the central-system computers.

Program control is required even though each computer is equipped with the MPX-32 operating system because the individual copies of MPX-32 operate only within the confines of their own CPUs and are ignorant of one another's existence. Thus, a higher-level supervisor is needed in order to stage operations involving cooperation between computers.

Program control consists of several somewhat independent pieces to handle intercomputer cooperation that involves the console operating stations, terminal operating stations, and the monitors for event-driven programs and for system-status and error conditions.

COUNTDOWN TO FUSION

The entire on-line CICADA system operates in an integrated fashion during and immediately after a 5-min. countdown that encompasses each pulse of the reactor. This countdown is an event-driven system that governs the operation of the Tokamak, the creation of the plasma, and its pulsing with ohmic heating and neutral-beam injections.

Event-driven tasks include the start of cycle, the arming of physical systems, the start of discharge, the acquisition of data, and the activation of other tasks that set cyclic status flags. The programmed sequencing of hardware and software components during countdown is achieved by a combination of an encoded facility clock system and the software program control system.

The facility clock consists of a 1-MHz biphase clock, Manchester-encoded with ASCII characters. The clock signal is distributed along the CAMAC highway to all crates, which contain timing modules that recognize codes and activate a countdown timer. This, in turn, issues a trigger pulse at the preprogrammed times that can activate timebase generators, timed gates, timing and sequencing modules, and timed-interval counters.

The program control system provides scheduling of tasks in relation to facility clock events and to software events such as input-data availability, alarm conditions, faults, or completion of other software processes.

The countdown begins with the start of count and is terminated 5 min. later with the end of count that succeeds the pulse. During the pulse, diagnostic data is acquired; after the pulse, more event-driven tasks generate diagnostic displays that show the results achieved. These are transmitted throughout the entire facility over the 70-channel TV comfort displays.

For approximately 4.5 min. after the start of count, CICADA analyzes data coming from the previous pulse and decides how to structure the upcoming pulse. The parameters for this pulse are posted in shared memory.

At 4.5 min. into the countdown, the start of freeze begins, freezing all interactive setting of parameters. From this moment, nothing can be adjusted manually. The directions posted in the machine state in shared memory are now automatically validated, and if they meet approval the system goes into the start sequence.

During this state, all parameters that have been predefined for the pulse are automatically set up and systems are armed during a sequence of events that set off the series of tasks in each computer.

Now a prepulse check, or snapshot of all control and monitoring points, is stored in shared memory in the machine-pulse database, an event that sets off reading of critical monitors and a check for required states. A failure would inhibit discharge and reprogram the power supplies and other systems.

If the check passes muster, the system goes into start of pulse. This event starts up the power supplies that energize the magnetic fields. After the specified currents are attained to control the magnetic fields, a start-of-discharge event is generated. This event controls the plasma-current sequences and the neutral-beam-injection process.

After 1 s, end of plasma is declared, shutting down the plasma systems and stopping the acquisition of plasma data but starting the reduction of data and the creation of about 100 waveforms that contain 500 to 800 words each. These graphs contain five predefined areas of interest that may be magnified.

The waveform data is stored in shared memory, where it is available to all user application programs as well as the diagnostic display program in the console computer. This program utilizes predefined templates on which the wave data is plotted by the comfort display system, which transmits the graphics over TV channels immediately after the pulse.

EXPANSION AND EVOLUTION

There are many specifications in putting together the hardware and software for this versatile and powerful system. Important, of course, was the high throughput that the superminis provide, working in conjunction with a high-speed bus that allows 26.67 Mbytes/s of internal data transfer. These performance characteristics

Toward Controlled Fusion

To appreciate fully the role of the CICADA computer system in controlling the Tokamak Fusion Test Reactor (TFTR), you must understand the concept of nuclear fusion and how it will be achieved on the reactor. This is the first experimental device capable of achieving a break-even fusion demonstration, analogous to the historic fission power demonstration carried out by Enrico Fermi under Stagg Field at the University of Chicago in December of 1942.

The theoretical basis for both nuclear fission and fusion became apparent in the 1930s after the discovery that, in general, the mass of an element can be converted into energy following Einstein's famed equation, $E = mc^2$, where energy equals mass times the square of the speed of light.

The release of nuclear energy can be achieved either by the splitting of heavy elements into lighter ones (fission) or by the fusion of the lightest elements into slightly heavier ones. Both processes would result in a loss of mass and its conversion into energy, but fusion would produce by far the greatest release of energy.

In 1950, Soviet physicists conceived of taking the hydrogen atoms in a plasma—a fourth state of matter where all the atoms are ionized by tremendous heat—and confining them in a doughnut-shaped magnetic container. Here they could be confined long enough to subject them to tremendous temperatures, which would energize them sufficiently to overcome their repelling forces and allow the blinding forces in the nuclei to fuse them.

The key engineering task in such a design is to keep the ions from hitting the walls of the container, where they would lose the critical heat needed for fusion.

The vessel, named Tokamak after the Russian-language acronym for toroidal magnetic container, has been the principal focus of experiments at the Princeton Plasma Physics Laboratory and is the design principle that has shown the greatest promise in experiments over the last three decades.

In the new TFTR at Princeton, NJ, the break-even experiment will be made through the fusion of deuterium (one proton and one neutron) and tritium (two protons and two neutrons) and one neutron, the latter containing most of the released energy in the form of excitation.

Unlike some other possible fusion reactions, this one produces significant radiation, but of a much shorter half-life and in a much smaller quantity than in the case of fission. This process is being attempted first because it is far easier to achieve than other potential fusion reactions. To attain a breakeven number of fusion reactions with a deuterium-tritium plasma, it must be confined for 1 second at a temperature of 100 million degrees Celsius and at a density of 3×10^{14} atoms/cm^3 (about one-millionth the density of the atmosphere at sea level).

To accomplish this in the TFTR, formidable physical systems must be coordinated and brought to bear. The plasma must be successfully confined under a powerful toroidal magnetic field (whose lines of force are parallel to the doughnut-shaped container) at a magnetic force of up to 52 KGauss, more than 52,000 times the force of the earth's magnetic field.

There must also be a secondary poloidal (polar) magnetic field whose lines of force intersect the torus in vertical planes. Yet a third set of magnetic equilibrium fields is needed to fine tune the position of the plasma within the vacuum vessel, which measures 25.3 ft across and 7.9 ft high and weighs 80 tons.

The plasma is heated in two stages, beginning with heat generated from ohmic resistance to a current of 2.5 MA traveling in the plasma itself and capable of heating it to 20 million degrees Celsius. But the principal boost to temperatures between 50 million and 100 million degrees Celsius will come from the injection of neutral-particle beams into the plasma from four high-energy particle-beam lines that transfer their energy to the plasma through collisions.

The electrical system required to boost the plasma to these staggering temperatures must itself be formidable, capable of generating a pulse of 950 MW. To meet these requirements, electrical energy will be removed at a low rate of about 30 MW from the 13,800V electrical-utility system in the form of rotational energy in two large fly-wheels. During a pulse, the stored rotational energy will be converted by two large generator sets into the required ac electrical energy.

Formidable systems are also required for the creation of the Tokamak vacuums, cyrogenics, and the handling of radioactive tritium.

The initial TFTR experimental program this year will study the ohmic heating and compression of large-current hydrogen and deuterium plasmas. Next year, high-powered neutral-beam heating experiments will begin with deuterium plasmas. These are expected to reach approximate deuterium-tritium equivalent fusion break-even conditions toward the end of 1985. And in 1986, the process will be repeated with a demonstration with actual deuterium-tritium plasmas that achieve the break-even level of fusion power. ■

Fig 3 This console operating system color display of the Tokamak vacuum pumping system is one of hundreds of displays that allow for manual control of the Tokamak by use of a cursor and tracker ball.

were needed because the system's real-time and on-line functions require very fast processor times and because of the large amount of data transferred into and out of the computers.

As a corollary, very-high-speed intercomputer data-transfer links were also required along with the capability to provide high-performance network-linking software. A high-performance, multiported shared-memory system was also needed, along with readily expandable internal memory. And, of course, versatile I/O capability was a necessity.

But perhaps the main requirement was the ability to migrate easily to new technology to allow for expansion through additional processors. In a system designed for complex scientific experimentation, the needs change dynamically virtually every week—and even more dramatically over a longer period.

In 1976, for example, it was anticipated that CICADA would have to handle no more than 10,000 control and monitoring points. Since that time, this requirement has doubled. The system was originally planned for 12 console operating stations; now there are 20, and CICADA will eventually require 26. Also, the number of interactive terminals will almost triple in two years, from 12 to 32.

