Global Market Insight for Information Technology Companies

1995 RESEARCH PROGRAMS

From semiconductors to systems, software to services, telecommunications to document management, Dataquest's scope of expertise provides clients with a clear view of the relationships between information technology segments – relationships that can have a profound impact on making strategic business decisions.

Computer Systems and Peripherals

Computer Systems
- Client/Server Computing *Worldwide*
- Advanced Desktop, Workstation, and Client/Server Distribution Channels
- Computer and Client/Server Systems *Europe*
- Workgroup Servers *Europe*
- UNIX and Open Systems *Europe*
- Computer Systems *Japan*

Workstations
- Advanced Desktop and Workstation Computing *Worldwide*
- Advanced Desktop and Workstation Quarterly Statistics *Worldwide*
- Workstations *Europe*
- Workstations Quarterly Statistics *Europe*

Personal Computing
- Personal Computers *Worldwide*
- *European* PC Strategic Service
- European PC Market Update
- Personal Computers *Asia/Pacific*
- PCMCIA Systems and Peripherals *Worldwide*
- Mobile Computing *Worldwide*
- PC Quarterly Statistics *United States*
- PC Quarterly Statistics *Europe*
- PC Quarterly Statistics *Japan*
- PC Distribution Channels *Worldwide*
- PC Distribution Channels *Europe*
- Network Distribution Channels *Europe*
- PC Distribution Quarterly Statistics
- PC Tracking *Asia/Pacific*

Computer Storage
- Removable Storage *Worldwide*
- Optical Disk Drives *Worldwide*
- Optical Disk Drives *Europe*
- Rigid Disk Drives *Worldwide*
- Rigid Disk Drives *Europe*
- Tape Drives *Worldwide*
- Tape Drives *Europe*

Graphics
- Graphics and Displays *Worldwide*

Online, Multimedia, and Software

Emerging Technologies
- Multimedia *Worldwide*
- Online Strategies *Worldwide*

Business Productivity
- Client/Server Software *Worldwide*
- Personal Computing Software *Worldwide*
- *European* Personal Computing Software
- *European* Personal Operating Systems
- Workgroup Computing *Worldwide*
- *European* Workgroup Computing

Technical Applications
- AEC and GIS Applications *Worldwide*
- Electronic Design Automation *Worldwide*
- Mechanical CAD/CAM/CAE *Worldwide*
- CAD/CAM/CAE/GIS *Europe*
- CAD/CAM/CAE/GIS *Asia/Pacific*

Services

Customer Services
- Customer ServiceTrends *North America*
- European Customer ServiceTrends

Professional Services
- Professional ServiceTrends *North America*
  - Systems Integration and Applications Development
  - Consulting and Education
  - Systems Management
- Professional ServiceTrends: Vertical Market Opportunities *North America*
- European Professional ServiceTrends

Sector Programs
- System Services *North America*
- *European* Desktop Support Services
- Network Integration and Support Services *North America*
- *European* Network Integration & Support Services
- Software Services *North America*
- Strategic Service Partnering *North America*

Dataquest

1995 RESEARCH PROGRAMS

Document Management

Emerging Technologies
- Digital Documents *North America*
- Colour Products *Europe*
- Printer Distribution *Europe*

Copiers
- Copiers *North America*
- Copiers *Europe*

Facsimile
- Facsimile *North America*

Printers
- Printers *North America*
- Printers *Europe*
- Printer Distribution Channels *Europe*
- Printer Tracking *Asia/Pacific*

Semiconductors

Regional Markets
- Semiconductors *Worldwide*
  - Semiconductor Industry Snapshot
- Semiconductors *Europe*
- Semiconductors *Asia/Pacific*
- Semiconductors *Japan*
- Strategic Tracking Services
  - China/Hong Kong
  - Taiwan
  - Korea
  - Singapore

Products
- ASICs *Worldwide*
- ASICs Product Planning *Europe*
- Memories *Worldwide*
- Memory Product Planning *Europe*
- Microcomponents *Worldwide*
- Microcomponent Product Planning *Europe*
- Discretes Product Planning *Europe*

Application Markets
- Semiconductor Application Markets *Worldwide*
- Semiconductor Application Markets *Europe*
- Semiconductor Application Markets *Asia/Pacific*
- Communications Semiconductor Applications
- PC Semiconductor Applications
- PC Teardown Analysis
- PC Production Monitor *Europe*
- Electronic Equipment Production Monitor *Europe*

Manufacturing
- Semiconductor Equipment, Manufacturing, and Materials *Worldwide*
- LCD Multi-client Study

Pricing Trends
- Semiconductor Procurement *Worldwide*

Telecommunications

Networking
- Networking *North America*
  - Local Area Networks *North America*
  - Digital WANs *North America*
  - Modems *North America*
- Networking *Europe*
  - ISDN *Europe*
  - Modems *Europe*
  - Local Area Networks *Europe*
  - Private Digital WAN *Europe*
- Quarterly Market Watch—Intelligent Hubs

Voice
- Voice Communications *North America*
  - Voice Processing *North America*
  - Integrated Calls Centers *North America*
  - Premise Switching *North America*
- Voice Communications *Europe*
  - Voice Processing *Europe*
  - Call Centres *Europe*
  - Telephones *Europe*
  - PBX/KTS Systems *Europe*

Public
- Public Network Equipment & Services *Worldwide*
- Public Network Equipment & Services *North America*
  - Public Network Equipment *North America*
  - Public Network Services *North America*
- Public Network Equipment & Services *Europe*
  - Public Network Equipment *Europe*
  - Public Network Services *Europe*

Personal
- Personal Communications *North America*
- Personal Communications *Europe*
  - Infrastructure and Services *Europe*
  - Terminals *Europe*

Dataquest Catalyst Programs

Cross Technology insight for:
- Financial Service Providers
- Government Agencies
- IT Business Strategy and Planning Organizations
- Publishing, Media, and Consulting Firms
- IS and Purchasing Organizations
- IT Supporting Organizations

Emerging Markets

Central and Eastern European Markets
- Printers
- Computer Markets

Information Technology Market Opportunities
- China
- Asia/Pacific
- Latin America

Dataquest
a company of The Dun & Bradstreet Corporation

Boston Area
550 Cochituate Road
Framingham, MA 01701-9324
United States
Phone: (1) 508-370-5555
Fax: (1) 508-370-6262

London Area
Holmers Farm Way
High Wycombe, Buckinghamshire
HP12 4XH
United Kingdom
Phone: (44) 1494-422722
Fax (44) 1494-422742

Tokyo
Shinkawa Sanko Building
6th Floor
1-3-17, Shinkawa
Chuo-ku, Tokyo 104
Japan
Phone: (81) 3-5566-0411
Fax (81) 3-5566-0425

Taiwan
3/F, No. 87 Sung Chiang Road
Taipei
Taiwan, R.O.C.
Phone: (886) 2-509-5390
Fax (886) 2-509-6660

Corporate Headquarters
1290 Ridder Park Drive
San Jose, CA 95131-2398
United States
Phone: (1) 408-437-8000
Fax: (1) 408-437-0292

Paris
Tour Franklin, Cedex 11
94042 La Défense, Paris
France
Phone: (33) 1-41-25-16-00
Fax: (33) 1-41-25-16-16

Munich
Kronstadter Strasse 9
81677 Munich 80
Germany
Phone: (49) 89-930-9090
Fax (49) 89-930-3277

Asia/Pacific
1/fl China Hong Kong Tower
8-12 Hennessy Road
Hong Kong
Phone: (852) 527-7861
Fax (852) 520-0079

Korea
Trade Tower, Suite 3806
159 Samsung-dong
Kangnam-gu, Seoul 135-729
Korea
Phone: (82) 2-551-1331
Fax (82) 2-551-1330

MKTG 11/94 (PSB CORP-01)

Dataquest
a company of
The Dun & Bradstreet Corporation

251 River Oaks Parkway, San Jose, CA 95134-1913
(408) 468-8000 Fax (408) 954-1780

September 25, 1995

Errata

On page 15 of *PC Graphics Controllers: A Focused Analysis* (PSAM-WW-FR-9502), Figure 4-3 was mistitled. The correct title is "Bus Interface Forecast for Desktop Graphics Controllers." The figure is also titled incorrectly in the List of Figures. Dataquest regrets the error and apologizes for any inconvenience. Please insert this page into the inside front pocket of your PC Semiconductor Applications binder.

For further information, contact Industry Analyst Geoff Ballew at (408) 468-8676 or at gballew@dataquest.com.

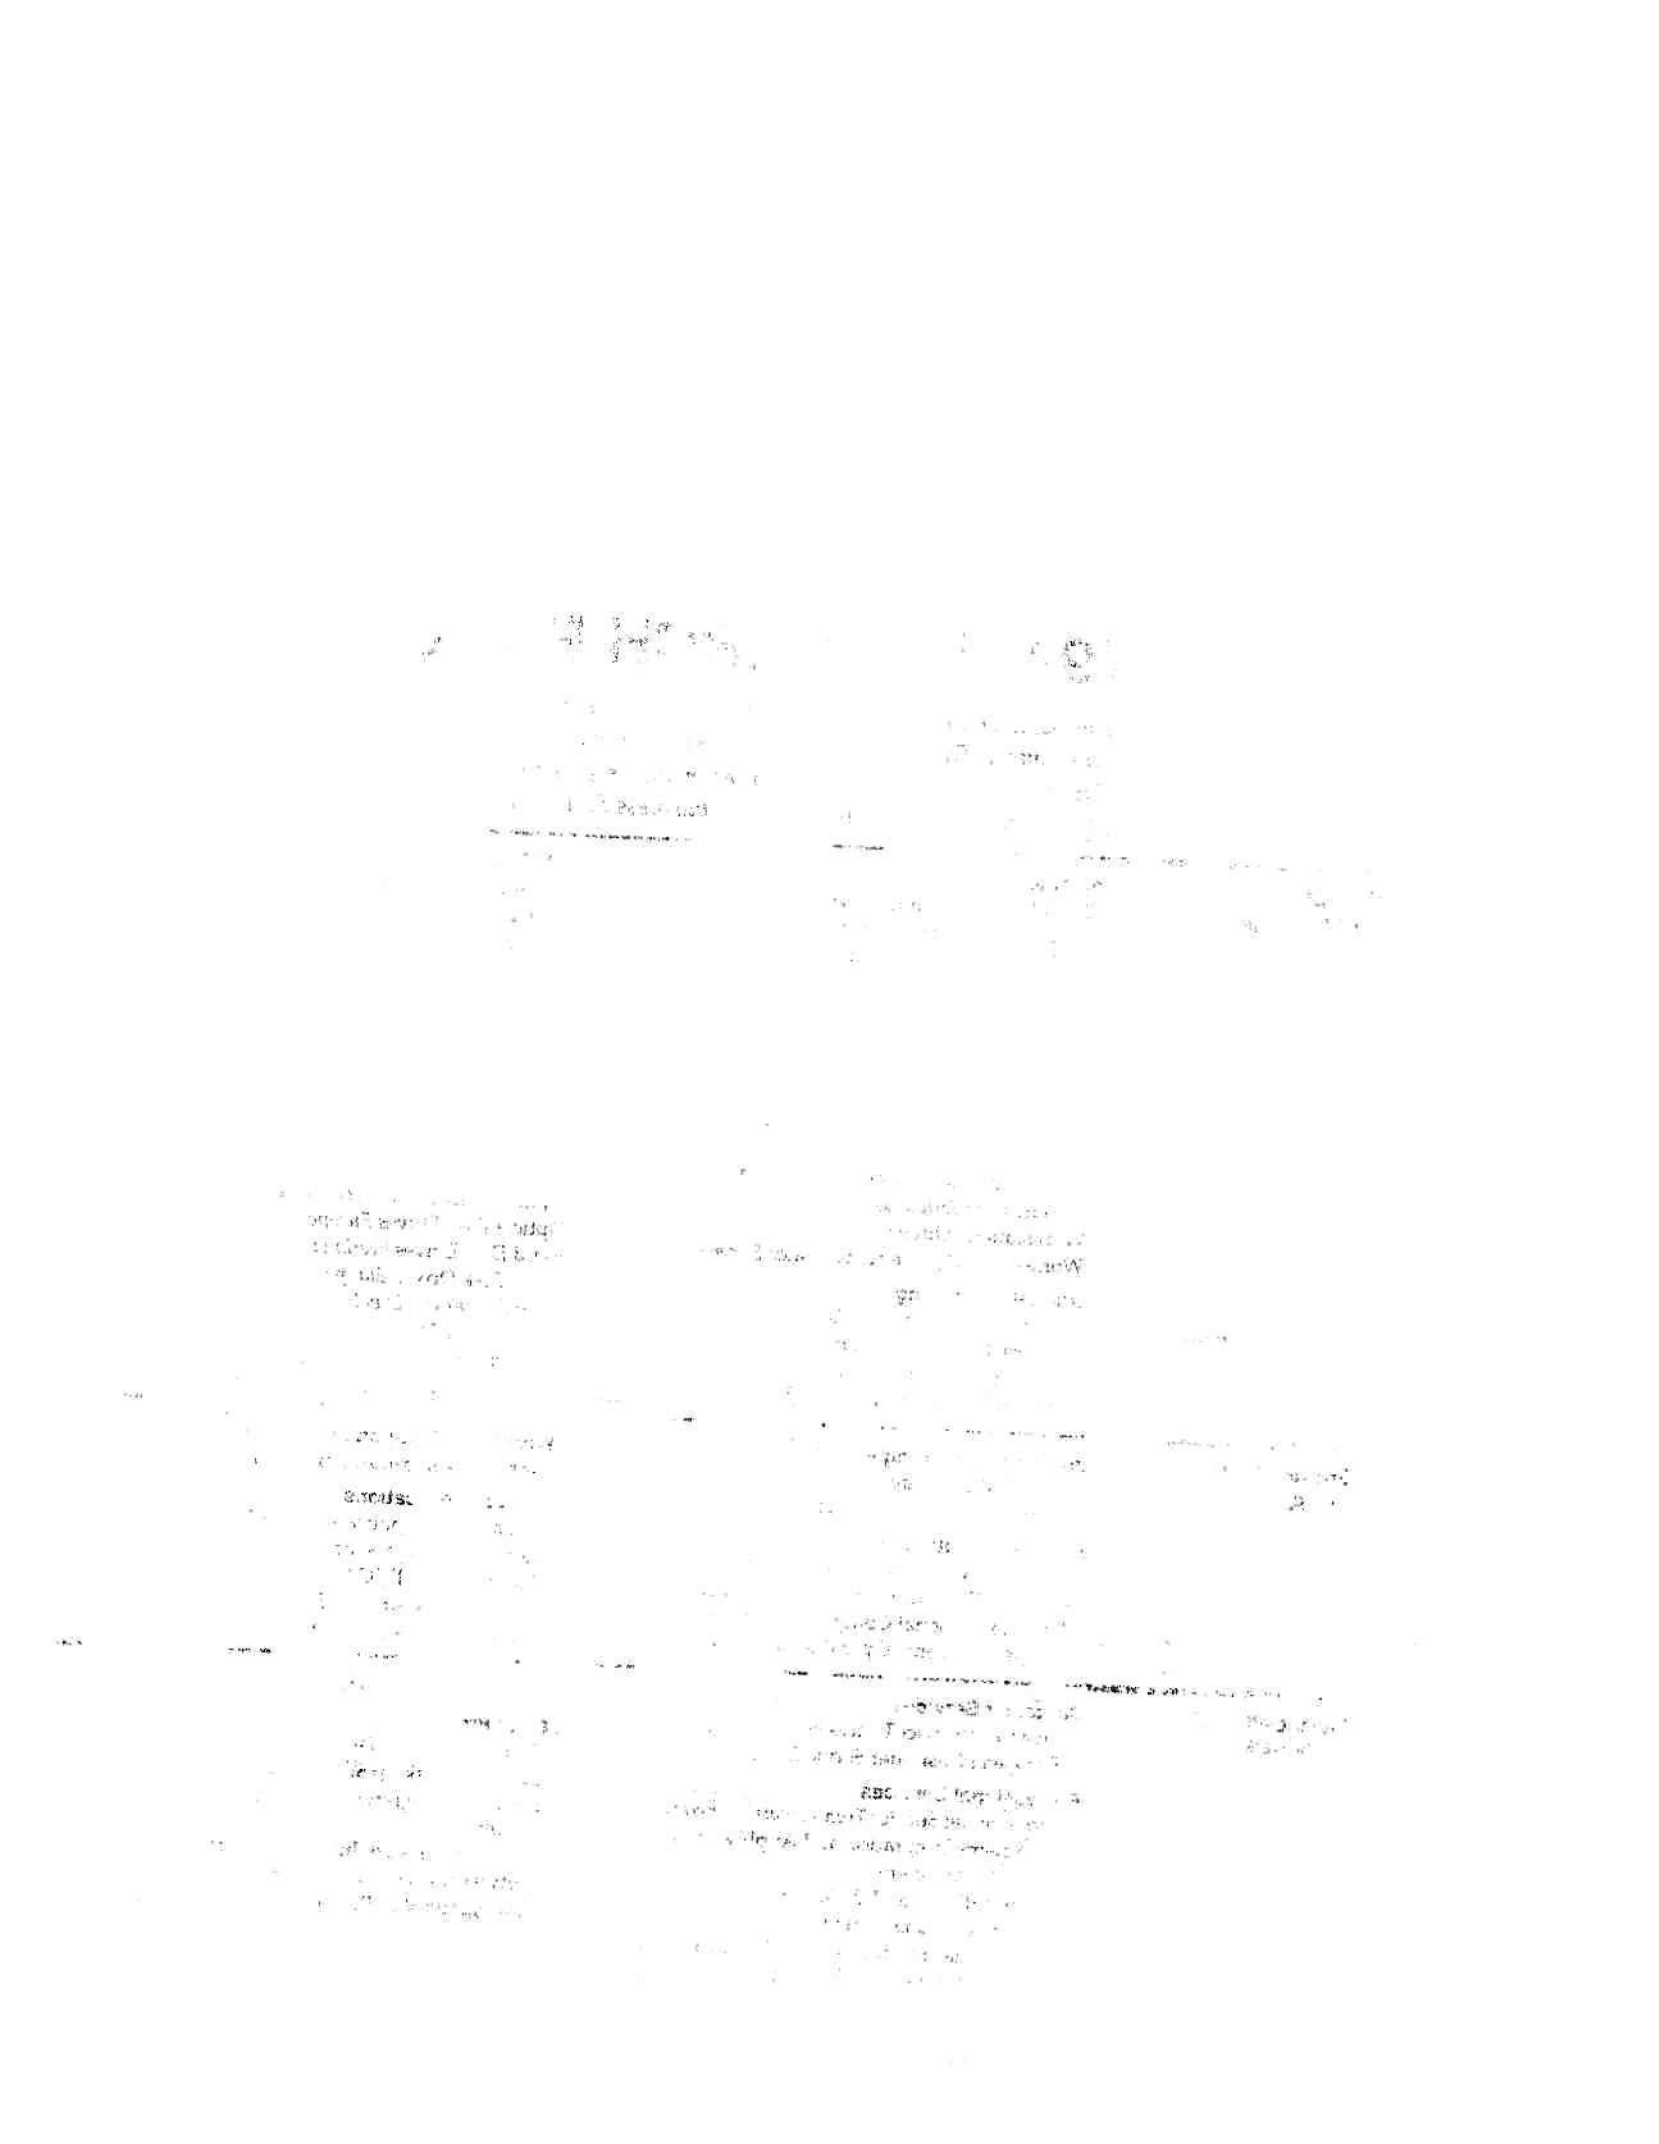

Global Market Insight for Information Technology Companies

1995
Publications Schedule

PC SEMICONDUCTORS AND APPLICATIONS

	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec
Perspective	❐	❐	■	❐	❐	■	❐	❐	■	❐	❐	■
Newsletters published on an event-driven basis												
Market Analysis	❐	❐	❐	❐	❐	❐	❐	❐	■	❐	❐	❐
Newsletters published on an event-driven basis												
User and Distribution Studies	❐	❐	❐	❐	❐	❐	❐	❐	❐	❐	❐	❐
Newsletters published on an event-driven basis												
Competitive Dynamics	❐	❐	❐	■	❐	❐	■	❐	❐	■	❐	■
Newsletters published on an event-driven basis												
Focus Studies	❐	■	❐	❐	■	❐	❐	❐	■	❐	❐	❐
Newsletters published on an event-driven basis												
Electronic Newsletters	❐	❐	❐	❐	❐	❐	❐	❐	❐	❐	❐	❐
Dataquest Alerts published on an event-driven basis												
Guides	❐	❐	❐	❐	❐	❐	❐	❐	❐	❐	■	❐

Dataquest

Program: PC Semiconductors and Applications
Program Manager: Jerry Banks

1995 Publications Schedule

PC Semiconductors and Applications

Complete Listing of 1995 Documents

Document Category	Document Type	Document Title	Arrival on Client Desk
Perspective	DQ Predicts	MPU Quarterly Forecast	March 1995
	DQ Predicts	MPU Quarterly Forecast	June 1995
	DQ Predicts	MPU Quarterly Forecast	September 1995
	DQ Predicts	MPU Quarterly Forecast	December 1995
Market Analysis	Market Trends	Data Processing Market Trends	September 1995
Competitive Dynamics	Product Analysis	MPU Complex in 1994	April 1995
	Product Analysis	Memory in PCs	July 1995
	Product Analysis	Controllers	October 1995
	Product Analysis	Overview	December 1995
Focus Studies	Focus Analysis	PC-Specific Standard Products	February 1995
	Focus Report	Graphics	May 1995
	Focus Report	Core Logic	September 1995
Newsletters	TBD	Special Topic	May 1995
	TBD	Memory	August 1995
	TBD	Communications	November 1995
Electronic NewsTakes	Fax Newsletter	Dataquest Alert	Event-Driven
Guides	Dataquest Guide	Semiconductors Market Definitions	November 1995

PSAMWW-ICON Qty: 1
Corporate Library
Dataquest Incorporated
1-9300
--INTERNAL DIST.--

Dataquest
a company of The Dun & Bradstreet Corporation

Corporate Headquarters
1290 Ridder Park Drive
San Jose, CA 95131-2398
United States
Phone: (1) 408-437-8000
Fax: (1) 408-437-0292

Boston Area
Nine Technology Drive
P.O. Box 5093
Westborough, MA 01581
United States
Phone: (1) 508-871-5555
Fax: (1) 508-871-6262

London Area
Holmers Farm Way
High Wycombe,
Buckinghamshire HP12 4XH
United Kingdom
Phone: (44) 1494-422722
Fax: (44) 1494-422742

Tokyo
Shinkawa Sanko Building
6th Floor
1-3-17, Shinkawa
Chuo-ku, Tokyo 104
Japan
Phone: (81) 3-5566-0411
Fax: (81) 3-5566-0425

Paris
Immeuble Defense Bergeres
345, av Georges Clemenceau
TSA 40002
92882 - Nanterre CTC Cedex 9
France
Phone: (33) 1-41-35-13-00
Fax: (33) 1-41-35-13-13

Munich
Kronstadter Strasse 9
81677 Munich 80
Germany
Phone: (49) 89-930-9090
Fax: (49) 89-930-3277

Taiwan
3/F, No. 87 Sung Chiang Road
Taipei 10428
Taiwan, R.O.C.
Phone: (886) 2-509-5390
Fax: (886) 2-509-5660

Korea
2506 Trade Tower
159 Samsung-dong
Kangnam-gu,
Seoul, Korea 135-729
Phone: (82) 2-551-1331
Fax: (82) 2-551-1330

Hong Kong
7/F China Underwriters Centre
88 Gloucester Road
Wanchai, Hong Kong
Phone: (852) 2-824-6168
Fax: (852) 2-824-6138

Singapore
105 Cecil Street
#06-01 The Octagon
Singapore 0106
Phone: (65) 227-1213
Fax: (65) 227-4607

Thailand
216/7 Moo 11,
Bangda Trad Road
Bangda, Pharkanong
Bangkok, Thailand
10260
Phone: (662) 361-6922
Fax: (662) 399-4883

Dataquest ALERT

251 River Oaks Parkway • San Jose • CA • 95134-1913 • Phone 408-468-8000 • Fax 408-954-1780

Intel Adopts Phoenix BIOS Firmware

In a surprise announcement, Intel and Phoenix Technologies revealed a technology alliance with broad and positive implications for both these companies and the industry at large. Intel will license PhoenixBIOS 4.0, the current version of the firmware that drives the I/O subsystem in many personal computers. This BIOS code plays a key role in shaping compatibility across PC platforms, and it is this compatibility, in turn, that has permitted the growth of the PC industry. Until now, Intel had been using BIOS code provided by American Megatrends Inc. (AMIBios), and there had been rumors that Intel was preparing its own BIOS as an alternative to the AMIBios and PhoenixBIOS. This announcement should kill those rumors.

