5

REVIEW & FORECAST
Part I

"The Gray Sheet"

EDP industry report
THE INTERNATIONAL DATA CORPORATION'S NEWSLETTER FOR EXECUTIVES CONCERNED WITH THE ELECTRONIC DATA PROCESSING INDUSTRY

D.C.BELLOMY, ASSOC.ED. JULY 1, 1985 • VOL.21, NO.4 P.J.MCGOVERN, PUBLISHER

IDC INTERNATIONAL DATA CORPORATION

VIEW FROM THE VALLEY: LARGE-SCALE SYSTEMS COMING OFF SLOWER GROWTH, FACING LOWER INCOME LEVELS IN 1985 DESPITE CONTINUING DEMAND

The shortest answer to the question about what has hit the computer industry is that there is no short answer. Issues differ in each segment of the market, with varying mixes of secular and cyclical trends, of accident and design. For large-scale systems, the dominant flavors in the bouillabaisse are cyclical and designed: a slowing tempo last year as the market for the 308X line and its rivals matured, and drops in demand this year during the transition to Sierra.

IDC forecasts for 1985 indicate a repeat of the weak sales during the 308X transition in 1981, when the majestic march of the mainframes was temporarily sidetracked with a 36.6% fall in large-system unit shipments and a 21.8% dip in revenue. However, financial officers will be smiling more starting next year:

- The $16.1 billion in large-scale revenue rung up in 1984 represented an increase of 11.1% over 1983, after a gain of 30.1% in 1983 over 1982 and 80.9% in 1982 over 1981.
- Large-scale shipments in 1985 should dip to 2,850 units, down 24.2% from 1984, while sector revenue should slip 15.5% to $13.6 billion.
- With mainframe shipments plummeting 45.7% from 2,533 units in 1984 to an estimated 1,375 in 1985, IBM large-scale revenues will be distinctly smaller-scaled at around $9.6 billion.
- Next year's new IBM large-scale shipments will probably only reach some 1,800 units, but since over three-quarters will be expensive Sierra boxes, the bottom line will be much healthier.
- Reflecting this turnaround, overall large-scale revenue of U.S. vendors should reach $18.3 billion next year, up 34.6% from 1985 and 13.7% from 1984.
- The installed base of large-scale systems should grow at a compound annual growth rate of 7.3% from 1984 through 1989.

All in all, during the decade of the Eighties (1980-89), there will have been approximately 33,270 large-scale shipments worth $151.3 billion.

EDP Industry Report is published in annual volumes of twenty-four issues each by International Data Corporation, Post Office Box 955, Five Speen Street, Framingham, MA 01701. Editor, David Moschella, Associate Editor, Donald Bellomy; Census Editors, Lloyd Cohen; Contributing Editors, Aaron Goldberg, Hank Janson, Ken McPherson, Ann Morley, Molly Upton, William F. Zachmann; Circulation Manager, Mary M. Trayte, Patrick J. McGovern, Chairman, International Data Group; Peter Rowell, President, International Data Corporation. Editorial ideas and information are researched and formulated by the editorial staff. Information and opinions in each issue are based on the best information available, but their accuracy and completeness cannot be guaranteed. Subscriptions: One year $395, two years, $705. Single copies priced as indicated on inside pages. Prices for multiple copies on request. Telex 95-1168 Telephone 617/872-8200.

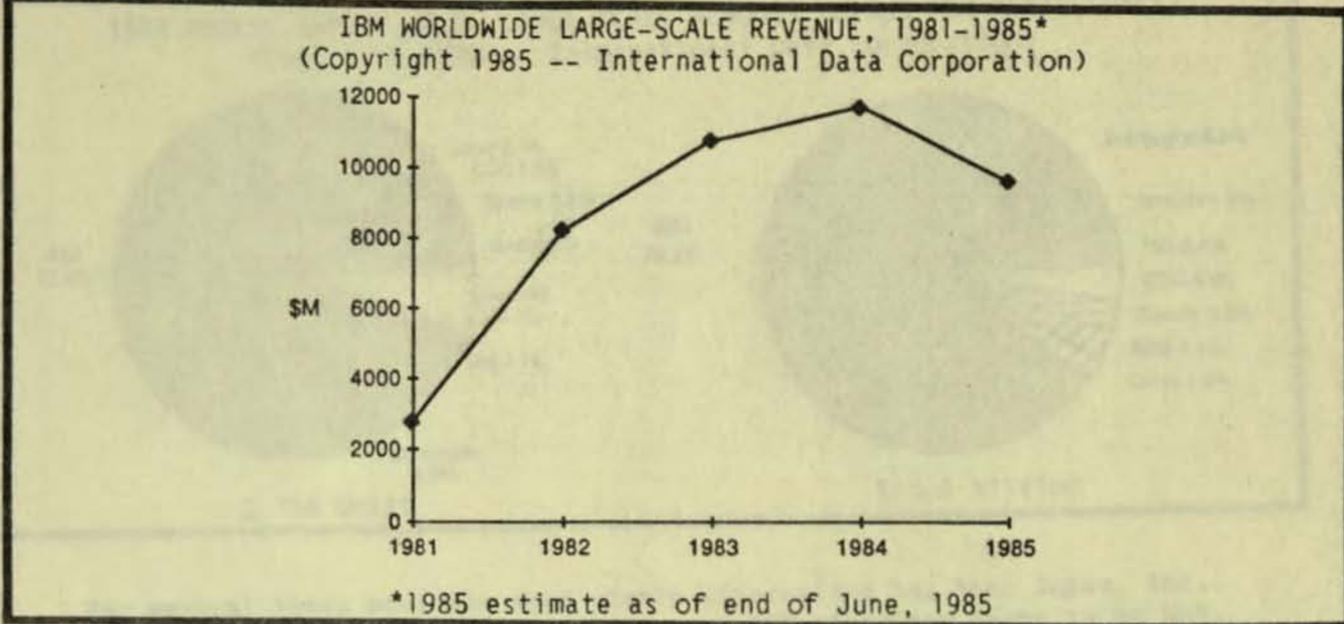

With profit margins high on large-scale systems, any growth engorges a company's coffers. A lot of money was made in 1984 by the manufacturers servicing this sector -- IDC estimates over $16 billion for 3,760 units plus peripherals split among eleven vendors. But different vendors got very different splits:

- IBM continued to lord it over other large-scale vendors, with 67.4% of mainframe computers sold and 76.1% of total value received. However, its large-scale sales growth did not quite match its rivals': revenue of $11.748 billion meant an increase of 9.2% over 1983.

- The PCM proportion of large-scale revenue stayed put at around 13%, with little movement in the rankings by either Amdahl or NAS.

- The Bunch registered impressive gains in mainframe revenue: up 31.4% over 1983. Noteworthy growth was registered by Burroughs (up 67.7% in large-scale sales) and Sperry (45.9%), but Honeywell and CDC trailed the industry average at 5.5% and 2.4% respectively.

Last year's figures represent the growing maturation of IBM's 308X family, just as some of Burroughs' and Sperry's spurt reflected delays in coming up with meaningful high-end competition. Announcement of IBM's 3090 generation of mainframes was probably timed to squeeze as much cash from its predecessor as possible. However, now that Sierra is official, squeezing much more money out looks like a blood and turnips proposition. For 1985 IBM should make fully half as much money on 308X upgrades ($930 million) as on new 308X systems.

As much as ever before, as IBM goes, so goes the large-scale computer industry. True, its 73.0% of large-scale system revenues (including peripherals) in 1984 seems a retreat of a few inches from 1983's 74.2%, but with that high a market share, its ups and downs will continue to define the health of the entire mainframe sector (compare the graph on the following page). Moreover, potential competition has yet to appear that could break IBM's full-nelson on mainframe money, particularly with the demise of Trilogy. But the other ten large-scale vendors still jockey for position to challenge Big Blue.

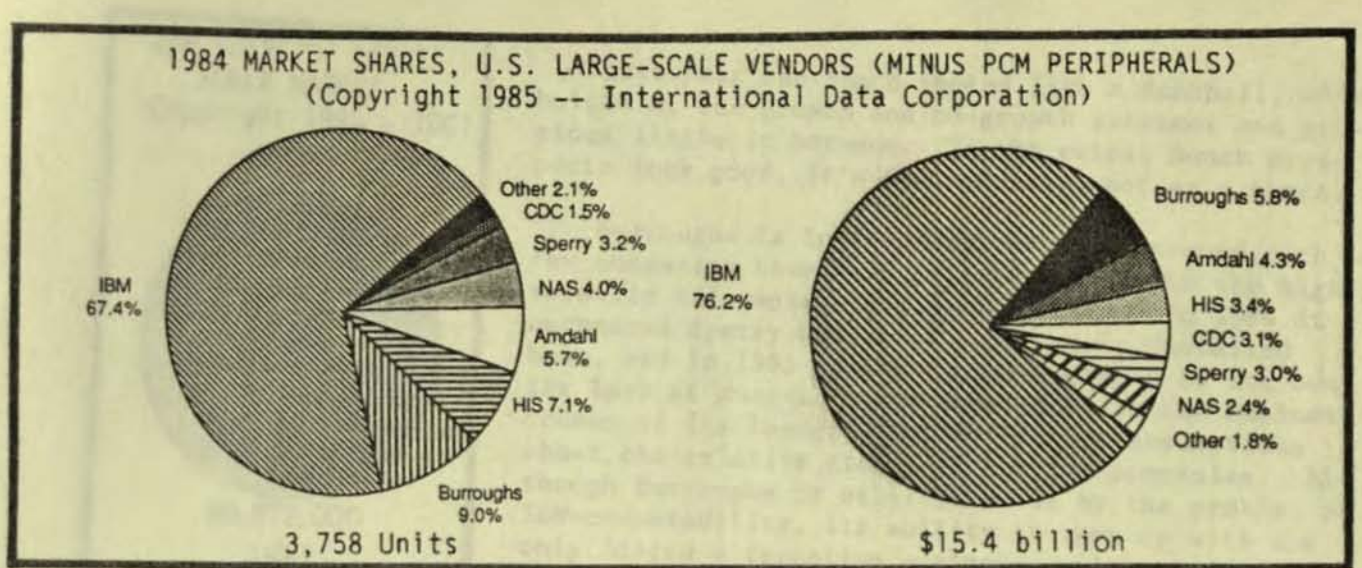

For several years now, the most viable alternative has been Japan, Inc., which in the U.S. means the PCMs. The purest PCM right now seems to be NAS, which is still completely dedicated to translating Hitachi engineering expertise into slightly less pricey IBM-compatible products. Its share of the market in 1984 is not overwhelming (2.3% of large-scale revenues), but the record revenue generated by NAS has just kept National Semi from plummeting far below the net loss of $2.7 million registered during the fiscal year just ended. NAS's prospects look fairly roseate right now, especially with the appearance of its AS/XL Sierra-class systems and double-density disks in 2Q 1986.

With 4.1% of large-scale revenue in 1984, Amdahl as usual did better than NAS. However, its activities so far in 1985 raise two questions about its future as a PCM. First there's the lack of activity in Sierra-class systems. The latest word, though, is that the Fujitsu-Amdahl combine will have a machine on the market in time to compete with NAS and IBM, but that its announcement will be delayed until shortly before delivery to avoid the delays and disappointed users that have harmed Amdahl's reputation in the past.

The second question is, plug-compatible with whom? Increasingly Amdahl has been cuddling up to AT&T and the Unix universe. On July 1 a multi-year agreement was announced whereby AT&T's Large Business Systems division will use Amdahl data communications products and systems. Later in the month word of close ties in Unix development is expected. It hardly matters whether AT&T is looking to pick up any of the Amdahl stock that Fujitsu doesn't own (it probably isn't) -- a working relationship is developing between Amdahl and AT&T, and hence between IBM plug-compatibility and Unix on the mainframe plane.

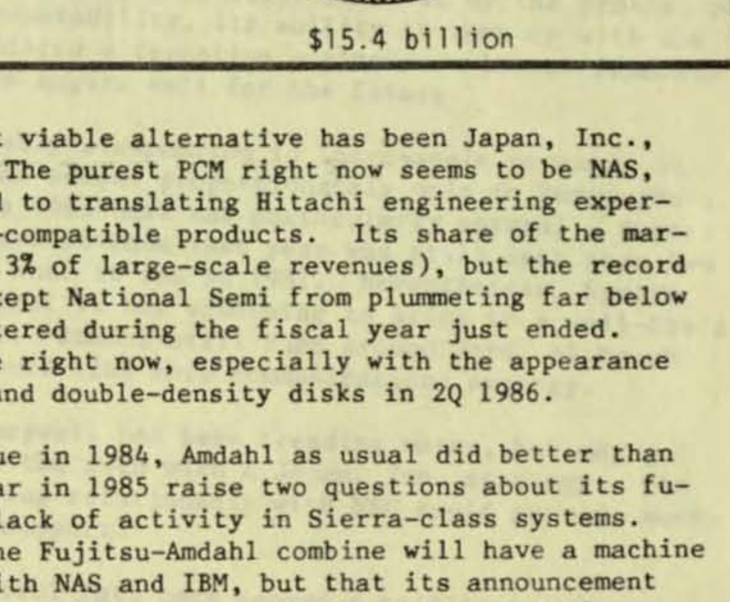

IBM SHARE OF LARGE-SCALE MARKET (Copyright 1985 - IDC)

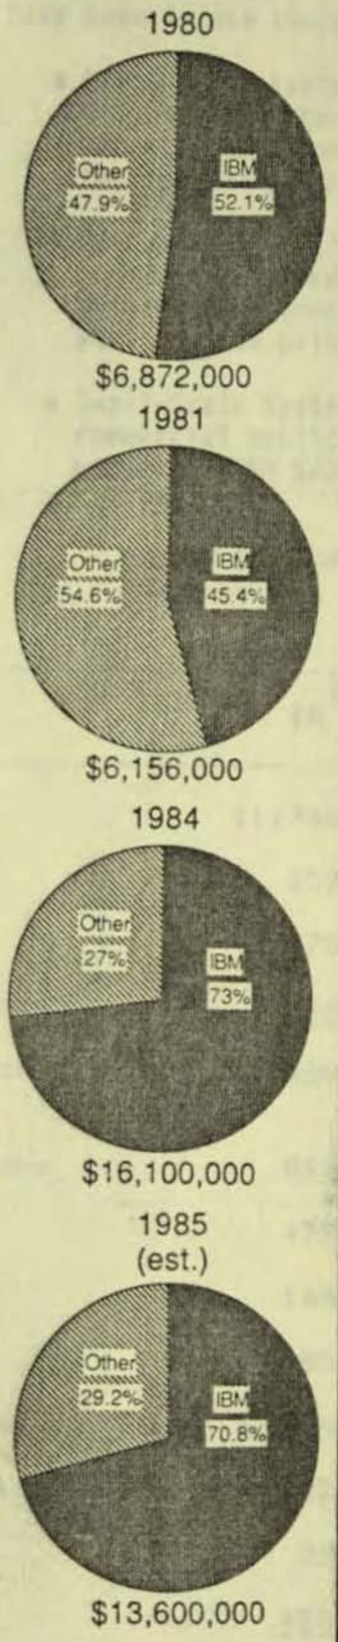

Last year the Bunch looked like a dumbbell, with bulges at the growth and no-growth extremes and precious little in between. To the extent Bunch prospects look good, it's individually, not as a bunch.

Burroughs is in the midst of a turnaround such as few companies have been able to achieve in the highly volatile information processing field. In 1984 it surpassed Sperry-Univac in large-scale installed base, and in 1985 sought to add Sperry's to its own. Its lack of success casts some doubt on the business acumen of its leaders, but not on its assumptions about the relative status of the two companies. Although Burroughs is still boxed in by the problem of IBM-compatability, its ability to come up with the only Sierra alternative engineered without Japanese knowhow augurs well for the future.

Sperry's expansion in large-scale revenues in 1984 was second proportionately only to Burroughs', but its 1983 base was particularly abysmal. Sperry mainframe revenue last year was still less than two thirds what it was in 1981. Nevertheless, Sperry feels that it has something to bring to a well-heeled partner. Rumors still race around Ford, although some would like more communications synergy.

Honeywell has been treading water, but may get back in the swim with a Japan, Inc. strategy. Its developing relationship with NEC could provide much-needed support.

Control Data once seemed a brilliant strategist for backing up its prowess in number-crunching with peripherals and services. Maybe it would have been if large-system stagnation were not accompanied by Japanese peripherals competition and a services slowdown. On July 1 CDC announced the sale of its Brokerage Transaction Services unit to ADP and promised additional "fundamental changes -- including divestiture of certain assets -- that will help us to solidify market positions." But it needs to do more than solidify -- it needs to find ways to grow again like Burroughs, or seek help like Sperry and Honeywell.

These, then, are IBM's rivals. The PCMs might gain a little ground during the Sierra turmoil, but it remains to be seen how different things will be when the dust settles. In 1981 the 308X transition boosted the market shares of IBM's competitors. Yet the next year IBM not only regained the market share it had lost, but immediately leapt to the 70+% position it has held ever since. That kind of domination is very hard to shake: despite its large-scale travails in 1985, IBM should still get 70% of sales.

ABOUT THE NEW DEFINITIONS

Long time readers of the Gray Sheet are reminded that new classifications have been in use since last year's Review and Forecast issues:

- Large-Scale systems. The very largest computers available from traditional mainframe and supercomputer vendors, large systems are expected to support over 128 users in many commercial environments and to bear average prices of more than $1 Million.
- Medium-Scale systems. Generally supporting 17-128 users in normal commercial environments, these include the IBM 4300 line, mid-range systems from Bunch companies, and the high end from DEC, DG, Prime, etc. System prices typically range from $100,000 to $1 Million.
- Small-Scale Systems. Generally ranging from 2-16 users in typical commercial applications and from $10,000-100,000 in price, these include the IBM S/36, DEC PDP-11 and Vax 730, and Altos-like systems.

VALUE OF WORLDWIDE SHIPMENTS OF LARGE-SCALE COMPUTERS
(Copyright 1985 -- International Data Corporation)

	1984		1983		1982		1981	
	$M	%	$M	%	$M	%	$M	%
IBM	11,748	73.0	10,755	74.2	8,211	73.7	2,796	45.4
Amdahl	659	4.1	585	4.0	344	3.1	281	4.6
NAS	375	2.3	335	2.3	225	2.0	122	2.0
PCM Periph.	686	4.3	701	4.8	816	7.3	700	11.4
IBM Base	13,468	83.7	12,376	85.4	9,596	86.2	3,899	63.3
Burroughs	894	5.6	533	3.7	177	1.6	531	8.6
CDC	479	3.0	468	3.2	394	3.5	370	6.0
Cray	165	1.0	125	0.9	122	1.1	91	1.5
DEC	85	0.5	174	1.2	276	2.5	65	1.1
Denelcor	4	0.0	2	0.0	6	0.1	0	0.0
Honeywell	522	3.2	495	3.4	206	1.8	458	7.4
NCR	25	0.2	0	0.0	24	0.2	0	0.0
Sperry	458	2.8	314	2.2	335	3.0	742	12.0
Non-IBM Base	2,632	16.3	2,111	14.6	1,540	13.8	2,257	36.7
TOTAL	16,100	100.0	14,487	100.0	11,136	100.0	6,156	100.0

LARGE-SCALE COMPUTER MARKET -- U.S. -BASED MANUFACTURERS

	Number of Systems Shipped	Cumulative Number in Use	$ Million Value Shipped	$ Billion Value in Use
		WORLDWIDE		
1980	2,380	14,170	6,890	60.3
1981	1,510	14,430	5,390	61.8
1982	2,320	15,180	11,060	68.3
1983	3,360	16,550	14,890	78.8
1984	3,760	18,150	16,090	91.1
1985	2,850	19,180	13,600	95.6
1986	3,580	20,610	18,300	104.1
1987	4,570	22,330	21,600	111.2
1988	4,720	24,200	22,100	120.0
1989	4,220	25,800	19,600	130.6
1990	5,000	27,700	24,300	140.7
		UNITED STATES		
1980	1,320	7,880	3,810	34.2
1981	890	8,290	3,130	36.3
1982	1,500	8,790	6,720	41.1
1983	1,830	9,460	8,340	46.3
1984	2,250	10,570	9,010	54.3
1985	1,530	11,100	7,300	57.1
1986	2,180	11,800	10,900	61.6
1987	2,790	12,800	13,100	65.3
1988	2,810	13,900	13,000	69.9
1989	2,420	14,800	10,900	76.2
1990	2,900	15,900	13,900	82.1
		INTERNATIONAL		
1980	1,060	6,290	3,080	26.1
1981	620	6,140	2,260	25.5
1982	820	6,390	4,340	27.2
1983	1,530	7,090	6,550	32.5
1984	1,510	7,580	7,080	36.8
1985	1,320	8,080	6,300	38.5
1986	1,400	8,810	7,400	42.5
1987	1,780	9,530	8,500	45.9
1988	1,910	10,300	9,100	50.1
1989	1,800	11,000	8,700	54.4
1990	2,100	11,800	10,400	58.6

LARGE-SCALE COMPUTER CENSUS, AS OF JAN. 1, 1985

(Copyright 1985 -- International Data Corporation)

Large-Scale Computer Systems. A large-scale system is either a general-purpose computer or a high-speed scientific computer with an approximate system price greater than $1 million. These machines are usually located in a centralized computer room. Representative examples include IBM 303X, 308X, Cray Computer products, and their competitors. Average purchase prices are based on system price including peripherals, except in the case of plug-compatible vendors Amdahl and National Advanced Systems, where only CPU value is counted.

NAME OF MANUFACTURER	COMPUTER MODEL	IDC SIZE CLASS	PURCHASE PRICE IN DOLLARS BASIC AVERAGE	DATE OF FIRST INSTALLATION	NUMBER INSTALLED IN U.S.	NUMBER INSTALLED OUTSIDE U.S.	TOTAL NUMBER INSTALLED
Amdahl	470V/5 II	L	2,350	09/77	18	10	28
	470V/6 II	VL	3,475	06/75	105	40	145
	470V/7,A,B	VL	2,542	09/78	80	60	140
	470V/7C	L	1,172	02/81	35	20	55
	470V/8	VL	2,592	09/79	140	85	225
	580/5840,50,60,67, 70,68	VL	3,500	08/82	250	150	400
	5868,5880	VL	0	--/85	0	0	0
	TOTAL				628	365	993
Burroughs	65/6700	L	3,377	11/69	80	70	150
	6800	L	1,824	06/77	230	90	320
	6900	L	2,373	11/80	118	62	180
	6925	L	1009	--/82	180	165	345
	7700	L	6,864	08/73	37	20	57
	7800	VL	5,396	11/78	152	42	194
	B-7900 FHK	VL	6,732	08/83	138	67	205
	TOTAL				935	516	1,451
Control Data Corporation	1604	L	1,865	01/60	1	0	1
	3600	L	2,246	06/63	4	2	6
	3800	L	2,737	12/65	10	0	10
	6200	L	1,945	11/70	0	0	0
	6400	L	2,902	05/66	30	9	39
	6500	L	3,944	12/67	12	10	22
	6600	VL	5,319	08/64	30	10	40
	6700	VL	6,401	07/70	2	0	2
	7600	VL	8,740	01/69	20	2	22
	CYBER 170-730	L	2,395	06/79	75	50	125
	CYBER 170-740	L	3,241	04/80	10	5	15
	CYBER 170-750	L	4,811	08/79	30	15	45
	CYBER 170-760	VL	6,147	04/79	40	8	48
	CYBER 170-855	L	4,040	11/82	20	15	35
	CYBER 170-865	VL	4,956	--/83	16	9	25
	CYBER 170-875	VL	6,497	--/82	16	9	25
	CYBER 170-875 DUAL	VL	9,209	--/83	3	0	3
	CYBER 172	L	2,138	07/75	12	30	42
	CYBER 173	L	3,553	09/75	7	18	25
	CYBER 174	L	4,453	09/75	10	10	20
	CYBER 175	VL	6,571	08/75	30	20	50
	CYBER 176	VL	8,048	01/78	15	3	18
	CYBER 203	VL	10,160	11/80	2	0	2
	CYBER 205	VL	11,735	01/81	22	7	29
	CYBER 73	L	2,803	03/72	20	25	45
	CYBER 74	VL	5,330	03/72	15	10	25
	CYBER 76	VL	8,426	03/72	5	6	11
	STAR 100	VL	12,240	12/74	0	0	0
	990,990 Dual	L	0	--/85	0	0	0
	TOTAL				457	273	730
Cray	1A, 1B	VL	8,836	04/76	10	7	17
	1S	VL	8,740	11/80	27	11	38
	M	VL	5,330	--/83	4	4	8
	X-MP/48	VL	14,000	06/84	0	0	0
	XMP/1	VL	6,000	08/84	4	1	5
	XMP22, 24	VL	10,659	--/83	12	6	18
	TOTAL				57	29	86
Denelcor	HEP	L	1,326	02/82	10	2	12
Digital Equipment Corp.	1080/90	L	1,419	06/75	450	255	705

LARGE-SCALE COMPUTER CENSUS, AS OF JAN. 1, 1985
(Copyright 1985 -- International Data Corporation)