The system has been reconfigured and expanded continually over the last six years. These changes were possible because of the advent of highly compatible computer families that retain essentially the same bus, I/O controllers, interchangeable memory modules, and other features. Thus, CICADA can now be reconfigured virtually daily as needs dictate.

The operational software for the system has also migrated from the RTM system for the 32/75 to two successively more sophisticated versions of MPX-32 for the later computers. The original operational system would support only ½-Mbyte of memory for any task, while the latest MPX-32 supports up to 16 Mbytes of memory, including up to 4 Mbytes of shared memory.

It became obvious very early that the growing size of CICADA's code required larger memory. To meet these needs, it was natural to move to ever-denser memory modules, switching from core to MOS, tripling memory in each of the central-system computers from 1 to 3 Mbytes. And with the next expansion of memory in the subsystem machines, memory in each computer will have quadrupled from ¼-Mbyte to 1 Mbyte. Shared memory has also been increased drastically, from ¼-Mbyte to 2.5 Mbytes, and will be going to 3 Mbytes.

Equally important, memory was increased without a change in the physical size of the system—in fact, sometimes reducing space from two cabinets to one cabinet—simply by switching from core to ever-denser MOS modules as these become available. A 1-Mbyte module now fits in the space of a 32-Kbyte core memory.

BOOSTING COMPUTER POWER

The need for more processing power is also increasing dramatically. Originally, the on-line configuration was obtained with only three computers—one central application-program computer, a console computer, and one subsystem computer to interface with CAMAC. As CICADA's on-line needs grew, the functions were split off among additional computers. Now another diagnostic computer and a second neutral-beam computer are to be added as TFTR operations go into full gear. Because of the modularity, these can be added quickly by links to shared memory and to the operations computer.

Also being upgraded is the program-preparation computer, which is currently part of the off-line data-analysis system. With the new processor, program preparation will be linked directly to CICADA on-line, so that when programs are developed they can be pushed directly down the interbus links to CICADA machines.

With the continuing migration of processors, execution speeds have been virtually quintupled—from 0.6 MIPS (million instructions per second) with the first 32/75 to 3 MIPS with the first 32/87s. And with the prospective use of IPUs in the 32/87s, speeds are expected to more than double, to 6.6 MIPS.

RELIABILITY AND REDUNDANCY

In dealing with a complex system of many networked processors, reliability is always an important factor. The network of 14 computers, as configured in Fall 1982, met rigid standards of acceptance testing, both in demonstration tests of the capabilities of the computers and operational software and in tests incorporating CICADA applications software under simulated operating conditions.

During a 30-day performance period, the total system provided an

availability of 91% of the scheduled operating time, and a subset of 12 computers provided an availability of over 99%. The 12-computer configuration was sufficient to perform all functions of the CICADA system. To replace a failed computer with a spare, it's only necessary to switch two twinax cables to the CAMAC link and reload the spare computer disk.

HIGH-SPEED INTERCOMPUTER LINKS

Of special interest are CICADA's high-speed data links between computers. For computers talking to each other, a custom software package called the intercomputer communications system (ICS) contains a protocol for computers networked over high-speed data (HSD) links, operating at 32 Mbytes/s. ICS is a compatible extension of the MPX-32 operating system, permitting real-time interactive communications among the computers in the CICADA network.

In this configuration, the nodes consist of a computer, MPX-32 operating system, and HSD interface—and the nodes are interconnected to each other through the ICS software. Any node can address only a directly-connected remote node. The ICS software for the network provides capabilities for three major functions:

- In the area of **remote intertask communication**, the software makes it possible to send a message to a task in a remote mode and activate a task there.
- In **remote file accessing**, the system allows an operator to tranfer a file to a remote node, request a file, or delete a file remotely.
- In a **remote operator function**, ICS enables transfer of data from memory to memory, allocation or deallocation of a communication link or output buffer, and allows link-identification parameters to be obtained.

In essence, the ICS software enables parallel direct access from any one computer to the controls of the other computers.

In spite of the fact that this is a networked system, CICADA achieves a high level of integration among all its components. User application programs, for example, have controlled access to all the resources of the system—whether device control and monitoring, data acquisition and storage, or console and terminal displays and graphics.

All CICADA application programs can run in any of the central-system computers and can access the resources of any subsystem computer. By calling service routines, the programs make use of the many resources provided.

Not least, assignments and controls are accomplished automatically by the system and are transparent to the user. In this way, CICADA looks to the user like a single tremendously powerful computer. ■

Robert Daniels, *Head of CICADA/ Computer Div. at the Princeton Plasma Physics Laboratory, Princeton, NJ, is an electrical engineer who has specialized in computer systems for high-energy physics devices, including particle accelerators.*

How do you keep your Industrial Control System from turning into a House of Cards?

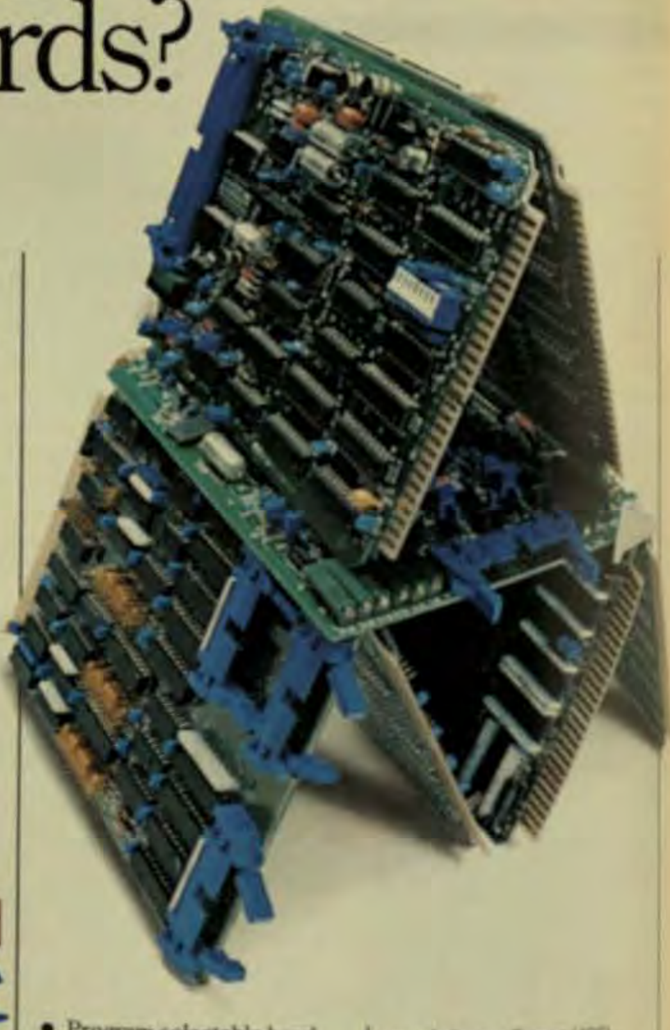

Simply specify Intersil Systems' STD BUS line.

There's no good reason why you should ever have to specify more functions than you really need. Or have to worry about control system integration problems.

Why?

Our STD BUS products take care of all that for you.

Intersil Systems STD BUS boards are designed to function as building blocks for your microcomputer-based system. You simply select the microprocessor, memory, I/O, peripheral controller, communications controller and special function cards you need. Then just plug them directly into the compact, bus-compatible card cage, in exactly the configuration you need. Add your selection of standard software and you're done.

Just think about the design flexibility and expandability that gives OEMs:

- You mix and match nearly three dozen different cards, customizing your microprocessor-based system to exactly match your application requirements.
- With Intersil Systems' STD BUS products, you specify either a custom or standard CPU that's just right for you—8085, Z-80 … enabling you to take advantage of readily available applications software and higher level languages.
- STD BUS allows concentration on control application problems, not integration problems. That saves you time. And makes your prices more competitive.
- Center-ground, multi-layer (4) board design, 100% 24 hour active burn-in and the highest MTBF rate let you design with proven quality products.
- There's even a development system available, off-the-shelf, that will enable you to develop, test and de-bug your system. And there's a designer's kit to fulfill your prototyping needs.

(Shown here)

And more configurability means more capabilities for systems users:

Intersil's breadth of products and features allows immense design flexibility for process control applications. For example:

- A full range of custom and networking interfaces.
- A complete line of industrial I/O.
- Easily upgradable modular configuration capabilities.
- Program selectable baud rates for synch/asynch communications.
- Small, dense packaging for space optimization.
- Complete customization services for any need a standard product can't fill.