Intel will acquire 6 percent of Phoenix's outstanding post-transaction shares for an investment of $10.9 million and has agreed not to increase its ownership of Phoenix beyond 19.9 percent for the next two years. This preserves Phoenix's independence, and allows it to work with other microprocessor vendors such as Cyrix and Advanced Micro Devices (AMD). These alternative x86 vendors need the support of BIOS developers like Phoenix. Intel's prospective entry into the BIOS business represented a serious threat to such developers. Shareholders of Phoenix Technologies, Cyrix, and AMD can breathe a sigh of relief on this account. (AMD's shareholders still need to hold their breath pending announcement of the results of a disappointing fourth quarter.)

Phoenix will open a development facility close to Intel's OEM Products and Services Division in Hillsboro, Oregon, and Intel can now cut back on its own efforts to staff an internal BIOS development group. BIOS writers need an intimate understanding of hardware and software issues and are extremely hard to recruit. This may have contributed to Intel's decision to outsource its BIOS technology from Phoenix.

The agreement should facilitate Intel's plan to continually introduce new technology into the personal computer market. Intel and Phoenix often duplicated each other's efforts as each added new features in support of initiatives like Plug and Play and Desktop Management Interface (DMI). This new approach leverages Phoenix's development efforts and eliminates a potential source of incompatibility down the road. For example, Intel will adopt the Phoenix/Microsoft/Compaq/National Semiconductor approach to host-side support for the forthcoming (in 1996) Universal Serial Bus (USB), a new method for connecting keyboards, mice, printers, modems, and other peripherals to personal computers. This means that developers of USB peripherals should not expect to encounter serious differences in the behavior of PCs based on PhoenixBIOS, at least compared with the BIOS on Intel motherboards.

Of course, a deal that creates winners may also create losers. In this case, the clear winners include Intel, Phoenix, system vendors (especially those that have developed a liking for Intel motherboards and PhoenixBIOS firmware), and alternative x86 MPU vendors, which should continue to expect support from a stronger Phoenix Technologies. The clear loser in this case is

American Megatrends Inc. (AMI), which had benefited greatly from Intel's success in driving its motherboards into more and more OEM accounts. AMI appears to have been outflanked by Phoenix in this matter, and we eagerly await its response.

By Nathan Brookwood
Microprocessor Analyst

Dataquest ALERT

251 River Oaks Parkway • San Jose • CA • 95134-1913 • Phone 408-468-8000 • Fax 408-954-1780

AMD Plugs Sixth-Generation Gap

Just when things were getting boring in the world of microprocessors, AMD livened them up with Friday's announcement that it would merge with tiny NexGen Microsystems. In a stock swap valued at over $850 million, AMD gains both the product designs to fill its new Fab 25 and the technologists to fuel further designs down the road. AMD will also gain access to the Pentium Pro (P6) market only a year after Intel ships its first production units and before Intel can position itself to deploy the Pentium Pro in the mass market. Intel has never before had to contend with any competitor having the design skills to create a high-end x86 CPU and the manufacturing skills to produce it in volume. All this could change with the AMD/NexGen merger, if AMD and NexGen can execute on their plan.

In Monday's joint press conference, AMD and NexGen explained the thinking that led to the merger, along with a few details about the Nx686 CPU they intend to introduce in the second half of 1996. The deal is clearly aimed at long-term benefits and has little impact on the near-term strategies of either company. AMD will continue to develop SSA/5-75 and K5 CPUs at its Austin, Texas, center and will redeploy those working on the K6 in Austin to the K5 effort or to the Nx686 effort (soon to be renamed "K6") at NexGen in Milpitas, California. Vinod Dham, who joined NexGen last spring after spearheading Intel's Pentium efforts, will lead the AMD sixth-generation campaign. The Nx586 will continue to be produced by IBM at its Essex Junction, Vermont, fab and will be integrated into AMD's marketing channels.

A midcourse correction to the Nx686 strategy, already planned by NexGen before the merger, will take on special significance. Until now, the Nx686, like its Nx586 predecessor, had a unique pinout, and this forced vendors using NexGen's CPU to use NexGen-specific chipsets, motherboards, and BIOS firmware. This clearly impacted market acceptance of the Nx586 and probably would have had the same effect on the Nx686. NexGen has finally acknowledged the need for external bus compatibility and plans to introduce the new K6 in a Pentium-compatible package—a change from the Nx686 discussed a few weeks ago at the Microprocessor Forum. This means the new K6 should plug into boards and chipsets designed to accommodate Pentium processors. AMD and NexGen even plan to increase dramatically the size of the on-chip level one cache to 64KB, compared with the 16KB of the Pentium Pro, to compensate for the slightly less-efficient nature of the Pentium bus compared with the high-tech GTL bus found on the Pentium Pro. Intel has been using the incompatibility of the Pentium Pro and Pentium sockets and buses to further its expansion into the systems business and in the process has been absorbing the profits that formerly went to infrastructure vendors, which provided chipsets and boards, and system vendors, which sold differentiated products. These infrastructure vendors are hurting, as anyone who has returned recently from Taiwan can testify. If AMD and NexGen are able to provide a competitive product in a meaningful time frame, in large quantities, and in an industry-standard package, it may represent the last hope for those accepting the challenge of competing with Intel.

The Nx686 will be fabbed in a 0.35-micron five-layer metal process at AMD's new Fab 25 in Austin. This plant is now running a four-layer metal process, and AMD is working to expand and add an additional level of interconnect. The Nx586 is manufactured by IBM at Essex Junction using a five-layer metal process.

In addition to a two-year head start on its K6 program, AMD also gains a team of veteran CPU designers who have learned the hard way what it means to produce and sell an x86-compatible CPU. Atiq Raza, who brought focus to NexGen's previously scattered efforts (remember when it was "NeverGen?") will become AMD's chief technical officer and brings a set of design skills and methodologies that nicely complement the expertise AMD has built up over the years in process technology. Vin Dham, a hard-driving, no-nonsense manager at Intel and more recently at NexGen, assumes overall responsibility for the program. Dana Krelle, a key NexGen marketer, also joined that company after years at Intel. The cultural and geographic distance between AMD Austin and Intel has limited

cross-fertilization between these two organizations, and this move clearly strengthens AMD's ability to compete in this arena.

By Nathan Brookwood
Microprocessor Analyst

Dataquest ALERT

251 River Oaks Parkway • San Jose • CA • 95134-1913 • Phone 408-468-8000 • Fax 408-954-1780

S3 Reports 300 Percent Growth in Earnings per Share

S3 Incorporated reported earnings per share of $0.20 for the third quarter of 1995, compared with $0.05 for the same period last year. Total revenue was $84.8 million for this quarter, up from $30.9 million a year ago. Gross margin improved, also, to 40.2 percent from 38.6 percent a year ago. A higher mix of 64-bit graphics products appears to have buffered S3's gross margin from the price erosion (and subsequent lower margin) of 32-bit products.

Manufacturing capacity is still a constraint and could limit unit growth. S3 is working to add capacity through a joint venture with United Microelectronics Corporation and Alliance Semiconductor, as well as new foundry agreements with NEC and Goldstar.

Dataquest Perspective

Strong demand for S3's graphics controllers and a portfolio of multimedia chips (MPEG decoders and audio codecs) provide a path for unit shipment growth. Every semiconductor company wants more dollars per PC, and S3 has leveraged its strength in graphics to expand into a broader range of multimedia products. The company has also taken a step toward the communications market with the acquisition of software company Floreat.

S3 has some interesting parallels with its competitor, Cirrus Logic. Both companies are technology leaders in the graphics market and have solid relationships with their customers. S3 and Cirrus compete aggressively with each other for the flagship accounts. One example of this competition is S3's design-win for some high-end Presario PCs from Compaq. Compaq is one of Cirrus' major accounts, so breaking into the Presario line was a coup for S3. PC OEMs are still designing graphics on the motherboard -- they demand graphics vendors who can provide volume shipments, bullet-proof software drivers, and proven technology. These requirements overwhelmingly favor the established graphics vendors for any motherboard-level designs.

The capacity issue is still a barrier to rapid growth, but S3 is striking deals to ensure continued access to manufacturing. Like its biggest competitor, S3 is pursuing multiple manufacturing relationships, including joint ventures and foundry relationships.

Demand for PCs is strong, and that spells demand for graphics and all the other ICs required to build those PCs. The challenge for S3 is to leverage its graphics design-wins into demand for its other multimedia products. The Cooperative Accelerator Architecture multimedia chipset will get its first real market test this Christmas. Sales of that chipset will demonstrate S3's ability to meet that challenge and diversify beyond its core graphics products.

By Geoff Ballew

Dataquest ALERT

1290 Ridder Park Drive • San Jose • CA • 95131-2398 • Phone 408-437-8000 • Fax 408-437-0292

VLSI Technology Says Two Chips Are Better than Four

VLSI Technology debuted the Lynx chipset for 586-class PCs on June 19. This chipset breaks new ground in terms of packaging and high level of integration. While other popular chipsets require four chips, the Lynx chipset requires only two. VLSI achieved this high level of integration by breaking with the quad flat pack (QFP) packaging tradition. The VL82C541 System Controller is sold in a ball grid array (BGA) package for the greater pin density. The VL82C543 ISA controller is sold in a QFP-208 package.

The Lynx chipset is positioned as a high-performance solution with support for EDO DRAM, multiple levels of concurrency across the MPU, DRAM, and expansion buses, and high PCI data rates for both short and extended transfers. An IDE interface is integrated into the VL82C541 System Controller and allows both IDE channels to be accessed simultaneously instead of multiplexing those signals. Other features include a two-wire interface that boosts performance when used with select VLSI graphics accelerators; there also are features that will allow ISA-based, legacy peripherals to run on the PCI bus.

This two-chip solution supports Pentium microprocessors as well as AMD's K5 family and Cyrix's M1 family of microprocessors. However, only Pentium processors can take advantage of the dual-processor features of the Lynx chipset. VLSI announced it would sample the chipset in August 1995, with volume production ramping sometime during the fourth quarter of 1995. Pricing for the two-chip solution is $35 per set in 10,000-unit volumes.

Dataquest Perspective

VLSI has created a feature-rich chipset with some compelling advantages over current chipset offerings. While competitors are pushing three- and four-chip solutions, VLSI has integrated all of their features into just two chips by using a high-pin-count (352) BGA package for the system controller. This allows system designers to save board space and enjoy the additional benefits that a BGA package has to offer, such as reduced coplanarity issues, as compared to QFP-300 packages and self-alignment during solder reflow. What will these designers do with the extra board space? We are eager to find out ourselves.

The BGA package does complicate the system board routing, compared to a QFP package, because of the multiple rows of pins. However, VLSI has chosen a Periphery BGA design so

this chip can be designed onto a board with only two signal layers. Some designers may choose a more expensive four signal-layer system card, but that is not required.

One other advantage VLSI has with the Lynx chipset is a special interface to select VLSI graphics accelerators. A system designer choosing a VLSI graphics controller may realize a gain of up to 15 percent in graphics performance. This appears to be a win/win situation where system OEMs enjoy higher levels of performance and VLSI enjoys some leverage from its chipset sales to enter a new business segment and increase its revenue per system.
By Geoff Ballew

Dataquest ALERT

1290 Ridder Park Drive • San Jose • CA • 95131-2398 • Phone 408-437-8000 • Fax 408-437-0292

Compaq and S3 Bring MPEG to the Motherboard in 1995

The Announcement

Compaq Computer Corporation and graphics chip provider S3 Incorporated June 13 jointly announced the availability of S3's new multimedia chipset and that Compaq would be using the new set in its consumer Presario line of products for the 1995 Christmas season. The chipset, called the Cooperative Accelerator Architecture, comprises an updated Trio graphics accelerator, an MPEG-1 decoder chip, and an audio DAC. Compaq has stated that the chips initially will target the mid- to high-end range of the consumer market--meaning PCs beginning at about $1,799. There will be an initial cost to the consumer of about $100 to put MPEG on the motherboard in 1995; Compaq plans to offer MPEG hardware as a feature across all consumer products by the second quarter of 1996.

The Chips

There are three chips in the set:

- The Trio64V+ graphics/video accelerator, $32 (10,000 units)
- The Scenic MX1 (for add-in cards) or Scenic MX2 (for motherboard or daughtercard) MPEG decoder, $35 (10,000 units)
- The Sonic/AD programmable audio DAC, $3.85 (10,000 units)

Systems that already have an S3 graphics accelerator with video capability can be upgraded with a "daughtercard" by the end user, for an expected cost of $100. Add-in cards featuring the new chipset are expected to cost about $250.

MPEG to Date

There are several implications of this announcement for board vendors, ISVs, and consumers. This announcement changes the potential for MPEG to finally catch on as *the* PC codec. The Open PC MPEG Consortium, cofounded last year by Jazz Multimedia and backed by both hardware and software companies including Microsoft and Sigma Designs, has been trying to promote MPEG despite initial resistance. Once Microsoft backed a Windows API, the DOS API

followed shortly after. To date, software developers have used Indeo and Cinepak and other soft codecs, including some proprietary ones, because MPEG lacked a critical mass of PCs with dedicated hardware for playback. Without an installed base for product, ISVs would not provide content for an incredibly small market. And because there was no compelling video content, board vendors claimed they could not sell boards to the consumer market. Dataquest numbers showed that the installed base for MPEG cards was barely 1 million by the end of 1994, further reinforcing developers' conceptions about the viability of MPEG, especially because cards were priced at about $350.

Initial Deployment of Hardware

Although the first implementation of MPEG will be on higher-end PCs, the potential for MPEG hardware is substantial and cannot be overlooked. Compaq is scheduled to ship this product during Christmas 1995, and several other key PC OEMs will be making similar announcements in the third and fourth quarters of this year with respect to MPEG hardware and/or software implementations. Although it is not clear if they too will be shipping product into the channels in time for this holiday season, MPEG hardware (or software) will be a given feature in 1996. If indeed Compaq is the first to ship, this may be a head start on the market.

Chip vendors have been working on integrated solutions combining acceleration and video playback for some time, but the potential for compelling video titles remained small with the mass market not buying add-in MPEG cards. ISVs will be able to develop titles using MPEG without haphazardly trying to guess the average installed base configuration. The general rule employed in development--write for the lowest common denominator--has kept video from being widely used for titles. In the next year, as all the OEMs migrate to hardware MPEG offerings in their products, ISVs will be ensured a base capable of playing their titles at full screen and 30 frames per second (fps) with little degradation and few artifacts.

What about the MPEG Board and Accelerator Card Vendors?

By 1996 board vendors will find it hard to compete with OEMs offering MPEG hardware on the motherboard that also includes video and graphics acceleration. The other factor is soft-MPEG. Microsoft recently announced it was licensing Mediamatics' soft MPEG algorithm for inclusion in future releases of its Windows operating system subsequent to Windows 95. With soft playback available, many consumers will not see the need to purchase an add-in card. Although hardware acceleration is preferable in terms of performance and Windows 95 will certainly test the boundaries because it is the first 32-bit operating system, there will remain only a small market for add-in cards in the retail arena. Hardware solutions will be implemented on the motherboard, and some add-in cards will be bundled at the OEM level. Furthermore, soft playback on some PC models may very well suffice for the average consumer.

The Effect on Motherboard Design

To implement the new architecture, the graphics accelerator must be an S3 video-enabled chip. This can lock out other graphics chip manufacturers that do not offer MPEG decoders of their own. OEMs may also choose an S3 graphics accelerator for a motherboard without MPEG

capability, in order to sell an MPEG-upgradable machine. Although S3 is pushing the Sonic audio DAC as part of its multimedia solution, this device only supports the MPEG chip, and the computer still must have a 16-bit sound card or the equivalent circuitry.

The Critical Hardware Influencer: The Ever-Changing Processor

Another factor influencing the adoption and deployment of MPEG in addition to Windows is the rapid rate of processor improvements between generations. The next chips from Intel--the P6, the P7, and beyond--will be able to handle much of the soft playback without constraining CPU bandwidth. If extra bandwidth is needed, the system will partition to any dedicated hardware. If such hardware exists on the motherboard itself, there is little need again for the add-in card.

Will This Jump-Start the MPEG Revolution?

S3 and Compaq are betting that MPEG will be the standard of choice for PC video, and are doing their best to make it so. By colaunching this chipset, and bringing Microsoft and other ISVs on the bandwagon, they may be able to finally give MPEG the market acceptance needed to take off.

Allen Leibovitch and Kathy Klotz

Dataquest ALERT

1290 Ridder Park Drive • San Jose • CA • 95131-2398 • Phone 408-437-8000 • Fax 408-437-0292

Compaq and S3 Bring MPEG to the Motherboard in 1995

The Announcement

Compaq Computer Corporation and graphics chip provider S3 Incorporated June 13 jointly announced the availability of S3's new multimedia chipset and that Compaq would be using the new set in its consumer Presario line of products for the 1995 Christmas season. The chipset, called the Cooperative Accelerator Architecture, comprises an updated Trio graphics accelerator, an MPEG-1 decoder chip, and an audio DAC. Compaq has stated that the chips initially will target the mid- to high-end range of the consumer market--meaning PCs beginning at about $1,799. There will be an initial cost to the consumer of about $100 to put MPEG on the motherboard in 1995; Compaq plans to offer MPEG hardware as a feature across all consumer products by the second quarter of 1996.

The Chips

There are three chips in the set:

- The Trio64V+ graphics/video accelerator, $32 (10,000 units)
- The Scenic MX1 (for add-in cards) or Scenic MX2 (for motherboard or daughtercard) MPEG decoder, $35 (10,000 units)
- The Sonic/AD programmable audio DAC, $3.85 (10,000 units)

Systems that already have an S3 graphics accelerator with video capability can be upgraded with a "daughtercard" by the end user, for an expected cost of $100. Add-in cards featuring the new chipset are expected to cost about $250.

MPEG to Date

There are several implications of this announcement for board vendors, ISVs, and consumers. This announcement changes the potential for MPEG to finally catch on as *the* PC codec. The Open PC MPEG Consortium, cofounded last year by Jazz Multimedia and backed by both hardware and software companies including Microsoft and Sigma Designs, has been trying to promote MPEG despite initial resistance. Once Microsoft backed a Windows API, the DOS API

followed shortly after. To date, software developers have used Indeo and Cinepak and other soft codecs, including some proprietary ones, because MPEG lacked a critical mass of PCs with dedicated hardware for playback. Without an installed base for product, ISVs would not provide content for an incredibly small market. And because there was no compelling video content, board vendors claimed they could not sell boards to the consumer market. Dataquest numbers showed that the installed base for MPEG cards was barely 1 million by the end of 1994, further reinforcing developers' conceptions about the viability of MPEG, especially because cards were priced at about $350.

Initial Deployment of Hardware

Although the first implementation of MPEG will be on higher-end PCs, the potential for MPEG hardware is substantial and cannot be overlooked. Compaq is scheduled to ship this product during Christmas 1995, and several other key PC OEMs will be making similar announcements in the third and fourth quarters of this year with respect to MPEG hardware and/or software implementations. Although it is not clear if they too will be shipping product into the channels in time for this holiday season, MPEG hardware (or software) will be a given feature in 1996. If indeed Compaq is the first to ship, this may be a head start on the market.

Chip vendors have been working on integrated solutions combining acceleration and video playback for some time, but the potential for compelling video titles remained small with the mass market not buying add-in MPEG cards. ISVs will be able to develop titles using MPEG without haphazardly trying to guess the average installed base configuration. The general rule employed in development--write for the lowest common denominator--has kept video from being widely used for titles. In the next year, as all the OEMs migrate to hardware MPEG offerings in their products, ISVs will be ensured a base capable of playing their titles at full screen and 30 frames per second (fps) with little degradation and few artifacts.

What about the MPEG Board and Accelerator Card Vendors?

By 1996 board vendors will find it hard to compete with OEMs offering MPEG hardware on the motherboard that also includes video and graphics acceleration. The other factor is soft-MPEG. Microsoft recently announced it was licensing Mediamatics' soft MPEG algorithm for inclusion in future releases of its Windows operating system subsequent to Windows 95. With soft playback available, many consumers will not see the need to purchase an add-in card. Although hardware acceleration is preferable in terms of performance and Windows 95 will certainly test the boundaries because it is the first 32-bit operating system, there will remain only a small market for add-in cards in the retail arena. Hardware solutions will be implemented on the motherboard, and some add-in cards will be bundled at the OEM level. Furthermore, soft playback on some PC models may very well suffice for the average consumer.

The Effect on Motherboard Design

To implement the new architecture, the graphics accelerator must be an S3 video-enabled chip. This can lock out other graphics chip manufacturers that do not offer MPEG decoders of their own. OEMs may also choose an S3 graphics accelerator for a motherboard without MPEG

capability, in order to sell an MPEG-upgradable machine. Although S3 is pushing the Sonic audio DAC as part of its multimedia solution, this device only supports the MPEG chip, and the computer still must have a 16-bit sound card or the equivalent circuitry.

The Critical Hardware Influencer: The Ever-Changing Processor

Another factor influencing the adoption and deployment of MPEG in addition to Windows is the rapid rate of processor improvements between generations. The next chips from Intel--the P6, the P7, and beyond--will be able to handle much of the soft playback without constraining CPU bandwidth. If extra bandwidth is needed, the system will partition to any dedicated hardware. If such hardware exists on the motherboard itself, there is little need again for the add-in card.

Will This Jump-Start the MPEG Revolution?

S3 and Compaq are betting that MPEG will be the standard of choice for PC video, and are doing their best to make it so. By colaunching this chipset, and bringing Microsoft and other ISVs on the bandwagon, they may be able to finally give MPEG the market acceptance needed to take off.

Allen Leibovitch and Kathy Klotz

Dataquest ALERT

1290 Ridder Park Drive • San Jose • CA • 95131-2398 • Phone 408-437-8000 • Fax 408-437-0292

Intel: Project Jump-Start

If Intel does not have a specific program with such a title, it is sure acting like it does. The company seems adamant not to let anything slow the penetration rate of its microprocessors. Intel is methodically isolating the critical paths that might limit the growth of a new x86 architecture and then systematically eliminating the obstacle.

In the Beginning ...

This trend might have originated back when Intel was trying to get the industry to embrace the PCI standard as the local bus of choice. Third-party PC core logic vendors were focusing their energies on the VESA, or VL, local bus because it had a jump-start on the market. Intel had a tough time convincing the core logic suppliers to support PCI. Intel took this effort onto its own shoulders because it believed that the basic architectural speed bottlenecks had to be eliminated in order for future PCs to get the full benefit of their higher-performance microprocessors. Intel believed that the VL local bus solution was just a stopgap measure and the next-generation PCs would need a higher-performance bus. As a consequence, Intel increased its focus into the PC core logic marketplace.

The Pentium Battle

An additional driver in this movement was the battle led by the PC behemoths IBM and Compaq to extend the life of the 486 rather than convert over to Pentium-based machines. This lack of support from the major PC companies for the Pentium caused Intel to take matters into its own hands. Intel added emphasis to its PC core logic program and added significant manufacturing capacity dedicated to building PC motherboards based on the Pentium microprocessor. Several PC OEMs took advantage of this and used Intel motherboards as the basis for their lines of Pentium machines. Foremost in this group are Dell, Packard Bell, and Gateway 2000. These companies were able to capitalize on the fact that IBM and Compaq were late realizing that the Pentium brand name was creating a tremendous consumer demand.

Next: The P6

Despite the fact that the Pentium demonstrated an unprecedented production ramp by shipping more than 5 million units in 1994, the production ramp actually was hampered by the lack of availability of 3V SRAMs used for cache memory in Pentium-based systems. Intel saw this SRAM problem as one that will continue for future generations of x86 product. The upshot of this was that cache memory availability was identified as an obstacle for growth for the P6 microprocessor. To eliminate this obstacle, Intel introduced the P6 with a dual-cavity package. One cavity housed the P6 processor and the other cavity housed a 256KB second-level cache. This does not mean that Intel intends to go into high-volume manufacture with this product; in fact, it would prefer that not be the case. It intends to make the transition to a conventional and lower-cost package as soon as it is convinced that there exists sufficient high-speed cache memory from third-party suppliers to meet the demand.

Dataquest Perspective

To accelerate the acceptance of new product architectures, Intel has invested substantial time and money into:

- PC core logic
- Motherboard manufacture
- Cache memory
- Dual-cavity package development

These are ongoing efforts, and we do not see the company backing off from this strategy. In fact, one could speculate, what does Intel perceive as the next obstacle or bottleneck? Here we can speculate, and the results of this speculation are quite interesting. The current device that is single-handedly slowing the rapid growth of PCs is the DRAM and its availability.