NAME OF MANUFACTURER	COMPUTER MODEL	IDC SIZE CLASS	PURCHASE PRICE IN DOLLARS BASIC AVERAGE	DATE OF FIRST INSTALLATION	NUMBER INSTALLED IN U.S.	NUMBER INSTALLED OUTSIDE U.S.	TOTAL NUMBER INSTALLED
HIS	DPS 8/70	L	3,816	05/80	269	223	492
	DPS 88/41	L	1,850	--/85	0	0	0
	DPS 88/42	L	2,950	--/85	0	0	0
	DPS 88/42 Dual	L	3,700	--/85	0	0	0
	DPS 88/81	VL	6,059	--/83	11	9	20
	DPS 88/82	VL	8,609	--/83	7	3	10
	DPS 88/82T	VL	9,600	--/85	0	0	0
	G-600	L	3,524	04/65	5	5	10
	G-6050/60	L	3,464	07/71	30	35	65
	G-6070/80	L	5,212	07/71	7	13	20
	G-6180	VL	5,144	02/74	2	0	2
	H-66 DPS, BC	L	3,759	07/78	150	209	359
	H-66/40/60	L	3,920	10/74	70	130	200
	H-66/80	VL	6,125	10/74	10	30	40
	H-68/80	VL	5,242	10/74	12	1	13
	H-68/DPS	VL	6,228	09/78	20	8	28
	H-8200	L	3,144	02/69	3	2	5
	TOTAL				596	668	1,264
IBM	3031	L	2,601	03/78	510	665	1,175
	3032	L	4,402	03/78	145	295	440
	3033 UAM	VL	7,770	03/78	800	430	1,230
	3033N	VL	4,110	01/80	220	255	475
	3033S	L	3,603	01/81	80	50	130
	3081D,G,K	VL	7,781	10/81	1,700	1,050	2,750
	3083 B,E,J	L	4,590	01/83	1,850	1,375	3,225
	3084	VL	11,500	01/84	210	110	320
	360/65	L	3,464	11/65	85	15	100
	360/67,75,85,9X, 370/195	VL	7,140	01/66	10	3	13
	370/155	L	3,464	02/71	50	20	70
	370/158	L	5,356	05/73	600	430	1,030
	370/165	VL	6,502	06/71	15	2	17
	370/168	VL	9,646	07/73	200	100	300
	7030	L	1,428	05/61	0	0	0
	7080	L	2,985	08/61	0	0	0
	7090/94	L	3,571	08/60	0	0	0
	TOTAL				6,475	4,800	11,275
NAS	AS/5	L	1,620	03/77	60	50	110
	AS/6	L	2,543	12/78	30	25	55
	AS/7000	L	1,751	04/80	50	42	92
	AS/7031	L	1,382	12/78	15	10	25
	AS/8023	L	699	04/84	4	1	5
	AS/8040	L	1,376	--/83	17	6	23
	AS/8050	L	1,793	--/83	10	5	15
	AS/8060	L	2,296	--/83	6	3	9
	AS/8083	L	3,506	10/84	2	1	3
	AS/9000	VL	4,582	12/80	41	32	73
	AS/9040	VL	2,132	10/82	20	17	37
	AS/9050	VL	2,739	10/82	25	5	30
	AS/9060	VL	3,784	08/82	56	19	75
	AS/9070	VL	4,423	12/82	8	2	10
	AS/9080	VL	5,596	12/82	14	3	17
	TOTAL				358	221	579
NCR	8650	L	1,020	11/81	3	15	18
	8665,8675,8685, 8695	L	3,000	--/84	4	2	6
	8670	L	1,326	05/82	0	7	8
	TOTAL				8	24	32
Sperry Univac	1100/40	L	4,789	07/75	24	26	50
	1100/80/81/82	L	5,806	03/77	250	188	438
	1100/83	VL	9,307	07/78	22	27	49
	1100/84	VL	10,200	11/78	47	26	73
	1100/90/91	L	7,000	--/84	30	38	68
	1106	L	2,047	12/69	30	50	80
	1107	L	2,691	10/62	0	0	0
	1108	VL	3,997	09/65	65	15	80
	1110	VL	6,662	06/72	5	10	15
	494	L	3,411	02/66	43	16	59
	SPECTRA 70/6	L	2,393	10/71	5	0	5
	SPECTRA 70/7	L	2,281	10/71	3	1	4
	SPECTRA 70/60	L	2,393	09/70	3	0	3
	SPECTRA 70/61	L	2,718	03/71	0	0	0
	TOTAL				527	397	924
Xerox	SIGMA 9	L	2,558	11/71	65	25	90
GRAND TOTAL					10,566	7,575	18,141

REVIEW & FORECAST
Part I

"The Gray Sheet"

EDP industry report

● THE INTERNATIONAL DATA CORPORATION'S NEWSLETTER FOR EXECUTIVES CONCERNED WITH THE ELECTRONIC DATA PROCESSING INDUSTRY

D.C.BELLOMY, ASSOC.ED. JULY 1, 1985 • VOL.21, NO.4 P.J.MCGOVERN, PUBLISHER

VIEW FROM THE VALLEY: LARGE-SCALE SYSTEMS COMING OFF SLOWER GROWTH, FACING LOWER INCOME LEVELS IN 1985 DESPITE CONTINUING DEMAND

The shortest answer to the question about what has hit the computer industry is that there is no short answer. Issues differ in each segment of the market, with varying mixes of secular and cyclical trends, of accident and design. For large-scale systems, the dominant flavors in the bouillabaisse are cyclical and designed: a slowing tempo last year as the market for the 308X line and its rivals matured, and drops in demand this year during the transition to Sierra.

IDC forecasts for 1985 indicate a repeat of the weak sales during the 308X transition in 1981, when the majestic march of the mainframes was temporarily sidetracked with a 36.6% fall in large-system unit shipments and a 21.8% dip in revenue. However, financial officers will be smiling more starting next year:

- The $16.1 billion in large-scale revenue rung up in 1984 represented an increase of 11.1% over 1983, after a gain of 30.1% in 1983 over 1982 and 80.9% in 1982 over 1981.
- Large-scale shipments in 1985 should dip to 2,850 units, down 24.2% from 1984, while sector revenue should slip 15.5% to $13.6 billion.
- With mainframe shipments plummeting 45.7% from 2,533 units in 1984 to an estimated 1,375 in 1985, IBM large-scale revenues will be distinctly smaller-scaled at around $9.6 billion.
- Next year's new IBM large-scale shipments will probably only reach some 1,800 units, but since over three-quarters will be expensive Sierra boxes, the bottom line will be much healthier.
- Reflecting this turnaround, overall large-scale revenue of U.S. vendors should reach $18.3 billion next year, up 34.6% from 1985 and 13.7% from 1984.
- The installed base of large-scale systems should grow at a compound annual growth rate of 7.3% from 1984 through 1989.

All in all, during the decade of the Eighties (1980-89), there will have been approximately 33,270 large-scale shipments worth $151.3 billion.

EDP Industry Report is published in annual volumes of twenty-four issues each by International Data Corporation, Post Office Box 955, Five Speen Street, Framingham, MA 01701. Entire contents© 1985 by International Data Corporation. **Editor,** David Moschella, **Associate Editor,** Donald Bellomy; **Census Editors,** Lloyd Cohen; **Contributing Editors,** Aaron Goldberg, Hank Janson, Ken McPherson, Ann Morley, Molly Upton, William F. Zachmann; **Circulation Manager,** Mary M. Trayte, Patrick J. McGovern, Chairman, International Data Group; Peter Rowell, President, International Data Corporation. Editorial ideas and information are researched and formulated by the editorial staff. Information and opinions in each issue are based on the best information available, but their accuracy and completeness cannot be guaranteed. **Subscriptions:** One year $395, two years, $705. Single copies priced as indicated on inside pages. Prices for multiple copies on request. Reproduction without written permission is completely forbidden. Telex 95-1168 Telephone **617/872-8200.**

IBM WORLDWIDE LARGE-SCALE REVENUE, 1981-1985*
(Copyright 1985 -- International Data Corporation)

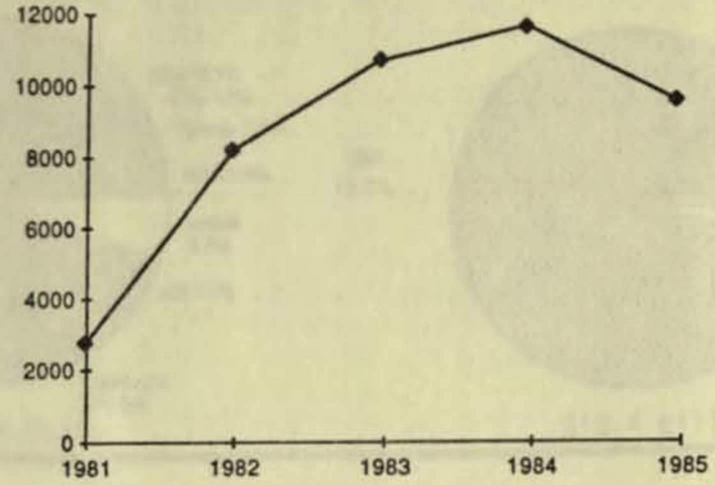

*1985 estimate as of end of June, 1985

With profit margins high on large-scale systems, any growth engorges a company's coffers. A lot of money was made in 1984 by the manufacturers servicing this sector -- IDC estimates over $16 billion for 3,760 units plus peripherals split among eleven vendors. But different vendors got very different splits:

- IBM continued to lord it over other large-scale vendors, with 67.4% of mainframe computers sold and 76.1% of total value received. However, its large-scale sales growth did not quite match its rivals': revenue of $11.748 billion meant an increase of 9.2% over 1983.
- The PCM proportion of large-scale revenue stayed put at around 13%, with little movement in the rankings by either Amdahl or NAS.
- The Bunch registered impressive gains in mainframe revenue: up 31.4% over 1983. Noteworthy growth was registered by Burroughs (up 67.7% in large-scale sales) and Sperry (45.9%), but Honeywell and CDC trailed the industry average at 5.5% and 2.4% respectively.

Last year's figures represent the growing maturation of IBM's 308X family, just as some of Burroughs' and Sperry's spurt reflected delays in coming up with meaningful high-end competition. Announcement of IBM's 3090 generation of mainframes was probably timed to squeeze as much cash from its predecessor as possible. However, now that Sierra is official, squeezing much more money out looks like a blood and turnips proposition. For 1985 IBM should make fully half as much money on 308X upgrades ($930 million) as on new 308X systems.

As much as ever before, as IBM goes, so goes the large-scale computer industry. True, its 73.0% of large-scale system revenues (including peripherals) in 1984 seems a retreat of a few inches from 1983's 74.2%, but with that high a market share, its ups and downs will continue to define the health of the entire mainframe sector (compare the graph on the following page). Moreover, potential competition has yet to appear that could break IBM's full-nelson on mainframe money, particularly with the demise of Trilogy. But the other ten large-scale vendors still jockey for position to challenge Big Blue.

1984 MARKET SHARES, U.S. LARGE-SCALE VENDORS (MINUS PCM PERIPHERALS)
(Copyright 1985 -- International Data Corporation)

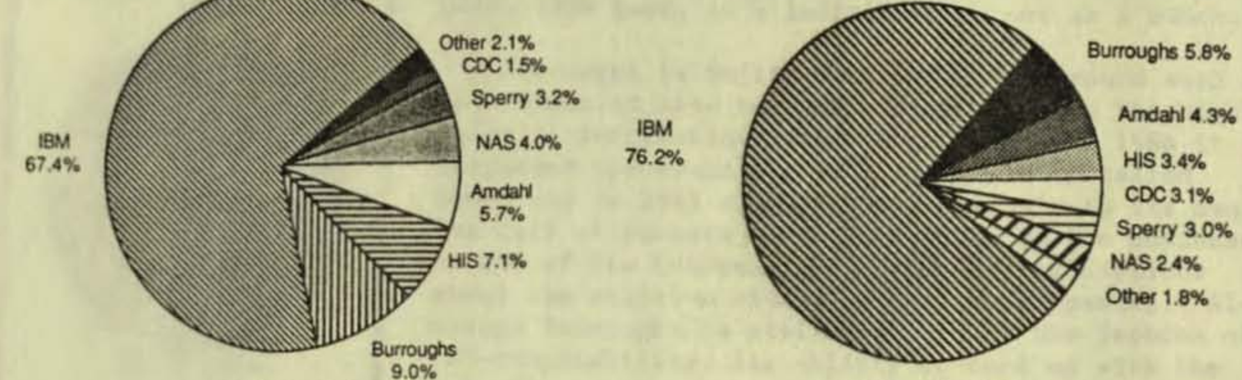

3,758 Units $15.4 billion

For several years now, the most viable alternative has been Japan, Inc., which in the U.S. means the PCMs. The purest PCM right now seems to be NAS, which is still completely dedicated to translating Hitachi engineering expertise into slightly less pricey IBM-compatible products. Its share of the market in 1984 is not overwhelming (2.3% of large-scale revenues), but the record revenue generated by NAS has just kept National Semi from plummeting far below the net loss of $2.7 million registered during the fiscal year just ended. NAS's prospects look fairly roseate right now, especially with the appearance of its AS/XL Sierra-class systems and double-density disks in 2Q 1986.

With 4.1% of large-scale revenue in 1984, Amdahl as usual did better than NAS. However, its activities so far in 1985 raise two questions about its future as a PCM. First there's the lack of activity in Sierra-class systems. The latest word, though, is that the Fujitsu-Amdahl combine will have a machine on the market in time to compete with NAS and IBM, but that its announcement will be delayed until shortly before delivery to avoid the delays and disappointed users that have harmed Amdahl's reputation in the past.

The second question is, plug-compatible with whom? Increasingly Amdahl has been cuddling up to AT&T and the Unix universe. On July 1 a multi-year agreement was announced whereby AT&T's Large Business Systems division will use Amdahl data communications products and systems. Later in the month word of close ties in Unix development is expected. It hardly matters whether AT&T is looking to pick up any of the Amdahl stock that Fujitsu doesn't own (it probably isn't) -- a working relationship is developing between Amdahl and AT&T, and hence between IBM plug-compatibility and Unix on the mainframe plane.

YEARLY CHANGE IN LARGE-SCALE SYSTEM REVENUE, 1980-1985 (EST.)
(Copyright 1985 - IDC)

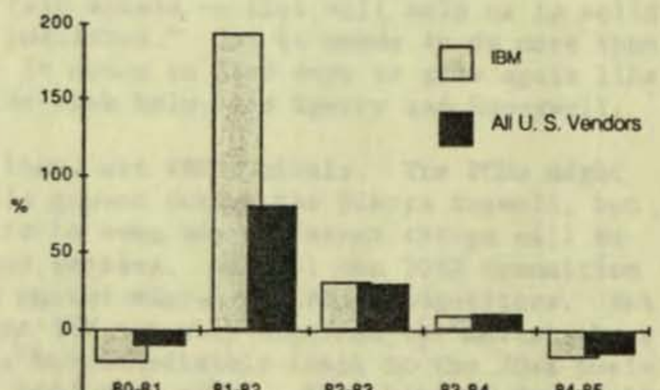

IBM SHARE OF LARGE-SCALE MARKET (Copyright 1985 - IDC)

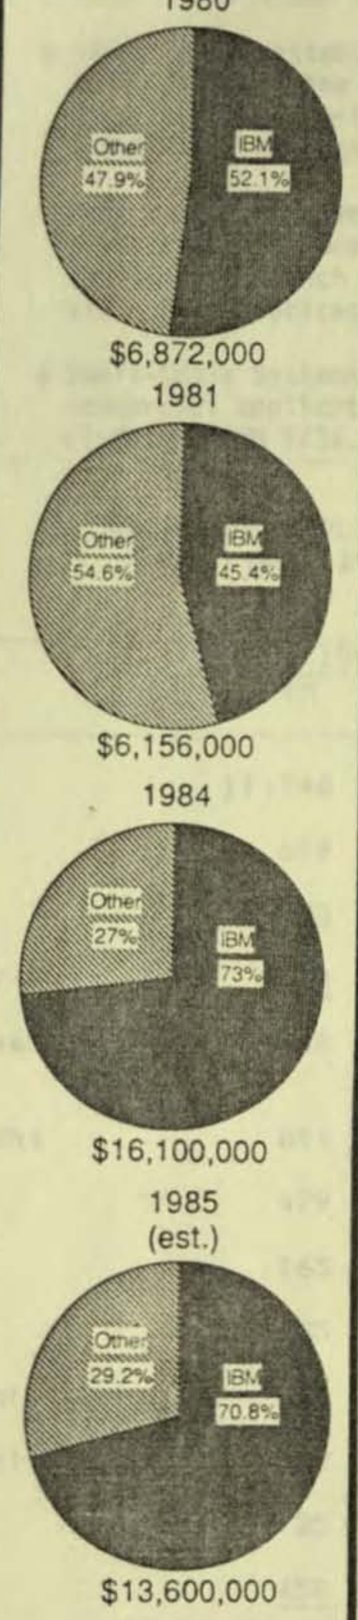

Last year the Bunch looked like a dumbbell, with bulges at the growth and no-growth extremes and precious little in between. To the extent Bunch prospects look good, it's individually, not as a bunch.

Burroughs is in the midst of a turnaround such as few companies have been able to achieve in the highly volatile information processing field. In 1984 it surpassed Sperry-Univac in large-scale installed base, and in 1985 sought to add Sperry's to its own. Its lack of success casts some doubt on the business acumen of its leaders, but not on its assumptions about the relative status of the two companies. Although Burroughs is still boxed in by the problem of IBM-compatability, its ability to come up with the only Sierra alternative engineered without Japanese knowhow augurs well for the future.

Sperry's expansion in large-scale revenues in 1984 was second proportionately only to Burroughs', but its 1983 base was particularly abysmal. Sperry mainframe revenue last year was still less than two thirds what it was in 1981. Nevertheless, Sperry feels that it has something to bring to a well-heeled partner. Rumors still race around Ford, although some would like more communications synergy.

Honeywell has been treading water, but may get back in the swim with a Japan, Inc. strategy. Its developing relationship with NEC could provide much-needed support.

Control Data once seemed a brilliant strategist for backing up its prowess in number-crunching with peripherals and services. Maybe it would have been if large-system stagnation were not accompanied by Japanese peripherals competition and a services slowdown. On July 1 CDC announced the sale of its Brokerage Transaction Services unit to ADP and promised additional "fundamental changes -- including divestiture of certain assets -- that will help us to solidify market positions." But it needs to do more than solidify -- it needs to find ways to grow again like Burroughs, or seek help like Sperry and Honeywell.

These, then, are IBM's rivals. The PCMs might gain a little ground during the Sierra turmoil, but it remains to be seen how different things will be when the dust settles. In 1981 the 308X transition boosted the market shares of IBM's competitors. Yet the next year IBM not only regained the market share it had lost, but immediately leapt to the 70+% position it has held ever since. That kind of domination is very hard to shake: despite its large-scale travails in 1985, IBM should still get 70% of sales.

ABOUT THE NEW DEFINITIONS

Long time readers of the Gray Sheet are reminded that new classifications have been in use since last year's Review and Forecast issues:

- Large-Scale systems. The very largest computers available from traditional mainframe and supercomputer vendors, large systems are expected to support over 128 users in many commercial environments and to bear average prices of more than $1 Million.
- Medium-Scale systems. Generally supporting 17-128 users in normal commercial environments, these include the IBM 4300 line, mid-range systems from Bunch companies, and the high end from DEC, DG, Prime, etc. System prices typically range from $100,000 to $1 Million.
- Small-Scale Systems. Generally ranging from 2-16 users in typical commercial applications and from $10,000-100,000 in price, these include the IBM S/36, DEC PDP-11 and Vax 730, and Altos-like systems.

VALUE OF WORLDWIDE SHIPMENTS OF LARGE-SCALE COMPUTERS
(Copyright 1985 -- International Data Corporation)

	1984		1983		1982		1981	
	$M	%	$M	%	$M	%	$M	%
IBM	11,748	73.0	10,755	74.2	8,211	73.7	2,796	45.4
Amdahl	659	4.1	585	4.0	344	3.1	281	4.6
NAS	375	2.3	335	2.3	225	2.0	122	2.0
PCM Periph.	686	4.3	701	4.8	816	7.3	700	11.4
IBM Base	13,468	83.7	12,376	85.4	9,596	86.2	3,899	63.3
Burroughs	894	5.6	533	3.7	177	1.6	531	8.6
CDC	479	3.0	468	3.2	394	3.5	370	6.0
Cray	165	1.0	125	0.9	122	1.1	91	1.5
DEC	85	0.5	174	1.2	276	2.5	65	1.1
Denelcor	4	0.0	2	0.0	6	0.1	0	0.0
Honeywell	522	3.2	495	3.4	206	1.8	458	7.4
NCR	25	0.2	0	0.0	24	0.2	0	0.0
Sperry	458	2.8	314	2.2	335	3.0	742	12.0
Non-IBM Base	2,632	16.3	2,111	14.6	1,540	13.8	2,257	36.7
TOTAL	16,100	100.0	14,487	100.0	11,136	100.0	6,156	100.0

LARGE-SCALE COMPUTER MARKET -- U.S. -BASED MANUFACTURERS

WORLDWIDE

	Number of Systems Shipped	Cumulative Number in Use	$ Million Value Shipped	$ Billion Value in Use
1980	2,380	14,170	6,890	60.3
1981	1,510	14,430	5,390	61.8
1982	2,320	15,180	11,060	68.3
1983	3,360	16,550	14,890	78.8
1984	3,760	18,150	16,090	91.1
1985	2,850	19,180	13,600	95.6
1986	3,580	20,610	18,300	104.1
1987	4,570	22,330	21,600	111.2
1988	4,720	24,200	22,100	120.0
1989	4,220	25,800	19,600	130.6
1990	5,000	27,700	24,300	140.7

UNITED STATES

	Number of Systems Shipped	Cumulative Number in Use	$ Million Value Shipped	$ Billion Value in Use
1980	1,320	7,880	3,810	34.2
1981	890	8,290	3,130	36.3
1982	1,500	8,790	6,720	41.1
1983	1,830	9,460	8,340	46.3
1984	2,250	10,570	9,010	54.3
1985	1,530	11,100	7,300	57.1
1986	2,180	11,800	10,900	61.6
1987	2,790	12,800	13,100	65.3
1988	2,810	13,900	13,000	69.9
1989	2,420	14,800	10,900	76.2
1990	2,900	15,900	13,900	82.1

INTERNATIONAL

	Number of Systems Shipped	Cumulative Number in Use	$ Million Value Shipped	$ Billion Value in Use
1980	1,060	6,290	3,080	26.1
1981	620	6,140	2,260	25.5
1982	820	6,390	4,340	27.2
1983	1,530	7,090	6,550	32.5
1984	1,510	7,580	7,080	36.8
1985	1,320	8,080	6,300	38.5
1986	1,400	8,810	7,400	42.5
1987	1,780	9,530	8,500	45.9
1988	1,910	10,300	9,100	50.1
1989	1,800	11,000	8,700	54.4
1990	2,100	11,800	10,400	58.6

MAST

MARKET ANALYSIS & STATISTICS

IDC OFFERS TWENTY YEAR HISTORICAL CENSUS DATA
FOR STRATEGIC ANALYSIS

International Data Corporation's (IDC) Processor Installation Census (PIC) file is a machine readable database offering model-by-model census data for the years 1964 through 1984 including historical installed base, shipment, and average system value data.

The File, designed for PC or mainframe use, includes data on more than 1,400 systems and over 300 vendors identified in IDC's Censuses from the past twenty years. The database places valuable information at your fingertips. With it you can:

- Assess market shares.
- Establish price bands.
- Analyze processor shipment rates
- Determine processor retirement rates.

U.S. LARGE-SCALE INSTALLED BASE 1966 - 1984

PIC #	CODE	MODEL	NEW IDC	DFI	1966	1967	1968	1969	1970	1971	1972	1973	1974	1975	1976	1977	1978	1979	1980	1981	1982	1983	1984
** BURROUGHS																							
53	2663	65/6700	L	11/01/69	0	0	0	4	16	32	45	60	80	103	125	155	160	145	140	137	127	120	88
56	2740	6800	L	06/01/77	0	0	0	0	0	0	0	0	0	0	0	17	65	140	280	310	305	270	230
57	2745	6900	L	11/01/80	0	0	0	0	0	0	0	0	0	0	0	0	0	0	34	83	105	115	110
58	2748	6925	L	06/30/84	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	12	100	180
54	2675	7700	L	08/01/73	0	0	0	0	0	0	2	3	5	14	23	31	37	49	48	45	38	37	37
55	2682	7800	VL	11/01/78	0	0	0	0	0	0	0	0	0	0	0	0	2	4	38	108	129	158	152
59	2779	B-7900 F-K	VL	08/01/83	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	20	138
** Subtotal **					0	0	0	4	16	32	47	63	85	117	148	203	264	338	540	683	716	820	935

The Processor Installation Census (PIC) File is the only IDC product that provides revised censuses of the installed base. The PIC File has been

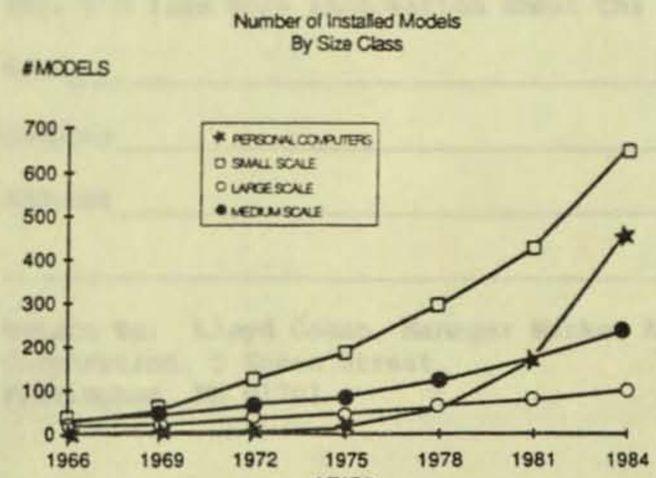

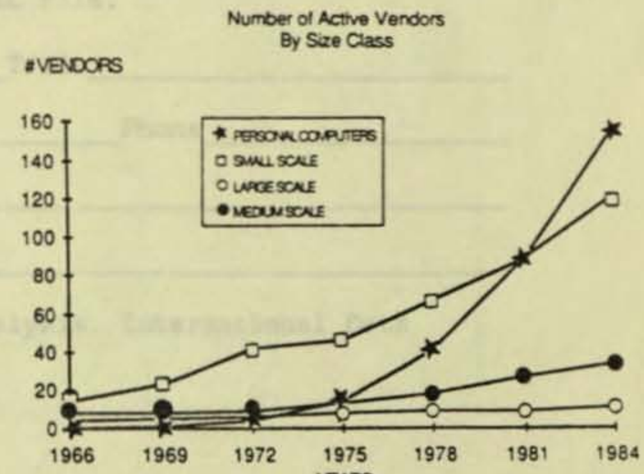

-Morethoroughly reevaluated and adjusted where necessary to reflect IDC's best current information and provides:

- A unified and consistent source for evaluating processor market trends and relative vendor performance.