Isn't it time you found out how easy it is to get your house in order? For systems needs just call us at (408) 743-4442. Or for off-the-shelf boards call your nearest Intersil distributor: ADM, Alliance, Anthem, Arrow, R.A.E. Industrial Electronics, Ltd., Quality Components, Schweber or Future Electronics. We'll help you get off the drawing board and on the Bus. On time. And on budget.

Intersil Systems

1275 Hammerwood Avenue, Sunnyvale, CA 94089 • (408) 743-4442

Networked Interactive Workstations Bring CAD Power to Plant Design

Dedicated software and responsive 3-D workstations with very large databases allow CAD technologies to be applied to the complexities of process-plant design.

by **James F. Callan,** Cadtrak Corp.

The problems involved in the design of large industrial facilities have resisted the application of computer-aided design (CAD) technology. However, a combination of dedicated software, high-resolution 3-dimensional graphics, and powerful, single-user workstations that can be networked now provides the benefits of data processing to this complex design task.

In this combination, clustered workstations can transmit data directly between one another, and each workstation communicates with the host computer. The host maintains a central database that describes the 3-D object being designed.

Each workstation takes on a design task, unencumbered by concerns about the amount of data to be handled locally or the expanse of that data in real-world coordinates. And each workstation is capable of off-line design tasks based on a copy of relevant portions of the master database. The workstation's database is augmented by design operations performed locally. At the end of an off-line design session, a file containing all changes to that database is sent back to the host to update its master database.

This approach can be applied to major projects that call for use of very large databases, many parts of which are undergoing design simultaneously (Fig 1). There is effectively no limit to the number of workstations that can be networked, the size of the central database, or the amount of data available for modification at a workstation.

Furthermore, tasks appropriate to a large computer, such as those calling for a heavy computational capability or very large storage, are performed in the host, and tasks appropriate to interactive workstations—those calling for an operator to wait for a response—are performed locally at the workstation. Each operator sees an immediate response to his commands, and time-consuming tasks are performed without an operator's presence.

PROCESS PLANT DESIGN

Cadtrak began by tackling the most difficult CAD application it could find—process-plant design. A single process plant—chemical or petroleum, for example—can occupy several square miles. It can contain up to 20,000 pipelines, each containing 30 or 40 distinct components in addition to the numerous structural supports, vessels, and electrical systems required.

At the same time, it's often necessary to maintain a resolution of 1 mm within the plant design in all three dimensions. Additionally, some plants, such as nuclear power facilities, place a premium on compactness and need to have all components placed in as dense a physical volume as possible.

Of course, there can be no interference between physical objects in the plant site, and allowance must be made for maintenance and operational access, and for paths to get components in and out during both construction and maintenance. A number of standards exist in this engineering task that govern the alignment of components, their connectivity, and the legality of components put together in combinations. Adherence to these standards imposes additional constraints upon the design.

Challenging as the design of such a plant is, it nonetheless must often be completed on very limited time schedules because the difference of only a few days in the startup of a plant can make a difference of millions of dollars to the plant owner. Consequently, efforts are significantly overlapped—hundreds of designers may be working simultaneously on a plant design.

These conflicting considerations lead to imperfect designs that contain errors often not discovered until the construction stage, when the required parts are not only unavailable, but may be half way around the world from the source. These errors can lead to substantial added expense—typically 2 to 5% of a plant's total installed cost—and delay the commencement of plant operation for many of those million-dollar days. And since every plant is unique, the design process must be started anew each time.

OUTMODED SOLUTIONS

For decades, firms specializing in the design of process plants have used plastic modeling for their design work—a physical scale model of the real plant, built to a scale on the order of ⅜ in. to the foot. These models are costly, fragile, awkward, and imprecise.

Over the last decade, turnkey CAD systems have become the norm for the design of many mechanical components. While CAD technology meets over 50% of the market in such design applications as integrated circuits, it has penetrated less than 5% of the market for process-plant design.

In turnkey systems, minicomputers typically serve as the hub of several terminals, each depending on the mini for all or most of its computation and storage facilities. This approach was necessary when these systems were first designed because the cost of terminal-resident storage and computing was prohibitive, but it imposes difficulties that are bothersome in some design functions and actually prohibitive in plant design.

A number of other problems with conventional CAD systems have prevented process-plant designers from taking full advantage of modern computer technology. The problems include:

- **Storage Limitations.** The database for a mechanical part, an integrated circuit, or a map can generally be held within the readily available storage limits of most minicomputers. In these applications, an individual portion can often be designed quite effectively independently of the rest. A mechanical part, for example, can be designed by itself even if it's part of a large mechanical system whose other components are temporarily ignored.

  In a large process plant, on the other hand, the impact of one designer's actions can severely affect another designer's operation. Pipelines extend the entire length of a plant and can be part of dozens of designers' areas. Because plants don't break down conveniently into easily manageable modules, on the few occasions where CAD technology has been used successfully in this area it has been restricted to very small, isolated portions of the design. The interface between regions designed with CAD systems is still a clumsy manual interface.
- **Simultaneous Designer Operation.** The tight time schedules required for many plant designs mean that more than four or five stations have to work on a database at one time. But, since a minicomputer can handle only four or five stations, additional minis are required, each with its own copy of the master database. As each group modifies its copy of the database, the several new and different masters that are created must be merged.
- **Proximity Requirements.** Accommodating physical distances between designers represents another problem. Since conventional CAD terminals depend so intimately on their host they must be very close to it physically. A large plant is often designed by many branch offices and subcontractors spread across the country or the world.
- **Expanse of Viewable Database.** At the terminals themselves, a viewing window into the database is typically restricted to the data held in the terminal's refresh memory. Since that memory cannot economically be much bigger than a screen's worth of data, a change in the viewing window entails a lengthy procedure for the user to read, compute, and load new memory contents. A plant designer wants fast access to highly detailed views of widely separated areas.
- **Picking the Wrong Element.** The process of pointing to displayed entities remains clumsy in many CAD implementations. Seconds often elapse while the system figures out which entity a designer has pointed to, and in crowded designs like those of dense nuclear plants, if that isn't the entity intended, the operator not only waits for his operation to be performed on the wrong item, he also has to endure a lengthy recovery procedure.

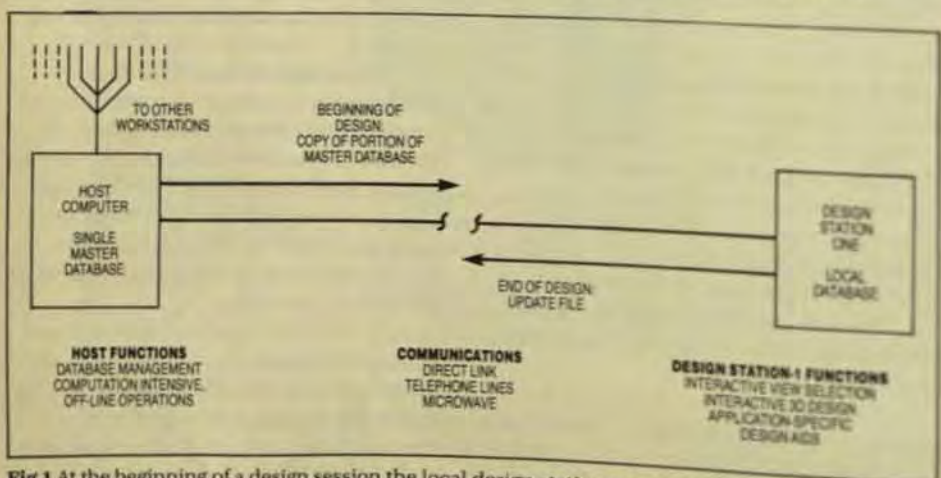

Fig 1 At the beginning of a design session the local design station requests a copy from the master database; at the end of the session all work performed is sent back to the host to update the master.

- **Unpredictable Response Time.** When several dumb terminals all tax a single minicomputer, the response time depends not only on how complicated an action is, but on what else the minicomputer has to do for the other operators. Not only is the response slow, it's unpredictable, and an operator is frustrated by not knowing how long it will take the computer to respond once he initiates an action.
- **Environmental Limitations.** Many CAD systems require a computer-room environment for their workstation—low ambient light, special power, tight temperature and humidity control, raised floors, and other considerations that designers find uncomfortable.
- **Screen-Image Quality.** While they have largely abandoned storage tubes, many CAD systems have substituted interlaced raster displays, which draw odd- and even-numbered scan lines alternately, with a refresh only 30 times per second. Interlaced displays are easier to build, and until very recently were the only way to support the high-resolution 1000 × 1000-pixel color displays needed to show tight detail.