486-based machines ship with a minimum 4MB of DRAM and Pentium-based machines ship with a minimum 8MB of DRAM. These are not optimum configurations for either type of machine. Both types of machine would ship with 16MB, or more, if the product were available. Basically speaking, the PC has such a voracious appetite for DRAM that, if the DRAM industry were to double its capacity overnight, today's PC would swallow every DRAM the industry could produce. This is based upon today's mix of 486- and Pentium-based machines using Windows 3.1 and not some exotic new architecture or a future operating system. Also, Windows '95 is speculated to need more than 8MB to run efficiently.

Will Intel view the general lack of availability of DRAMs as a stumbling block? Will it, or can it, do anything about this shortage? If Intel has taught us anything, it is that it is very creative at finding solutions to problems. It would not be a surprise to see Intel announce a DRAM alliance that involved the commitment of Intel dollars to develop DRAM capacity. As a minimum, this would be a boon to its internal motherboard group, which would be guaranteed supply of this precious commodity known as DRAMs. It might further allow Intel to ship motherboards, fully loaded with more than the basic memory requirements, which could mean that Intel motherboards would be able to support the Windows '95 operating system. This would help Intel ramp its microprocessors faster and provide a compelling advantage for PC OEMs to purchase motherboards from the giant microprocessor company.

We are not stating that Intel will do such a thing, but if the company's track record is any indication, it will not sit still and let outside market forces dictate the acceptance of an Intel microprocessor. Intel can afford to invest a lot of money to protect its revenue stream.

Jerry J. Banks
jbanks@dataquest.com

Dataquest ALERT

1290 Ridder Park Drive • San Jose • CA • 95131-2398 • Phone 408-437-8000 • Fax 408-437-0292

The Pentium Takes a Leap

Intel Announces a 120-MHz Pentium

Intel Corporation has introduced a new, higher-performance version of its Pentium processor. According to the company, this Pentium processor operates at 120 MHz, delivers 140 SPECint92 and 103 SPECfp92 of performance, and is available in volume now. This Pentium is the first volume microprocessor to be built using a 0.35-micron process technology. The new technology will allow the Pentium processor die to shrink to half its current size, which should translate into higher performance, lower power, and higher reliability.

Manufacturing Takes a Bow

This announcement is as much a manufacturing achievement as it is a performance statement. Intel's transition from 0.8 microns to 0.6 microns began in earnest less than 18 months ago, and such a rapid transition to the 0.35-micron geometry indicates that Intel is much further along on its manufacturing road map than earlier estimates. In fact, Intel was caught somewhat off guard by its own success. The company was not prepared to take full advantage of its good fortune, and this first product on the 0.35-micron process is not as efficient as it will be when a fully optimized design is deployed to production.

Time-to-Market Innovation

The Pentium being announced is little more than an optical shrink of the existing 100-MHz design. In its current incarnation the 120-MHz Pentium is nearly identical to the 100-MHz part in terms of overall die size and I/O pad placement. This is in spite of the fact that the actual microprocessor core is much smaller than the 100-MHz product (also known as the P54C). The reason for this is that Intel added a wide silicon boundary around the core that makes it the same size as the P54C die. Intel did this so that it could make use of existing wafer testing equipment and existing qualified packages. This approach allowed the company to qualify and bring to market a much higher-performance product in a much shorter time span than if it had waited to optimize the design, packaging, and test. Such a technique in no way sacrifices quality or reliability and allows Intel the time it needs to optimize and qualify the manufacturing flow of the next version of the product. When the optimized version of the 120-MHz Pentium is ready, its smaller die size should about double the total number of die per

wafer. This will provide a tremendous increase in capacity for the product, in addition to the higher performance it will provide.

Faster Pentiums on the Way

Implications of this announcement are many. First or all, the Pentium receives an early performance boost. Also, the path to future higher-performance Pentiums is now pulled in substantially. We estimate that Intel will be able to take the Pentium up to 180 MHz on this 0.35-micron BiCMOS process. This will roll out over the next 12 months as the 133-, 150-, and finally 180-MHz version.

The P6 also Benefits

The P6 will also benefit from this early technology introduction. When the P6 technology announcement was made, it was based on a 0.6-micron BiCMOS process. With this process Intel was estimating a 133-MHz product that could produce a benchmark rating of 200 SPECint92. With this new process technology, the P6 should be hitting the 200-MHz clock rate within the next 12 months. We estimate that this will increase the P6 benchmark rating to a little more than 300 SPECint92, which will put it in the performance range of most RISC microprocessors that will be shipping in volume in the same time frame. An additional benefit for the P6 is that it will more than likely be moved from its existing BiCMOS process to a more standard CMOS process. This will have the added benefit of higher yield, shorter manufacturing throughput times, lower power, and lower cost. We expect the Pentium to remain a BiCMOS product.

Manufacturing Rollout

The rollout of the process that started in Intel's Oregon facility known as D1 will next move to Intel's Fab 11 in New Mexico, which is expected to come online during the second half of 1995. The new Arizona fab known as Fab 12 will come online within six months of this occurrence. It is clear that maintaining a manufacturing leadership position is a very high priority for Intel. Intel's capital spending outlay was $2.4 billion in 1994; it is expected to be about $2.9 billion in 1995. Furthermore, the company's R&D budget was about $1.1 billion in 1994 and is expected to increase to $1.3 billion in 1995.

Price, Availability, and Positioning

The 120-MHz part initially will sell for $935 in 1,000-piece quantities and is available now. We expect this price to be rapidly reduced to make room at the high end for soon-to-be-announced faster versions of the Pentium, as well as the next-generation P6.

Intel is treading on new ground in terms of price and position. It is an enviable position to have such a plethora of high-performance products released in such a short time frame, but it calls for the development of a new price reduction model. Now that Intel has successfully introduced its next-generation microprocessor, the P6, in half the time of its predecessor

products, Intel's marketing team will be stretched to get the most out of each successive microprocessor without slowing the adoption of next-generation product. Conversely, it does not want the buying public to fall into the trap of always waiting for the next version, thus killing the market for today's product.

Competitors Must Try Harder

Intel's competitors all have fifth-generation x86 products with more sophisticated architectures than the Pentium. This typically allows them to execute more instructions in a single cycle and makes them more efficient. However, Intel is making rapid strides in the technology side that may erase this architectural advantage.

Intel is the world's largest semiconductor vendor. It is No. 1 in terms of microprocessor revenue and shipments, and it continues to outspend its competitors. With the money it is investing, the direction it is heading, and the speed at which it is moving, Intel will be a very formidable competitor for some time to come. Because it is the leader, the followers are bound to be able to add more sophisticated design techniques. The only problem is that the leader is now running very fast and may be very difficult to catch.

Jerry J. Banks
jbanks@dataquest.com

Courtesy of Dataquest's new semiconductor program entitled "Semiconductor Directions in PCs."

Dataquest ALERT

1290 Ridder Park Drive • San Jose • CA • 95131-2398 • Phone 408-437-8000 • Fax 408-437-0292

DRAM the Memories: Full Speed Ahead

NeoMagic Corporation today unveiled its long-awaited graphics solution for the portable marketplace. The fundamental difference between NeoMagic's approach and that of the other vendors in the graphics controller market is the integration of display memory on-chip with the controller. Yes, that means DRAM and logic on a single chip.

Dubbed the MagicGraph NM2070 series, this family of graphics controllers contains the following functions:

- SuperVGA graphics accelerator
- RAMDAC
- Frequency synthesizer
- VESA video I/F
- Power management
- LCD controller
- Display memory

In this first product series, display memory consists of 1MB of on-chip DRAM.

Although it did not disclose much in the way of process parameters, the company did disclose that it was making use of a 16Mb DRAM process—the die is "slightly larger than a 16Mb die and the package is a little more expensive." As Table 1 shows, although a 16Mb DRAM process may match today's advanced logic processes in terms of lithography, its interconnect capability is extremely limited. Most logic processes today have at least three layers of metal for interconnect. These additional metal layers are used to reduce die size and resistance and, as a result, increase performance. Furthermore, logic processes are tuned to maximize transistor speed. Although DRAM designs also attempt to maximize speed, the first priority of the process and device engineers is a functioning and reliable DRAM cell. As a consequence, it is difficult to build a high-speed logic design with a DRAM process.

Table 1

Nominal 16Mb DRAM Process

Lithography	Metal Layers
0.55 Micron	2

Source: Dataquest (March 1995)

NeoMagic claims to have been able to work through this dilemma, and the result is a high-speed, low-power logic design tightly coupled with on-chip DRAM.

Portable Market Feature Set: Power, Size, and Performance

The high-level integration, when combined with on-chip power management and a 3.3V power supply, results in a specified power consumption of less than 400mW. The company claims that a discrete solution made up of a controller and two DRAMs will consume nearly 1W. Such a savings will be welcomed by the portable computer community of OEMs that have a very miserly power budget with which to work.

Another feature of the high level of integration is that the company is able to provide a 128-bit memory interface to the controller while simultaneously reducing total system package count from three chips to one and pin count from about 270 to a single 176-pin thin quad flat pack (TQFP). This results in a reported board space savings of 60 percent. This is another powerful feature for the space-constrained portable world.

An additional benefit is that the 128-bit interface allows for a high-memory bandwidth. This wide bus helps offset the fact that the logic transistors are not optimized for performance. The result is performance levels reported to exceed those of discrete 64-bit desktop systems despite having a controller designed with transistors optimized for a DRAM process.

Dataquest Perspective

If NeoMagic can bring this chip to market, it will have succeeded where many others have failed. Combining logic with DRAM is not a new story by anyone's measure; designing a product that is manufacturable is. The benefits of integration are well known, well chronicled, and are at the very heart of the semiconductor industry's phenomenal growth. The battle cry of "smaller, cheaper, faster" has been heard from the lips of designers for more than three decades and shows no signs of letting up. If NeoMagic is successful in this product launch, it will no doubt encourage many others to try to follow suit with integrated products of their own. It is not yet clear whether NeoMagic wields enough intellectual property clout to discourage competitors from following its lead.

A critical issue not disclosed was that of price. The NeoMagic solution will only be successful in mainstream portable applications if it can compete with existing discrete solutions on a price basis as well as provide all the stated features. Selling the value of smaller board space, lower power, lower inventory, and higher quality, among other things, has been tried in the past and

is only successful in a handful of applications. Designers have grown accustomed to receiving the benefits of integration for free and will become very shortsighted if the cost of the new solution does not at least match that of the prior solution.

NeoMagic is focusing its efforts on the portable market, where it offers compelling reasons to buy. Although this market is much smaller than the desktop market, it is growing quickly, which allows the company to establish a presence in the market as a technology innovator. As its technology comes down the learning curve and its growth demands require a larger market, the company will have plenty of time to move into the high-volume world of desktop PCs.
Jerry J. Banks

jbanks@dataquest.com

This Dataquest Alert is distributed as a courtesy of Dataquest's new semiconductor program PC Semiconductor Application Markets Worldwide, which covers semiconductors in the PC. For information on this program contact your Dataquest representative or send an e-mail to the author at jbanks@dataquest.com.

Dataquest ALERT

1290 Ridder Park Drive • San Jose • CA • 95131-2398 • Phone 408-437-8000 • Fax 408-437-0292

PowerPC Goes Consumer

A New Spin for the PowerPC

IBM and Motorola have jointly announced the PowerPC 602 microprocessor. The PowerPC 32-bit instruction-set-compliant 602 is intended to compete in the low-cost, yet high-performance space of the consumer electronics market.

The 602 is a 32-bit architecture that can issue up to two instructions per cycle to its four parallel execution units: integer execution unit, floating point unit, branch unit, and load/store unit. It follows the lead of the 603e by offering an improved load/store unit that requires only one clock cycle for cache access. The floating point unit is optimized for single-precision operations, but double-precision operations are trapped for software emulation. The FPU is pipelined such that when the pipeline is full a single-precision multiply-add instruction can complete every clock cycle. The branch unit allow the 602 to perform branch speculation and instruction. The 602 includes a power management unit that, among other features, can dynamically power down each of the execution units when not in use on a per-cycle basis.

Instructions for the 602 can execute out of order. However, the instructions complete and write back in program order.

The 602 interface protocol allows multiple masters to compete for system resources through a central external arbiter. The 602 provides a three-state coherency protocol that supports the modified, exclusive, and invalid (MEI) cache states. This protocol is a compatible subset of the modified/exclusive/shared/invalid (MESI) four-state protocol and operates coherently in systems that contain four state caches. The 602 supports single beat and burst data transfers for memory accesses and memory-mapped I/O. A single multiplexed bus interface is used for transferring both 32-bit addresses and 64-bit data.

Table 1 shows the vital statistics of the 603e.

The 602 has already received public endorsements from The 3DO Company and from Matsushita Electric Industrial Co. Both companies will be including the 602 in future 3DO designs such as the 3DO Multiplayer from Matsushita and the M2 Accelerator from 3DO.

Table 1

Technical Specifications of the PowerPC 602

Technology	0.5µm CMOS, four-layer metal
Die Size	7.04 x 7.04mm (49 mm^2)
Number of Transistors	950,000
Instruction Cache	4KB (write back)
Data Cache	4KB (write back)
Performance	40 SPECint92 at 66 MHz (est.)
Signal I/O	98
Power Supply	3.3V ±0.3V
Power Dissipation	1.2W at 66 MHz
Packaging	Plastic quad flat pack (144 pins)

Source: IBM/Motorola

Dataquest Perspective

Although the 602 has not yet seen silicon, the estimated specifications do not seem unreasonable and should be quite achievable. And given the track record of the PowerPC alliance for working first-time silicon, we expect the 602 to follow suit, as it does not appear to be attempting any radical changes in the basic PowerPC architecture.

The power performance ratio of the 602 does not match that of other recent announcements of companies or alliances going after the same marketplace. However, with the high profile 3DO design wins for the 602, the PowerPC family becomes one of the broadest families of microprocessors in the industry, with design wins ranging from the consumer game market to notebooks, desktop computers, and all the way to technical workstations. With such a broad range of applications in the fold of the PowerPC alliance, it would seem a sure bet that at least some of these designs will result in high volumes for the alliance.

Jerry J. Banks

Dataquest ALERT

1290 Ridder Park Drive • San Jose • CA • 95131-2398 • Phone 408-437-8000 • Fax 408-437-0292

PowerPC 603 Gets a Face-Lift

Minor Modifications Provide Significant Performance Boost

IBM and Motorola have jointly announced the PowerPC 603e microprocessor. This is a much improved version of the 603 architecture that has not yet seen its way into volume manufacture, and for a variety of reasons may not.

The basic architecture of the 603e remains the same. It is a 32-bit architecture that can issue three instructions per clock to five parallel execution units: integer execution unit, floating point unit, branch unit, system unit, and load/store unit. The primary architectural difference between the two processors is that the 603e has doubled the sizes of both the on-chip data and instruction primary caches. Table 1 shows the vital statistics of the 603e.

Table 1

Technical Specifications of the PowerPC 603e

Technology	0.5µm CMOS, four-layer metal
Die Size	8.4 x 11.67mm (98 mm^2)
Number of Transistors	2.6 million
Performance	120 SPECint92 and 105 SPECflt92 and 100 MHz
Signal I/O	165
Power Supply	3.3V ±5 percent
Power Dissipation	3W at 100 MHz
Packaging	Ceramic quad flat pack (240 pins) Ball grid array (16 x 16)

Source: IBM/Motorola

The dramatic performance improvement of the 603e stems from three basic changes in the design. One is the increase in size of both on-chip caches from 8KB each to 16KB each. This doubling of on-chip cache will result in a significantly higher "hit ratio" when the processor is looking at the on-chip cache for its next instruction and/or data, before having to look off-chip to either a slower secondary cache or an even slower yet main memory location. The second is a modified load/store unit that eliminates a full cycle, when compared to the 603. The

load/store unit in the original unit requires two cycles to execute a cache access, whereas the modified unit on the 603e requires only a single cycle. Both of these features provide a significant performance boost, but the 603 had some physical limitations limiting its clock speed.

The joint Apple, IBM, and Motorola PowerPC design group in Somerset, New Yorkperformed extensive analysis of the physical layout and timing of the original 603 and was able to identify some critical speed paths that prevented the design from operating at higher clock speeds. Once these were identified and eliminated, the path was open to increase the clock speed of the device. This combined effort resulted in a dramatically improved product with a SPECint to frequency ratio of 120 percent.

The 603e has also become a little more user-friendly, from a system perspective. The internal clock multiplier is programmable in one-half-step increments from 1x to 4x. A beneficial example of this finer granularity is evident in 40-MHz motherboard designs that can now be upgraded to the 100-MHz 603e by selecting the 2.5x clock multiplier. Such a feature allows the system designer more flexibility in optimizing performance and system cost. It also allows existing designs based on the 80-MHz 603 to be easily upgraded to the 100-MHz 603e with minimal system implications.

Dataquest Perspective

The 603e should see immediate success in applications that were originally targeted to use the original 603 product. It is no secret that system vendors were less than satisfied with the original 603's performance. The 603e's performance rating of 120 SPECint92 at a typical power consumption of 3W is quite respectable and should translate into a very competitive notebook platform. At the very least we would expect to see Apple announcing Power Book notebooks based upon the 603e in the very near future. IBM has announced that it will use the 603e in a mobile product, but it is unclear what success such a product will have in the market unless it will run a Macintosh operating system.

Jerry J. Banks

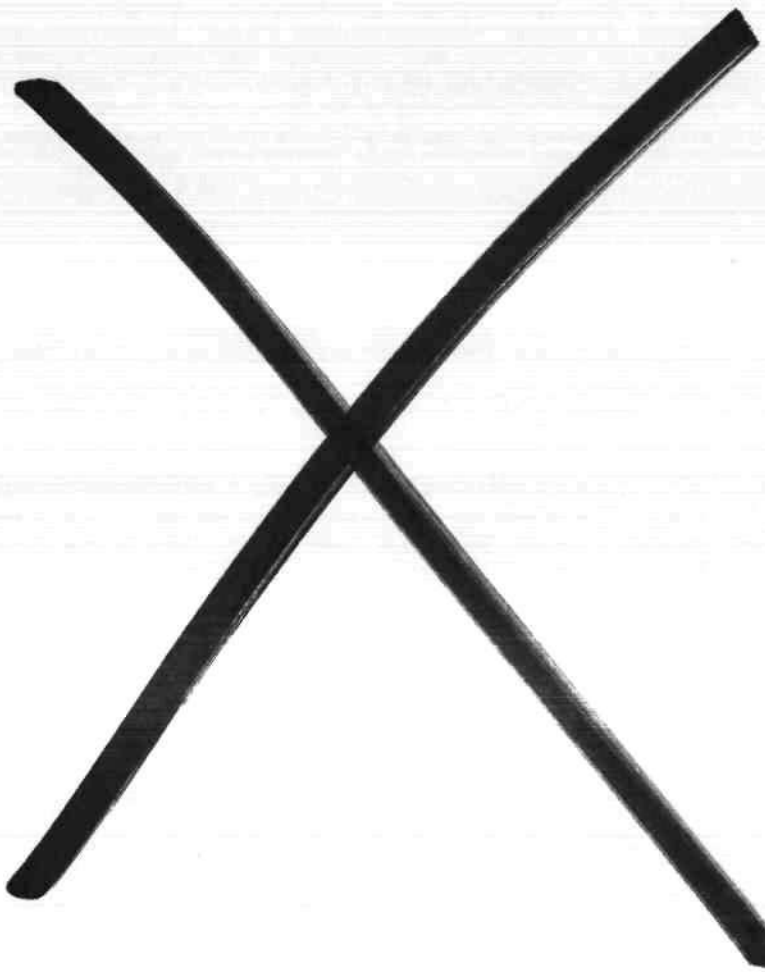

Dataquest

Perspective

PC Directions in Semiconductors

Technology Analysis

1 + 1 = 1? Integration Propels Modem and Audio Functions on a Collision Course

Abstract: *Modem and audio functions are becoming standard features on PCs. These features are checklist items for the consumer and small office/home office (SOHO) markets and are making inroads into the corporate PC market, as well. This document provides forecasts for both the modem and audio chipset markets and discusses the collision course that these two technologies are taking. Synergy between these two markets is driving integration at both the board and chip level, creating a variety of hardware options for adding these and other features to PC designs.*
By Geoff Ballew

Major Shift in Modem Chipset Market

PC user demand for connectivity is pushing modems closer to becoming a standard feature for PCs. The transition of home PCs from computing tools to communications tools, as well as changes in the business world, are driving this demand. Modems make new sources of information accessible by connecting personal computers to other computers over vast public and private networks. Figure 1 and Table 1 show Dataquest's forecast for worldwide shipments of PCs and modem chipsets. Modem chipsets are defined as one or more chips that provide the modem functions.

Note that the five-year compound annual growth rate (CAGR) for modem chipsets is lower than that for PCs. This does not mean that fewer PCs will have modem functions; in fact, quite the opposite is true. It does highlight a shift in how modem functions are sold to the end user, which in turn affects the market dynamics for modem chipsets.

Dataquest

Program: PC Directions in Semiconductors
Product Code: PSAM-WW-TA-9502
Publication Date: December 11, 1995
Filing: Competitive Dynamics

Figure 1
Worldwide Forecast for PC and Modem Chipset Markets

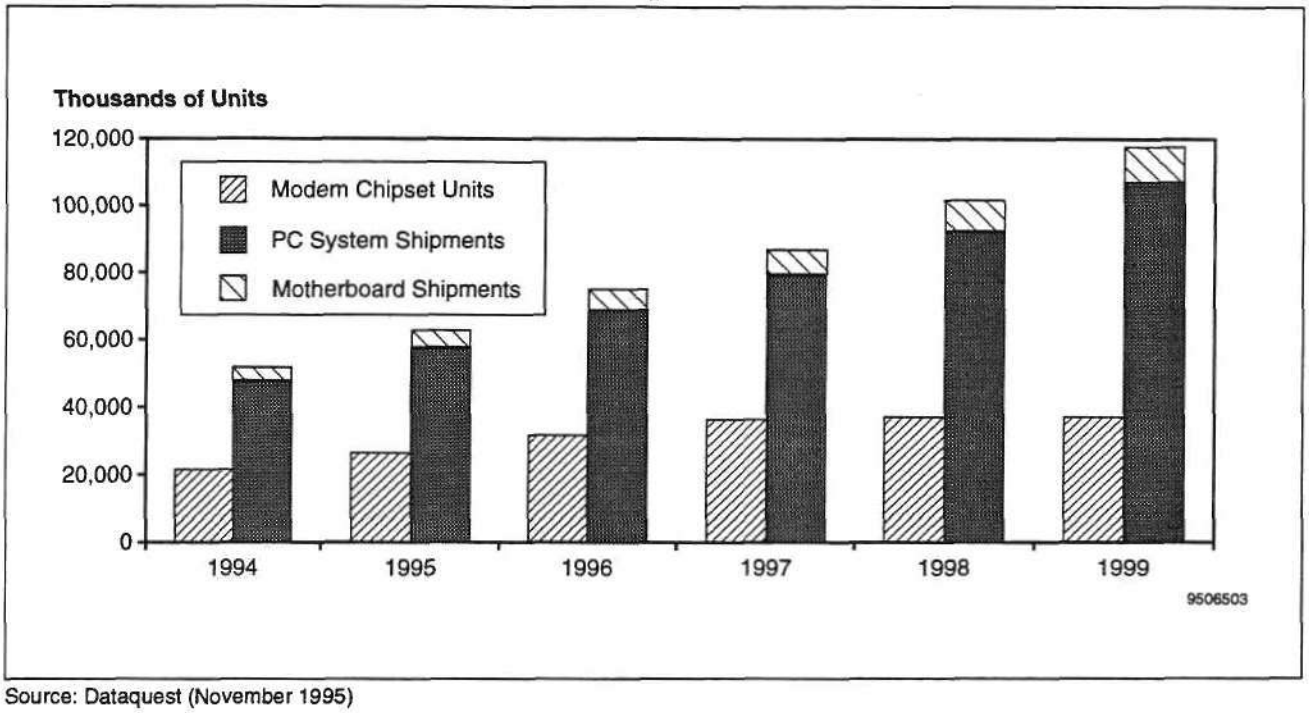

Source: Dataquest (November 1995)

Table 1
Worldwide Forecast for PC and Modem Chipset Markets

	1994	1995	1996	1997	1998	1999	CAGR (%) 1994-1999
Modem Chipset Units (K)	21,531	26,489	31,733	36,396	37,211	37,166	11.5
Modem Chipset ASP ($)	26.3	22.2	21.0	20.3	20.2	20.7	-4.7
Modem Chipset Revenue ($M)	566	587	667	738	751	770	6.3
PC System Shipments (K)	47,876	57,564	68,712	79,313	92,355	107,000	17.4
Motherboard Shipments (K)	4,404	5,476	6,738	7,992	9,582	11,190	20.5
Total PC Unit Shipments (K)	52,280	63,040	75,450	87,305	101,937	118,190	17.7
Ratio of Modem Chipsets to Total PCs (%)	41.2	42.0	42.1	41.7	36.5	31.4	

Source: Dataquest (November 1995)

Modem sales today are a combination of retail sales and sales to PC OEMs. Over the next four years, modem sales to PC OEMs will grow faster than retail sales and will be the primary driver for market growth in 1998 and 1999. For the purposes of this document, retail sale is defined as the sale of a modem directly to an end user, whether it is an internal or an external modem. The market for external modems is almost entirely a retail market; PC OEMs that provide modem functions in their PCs prefer add-in card modems.