- Ready to use data. NO data entry is required.

Since 1964, IDC has been regularly publishing its annual processor census. IDC has expanded the number of censuses it annually publishes as new processor markets have emerged. IDC censuses now provide system by system coverage of Large Scale, Medium Scale, Small Scale, and Personal Computer processor marketplaces. IDC is the only computer market research and consulting research firm that publishes this data on EVERY U.S. based computer manufacturer on a model-by-model basis.

The 1964-1984 PIC file with complete 1984 census data is available now. Annual updates will be available each year during the month of February. Having this timely and accurate data will provide a great benefit to both research and strategic planning work efforts. Each of the charts and tables on the front of this page were derived easily from the PIC file.

The media and recording formats include:

- ASCII files on IBM PC-Compatible diskettes.

- ASCII files on standard 1/2", 1600 BPI magnetic tape.

- EBCDIC files on standard 1/2", 1600 BPI magnetic tape.

- dBase III* files on IBM PC-Compatible diskettes.

For additional information on how the PIC file can assist you in your planning and research activities please call Lloyd Cohen at (617) 872-8200 or return the form below.

YES, I'd like more information about the PIC File.

Name________________________________ Title______________________

Company__________________________________ Phone________________

Address____

____

Return to: Lloyd Cohen, Manager Market Analysis, International Data Corporation, 5 Speen Street, Framingham, MA 01701

LARGE-SCALE COMPUTER CENSUS, AS OF JAN. 1, 1985

(Copyright 1985 -- International Data Corporation)

Large-Scale Computer Systems. A large-scale system is either a general-purpose computer or a high-speed scientific computer with an approximate system price greater than $1 million. These machines are usually located in a centralized computer room. Representative examples include IBM 303X, 308X, Cray Computer products, and their competitors. Average purchase prices are based on system price including peripherals, except in the case of plug-compatible vendors Amdahl and National Advanced Systems, where only CPU value is counted.

NAME OF MANUFACTURER	COMPUTER MODEL	IDC SIZE CLASS	PURCHASE PRICE IN DOLLARS BASIC AVERAGE	DATE OF FIRST INSTALLATION	NUMBER INSTALLED IN U.S.	NUMBER INSTALLED OUTSIDE U.S.	TOTAL NUMBER INSTALLED
Amdahl	470V/5 II	L	2,350	09/77	18	10	28
	470V/6 II	VL	3,475	06/75	105	40	145
	470V/7,A,B	VL	2,542	09/78	80	60	140
	470V/7C	L	1,172	02/81	35	20	55
	470V/8	VL	2,592	09/79	140	85	225
	580/5840,50,60,67, 70,68	VL	3,500	08/82	250	150	400
	5868,5880	VL	0	--/85	0	0	0
	TOTAL				628	365	993
Burroughs	65/6700	L	3,377	11/69	80	70	150
	6800	L	1,824	06/77	230	90	320
	6900	L	2,373	11/80	118	62	180
	6925	L	1009	--/82	180	165	345
	7700	L	6,864	08/73	37	20	57
	7800	VL	5,396	11/78	152	42	194
	B-7900 FHK	VL	6,732	08/83	138	67	205
	TOTAL				935	516	1,451
Control Data Corporation	1604	L	1,865	01/60	1	0	1
	3600	L	2,246	06/63	4	2	6
	3800	L	2,737	12/65	10	0	10
	6200	L	1,945	11/70	0	0	0
	6400	L	2,902	05/66	30	9	39
	6500	L	3,944	12/67	12	10	22
	6600	VL	5,319	08/64	30	10	40
	6700	VL	6,401	07/70	2	0	2
	7600	VL	8,740	01/69	20	2	22
	CYBER 170-730	L	2,395	06/79	75	50	125
	CYBER 170-740	L	3,241	04/80	10	5	15
	CYBER 170-750	L	4,811	08/79	30	15	45
	CYBER 170-760	VL	6,147	04/79	40	8	48
	CYBER 170-855	L	4,040	11/82	20	15	35
	CYBER 170-865	VL	4,956	--/83	16	9	25
	CYBER 170-875	VL	6,497	--/82	16	9	25
	CYBER 170-875 DUAL	VL	9,209	--/83	3	0	3
	CYBER 172	L	2,138	07/75	12	30	42
	CYBER 173	L	3,553	09/75	7	18	25
	CYBER 174	L	4,453	09/75	10	10	20
	CYBER 175	VL	6,571	08/75	30	20	50
	CYBER 176	VL	8,048	01/78	15	3	18
	CYBER 203	VL	10,160	11/80	2	0	2
	CYBER 205	VL	11,735	01/81	22	7	29
	CYBER 73	L	2,803	03/72	20	25	45
	CYBER 74	VL	5,330	03/72	15	10	25
	CYBER 76	VL	8,426	03/72	5	6	11
	STAR 100	VL	12,240	12/74	0	0	0
	990,990 Dual	L	0	--/85	0	0	0
	TOTAL				457	273	730
Cray	1A, 1B	VL	8,836	04/76	10	7	17
	1S	VL	8,740	11/80	27	11	38
	M	VL	5,330	--/83	4	4	8
	X-MP/48	VL	14,000	06/84	0	0	0
	XMP/1	VL	6,000	08/84	4	1	5
	XMP22, 24	VL	10,659	--/83	12	6	18
	TOTAL				57	29	86
Denelcor	HEP	L	1,326	02/82	10	2	12
Digital Equipment Corp.	1080/90	L	1,419	06/75	450	255	705

LARGE-SCALE COMPUTER CENSUS, AS OF JAN. 1, 1985
(Copyright 1985 -- International Data Corporation)

NAME OF MANUFACTURER	COMPUTER MODEL	IDC SIZE CLASS	PURCHASE PRICE IN DOLLARS BASIC AVERAGE	DATE OF FIRST INSTALLATION	NUMBER INSTALLED IN U.S.	NUMBER INSTALLED OUTSIDE U.S.	TOTAL NUMBER INSTALLED
HIS	DPS 8/70	L	3,816	05/80	269	223	492
	DPS 88/41	L	1,850	--/85	0	0	0
	DPS 88/42	L	2,950	--/85	0	0	0
	DPS 88/42 Dual	L	3,700	--/85	0	0	0
	DPS 88/81	VL	6,059	--/83	11	9	20
	DPS 88/82	VL	8,609	--/83	7	3	10
	DPS 88/82T	VL	9,600	--/85	0	0	0
	G-600	L	3,524	04/65	5	5	10
	G-6050/60	L	3,464	07/71	30	35	65
	G-6070/80	L	5,212	07/71	7	13	20
	G-6180	VL	5,144	02/74	2	0	2
	H-66 DPS, BC	L	3,759	07/78	150	209	359
	H-66/40/60	L	3,920	10/74	70	130	200
	H-66/80	VL	6,125	10/74	10	30	40
	H-68/80	VL	5,242	10/74	12	1	13
	H-68/DPS	VL	6,228	09/78	20	8	28
	H-8200	L	3,144	02/69	3	2	5
	TOTAL				596	668	1,264
IBM	3031	L	2,601	03/78	510	665	1,175
	3032	L	4,402	03/78	145	295	440
	3033 UAM	VL	7,770	03/78	800	430	1,230
	3033N	VL	4,110	01/80	220	255	475
	3033S	L	3,603	01/81	80	50	130
	3081D,G,K	VL	7,781	10/81	1,700	1,050	2,750
	3083 B,E,J	L	4,590	01/83	1,850	1,375	3,225
	3084	VL	11,500	01/84	210	110	320
	360/65	L	3,464	11/65	85	15	100
	360/67,75,85,9X, 370/195	VL	7,140	01/66	10	3	13
	370/155	L	3,464	02/71	50	20	70
	370/158	L	5,356	05/73	600	430	1,030
	370/165	VL	6,502	06/71	15	2	17
	370/168	VL	9,646	07/73	200	100	300
	7030	L	1,428	05/61	0	0	0
	7080	L	2,985	08/61	0	0	0
	7090/94	L	3,571	08/60	0	0	0
	TOTAL				6,475	4,800	11,275
NAS	AS/5	L	1,620	03/77	60	50	110
	AS/6	L	2,543	12/78	30	25	55
	AS/7000	L	1,751	04/80	50	42	92
	AS/7031	L	1,382	12/78	15	10	25
	AS/8023	L	699	04/84	4	1	5
	AS/8040	L	1,376	--/83	17	6	23
	AS/8050	L	1,793	--/83	10	5	15
	AS/8060	L	2,296	--/83	6	3	9
	AS/8083	L	3,506	10/84	2	1	3
	AS/9000	VL	4,582	12/80	41	32	73
	AS/9040	VL	2,132	10/82	20	17	37
	AS/9050	VL	2,739	10/82	25	5	30
	AS/9060	VL	3,784	08/82	56	19	75
	AS/9070	VL	4,423	12/82	8	2	10
	AS/9080	VL	5,596	12/82	14	3	17
	TOTAL				358	221	579
NCR	8650	L	1,020	11/81	3	15	18
	8665,8675,8685, 8695	L	3,000	--/84	4	2	6
	8670	L	1,326	05/82	0	7	8
	TOTAL				8	24	32
Sperry Univac	1100/40	L	4,789	07/75	24	26	50
	1100/80/81/82	L	5,806	03/77	250	188	438
	1100/83	VL	9,307	07/78	22	27	49
	1100/84	VL	10,200	11/78	47	26	73
	1100/90/91	L	7,000	--/84	30	38	68
	1106	L	2,047	12/69	30	50	80
	1107	L	2,691	10/62	0	0	0
	1108	VL	3,997	09/65	65	15	80
	1110	VL	6,662	06/72	5	10	15
	494	L	3,411	02/66	43	16	59
	SPECTRA 70/6	L	2,393	10/71	5	0	5
	SPECTRA 70/7	L	2,281	10/71	3	1	4
	SPECTRA 70/60	L	2,393	09/70	3	0	3
	SPECTRA 70/61	L	2,718	03/71	0	0	0
	TOTAL				527	397	924
Xerox	SIGMA 9	L	2,558	11/71	65	25	90
GRAND TOTAL					10,566	7,575	18,141

REVIEW & FORECAST
Part II

"The Gray Sheet"

EDP industry report
THE INTERNATIONAL DATA CORPORATION'S NEWSLETTER FOR EXECUTIVES CONCERNED WITH THE ELECTRONIC DATA PROCESSING INDUSTRY

D.C.BELLOMY, ASSOC. ED. July 31, 1985 • VOL.21,No.5,6 P.J.MCGOVERN, PUBLISHER

IDC INTERNATIONAL DATA CORPORATION

ROLLERCOASTER REVENUE CLIMBS AND SLIDES UNSETTLE STOMACHS FOR MEDIUM-, SMALL-SCALE SYSTEM VENDORS

You can't always tell from watching cable television, but most Americans are still Puritans at heart who believe that whatever they suffer now is retribution for having enjoyed themselves too much in the past. No matter what the economists said during the Great Depression, a lot of people were sure they were getting what they deserved for making too much money and drinking too much illegal booze during the Roaring Twenties. These days some medium- and small-scale vendors, especially the former, are feeling that they're getting what they deserve for having had so much fun, and made so much money, in 1984.

- Medium-scale shipments were up 59% in units over 1983, and 40% in dollar value.
- Small-scale shipments gained 22.7% in units, although declining prices kept the growth in dollar value of shipments to 13.3%.
- The installed base of medium-scale systems grew 30%, and of small-scale systems 20%.

But then the drought began. In case anybody hasn't noticed, 1985 looks vastly different from its immediate predecessor:

- The value of medium-scale shipments should climb a paltry 3.7%, while unit shipments will likely be marginally down .3% from 1984.
- Small-scale shipments should climb at a slower 8.7% rate this year, with the dollar value gliding up 10.8%.

IDC anticipates that the worst will be over by the end of the year, but that declining prices will put a cap on how much revenue can be garnered from the two sectors over the next half decade:

- Between 1984 and 1989, medium-scale shipments should grow at an annual rate of 12.9%, with the value of shipments increasing at 10.1%.
- Over the same five years, small-scale shipments should leap forward at 14.9% per annum, with revenue almost keeping pace at 13.7%.

EDP Industry Report is published in annual volumes of twenty-four issues each by International Data Corporation, Post Office Box 955, Five Speen Street, Framingham, MA 01701. Entire contents© 1985 by International Data Corporation. Editor, David Moschella, Associate Editor, Donald Bellomy; Census Editors, Lloyd Cohen; Contributing Editors, Aaron Goldberg, Hank Janson, Ken McPherson, Ann Morley, Molly Upton, William F. Zachmann; Circulation Manager, Mary M. Trayte, Patrick J. McGovern, Chairman, International Data Group; Peter Rowell, President, International Data Corporation. Editorial ideas and information are researched and formulated by the editorial staff. Information and opinions in each issue are based on the best information available, but their accuracy and completeness cannot be guaranteed. Subscriptions: One year $395, two years, $705. Single copies priced as indicated on inside pages. Prices for multiple copies on request. Reproduction without written permission is completely forbidden. Telex 95-1168 Telephone 617/872-8200.

Medium-Scale Tedium

"However, demand for intermediate systems in the U.S. remains soft." Thus spake John Akers for IBM's second quarter, but he could have been speaking for virtually the entire industry. The 4% rise in demand for new shipments that IDC is projecting for 1985 (and even that requires more sales, soon) is being reflected in some noticeably weak performances by last year's high flyers:

- Wang suffered an 11.0% slide in revenue from the second quarter of 1984 and a loss of $109.0 million, its first since it entered the computer market. Through 1984 Wang revenue was up 35.0%, while earnings grew 29.1%.

- With revenue down 6% from a year ago, Data General posted an $8.3 million loss for the quarter. During FY1984 DG was up 40% in revenue and 260% in earnings.

- Tandem eeked out a 1.6% increase in revenue that translated into a 74.2% dip in earnings. In FY84 it jumped 27.3% in revenue and 39.3% in earnings.

- Prime revenue gained 16.0% over 1984's second quarter, but earnings inched up only 2.3%. Last year revenue ballooned 24.4%, while earnings boomed 83.7%.

Nowhere in the industry is soul-searching over the origins of the current downturn more intense than in the medium-scale market. Unlike the situation in the large-scale segment, there are relatively few product cycle problems -- certainly none that could have an impact on other companies the way waiting for Sierra has. Unlike the case with the PC segment, there is little evidence of temporary saturation or miscalculation of key submarkets.

Groping for explanations, competitors in the field hope that the disease is not systemic. Vendors with a presence on the factory floor like DG tend to blame the economy, which is still weak-to-dismal for traditional manufacturing industries. Companies with a strong office-automation profile can take comfort from the probability that at least some of the slowness represents indecisiveness on the part of office managers about where to go next after last year's buying binge. Presumably they'll be buying again someday, maybe soon.

But what if the fault lies not in the vendors' stars but in themselves that they are underlings? What if a significant portion of the downturn stems from problems endemic to the medium-scale market that will not go away any time soon, if ever, like encroachment from networked PCs and supermicros? What if vendors were basking in the glow of 1984's long shimmering summer and failed to prepare adequately for changing markets and accelerating technology curves?

The encroachment question may not be answerable except in retrospect a couple of years from now, but answers to the last question will begin to become clearer over the next six months. It will not be easy for a company like Wang to recover a position anywhere close to what it has achieved in the past. An Wang may have to remake his company's image as dramatically as when he moved from calculators to word processing. As of July 25 Wang is also finally acknowledging the necessity to support the domination of the officeplace by the IBM PC, something that DG and Digital have also been forced to do.

1984 MARKET SHARES OF WORLDWIDE INSTALLED BASE, MEDIUM-SCALE SYSTEMS
(Copyright 1985 -- International Data Corporation)

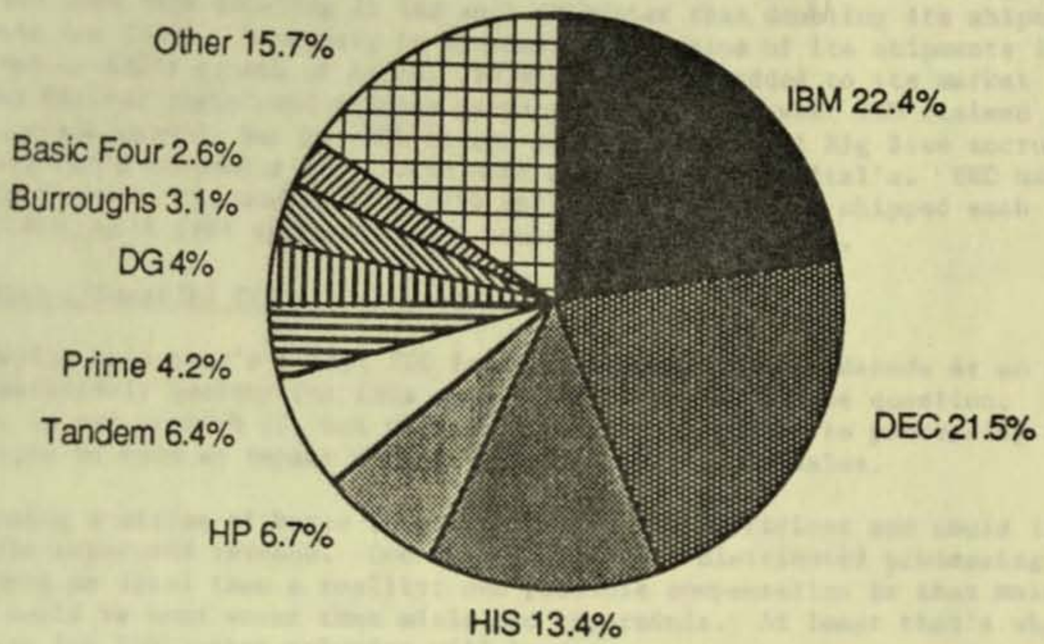

138,359 SYSTEMS

1984 MARKET SHARES OF U.S. MEDIUM-SCALE SHIPMENTS
UNITS (Left); DOLLARS (Right)
(Copyright 1985 -- International Data Corporation)

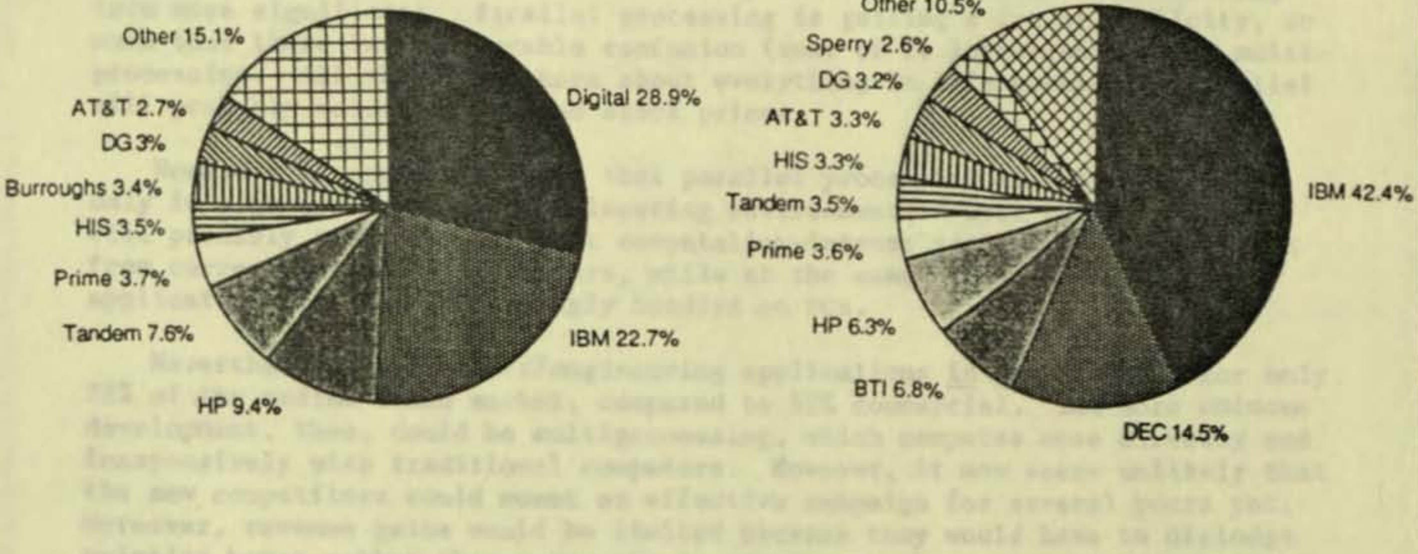

23,125 UNITS

$8.5 BILLION

IBM and DEC continue to dominate the mid-scale market, with DEC winning out on number of systems shipped and IBM cleaning up on revenue earned. Both increased in 1984 at rates faster than the industry average for revenue growth, with IBM more than doubling it (as well as better than doubling its shipment revenues for 1983). Moreover, by increasing the value of its shipments 114.5% compared to DEC's growth of 63.1%, IBM significantly added to its market share. In 1983 Digital registered revenue gains worth 85.8% of what IBM claimed in the medium-scale market, but in 1984 it got only 65.3% of what Big Blue accrued. However, IBM's shipments are always more cyclical than Digital's. DEC has been able to increase the number of 11/750 and up systems it has shipped each year since 1980, with 1984 shipments up a healthy 61.7% over 1983.

New Multi-, Parallel Processing Options

Despite this year's slump, IDC feels that the next half-decade or so should be comparatively healthy for this segment of the market. The question, though, is not so much if, but when alternative approaches to processing data will begin to have an impact on traditional medium-scale sales.

Running a series of boxes in tandem can be cost-efficient and could leach away some supermini revenue. One drawback is that distributed processing is still more an ideal than a reality; one possible compensation is that mainframes could be hurt worse than minis and superminis. At least that's what DEC hopes its VAXCluster solution will lead to.

More serious threats will likely come from multiprocessing and parallel processing, which would distribute the processing among chips rather than different systems. The multiprocessing machines from Encore produced much of what little enthusiasm there was at this year's NCC. From 2 to 20 processors can execute multiple jobs simultaneously. Meanwhile other companies, including Sequent and Alliant, are coming out with systems that they call true parallel processing, which can speed up data manipulation exponentially by running the same program simultaneously on different processor chips.

There are long-term and short-term implications for both, with the long-term more significant. Parallel processing is getting a lot of publicity, so much that there is considerable confusion (some of it intentional) with multiprocessing. All those metaphors about everything in life working in parallel will probably raise capital and stock prices.

However, it appears for now that parallel processing is cost-efficient only in high-end scientific/engineering environments. Over the long run it will probably take over the most computation-intense segments of CAD and CAE from current medium-scale vendors, while at the same time less involved applications will be increasingly handled on PCs.

Nevertheless, scientific/engineering applications in toto account for only 28% of the medium-scale market, compared to 52% commercial. The more ominous development, then, could be multiprocessing, which competes more directly and inexpensively with traditional computers. However, it now seems unlikely that the new competitors could mount an effective campaign for several years yet. Moreover, revenue gains would be limited because they would have to dislodge existing boxes rather than enter virgin territory, the way that minis and PCs did in their glory days. And there's still time for traditional vendors to hop on the bandwagon if they choose.