  However, the real refresh rate of any interlaced display is only 30 Hz or less, which presents an unpleasant choice between the use of a long-persistence phosphor that smears when it moves, or eye-straining flicker. Most systems compromise by using phosphors that both smear and flicker.

While such problems are found across the spectrum of CAD applications, they've proven particularly nettlesome in the conservative plant-engineering industry, where high technology is still widely regarded as suspect and foreign.

The only substantial role that CAD systems have played in plant engineering is in such 2-dimensional applications as annotations for drawings and the preparation of piping and instrumentation diagrams. No conventional CAD system has been applied in a large-scale way to 3-D plant design.

THE PLANT DESIGN MANAGEMENT SYSTEM

A comprehensive software package, the Plant Design Management System (PDMS), has been created for execution on supermini or mainframe computers. PDMS has the capacity to hold and maintain databases as large and as rich as they need to be for even the most complex of plants.

Included are facilities to detect interferences not only between hard physical objects but also with maintenance access-ways and other volumes of the plant that need to be clear. It even includes the rules and standards of plant design, primarily in the area of pipework—the most difficult and taxing aspect. Since PDMS can also completely validate a design prior to construction, a plant designed with this system is more likely to be built on time and on budget.

PDMS does circumvent many of the problems that conventional CAD system architecture has had in addressing plant design. For example, its 32-bit floating database permits plants miles long to be held to 1 mm of resolution. It has a single, multiple-user database that can support terminals placed nearly anywhere, providing access via telecommunications.

But despite the enormous advantages PDMS has over conventional approaches to computerized plant design, it hasn't achieved the widespread usage it deserves because its I/O is comparatively primitive—via simple dumb-terminal keyboard entry (and, recently, digitizer entry). Pictorial feedback was provided only in the form of storage tubes and their equivalents.

Response was often unbearably slow, and no use was made of color, sophisticated interactive devices, picture dynamics, or any of the other modern computer-graphics techniques that conventional CAD systems mastered long ago. Many potential users simply refused to try the system at all. Others found its possible benefits great enough to overcome its

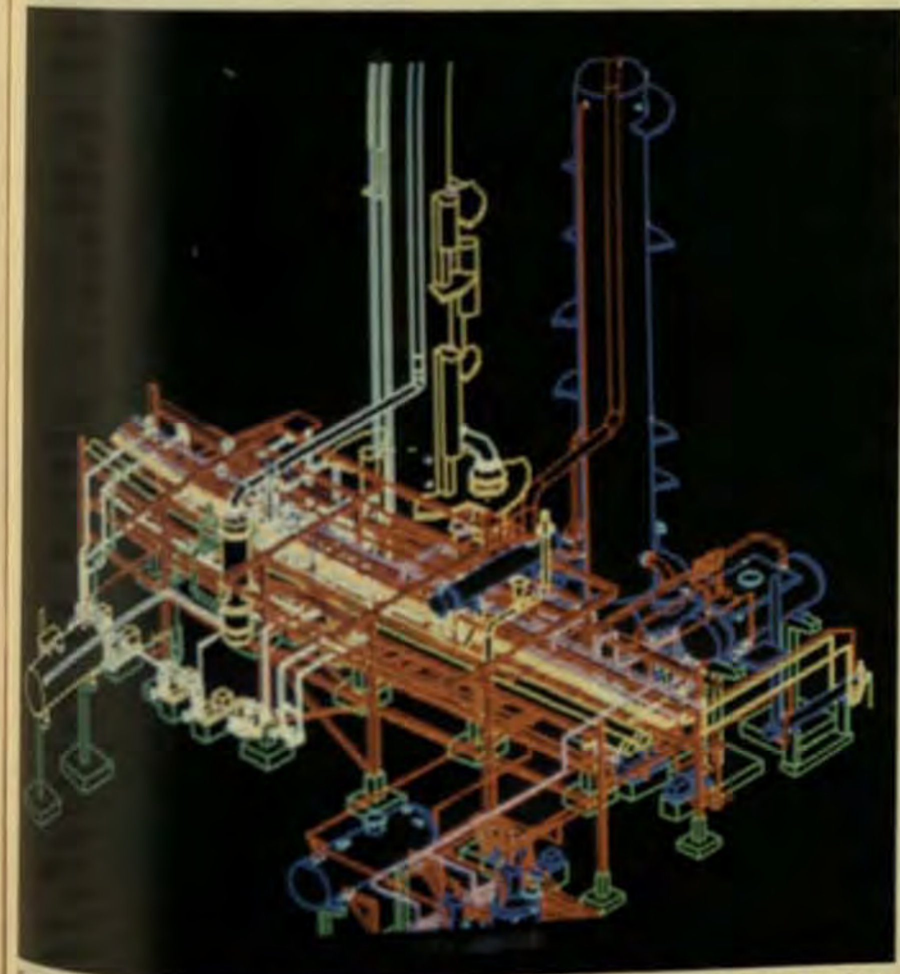

Fig 2 Views are selectively depicted in full color, and can be interactively panned, zoomed, and exchanged for other views.

lack of user-friendliness, but wound up using the system just to verify designs, because they felt it was too slow and unfriendly for design creation.

A CAD WORKSTATION

Cadtrak's Design Station One (DS-1), an engineering workstation that can serve as a node in a network of design stations talking to PDMS, extracts local databases from the master database resident within PDMS and communicates them to local storage within the workstation.

That database, together with hidden-line views of it prepared by PDMS, are stored on the local Winchester disk of a DS-1 when a design session begins. Thereafter, these views are selectively brought onto the display, where they are depicted in full color (Fig 2) and can be panned, zoomed, and exchanged for other views.

The workstation's interactive devices add or change 3-D components that represent pieces of pipework in the local 3-D database. At the end of a session, the workstation updates the PDMS master with all additions and changes by telling PDMS how to update its master—not by sending a revised version of the original. DS-1 has the local storage and processing power required to be a user-friendly front-end to PDMS and a standalone station for pipework design.

The portion of the database extracted from PDMS is still in 32-bit floating point form to permit a full-scale numerical model of the plant to be housed locally. And since the database and drawings are virtualized, the only real limit on the amount of information that can be stored locally is the size of the local disk.

This workstation also addresses other weaknesses of current CAD terminals. A virtual-raster technique, for example, permits immediate access to any part of the database at any scale. The viewing window can be panned over the full expanse of the database and zooming produces high-resolution closeups at any scale.

In addition, a mark symbol is depicted over the nearest item to the cursor, and in a different color so that the designer can easily identify any item on which he might choose to operate. As the cursor slides over a view, the mark symbol dances from item to item in real time. And since this workstation is for single users, the designer no longer has to share processing power with others, and the response time is predictable.

This workstation also contains local software to assure that only legal components are placed on a pipeline, that every component on a line is aligned with the others, and that components are connected precisely. Multiple-viewport capability provides plan and elevation views simultaneously so that the designer can see his 3-D design changes as he makes them.

With this system, a network of workstations can be put together, all accessing the same master database to address complex design operations. Everyone sees the most recent state of the design. Local-area networking permits two nearby workstations to communicate directly with one another without affecting the host, but workstations can be located virtually anywhere.

FUTURE APPLICATIONS

This workstation can be applied to any graphics application. Not only does it have such capabilities as instant access at high resolution to any part of an arbitrarily large database, it allows distributed processing to be applied to CAD applications. Many of these can benefit from a configuration in which a substantial database is held in a computer and many workstations can take copies of that database for local work and return only a file of changes to the host.

In the future, other difficult applications no doubt will be addressed. Major map-making applications, major electrical-system designs, and major mechanical-system designs can all benefit by an approach that sorts functions into those that should be performed by a big computer and those that should be done by highly responsive intelligent workstations. ■

James F. Callan, *director of marketing for Cadtrak Corp., is responsible for its market research and planning, and for communications and promotional programs for Cadtrak's products. He has a BA and an MA from the State Univ. of New York, Buffalo.*

BUSINESS OPPORTUNITY

Exclusive franchise in America's most profitable and dynamic industry is being offered for the first time in your area. International company will place qualified individual in "Turn Key" business, train people, provide inventory, finance your customers, and pay you thousands of dollars "up front" on orders where your customers pay only on future energy savings. Existing customers of our franchisees reads like "Who's Who" of Fortune 500.

If you qualify, you will be flown to Los Angeles for a tour of installations and personal interview. Minimum investment of $29,500 cash required. Call president at 1-800-323-6556, ext. R-137.

FEDERAL ENERGY SYSTEMS, INC.