Retail sales have been robust for two reasons. First, the increasing popularity of online services and access to the Internet has caused many new users to purchase modems for their existing PCs, which did not have modems installed. Second, many modem users upgrade their modems more frequently than their PCs because modem speeds have been increasing and the cost to upgrade is relatively small. Modem speeds are slow compared to other data transfer rates in even a low-end PC, and modem users are sensitive to that performance bottleneck.

This market is changing, however. Dataquest believes that the retail sales of modems directly to end users will decline, and that sales to PC OEMs will drive market growth at the end of this decade. There are two reasons for this. First, modems are now standard features for PCs targeted at the consumer and small office/home office (SOHO) markets, so it is unnecessary for the buyer to purchase a modem directly. Second, modem speeds are not increasing as fast as they have previously, which reduces the incentive for users to buy a new modem. Dataquest expects modem speeds to peak at 28.8 Kbps, the current standard known as V.34, for communication over standard telephone lines.

Without faster modems, the retail upgrade market will decline gradually as the installed base becomes saturated with V.34 modems. The addition of digital simultaneous voice and data (DSVD) to V.34 modems in 1996 will help postpone the decline until 1998. DSVD will allow modem users to transfer voice and data simultaneously over a single telephone line. Although some techniques for this exist today, they are less sophisticated than DSVD and involve switching between voice and data modes rather than processing both streams simultaneously. Not every modem user will have a V.34 modem, but those users willing to upgrade will not find new modems compelling if they do not have significantly enhanced features. Other users may be content with V.32bis modems, which operate at 14.4 Kbps.

Users who demand greater performance may upgrade to digital modems using integrated service digital network (ISDN) lines instead of standard analog telephone lines (plain old telephone service—POTS). ISDN modems and services are available in some locations today, but cost and availability issues will delay widespread adoption. Dataquest forecasts that 11 million ISDN lines will be in use by the year 2000, so ISDN modems will remain high-end products through that time. In other words, analog modems are expected to dominate the market at least through the end of this decade. Figure 2 shows the market positioning of modem standards through 1999.

Figure 2
Market Positioning of Modem Standards

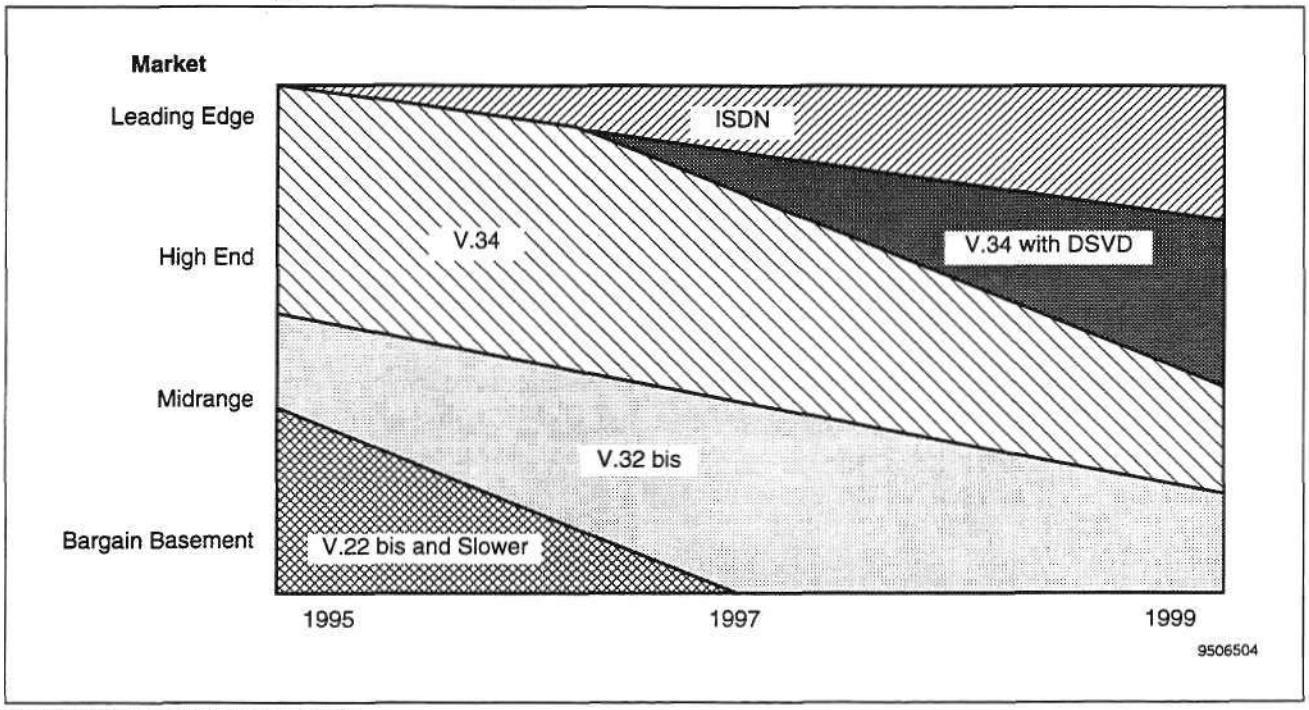

Source: Dataquest (November 1995)

Modem chipsets are rarely implemented on the motherboard because of Federal Communications Commission (FCC) certification issues. Modems are considered FCC Class A devices, while motherboards are generally FCC Class B devices. However, if a designer adds a modem chipset to a motherboard design, the motherboard becomes a Class A device and must meet the more stringent Class A specification. For this reason, modem chipsets are typically implemented on add-in cards or as external modems.

The Audio Market Marches On

Useful? Maybe. Fun? Yes. Either way, audio features are more common in PCs today than ever before. With capabilities ranging from "business" audio (that is, no Sound Blaster compatibility for games) to high-end wavetable solutions, PCs just want to be heard these days. Audio capabilities have already become standard features for PCs targeted at the consumer and home office markets. The corporate market is a tougher nut to crack because many information systems managers still question the need for audio features in the workplace. Notebook PCs are one exception to this rule; the target market is the corporate world, but audio features are becoming common. Of course, Apple Macintosh computers have always had built-in audio capabilities, whether they were selling into consumer or business accounts. Figure 3 and Table 2 show the forecast for worldwide shipments of PCs and PC audio chipsets. Audio chipsets are defined as one or more chips that provide the audio functions. Note that 1998 is a milestone; it is the first year that Dataquest believes more computers will be configured with audio features than without.

Figure 3
Forecast for PCs and PC Audio Chipsets

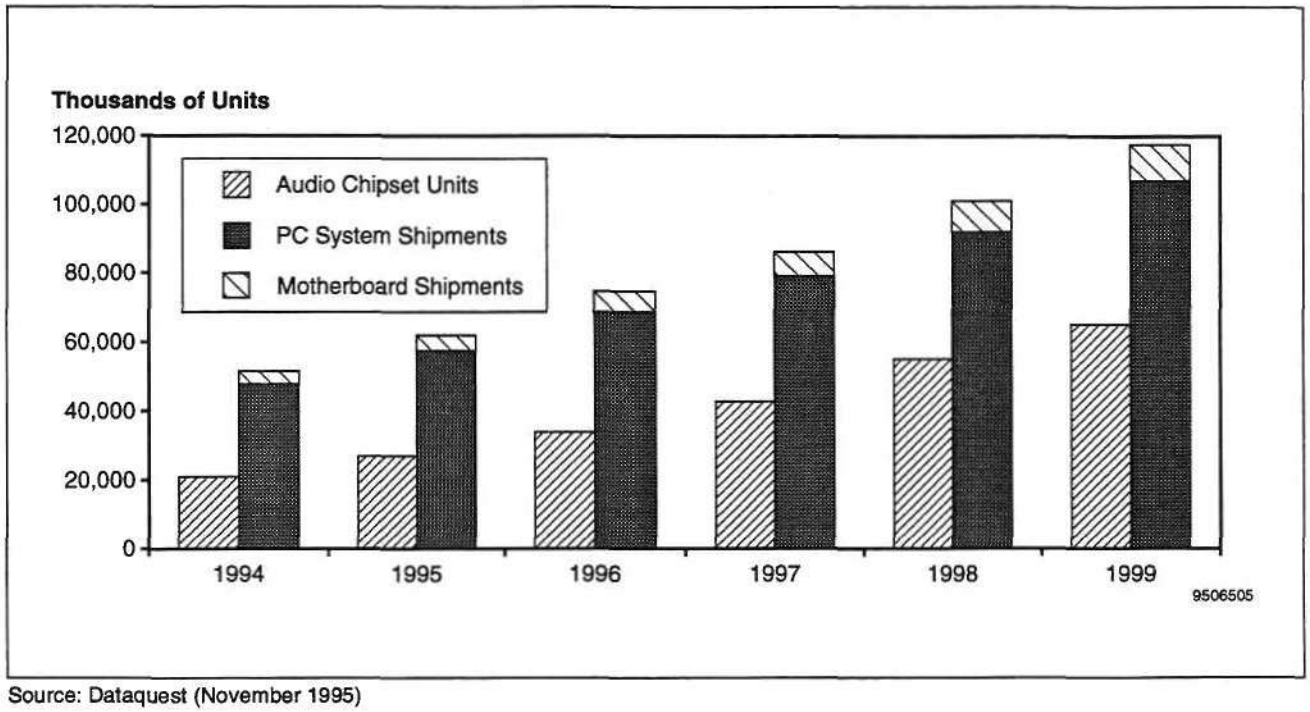

Source: Dataquest (November 1995)

Table 2
Forecast for PCs and PC Audio Chipsets

	1994	1995	1996	1997	1998	1999	CAGR (%) 1994-1999
Audio Chipset Units (K)	20,996	26,978	33,889	42,762	55,173	65,184	25.4
Audio Chipset ASP ($)	16.14	15.00	13.64	12.48	11.55	10.56	-8.1
Audio Chipset Revenue ($M)	338.9	404.7	462.2	533.7	637.2	688.3	15.2
PC System Shipments (K)	47,876	57,564	68,712	79,313	92,355	107,000	17.4
Motherboard Shipments (K)	4,404	5,476	6,738	7,992	9,582	11,190	20.5
Total PC Unit Shipments (K)	52,280	63,040	75,450	87,305	101,937	118,190	17.7
Ratio of Audio Chipsets to Total PCs (%)	40.2	42.8	44.9	49.0	54.1	55.2	

Source: Dataquest (November 1995)

The PC audio market is a mix of add-in cards and motherboard implementations. For 1995, the audio add-in card market represents about half of the total demand for audio chips, while motherboard implementations account for the other half. Motherboard implementations are heavily skewed toward Macintosh and notebook PCs, but that is changing. Multimedia PCs with x86 microprocessors will increasingly have audio functions implemented on the motherboard. Dataquest expects motherboard implementations to represent over 70 percent of the demand for audio chipsets by 1999, with add-in card implementations absorbing the remainder. In that year, 40 percent of PCs and motherboards sold will have audio functions integrated onto the motherboard, compared to only 20 percent in 1995.

The CAGR for audio chipset shipments is substantially higher than that for total PC unit shipments (25.4 percent versus 17.7 percent). This highlights the migration of audio features into a higher percentage of PCs sold.

Multifunction Solutions Will Dominate

As demand grows for audio and modem functions in PCs, these two feature sets are propelled on a collision course. The synergy between the modem and audio markets is tied to common signal-processing elements as well as competition for system resources. If a PC OEM buys an audio card and an add-in card modem, it is purchasing some semiconductor devices twice; both cards require their own codecs and amplifiers.

Competition for resources is also an issue because the modem chipset vendors have added call-processing features such as voice mailboxes and speakerphone capabilities to their chipsets. If a user wants to answer a telephone call using the modem as a speakerphone, then the modem needs control of the speakers as well as control of the microphone. But the speakers and microphone are plugged into the sound card, and nobody wants to pay for a second set of speakers and microphone for the modem. The situation is worse if the user happens to be playing a game or an audio CD and is actually using the speakers at the time the telephone rings.

Multifunction cards are the trend of the future. Board vendors have taken the first integration step by creating dual-function cards that use two chipsets—one audio chipset and one modem chipset. Figure 4 shows the layout of one of these combination cards.

This example is an Aztech audio card with a Rockwell chipset for the modem functions and other telephony features. All of these functions can be integrated into a single chipset, but there is little overlap between the audio chipset vendors and the modem chipset vendors. Rockwell has over 60 percent of the modem chipset market but does not participate in the audio chipset market. Aztech designs its own chipsets for the audio market but does not design modem chips. Two exceptions are Creative Laboratories and Cirrus Logic; both of these companies sell audio products as well as modem products. Cirrus' products are sold at the chip level, while Creative sells at both the chip level and the board level. Either one of those companies could conceivably produce a dual-function chipset that provides both feature sets by integrating its current products.

One company has already introduced a dual-function chipset. The Quartet chipset from Sierra Semiconductor provides audio, modem, fax, and speakerphone functions in a single chipset.

Figure 4
Board Layout for a Dual-Function Audio and Modem Card

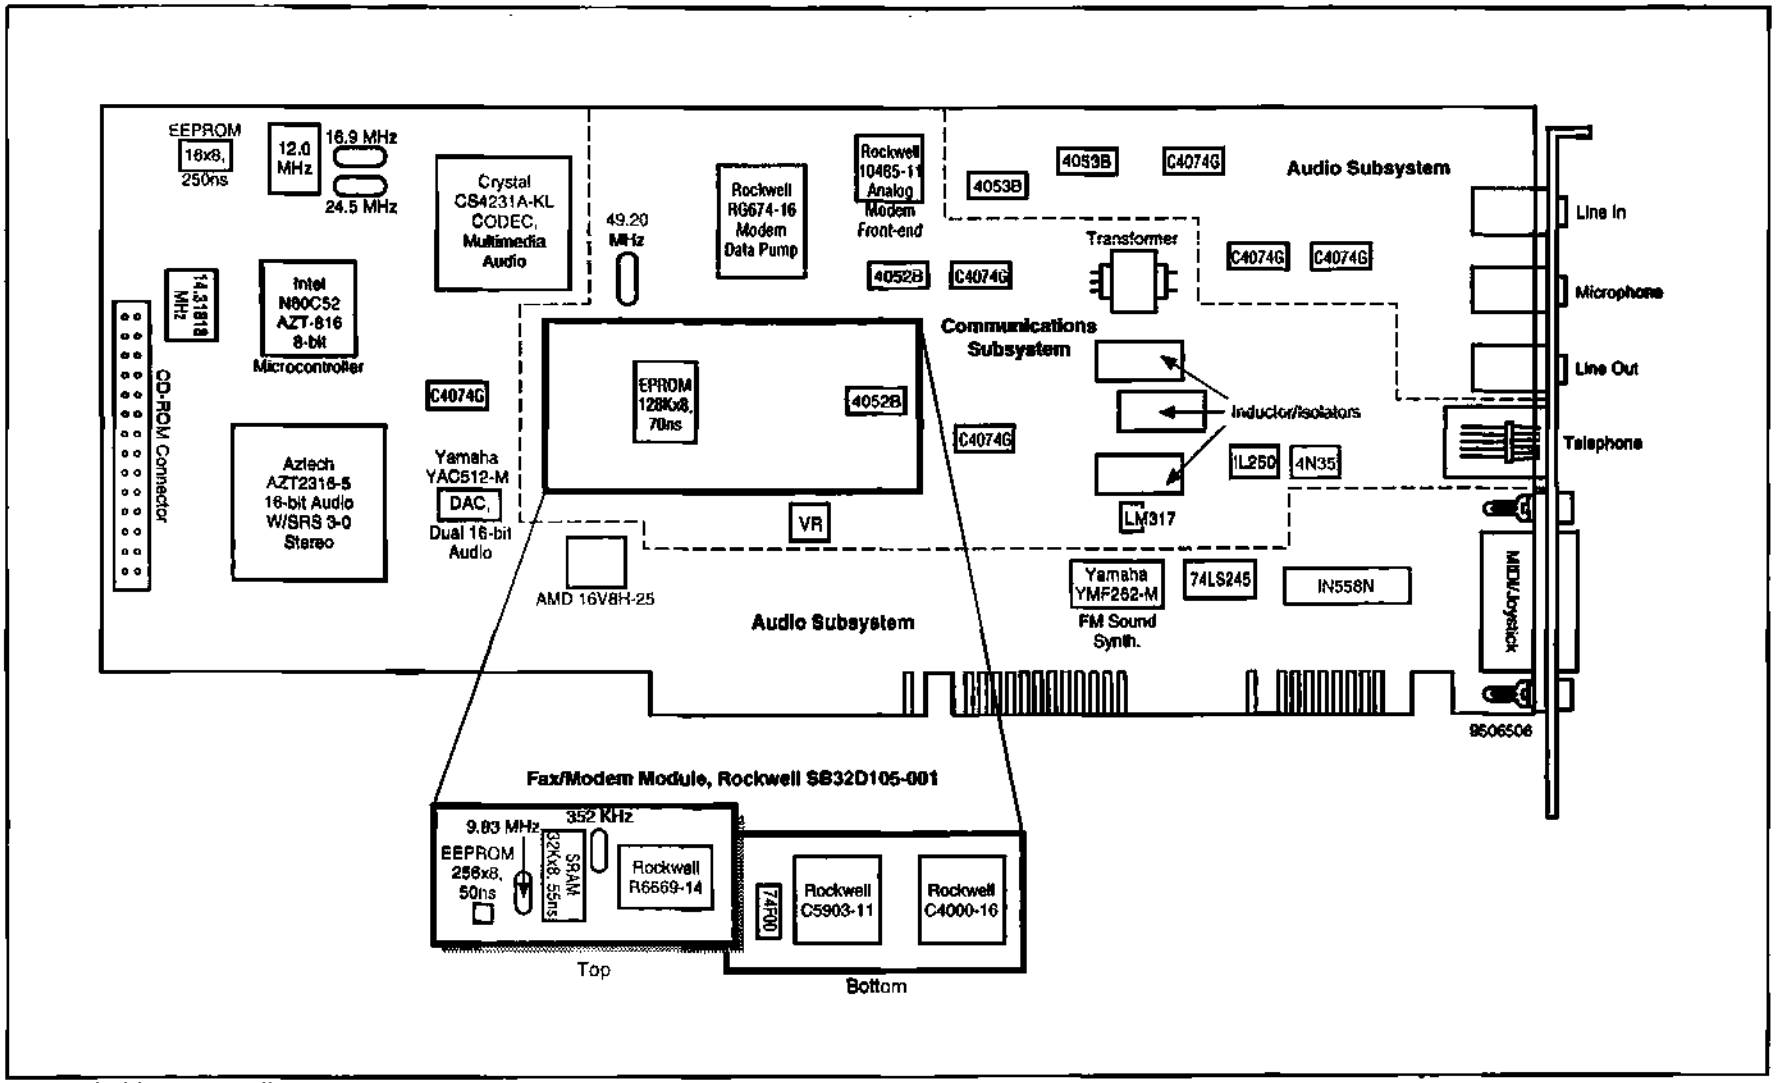

Source: Dataquest (November 1995)

There are also alternatives to dedicated chipsets. Dedicated chipsets are defined as chipsets in which device compatibility and compliance to standards is designed at the hardware level. Several companies have stepped away from fixed-function designs and are using special processors that run software to provide specific functions and device compatibility. Two examples are known as digital-signal processors (DSPs) and media engines. Figure 5 shows the current application areas for media engines and DSPs.

Media engines and DSPs are more expensive than fixed-function chipsets, but they are also more flexible. PC OEMs and end users can add new functions simply by adding software rather than adding silicon. That level of flexibility could make them less expensive than buying several fixed-function chipsets. Media engines are generally more powerful than DSPs and can handle some graphics and motion-video processing as well as offload some tasks from the CPU. Three companies promoting media engine products are Philips Semiconductor, Chromatic Research, and MicroUnity. Philips and Chromatic are targeting the PC market for their products, which are named Trimedia and Mpact, respectively. MicroUnity appears to be targeting non-PC consumer devices such as digital set-top boxes. One example of a DSP-based solution is the Mwave chip from IBM. Mwave chips are shipping in some IBM PCs. Figure 6 shows a road map for hardware implementation trends.

Figure 5
Function Map for Alternative Hardware Solutions

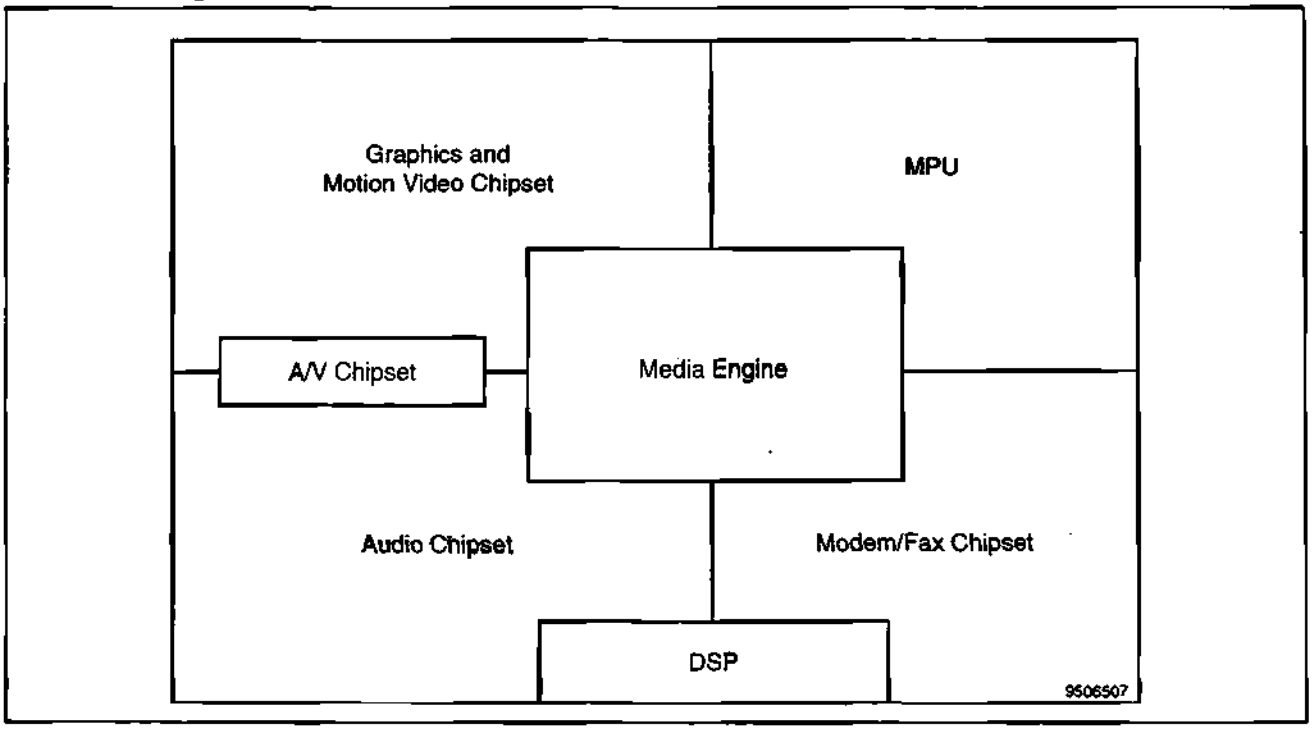

Source: Dataquest (November 1995)

Figure 6
Hardware Implementation Trends

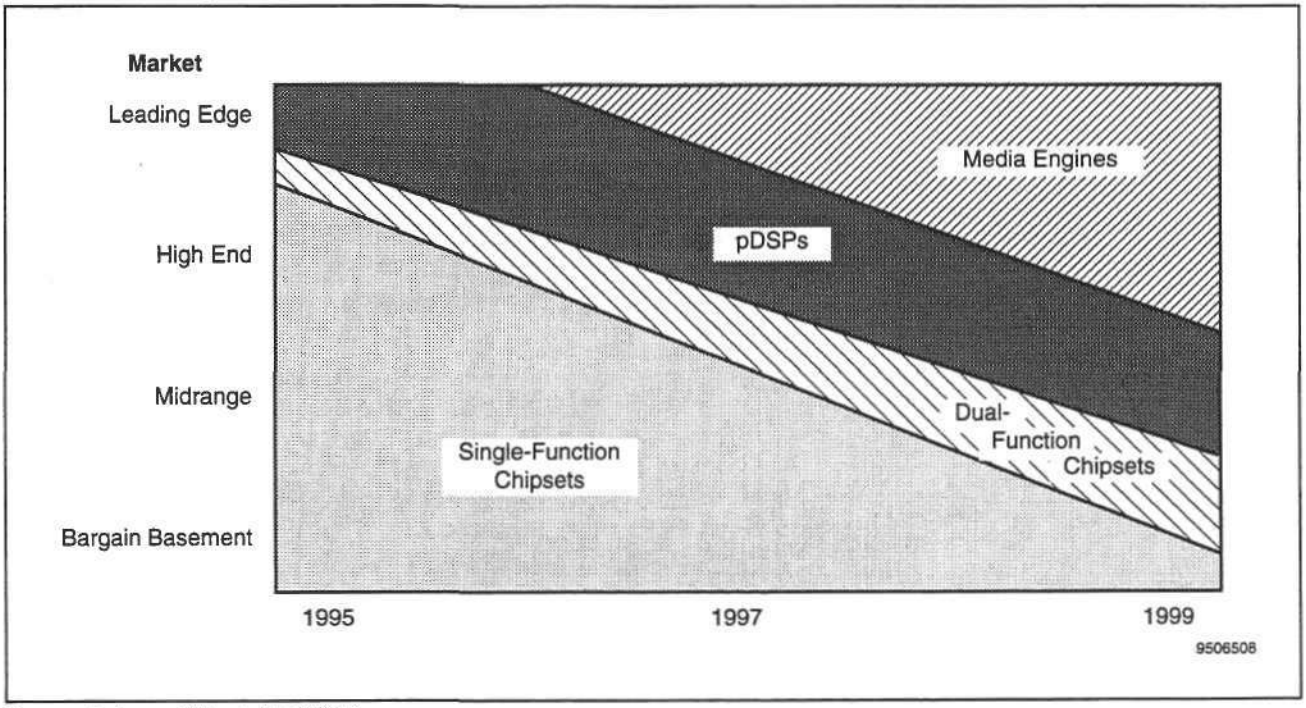

Source: Dataquest (November 1995)

Dedicated chipsets are the least expensive solutions today, so alternative solutions need to be cost-effective on a price-per-function basis to gain significant market share.