1984 MARKET SHARES OF WORLDWIDE INSTALLED BASE, SMALL-SCALE SYSTEMS
(Copyright 1985 -- International Data Corporation)

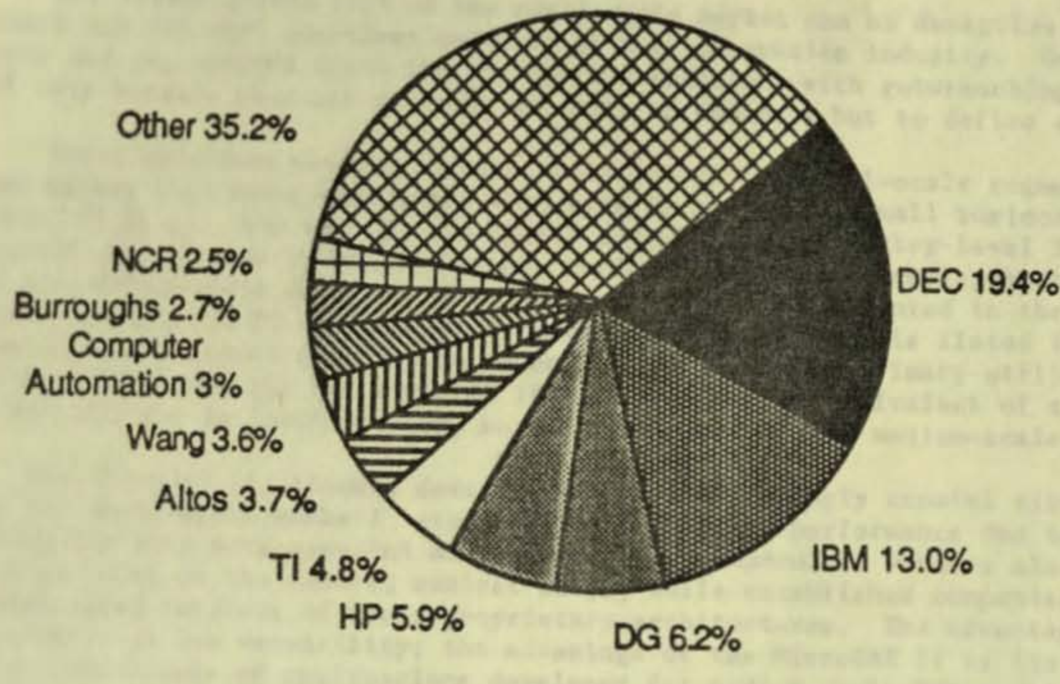

1,861,713 UNITS

1984 MARKET SHARES OF U.S. SMALL-SCALE SHIPMENTS
UNITS (Left); DOLLARS (Right)
(Copyright 1985 -- International Data Corporation)

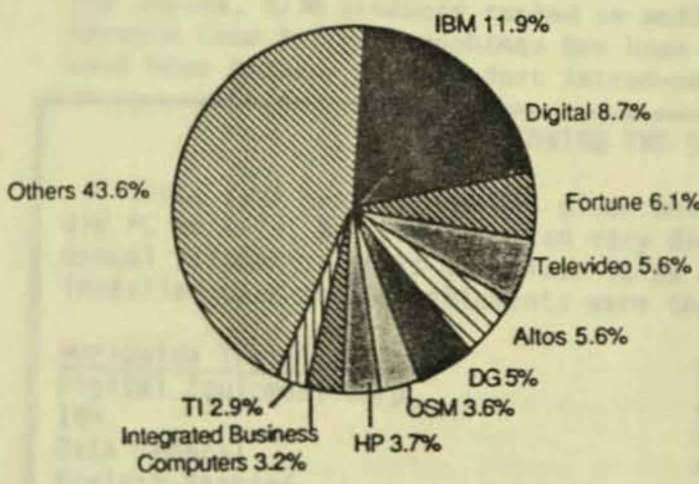

255,540 UNITS

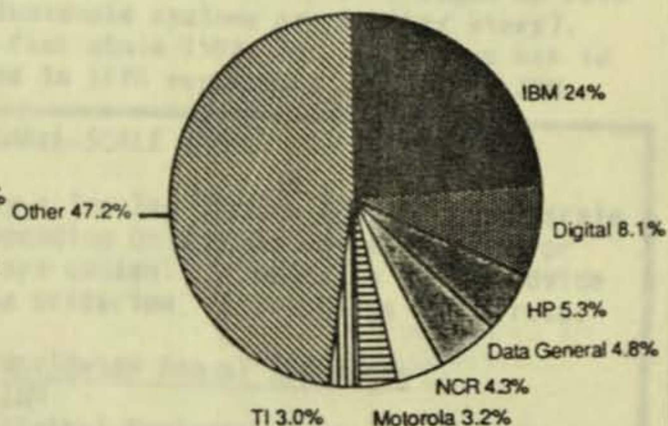

$4.9 BILLION

Bigger Bangs for the Buck in the Small-Scale Market

The steady growth rate of the small-scale market can be deceptive, for there are few more turbulent segments of the information industry. Companies come and go, vendors crack the Top Ten and slide out with gutwrenching speed, if only because they are striving not only to dominate but to define a market.

There sometimes seem to be more varieties to the small-scale segment of the market than Heinz has soups, what with 16-bit minis, small business systems from IBM et al., the new 16- and 32-bit supermicros, and entry-level 32-bit superminis. To add to the confusion, there is conceptual overlap both with the PC and medium-scale markets. Like supermicros that are counted in the small-scale market, the PC-AT can support several terminals, but is listed in IDC's census as a personal computer based on what is still its primary utilization. At the other end, the new MicroVAX II is virtually the equivalent of a 780 in power, but not in functionality, so can't be counted as a medium-scale machine.

The MicroVAX II, though, does point in an increasingly crucial direction for the small-scale market: stunning gains in price/performance due to falling prices for ever more powerful microprocessors. Supermicro vendors aim for open architectures on the latest, sexiest chips, while established companies pack miniaturized versions of their proprietary architectures. The advantage of the supermicro is its versatility; the advantage of the MicroVAX II is its access to the shelfloads of applications developed for medium-scale DEC machines (All-in-1 is now available). Whether new supermicro or established-vendor-on-a-chip, the microprocessor-based small-scale system offers a low cost/high performance combination that is ideally suited to specialized marketing by VARs.

IBM has maintained its position in the small-scale market almost solely on the basis of System/36 sales. In one year S/36 revenue jumped from 15% of the dollar value of IBM's small-scale shipments to over two-thirds. Much of this growth has been siphoned from lower-end S/38 sales, but IBM need not weep: in two years S/36 raked in more money than small-scale S/38 models brought in five (of course, S/38 products ranked as medium-scale systems are another story). Revenue from Series 1 machines has been flat since 1982, but for income not to have been falling on a product introduced in 1976 represents a gain for IBM.

TOP-TENNING THE SMALL-SCALE MARKET

Since last year EDP/IR has given only a Top Ten listing for the small-scale and PC markets, but rankings can vary depending on whether installed base or annual shipments, units or dollar value are chosen. If units in the worldwide installed base or 1984 shipments were the criterion, the rankings would read:

Worldwide Installed Base	Worldwide Annual Shipments
Digital Equipment Corp.	IBM
IBM	Digital Equipment Corp.
Data General	Altos Computer
Hewlett-Packard	Fortune Systems
Texas Instruments	Televideo
Altos Computer	Data General
Wang	Hewlett-Packard
Computer Automation Inc.	OSM Computer
Burroughs	Texas Instruments
NCR	NCR

Unix Percolates Through the Mid-Markets

Burroughs and Sperry may not have been able to agree on merger terms, but they do think alike on some questions. They came up with very similar strategies at this year's NCC in Chicago, notably their decisions to emphasize Unix.

In the large-scale and PC market sectors, where software compatibility is a critical issue, vendors brag about their technical ties to IBM products. In the two mid-markets of medium- and small-scale systems, on the other hand, the trend is to talk about Unix. The comparative dearth of commercial programs written for the Unix operating system is not as detrimental in mid-market sectors, where programmers abound who find Unix relatively easy to write for. Also, since computer professionals tend to run small- and medium-scale systems, what many feel to be Unix's user unfriendliness need not be a major drawback. Meanwhile Unix's portability can be incorporated in users' networking plans.

Deciding to accentuate Unix, though, need not mean taking the same approach to it. There are essentially two attitudes towards Unix on the part of major vendors: support, more or less grudging, and endorsement, more or less enthusiastic. The chief reason for vendors' diffidence is their commitment to proprietary operating systems, which they often feel are superior to Unix in the markets in which they compete. In any event, sales have been made on the basis of alleged superiority to others' (typically IBM's) operating systems, and the established customer base must be supported technically and psychologically.

For all of the references to Unix in Burroughs' NCC booth, it still seems to fall into the more-or-less-grudging-support category, especially if one starts to talk with the company's marketing personnel. Wang now has IN/ix, billed as the first in a series of Unix offerings bridging its PC and VS lines, but you would hardly know it from reading the company's press releases. The most reluctant support for Unix probably comes from IBM, which is loath to add asteriks to its announcements noting that Unix is a trademark of AT&T. Nevertheless, Big Blue's new IX/370 offers Unix as a guest system for Unix users who want to use a System/370 machine, including any in the 43XX medium-scale line, and don't mind not being able to exchange files with regular IBM systems.

Despite the importance of its large customer base, Sperry is gambling on the more-or-less-enthusiastic-endorsement approach to win new sales. The single message that Sperry wanted visitors to NCC to take away was the company's "Focus on Unix." Data General is touting its two Unix operating systems as the strongest Unix offering around. About 20% of VAX machines currently run on Unix, with some at DEC estimating that the MicroVAX could be 40-50% Unix.

Why all this motion and emotion? What is generating fervor from some, and at least reluctant recognition from the rest? One factor is the possiblity, however remote, of using IBM's own proprietary operating systems to rope it off from the new industry standard. In a related development, American vendors are eying the European market, where in February the seven leading indigenous minicomputer companies joined to standardize their software on Unix.

A second factor of even more immediate consequence is the federal government's interest in sponsoring, and maybe imposing, an industry standard. Although different agencies and projects still vary on whether they want mere Unix support or full endorsement, the trend is clear: Unix is now a standard RFP (Request for Proposal) requirement for government contracts.

MEDIUM-SCALE COMPUTER MARKET -- U.S.-BASED MANUFACTURERS
(Copyright 1985 -- International Data Corporation)

	Number of Systems Shipped	Cumulative Number in Use	$ Million Value Shipped	$ Billion Value in Use
		WORLDWIDE		
1980	16,100	55,700	7,300	34.9
1981	19,200	68,000	9,600	39.0
1982	24,100	86,900	9,200	45.0
1983	23,800	106,300	9,600	49.7
1984	37,900	138,300	13,900	58.0
1985	37,800	168,200	13,900	65.1
1986	44,800	203,600	15,500	75.4
1987	52,800	242,900	17,300	86.9
1988	61,400	288,200	19,400	100.0
1989	69,600	337,600	21,700	114.4
1990	77,900	390,200	24,100	129.3
		UNITED STATES		
1980	9,300	30,200	4,500	17.2
1981	11,100	39,000	5,800	21.1
1982	14,100	50,500	5,600	24.5
1983	14,700	62,900	5,900	28.2
1984	23,100	82,800	8,300	33.3
1985	21,500	100,500	8,400	37.0
1986	25,900	122,400	9,200	43.5
1987	30,900	146,700	10,300	50.7
1988	36,300	175,100	11,500	58.9
1989	41,400	206,000	12,800	67.8
1990	46,300	239,000	14,200	77.1
		INTERNATIONAL		
1980	6,800	25,500	2,800	17.7
1981	8,100	29,000	3,800	17.9
1982	10,000	36,400	3,600	20.5
1983	9,100	43,400	3,700	21.5
1984	14,800	55,500	5,100	24.7
1985	16,300	67,700	5,500	28.1
1986	18,900	81,200	6,300	31.9
1987	21,900	96,200	7,000	36.2
1988	25,100	113,100	7,900	41.1
1989	28,200	131,600	8,900	46.6
1990	31,600	151,200	9,900	52.2

WORLDWIDE

1980	197,800	858,800	7,700	34.1
1981	212,500	1,042,500	8,800	40.8
1982	266,800	1,280,000	9,100	49.2
1983	313,600	1,543,700	9,800	56.9
1984	384,700	1,853,700	11,100	65.9
1985	418,200	2,190,300	12,300	74.8
1986	496,400	2,582,700	14,200	84.2
1987	582,700	3,025,300	16,400	96.1
1988	675,400	3,517,600	18,600	109.8
1989	770,700	4,051,600	21,000	125.0
1990	865,600	4,636,400	23,600	140.8

UNITED STATES

1980	122,400	566,500	4,600	20.4
1981	127,800	679,100	5,400	23.9
1982	172,800	828,700	5,700	29.4
1983	202,200	992,100	6,100	34.4
1984	262,000	1,201,800	6,900	40.6
1985	281,400	1,431,100	7,700	46.5
1986	337,800	1,705,200	9,000	53.1
1987	398,600	2,014,600	12,000	60.7
1988	465,600	2,357,100	12,000	69.4
1989	533,600	2,731,400	13,600	79.4
1990	600,100	3,145,800	15,400	89.6

INTERNATIONAL

1980	75,400	292,300	3,100	13.7
1981	84,700	363,400	3,400	16.9
1982	94,000	451,300	3,400	19.8
1983	111,400	551,600	3,700	22.5
1984	122,700	651,900	4,200	25.3
1985	136,800	759,200	4,600	28.3
1986	158,600	877,500	5,200	31.1
1987	184,100	1,010,700	5,900	35.4
1988	209,800	1,160,500	6,600	40.4
1989	237,100	1,320,200	7,400	45.6
1990	265,500	1,490,600	8,200	51.2

A gradual transformation of at least segments of the medium-scale market is therefore underway. IDC surveys indicate that the 10% of supermini sites currently running Unix on at least one of their machines should rise to almost a quarter of medium-scale sites by the end of 1986. To meet this rising demand, about half of all packaged software vendors plan to be offering Unix programs by the end of this year, about double the percentage at yearend 1984.

The changes may also lead to a rearrangement of the medium-scale hierarchy, although the jury is still out on the question. Speculation centers on AT&T, which has hitched its wagon to the Unix star far more singlemindedly than even Sperry. The communications giant garnered much publicity at the end of June when the National Security Agency chose its 3B systems as the centerpiece of a contract that should net AT&T almost $1 billion for sales and services. Certainly the deal lends credibility to AT&T, which with one stroke of the pen will increase its January 1, 1985 installed base by about 22%, and to Unix, which has garnered its gaudiest government seal of approval to date.

On the other hand, the increase to AT&T's installed base just by adding up to 250 3B systems also emphasizes its failure so far to penetrate non-BOC commercial markets. The government's interest in Unix may represent AT&T's best hope to establish itself as a major player in the computer industry. Even its NSA success, however, may not translate easily into other government contracts, since the secretive intelligence agency has unique priorities.

For the immediate future, last year's "Review and Forecast" assessment (EDP/IR, Vol. 20, No. 15, 16, December 12, 1984) remains valid: not medium-scale but "small-scale systems are likely to be the first major Unix success story." Vendors in this segment of the market tend to be smaller and younger, which means that they typically lack long-standing customer bases wedded to proprietary chips and/or operating systems.

As a full-fledged member of the Bunch, NCR looked for a time to be an exception. Besides its successful Tower line of Unix supermicros, which in an OEMed transfiguration serves as a linchpin of the Sperry Unix group, NCR was working on a Unix-based 32-bit proprietary processor that could have extended the lifespan of its 9300 line. It has now given up the endeavor as too costly, however, and seems by default to have anointed the open-architecture Tower system, based on commercially available Motorola chips, as its royal family.

NCR's failure to develop a proprietary chip may be a blessing in disguise. The Unix small-scale success story is based in large part on cost/performance advantages, which in turn depend on the wide availability of Unix microprocessors. The Tower team knew what it was doing when it tapped Motorola.

In the meantime, small-scale Unix vendors are competing vigorously, with enough of them winning sales to raise the hopes and stoke the efforts of all the others. Last year's sweepstakes winner for start-ups in this sector was Altos, which ascended to the top-ten list of small-scale vendors based on the value of 1984 shipments, while Tower sales were a major contributor to NCR's rise from sixth place to fourth place.

This year there is no end to would-be successors to the crown. At NCC a giant inflated frog beside the Charles River Data Systems booth represented virtually the only aggressive marketing at a staid convention and punctuated its promise to become "the largest shipper of 68020-based systems in 1985."

MEDIUM-SCALE COMPUTER CENSUS

(Copyright 1985 -- International Data Corporation)

Medium-Scale computer systems generally support 17 to 128 users in a normal commerical environment. Included are the IBM 4300 line, the mid-range systems from the Bunch, and the high end of lines from companies such as DEC, DG and Prime. System prices tend to fall in the $100,000 to $1 million range.

NAME OF MANUFACTURER	COMPUTER MODEL	SIZE CLASS	PURCHASE PRICE IN DOLLARS (000) BASIC AVERAGE	DATE OF FIRST INSTALLATION	NUMBER INSTALLED IN U.S.	NUMBER INSTALLED OUTSIDE U.S.	TOTAL NUMBER INSTALLED
AT&T	3820	M	330	4Q/83	200	20	220
	3820 Duplex	M	600	4Q/83	670	20	690
	3B5/100,/200,/300	M	125	03/84	150	90	240
				TOTAL	1,020	130	1,150
Auragen Systems Corp.	4000	M	135	01/83	10	15	25
Basic Four Corporation	700	M	130	08/77	100	95	195
	730	M	124	06/78	525	590	1,115
	MAI 8000 Series (810)	M	150	04/82	1,103	1,214	2,317
				TOTAL	1,728	1,899	3,627
BTI	8000	M	175	05/81	51	10	61
Burroughs	2500	M	499	02/67	1	2	3
	2700	M	530	08/72	17	15	32
	2800	M	471	12/76	160	100	260
	2900	M	533	07/80	385	225	610
	2925	M	369	03/83	450	350	800
	3500	M	912	05/67	75	50	125
	3700	M	1,230	11/72	100	105	205
	3800	M	705	--/77	160	35	195
	3955	M	778	09/81	205	45	250
	45/4700	M	1,560	10/71	90	45	135
	4800	M	1,440	09/76	315	110	425
	4925	M	500	2Q/84	100	40	140
	4955	M	1,805	08/83	200	105	305
	55/5700	M	1,680	03/63	0	0	0
	5920	M	448	--/82	138	80	218
	5930	M	579	01/81	205	245	450
	A3-D,A3-F	M	270	10/84	20	10	30
	A3-K	M	250	4Q/85	0	0	0
	A9-B	M	0	--/85	0	0	0
	A9-D	M	0	--/85	0	0	0
	A9-F	M	1,200	2Q/84	85	85	170
				TOTAL	2,706	1,647	4,353
Cambex	1636-1	M	100	10/81	6	0	6
	1636-10	M	99	01/83	5	0	5
	1641	M	178	10/80	9	15	24
	1641-11	M	181	--/83	5	9	14
	1651	M	213	--/82	7	7	14
	1651-11	M	256	--/83	4	1	5
				TOTAL	36	32	68
Control Data Corp	180-840	M	760	03/85	0	0	0
	180-850	M	1,115	03/85	0	0	0
	180-860	M	1,575	03/85	0	0	0
	180-990	M	0	06/85	0	0	0
	31/3150	M	419	12/64	10	5	15
	3170	M	970	12/70	5	2	7
	3200	M	625	05/64	6	5	11
	3300	M	1,491	09/65	40	15	55
	3400	M	1,019	11/64	0	0	0
	3500	M	1,341	01/69	20	10	30
	830, 830 Dual	M	395	--/84	2	0	2
	Cyber 170-720	M	1,625	05/79	50	85	135
	Cyber 170-810	M	250	07/84	12	8	20
	Cyber 170-815	M	398	01/83	20	35	55
	Cyber 170-825	M	949	03/82	25	30	55
	Cyber 170-835	M	1,940	03/82	10	15	25
	Cyber 170-845	M	1,632	12/83	50	40	90
	Cyber 171	M	1,184	08/77	3	32	35
	Cyber 71	M	1,111	07/76	0	0	0
	Cyber 72	M	1,810	10/71	4	12	16
	G-20	M	830	04/61	1	0	1
	Omega 480-1	M	376	06/77	20	0	20
	Omega 480-2	M	391	08/78	30	3	33
	Omega 480-3	M	448	01/80	30	2	32
				TOTAL	338	299	637
Convex Computer	C-1	M	890	11/84	4	0	4
Data General	Eclipse C/300, 330	M	150	02/75	723	175	898
	Eclipse C/350	M	180	11/78	560	190	750
	Eclipse M/600	M	215	05/78	235	90	325
	Eclipse MV/10000	M	520	--/83	460	130	590
	Eclipse MV/6000	M	200	02/82	365	150	515
	Eclipse MV/8000	M	300	10/80	930	540	1,470
	Eclipse S/250	M	150	11/78	645	345	990
	MV/10000 SX	M	550	03/85	0	0	0
				TOTAL	3,918	1,620	5,538

MEDIUM-SCALE COMPUTER CENSUS AS OF JANUARY 1, 1985
(Copyright 1985 -- International Data Corp.)

NAME OF MANUFACTURER	COMPUTER MODEL	SIZE CLASS	PURCHASE PRICE IN DOLLARS (000) BASIC AVERAGE	DATE OF FIRST INSTALLATION	NUMBER INSTALLED IN U.S.	NUMBER INSTALLED OUTSIDE U.S.	TOTAL NUMBER INSTALLED
Digital Equipment Corp.	1040/50	M	694	10/67	35	15	50
	1060/70	M	880	06/72	60	30	90
	2040/50	M	756	06/76	50	25	75
	2060	M	1,013	07/79	265	135	400
	PDP-15 XVM	M	82	02/70	350	127	477
	PDP-6	M	400	10/64	0	0	0
	VAX 11/785	M	440	--/84	600	100	700
	VAX 8600	M	0	--/85	0	0	0
	VAX-11/750	M	100	--/80	9,300	7,040	16,340
	VAX-11/780	M	320	05/78	6,700	4,650	11,350
	VAX-11/782	M	600	--/82	190	120	310
				TOTAL	17,550	12,242	29,792
Digital Scientific Corp.	Meta 4/5030	M	150	06/79	8	1	9
Elxsi	Elxsi 6400	M	475	12/83	32	8	40
Flexible Computer	Flex/32	M	126	--/85	0	0	0
Formation	F/4000	M	134	08/81	53	22	75
Global U.S.I.	MBO/3	M	195	01/79	12	1	13
	MBO/30	M	89	08/81	21	0	21
	MBO/30E	M	117	12/81	1	0	1
	MBO/31	M	140	06/80	70	5	75
	MBO/4	M	310	12/78	20	2	22
	MBO/43	M	254	03/81	55	5	60
	USE 44	M	210	07/84	3	2	5
	USX 40	M	150	07/84	25	0	25
	USX 43	M	185	07/84	19	0	19
	USX 46	M	250	07/84	3	0	3
				TOTAL	229	15	244
Gould Computer Systems	32/75,/77	M	150	03/78	1,140	279	1,419
	32/87	M	400	11/81	332	80	412
	32/97	M	450	--/84	80	20	100
	Concept 32/67	M	250	04/83	200	50	250
	Powernode 6000	M	250	02/84	10	5	15
	Powernode 9000	M	450	06/84	7	3	10
				TOTAL	1,769	437	2,206
Harris	H100	M	150	01/71	755	33	788
	H1000	M	525	06/84	40	10	50
	H300	M	200	09/81	85	17	102
	H500	M	250	01/79	375	81	456
	H700	M	175	03/83	401	84	485
	H80	M	100	02/81	80	28	108
	H800 Series	M	275	03/80	555	65	620
	Mind	M	140	08/82	180	85	265
	Series 100	M	190	10/75	35	1	36
	Series 1X5	M	175	09/77	85	15	100
	Series 200	M	240	10/75	48	1	49
	Series 5X0	M	300	12/78	88	8	96
	Slash 7	M	180	10/75	5	0	5
				TOTAL	2,732	428	3,160
Hewlett-Packard	3000 I, II	M	200	11/72	350	250	600
	3000 III	M	200	06/78	2,220	970	3,190
	3000-48	M	194	12/83	1,700	1,340	3,040
	3000-64	M	325	01/82	945	696	1,641
	3000-68	M	375	12/83	480	380	860
				TOTAL	5,695	3,636	9,331
HIS	(GE) 3010/3012	M	90	05/70	3	1	4
	116/316/416/516	M	62	10/66	140	90	230
	6/43,47,53,57,48, 54,74,76	M	177	01/77	7,150	5,721	12,871
	6/75	M	155	30/83	210	180	390
	6/92,96	M	552	10/81	177	246	423
	6/95	M	220	30/83	140	110	250
	Comshare/Xerox	M	612	06/79	13	0	13
	DPS 7	M	512	--/81	130	687	817
	DPS 8/20	M	563	04/80	40	15	55
	DPS 8/44	M	881	01/80	65	45	110
	DPS 8/44 Dual	M	1,224	--/83	11	3	14
	DPS 8/47	M	510	03/83	70	35	105
	DPS 8/49	M	816	03/83	60	35	95
	DPS 8/52	M	2,034	11/80	91	86	177
	DPS 8/62	M	2,401	12/81	28	9	37
	G-200	M	425	04/61	25	5	30
	G-400	M	700	05/64	45	45	90
	G-6023/25	M	1,490	11/73	5	6	11
	G-6030/40	M	1,900	06/71	5	40	45
	Gamma-30	M	425	02/62	0	0	0
	H-1015/2015	M	805	02/71	5	4	9
	H-1200	M	629	02/66	50	15	65
	H-125	M	365	12/67	2	10	12
	H-2040	M	595	04/72	30	25	55
	H-2050/60	M	920	04/72	60	40	100
	H-2070	M	1,600	05/72	10	0	10
	H-2200	M	862	01/66	20	5	25
	H-3200	M	1,225	02/70	25	10	35
	H-4/1400	M	560	12/61	3	2	5
	H-4200	M	1,920	09/68	0	0	0
	H-64/300	M	413	04/79	90	810	900
	H-66/05/07	M	920	10/76	35	70	105
	H-66/10/17	M	1,504	06/75	23	60	83
	H-66/20/27	M	2,025	11/74	50	100	150
	H-66/440/520	M	932	12/79	70	35	105
	H-8/1800	M	1,500	12/60	2	0	2
	H-Level 64	M	605	10/74	208	900	1,108
				TOTAL	9,091	9,445	18,536

MEDIUM-SCALE COMPUTER CENSUS AS OF JANUARY 1, 1985
(Copyright 1985 -- International Data Corp.)