Suite 200, 336 N. Foothill Road, Beverly Hills, CA 90210

THIS IS NOT AN OFFERING TO SELL

Advertiser's Index

This advertiser's index is an additional service provided by CTR. The Publisher assumes no liability for omissions or errors.

SAVE $10! ... and own *your own* subscription ... TODAY!

Please enter my **new** subscription to COMPUTER TECHNOLOGY REVIEW for one year. Since you offer a multiple payment option, here's my choice to pay:

☐ I am only paying $40.00 by enclosing my check for one year subscription (20% off the regular rate of $50.00).

☐ Charge my: VISA # ______ MASTER CARD # ______ AMERICAN EXPRESS # ______ Expiration Date ______

☐ Bill me
☐ Bill my company } *Regular full rate of $50.00 applies*
Add $30.00 for mail outside the U.S.A.

Signature ______ Date ______
(Please print or affix peel-off label from front cover)
Name ______
Title ______ Div./Dept. ______
Company ______
Address ______
City ______ State/Country ______ Zip Code ______

My job function is: (check only one)
a. ☐ Corporate Management
b. ☐ Marketing/Sales Management
c. ☐ Engineering Management
d. ☐ Systems/Hardware Design Engineer
e. ☐ Systems Integrator
f. ☐ Software Engineer/Programmer
g. ☐ DP Management
h. ☐ Other (specify) ______

My company is: (check only one)
a. ☐ Mfr. of computer-based systems
b. ☐ Systems House/Integrator
c. ☐ Hardware OEM/Distributor/Dealer
d. ☐ Software House/Packager
e. ☐ Mfr. of Computers/Peripherals
f. ☐ Volume End User of Computers/Peripherals
g. ☐ Consultant/Educator
h. ☐ Government/Military
i. ☐ Other (specify) ______

I buy/specify/influence the purchase of: (check all that apply)
a. ☐ Minicomputers
b. ☐ Microcomputers
c. ☐ CRT Terminals
d. ☐ Intelligent Terminals
e. ☐ Graphic Terminals
f. ☐ Peripheral Controllers
g. ☐ Printers
h. ☐ Plotters/Digitizers
i. ☐ Tape Drives
j. ☐ Floppy Disk Drives
k. ☐ Hard Disk Drives
l. ☐ Add-In/Add-On Memories
m. ☐ Modems and Data Comm Equipment
n. ☐ Software Packages
o. ☐ Word Processing/Small Bus. Systems
p. ☐ Computer Supplies (Disk Packs, Diskettes, Forms etc.)

The following mini/microcomputers are installed at this company:

	Manufacturer	Model	Quantity
1.			
2.			
3.			

SAVE $10! ... and own *your own* subscription ... TODAY!

Please enter my **new** subscription to COMPUTER TECHNOLOGY REVIEW for one year. Since you offer a multiple payment option, here's my choice to pay:

☐ I am only paying $40.00 by enclosing my check for one year subscription (20% off the regular rate of $50.00).

☐ Charge my: VISA # ______ MASTER CARD # ______ AMERICAN EXPRESS # ______ Expiration Date ______

☐ Bill me
☐ Bill my company } *Regular full rate of $50.00 applies*
Add $30.00 for mail outside the U.S.A.

Signature ______ Date ______
(Please print or affix peel-off label from front cover)
Name ______
Title ______ Div./Dept. ______
Company ______
Address ______
City ______ State/Country ______ Zip Code ______

My job function is: (check only one)
a. ☐ Corporate Management
b. ☐ Marketing/Sales Management
c. ☐ Engineering Management
d. ☐ Systems/Hardware Design Engineer
e. ☐ Systems Integrator
f. ☐ Software Engineer/Programmer
g. ☐ DP Management
h. ☐ Other (specify) ______

My company is: (check only one)
a. ☐ Mfr. of computer-based systems
b. ☐ Systems House/Integrator
c. ☐ Hardware OEM/Distributor/Dealer
d. ☐ Software House/Packager
e. ☐ Mfr. of Computers/Peripherals
f. ☐ Volume End User of Computers/Peripherals
g. ☐ Consultant/Educator
h. ☐ Government/Military
i. ☐ Other (specify) ______

I buy/specify/influence the purchase of: (check all that apply)
a. ☐ Minicomputers
b. ☐ Microcomputers
c. ☐ CRT Terminals
d. ☐ Intelligent Terminals
e. ☐ Graphic Terminals
f. ☐ Peripheral Controllers
g. ☐ Printers
h. ☐ Plotters/Digitizers
i. ☐ Tape Drives
j. ☐ Floppy Disk Drives
k. ☐ Hard Disk Drives
l. ☐ Add-In/Add-On Memories
m. ☐ Modems and Data Comm Equipment
n. ☐ Software Packages
o. ☐ Word Processing/Small Bus. Systems
p. ☐ Computer Supplies (Disk Packs, Diskettes, Forms etc.)

The following mini/microcomputers are installed at this company:
Manufacturer / Model / Quantity
1. ______
2. ______
3. ______

SAVE $10! ... and own *your own* subscription ... TODAY!

Please enter my **new** subscription to COMPUTER TECHNOLOGY REVIEW for one year. Since you offer a multiple payment option, here's my choice to pay:

☐ I am only paying $40.00 by enclosing my check for one year subscription (20% off the regular rate of $50.00).

☐ Charge my: VISA # ______ MASTER CARD # ______ AMERICAN EXPRESS # ______ Expiration Date ______

☐ Bill me
☐ Bill my company } *Regular full rate of $50.00 applies*
Add $30.00 for mail outside the U.S.A.

Signature ______ Date ______
(Please print or affix peel-off label from front cover)
Name ______
Title ______ Div./Dept. ______
Company ______
Address ______
City ______ State/Country ______ Zip Code ______

My job function is: (check only one)
a. ☐ Corporate Management
b. ☐ Marketing/Sales Management
c. ☐ Engineering Management
d. ☐ Systems/Hardware Design Engineer
e. ☐ Systems Integrator
f. ☐ Software Engineer/Programmer
g. ☐ DP Management
h. ☐ Other (specify) ______

My company is: (check only one)
a. ☐ Mfr. of computer-based systems
b. ☐ Systems House/Integrator
c. ☐ Hardware OEM/Distributor/Dealer
d. ☐ Software House/Packager
e. ☐ Mfr. of Computers/Peripherals
f. ☐ Volume End User of Computers/Peripherals
g. ☐ Consultant/Educator
h. ☐ Government/Military
i. ☐ Other (specify) ______

I buy/specify/influence the purchase of: (check all that apply)
a. ☐ Minicomputers
b. ☐ Microcomputers
c. ☐ CRT Terminals
d. ☐ Intelligent Terminals
e. ☐ Graphic Terminals
f. ☐ Peripheral Controllers
g. ☐ Printers
h. ☐ Plotters/Digitizers
i. ☐ Tape Drives
j. ☐ Floppy Disk Drives
k. ☐ Hard Disk Drives
l. ☐ Add-In/Add-On Memories
m. ☐ Modems and Data Comm Equipment
n. ☐ Software Packages
o. ☐ Word Processing/Small Bus. Systems
p. ☐ Computer Supplies (Disk Packs, Diskettes, Forms etc.)

The following mini/microcomputers are installed at this company:
Manufacturer / Model / Quantity
1. ______
2. ______
3. ______

IF YOU'RE GOING TO MAKE IT DEC-COMPATIBLE, IT BETTER BE BETTER.

Some people think using anything other than Digital disk systems with Digital CPUs is nothing but trouble.

We think differently.

We know there's no sense making something DEC®-compatible unless you can make it better. Especially in the three critical areas of compatibility, reliability, and service.

That's why, while our floppy systems, Winchesters, and controllers start out being fully DEC-compatible, it's our advanced technology and design techniques that take them beyond mere compatibility.

What's more, they're more reliable. We use Winchesters, which are inherently reliable, the best floppy drives made, significantly fewer parts, and extensive burn-in procedures to make sure.

Even so, all this would mean nothing without better service and support. That's why we offer DSD exclusives like on-board self-diagnostics. A hotline to our remote diagnostic technicians. And a twenty-four hour or less module swap that just might change the way you think about the need for conventional service. Not to mention its cost.

We also offer one of the most comprehensive—yet inexpensive—extended service agreements in the industry.

Finally, all these products are priced economically, making them that much more sensible to own.

We do all this because we want you to buy our disk systems with as much confidence as you buy Digital's. And after all, why shouldn't you? We not only make them DEC-compatible, we make them better.

For complete information about our floppy systems, Winchesters or controllers, just contact us.