Dataquest Perspective

PCs are being used more frequently as communications tools, and this shift is driving the market for both audio and modem chipsets. Semiconductor vendors are adding value to their products by integrating multichip solutions for audio and modem functions onto a single chip as well as pushing performance levels to new heights. Audio and modem functions have already been integrated at the board level by a couple of vendors, and that trend will continue until those functions are integrated into a single dual-function chip. At the same time, those dedicated chipsets will be challenged by alternatives that use generic processing capability and provide functional compatibility through software.

Media engines could take significant market share from dedicated chipsets in 1997 or 1998, but cost of implementation will continue to be an issue and may delay broad-based acceptance. DSPs have a lower price point than media engines and could challenge dedicated chipsets in 1996, but cost remains an issue. If those solutions can handle digital modem functions at full ISDN speeds, they could dominate the top end of the modem market because of their additional functionality. Otherwise, the cost issue will determine how quickly the market shifts from dedicated chipsets to new alternatives.

PSAM-WW-TA-9502 Qty: 1
Mr. Kevin McClure
Dataquest Incorporated
.
.,
--INTERNAL DIST.--

For More Information...
Geoff Ballew, Industry Analyst ..... (408) 468-8676
Internet address ..... gballew@dataquest.com
Via fax ..... (408) 954-1780

Dataquest
a company of The Dun & Bradstreet Corporation

Perspective

Semiconductor Directions in PCs

Research Brief

A Review of the International Solid State Circuits Conference (ISSCC)

Abstract: *The IEEE's recent annual International Solid State Circuits Conference, ISSCC 1995, is a showcase of the R&D efforts of the leaders in the semiconductor industry. This year's conference was a fascinating show heralding the "giga" era. Gigabit DRAMs took center stage, while gigabaud communications, gigaflop DSPs, and giga-instructions-per-second CPUs will fill out the rest of the information superhighway chip complement. Intel announced the next-generation x86 microprocessor at this San Francisco conference. Many other extremely interesting innovations in memories, communication, data conversion, and even neural networks were presented, along with papers selected on their value, innovation, and the fact that a device was made to work in a laboratory environment.*
By Jim Handy (Memories Worldwide) and Jerry Banks (Microcomponents Worldwide)

The Conference Is Over, but Many Interesting Memories Linger On

DRAM

In light of the new era of gigabit DRAMs (2 to the 30th power, or 1,073,741,824 bits), NEC, Hitachi, and even Toshiba (the last of which was the only one of the three to not show a gigabit device) had evening parties to show off the cream of their research staffs to the press and analysts. At the NEC showing, one general manager waved in front of the audience a tie bar built around a gigabit die that looked to be almost two inches on each side. NEC's approach to producing this behemoth is to manufacture a wafer full of 256Mb DRAMs, find four adjacent working devices, and saw the wafer in such a way as to leave these four together, effectively making a 1Gb monolithic device. Surprisingly, NEC does not call this wafer-scale integration.

Dataquest

Program: Semiconductor Directions in PCs
Product Code: PSAM-WW-RB-9502
Publication Date: May 29, 1995
Filing: Perspective

A more interesting trend in the show was the embracing of the concept of merging ASIC and DRAM technologies. There was a workshop about this subject the day before the conference, a paper by Toshiba showing an "Embedded DRAM" to be used with a sea of gates, and a position taken in the "DRAMs for Graphics" evening panel that the best choice for future graphics controllers would be to combine the frame buffer with the graphics accelerator ASIC.

What would drive this? Certain DRAM applications, such as MPEG decoders for set-top boxes and multimedia, and graphics accelerator cards, use a relatively fixed amount of DRAM on the order of 24Mb. This is limited by the resolution of the display, a statistic that does not change much over an extended period. Soon (before the turn of the century), 24Mb will be less than half of the density of a standard 64Mb DRAM, leading researchers to wonder how to use the rest of the die. It certainly looks as though this approach could catch on, but we still do not see too much SRAM added to today's ASICs (with rare exceptions, such as Sony's cache chip). So we cannot truly tell whether the DRAM/ASIC approach will be well accepted down the road.

One evening session promised to be an honest debate about the best DRAM architecture to use when implementing video systems. The session promised so much that it played to a packed auditorium. Unfortunately, six of the seven panelists were from semiconductor companies whose corporate position advocated one architecture or another, with only one company being an actual user of the devices, so the session ended up looking like the sort of posturing usually seen in trade shows. The sentiment we heard from audience members we later encountered was one of disappointment.

SRAM

Three extremely fast SRAMs were revealed by Toshiba (with Hewlett-Packard), NEC, and Hitachi, spanning access times from 3ns to 1ns. The scariest was a wave-pipelined device of Hitachi's that used a multiphase phase-locked loop incorporating a voltage-controlled oscillator made out of a rectangular array of 96 amplifiers. This design actually came from another ISSCC paper presented in an earlier year.

Flash Memory

As usual, everybody in the industry seemed to be ready to show off some new flash memory design. This is encouraging, but it would be refreshing to see some real activity in the market, enough to pose a threat to the stranglehold Intel and AMD have on flash, with a combined market share of almost 90 percent.

Seven flash papers were presented by Intel, Mitsubishi, Hitachi, Matsushita (with SunDisk), Toshiba, and, for the first time, Samsung. All but one featured 3.3V operation, indicating the resounding agreement that flash is destined for portable operation. Intel presented its experimental approach of storing multiple bits on a single cell, thereby multiplying the bit storage and slashing the per-bit price. Two of the remaining six papers dealt with 16Mb designs, and the rest with 32Mb designs, showing a sincere desire on the part of the flash manufacturers to get into business on the highest rational densities.

Ferroelectrics

Two papers (one from Rohm and one from Matsushita, Symetrix, and the University of Colorado) focused on the use of ferroelectrics. It surprised us a little that the 1Gb DRAMs presented by Hitachi and NEC in the DRAM session did not use ferroelectric dielectrics, as we have been predicting for years. The ferroelectric papers were about techniques to design nonvolatile memories using this technology. One offered a lifetime of 10^{12} cycles (1 trillion), while the other offered 10^{13} (10 trillion), highlighting the major strength of ferroelectrics over flash of a much higher number of read/write cycles. Today's flash devices offer about 10^{6} (1 million) cycles, or a 10-millionth as many.

The P6: Not Just Another x86

Intel's new microprocessor, now known as the P6, is proof positive that Intel has truly halved the development time of succeeding generations of x86 microprocessors. The P6 comes just two years after the Pentium. All prior generations of x86 microprocessors had taken four to five years between generations.

Intel does not call its approach to the newly revealed P6 CPU a multichip module, although the difference between Intel's approach and an MCM was utterly lost on us, as well as on others in the audience. The P6 uses a CPU chip with level-one caches similar in size and construction to those on the Pentium, then augments this with a 256KB secondary cache SRAM die mounted in the same package. Lots of fancy new architectural twists, many borrowed from the more advanced variations of the i960, have been added to push the performance significantly above that of the Pentium. The processor is designed to be connected to as many as three other CPUs on the same bus, supporting tightly coupled multiprocessing architectures. (Ironically enough, Intel Scientific Computers only uses loosely coupled processor arrays, so this twist will not be of any help to it.)

The heart of the architecture is what Intel refers to as "Dynamic Execution," the unique combination of three processing techniques the P6 uses to speed up software—multiple branch prediction, data-flow analysis, and speculative execution:

- Multiple branch prediction: First, the processor looks multiple steps ahead in the software and predicts which branches, or groups of instructions, are likely to be processed next. This increases the amount of work fed to the processor.
- Data-flow analysis: Next, the P6 analyzes which instructions are dependent on each other's results, or data, to create an optimized schedule of instructions.
- Speculative execution: Instructions are then carried out speculatively, and possibly out of order, based on this optimized schedule, keeping all the chip's superscalar processing power busy and boosting overall software performance.

The P6 is very fast, will rapidly scale to higher frequencies, and is built for high-volume manufacture. It uses a fine-grain 11-stage pipeline in which no single stage is required to perform a significant amount of work. This will allow Intel to rapidly increase the performance by simply increasing the clock. We believe that the present product announcement of 200 SPECint92 at 133 MHz is quite conservative. We would not be too surprised to see the P6 approach 400 SPECint92 in 12 to 18 months.

A detailed manufacturing and capacity analysis gives us reason to believe that Intel is bringing enough fab capacity online to support a production ramp for the P6 that could exceed that of the rapidly ramping Pentium.

This basically is a very fast scalable design. It is intended for the mainstream desktop and will be a mainstream notebook processor in 1997. Intel has the manufacturing muscle in place, and the closest significant competitor is still a generation behind.

Dataquest Perspective

As always, memories played an important role at the ISSCC, and we expect them to continue to do so. The job of the Dataquest analysts, then, is to sort out which ones should enjoy success, and to try to determine when and how their success will be attained. This will figure into any long-term prognostications we will make in future publications.

All in all, it was a fascinating show. There were lots of stunning advancements, like the mere fact that the 1Gb DRAM was introduced two—not three—years after the first paper heralding a 256Mb DRAM, or NexGen's use of IBM's five-layer metallization, along with C4 bonding and more precise layout tools, to reduce its die size by more than 40 percent without reducing process geometries.

For More Information...

Jerry J. Banks, Director/Principal Analyst ..... (408) 468-8677
Internet address ..... jbanks@dataquest.com
Via fax ..... (408) 954-1780

Perspective

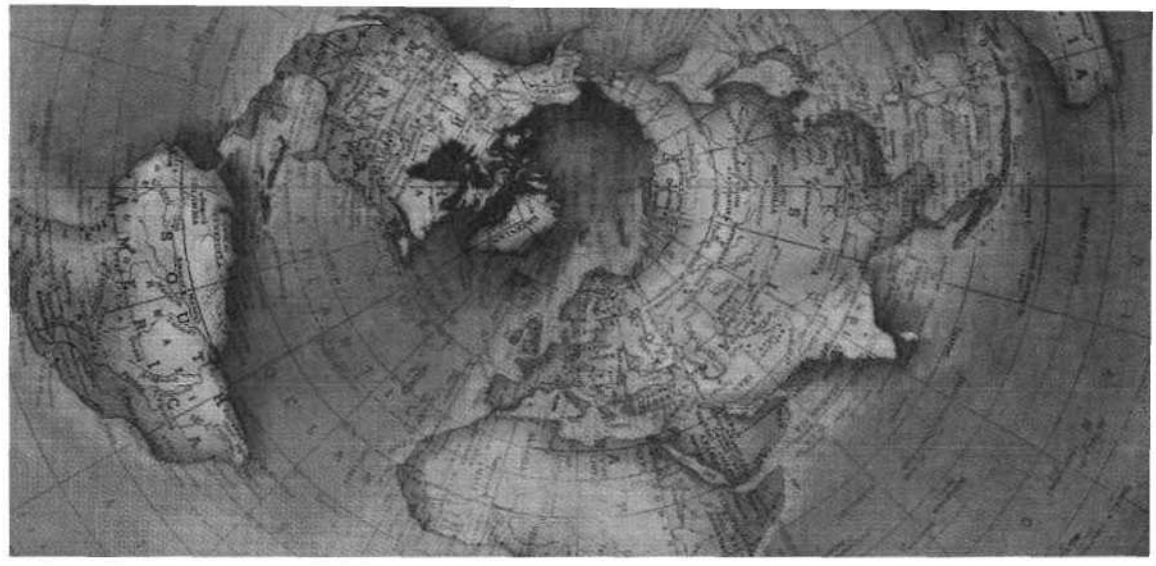

Semiconductor Directions in PCs

Research Brief

1994 PC Core Logic Market Share Leadership—Uneasy Lies the Head That Wears the Crown

Abstract: *In this Research Brief, Dataquest lists the top 10 vendors in the market for PC core logic chipsets. We also analyze the various competitors and provide our expectations for success in 1995.*
By Jerry Banks

Still No. 1: VLSI Technology

A familiar name is at the top of the list of PC core logic market leaders, according to a recent Dataquest survey. VLSI Technology is once again No. 1 in terms of overall PC core logic unit shipments (see Table 1). In 1994 VLSI held off all challengers vying for the No. 1 position in total PC core logic chips. Although the company's market share fell from 21 percent in 1993 to 17.9 percent in 1994, VLSI was still able to hold off the hard charge of Taiwan-based Silicon Integrated Systems (SIS). In 1994 SIS grew at a pace nearly double that of the PC core logic industry, leaping over both United Microelectronics (UMC) and Opti. At its current pace, SIS should overtake VLSI in 1995 in terms of total PC core logic chipsets shipped. Both UMC and Opti grew in the 20 percent range and were able to hold off the challenge of Intel, whose meteoric rise may propel it above both SIS and VLSI into the No. 1 position before 1995 is finished. The only company to outperform Intel in terms of growth was the Cirrus Logic subsidiary, Pico Power. Its 469 percent growth rate, based entirely upon the company's notebook product offerings, propelled the company from the No. 10 position to No. 8.

Dataquest

Program: Semiconductor Directions in PCs
Product Code: PSAM-WW-RB-9501
Publication Date: May 29, 1995
Filing: Perspective

Table 1
Top 10 Core Logic Chipset Producers, Worldwide (Thousands of Units)

1994 Rank	1993 Rank	Company	1993 Total	1994 Total	Percentage Change	Market Share (%)
1	1	VLSI	7,392	9,047	22.4	17.9
2	4	Silicon Integrated Systems	4,652	8,600	84.9	17.0
3	2	UMC	6,800	8,200	20.6	16.2
4	3	Opti	6,344	7,775	22.6	15.4
5	7	Intel	1,500	5,300	253.3	10.5
6	5	Acer	3,900	4,200	7.7	8.3
7	6	ACC Microelectronics	1,983	2,563	29.2	5.1
8	10	Cirrus/Pico Power	325	1,850	469.2	3.7
9	8	Chips & Technologies	1,400	1,700	21.4	3.4
10	9	Symphony	410	860	109.8	1.7
		Others	800	535	-33.1	1.1
		Total Shipments	35,506	50,630	42.6	

Source: Dataquest (May 1995)

Growth Adjusted

The 43 percent market growth is skewed somewhat by the rapid transition of PC OEMs, Compaq and IBM, switching from the use of custom PC core logic to standard, off-the-shelf PC core logic. As was estimated in an earlier report from Dataquest's Semiconductor Directions in PCs program, the use of custom PC core logic in desktop PCs fell from 15 percent in 1993 to 9 percent in 1994 and is expected to fall to less than 5 percent in 1995, while the use of custom PC core logic in notebook PCs fell from 41 percent in 1993 to 32 percent in 1994 and is expected to approach 15 percent in 1995. We expect this trend to continue into the future because major OEMs are rapidly disbanding, or significantly reducing, internal PC core logic R&D efforts. Because Dataquest does not include the captive PC core logic shipments in its market share calculations, the growth of the overall PC core logic market is overstated by about 6.3 percent for 1994. Adjusting for this, the overall market for PC core logic grew by 37 percent. The 43 percent growth figure reflects the average growth experienced by the merchant market suppliers that were able to replace the custom products previously used, in addition to the growth of the PC and motherboard upgrade markets.

Dataquest Perspective

PC core logic chipsets are as critical to the overall performance of the PC as any other component in the system. The PC core logic chips are the interface between the MPU complex (microprocessor and cache memory) and the rest of the PC. These chips manage the system resources that make up the PC. Many of the critical performance bottlenecks—such as the cache algorithms, DRAM refresh, memory management, power management, expansion bus control, local bus control, and DMA control—are handled directly by the PC core logic chips. A poorly designed chipset architecture will severely cripple a PC's performance, completely independent of the microprocessor used in the system.

The PC core logic market is undergoing significant changes. In order to compete in this arena it is essential to understand where the market has been, how it has changed, the dynamics affecting that change, and where the PC core logic market is going. In an upcoming publication, Dataquest will provide further details on this dynamic market. This report will start with a historical backdrop featuring the players, the product features, desktop versus notebook feature requirements, market size, and market share. It then will work forward to the present, where it discusses the issues changing the market, and finally will take a look at the future in terms of market size and feature set, and then give the foundation of what it will take to win in this highly competitive market.

The dynamics of this market are such that, if one of today's leaders loses track of the market for the briefest of moments, the entire makeup of the leader board shown in Table 1 will change. It is essential to track this market closely and measure competitors' progress and determine which of the players is best positioned for the future and why.

For More Information...

Jerry Banks, Director/Principal Analyst ..... (408) 468-8677
Internet address ..... jbanks@dataquest.com
Via fax ..... (408) 954-1780

Perspective

Semiconductor Directions in PCs

Dataquest Predicts

Future Trends for DRAM Interface in PCs

Abstract: *New types of DRAM are necessary to alleviate performance bottlenecks associated with memory in PCs. Current techniques for increasing bandwidth create new problems and increase system cost by requiring larger memory increments. This document describes a number of current and future types of DRAM and explains the advantages and disadvantages of using these memory technologies in PCs. Also, a discussion of memory granularity (what it is and why it is important) is included. PC OEMs are concerned about availability and multivendor sourcing, so the ultimate success of any one of these memory technologies depends upon general logistics as well as price and performance. All of these issues are considered for Dataquest's road map of DRAM I/O structures in PCs through the year 2000.*
By Geoff Ballew

PC Designers Need Faster DRAM I/O

Personal computer designers are demanding greater DRAM bandwidth to keep pace with performance increases for microprocessors and graphics controllers. Memory bandwidth is simply the capacity to move data in and out of the memory bank. It is often measured in megabyte per second, so a higher number indicates better performance. PC designers want the improved bandwidth for as little additional cost as possible. Increased bandwidth is important for both the main memory and the graphics memory because these areas are becoming performance bottlenecks. CPU clock speeds continue to rise dramatically, so main memory needs to deliver both data and instructions at higher rates. Also, graphics standards are moving to higher refresh rates as well as more demanding types of information, such as motion video and 3-D graphics. These trends create a widespread need for greater memory bandwidth.

Dataquest

Program: Semiconductor Directions in PCs
Product Code: PSAM-WW-PD-9502
Publication Date: August 28, 1995
Filing: Perspective
(For Cross-Industry, file in the Semiconductors Volume 3 of 3 binder behind the **Semiconductor Directions in PCs** name)

One technique already employed is simply using a wider memory bus. Pentiums are 32-bit microprocessors, but have a 64-bit memory bus. Many graphics controllers also have a 64-bit bus. This technique creates additional problems because of memory granularity. For a thorough discussion of memory granularity, refer to the last section of this document. This section reprints an October 31, 1994, article on memory granularity by Dataquest Principal Analyst Jim Handy.

The solution is to use new and faster DRAM architectures. There are many contenders at this time, but the field will narrow quickly as the industry chooses the next standard. Figure 1 shows a road map for the acceptance of new types of DRAM for different applications.

Figure 1
DRAM Architecture Trends in PCs

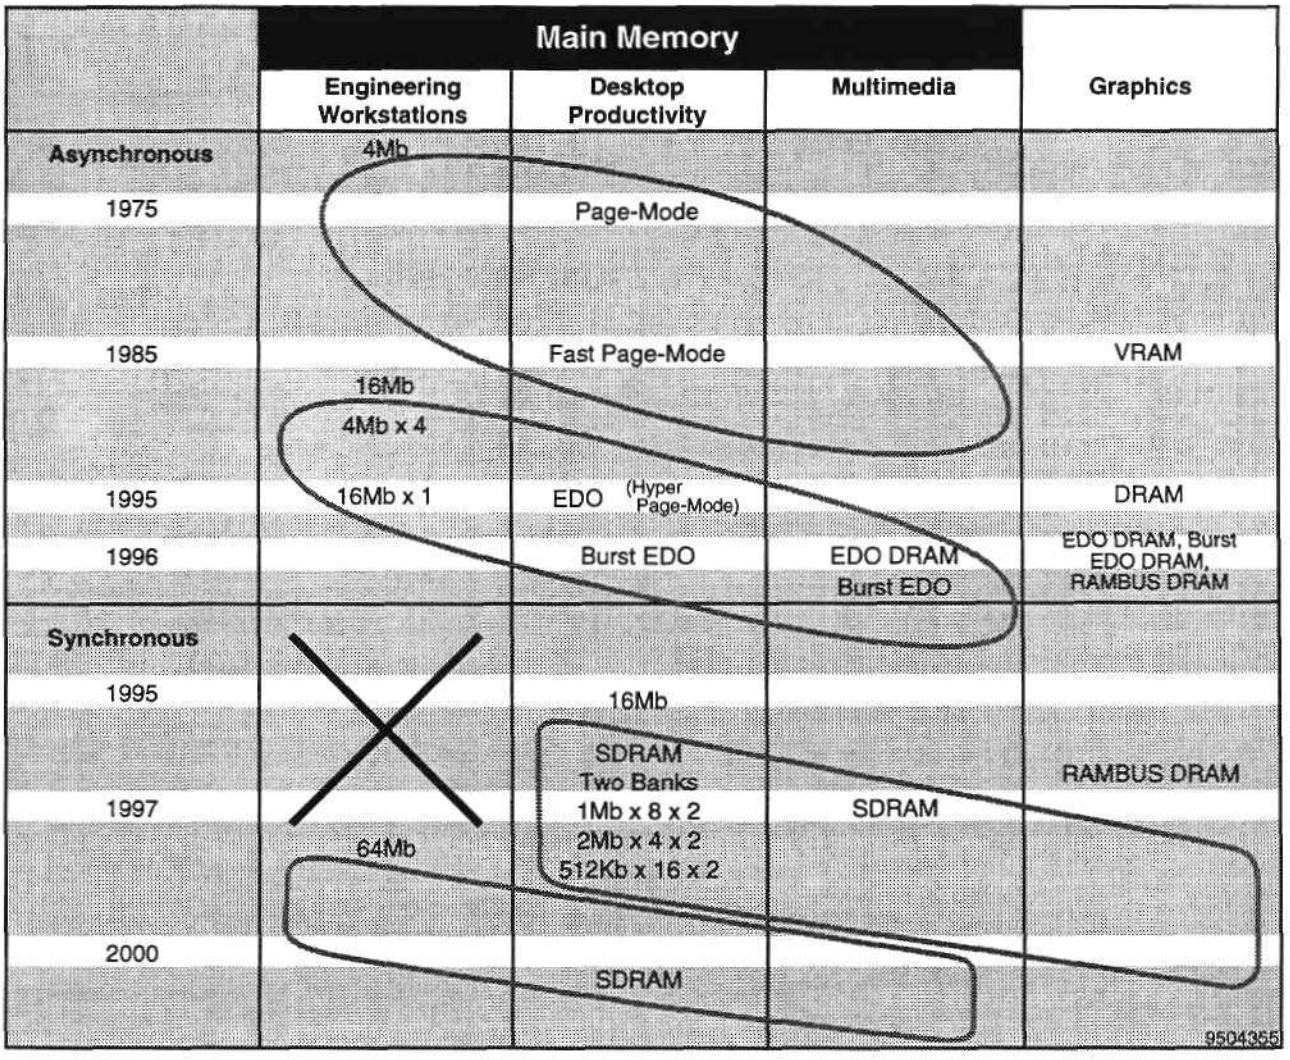

Source: Dataquest (August 1995)

The Contenders

Fast Page-Mode DRAM

Fast page-mode DRAM is currently the mainstream choice for DRAM in PCs, but its popularity is forecast to fade over the next two years. It is an asynchronous design without any burst features. Fast page-mode memory simply does not provide the bandwidth that designers require. Dataquest expects fast page-mode memory to fade quickly as the transition from 4Mb devices to 16Mb devices continues.