NAME OF MANUFACTURER	COMPUTER MODEL	SIZE CLASS	PURCHASE PRICE IN DOLLARS (000) BASIC AVERAGE	DATE OF FIRST INSTALLATION	NUMBER INSTALLED IN U.S.	NUMBER INSTALLED OUTSIDE U.S.	TOTAL NUMBER INSTALLED
IBM	1410	M	800	11/61	4	1	5
	1460	M	460	10/63	5	1	6
	360/30	M	528	06/65	180	60	240
	360/40	M	895	06/65	200	65	265
	360/44	M	880	07/66	15	5	20
	360/50	M	1,550	08/65	90	30	120
	370/125	M	652	06/78	90	240	330
	370/135	M	960	05/72	180	165	345
	370/138	M	997	11/76	350	560	910
	370/145	M	1,871	07/71	150	155	305
	370/148	M	1,700	03/77	400	430	830
	4331-2/11	M	394	03/82	2,200	1,675	3,875
	4341-1/9/10	M	830	10/79	4,200	4,600	8,800
	4341-2/11/12	M	969	03/81	4,000	2,400	6,400
	4361-4,-5	M	465	--/84	1,200	1,100	2,300
	4381-2, 1	M	900	--/84	1,900	500	2,400
	7010	M	1,200	10/63	0	0	0
	7040/44	M	1,439	06/63	0	0	0
	7070/74	M	1,299	03/68	0	0	0
	System 38-7	M	405	06/82	1,000	980	1,980
	System 38-8	M	408	07/83	1,000	860	1,860
				TOTAL	17,164	13,827	30,991
IPL	4436	M	176	11/80	16	0	16
	4443	M	223	03/80	35	0	35
	4445	M	385	12/82	3	0	3
	4446	M	313	12/81	115	16	131
	4460	M	400	--/83	22	3	25
				TOTAL	191	19	210
McDonnell Douglas Computer Systems	9000	M	150	--/81	510	90	600
	9100	M	175	01/85	0	0	0
	9208	M	300	10/84	50	0	50
				TOTAL	560	90	650
MDS Qantel	64 Series	M	225	04/83	230	45	275
Modular Computer Systems	Classic 32/85	M	306	--/84	15	20	35
	Classic 7870	M	159	08/78	410	155	565
	Classic II/75	M	156	05/82	233	122	355
	Modcomp IV	M	100	05/74	339	130	469
				TOTAL	997	427	1,424
Motorola/Four Phase	5000 Series	M	200	02/84	50	30	80
	System 311	M	265	02/81	145	15	160
	System 312	M	320	12/81	52	8	60
	Two PI V/32	M	125	06/78	80	6	86
				TOTAL	327	59	386
NAS	7020/30	M	836	01/79	0	6	6
	AS/3	M	624	11/78	20	25	45
	AS/3000	M	331	01/80	35	0	35
	AS/4	M	916	06/77	10	4	14
	AS/5000	M	608	02/80	70	50	120
	AS/6100	M	474	10/82	3	5	8
	AS/6620	M	373	--/83	52	22	74
	AS/6630	M	496	--/83	21	70	91
	AS/6650	M	618	--/83	53	38	91
	AS/6660	M	475	10/84	2	0	2
				TOTAL	266	220	486
NCR	315 RMC	M	500	09/65	0	0	0
	8560	M	745	11/77	35	23	58
	8565	M	585	04/79	90	145	235
	8565 II	M	469	04/82	75	94	169
	8570	M	953	05/76	80	75	155
	8575	M	729	04/79	95	70	165
	8575 II	M	548	05/82	59	68	127
	8580	M	1,734	12/78	4	8	12
	8585	M	1,472	08/79	30	50	80
	8585 II	M	585	05/82	34	33	67
	8595 II	M	742	05/82	36	37	73
	8635,8645,8655	M	1,300	--/83	47	29	76
	Century 200	M	367	09/69	25	28	53
	Century 201	M	590	09/69	45	48	93
	Century 251	M	1,005	11/73	5	8	13
	Century 300	M	1,426	02/72	26	32	58
				TOTAL	686	748	1,434

MEDIUM-SCALE COMPUTER CENSUS AS OF JANUARY 1, 1985
(Copyright 1985 -- International Data Corp.)

NAME OF MANUFACTURER	COMPUTER MODEL	SIZE CLASS	PURCHASE PRICE IN DOLLARS (000) BASIC AVERAGE	DATE OF FIRST INSTALLATION	NUMBER INSTALLED IN U.S.	NUMBER INSTALLED OUTSIDE U.S.	TOTAL NUMBER INSTALLED
Norsk Data North American	ND-500 Series	M	220	10/84	50	530	580
Perkin-Elmer	3200MPS	M	246	--/83	100	18	118
	3240	M	159	10/79	348	94	442
	3250/XP	M	228	02/82	277	148	425
	8/32	M	206	06/75	879	286	1,165
	Cadam	M	220	05/78	75	5	80
				TOTAL	1,679	551	2,230
Prime	500	M	330	06/76	40	40	80
	550	M	150	02/79	950	625	1,575
	550 II	M	110	07/81	565	265	830
	650	M	165	02/79	80	55	135
	750	M	270	07/79	1,270	450	1,720
	850	M	370	07/81	305	120	425
	9650	M	146	02/84	100	50	150
	9750	M	252	07/84	150	75	225
	9950	M	520	08/83	455	245	700
	9955	M	600	01/85	0	0	0
				TOTAL	3,915	1,925	5,840
Pyramid Technology Corp.	90MX	M	0	01/85	0	0	0
	90X	M	300	--/84	80	20	100
				TOTAL	80	20	100
Sperry Univac	1100/10	M	2,094	04/76	50	70	120
	1100/20	M	2,850	04/75	2	10	12
	1100/60	M	1,941	03/80	600	1,017	1,617
	1100/70/71	M	500	--/84	137	193	330
	301	M	295	10/60	2	0	2
	3301	M	1,021	06/64	8	0	8
	418 I, II	M	450	06/63	28	25	53
	418 III	M	1,700	06/69	0	20	20
	491/492	M	1,150	12/61	0	0	0
	501	M	695	06/59	0	0	0
	90/40	M	765	06/78	85	130	215
	90/60	M	1,184	01/74	90	100	190
	90/70/9700	M	1,275	08/72	4	15	19
	90/80	M	3,254	12/76	80	50	130
	9400/80	M	530	05/69	80	50	130
	III	M	1,134	08/62	5	1	6
	Spectra 70/2	M	620	07/71	0	0	0
	Spectra 70/3	M	1,040	10/71	0	0	0
	Spectra 70/35	M	634	01/67	11	2	13
	Spectra 70/45	M	1,181	11/65	25	0	25
	Spectra 70/46	M	1,963	11/68	0	0	0
	Spectra 70/55	M	1,800	11/68	0	0	0
				TOTAL	1,207	1,683	2,890
Stratus Computer	FT 200	M	200	02/82	187	67	254
	FT 240	M	220	03/84	4	1	5
	XA 400	M	275	10/84	8	3	11
	XA 440	M	300	03/84	6	1	7
	XA 600	M	350	05/84	27	14	41
				TOTAL	232	86	318
Synapse	N+1	M	600	12/82	8	0	8
Tandem	Non Stop, Nonstop II, TXP	M	170	05/76	6,475	2,360	8,835
Wang	VS-300	M	310	--/85	0	0	0
	VS-100	M	270	01/81	1,026	730	1,756
	VS-90	M	166	01/82	670	320	990
				TOTAL	1,696	1,050	2,746
Xerox	550/560	M	631	11/74	14	1	15
	940	M	1,040	04/66	25	0	25
	Sigma 5	M	585	08/67	25	3	28
	Sigma 6	M	1,030	11/70	10	5	15
	Sigma 7	M	1,170	12/66	15	2	17
	Sigma 8	M	920	09/71	0	0	0
				TOTAL	89	11	100
				GRAND TOTAL	82,822	55,537	138,359

SMALL-SCALE COMPUTER CENSUS

Census of top ten small-scale computer vendors, by dollar value of 1984 worldwide shipments. Small-scale computer systems generally range from 2-16 users in typical commercial applications. Representative systems include the IBM S/36, DEC PDP-11 and Vax 730, and systems such as those from Altos. Average systems prices are normally in the $10,000 to $100,000 range.

Name of Manufacturer	Computer Model	New Size Class	Average Purchase Price ($K)	Date of 1st Install.	Number Installed in U.S.	Number Installed Outside U.S.	Total Number Installed
IBM	1130	S	120	02/66	345	185	530
	1401	S	118	09/60	180	20	200
	1401 G-H	S	106	05/64	0	0	0
	1440	S	227	04/63	8	3	11
	1620	S	130	09/60	12	5	17
	360/20	S	136	12/65	190	70	260
	360/22	S	185	07/71	5	3	8
	360/25	S	391	10/68	12	40	52
	370/115	S	480	04/74	70	320	390
	4321, 4331-1	S	229	03/82	4,000	3,000	7,000
	8100	S	143	08/79	7,000	5,500	12,500
	Series 1	S	31	11/76	46,000	18,500	64,500
	SYS/23 Datamaster	S	13	07/81	25,700	9,000	34,700
	System 3-10	S	123	01/70	1,575	500	2,075
	System 3-12	S	184	03/78	500	520	1,020
	System 3-15	S	341	03/74	1,600	810	2,410
	System 3-4	S	62	06/76	40	20	60
	System 3-6	S	76	12/70	175	45	220
	System 3-8	S	107	06/75	310	410	720
	System 38-3	S	149	07/80	850	690	1,540
	System 38-4	S	160	08/81	1,800	1,625	3,425
	System 38-5	S	298	07/80	4,550	3,550	8,100
	System 38-6	S	120	06/84	1,200	75	1,275
	System/32	S	42	03/75	3,600	3,500	7,100
	System/34	S	65	12/77	33,000	25,500	58,500
	System/36	S	57	10/83	20,000	16,000	36,000
				TOTAL	152,722	89,891	242,613
Digital Equipment Corp.	2020	S	333	07/78	230	190	420
	MICRO PDP #11/SV	S	6	06/85	0	0	0
	MICRO PDP-11/73	S	25	12/84	24	0	24
	PDP 11/03,23 LSI-11	S	10	06/75	76,300	45,000	121,300
	PDP-11/04,05,10,15, 20	S	17	03/70	32,480	18,170	50,650
	PDP-11/24	S	37	10/81	7,760	6,930	14,690
	PDP-11/34, 35, 40	S	40	01/75	32,505	24,000	56,505
	PDP-11/44	S	60	06/79	17,200	10,295	27,495
	PDP-11/45, 50, 55	S	91	04/72	3,800	1,180	4,980
	PDP-11/60	S	75	06/77	1,185	1,085	2,270
	PDP-11/70	S	175	07/75	8,270	3,650	11,920
	PDP-11/84	S	20	06/85	0	0	0
	PDP-12	S	46	04/69	200	15	215
	PDP-8/a, e, f, m	S	15	04/71	42,000	20,900	62,900
	PDP-8/1, 1, s	S	27	05/65	500	300	800
	PDP-9,9L	S	50	12/66	0	5	5
	VAX-11/725	S	30	10/83	650	450	1,100
	VAX-11/730	S	42	06/82	3,200	2,350	5,550
				TOTAL	226,304	134,570	360,824
Hewlett-Packard	1000-A/600	S	20	02/82	7,820	4,850	12,670
	1000-A/700	S	30	02/82	1,414	986	2,400
	1000-A/900	S	50	02/82	2,605	1,735	4,340
	1000-E (MX-E)	S	41	01/77	15,363	15,337	30,700
	1000-F	S	56	09/78	5,961	4,426	10,387
	1000-L, XL	S	10	06/79	1,830	1,560	3,390
	1000-M (MX-M, K)	S	26	06/74	8,556	6,687	15,243
	2100A	S	15	06/71	2,520	1,810	4,330
	2100S	S	36	08/73	635	810	1,445
	2114/15/16	S	29	11/66	1,110	205	1,315
	260-15	S	0	06/85	0	0	0
	260-24	S	0	06/85	0	0	0
	260-55	S	0	06/85	0	0	0
	300	S	58	09/79	190	280	470
	3000-30	S	80	10/79	1,067	455	1,522
	3000-33	S	122	01/79	728	300	1,028
	3000-40	S	100	11/81	1,686	1,274	2,960
	3000-42	S	113	12/83	543	428	971
	3000-44	S	170	03/81	2,790	760	3,550
	9020, 9030, 9040	S	35	01/83	3,000	2,000	5,000
	250	S	28	09/78	2,319	4,759	7,078
	3000-37	S	35	12/84	118	0	118
	3000-39	S	50	01/83	925	540	1,465
				TOTAL	61,180	49,202	110,382
NCR	315	S	164	05/62	0	0	0
	8130, 40, 50	S	23	02/77	925	3,580	4,505
	8200,30/31,50/51	S	55	09/74	5,800	6,870	12,670
	8270	S	93	11/79	435	640	1,075
	8350	S	164	01/77	20	2	22
	8370	S	287	07/78	4	0	4
	8410	S	160	02/79	105	180	285
	8430	S	264	06/78	75	210	285
	8450	S	386	11/77	110	115	225
	8455	S	275	04/79	192	350	542
	8535 II	S	234	04/82	55	80	135
	8545 II	S	336	04/82	39	70	109
	8550	S	633	04/76	125	90	215
	8555	S	442	04/79	220	215	435
	8555 II	S	426	04/82	240	275	515
	9010	S	18	04/81	184	2,800	2,984
	9020	S	29	03/81	3,400	6,600	10,000
	9040	S	122	05/81	375	1,430	1,805
	9050	S	153	05/81	250	350	600
	9300	S	60	06/83	3,000	2,000	5,000
	Century 100	S	167	09/68	45	40	85
	Century 101	S	214	10/72	145	180	325
	Century 151	S	308	02/75	50	60	110
	Century 50	S	101	02/71	25	0	25
	Century 75	S	122	07/76	12	15	27
	Tower Series	S	29	06/83	2,000	1,000	3,000
	9100	S	23	06/84	200	1,400	1,600
				TOTAL	18,111	28,552	46,663

SMALL-SCALE COMPUTER CENSUS AS OF JANUARY 1, 1985
(Copyright 1985 -- International Data Corp.)

Name of Manufacturer	Computer Model	New Size Class	Average Purchase Price ($K)	Date of 1st Install.	Number Installed in U.S.	Number Installed Outside U.S.	Total Number Installed
Data General	DCC-112	S	12	08/70	50	40	90
	DCC-116	S	10	11/71	3,500	900	4,400
	DCC-416/616	S	7	09/75	20	0	20
	DCC-Model 5	S	6	02/76	15	0	15
	Desktop Generation 20	S	10	07/83	7,000	500	7,500
	Desktop Generation 30	S	12	07/83	700	150	850
	Eclipse 280	S	92	06/83	820	430	1,250
	Eclipse C/150	S	122	12/78	510	300	810
	Eclipse MV/4000	S	71	12/82	2,005	975	2,980
	Eclipse S100, 130, 140, AP/130	S	31	02/75	7,110	3,605	10,715
	Eclipse S200, 230	S	56	02/75	4,317	1,519	5,836
	Micronova System	S	9	11/76	3,500	1,370	4,870
	MP 100, 200, MPT 100	S	8	05/79	1,040	280	1,320
	Nova	S	12	02/69	50	10	60
	Nova 1200 Series	S	11	12/70	3,100	1,700	4,800
	Nova 2 Series	S	10	10/73	6,800	2,300	9,100
	Nova 3/12	S	18	01/76	17,070	5,255	22,325
	Nova 3/4	S	10	01/75	3,405	1,070	4,475
	Nova 3/8	S	31	10/76	2,345	2,145	4,490
	Nova 4/C, S, X	S	20	03/79	14,700	9,700	24,400
	Nova 800/820	S	40	03/71	200	110	310
	Nova 830/840	S	44	07/73	100	90	190
	S 20	S	12	06/82	1,490	740	2,230
	S120	S	20	06/82	1,110	720	1,830
	Supernova SC	S	27	04/70	0	2	2
				TOTAL	80,957	33,911	114,868
Texas Instruments	200	S	9	08/82	270	200	470
	300	S	14	01/83	6,200	3,500	9,700
	600	S	40	08/82	4,900	2,425	7,325
	800	S	54	01/83	1,200	570	1,770
	960-A/B	S	30	12/71	2,750	1,860	4,610
	980-A/B	S	40	10/71	2,600	1,800	4,400
	990/10	S	23	03/76	21,815	9,895	31,710
	990/12	S	52	11/79	4,230	2,885	7,115
	990/4	S	8	06/76	8,250	2,930	11,180
	990/5	S	10	03/79	7,685	3,190	10,875
				TOTAL	59,900	29,755	89,155
Wang	2200	S	19	04/73	10,900	9,370	20,270
	2200 SVP, LVP	S	20	09/79	9,130	7,195	16,325
	2200 VP, MVP	S	32	09/76	11,463	9,112	20,575
	VS-15	S	29	08/83	500	300	800
	VS-25	S	46	12/82	1,425	1,340	2,765
	VS-45	S	87	10/82	1,240	1,030	2,270
	VS-80	S	106	07/77	2,435	1,930	4,365
	VS-85	S	73	04/84	300	200	500
				TOTAL	37,393	30,477	67,870
Burroughs	1710	S	132	08/72	155	195	350
	1720	S	350	02/73	100	95	195
	1815/25/30	S	130	05/77	170	305	475
	1855/60/70/85	S	271	05/77	710	225	935
	1905	S	106	04/80	245	140	385
	1955/1985	S	208	04/80	2,155	1,645	3,800
	1990	S	120	11/83	675	875	1,550
	2/3/500	S	317	11/61	30	15	45
	8700	S	68	03/73	1,400	1,200	2,600
	880, 890	S	30	11/76	11,440	17,160	28,600
	8800	S	51	05/77	5,100	2,750	7,850
	8900	S	78	08/80	2,595	1,305	3,900
	XE 520,550	S	70	03/84	180	120	300
				TOTAL	24,955	26,030	50,985
Sperry Univac	1004	S	49	02/63	25	10	35
	1005	S	91	04/66	20	0	20
	1050	S	380	09/63	30	8	38
	5201	S	13	09/68	0	0	0
	620/L+100C	S	12	02/71	110	320	430
	620/L-100	S	14	06/72	880	420	1,300
	620/a	S	75	12/65	0	0	0
	620/F,F-100	S	12	02/71	15	20	35
	6201	S	29	12/65	25	60	85
	90/25	S	199	06/77	20	50	70
	90/30	S	428	02/75	590	960	1,550
	9200	S	104	06/67	300	200	500
	9300	S	212	09/67	200	100	300
	BC-7	S	46	04/77	2,160	1,170	3,300
	B622/1	S	20	12/69	5	0	5
	S/80-3-4	S	145	12/81	500	710	1,210
	S/80-5-6	S	216	12/82	320	350	670
	Spectra 1600	S	189	01/71	12	0	12
	Spectra 70/15	S	203	09/64	0	0	0
	Spectra 70/25	S	343	09/65	0	0	0
	SS 80/90	S	357	08/58	0	0	0
	Sys/80-8	S	440	06/83	140	210	350
	UTS 4020	S	26	06/81	1,650	650	2,300
	UTS 4040	S	37	06/81	270	370	640
	V-72	S	45	02/74	770	370	1,140
	V-73	S	45	11/72	540	285	825
	V-74	S	80	05/74	50	15	65
	V-75	S	80	08/75	10	5	15
	V-76	S	16	01/76	675	260	935
	V-77-200	S	7	01/76	860	330	1,190
	V-77-400	S	28	12/77	275	180	455
	V-77-500	S	60	03/81	95	95	190
	V-77-600	S	66	03/77	1,220	490	1,710
	V-77-700	S	92	03/81	50	82	132
	V-77-800	S	97	03/79	757	524	1,281
	Mapper 10, Sys.11	S	200	—/84	66	78	144
	Mapper 5	S	50	—/84	300	300	600
				TOTAL	12,940	8,622	21,562
Altos Computer	186/486	S	5	07/84	1,000	200	1,200
	8000 Series	S	9	06/77	19,877	12,674	32,551
	ACS 586/986 Series	S	11	01/83	11,700	4,800	16,500
	ACS 68000	S	15	10/82	700	400	1,100
	ACS 8600-12, 14	S	14	01/82	1,000	400	1,400
	Series 5/580-20, 40, E Series	S	8	11/81	10,500	6,500	17,000
				TOTAL	44,777	24,974	69,751

REPRINTED FROM
A McGRAW-HILL PUBLICATION
Electronics®

New system manages hundreds of transactions per second

New system manages hundreds of transactions per second

Parallel data paths, pipelining, large cache memory, and 32-bit hardware combine to increase transaction system performanc

by Robert Horst and Sandra Metz, *Tandem Computers Inc., Cupertino, Calif.*

☐ Computer systems for on-line transaction processing have a unique set of requirements that pose an enormous challenge to designers. These systems have to be fault-tolerant, expandable through the addition of modules, and able to process multiple transactions at a reasonable cost, while maintaining data integrity. The coming generation of transaction-processing systems must also address a fast-growing need for very high-volume applications that require the processing of more and more transactions per second.

Designed to handle very high-volume transaction processing, the 32-bit NonStop TXP system reaches two to three times the speed of the NonStop II system it supercedes, while retaining complete software compatibility. Without reprogramming, a TXP system can grow from a single system containing from 2 to 16 processors, to a local cluster of up to 224 processors linked with fiber-optic cables, to a worldwide network of up to 4,080 processors.

Many of the problems in designing the TXP processor had already been solved in the NonStop II processor and system design. The NonStop II extended the instruction set of the NonStop 1+ system to handle 32-bit addressing but did not efficiently support that addressing in hardware. The existing 5-megabyte input/output bus and 26-megabyte Dynabus, Tandem's proprietary bus structure, had more than enough bandwidth to handle a processor with two to three times the performance. The existing packaging had an extra central-processing-unit card slot for future enhancements, and the existing power supplies could be reconfigured to handle a higher-power CPU.

The main problems involved designing a … micro-architecture that would efficiently support th… 2-bit instructions at much higher speeds, with only …% more printed-circuit-board real estate and an ex…g back-plane. This involved eliminating some feature… at were not critical to performance and finding creat… ways to save area on the pc board, including clever … of programmable array logic and an unusual multile… control-store scheme.

Since the new TXP processor was to be c…ct-code-compatible with the Nonstop II system yet ha… significant price-performance advantage, it was ex…ed that soon after announcement much of the compa… produc-

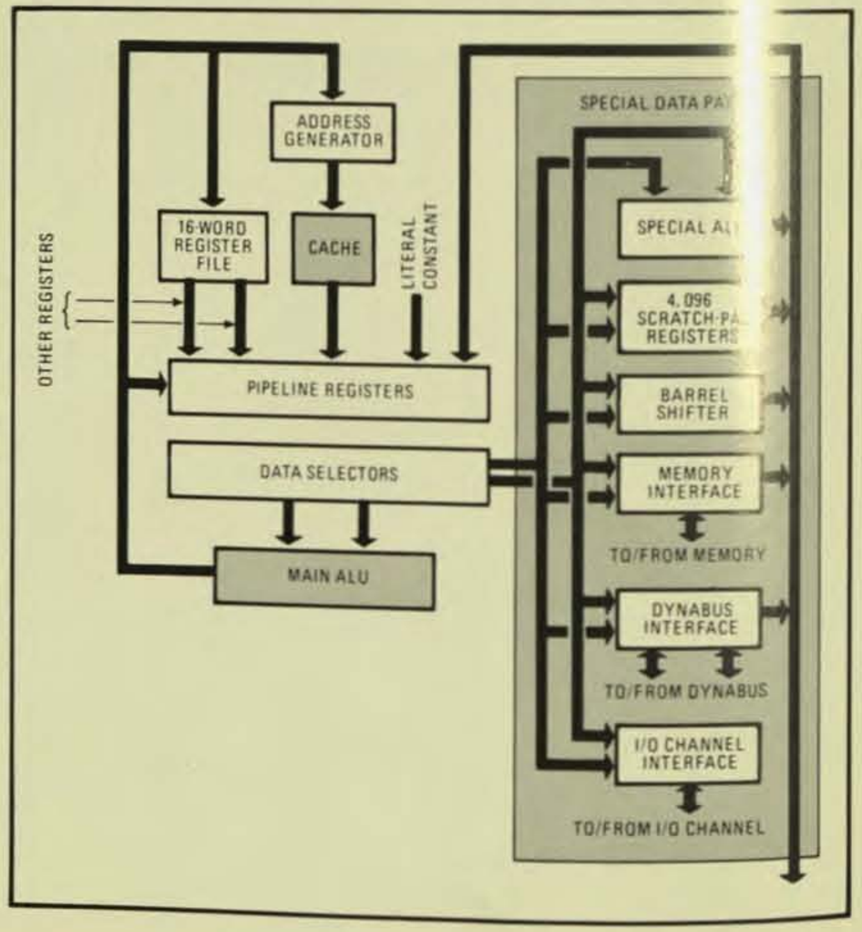

1. Parallel data paths. The NonStop TXP's architecture lets the main arithmetic and logic unit operate in parallel with either a special ALU, one of 4,096 scratch-pad registers, a barrel shifter, the memory interface, the Dynabus interface, or the input/output channel.

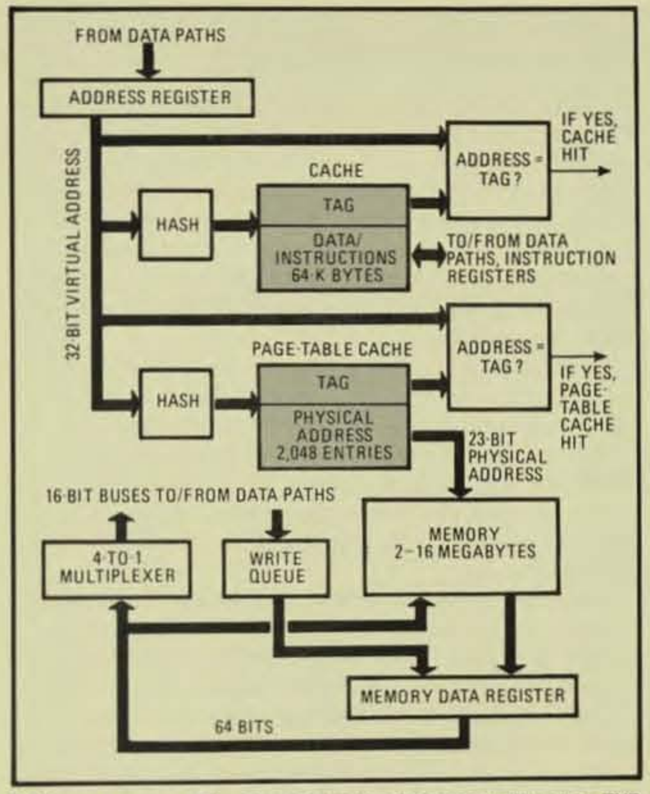

3. Memory access. The simple but extensive organization of the TXP cache provides an average hit ratio of over 96%. With a cache hit, the data is read out of the cache in 83 nanoseconds. When the data requested is not in cache, a cache miss results and the 64-bit-wide access to memory speeds the cache refill.

and it can execute load and branch instructions, which are frequently used, in only three clock cycles (250 ns).