Corporate Headquarters: 2241 Lundy Avenue, San Jose, CA 95131. Eastern Region Sales and Service: Norwood, MA, (617) 769-7620. Central Region Sales: Dallas, TX, (214) 980-4884. Western Region Sales: Santa Clara, CA, (408) 727-3163.

DATA SYSTEMS DESIGN

INTERNATIONAL SALES: Australia 03/544 3444; Belgium and Luxembourg 02/7209038; Canada 416/625 1907; Denmark 02/63 22 33; Finland 90/88 50 11; France 03/411 5454; Hong Kong and Peoples Republic of China 03/696231; Israel 52-32444; Italy 02/4047648; Japan, Osaka 06/323 1707, Tokyo 03/345 1411; Netherlands 02977-22456; New Zealand 04/693 008; Norway 02/78 94 60; Singapore, Malaysia, and Indonesia 2241077; Spain 01/433 2412; Sweden 08/38 03 70; Switzerland 01/741 41 11; United Kingdom 7073/34774; West Germany and Austria 089/1204-0; Yugoslavia 61/263 261

DEC® is a registered trademark of Digital Equipment Corporation.

Architecture Determines Cost and Reliability of Fault-Tolerant Design

Different architectures are available to achieve fault-tolerant computing. Cost-effectiveness and reliability are now available for commercial environments.

by **Dennis McEvoy and Sandra Metz,** Tandem Computers Inc.

Fault tolerance, or the ability of computers to recover automatically from failures and continue processing, is becoming a critical element in systems used in an on-line, interactive mode. There are several fault-tolerant system architectures and configurations available today, each optimized for the needs of different high-reliability markets. This article will discuss several ways to achieve fault tolerance in computing systems, and the relative merits of each.

A classical system architecture typically consists of a centralized processor, an I/O channel, and controllers to manage such peripherals as terminals and disks. Because a failure in any one of these components can take the entire system off-line, this configuration is acceptable only for applications that don't require continuous system availability.

Applications that demand continuous availability typically require some kind of fault-tolerant computing system. Such a design must strategically duplicate components in an efficient and cost-effective manner. The methods of duplication—and the resulting levels of fault tolerance—vary dramatically and are suitable for a wide range of applications.

SWITCHED BACKUP

A very basic approach to fault-tolerant computing is to duplicate the central-processing unit. Fig 1 shows a dual-processor system in which one processor acts as a backup for the other. The processors are connected through a switch, and either one can communicate with the other components to control the entire system.

This standby configuration is known as switched backup and provides the benefit of fault tolerance in the case of a single processor failure. Switched backup is one of the most commonly used configurations for large, on-line mainframe applications.

Several problems are associated with switched backup. One involves every company's biggest concern—dollars. Since the second processor is a backup, it remains idle until the first processor fails. If the backup happens to be a $2 million mainframe, a fairly expensive piece of hardware will be idle until a failure occurs in the on-line system.

However, to gain better return on your investment, you can use the backup system for program development while the on-line system is still functional. In that case, however, when the on-line processor fails, the development work on the backup processor must be halted so that the backup system can be switched on-line.

Another problem with a switched backup is that it provides fault tolerance for only one situation—a single processor failure. An I/O-channel failure between the bus switch and disk controller, for example, would quickly turn both $2 million mainframes castors up. The bus switch itself is another single point of system failure.

Architecture Determines Cost and Reliability of Fault-Tolerant Design

Different architectures are available to achieve fault-tolerant computing. Cost-effectiveness and reliability are now available for commercial environments.

by **Dennis McEvoy and Sandra Metz,**
Tandem Computers Inc.

Fault tolerance, or the ability of computers to recover automatically from failures and continue processing, is becoming a critical element in systems used in an on-line, interactive mode. There are several fault-tolerant system architectures and configurations available today, each optimized for the needs of different high-reliability markets. This article will discuss several ways to achieve fault tolerance in computing systems, and the relative merits of each.

A classical system architecture typically consists of a centralized processor, an I/O channel, and controllers to manage such peripherals as terminals and disks. Because a failure in any one of these components can take the entire system off-line, this configuration is acceptable only for applications that don't require continuous system availability.

Applications that demand continuous availability typically require some kind of fault-tolerant computing system. Such a design must strategically duplicate components in an efficient and cost-effective manner. The methods of duplication—and the resulting levels of fault tolerance—vary dramatically and are suitable for a wide range of applications.

SWITCHED BACKUP

A very basic approach to fault-tolerant computing is to duplicate the central-processing unit. Fig 1 shows a dual-processor system in which one processor acts as a backup for the other. The processors are connected through a switch, and either one can communicate with the other components to control the entire system.

This standby configuration is known as switched backup and provides the benefit of fault tolerance in the case of a single processor failure. Switched backup is one of the most commonly used configurations for large, on-line mainframe applications.

Several problems are associated with switched backup. One involves every company's biggest concern—dollars. Since the second processor is a backup, it remains idle until the first processor fails. If the backup happens to be a $2 million mainframe, a fairly expensive piece of hardware will be idle until a failure occurs in the on-line system.

However, to gain better return on your investment, you can use the backup system for program development while the on-line system is still functional. In that case, however, when the on-line processor fails, the development work on the backup processor must be halted so that the backup system can be switched on-line.

Another problem with a switched backup is that it provides fault tolerance for only one situation—a single processor failure. An I/O-channel failure between the bus switch and disk controller, for example, would quickly turn both $2 million mainframes castors up. The bus switch itself is another single point of system failure.

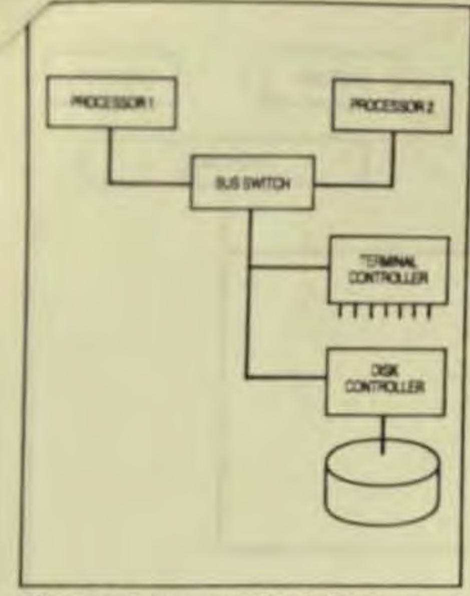

Fig 1 A very basic approach to fault-tolerant computing is switched backup, which duplicates the central-processing unit. A switch connects two processors so that either can communicate with the other system components and control the computer as a whole.

TRIPLE MODULAR REDUNDANCY

A higher level of fault tolerance is achieved with a configuration known as triple modular redundancy. Like switched backup, triple modular redundancy is a hardware-only approach, but it goes a step further by providing three copies of every piece of system hardware (Fig 2).

The three processors are lock-stepped together and simultaneously run the same instruction stream. A piece of hardware called a voter receives and compares the output from all three processors. If one output doesn't match the others, the voting hardware accepts the output from the two matching processors and passes it along to the appropriate interface.

This approach has been used successfully for many years in such areas as space research and nuclear power plants, where computer failure is unacceptable. In the commercial environment, however, triple modular redundancy poses two major problems: tripled cost and complex design. A single clock must run the entire system, and the clock must be implemented so that it can't possibly fail. This is a technically complex requirement.

These limitations weren't a stumbling block in applications like the U.S. space program, where literally billions of dollars were riding on the fact that the computer would be available 100% of the time.

COMPARISON LOGIC

A similar approach is to have four processors running the same instruction stream (Fig 3). The processor outputs are paired with each other and compared. If a mismatch occurs, the pair containing the mismatch is shut down. The system continues running with the remaining pair of processors.

Although similar to triple modular redundancy, such comparison logic has the advantage of being easier to implement than voter logic, and therefore is less expensive to design and maintain. However, the financial problems of redundant hardware all running a single instruction stream and performing identical operations still remain.

Also, comparison logic causes a good processor to be taken off-line along with the problem unit, thus wasting a resource that could be working with the system. And, again, the clock is the weakest link in the system, a single point at which the entire system could fail.

A MULTICOMPUTER SYSTEM

A different approach to fault-tolerant computing, involving a unique, integrated hardware/software architecture, was taken by Tandem Computers. Tandem's NonStop system was designed specifically for the on-line transaction-processing marketplace. This arena requires a system architecture that's not only fault-tolerant and easily expandable, but also maintains a price/performance ratio suitable for commercial transaction-processing applications.