VRAM

VRAM is a type of dual-ported DRAM. Its main application has been graphics memory because of the need for the graphics controller to update the frame buffer without interrupting the screen refresh. VRAM sells for a premium of roughly 60 percent over standard DRAM and is therefore cost-prohibitive. Standard DRAM replaced VRAM in the low- and midrange graphics applications in 1993 because of the cost differential.

High-performance applications are still using VRAM because of the dual-port feature. VRAM has one parallel port and one serial port, but the parallel port cannot access the memory core while the serial buffer is being reloaded. Pixel data for an entire scan line is moved to the serial port's line buffer with a single request. Once that request is fulfilled, the memory core is accessible through the parallel port until the next line of pixel data is needed. This makes the memory core accessible to the graphics controller a higher percentage of the time than is possible with standard DRAM. The result is more time for the graphics controller to update the frame buffer without disrupting the screen refresh.

EDO DRAM

Extended data out DRAM (EDO DRAM) is already shipping from a number of suppliers and may account for as much as 20 percent of the 1Mbx4 devices in 1995. EDO DRAM is also referred to as hyper-page-mode DRAM and is an asynchronous memory architecture. Dataquest expects EDO DRAM to replace fast page-mode DRAM over the next two years, with almost all production capacity converted from fast page-mode to EDO by 1997. EDO DRAM will be particularly strong at the 16Mb device level, but is likely to lose some ground to synchronous DRAM (SDRAM) in 1998 and be almost completely replaced by the year 2000. Dataquest does not expect EDO DRAM to compete effectively at the 64Mb device level. Bus width requirements combined with higher bit densities will favor synchronous memory devices.

Burst EDO DRAM

Burst EDO DRAM is a variant of EDO DRAM championed by Micron Technology. Dataquest believes that if burst EDO DRAM is accepted by the marketplace it could delay the acceptance of synchronous DRAM by alleviating the current bandwidth problem. Additional suppliers are critical to the success of burst EDO DRAM because OEMs may hesitate to commit to a single-source DRAM technology. As a long-term solution, burst EDO DRAM does not address the bus width issue, so it is unlikely to compete effectively at the 64Mb level.

Synchronous DRAM

Dataquest believes that synchronous DRAM (SDRAM) will be the long-term winner in the DRAM market. SDRAM's acceptance has started at the 16Mb level, but shipments will be a small part of the market until 1997 and will then grow quickly. Figure 2 shows this trend. Expect SDRAMs to ship in dual-bank packages in the near term with configurations such as 1Mbx8x2, 2Mbx4x2 and 512Kbx16x2. The dual-bank packages allow data to be interleaved between the two banks. Dataquest believes that SDRAM will be the standard architecture at the 64Mb device level because of bandwidth and bus width issues. Synchronous DRAM will allow designers to use a much narrower interface because the burst characteristics of SDRAM can be used to emulate a wide bus. The minimum amount of memory a system requires will be reduced by moving back to narrower buses because a single DRAM could fill the entire bus. For example, a single x16 DRAM would meet the minimum requirements for a 16-bit bus, compared to four devices for today's 64-bit Pentium memory bus. OEMs could configure their PCs with as much or as little memory (down to a single device) as consumers are willing to pay for.

Figure 2
DRAM Unit Volume by Interface Type

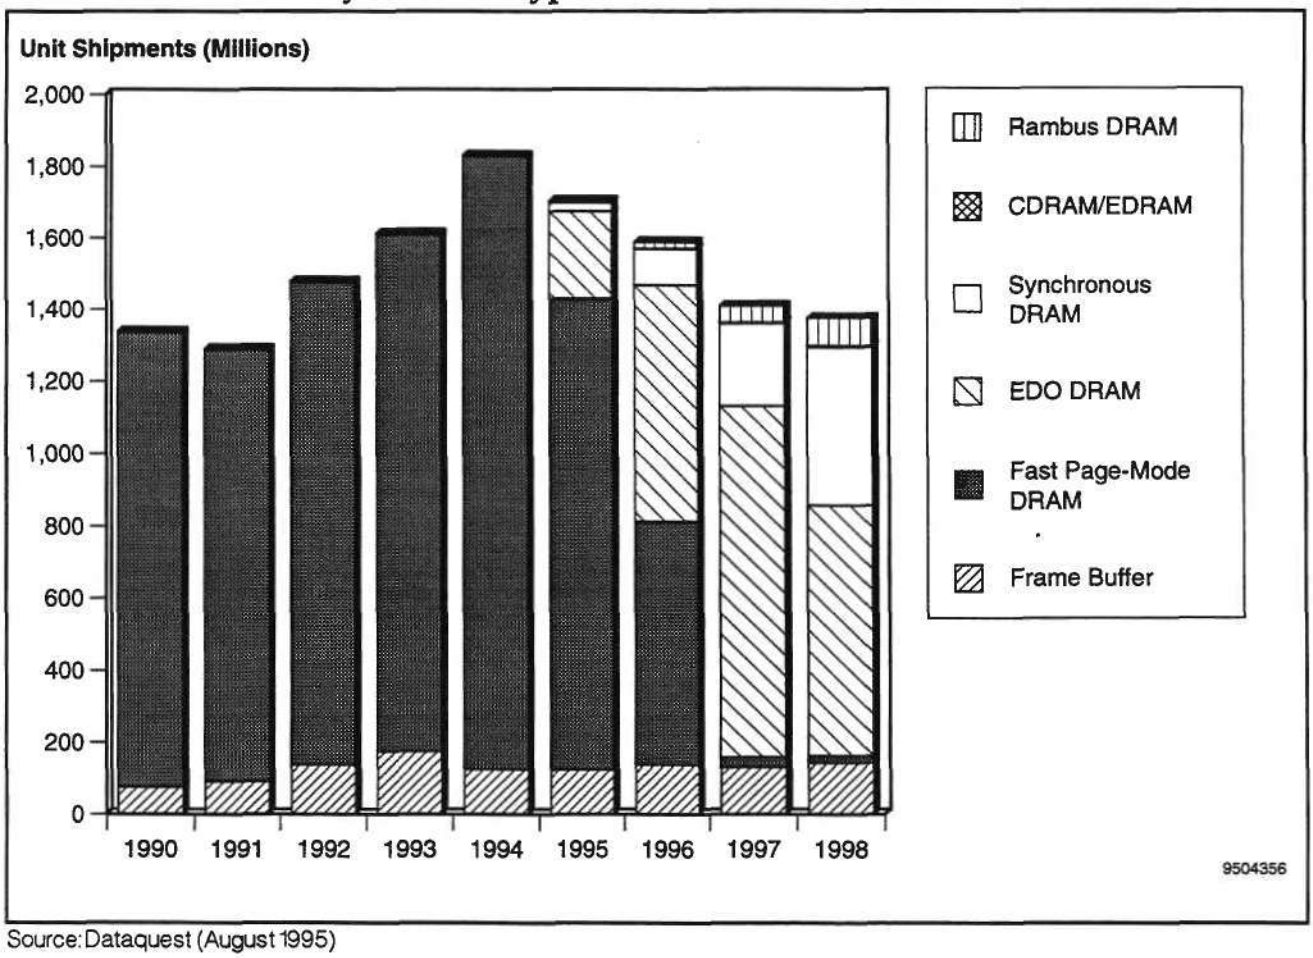

Source: Dataquest (August 1995)

An additional benefit from using SDRAM for PC main memories is the opportunity to eliminate the secondary cache SRAM. Some low-end PCs today are designed without secondary cache memory, which reduces their overall performance because of the relatively slow access times for asynchronous DRAM. Synchronous DRAM interfaces can accept addresses at the rate of one or two per clock cycle, which means faster access to a random address when compared to asynchronous memory. These faster access times may be enough of an improvement that cache SRAM will not be attractive for low-end PCs on a price-per-performance basis. Dataquest does not expect secondary caches to disappear, though. Midrange and higher-performance PCs will most likely continue to have caches, as well as some low-end PCs.

Rambus DRAM

Rambus DRAM (RDRAM) is a type of synchronous DRAM. Rambus has design-wins already from Nintendo, Silicon Graphics, and Cirrus Logic, but many designers have adopted a wait-and-see attitude toward RDRAM. Additional design-wins in the near term would greatly boost designer confidence and might speed the acceptance of other types of synchronous DRAM. Dataquest believes that RDRAM's biggest competitor in the next few years to be JEDEC-standard SDRAM.

What's All This "Granularity" Stuff, Anyhow?

The term "granularity" is a new one to DRAM marketers, and its meaning can be a little difficult to grasp. This section will try, in the simplest possible terms, to explain the term in such a way that it not only will make sense, but will also stay with the reader.

Bits versus Bytes

First of all, let us look at two different conventions: the one used for chips versus the one used for computers. Chip suppliers and users talk about the number of bits in a memory chip. The current hot chip is the 4-megabit (Mb) density, and the upcoming generation is the 16Mb density.

System-level memory users measure memory size in megabytes (MB). SIMMs (single in-line memory modules) provide a way to add a lot of this memory to systems. In the recent past, SIMMs came in 1MB sizes. The current standard is the 4MB SIMM. Today, the base level of main memory shipped in most PCs is 4MB.

Converting from bits to bytes is easy: a byte is eight bits. "Byte" is just a slurred way of saying "by eight," the standard width for main memory systems 25 years ago. To derive bytes from bits, just divide the bits by eight to get bytes, then use an uppercase "B."

Table 1 shows some current densities of DRAMs (in bits) and the corresponding number of bytes held by each chip.

Now that we have a way of translating the chip's density to the amount it adds to the system, let us see the number of DRAM chips required to make a system or a SIMM. Table 2 shows how many chips it takes for the densities shown in Table 1 to make either a 1MB or a 4MB SIMM or main memory.

Table 1
Converting DRAM Sizes from Bits to Bytes

Bits	Bytes
1Mb	128KB
4Mb	512KB (half megabyte)
16Mb	2MB

Source: Dataquest (October 1994)

Table 2
Chips Required to Build Certain Size Main Memories

Memory Size	Number of 1Mb Chips	Number of 4Mb Chips	Number of 16 Mb Chips
1MB	8	2	*
2MB	16	4	1
4MB	32	8	2
8MB	64	16	4
16MB	128	32	8
Bytes per Chip	128KB	512KB	2MB

*A 1MB memory cannot be built out of a 16Mb (2MB) chip.
Source: Dataquest (October 1994)

Processor Bus Widths

Now let us change to a completely different side of the argument. Different processors have different data bus widths. Let us avoid going into a technical deep end. Just remember that the number of wires carrying data into and out of the processor must match the number of wires running into and out of the DRAMs. This does not mean every DRAM, but rather that the number of wires coming from the DRAM array must match the number from the CPU.

Table 3 shows a long history: the number of input-output (I/O) pins on the data bus of x86 processors.

Table 3
Data Bus Widths of x86 Processors

Model	Data Bus Width (Bits)
8086	16
8088	8
80286	16
80386DX	32
80386SX	16
486 (All)	32
Pentium	64

Source: Dataquest (October 1994)

We will limit the rest of the article to the 486 and Pentium processors, with their 32-bit and 64-bit data buses. The reason that bus widths have been gradually increasing is that more data can be processed if more bytes can be gotten into and out of the processor. Doubling the bus width doubles the traffic through this potential bottleneck. This is a somewhat expensive way to increase the throughput of the processor, so designers move from one width to another with some hesitation.

Table 2 shows that a 4MB main memory can be built of eight half-megabyte (512KB or 4Mb) chips. For the 486 we will need 32 data bus pins, so these data or I/O pins will be spread out among the eight DRAM chips. The DRAM we will want to use has four I/O pins, since 4 x 8 = 32. This fits well with the most popular current DRAM chip, the 1Mx4 DRAM. For a 4MB main memory on a Pentium system we still need eight 4Mb DRAM chips, but to satisfy a 64-bit bus, these chips must have 64/8 = 8 I/O pins. We will explain the difficulty in this approach shortly, but for the meantime make a mental note to look for advertisements for Pentium systems offering 4MB main memories. It is much easier to find memories starting at twice that size, and these memories can be built out of 16 1Mx4 DRAM chips (16 x 4 = 64).

DRAM Organizations

We divert briefly for a history lesson. All DRAMs are presented in multiple organizations. For example, the 1Mb density shipped either as a 1Mx1 or as a 256Kx4. This means that the 1Mx1 version had a single data bit available at a time (on a wire to the outside world). By contrast, the 256Kx4 version had four bits available. Because both contained the same amount of bits, the four-bit version (256Kx4) had one quarter as many locations for locating or addressing four bits within the chip. In a similar vein, the 4Mb density comes today as a 4Mx1 or a 1Mx4, with limited shipments being made in the 512Kx8 and 256Kx16 organizations.

The 16Mb chip is available in limited volume as a 16Mx1, a 4Mx4, a 2Mx8, and a 1Mx16. More data bits in the memory chip (the number after the "x") mean a higher pin count and a bigger device package. Figure 3 illustrates the organization of key 16Mb DRAMs.

Ever since the development of the first 1Kb DRAMs, there have only been two widths offered: x1 and x4. Nobody ever asked for wider DRAMs. Even for the current generation of the 4Mb density, the highest-volume parts, and the parts first developed, were the 1Mx4 and the 4Mx1 organizations. DRAM manufacturers have always introduced only these two products, and testing and I/O design have never been challenging since DRAM bus widths never changed. Suppliers introduced the first 16Mb density DRAMs in the same way as their predecessors—in the 16Mx1 and 4Mx4 organizations. If suppliers did any differently, they would encounter new and difficult problems. Potential problems range from power dissipation through noise, speed, die size, and pin count to tester throughput.

The Problem

Now back to the number of chips used to build a system. Remember that a 4MB main memory in a Pentium system would require the use of eight 4Mb DRAM chips? Well, this is not at all a commodity part, and although the

Figure 3
Organizations of 16Mb DRAMs

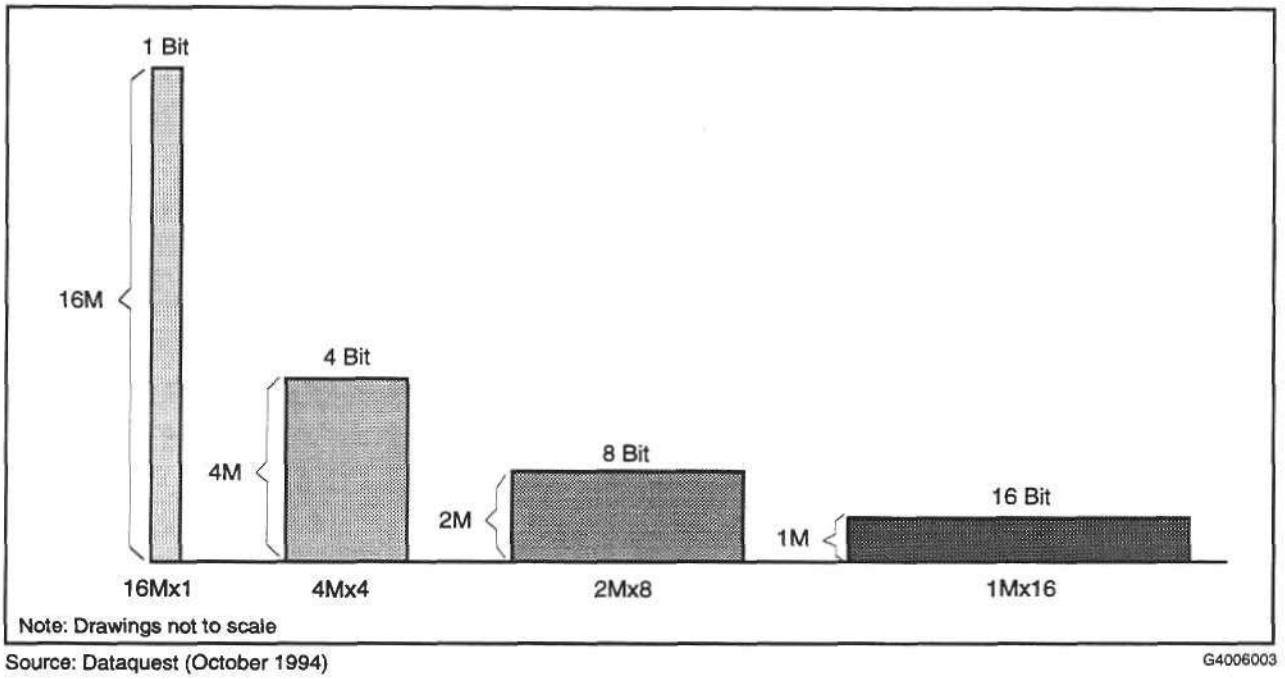

Source: Dataquest (October 1994) G4006003

device does exist (its organization is 512Kx8), it is hard to find and commands a price premium.

To make matters worse, let us move on to the 2MB DRAM chip, more conventionally called the 16Mb density. Just two 16Mb devices (properly configured) make a 4MB main memory. So on the 486 system, with its 32-bit bus, the DRAM chips required each need to be 16 bits wide. These chips are the ones that promise to be the most popular in the near term. Unfortunately, their current availability is slim, so heavy use of the 1Mx4 should continue through next year.

Figure 4 illustrates how a 1Mx32 SIMM (based on eight 1Mx4 DRAMs) is connected to a 32-bit data bus like the 486 processor.

Once again, the Pentium poses a new brand of problem. The Pentium's 64-bit bus needs a split between two chips in a 4MB system built of 2MB chips. Each chip would need to be 32 bits wide, which is not a device that anybody has on the drawing boards yet. Of course, 16Mb chips of half that width, the 1Mx16, easily build into 8MB main memory for Pentium systems.

Granularity

So what is the definition of the term "granularity"? Granularity refers to the incremental size of main memory system users want to add to their system. This also means the minimum configuration with which a system ships. These two meanings are the same. Today's granule is 4MB. Some argue that this is set by price. The current cost for 4MB is about $100.

Figure 4
Connecting 1Mx4 DRAMs to a 32-Bit Bus

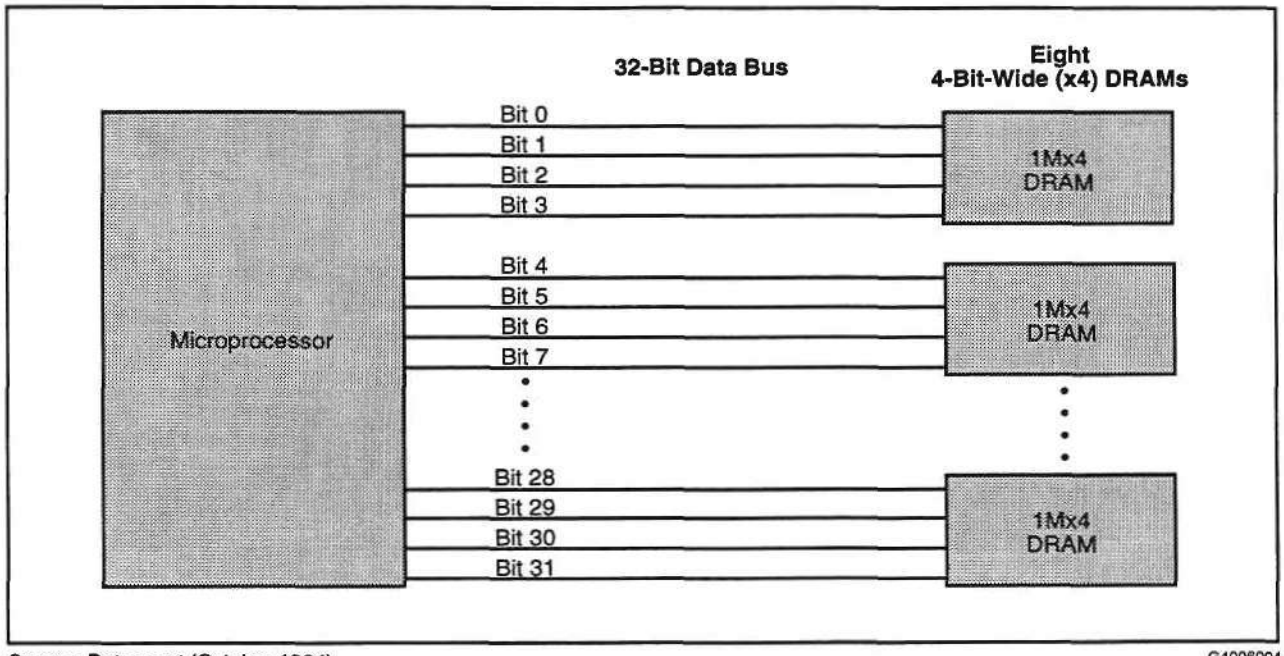

Source: Dataquest (October 1994)

G4006004

This argument works toward the granule being the amount of memory that $100 buys. Whenever 8MB starts to cost this much, this camp will argue that the popular granule size will be 8MB. This will cause the 1Mx16 version of the 2MB DRAM to lose market to the 2Mx8 version. Dataquest anticipates such a shift to occur in 1996.

So far all of the arguments center around the PC. What about other forms of computing systems? The reason for the PC focus is that PCs account for a growing share of the DRAMs consumed. The portion exceeds 60 percent of DRAM consumption. Where the PC market goes, the DRAM market follows.

Still, there are applications that use the x4 and x1 versions of the 2MB DRAM. Most workstations and file servers still use 32-bit processors, and if we turn our math around a bit, we can see that these chips can build reasonably sized main memories for such systems. A 32-bit bus built of 4Mx4 2MB DRAMs would need to use eight chips, making the main memory size eight times 2MB, or 16MB. This is a good granule for a workstation. A 32-bit bus built of 16Mx1 2MB DRAMs would need to use 32 chips, making the main memory size 32 times 2MB, or 64MB. This immediately puts us into the realm of heavyweight computers like superworkstations and midrange computers. You can see that the market becomes more and more limited the narrower the device becomes.

Dataquest Perspective

Memory bandwidth is a performance bottleneck in today's PCs. The current technique to improve bandwidth has been simply to use a wider bus, but this creates additional problems by increasing the minimum amount of memory a system requires, as well as increasing the upgrade increment. These issues are tied to memory granularity, as discussed above. New types of DRAM are necessary to alleviate the bandwidth issue without requiring wider and wider buses. Dataquest expects the solution to come in two stages.

Stage one is a shift from fast page-mode DRAM to EDO DRAM and has already started. This stage only partly addresses the bandwidth issue, because EDO DRAM is an asynchronous design. EDO DRAM will most likely be the largest-selling type of DRAM at the 16Mb level, but will lose popularity at the 64Mb level because of the difficulty in manufacturing memory devices with wide interfaces. Wide buses and memory granularity issues will continue to be a problem until stage two.

Stage two is a shift to synchronous DRAM or variations of SDRAM like Rambus DRAM. Dataquest expects stage two to begin in 1997, with the larger share of the DRAM market shifting to synchronous DRAM by the year 2000. Synchronous DRAM is expected to be the dominant DRAM architecture at the 64Mb device level.

The timing and specific directions of these shifts may be more dependent upon price and multiple vendor sourcing than upon strict device technology. PC OEMs need improved performance, but are working with narrow margins and high DRAM prices right now. Semiconductor vendors need to assure PC OEMs that parts shortages or single-source pricing will not be issues. Commodity pricing and multivendor sourcing will speed acceptance of new DRAM architectures.

For More Information...

Geoff Ballew, Industry Analyst ..... (408) 437-8676
Internet address ..... gballew@dataquest.com
Via fax ..... (408) 437-0292

Perspective

PC Semiconductor Applications Worldwide

Dataquest Predicts

The PC Processor: What's Hot, What's Not

Abstract: *This report discusses the hottest issues taking place in the PC microprocessor market. Included is an analysis of the quarterly shipment history and an eight-quarter forecast for the 486, Pentium, and P6 microprocessors. We discuss the major players as well as the latest status on their various PC microprocessor product offerings. This report also analyzes the ever-shrinking product life cycles of the x86 microprocessors and the ever-shrinking performance gap between the RISC and CISC microprocessor architectures.*
By Jerry Banks

Setting the Stage—1994: The Year of the 486, the Quarter of the Pentium

Although 1994 saw record shipments of 486 microprocessors, the phenomenal acceptance of the Pentium in the fourth quarter of 1994 provides an indication that the 486 has already seen its peak. More than 43 million units of 486 microprocessors were shipped, with the 486DX2 family making up the bulk of those shipments (see Figure 1).With the advent of the clock-multiplier technology for x86 microprocessors, single-clock variations of the 486 microprocessor began a rapid falloff beginning in the first quarter of 1994 (see Figure 1). This rapid decline continued throughout the year. We expect it to accelerate in 1995.