Each NonStop TXP processor has a 64-K-byte cache that holds both data and code. A 16-processor NonStop TXP system has a full megabyte of cache memory. To determine the organization of the cache, a number of measurements were performed on a NonStop II system using a specially designed hardware monitor (see "Hardware-performance monitor helps optimize design," p. 149). The measurements showed that higher cache hit ratios resulted with a large, simple cache (directly mapped) than with a smaller, more complex cache (organized as two- or four-way associative). Typical hit ratios for transaction processing on the NonStop TXP system are in the range of 96% to 99%.

Cache miss

Cache misses are handled in a firmware subroutine rather than by the usual method of adding a special state machine and dedicated data paths for handling a miss. Because of the large savings in cache hardware, the cache can reside on the same board as the primary data paths; keeping these functions proximal reduces wiring delays and contributes to the fast 83.3-ns cycle time.

The cache is addressed by the 32-bit virtual address rather than by the physical address, thus eliminating the extra virtual-to-physical translation step that would otherwise be required for every memory reference. The virtual-to-physical translation, which is needed for refilling the cache on misses and for storing through to memory, is handled by a separate page table cache that holds mapping information for as many as 2,048 pages of 2-K bytes each (Fig. 3).

A cache memory by itself does not necessarily boost a processor's performance significantly. It is of little use for the cache to provide instructions and data at a higher rate than the rest of the CPU can process. In the TXP processor, the cache's performance was tuned to provide instructions and data at a rate consistent with the enhancements to instruction processing provided by increased pipelining and parallelism.

32 bits and more

The two concerns related to a system's word length are capability and performance. The NonStop TXP system has 32-bit virtual addressing built into the hardware, so is capable of addressing a gigabyte of virtual memory. In addition, the TXP processor can manipulate 32 bits of data at a time through its dual 16-bit data paths. Thus the 32-bit NonStop TXP system has the additional advantage of being able to run software that was originally written for the 16-bit NonStop II system; both systems have been provided with instructions that can operate on 8-, 16-, 32-, and 64-bit data types.

In transaction processing, measurements of instruction frequencies show that data-movement instructions (loads, stores, and moves) occur much more frequently than 32-bit arithmetic instructions. For this reason, the NonStop TXP system is optimized to handle data movement by providing 64-bit access to main memory and 32-bit buses and address registers to make memory addressing as efficient as possible.

The NonStop TXP processor was implemented on four large pc boards using high-speed FAST logic, PALs, and high-speed static RAMs. The CPU's logical and physical partitioning was carefully controlled to ensure that the machine's basic cycle time would not be slowed by long propagation delays. The four CPU boards are:

- SQ: containing the control store and sequencing logic.
- CC: containing the I/O channel and various special modules.
- IP: holding the main data paths and cache.
- MC: providing the memory interface, barrel shifter, and interprocessor bus interface.

Each CPU module also has from one to four memory boards. On the initial release, each memory board contains 2 megabytes of error-correcting memory implemented with 64-K dynamic RAMs. A 16-processor NonStop TXP system can therefore contain up to 128 megabytes of physical memory.

The NonStop TXP system was designed to be easy to manufacture and efficient to test. Data and control registers were implemented with shift registers configured into several serial-scan strings. The scan strings are of value in isolating failures in field-replaceable units. This serial access to registers also makes board testing much faster and more efficient because the tester can directly observe and control many control points. A single custom tester was designed for all four CPU boards and for the memory-array board as well.

The NonStop TXP system is the first product to be

New system manages hundreds of transactions per second

Parallel data paths, pipelining, large cache memory, and 32-bit hardware combine to increase transaction system performance

by Robert Horst and Sandra Metz, *Tandem Computers Inc., Cupertino, Calif.*

□ Computer systems for on-line transaction processing have a unique set of requirements that pose an enormous challenge to designers. These systems have to be fault-tolerant, expandable through the addition of modules, and able to process multiple transactions at a reasonable cost, while maintaining data integrity. The coming generation of transaction-processing systems must also address a fast-growing need for very high-volume applications that require the processing of more and more transactions per second.

Designed to handle very high-volume transaction processing, the 32-bit NonStop TXP system reaches two to three times the speed of the NonStop II system it supercedes, while retaining complete software compatibility. Without reprogramming, a TXP system can grow from a single system containing from 2 to 16 processors, to a local cluster of up to 224 processors linked with fiber-optic cables, to a worldwide network of up to 4,080 processors.

Many of the problems in designing the TXP processor had already been solved in the NonStop II processor and system design. The NonStop II extended the instruction set of the NonStop 1+ system to handle 32-bit addressing but did not efficiently support that addressing in hardware. The existing 5-megabyte input/output bus and 26-megabyte Dynabus, Tandem's proprietary bus structure, had more than enough bandwidth to handle a processor with two to three times the performance. The existing packaging had an extra central-processing-unit card slot for future enhancements, and the existing power supplies could be reconfigured to handle a higher-power CPU.

The main problems involved designing a new microarchitecture that would efficiently support the 32-bit instructions at much higher speeds, with only 33% more printed-circuit-board real estate and an existing backplane. This involved eliminating some features that were not critical to performance and finding creative ways to save area on the pc board, including clever uses of programmable array logic and an unusual multilevel control-store scheme.

Since the new TXP processor was to be object-code-compatible with the Nonstop II system yet have a significant price-performance advantage, it was expected that soon after announcement much of the company's production would have to shift quickly from the NonStop II system to the TXP system. This required that efficient board-testing procedures be in place by the time the product was announced and precluded the use of traditional functional board testers, which need months of programming after the design is finished. Instead, scan logic was designed into the processor and a scan-based board-test system using pseudorandom test vectors was developed.

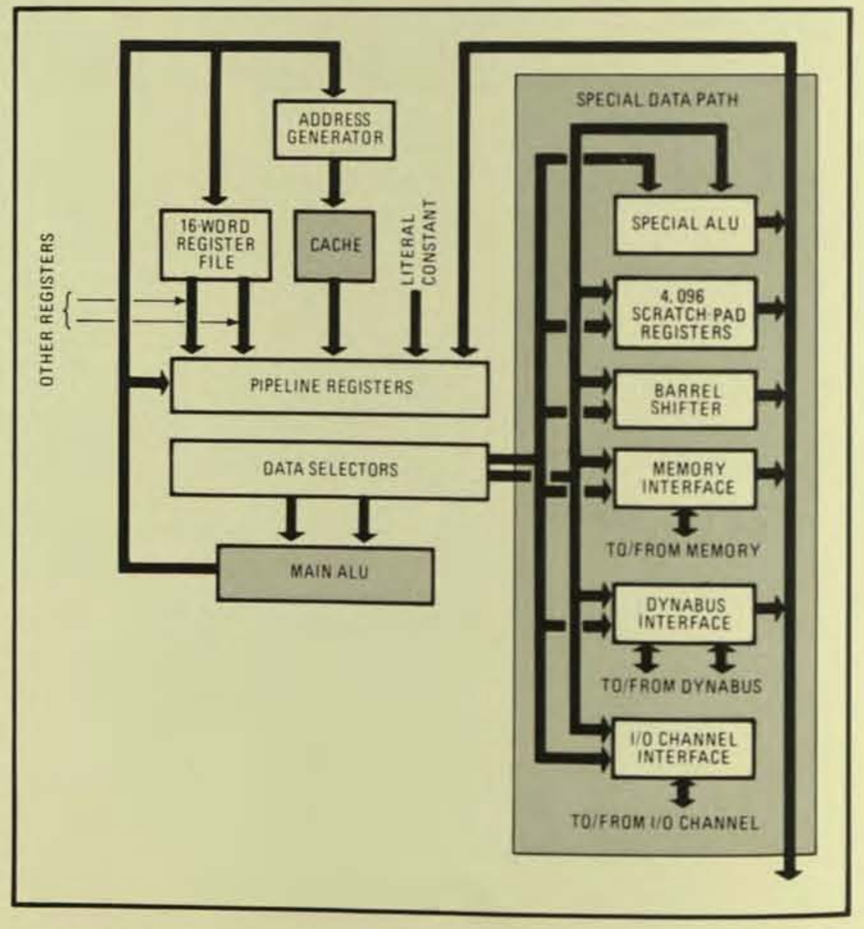

1. Parallel data paths. The NonStop TXP's architecture lets the main arithmetic and logic unit operate in parallel with either a special ALU, one of 4,096 scratch-pad registers, a barrel shifter, the memory interface, the Dynabus interface, or the input/output channel.

TABLE 1: COMPARE BYTE INSTRUCTIONS (INNER LOOP)

Clock cycle	NonStop TXP: Main ALU	NonStop TXP: Special ALU	Traditional architecture
1	extract byte 1	extract byte 2	extract byte 1
2	compare bytes	—	extract byte 2
3	(repeat)	(repeat)	compare bytes
4	—	—	(repeat)

TABLE 2: DYNABUS-RECEIVE MICROCODE INSTRUCTIONS (INNER LOOP)

Clock cycle	NonStop TXP: Main ALU	NonStop TXP: Special data path	Traditional architecture
1	compute checksum on previous word	read next word from bus queue	compute checksum on previous word
2	address next memory location	write data to cache and memory	read next word from bus queue, increment address
3	(repeat)	(repeat)	write data to cache and memory
4	—	—	(repeat)

Performance improvements

The performance improvements in the NonStop TXP system were attained through a combination of advances in architecture and technology. The NonStop TXP architecture uses dual 16-bit data paths, three levels of macroinstruction pipelining, 64-bit parallel access from memory, and a large cache (64 kilobytes per processor). Additional performance gains were obtained by increasing the hardware support for 32-bit memory addressing.

The machine's technology includes 25-nanosecond programmable array logic, 45-ns 16-K static random-access-memory chips, and Fairchild Advanced Schottky Technology (FAST) logic. With these high-speed components plus a reduction in the number of logic levels in each path, a 12-megahertz (83.3 ns per microinstruction) clock rate could be used.

The system's dual-data-path arrangement increases performance through added parallelism (Fig. 1). A main-arithmetic-and-logic-unit operation can be performed in parallel with another operation done by one of several special modules. Among them are a second ALU that performs both multiplications and divisions, a barrel shifter, an array of 4,096 scratchpad registers, an interval timer, and an interrupt controller. Other modules provide interfaces among the CPU and the interprocessor bus system, I/O channel, main memory, and a diagnostic processor.

The selection of operands for the main ALU and the special modules is done in two stages. In the first, data is accessed from the dual-ported register file or external registers and placed into two of the six registers. During the same cycle, the other four pipeline registers are loaded with cache data, a literal constant, the results of the previous ALU operation, and the result of the previous special-module operation.

In the next stage, one of the six pipeline registers is selected for each of the main ALU inputs and one for each special-module operand. Executing the register selection in two stages, so that the registers can be two-rather than four-ported, greatly reduces the cost of multiplexers and control storage, while the flexibility in choosing the required operands is unimpaired.

Some examples of the way microcode uses the parallel data paths are shown in Tables 1 and 2. The first example shows the inner loop of the compare-bytes instruction. Each of the dual ALUs in the TXP system extracts one byte; then the extracted bytes are compared. This operation takes two clock cycles on the TXP system but would require three if the extract operations could not be done simultaneously.

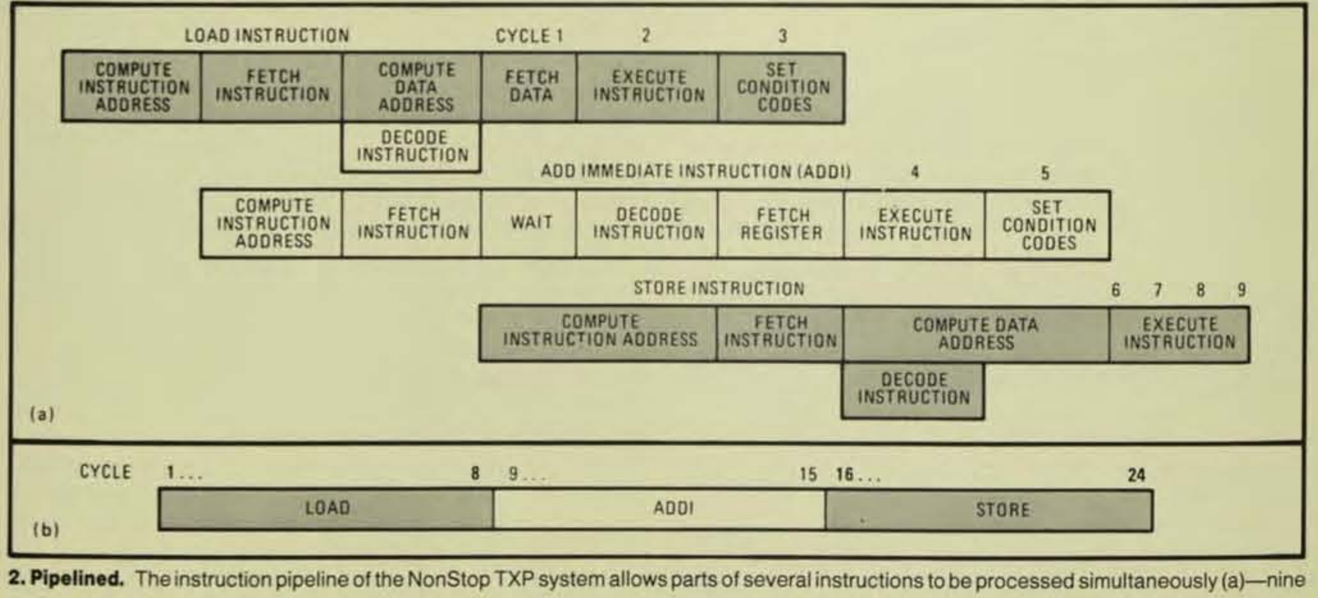

2. Pipelined. The instruction pipeline of the NonStop TXP system allows parts of several instructions to be processed simultaneously (a)—nine cycles are required to execute three typical instructions. Without pipelining (b), 24 clock cycles would be required.

Hardware-performance monitor helps optimize design

While new architectural concepts were being developed for the TXP system, a hardware-performance monitor was built to record measurements of the software-compatible NonStop II processor. Xplor consists of two large Wire-Wrap boards plus a small board to interface to the processor under test. It has approximately 800 Schottky TTL components and took more than two years to develop.

This general-purpose tool is capable of capturing 64 bits of data every 100 nanoseconds and reducing that data to usable form. The 256 kilobits of internal memory can be configured in many different word lengths to record, for instance, a 64-bit count of 4,096 different events, a 32-bit count of 8,192 different events, or a single flag for 256-K events. In addition, Xplor has programmable state machines with which data can be captured based on complex sequences of events; it includes hardware for the emulation of various cache organizations.

Two different Xplor configurations were developed to gather data for the TXP processor. The first was an instruction histogram measurement that records the frequency with which each instruction occurs, the percentage of time spent in each instruction, and the average number of code and data reads and writes performed by each instruction. The data is recorded in 64-bit counters, so in effect an unlimited amount of real-time data can be taken before the counters overflow.

The second Xplor configuration monitors memory addresses and emulates the tag store of a cache. Hit ratios for many different cache organizations can be determined by varying the effective cache size, associativity (one-, two, or four-way), block size, and replacement algorithm. Because the data is taken in real-time and reduced on-line, the hit-ratio measurements are much more accurate than the traditional technique, in which short address traces are recorded on tape for later analysis. This is especially important in transaction processing, since a large amount of process switching takes place; some individual transactions can last several seconds, during which millions of memory references take place.

Once the measurement methods were working, Xplor was attached to an eight-processor NonStop II system. A typical transaction-processing benchmark was brought up on the system, and transactions then were generated by another system, running software that simulated users at a number of terminals. At that point, histogram and cache measurements were taken for several of the central processing units.

The results of the histogram measurements helped determine some of the data-path widths and organizations for the TXP processor. Once the most frequently executed instructions were known, the design was modified to provide more hardware support for them. Since the measurements distinguished different paths through some instructions, tradeoffs could be made in the microcode to make the frequent cases faster.

The results of the cache measurements brought about some major changes in the original cache organization. In one measurement, the hit ratio went from 97% for the original cache to 99% for the final one, for an overall CPU performance gain of over 15%.

The dual 16-bit data paths tend to require fewer cycles than a single 32-bit path when manipulating byte and 16-bit quantities and slightly more cycles when manipulating 32-bit quantities. A 32-bit add takes two cycles rather than one, but the other data path is free to use the two cycles to perform either another 32-bit operation or two 16-bit operations.

Time disadvantage

The time disadvantage in performing a single 32-bit operation is partially offset by the cycle-time advantage for 16- versus 32-bit arithmetic (32-bit arithmetic requires more time for carry propagation). Measurements of transaction-processing applications have shown that the frequencies of 32-bit arithmetic are insignificant relative to data-movement and byte-manipulation instructions, which are handled more efficiently by the dual data paths than by a single 32-bit data path. Most instructions have enough parallelism to let the microcode make effective use of both data paths.

To control the large amount of parallelism in the NonStop TXP system processor, a wide control-store word is required. The effective width of the control store is over 100 bits. To reduce the number of RAMs required, the control store is divided between a vertical control store of 8-K 40-bit words and a horizontal control store of 4-K 84-bit words. The vertical control store controls the first stage of the microinstruction pipeline and includes a field that addresses the horizontal control store, whose fields control the pipeline's second stage. Lines of microcode that require the same or similar horizontal controls can share horizontal-control-store entries.

Unlike microprocessor-based systems that have microcode fixed in read-only memory, the NonStop TXP system microcode is implemented in RAM, so it can be changed along with normal software updates and new performance-enhancing instructions can be added.

The NonStop TXP processor uses three-stage pipelining for both macro- and microinstructions. Figure 2 illustrates the operation of the macroinstruction pipeline for a sequence of three instructions. The first is a load instruction that loads a word into the hardware stack. The second is an add immediate instruction that adds a constant to a register on the hardware stack, and the third is a final store, which stores the result in memory.

With no pipelining, this sequence would require 24 (8+7+9) clock cycles to execute, but because the prefetch and part of the execution of each instruction can be overlapped with previous instructions, the actual execution time is just 9 (3+2+4) clock cycles. Because instructions are pipelined, the TXP processor can execute its fastest instructions in just two clock cycles (167 ns),

MIPS and transactions per second

Determining relative performance among computer systems has never been an easy task. The often-quoted millions-of-instructions-per-second rate is intended as a way to compare basic central-processing-unit-hardware performance. Comparisons are also made on the basis of benchmarks. CPU-intensive benchmarks measure the performance of the CPU hardware and compiler; more extensive benchmarks measure the entire system performance—including the hardware, compiler, operating system, and data-base-management system. In general, the more extensive benchmarks give a more accurate prediction of actual system performance.

Each of the various measurement techniques has pitfalls. The MIPS rate is perhaps the least accurate way to compare systems. One reason is that there is no easy way to relate the power of one instruction set to another. In addition, vendors vary in the way they measure MIPS: some use it for the speed of the fastest instructions, others measure the speed of the most frequently executed instructions, and still others measure the speed of a "typical" mix of instructions. According to these definitions, each NonStop TXP processor is 6, 4, or 2 MIPS, respectively.

developed using Tandem's proprietary computer-aided-design system. The CAD system's capabilities for logic entry, logic simulation, and automated pc-board routing were instrumental in reducing the design time. While most high-performance CPUs require four to five years to develop, the NonStop TXP processor took just 2½ years—six months to complete a written specification, one year to construct a working prototype, and another year to reach volume production.

Performance measurement

Some simple benchmark programs have recently become popular in measuring performance (see "MIPS and transactions per second," p. above). One is the Puzzle benchmark, which is a CPU-intensive program to solve a three-dimensional puzzle. Execution times for Puzzle can vary widely for the same machine, depending on whether the program accesses arrays through subscripts or pointers and whether frequently used variables are assigned to registers. Versions of the Puzzle benchmark with pointers and registers were used to compare relative performance for a TXP processor.

Puzzle was written in TAL (transaction application language, the company's system-programming language); the execution time, using a single TXP processor, was measured at 1.67 s. This compares with 4 s on a VAX-11/780 for Puzzle written in C.[1] Because Puzzle does not measure such system features as support for virtual memory, I/O bandwith, and the ability to do fast context switching, a standard benchmark for comparing transaction-processing systems is still needed.

One transaction-processing benchmark has been developed by a third party, however. The U. S. Public Health Service ran an extensive benchmark in 1981 to determine which system to select for a large on-line medical-information system.[2] In that study, a 15-processor Tandem NonStop system running a 1981 version of Tandem's Encompass DBM system performed the benchmark at a rate of 4.5 transactions/s. An International Business Machines Corp. System 370/168-3 running version 3 of the Adabas DBM system performed the same benchmark at 2 transactions/s.

This benchmark gives a data point for comparisons between Tandem and IBM systems. A 15-processor NonStop system performs the Public Health Service benchmark 2.25 times as fast as an IBM 370/168-3. Though it would be desirable to compare the TXP system directly to one of IBM's newest systems, such as the IBM 4381-2, no competitive benchmarks have been published. However, comparisons of the MIPS rate of different processors within a single family are fairly accurate and can be used to extrapolate to newer systems.

According to market research performed by the Gartner Group,[3] the IBM 4381-2 is rated at 2.7 MIPS, compared with the older IBM 370/168-3's 2.4 MIPS rating—a ratio of 1.125 : 1. Company tests have shown the NonStop TXP to have a MIPS rate approximately three times that of the NonStop processor. The extrapolation of the Public Health Service benchmark performance to the two newer systems is shown in Table 3.

Unlike many shared-memory multiprocessor systems, Tandem systems provide linear growth in transaction-processing power as the system expands. A single system can include up to 16 processors, and clusters with as many as 224 NonStop TXP processors may be configured with Tandem's fiber-optic link. Clusters with up to 60 processors are currently in operation, and their users have verified the linear-performance growth within a cluster of this size.