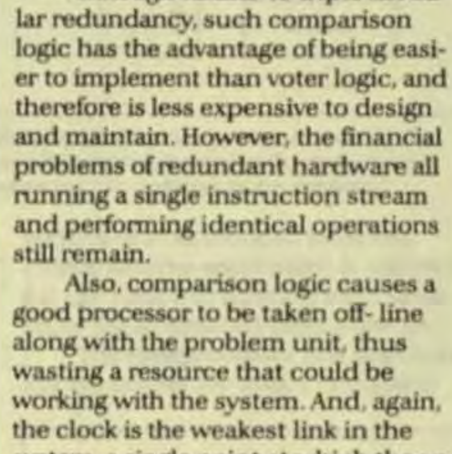

Fig 2 A higher level of fault tolerance is achieved with triple modular redundancy, which provides three copies of every piece of system hardware. The three processors are lock-stepped together and voter logic receives and compares the output from all three.

This architecture combines hardware and software so that there's no single point of failure (Fig 4). Components are duplicated but are not redundant. Each processor runs its own instruction stream, so no backup components sit idle until a failure occurs.

Fault tolerance is achieved by keeping each processor in constant communication with the others in the system via a dual high-speed interprocessor system bus. If one processor fails, its workload will automatically be absorbed by the remaining processors.

Each processor has its own I/O channel and each I/O controller is dual ported. Thus, if one processor fails, another can still use that controller. Since disks are physically duplicated, with identical data stored on each, even a disk crash can't bring down the system. The duplicate disks pay for themselves by ensuring complete system fault tolerance and improved performance. One disk performs seeks on the inner portion of its surface while the other disk performs the same seeks on its outer portion. This provides a significant performance improvement in seek times.

The software architecture is based on independent processors and independent programs within them, all of which communicate through a message system that is an integral part of the operating system. This operating system sees all programs and data transfers as communications distributed over several processors.

Programs can access any device anywhere in the system, even those not physically connected to the processor running the program. Conversely, each program is unaffected by the processor on which it runs. The operating system sees all physical resources as logical files. Only the message-routing part of the operating

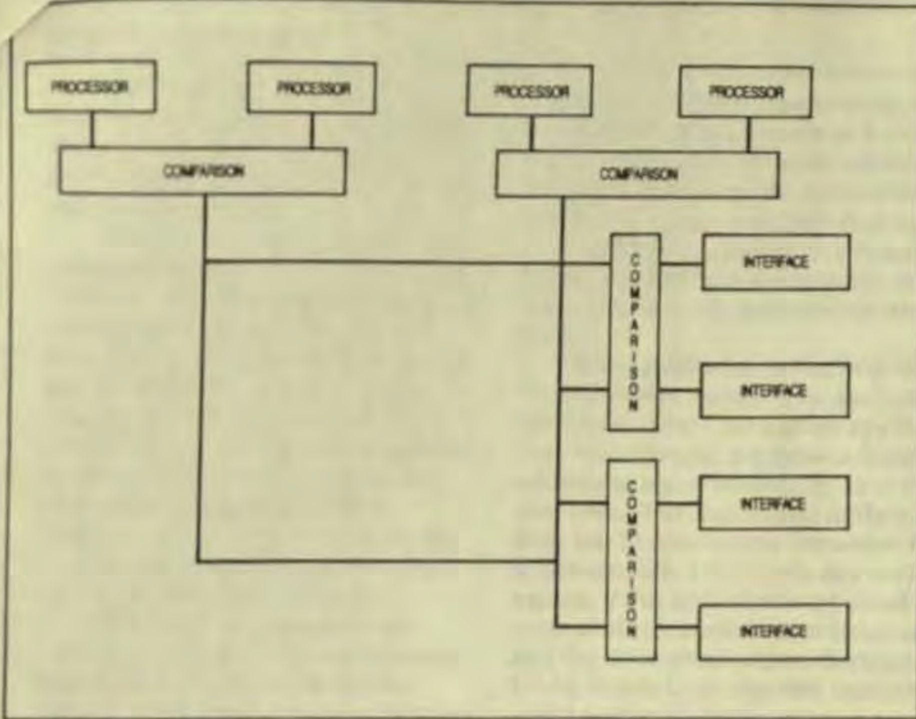

Fig 3 An approach similar to triple modular redundancy is to have four processors running the same instruction stream, with their output paired and compared. If a mismatch occurs, the pair containing the mismatch is shut down and the remaining pair continues to run the system.

system knows the geographic locations of resources, so data can be rerouted and resources dynamically reallocated during a failure.

Software fault tolerance is improved by distributing the operating system across all the processors in the system. If one processor fails, its job will be picked up by those remaining. In a system where three or four copies of the hardware are lockstepped together, a severe software bug could crash all the processors simultaneously. Past experience shows that software bugs severe enough to crash a system are typically related either to timing or to the placement of data in memory. In a NonStop system, however, a software failure will probably not crash more than one processor because the timing circumstances and exact placement of data in memory are unique to each processor.

SYSTEM EXPANSION

On-line transaction processing is characterized by the need for easy system expansion. For example, if a banking application initially involves control of 50 automated teller machines and the application is successful, the bank will typically wish to expand the automated-teller network.

The interprocessor bus system is the key to easy hardware expansion because each processor module has its own memory, I/O channel, and bus controller. To expand this system, additional processor modules are simply added to the bus. The message system keeps track of all system resources, so the system software load is automatically allocated to the appropriate hardware. The system can then be fine tuned for optimum performance.

It's possible to have a modularly expandable multiprocessor system where the processors are linked through a shared memory. A shared-memory architecture provides fast data access to all the processors and provides additional processing power. But the additional processing power obtained doesn't increase linearly.

For example, a second processor doubles your investment and provides only 1.7 to 1.8 times the power of a single CPU because of contention for the shared memory and associated shared resources. A third unit provides approximately 2.4 times the power for three times the cost. By the time the fourth processor is added, it's possible to reach a state of diminishing returns.

One way to get around the contention problem is to locate a cache memory in front of each CPU. Each

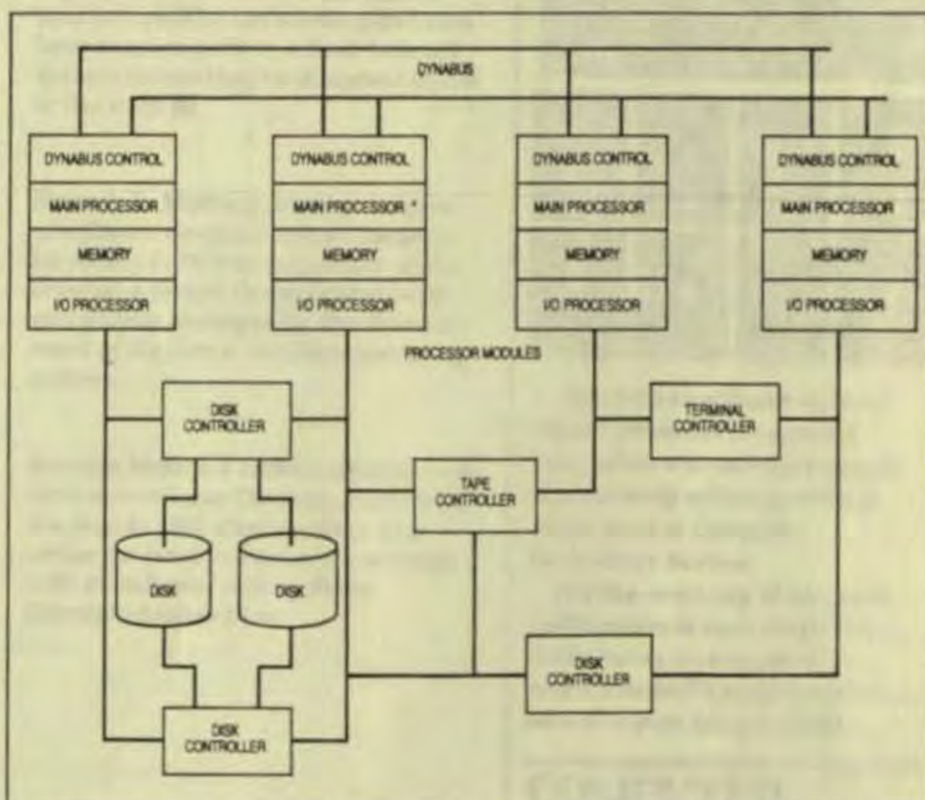

Fig 4 In a multicomputer system architecture, components are duplicated but not redundant. Each processor runs its own instruction stream and none of the backup components sits idle waiting for a failure.

cache still has to fetch instructions from memory though, so that while the problem may be minimized, the limits of the shared memory will be reached eventually.