In conjunction with the 486 unit shipment life cycles, the 486 average selling price has been in a state of decline. The 486DX2 has declined to the point where it sits on top of the single-clocked 486DX, thus all but eliminating the original 486DX as a product category. The 486SX also has fallen off rapidly. Some high-volume purchasers were able to purchase a clock-doubled version of the 486SX for $50 or less by the end of 1994. This is clearly an area where Intel, Advanced Micro Devices (AMD), and Cyrix do not want to participate: Intel because of its need to protect its high margins, AMD because of its limited fab capacity and need to maximize revenue, and

Dataquest

Program: PC Semiconductor Applications Worldwide
Product Code: PSAM-WW-PD-9501
Publication Date: April 24, 1995
Filing: Perspective

Figure 1
486 Series Unit Shipment History

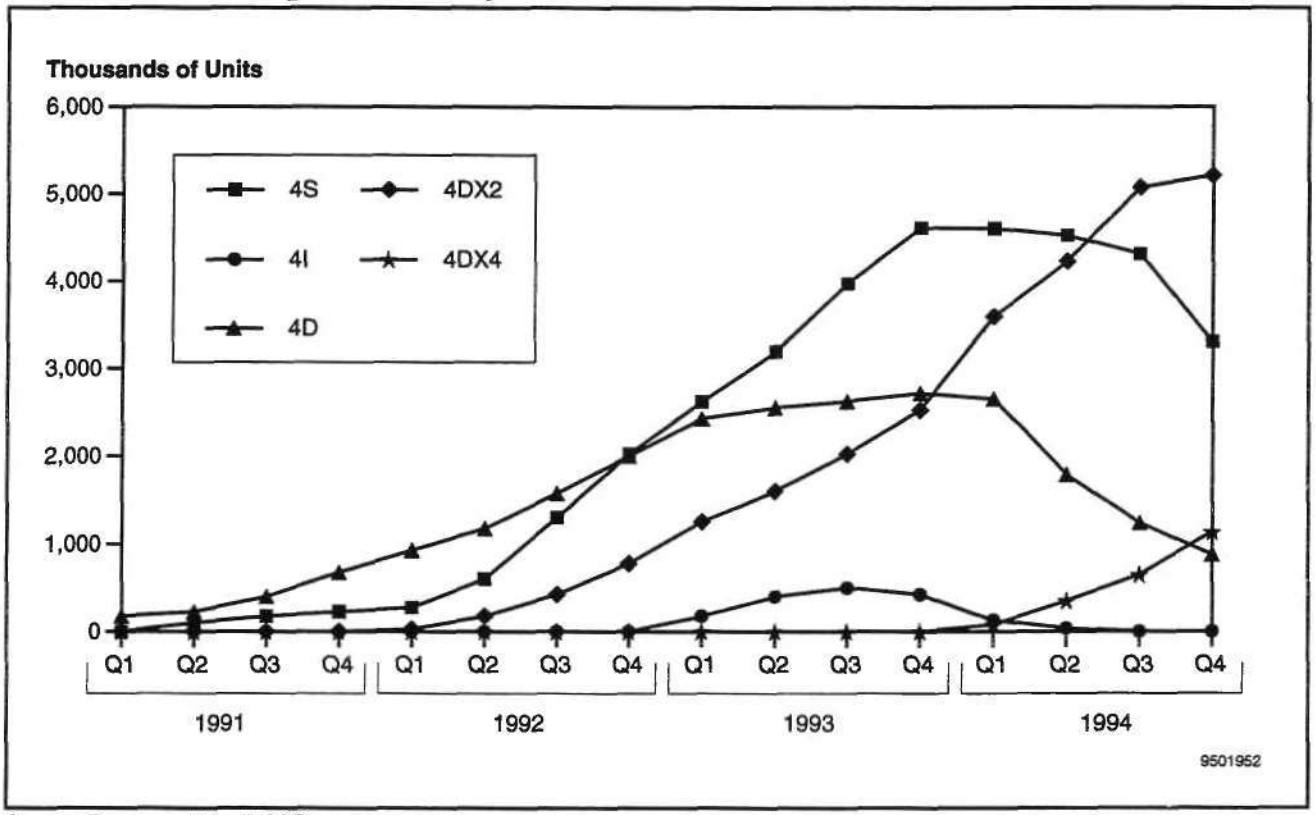

Source: Dataquest (April 1995)

Cyrix because designing and manufacturing small, low-cost die is not its forte. Perhaps only Texas Instruments has the ability and/or desire to continue to drive this low-cost market.

When Is a "Flaw" Not a "Flaw"?

The Pentium processors' success seemed anything but assured as late as December 1994. The now infamous Pentium "flaw" definitely caused a few Pentium users a lot of grief, but more than a few Intel employees worked many hours and lost much sleep during the fourth quarter of 1994 trying to resolve this potentially disastrous issue. The history of how this incident initially was mishandled by Intel has been well documented. But despite Intel's earlier mishandling, it did come clean with its solution and most, if not all, parties are now quite satisfied with Intel's final response. The world also learned that it is worth at least $475 million to ensure the design integrity of a microprocessor. Despite the uproar in the media and on the Internet, Intel still was able to ship 5 million Pentium processors in 1994, with the bulk of that coming in the fourth quarter (see Figure 2). This tremendous acceptance was driven by a user community consumed with the desire for more performance. Few, if any, seemed concerned with all the flap about a flaw that the average layperson could not understand. Perhaps the "Field of Dreams" theory espoused by Intel is true: "Build it and they will come." The Pentium is the most successful microprocessor in history at this stage of its life cycle.

Figure 2
Pentium Shipment History

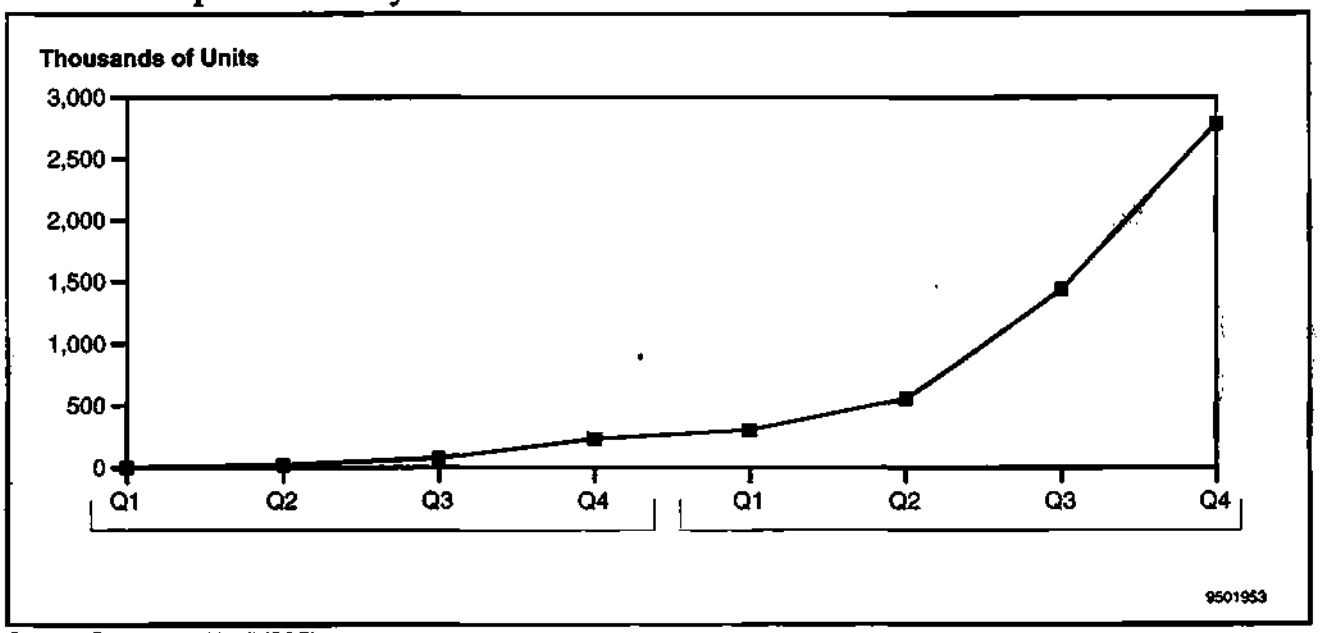

Source: Dataquest (April 1995)

Although the second generation of the Pentium (referred to as the P54C) began to rapidly gain on the first-generation product (also known as P5), the P5 still shipped more than half the total Pentium volume in the year.

Dataquest Predicts

The 586-series microprocessor, led by Intel's Pentium processor, will overtake the 486 as the processor of choice in 1995 and will continue as the processor of choice through 1996. It will ship in record unit volumes during that time frame, with shipments peaking in late 1996.

What Does the Future Hold?

The clock-tripled 486 series (also known as the DX4) is the new growth category within the 486 series microprocessor. This variation of the 486 is providing 100-MHz performance with a 33-MHz bus, thus carrying on the tradition of maintaining the use of low-cost 33-MHz motherboard ICs while milking more performance out of the microprocessor. This allows the MPU manufacturer to increase prices for the added value of 100-MHz performance without impacting the end system cost. All of the necessary system cost reductions are coming out of the other motherboard components that continue down on the 33-MHz learning curve (the non-MPU exception to this is DRAM, which is in the midst of a continuing shortage situation). The only alternative to the motherboard left for the non-MPU and non-DRAM suppliers is to move up in the performance realm of the Pentium and/or integrate more functions in order to add sufficient value to stem the rapid price erosion.

As is seen in Table 1 and Figure 3, only the clock-tripled 486 products are expected to show any growth in 1995. We expect the growth to tail off by late 1995, with a rapid decline continuing into 1996. Although the 486 is in decline, it will still ship nearly 33 million units in 1995 and nearly 8 million units in 1996.

Table 1
486 Quarterly Forecast (Thousands of Units)

	Q1/95	Q2/95	Q3/95	Q4/95	Q1/96	Q2/96	Q3/96	Q4/96
4S	1,537	929	650	493	286	155	69	24
4D	589	505	528	275	49	25	12	0
4D2	4,853	4,167	3,121	2,159	1,116	536	226	201
4DX4	1,990	2,974	4,054	4,088	3,830	3,073	2,254	1,414

Source: Dataquest (April 1995)

Figure 3
486 Quarterly Forecast

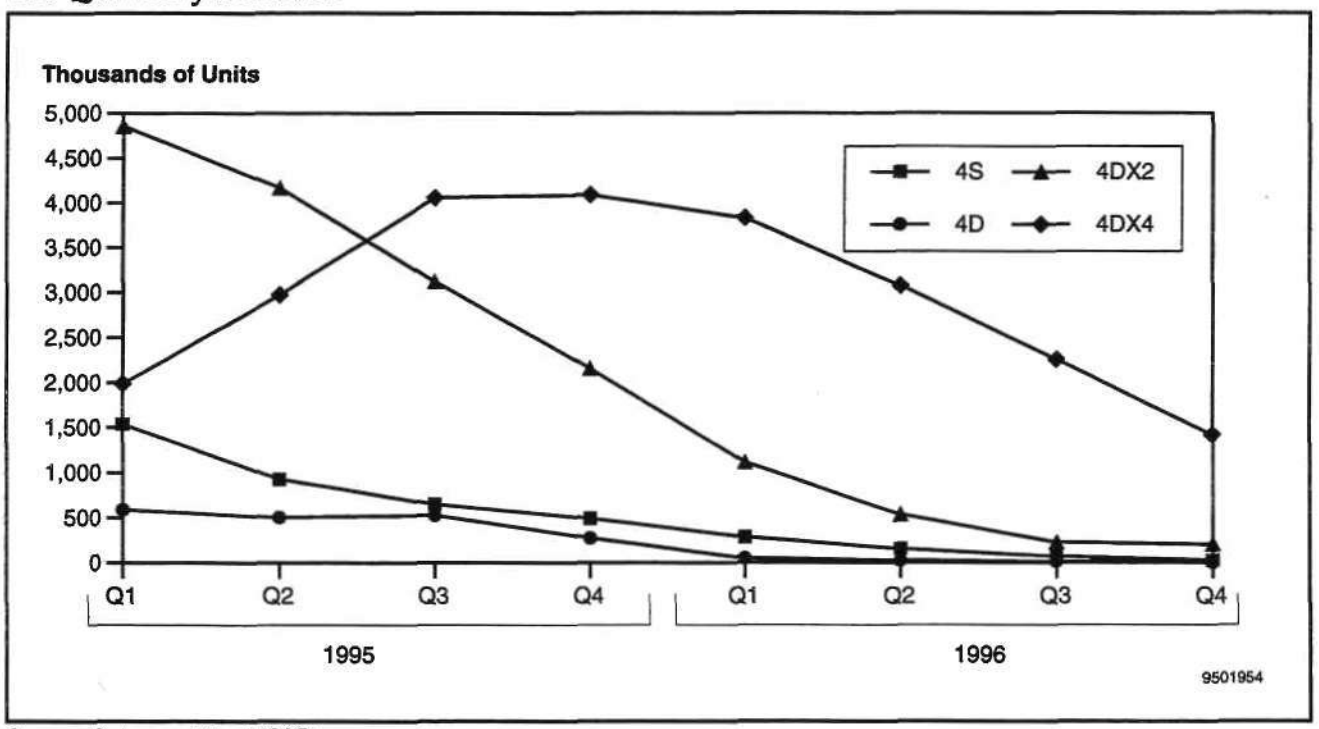

Source: Dataquest (April 1995)

Pentium Futures

The Pentium is on an unprecedented rise. Despite its shaky fourth quarter in 1994 (from a public relations viewpoint), the Pentium seems destined to break all previous shipment records. It was often speculated during the fracas caused by the flaw that the Pentium processor's reputation was permanently damaged and that the production ramp rate for the Pentium would be severely impacted. The latter speculation, as it turns out, is exactly on the mark, although the impact is not quite what the doomsayers (Dataquest was not in this group) were predicting. The Pentium shipped 5 million units in 1994 and is on a run rate to exceed 25 million 1995—what an impact the flaw had on the Pentium's run rate! The 5 million unit shipments in its first full year of production is a world's record, and the run rate it is on for 1995 indicates it will break all previous records in its second full year of production as well. Once again the old adage is proven: The only bad publicity is no publicity. The Pentium is on an unprecedented volume ramp. We are sure that Intel would have preferred that the whole issue had never come up, but one could argue that the current run rates would not

have been possible without the massive manufacturing effort Intel put together at the end of last year. This accelerated effort was organized as part of a recovery plan to get the fixed Pentium into production as fast as possible and develop replacement products for the potential returns of the flawed product. Without this monumental effort, the company would probably still be on its prior plan to ship only in the low-20 million-unit range (which would still have been a world record run rate).

So, Intel once again lands on its feet. Was this luck or planning? Having a problem with your flagship product dragged across all forms of media all over the world is not exactly a lucky thing to have happen. Yet it is also difficult to give Intel a lot of planning credit. What is remarkable is how quickly the company reacted. Nearly the entire microprocessor group was mobilized to support the effort to resolve the problem, and the management team reacted quickly to modify its position until it finally came up with a very acceptable resolution in a relatively short time.

The Pentium Clones Arrive

Although the expected shipment of the upcoming K5 microprocessor from AMD is being touted as the first real competitor to the Pentium processor, the first company to ship a Pentium-class processor outside of Intel is the oft-maligned Nexgen. This eight-year-old company has always had the dream to be the first with the next-generation x86 microprocessor. For nearly its entire existence, the company has been thwarted in its efforts by Intel's product introductions and then by AMD's entry in the market as the only real x86 competitor. Nexgen was constantly resetting itself and looking to the future.

The time for joking apparently has ended. With its announcement of the Nx586 in 1994, Nexgen became the first non-Intel manufacturer of a Pentium-class product. In fact, with the first quarter of 1995 now being finished, Nexgen still holds that title. The company still needs to add a floating-point unit to its current product because the current coprocessor strategy is giving some doubts to the OEM community. The recent announcement of an agreement between Compaq and Nexgen should help Nexgen gain some credibility with more tier-one and tier-two PC OEMs. The next biggest hurdle that Nexgen must overcome is the timely introduction of a 686-series product. It took eight years for this product—can Nexgen match or beat today's current two-year development cycle for next-generation products? Only time will tell, but so far the word is that it is well on its way.

AMD's K5 is coming. Our best estimates are that it will start shipping in the second half of 1995 and will contribute at least 300,000 units to the Pentium series shipments for 1995. The only real obstacle holding up the K5 appears to be manufacturing capacity. The company's long-awaited Fab 25 in Austin, Texas, will give it a much-needed boost. This 8-inch, or 200mm, submicron facility is expected to propel AMD into a much more competitive position. The facility will be able to turn out 486s by the millions and should be sufficient to launch the K5 effort in earnest. The K5 is a more efficient architecture than Intel's Pentium and should provide up to a 30 percent performance edge at similar clock frequencies. This performance advantage does pose a slight marketing challenge for the company. The company historically has introduced an x86 product with a higher clock

rate than a similar Intel product. It would then offer this product at the same price as the lower-frequency Pentium product. This higher-performance, lower-price strategy was quite successfully employed by AMD in the 286-, 386-, and 486-generation products. The problem it faces now is that the earlier-generation products used clock frequency as the measure of performance and the K5 achieves its performance advantage at the same clock frequency as Intel's Pentium. It will be a difficult educational process to convince buyers that the AMD part has superior performance when they have been trained to look for clock frequency as the metric for performance. One option is the route taken by Nexgen. Nexgen claims that its performance operating at 93 MHz is the same as the Pentium operating at 100 MHz. So, it markets its 93-MHz products as Nx586-100. This option is not perceived by all as the most optimum solution. For one thing, the buying public may perceive this as deceptive. For another, using nonstandard clock frequencies poses new challenges for designers. It will be interesting to see how AMD overcomes this marketing problem created by engineering innovation.

The third expected entrant into the Pentium clone wars is Cyrix. The introduction of the much-heralded M1 microprocessor was overshadowed by its monstrous die size. The M1 measures 19.3 x 20.2mm. The only known die larger is the experimental digital signal processor chip from Analog Devices, which contains 4Mb of SRAM and was developed for MIT. The M1's die size is not suited for volume manufacture. The company has launched a massive die reduction effort that should get it into a competitive position when the product is expected to sample in late 1995. In the meantime, Cyrix has scaled down the original M1 architecture to get a product to market sooner. This new product is known as the M1-SC and is not a Pentium-class product. It will be targeted at DX4 sockets. The 100-MHz version is expected to provide the performance of a 75-MHz Pentium. The M1 should start shipping sometime late in the second quarter. Cyrix has stated that it will stop the manufacture of all current-generation 486 products by the fourth quarter of 1995 and devote all of its manufacturing capacity to the M1-SC and M1 when it is available. As more of a marketing statement, Cyrix has reclassified the M1 as a 686-class processor than the 586 or Pentium class at which it was originally announced. Our position is that we will leave the M1 as part of our 586-series forecast.

1995: Pentium Growth Explosion

The widespread user acceptance of Pentium-based PCs is driving the unit volumes for this class of product to record levels (see Figure 4 and Table 2). Intel will represent the lion's share of this volume, but the introduction of 586-series product by AMD, Cyrix, and Nexgen also will help fuel this growth phenomenon. Our forecast calls for 29 million 586-series product to ship in 1995. This explosive growth is expected to continue into 1996, and the only event that can slow the 586-series growth expectation is the production of the next-generation 686 series, which already has been kicked off with Intel's P6 product announcement. For a comprehensive look at the P6, please refer to Dataquest's Focus Report entitled, *P6: A Summary for Systems Manufacturers* (NGPC-WW-FR-9501, published April 10, 1995).

Figure 4
Pentium Forecast

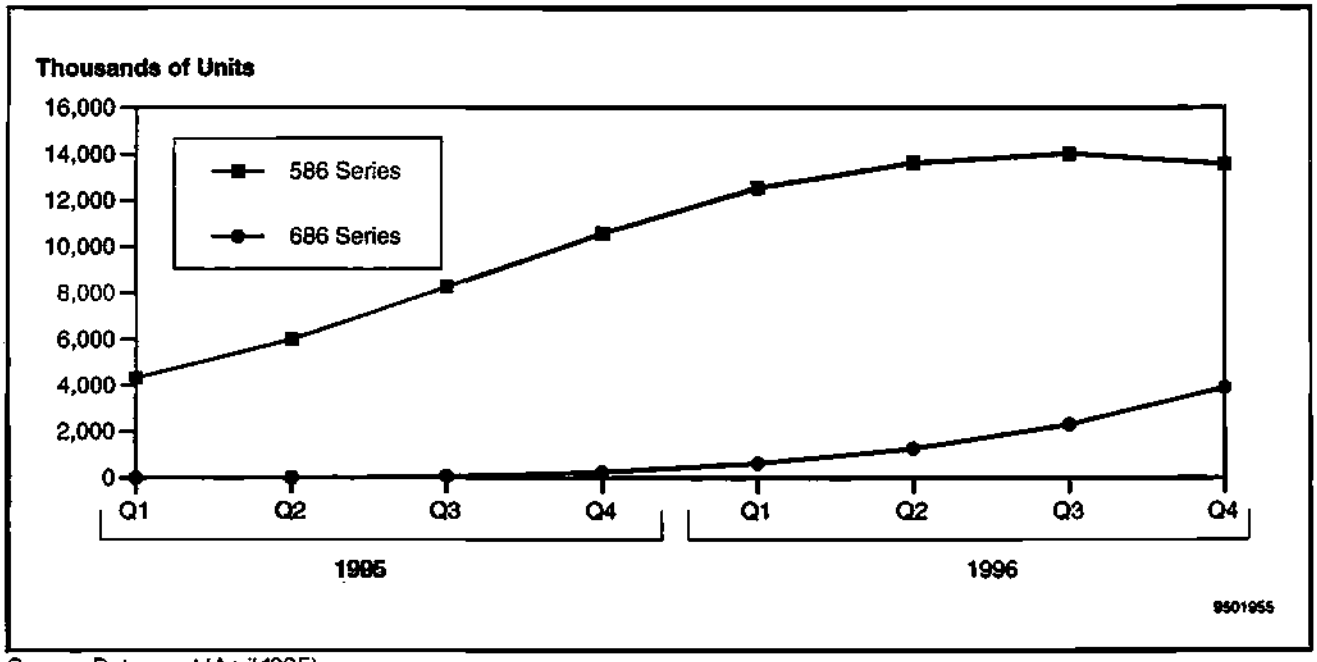

Source: Dataquest (April 1995)

Table 2
Pentium Forecast (Thousands of Units)

	Q1/95	Q2/95	Q3/95	Q4/95	Q1/96	Q2/96	Q3/96	Q4/96
586 Series	4,328	5,986	8,238	10,588	12,564	13,622	14,023	13,560
686 Series	0	0	50	250	625	1,250	2,290	3,885

Source: Dataquest (April 1995)

Uncharted Territory

The introduction of the P6 microprocessor a mere two years after the introduction of the Pentium establishes a new benchmark in microprocessor introductions. The speed of the technology treadmill has doubled. Viewed another way, the time to develop a next-generation microprocessor has been cut in half. This puts tremendous pressure on all microprocessor vendors with an eye on the desktop that Intel now dominates. Most, if not all, of these competitors already had established design flows that would allow them to either maintain a performance lead or surpass Intel in the near future. Intel has just raised the bar, and it is not at all clear whether all the current competitors can match this effort.

The P6 also has introduced new architectural techniques that until now had been relegated to the realm of RISC microprocessors. The P6 makes use of branch prediction, out-of-order execution, and speculative execution, has six execution units, and is superpipelined with a fine-grain strategy that will allow rapid scaling of the processor clock. Traditional thinking would have the current performance gap between RISC microprocessors and CISC microprocessors rapidly increasing over time (see Figure 5).