The largest IBM mainframe today is the IBM 3084, which is rated at approximately 23 MIPS. Extrapolation from the benchmark data suggests that the performance of a cluster of 224 TXP processors is on the order of 10 times as powerful as IBM's top-of-the-line 3084 processor. □

TABLE 3: TANDEM VERSUS IBM PERFORMANCE COMPARISONS

	U.S. Public Health Service benchmark: results (transactions per second)	USPHS benchmark: extrapolated results* (transactions per second)
IBM 370/168-3	2	—
Tandem NonStop 15-processor system	4.5	—
IBM 4381-2	—	2.25
Tandem NonStop TXP 3-processor system	—	2.7

*Not actual measurements

References
[1]Malcolm A. Gleser, Judith Bayard, and David D. Lang, "Benchmarking for the Best," Datamation, May 1981.
[2]Computer Architecture News, 10 : 1, March 1982, p. 29.
[3]Gartner Group Inc., Stamford, Conn., market research surveys.

coast #500091

A McGRAW-HILL PUBLICATION
THREE DOLLARS

Electronics

THE WORLDWIDE TECHNOLOGY WEEKLY
APRIL 14, 1986

TANDEM MAKES A GOOD THING BETTER

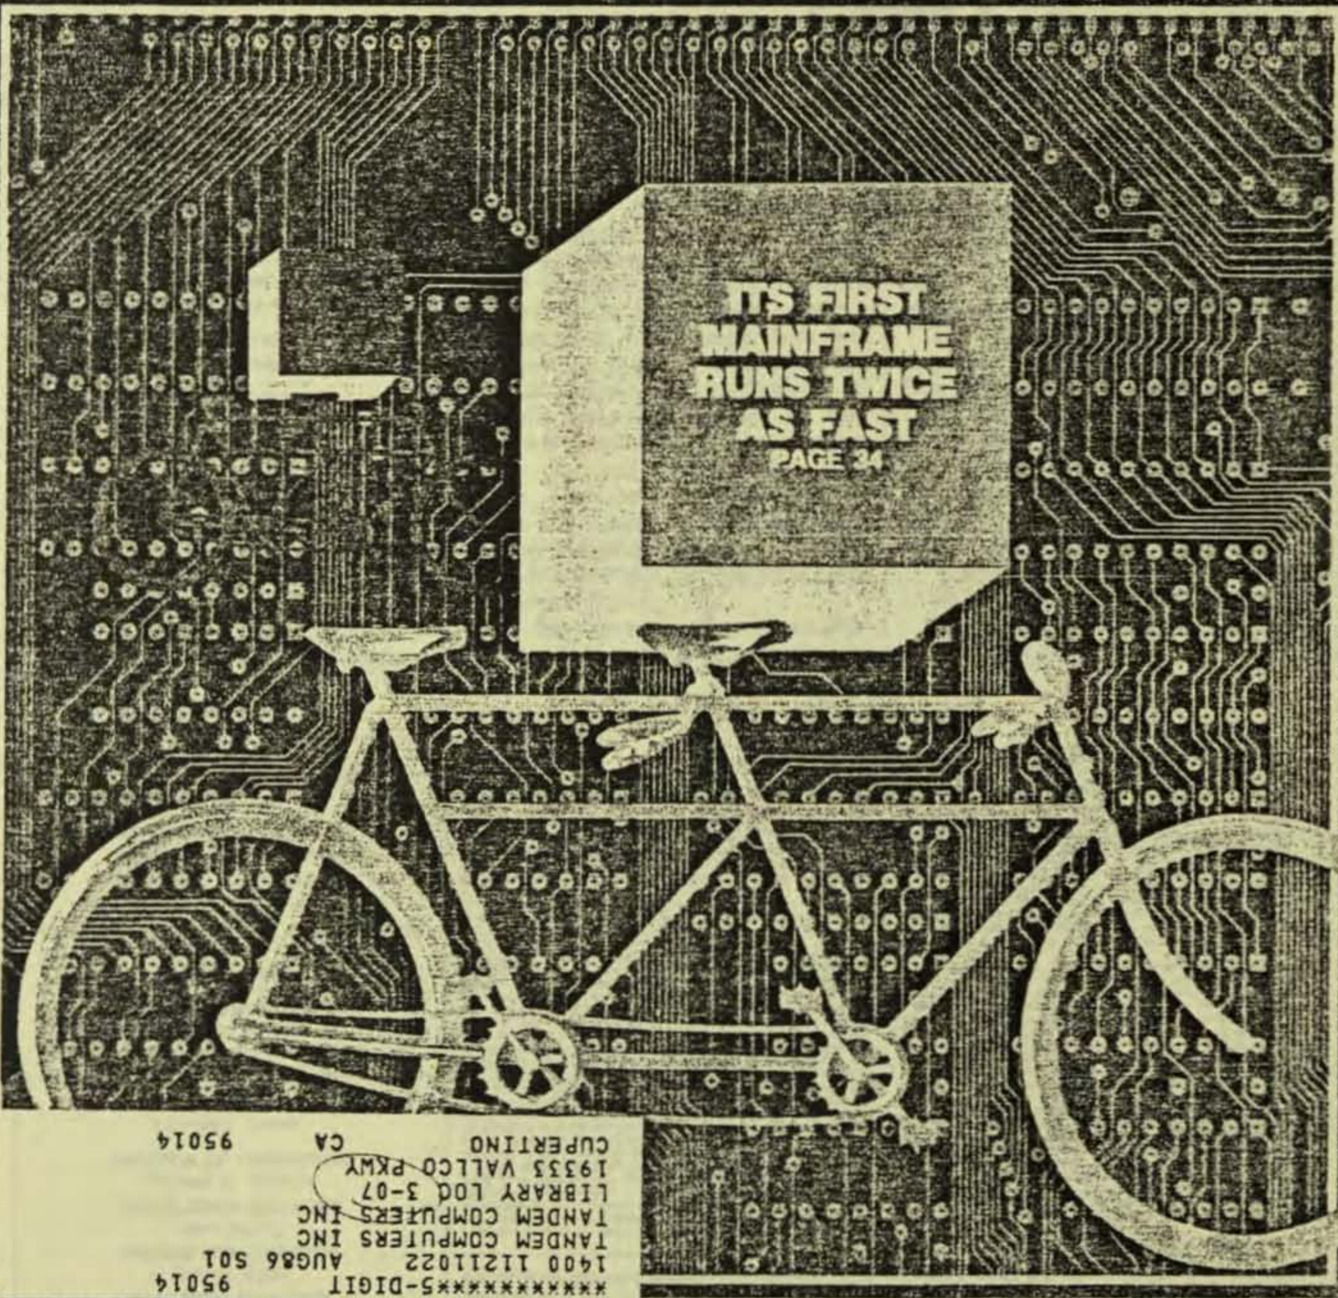

Tandem is a company that's "proving to be as rugged, and maybe even as resistant to failure, as its fault-tolerant computers," says Jonah McLeod about Tandem Computers Inc. "It has managed to prosper in a business that is getting tougher by the day."

McLeod, our test and measurement editor in Palo Alto, is the author of this week's cover story (p. 34) on the first Tandem machine that can be considered a mainframe, the NonStop VLX. Completing the Tandem coverage package, the company and its strategy are explored by Palo Alto bureau manager Cliff Barney in a Probing the News that begins on p. 39, following Jonah's piece.


JONAH McLEOD: Finds an impressively rugged company.

Jonah says that the ironic thing about Tandem's ruggedness is that it may be absorbed unintentionally from Tandem's archrival in the fault-tolerant computer business, IBM. For example, both lead designers on the VLX system worked for IBM before going to Tandem. Al McBride, technology director at Tandem, was responsible for the new system's macrocell arrays, and John Beirne, engineering manager, was in charge of the hardware work. Actually, the way things are set up at Tandem, McBride's title puts him one managerial level above Beirne.

But the corporate culture of the company tends to make such distinctions irrelevant, according to Jonah. The way Beirne explains it, even though McBride is technically in a higher position, "technology people are a resource shared by everybody in the company. It's just like at IBM, where the technology guys are out there helping the systems people solve the problems."

Editing material in New York gathered from around the world is a task that often requires the skill of a surgeon. But sometimes the scalpel slips, and you can be sure that the editors hear about it—often in the form of a deft jab that cuts neatly and quickly.

That's what happened last week when Charlie Cohen, in Tokyo, discovered an editing gaffe in one of his stories before it was printed. Charlie's reminder about the subtleties of the language is a classic of the genre:

"Did you ever hear Harrison Salisbury's story of how he wrote a dispatch from Moscow [when he was the *New York Times*'s correspondent] stating that he stood on the reviewing stand a stone's throw from Stalin? Evidently, the censor didn't understand the idiom and after receiving an explanation changed the text to read as follows:

" 'I stood on the reviewing stand close to Stalin. I threw no stones.' "

You might call the whole incident a bull's-eye for Charlie.

April 14, 1986 Volume 59, Number 15
94,985 copies of this issue printed

Electronics (ISSN 0883-4989). Published weekly by McGraw-Hill Inc. Founder: James H. McGraw 1860-1948. Publication office 1221 Avenue of the Americas, N.Y., N.Y. 10020; second class postage paid at New York, New York and additional mailing offices. Postage paid at Montreal, P.Q. Registration Number 9034.

Executive, editorial, circulation, and advertising addresses: Electronics, McGraw-Hill Building, 1221 Avenue of the Americas, New York, N.Y. 10020. Telephone (212) 512-2000. Teletype 12-7960 TWX 710-581-4879. Cable address: M C G R A W H I L L N E W Y O R K.

Subscriptions limited to professional persons with active responsibility in electronics technology. No subscriptions accepted without complete identification of subscriber name, title or job function, company or organization, and product manufactured or services performed. Based on information supplied, the publisher reserves the right to reject nonqualified requests. Subscription rates: the United States and possessions $32 one year, $55 two years, $76 three years; company addresses and company libraries $40 one year, $67 two years, $80 three years; Canada and Mexico $34 one year, $57 two years, $78 three years; Europe $50 one year, $85 two years, $125 three years; Japan, Israel, and Brazil $85 one year, $140 two years, $200 three years; Australia and New Zealand $95 one year, $170 two years, $240 three years, including air freight; all other countries $50 one year, $85 two years, $125 three years. Limited quota of subscriptions available at higher-than-basic rate for persons allied to field served. Check with publisher for these rates. Single copies $3.00. Please allow four to eight weeks for shipment.

Officers of McGraw-Hill Information Systems Company: President: Richard B. Miller; Executive Vice Presidents: Frederick P. Jannot, Construction Information Group; Russell C. White, Computers and Communications Information Group; J. Thomas Ryan, Marketing and International. Senior Vice Presidents: Francis A. Shinal, Controller; Robert C. Violette, Manufacturing and Technology. Senior Vice Presidents and Publishers: Laurence Altman, Computers and Communications; Harry L. Brown, Computers and Communications; David J. McGrath, Construction. Group Vice President: Peter E. McCuen, Communications. Vice Presidents: Fred O. Jensen, Planning and Development; Margaret L. Drayer, Human Resources.

Officers of McGraw-Hill Inc.: Harold W. McGraw Jr., Chairman; Joseph L. Dionne, President and Chief Executive Officer; Robert N. Landes, Executive Vice President and Secretary; Walter D. Serwatka, Executive Vice President and Chief Financial Officer; Shel F. Asen, Senior Vice President, Manufacturing; Robert J. Bahash, Senior Vice President, Finance and Manufacturing; Ralph R. Schulz, Senior Vice President, Editorial; George R. Elsinger, Vice President, Circulation; Ralph J. Webb, Vice President and Treasurer.

Subscribers: The publisher, upon written request to our New York office from any subscriber, agrees to refund that part of the subscription price applying to copies not yet mailed. Please send change-of-address notices or complaints to Fulfillment Manager; subscription orders to Circulation Manager, Electronics, at address below. Change-of-address notices should provide old as well as new address, including zip codes. Attach address label from recent issue. Allow one month for change to become effective. Subscriber Service: call (201) 356-7980, 9 a.m. to 4 p.m. EST.

Postmaster: Please send address changes to Electronics, att. Fulfillment Manager, CN 808, Martinsville, N.J. 08836.

Editorial department phones: Administration (212) 512-2645, News and New Products (212) 512-2685, Technology (212) 512-2666. Bureaus: Boston (617) 262-1160, Chicago (312) 751-3811, Dallas (214) 458-2400, Los Angeles (213) 480-5234, New York (212) 512-3322, San Francisco/Palo Alto (415) 968-2712, Washington (202) 463-1650, Frankfurt 72-5566, London 493-1451, Paris 42-89-03-80, Tokyo 581-9816. **Business departments:** (212) 512-3469 (Business departments follow the Advertisers' Index).

APRIL 14, 1986 VOLUME 59, NO. 15

Electronics

Contents

TANDEM MAKES A GOOD THING BETTER

ITS FIRST MAINFRAME RUNS TWICE AS FAST

In the competitive world of computer technology, a company is only as good as its latest design. And on that benchmark, Tandem Computers Inc. gets high marks with its first mainframe-class machine, the NonStop VLX. The new model not only lops 30% off the per-transaction cost of the two-year-old NonStop TXP superminicomputer, but it also rests on a hardware base that is three times more reliable, the company says (see related story, p. 39).

Designed for heavy-duty on-line transaction processing in such areas as airline reservations, banking, computer-integrated manufacturing, and telecommunications switching, the NonStop VLX executes 12 million to 48 million instructions per second. The Cupertino, Calif., company says that's roughly twice the performance of the TXP for about the same price. It attributes the higher price/performance ratio and reliability primarily to the extensive use of the MCA2800ALS, an emitter-coupled-logic gate array that serves as a building block in critical parts of the central processing unit (see "Easy-design features make macrocell a hit," p. 35). Other performance hikes come from streamlined instruction execution, reorganized cache memory, and a faster interprocessor bus.

Although speed was the paramount concern in designing the VLX, its developers also concentrated on fault tolerance and compatibility with the previous-generation system. The NonStop configuration of both the TXP and VLX provides dual paths to every element in the system. If one path fails, a second is available to make the connection. If one processor fails, another assumes its workload. Although this duplication slows throughput to a small extent, it guarantees that failures will not affect system operation.

When all units are functioning, they carry their full share of the processing load—there are no idle spares. This fault-tolerant architecture, which remains unchanged on the new VLX system, can be described as a loosely coupled parallel-processing system with distributed, non-shared memory.

One benefit of the distributed processing architecture is that it does not require one large central processor running at the highest possible clock speed. "Our design was at a point where we needed a faster central processor—but not the fastest possible," says Al McBride, Tandem's technology director. "We could get higher total system performance from the parallelism of the system architecture."

The parallel architecture allows Tandem to be more conservative in processor design than manufacturers of high-speed CPUs. For example, the company limited the number of circuits implemented on a gate array to 2,000 out of the 2,800 available gates. This meant that the macrocell arrays were easier to design and yielded a more reliable system.

The minimum VLX system consists of four CPUs, but the architecture can accommodate as many as 16. Thanks to the macrocell technology it developed jointly with Motorola Inc., Tandem was able to implement a two-board processor that executes 3 mips. By comparison, the TXP's 1.5-mips CPU fit on four printed-circuit boards using medium-scale-integration TTL chips.

The denser macrocell array chips allowed each VLX processor board to contain about three times as many gates. Much of the VLX's speed improvement came from packing more logic onto fewer chips and boards, which reduces the number of interchip and interboard signals.

CIRCUITS TO SPARE

"We went from approximately 53,000 gates on four logic boards to approximately 85,000 on two," says John Beirne, engineering manager for VLX hardware development. "The macrocell arrays allow up to 2,800 circuits per array. But 2,000 to 2,300 is a more reasonable count for this circuit."

The mainframe's CPU holds 21 types of custom macrocell arrays (Fig. 1). In total, there are 33 different macrocell arrays in the Tandem two-board processor set. Of these, 31 are on the two logic boards that make up the CPU and two are on the accompanying memory boards. These include clever macrocell array designs that speed instruction execution from cache and branch operations.

Two critical elements of the new CPU that improve its overall speed in these areas are the instruction unit and execution unit (Fig. 2). The instruction unit consists of 10 macrocell arrays. One is a chip that makes a four-stage pipeline for simultaneously processing four instructions obtained from the

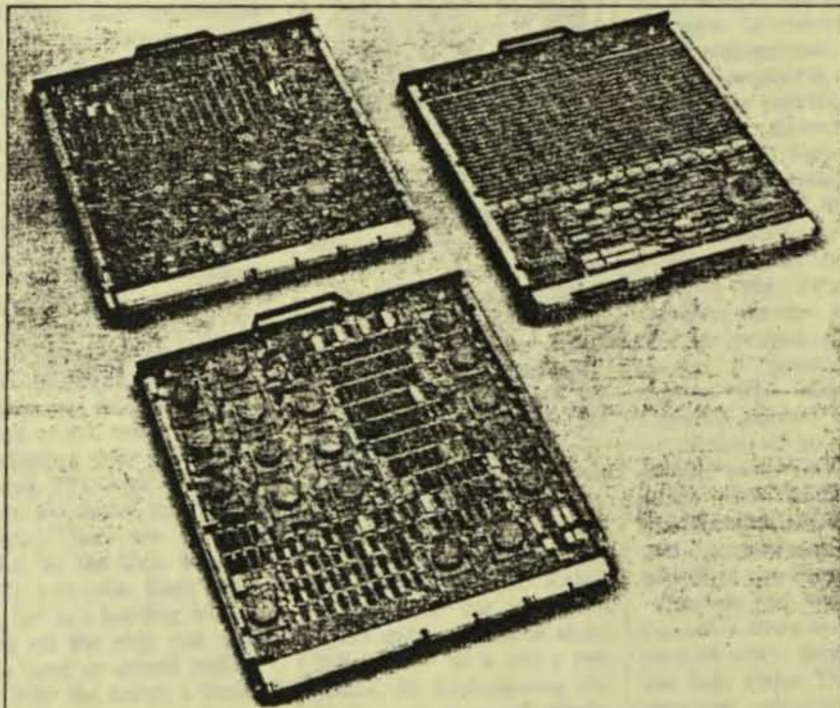

1. ARRAYS ABOUND. The 31 gate arrays in the CPU's two boards (left) make high-density logic circuitry. Two arrays and 256-K chips populate the 8-megabyte memory board (right).

64-K static-random-access-memory cache. The processor can fetch one instruction, decode a second, preprocess a third, and execute a fourth.

Two more macrocell arrays build a displacement adder that does address arithmetic for prefetching operands. This unit adds a displacement number to a virtual address to determine its physical address on a board. When the instruction unit prefetches instructions, address processing occurs in parallel with execution to anticipate the next instruction address.

In the older TXP, microcode-carried-out address calculation used discrete logic chips. In both systems, the processor instruction set holds 220 machine instructions that handle such jobs as stack operations, integer and decimal arithmetic, and byte-oriented functions. In addition, 43 other instructions perform scientific calculations.

The instruction set is implemented by microinstructions in the control store, which users can use to implement new instructions or improve existing ones. Two of the 10 instruction-unit macrocells handle addressing of the control store.

Another macrocell array in the CPU is dedicated to branch control. It examines the conditions of all branch (or jump) instructions being executed. If the conditions indicate a jump, the array helps determine the next address to be accessed. This look-ahead capability tests the jump condition prior to execution of the jump instruction. By knowing that a jump is imminent, the contents of the cache can be flushed and reloaded with the contents of the new jump address and the subsequent addresses in this new sequence of instructions. Anticipating the jump can shave microseconds off an operation.

Four macrocell arrays in the address-translation unit cut the cache-fill time in half when the cache must be flushed and reloaded, as when a jump is executed. Tandem says this alone contributes several percentage points to the performance improvement of the new-generation system.

The cache itself got a speed boost. Cache cuts the time required to access data and instructions from slower main memory. During operation, several instructions following the one being executed have been prefetched into the cache. For example, when there is a program in memory that the computer needs to execute, the processor fetches the program instructions from the high-speed cache instead of directly from main memory, thus reducing the amount of time to get an instruction. Besides being slower than cache memory, main memory is located off the processor board, which delays fetches from it even more.

New cache-hashing algorithms ensure a higher number of cache hits—that is, that needed instructions are in cache rather than in main memory. "Improving the performance of the

EASY-DESIGN FEATURES MAKE MACROCELL A HIT

The best-seller in Motorola Semiconductor Products Inc.'s cell library is the MCA2800-ALS. The emitter-coupled-logic macrocell array, the product of a joint development effort between the Phoenix, Ariz., Motorola Inc. subsidiary and Tandem Computers Inc., is a critical component in Tandem's fault-tolerant VLX NonStop computer.

Built with an advanced process called Mosaic II (Motorola oxide-isolated, self-aligned integrated circuits), the array outperforms discrete 100K ECL chips. With a single 5-V power supply, it runs at 125 MHz with a typical 600-ps gate delay.

Mosaic II is a bipolar process that uses oxide isolation, which produces much smaller transistors than junction isolation. Some of the MCA-2800ALS speed improvement results from compressing the 2,000 or so gates down to an area no larger than 100 mils^2.

"One significant contribution Tandem made to the macrocell array is making it easier to design with," says John Carey, Motorola's merchandising manager. An especially attractive feature to logic designers is the TTL input/output capability Tandem added to the macrocell. TTL I/O cells, not found on other any ECL macrocell arrays as yet, make it easy for designers who are adept at using TTL logic to incorporate the device into their designs. There are 120 signal pins on the chip, hence 120 TTL I/O cells. Each I/O cell is tied to a bonding wire coming off the chip and can be an input or output buffer.

With the array, a designer does not have to learn new design techniques to create his final design on the chip because the library of macro functions is very similar to existing discrete TTL circuits. Using a gate array with TTL I/O pins affords other benefits as well. For example, the designer does not have to use controlled-impedance boards, a must for ECL designs.

All 120 I/O cells can be used. "It is one of the more dense I/O structures on the market," says Carey. "There are no restrictions concerning which pins can be an input, output, or bidirectional. Any one can be an input, output, or a tristate [high-impedance state] cell." Most arrays impose restrictions on which pins can be used.

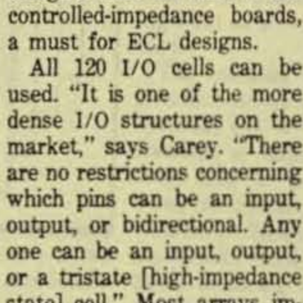

BEST SELLER. The MCA2800-ALS comes in a pin-grid array with a heat sink on the back.

The I/O cell is a significant addition to the Motorola library, but Tandem didn't stop there. The Cupertino, Calif., company changed the components making up the macrocell, which made it possible to implement a circuit element, such as a flip-flop, more efficiently. Tandem also improved the implementation of plain NOR and NAND gates inside the macrocell. These gates constitute about 20% to 25% of a chip's real estate. By implementing the NOR gates more effectively, for example, the designers achieved a 20% chip-density improvement.

For its NonStop mainframe, "we changed the distribution of the types of resistors and transistors in the macrocell," says Al McBride, Tandem's technology director. "We made changes to both the simple and complex functions to make sure we used a large percentage of the 2,800 circuits on each array."

Tandem also trimmed the power consumed by individual gates by changing the output resistor's rating. ECL internal structures have an emitter-output follower circuit. The output signal is dropped across the emitter resistor.

A variety of values are possible with this resistor. If the cell is used in a relatively low-speed data path, for example, a designer can reduce the power of the circuits in this path to reduce system power consumption.

"We changed the operating point of the current gates in the chip's internal ECL," says McBride. Very high-performance ECL products dissipate more than 8 W. Tandem needed to have no more than 5 W in worst case. "We changed the output-switch operating current. As current flow is lowered, the performance, is necessarily decreased. The output-switch operating current is a value we traded off to reach our design goal," McBride explains. The tradeoff gave Tandem the necessary system performance without having to use liquid cooling.

Tandem and Motorola also created a three-level ceramic pin-grid-array package that has heat sinks. This further improves reliability of the chip. Reducing the operating point by 10° to 15°C doubles the chip's life.

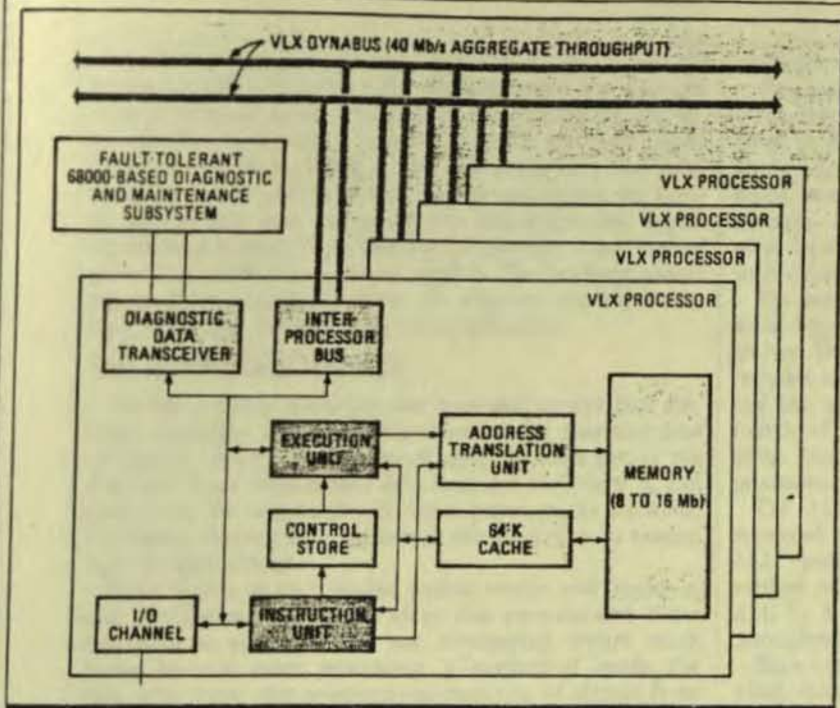

2. BASIC SET. A starter VLX system contains four identical CPUs, each with one or two memory boards, tied together by the Dynabus. A system can be expanded to 16 processors.

cache-hashing algorithms [the way cache is organized for quick retrieval] is an important means of getting the hit rate up," says Beirne. "We took benchmarks on our early designs and tried various hashing algorithms. By running the benchmarks in a simulator to determine hit rates early in the design process, we knew what kind of performance we were going to get with the architecture."

Like the displacement adder, the address-translation unit converts virtual addresses to physical addresses. There are two identical arrays on the memory board for moving 64-bit data words (8 bytes wide) to and from memory. Two macrocell arrays located between memory and cache provide error-correction codes.

The loadable control store, which contains the microinstructions that tell the instruction unit how to execute its instruction set, been made even faster with a dual interleaved design. It consists of fast SRAM that supplies microinstruction words to both the execution and instruction units.

Two more arrays control the addressing of microinstructions in the control store and four chips receive the output of the control store and put it into the three-stage pipeline of microinstructions. The four chips are identical, each one fourth of a total data flow path to the logic that executes the instructions.

TWIN MICROINSTRUCTION BANKS

In the control store is a unique design feature that contributes significantly to improving the operating speed of the macrocell arrays in the instruction unit. It consists of two separate banks of SRAM, each containing identical copies of the microinstructions.

A conventional interleaving scheme divides the microinstructions into two halves, even-address microinstructions in one bank and odd-address ones in the other. The state machine—or instruction unit, in this system—would execute one microinstruction from bank A, the next from bank B, the third from bank A, and so on, because the cycle time of the instruction unit is faster than the RAM access time.

In Tandem's implementation, the instruction unit fetches a microinstruction from bank A, and as it executes, fetches the next microinstruction from bank B. One reason for this approach is that microinstructions are of variable length. Thus microinstruction A may be 1 cycle long and microinstruction B could be 3 cycles long. Having duplicate copies of the same microinstructions in each bank affords the most amount of overlap between microinstruction fetches, hence allowing a greater increase in speed over the conventional interleaving.

A fallout of the duplicated banks of microinstruction store is increased reliability, because a soft error in one bank can be repaired by loading the suspect bank with the known-good bank. In addition, if one of the two banks has a hard failure, the processor can continue operating—but at slower speed—using only one bank.

The company had to use the same instruction set as in its previous-generation systems, but it had some flexibility and freedom in improving the microinstructions used to realize the macroinstructions, or common machine instructions. Tandem's NonStop system instructions can require from one to five microinstructions. Reducing the number of microinstructions needed for one instruction required some additional logic. But Tandem decided that the increased performance resulting from faster execution of instructions was significant enough to warrant the extra logic.

"From what our early benchmarks told us about our existing computer architecture, we discovered which instructions to optimize," says Beirne. "We plotted histograms of instruction usage to see how much time it requires to execute each instruction. Looking at instructions that were executed the most told us where to look for the greatest savings in instruction-execution time. From this data we were able to perform analysis which would tell us we would get so much improvement in performance by adding logic to improve instruction-execution performance. We made the changes, then reran the simulations to see that the benchmarks improved."

OLD BUS, NEW PROTOCOL

Another increase in speed is provided by using a new bus protocol on the existing Dynabus. The previous bus protocol had radial clock distribution, which requires costly cabling. The new protocol uses a double-clocking scheme in which the clocks are distributed with the data. This method automatically reduces the amount of skew between the clock and data, thus allowing the system to more tightly compress the interval of data transmission.