A multicomputer software architecture eliminates such system bottlenecks as memory contention because the operating system runs in each processor, and the only communication between CPUs occurs when a user program requests I/O from another processor. This communication is done over the system bus, which runs at an aggregate rate of 26 Mbytes/s. Since the processor memory runs at 5 Mbytes/s, bus contention doesn't cause a problem.

Like the shared-memory approach, multicomputer architecture is modularly expandable. But because interprocessor communication occurs in the system bus, memory contention is eliminated along with the potential single-point failure of a shared memory. Therefore, a dual-processor system yields a full two processors' worth of computing power, three units yield three processors' worth of power, and so forth. There's virtually no performance performance degradation associated with expandability.

Another advantage of multicomputer architecture is the networking scheme inherent in a single system. Since the message system already keeps track of the physical location of all resources, a configuration can be expanded into a network of systems without any changes in existing software. The message system keeps track of each system as easily as it keeps track of each processor, and the user doesn't require any additional programming.

This networking scheme also contains a distributed database capability so data files can be distributed among a network of systems, while the message system automatically keeps track of their locations. The user doesn't need to know where a specific piece of data resides in order to access it.

CHOOSING A FAULT-TOLERANT SYSTEM

The choice of a fault-tolerant system should depend on the application. Some require 100% availability, but only during critical time periods. An automated betting system at a race track, for instance, must be available on race days, but could be serviced at night. In contrast, a system that monitors medical equipment in a hospital must be available continuously, even during periods of maintenance and repair.

You can determine the degree of fault tolerance required by analyzing the cost of failure for a given application. For example, suppose a system advertised a guaranteed up-time of over 99%. That may sound fairly reliable, but 1% downtime translates into approximately a third of a day each month. If the application involved is control of automated teller terminals and the downtime occurs during a Friday lunch hour, the 99% uptime won't soothe the angry customers. The cost of failure in this case requires 100% system availability.

Other considerations include the intial system cost, ease and cost of expansion and modification, and the vendor-supplied system software. Also, if a database management system is needed for the application, you should choose a vendor who can supply one because it's difficult to add features like networking and database management to a fault-tolerant system unless they're designed into it at the start. ■

Dennis L. McEvoy *is vice president of software development at Tandem. He joined in 1974 as a member of the original software development team and project manager for the development of the firm's NonStop operating system.*

Sandra Metz *is a technical public relations specialist at Tandem. She joined the firm in 1982 after working as a writer for Hewlett-Packard and holds a BS in technical writing from Carnegie-Mellon Univ.*

ALL IN ONE. ALL FOR $50

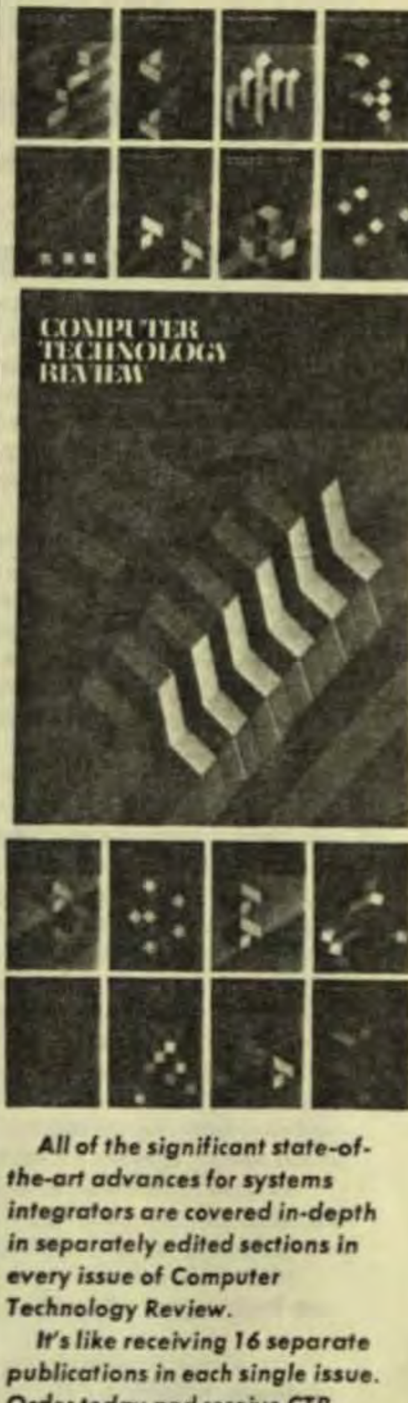

All of the significant state-of-the-art advances for systems integrators are covered in-depth in separately edited sections in every issue of Computer Technology Review.

It's like receiving 16 separate publications in each single issue. Order today and receive CTR and it s updated supplements for an entire year for only $50.00.

COMPUTER TECHNOLOGY REVIEW

The Systems Integration Sourcebook

924 Westwood Boulevard, Suite 650
Los Angeles, California 90024

TANDEM COMPUTERS
TANDEM COMPUTERS
TANDEM COMPUTERS

CORPORATE INFORMATION CENTER

An Interesting Application Controlled By
A Tandem NonStop™ 16 Transaction-Oriented
Computing System

As Reprinted From:
COMPUTERWORLD

Handles 1,300 Users in 46 States

Library Net Puts Minis at Front and Back Ends

COLUMBUS, Ohio – More than 1,300 libraries in 46 states have cut the cost of cataloging materials and speeded the placement of new books on their shelves by subscribing to a minicomputer-based communications network operated here by OCLC, Inc.

Each week, from about 2,000 remote on-line terminals, libraries of all types – college, public, specialized and school – catalog more than 200,000 book titles in OCLC's data base, which holds some 3.6 million records of materials like books, maps, music scores, sound recordings, manuscripts, audio-visual materials, magazines and journals.

Using the contents of this data base the hierarchical network produces machine-readable cataloging tapes and computer-generated accession lists. It also custom-prints about 1.7 million catalog cards each week for the non-profit organization's customers, according to Don Keates, manager of OCLC's system department.

Front and Back Minis

For front-end supervision, the nationwide network uses four TC-16 Nonstop minicomputers from Tandem Computers, Inc. Together, these front-end processors control incoming communications from more than 65 leased multidrop lines, Keates said.

Back-end data transfer, meanwhile between OCLC's 12G-byte data base and four Xerox Corp. Sigma 9 applications processors depends on 10 Tandem TC-16 minis, Keates added.

All 14 front-end and back-end processors in the OCLC network incorporate 484K-byte main memories, address up to 256 peripherals and access a 6G-byte Tandem disk drive system installed to duplicate the data base and thus ensure system integrity.

Each processor also uses Guardian, Tandem's virtual operating system that reportedly permits parallel processing in separate CPUs and multiprogramming in one CPU.

"Tandem was a logical choice for us," Keates said in explaining why OCLC picked its current configuration instead of rival systems. "Tandem offers the only commercially available uninterruptible computer system. Its systems can be configured with up to 16 interlinked multiprocessors, and the company supplies a data base record manager that gives us access to files and duplicates the data base."

Prospective Applications

When OCLC completes a planned expansion of its existing configuration, the firm will be able to automate several library activities in addition to the ones already on-line. Some of these prospective applications will include interlibrary loan management, materials acquisition, subject searches, remote catalog access and circulation control.

Plans are also under way to increase the front-end supervisor network to eight Tandem computers and to convert the network's CCITT X.25 standard packet protocol to communicate with other library systems throughout the country, Keates said.

To ensure continued data access during systems failure, OCLC has provided each of its peripherals with two paths to the communications network. If one of the paths fails, the configuration automatically activates the back-up path, Keates explained.

Interprocessor and data transfers are controlled independently in the OCLC network. A dual communications bus transfers data between processors at 13M byte/sec, while a block multiplexer channel, one for each CPU, oversees 4M byte/sec data transfers.

Because of the redundant access paths, OCLC operators can remove, replace or add any peripheral without interrupting network operations, the systems department manager explained.

To further ensure data base integrity, the network uses Tandem's Enscribe data base record manager as part of the Guardian operating system, Keates added. Enscribe reportedly provides a data definition language and a cache buffering scheme that reserves part of the network's main memory as a cache buffer.

With the network's software, applications programmers can write communications programs without knowing which processor will eventually run them. Programmers can also communicate with I/O devices or with other programs without knowing the physical addresses of the devices connected to the network.

As a result, the Tandem software simplifies the usually complex task of writing multiprogramming and multiprocessing application software, Keates noted.

Copyright 1978 by CW Communications/Inc., Newton, Mass. 02160