The clock and data slow down the same amount over a longer length—if the data arrives later, so does the clock. Once transmission begins, the interval between data bursts can be very tight.

The new protocol allows Tandem to extend the length of the bus as well as increase its speed from 13 to 20 megabytes/s per bus. With two buses in the system, the aggregate bus transfer rate went from 26 to 40 megabytes/s. "We have not seen a case where bus speed is a bottleneck," says Beirne. "However, we feel that the improvement positions us well for future processors as well as allowing heavier loading on the VLX processor. It gives us more margin."

When the system is operating, packets of transaction information move to and from the CPU through its I/O channel and the Dynabus. A high-speed bus on the CPU connects the Dynabus, I/O channel, and diagnostic data transceiver (a microprocessor that automatically monitors the CPU) with the execution and instruction units. Packets enter a one-packet-wide input queue in the Dynabus, diagnostic data transceiver, or I/O channel, depending on which is active, and their arrival generates an interrupt flag to the instruction unit. The interrupt causes the instruction unit to begin processing the packet. During this time, the processor's instruction unit begins fetching instructions from memory to determine what kind of processing the incoming packet requires. The incoming packet enters the execution unit when the program running in the instruction unit executes a Receive instruction.

ONE MACROCELL PER BUS

On the Dynabus board are two macrocell arrays that execute a sequence of prescribed instructions for receiving data off the bus. There is one macrocell array for each part of the dual bus. These chips receive data from the bus, check it, and pass it into the one-packet-wide input queue on the Dynabus. Previously, discrete logic performed this operation by executing microinstructions.

Functional logic that handles packet receipt and transmission is in the same macrocell array that receives and transmits data to and from the bus. Processing occurs much faster because more processing is performed inside the chip, with fewer chip crossings—movements of signals from one logic chip to another on a pc board.

Two macrocell arrays make up the I/O channel and perform a function similar to the Dynabus. The two arrays replace the 70 or so discrete TTL components comprising the earlier system's I/O function. Improved reliability and a reduced chip count were the main benefits of using macrocells in this instance. Improving performance was a secondary consideration because the system was required to remain compatible with existing TXP I/O channels.

The execution unit holds seven macrocell arrays: four arithmetic-logic-unit slices, two register-file slices, and a barrel shifter. The ALU is a slice of the execution data flow path. It includes all the registers, parity-prediction logic, multiplexers, and the data path. Each of the four arrays represents one fourth of a 32-bit-wide ALU, each identical eight-bit-wide units. Each operates on eight bits of the total data word being processed.

The ALU follows the strategy of creating one common macrocell and using it four times, rather than partitioning the ALU function into several distinct functions. The former method reduced the number of unique array designs for the ALU by a fourth. This strategy was used wherever possible throughout the processor and memory boards.

Each two-board CPU can have one or two memory boards, which store 8 megabytes each, for a total of 16 megabytes per

CHANGING COMPUTER FAULT SIMULATION AND REPAIR

The conventional wisdom in fault diagnosis and repair says wait until the problem occurs, then begin. But too often, the fault is a transient failure that cannot be recreated. So in the new NonStop VLX system from Tandem Computers Inc., the TMDS (Tandem Maintenance Diagnostic Subsystem) constantly monitors the system environment: processors, interprocessor bus, and tape subsystem. Microprocessor-based sensors scattered throughout the system can detect a fault as it occurs.

There are microprocessors throughout the system. The two system bus controllers contain the 68000-based remote maintenance interface. Hall-effect sensors inside the fans monitor intake and exhaust temperatures. On power supplies, TMDS can measure the actual analog output level. "These microprocessors collect this information and feed it into the remote maintenance interface and it is then fed back into TMDS, where it is analyzed constantly," says software manager Jamie Allen.

If a fault occurs, say, on one of the main processors, its 6809 microprocessor captures its entire state, 4,000 bits (500 bytes) of information called the event signature, and stores it on disk. It stops the processor clocks within one cycle and captures all the registers and states of the parity checkers throughout the machine. "We do parity checking across the control lines," says Allen.

Expert-system techniques are used to perform the fault analysis. In 90% of the cases, the program isolates problems down to field-repairable units, such as circuit boards and disk-drive modules. Eventually, the TMDS designers expect

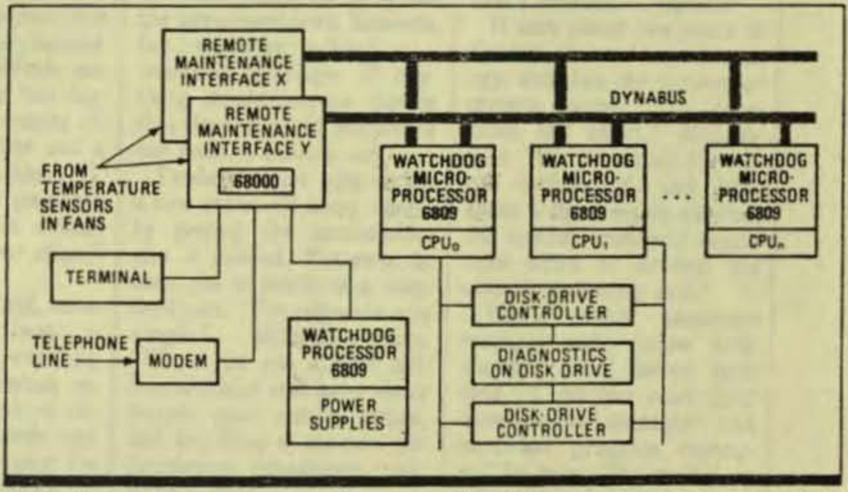

WATCHDOGS. Microprocessors scattered around the VLX system gather maintenance data.

to achieve 100% accuracy.

There are other fault analyzers in the system. A general-purpose program called the Mother Fault Analyzer is written in Lisp using MRS, a rule-based language that sits on top of Lisp. Developed at Stanford University, MRS is similar to Prolog.

Because it is rule-based, the general-purpose analyzer knows nothing about the events. The program interprets the rules against the event to determine what to do next. It may do nothing, for example, if it reasons that an event is of no importance, such as a corrected soft memory error. The fault analyzer can make some repairs itself.

A user can trace the entire analysis process, including results, at a terminal running a program called Problem Reporter. A Tandem engineer can also perform the analysis over the phone. Remotely, he can perform maintenance operations such as measuring the power-supply voltage, checking the operation of fans, adjusting power-supply margins, and adjusting the clock frequency up or down by 5%.

"We can come in over a modem and perform the analysis and in some cases actually perform the repair, especially if the cure is to rebalance the system, reload a processor that had an intermittent error, or patch a software bug," says Allen.

The remote capability addresses one of the most troublesome parts of the maintenance problem, the "no fault found" service call caused by intermittent faults. With this system, the actual event is captured without having to duplicate the trouble after it has gone. □

processor set. The company points out that memory capacity can be expanded up to 256 megabytes as higher-density megabit RAM chips become available.

Gathering maintenance data has been speeded up, too. Two macrocell arrays inside the diagnostic data transceiver can capture every internal state of the VLX processor board in one clock cycle. Once the data is grabbed, it is shifted serially to a separate maintenance processor. Another macrocell array controls the collection of reliability data. There are eight strings of reliability data collected from each board in the VLX system.

MICROPROCESSOR-MACROCELL INTERACTION

One chip in the diagnostic data transceiver controls all the scanning for data and multiplexes it into a serial bit stream that goes to a 6809 microprocessor on-board the diagnostic transceiver. From the 6809, the information is routed to a separate 68000-based maintenance processor. The second macrocell interfaces the 6809 to the execution and instruction units. "Here is a case where a single microprocessor is used but its related support circuits are put into one custom macrocell array," says Beirne.

Another way macrocell arrays improve processing speed is by allowing the system designer to concentrate on maximizing the performance of circuits that have the greatest impact on the overall system processing speed. Tandem's analysis revealed that about 50% of a computer's operating cycle is spent either in the cache or control store.

Another 20% of the total operating cycle is spent moving data between chips: the output of one logic chip is routed into the next. Gate delays and travel time between chips combine to slow computer performance. Finally, 30% of a computer's operating cycle is spent in the logic of any given chip in a computer design.

Thus in designing VLX, the company spent much of its design effort improving the cache and control store and using high-speed ECL internal chips to speed processing. But they chose to compromise speed for simplified board design in the 20% of the total processor time spent moving data between chips.

Up until recently, ECL macrocells all came with ECL I/O cells. These cells require special interfacing translations if the macrocell I/O is to be connected with TTL circuits. They also require specialized pc boards. All of these requirements conspire to make ECL difficult to design with.

Tandem chose to change the macrocell array so that it offered a TTL I/O cell instead of ECL. "It was a design tradeoff that paid off," says McBride. "We did not have to use controlled-impedance circuit boards, which would be required if using full-ECL gate arrays."

In addition, the company is able to build its boards using many off-the-shelf VLSI and LSI RAMs and microprocessors without having to redesign the entire computer from the ground up. Tandem's designers could have used ECL arrays throughout. But then they would not have been able to use non-ECL VLSI and LSI and the whole processor would have cost more.

A TTL I/O buffer interface on a macrocell chip is slower than ECL I/O buffers typically found on ECL macrocell arrays. "ECL is about two times as fast as TTL, but I/O only affects 20% of the total performance of the computer, so you're talking about only a 10% effect on the total cycle time of the computer," says McBride. In addition, implementing macrocells with ECL I/O could raise the price of the design considerably.

Tandem's design resulted from the fact that they could get enough system performance improvement even using TTL I/O, but TTL I/O allows the chips to be interfaced with all general-purpose commodity parts, RAMs, microprocessors and other non-ECL circuits. It gave Tandem a better system solution in terms of other components on the board as well as the pc-board technology. □

THREE YEARS AND A MILLION DOLLARS LATER...

Tandem's NonStop mainframe project was one program that lived up to its name—technology director Al McBride can attest to that. "My first day, when I entered an empty office, I had $1 million and a note from Jim [James G. Treybig, Tandem's president and chief executive officer] saying turn this into chips," he recalls.

The Cupertino, Calif., company's VLX project began in the fall of 1983 and was fully staffed to meet product requirements by March of the following year. McBride was charged with developing the semiconductor technology.

The first order of business was finding a supplier that could produce high-performance semicustom chips for the VLX, and developing a symbiotic relationship was the key. "When we went into the agreement with Motorola Inc., we were looking at a nine-year marriage. If anything should happen during that time, it could jeopardize our product coming out."

Tandem would help define a new macrocell array, thereby getting the semiconductors it needed. Motorola, in turn, got to sell it as a standard part. "The rationale was simple," McBride says. "When you are a $100 million company and have never bought gate arrays before, and you have a one-man circuit-design department, you make a deal."

McBride worked at IBM Corp. for 15 years before joining Tandem. Since the early 1970s, he has worked on microprocessor and microcomputer chips, some for IBM's Personal Computer.

JOHN BEIRNE

AL McBRIDE

It took about two years to develop the silicon technology, and then the system architects began. "We dovetailed the effort," McBride says. "We developed the silicon technology, and with about a three-month overlap, the system architects began their effort to develop the central processing unit."

That's when hardware manager and fellow IBM alumnus John Beirne took over. "I had two roles: CPU development manager and hardware program manager," he says. "We staffed up a team that designed the macrocell chips for the CPU."

Before coming to Tandem, Beirne, an energetic young man with engineering degree and MBA, spent seven years with IBM. "I was a technical leader and led a software project and worked with advanced chip technologies," he says. "One reason I got on the VLX project was my experience with very large-scale integration."

Though McBride developed the chip technology separately from the actual system, there was considerable interplay. "If we had an idea for a circuit, we could evaluate it within a day or two against what impact it would have on system performance."

PROBING THE NEWS

TANDEM'S OLD ARCHITECTURE PAYS OFF IN NEW MARKETS

IS IT A 'THREATENING COMPETITOR' TO IBM IN TRANSACTION PROCESSING?

by Clifford Barney

SUNNYVALE, CALIF.

Tandem Computers Inc. is a true original in a world of look-alikes. Though all computer makers strive mightily for product differentiation, they tend to come up with variations on a very few themes. Tandem has parlayed a unique computer architecture into an almost unmatchable position in the specialized world of transaction processing. The fault-tolerant design is now offering a new payoff: a powerful role in distributed computing and networking, two of the high-growth markets of the 1980s.

"The world is moving toward on-line processing," says Tandem president James G. Treybig. "Businesses are distributed everywhere: branch banks, retail stores, sales offices, point-of-sale operations. If they can understand what is happening instantaneously, they can keep inventories low and provide better services."

The company claims to have 60% of the automatic bank teller market—where it goes head to head with IBM Corp.—and a similar share of the electronic funds-transfer business between banks. Every major oil company uses Tandem equipment for credit-card transactions, Treybig says, and Tandem machines are used in 15 major stock exchanges, including the New York Stock Exchange. And the company is about to sign a contract for total automation of a major U. S. airline. Tandem hopes for even more market penetration with its new NonStop VLX, a cheaper and more reliable successor to its two-year-old NonStop TXP superminicomputer (see related story, p. 34).

Transaction processing will account for $1.8 billion in sales for U. S. computer manufacturers this year, according to a recent report by Frost & Sullivan Inc., the New York market research firm, and will grow to $4.7 billion by 1990. At present, two thirds of the market is from on-line processing, such as automatic teller machines. But office automation, now only 6%, will triple its market share by the end of the decade, and industrial process control will jump from 20% to 25%, the report forecasts. Tandem is even more bullish about the future of on-line processing. Dennis P. McEvoy, vice president of software development, says the market already runs well beyond $10 billion annually, the exact size depending on how the categories are chosen and who is doing the counting.

Last year, Tandem recorded $624.1 million in sales, a sum over 17% higher than fiscal 1984's $532.6 million and a reasonable figure in a miserable year for computer makers. The outlook is tougher this year: first-quarter sales increased only 6.5% over the comparable period in 1985.

"How well we do in 1986 depends on how well the computer market does in the U. S. If it doesn't pick up, we will be hard-pressed to have a better year than last year," Treybig says. International business—mostly to European countries—is strong, he adds, but domestic sales remain soft. "We did relatively well last year," Treybig says. "At least all of our people"—5,500 people at five manufacturing sites, 100 sales offices worldwide, and 19 subsidiary companies—"kept their jobs."

Tandem's architecture is a form of parallel computing. But instead of breaking down one large problem into many small ones, as in most parallel systems, the Tandem approach starts with many small problems and processes them very quickly. Hence its suitability for distributed computing.

TAILOR-MADE. The location of a data base or a peripheral is immaterial in a Tandem system, which makes it well-suited for networking. One process communicates with another through packet-switched messages without regard for physical location. The process of creating a network is therefore straightforward, requiring less input/output buffering than with conventional architectures.

Its ability to offer networks makes Tandem one of the first mainframe manufacturers to be able to challenge networking companies. Through product development, it is already beginning to stress value-added networks, and, with Rockwell International's Switching Division, it has developed an integrated communications and computer system for telemarketing.

Tandem was founded in 1974 to make computers for financial institutions, manufacturers, transportation companies, and others who needed the continuous processing of multiple events rather than batch processing of data. Instead of building hardware redundancy into its system, Tandem designed a message-based architecture, which allows all parts of the system to operate independently. Because the Tandem system can shut down gracefully, nodes can also be added easily. And because all this equipment must be continuously

TREYBIG: 'We are far ahead.'

PROBING THE NEWS

TANDEM'S OLD ARCHITECTURE PAYS OFF IN NEW MARKETS

IS IT A 'THREATENING COMPETITOR' TO IBM IN TRANSACTION PROCESSING?

by Clifford Barney

SUNNYVALE, CALIF.

Tandem Computers Inc. is a true original in a world of look-alikes. Though all computer makers strive mightily for product differentiation, they tend to come up with variations on a very few themes. Tandem has parlayed a unique computer architecture into an almost unmatchable position in the specialized world of transaction processing. The fault-tolerant design is now offering a new payoff: a powerful role in distributed computing and networking, two of the high-growth markets of the 1980s.

"The world is moving toward on-line processing," says Tandem president James G. Treybig. "Businesses are distributed everywhere: branch banks, retail stores, sales offices, point-of-sale operations. If they can understand what is happening instantaneously, they can keep inventories low and provide better services."

The company claims to have 60% of the automatic bank teller market—where it goes head to head with IBM Corp.—and a similar share of the electronic funds-transfer business between banks. Every major oil company uses Tandem equipment for credit-card transactions, Treybig says, and Tandem machines are used in 15 major stock exchanges, including the New York Stock Exchange. And the company is about to sign a contract for total automation of a major U. S. airline. Tandem hopes for even more market penetration with its new NonStop VLX, a cheaper and more reliable successor to its two-year-old NonStop TXP superminicomputer (see related story, p. 34).

Transaction processing will account for $1.8 billion in sales for U. S. computer manufacturers this year, according to a recent report by Frost & Sullivan Inc., the New York market research firm, and will grow to $4.7 billion by 1990. At present, two thirds of the market is from on-line processing, such as automatic teller machines. But office automation, now only 6%, will triple its market share by the end of the decade, and industrial process control will jump from 20% to 25%, the report forecasts. Tandem is even more bullish about the future of on-line processing. Dennis P. McEvoy, vice president of software development, says the market already runs well beyond $10 billion annually, the exact size depending on how the categories are chosen and who is doing the counting.

Last year, Tandem recorded $624.1 million in sales, a sum over 17% higher than fiscal 1984's $532.6 million and a reasonable figure in a miserable year for computer makers. The outlook is tougher this year: first-quarter sales increased only 6.5% over the comparable period in 1985.

"How well we do in 1986 depends on how well the computer market does in the U. S. If it doesn't pick up, we will be hard-pressed to have a better year than last year," Treybig says. International business—mostly to European countries—is strong, he adds, but domestic sales remain soft. "We did relatively well last year," Treybig says. "At least all of our people"—5,500 people at five manufacturing sites, 100 sales offices worldwide, and 19 subsidiary companies—"kept their jobs."

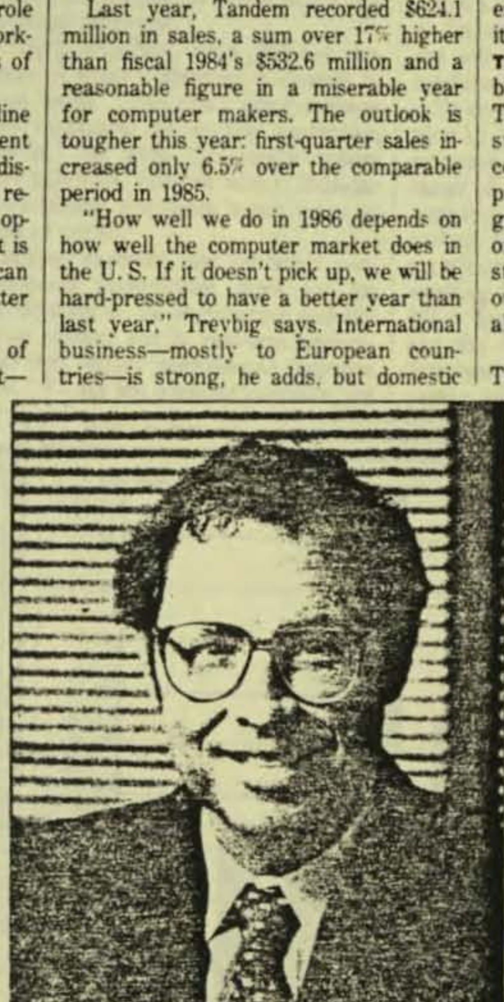

TREYBIG: 'We are far ahead.'

Tandem's architecture is a form of parallel computing. But instead of breaking down one large problem into many small ones, as in most parallel systems, the Tandem approach starts with many small problems and processes them very quickly. Hence its suitability for distributed computing.

TAILOR-MADE. The location of a data base or a peripheral is immaterial in a Tandem system, which makes it well-suited for networking. One process communicates with another through packet-switched messages without regard for physical location. The process of creating a network is therefore straightforward, requiring less input/output buffering than with conventional architectures.

Its ability to offer networks makes Tandem one of the first mainframe manufacturers to be able to challenge networking companies. Through product development, it is already beginning to stress value-added networks, and, with Rockwell International's Switching Division, it has developed an integrated communications and computer system for telemarketing.

Tandem was founded in 1974 to make computers for financial institutions, manufacturers, transportation companies, and others who needed the continuous processing of multiple events rather than batch processing of data. Instead of building hardware redundancy into its system, Tandem designed a message-based architecture, which allows all parts of the system to operate independently. Because the Tandem system can shut down gracefully, nodes can also be added easily. And because all this equipment must be continuously available to its users, Tandem made it fault-tolerant. As a result, any given fault can be quickly isolated from the rest of the system.

So well has Tandem succeeded in establishing a reputation for fault tolerance that the company's other big strength—expandability—tends to be overlooked in the industry, and Tandem is not generally seen as a leader in technology. Says McEvoy, "If we were a new company and announced the VLX [Tandem's new top-of-the-line product] . . . everyone would say that we had outstanding technology. But we didn't start three years ago with venture capital, we started 11 years ago, and people tend to say we have older technology."

EXPANDABLE TO 48 MIPS. The new machine is linearly expandable from 3 million instructions per second to 48 mips, all completely transparent to the user, with no special programming. It can be "put into a distributed network with the same programs, with alternate routing in case of failures," says McEvoy. It boasts state-of-the-art ECL gate arrays and a special disk architecture for parallel throughput, he says.

To compensate for its lack of visibility, Tandem is beginning to take the wraps off some of its technology. With its own semiconductor facility and computer-aided design tools, Tandem worked directly with Motorola Inc. on the design of the ECL gate array for the VLX, thus shaving weeks off turnaround time. This circuit is only the beginning. A CMOS processor is in the wings, and Al McBride, director of very large-scale integration, hints that a new generation of processors awaits only the development of triple-layer metalization for the fabrication of very dense "sea of gates" arrays.

McEvoy claims that Tandem's networking software is also years ahead of the competition. No other distributed data base can run asynchronously and concurrently in multiple memories, he says; Tandem got a jump because it already runs a distributed data base in a single system.

One criticism sometimes leveled at Tandem machines is that the message-based architecture means higher system overhead for users. But McEvoy says that it also results in the lowest cost per transaction and the cheapest route to modular expansion.

"If you stay with a single box," he says, "you are limited to the 20% to 30% performance gain made each year by the technology." With the Tandem system, linking processors results in a near-linear increase in performance, the company maintains.

Tandem's success forced IBM to offer fault-tolerant systems

That's important, says McEvoy, because it gives customers a chance to start small. "You don't always know how successful an on-line application will be," he adds. "Automatic tellers have taken off, videotex hasn't. Modular expandability lets users invest more as demand grows."

McEvoy concedes that Tandem's message-based architecture rules out its participation in the engineering and scientific market and in real-time applications that require microsecond response time. "We are not compute intensive, we are I/O intensive," he says. "You can call it special-purpose computing, but it's applicable to a good third of the business."

On the business side, says Treybig, Tandem's major achievement is its performance against IBM. "We forced them to buy someone else's computer," he says proudly, referring to IBM's use of equipment made by Stratus Computer Ltd., Natick, Mass., in some of its transaction-processing applications. In addition, Treybig adds, AT&T had to come out with a hardware-redundant machine, the 3B20, to compete with Tandem in on-line computing.

'THREATENING.' Indeed, says Stephen K. Smith, an analyst at Paine Webber Inc., New York, Tandem is a "threatening competitor" to IBM in transaction processing, a market IBM has traditionally dominated. IBM has mounted a significant effort to stop Tandem from taking away business, Smith says.

But Tandem's networking ability will keep the company a jump ahead of Stratus, IBM's supplier, says Michael Murphy of the *California Technology Stock Letter*. "They are clearly the leader [in transaction processing] and they can sell a network, not just one or two machines," Murphy says.

To broaden its markets, Tandem is beginning to introduce message-handling software. It is beginning to offer a range of information-transfer programs, including electronic mail, data communications, facsimile, and microcomputer-to-mainframe communications. Its networking software supports both IBM's LU.6 Systems Network Architecture and the International Organization for Standardization's open-systems interconnection reference model.

Tandem also sees opportunities in factory automation and was one of the first companies to support the General Motors Corp.-sponsored Manufacturing Automation Protocol. Late last month, Tandem bought a small piece of Triplex, a Torrance, Calif., maker of fault-tolerant programmable logic controllers, and the two will jointly produce a MAP system. "We are far ahead of anyone on networking software and distributed data bases," Treybig says in summary. "Others build networks on top of individual processors. Our lowest-level operating system is a network. That is unique." □

TANDEM'S SALES HAVE TRIPLED . . .

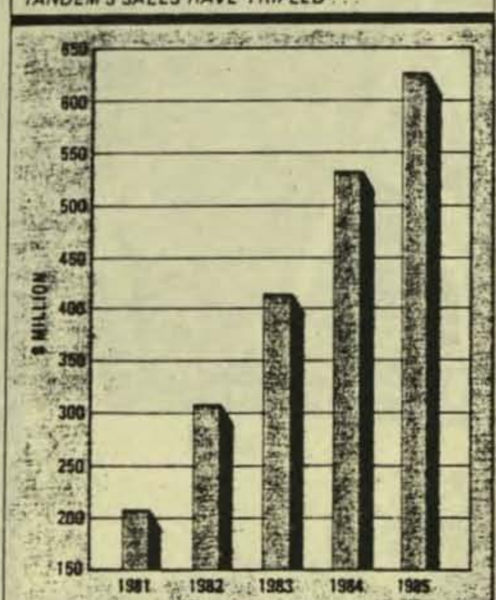

. . . BUT PROFITS HAVE RISEN MORE SLOWLY