Figure 5
RISC versus CISC: Traditional View

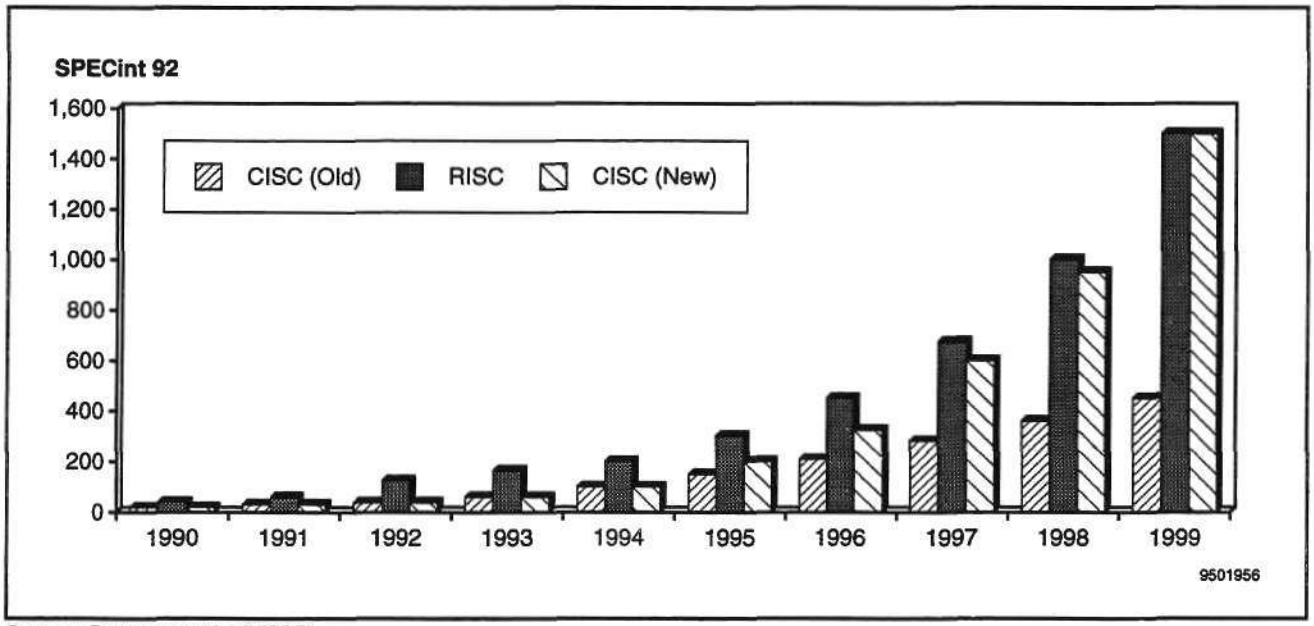

Source: Dataquest (April 1995)

The performance gains of the P6 and future expectations for P7 and the continuing work being performed by the Intel-HP alliance for the P8 give reason to believe that the performance advantage of RISC microprocessors will be short-lived. The first P6 product will be introduced at a clock frequency of 133 MHz and an integer benchmark rating of 200 SPECint92. This clearly places Intel at the top of the x86 performance charts. The P6 also is beginning to make serious inroads into the RISC camp on the basis of integer performance. We should see a P6 integer benchmark of about 300 SPECint'92 by this time next year, and the x86 performance ramp should bring it onto a par with the highest-performance RISC microprocessors by the end of the decade.

New Marketing Challenges

When the first P6 products begin shipping in volume in the third quarter of 1995, it will mark the first year we will see three distinct generations of x86 microprocessors having significant impact on the PC market in the same year. With each distinct generation spawning higher-performance derivatives, price and positioning strategies have to be completely reworked. As is shown in Figure 6, each successive generation of x86 is reaching volume production faster than the prior generation. In addition to reaching volume production status sooner, the production ramp itself is steeper for each succeeding generation.

Marketing teams that were learning how to deal with this shortening of time to production are now seeing the product retirement rate also begin to accelerate. With the two-year development cycle in place, we will see much steeper declines after the microprocessors' peak shipment year, which will result in much shorter life cycles. Managing product price and positioning in this new scenario takes on new challenges. How long can a company extract revenue out of a product without stepping too hard on the prior generation while not slowing the acceptance of the new generation? How

Figure 6
x86 Product Life Cycles

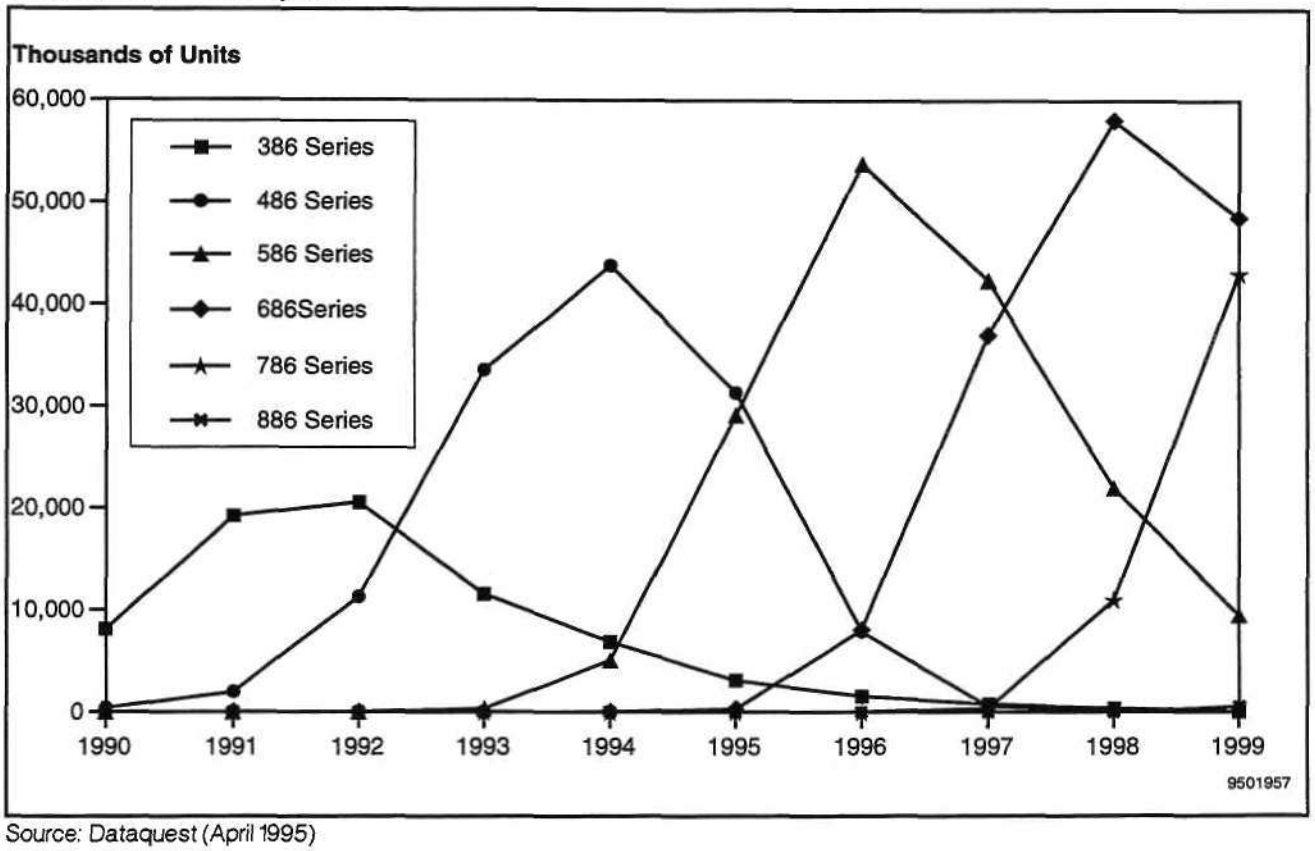

Source: Dataquest (April 1995)

does a company introduce a new-generation product without making the existing revenue stream obsolete? These are questions that have plagued the electronics industry for well more than a decade (how many remember the "Osborne Effect"?) and are becoming more complex. Product marketing has been taking on an increasingly important role as companies strive to maximize revenue without sacrificing the future. This balancing act, which has never been an exact science, is becoming less exact all the time and requires careful management. Product marketeers that understand the market, the competitive forces, and their companies' capabilities will become critical assets to any corporation.

For More Information...

Jerry Banks, Director/Principal Analyst ..... (408) 437-8677
Internet address ..... jbanks@dataquest.com
Via fax ..... (408) 437-0292

Market Analysis

Dataquest

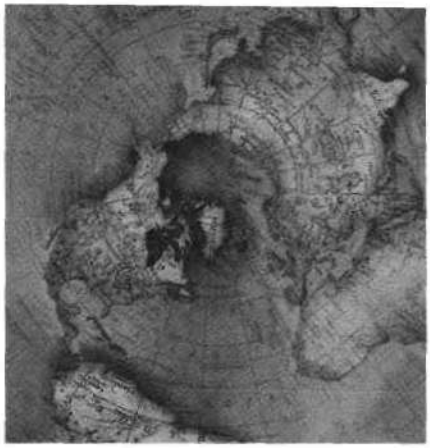

Dataquest

Final 1994 Compute Microprocessor Unit Shipments

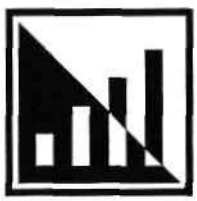

Market Statistics

Program: Semiconductor Directions in PCs
Product Code: PSAM-WW-MS-9502
Publication Date: July 31, 1995
Filing: Market Analysis

Final 1994 Compute Microprocessor Unit Shipments

Market Statistics

Program: Semiconductor Directions in PCs
Product Code: PSAM-WW-MS-9502
Publication Date: July 31, 1995
Filing: Market Analysis

Table of Contents

List of Tables

Note: All tables show estimated data.

Final 1994 Compute Microprocessor Unit Shipments

This document provides microprocessor unit shipment data for all compute applications. Compute functions involve user-changeable software rather than fixed, inaccessible software. Examples of compute applications include PCs, workstations, and home video hardware. Embedded functions involve software that is fixed and defines the functionality of the device. Embedded applications include printers and fax machines.

The data in this report is broken out by CISC and RISC architectures. CISC processors dominate the compute microprocessors market at 94.9 percent, including the x86 family and the 68K family. RISC microprocessors make up 5.1 percent of the total compute microprocessor market. RISC compute microprocessors currently shipping include SPARC, MIPS, PowerPC, PA-RISC, and Alpha.

The word width of a microprocessor is determined by the maximum word width of the software that can be run by the architecture. This is defined by the computed accuracy of a single ADD instruction. If an architecture can perform a 32-bit ADD in a single instruction, then it is a 32-bit architecture. This is independent of the input/output (I/O) bus width or the width of the integer unit. Examples of 32-bit MPUs include 80386, 80486, and 68000.

The data provided in Tables 1 through 7 captures shipments of all processor vendors with sales into the merchant market and excludes shipments of vendors with strictly captive sales.

Table 1
Ranking of Worldwide Shipments of Compute Microprocessors, by Company (Thousands of Units)

1994 Rank	1993 Rank	Company	1993 Units	1994 Units	Percent Change	1994 Market Share (%)
1	1	Intel	28,045	38,525	37.4	64.7
2	2	Advanced Micro Devices	6,790	8,210	20.9	13.8
3	3	Motorola	5,444	4,330	-20.5	7.3
4	5	IBM	947	3,515	271.2	5.9
5	6	Cyrix	900	2,382	164.7	4.0
6	4	Texas Instruments	1,475	1,940	31.5	3.3
7	9	Fujitsu	78	180	130.8	0.3
8	7	Chips & Technologies	126	147	16.7	0.2
9	8	NEC	95	120	26.3	0.2
10	11	Toshiba	61	74	21.3	0.1
11	12	Digital Equipment Corporation	30	65	116.7	0.1
12	10	Integrated Device Technology	64	64	0	0.1
13	NM	UMC	0	25	NA	*
14	13	Weitek	30	10	-66.7	NA
NM	14	Cypress Semiconductor	25	0	-100.0	NA
NM	15	LSI Logic	25	0	-100.0	NA
NM	16	Performance Semiconductor	4	0	-100.0	NA
		Total MPU	44,139	59,587	35.0	100.0

Note: All compute microprocessors are 32-bit and greater.
*Calculated value is less than 0.1 percent
NA = Not available
NM = Not meaningful
Source: Dataquest (July 1995)

Table 2
Worldwide Shipments of Compute Microprocessors—CISC and RISC (Thousands of Units)

Company	1994 Compute MPU	CISC	RISC	Percentage CISC	Percentage RISC
Total Market	59,587	56,557	3,030	94.9	5.1
North American Companies	59,188	56,532	2,656	95.5	4.5
Advanced Micro Devices	8,210	8,210	0	100.0	0
Chips & Technologies	147	147	0	100.0	0
Cyrix	2,382	2,382	0	100.0	0
Digital Equipment Corporation	65	0	65	0	100.0
IBM	3,515	1,515	2,000	43.1	56.9
Integrated Device Technology	64	0	64	0	100.0
Intel	38,525	38,440	85	99.8	0.2
Motorola	4,330	4,298	32	99.3	0.7
Texas Instruments	1,940	1,540	400	79.4	20.6
Weitek	10	0	10	0	100.0
Japanese Companies	374	0	374	0	100.0
Fujitsu	180	0	180	0	100.0
NEC	120	0	120	0	100.0
Toshiba	74	0	74	0	100.0
Asia/Pacific Companies	25	25	0	100.0	0
UMC	25	25	0	100.0	0

Source: Dataquest (July 1995)

Table 3
Each Company's Shipments of 32-Bit and Greater Compute Microprocessors to the World (Thousands of Units)

Architectural Word Width	1993 Units	1994 Units	Percent Change	1994 Market Share (%)
32-Bit	43,885	59,264	35.0	99.5
North American Companies	43,807	59,059	34.8	99.1
Advanced Micro Devices	6,790	8,210	20.9	13.8
Chips & Technologies	126	147	16.7	0.2
Cypress Semiconductor	25	0	-100.0	0
Cyrix	900	2,382	164.7	4.0
IBM	947	3,515	271.2	5.9
Intel	28,045	38,525	37.4	64.7
LSI Logic	25	0	-100.0	0
Motorola	5,444	4,330	-20.5	7.3
Texas Instruments	1,475	1,940	31.5	3.3
Weitek	30	10	-66.7	0
Japanese Companies	78	180	130.8	0.3
Fujitsu	78	180	130.8	0.3
Asia/Pacific Companies	0	25	NA	0
UMC	0	25	NA	0
>32-Bit	254	323	27.2	0.5
North American Companies	98	129	31.6	0.2
Digital Equipment Corporation	30	65	116.7	0.1
Integrated Device Technology	64	64	0	0.1
Performance Semiconductor	4	0	-100.0	0
Japanese Companies	156	194	24.4	0.3
NEC	95	120	26.3	0.2
Toshiba	61	74	21.3	0.1
Total MPU	44,139	59,587	35.0	100.0

NA = Not available
Source: Dataquest (July 1995)

Table 4
Ranking of Each Company's Shipments of 32-Bit and Greater Compute CISC Microprocessors to the World (Thousands of Units)

1994 Rank	1993 Rank	1993 Company	1994 Units	Percent Units	1994 Market Change	Share (%)
1	1	Intel	27,975	38,440	37.4	68.0
2	2	Advanced Micro Devices	6,790	8,210	20.9	14.5
3	3	Motorola	5,373	4,298	-20.0	7.6
4	5	Cyrix	900	2,382	164.7	4.2
5	4	Texas Instruments	1,200	1,540	28.3	2.7
6	6	IBM	872	1,515	73.7	2.7
7	7	Chips & Technologies	126	147	16.7	0.3
8	NM	UMC	0	25	-100.0	*
		Total MPU	43,236	56,557	30.8	100.0

*Calculated value of less than 0.1 percent
NM = Not meaningful
Source: Dataquest (July 1995)

Table 5
Ranking of Each Company's Shipments of 32-Bit and Greater Compute RISC Microprocessors to the World (Thousands of Units)

1994 Rank	1993 Rank	1993 Company	1994 Units	Percent Units	1994 Market Change	Share (%)
1	4	IBM	75	2,000	2,566.7	66.0
2	1	Texas Instruments	275	400	45.5	13.2
3	3	Fujitsu	78	180	130.8	5.9
4	4	NEC	95	120	26.3	4.0
5	6	Intel	70	85	21.4	2.8
6	8	Toshiba	61	74	21.3	2.4
7	9	Digital Equipment Corporation	30	65	116.7	2.1
8	7	Integrated Device Technology	64	64	0	2.1
9	5	Motorola	71	32	-54.9	1.1
10	10	Weitek	30	10	-66.7	0.3
NM	11	Cypress Semiconductor	25	0	-100.0	NA
NM	12	LSI Logic	25	0	-100.0	NA
NM	13	Performance Semiconductor	4	0	-100.0	NA
		Total MPU	903	3,030	235.5	100.0

NA = Not available
NM = Not meaningful
Source: Dataquest (July 1995)

Table 6
Shipments of Each Compute Microprocessor Family to the World
(Thousands of Units)

Family	1993 Units	1994 Units	Percent Change	1994 Market Share (%)	Type
x86	37,863	52,259	38.0	87.7	CISC
Intel	27,975	38,440	37.4	64.5	CISC
Advanced Micro Devices	6,790	8,210	20.9	13.8	CISC
Cyrix	900	2,382	164.7	4.0	CISC
Texas Instruments	1,200	1,540	28.3	2.6	CISC
IBM	872	1,515	73.7	2.5	CISC
Chips & Technologies	126	147	16.7	0.2	CISC
UMC	0	25	NM	0	CISC
68K	5,373	4,298	-20.0	7.2	CISC
Motorola	5,373	4,298	-20.0	7.2	CISC
68EC040	918	2,622	185.6	4.4	CISC
68030	3,439	1,200	-65.1	2.0	CISC
68040	1,016	475	-53.2	0.8	CISC
68060	0	1	NM	0	CISC
SPARC	433	590	36.3	1.0	RISC
Texas Instruments	275	400	45.5	0.7	RISC
Fujitsu	78	180	130.8	0.3	RISC
Weitek	30	10	-66.7	0	RISC
Cypress Semiconductor	25	0	-100.0	0	RISC
LSI Logic	25	0	-100.0	0	RISC
MIPS	224	258	15.2	0.4	RISC
NEC	95	120	26.3	0.2	RISC
Toshiba	61	74	21.3	0.1	RISC
Integrated Device Technology	64	64	0	0.1	RISC
Performance Semiconductor	4	0	-100.0	0	RISC
PowerPC	75	2,000	2,566.7	3.4	RISC
IBM	75	2,000	2,566.7	3.4	RISC
ALPHA	30	65	116.7	0.1	RISC
Digital Equipment Corporation	30	65	116.7	0.1	RISC
Other RISC	141	117	-17.0	0.2	RISC
Intel	70	85	21.4	0.1	RISC
Motorola	71	32	-54.9	0.1	RISC
Total MPU	44,139	59,587	35.0	100.0	

NA = Not available
NM = Not meaningful
Source: Dataquest (July 1995)

Table 7
Shipments of Compute x86 Microprocessors to the World (Thousands of Units)

Product	Architectural Word Width	1993 Units	1994 Units	Percentage Change	Market Share (%)
i486SX	32	11,850	13,295	12.2	25.4
i486DX2	32	6,700	12,802	91.1	24.5
Pentium	32	325	5,040	1,450.8	9.6
i486DX	32	7,750	4,135	-46.6	7.9
AM486DX2	32	150	3,155	2,003.3	6.0
AM386DX	32	4,378	2,550	-41.7	4.9
i486DX4	32	0	2,002	NA	3.8
486SX	32	872	1,515	73.7	2.9
i486SX2	32	0	1,000	NA	1.9
AM486DX	32	400	930	132.5	1.8
AM386SX	32	1,863	910	-51.1	1.7
80486DX	32	25	878	3,412.0	1.7
80486SX	32	250	847	238.8	1.6
TMS486SLC	32	750	740	-1.3	1.4
TMS486DLC	32	400	700	75.0	1.3
80486SLC	32	525	600	14.3	1.1
AM486S	32	0	600	NA	1.1
i486SL	32	1,350	166	-87.7	0.3
80386DX	32	90	105	16.7	0.2
TMS486S	32	50	100	100.0	0.2
AM486DX3	32	0	65	NA	0.1
80486DLC	32	100	57	-43.0	0.1
80386SX	32	36	42	16.7	0.1
UM386	32	0	25	NA	0
Total MPU	37,863	52,259	NA	100.0	

NA = Not available
Source: Dataquest (July 1995)

For More Information...

Sue Kapoor, Market Research Analyst.....(408) 437-8505
Internet address .....skapoor@dataquest.com
Via fax.....(408) 437-0292

Dataquest®

a company of
The Dun & Bradstreet Corporation

DATAQUEST WORLDWIDE OFFICES

NORTH AMERICA

Worldwide Headquarters

Dataquest Incorporated
1290 Ridder Park Drive
San Jose, California 95131-2398
United States
Phone: 1-408-437-8000
Facsimile: 1-408-437-0292

Dataquest Incorporated

Nine Technology Drive
P.O. Box 5093
Westborough, Massachusetts 01581-5093
United States
Phone: 1-508-871-5555
Facsimile: 1-508-871-6180

Dataquest Global Events

3990 Westerly Place, Suite 100
Newport Beach, California 92660
United States
Phone: 1-714-476-9117
Facsimile: 1-714-476-9969

Sales Offices:
- Washington, DC (Federal)
- New York, NY (Financial)
- Dallas, TX

LATIN AMERICA

Research Affiliates and Sales Offices:
- Buenos Aires, Argentina
- Sao Paulo, Brazil
- Santiago, Chile
- Mexico City, Mexico

EUROPE

European Headquarters

Dataquest Europe Limited
Holmers Farm Way
High Wycombe, Bucks HP12 4XH
United Kingdom
Phone: +44 1494 422 722
Facsimile: +44 1494 422 742

Dataquest Europe SA

Immeuble Défense Bergères
345, avenue Georges Clémenceau
TSA 40002
92882 - Nanterre CTC Cedex 9
France
Phone: +33 1 41 35 13 00
Facsimile: +33 1 41 35 13 13

Dataquest GmbH

Kronstadter Strasse 9
81677 München
Germany
Phone: +49 89 93 09 09 0
Facsimile: +49 89 93 03 27 7

Sales Offices:
- Brussels, Belgium
- Kfar Saba, Israel
- Milan, Italy
- Randburg, South Africa
- Madrid, Spain

JAPAN

Japan Headquarters

Dataquest Japan K.K.
Shinkawa Sanko Building
6th Floor
1-3-17, Shinkawa
Chuo-ku, Tokyo 104
Japan
Phone: 81-3-5566-0411
Facsimile: 81-3-5566-0425

ASIA/PACIFIC

Asia/Pacific Headquarters

7/F China Underwriters Centre
88 Gloucester Road
Wan Chai
Hong Kong
Phone: 852-2824-6168
Facsimile: 852-2824-6138

Dataquest Korea

Suite 2407, Trade Tower
159 Samsung-dong, Kangnam-gu
Seoul 135-729
Korea
Phone: 822-551-1331
Facsimile: 822-551-1330

Dataquest Taiwan

11F-2, No. 188, Section 5
Nan King East Road
Taipei
Taiwan, R.O.C.
Phone: 8862-756-0389
Facsimile: 8862-756-2663

Dataquest Singapore

105 Cecil Street
#06-01/02 The Octagon
Singapore 0106
Phone: 65-227-1213
Facsimile: 65-227-4607

Dataquest Thailand

12/F, Vanissa Building
19 Soi Chidlom
Ploenchit Road
Patumwan, Bangkok 10330
Thailand
Phone: 662-655-0577
Facsimile: 662-655-0576

Research Affiliates and Sales Offices:
- Melbourne, Australia
- Beijing, China

Dataquest

a company of
The Dun & Bradstreet Corporation

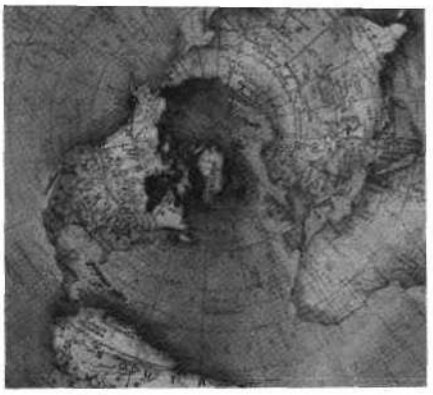

Dataquest

PC Semiconductor Application Markets

Market Trends

Program: Semiconductor Directions in PCs
Product Code: PSAM-WW-MT-9502
Publication Date: November 13, 1995
Filing: Market Analysis

PC Semiconductor Application Markets

Market Trends

Program: Semiconductor Directions in PCs
Product Code: PSAM-WW-MT-9502
Publication Date: November 13, 1995
Filing: Market Analysis

Table of Contents

List of Figures

List of Tables

List of Tables (Continued)

Chapter 1
Introduction and Methodology

This document provides reference information and analysis about the personal computer application market for semiconductors. Specific areas of information include:

- PC system market size (in production terms) in revenue, units, and average selling price
- PC system market and product feature trends
- Hardware architecture trends and semiconductor device opportunities
- Semiconductor content and market forecast
- Listings of key OEMs

The information in this report is gathered from both primary and secondary sources. Primary sources include surveys and interviews of industry vendors and customers, as well as analyst knowledge and opinion. Some of the primary sources include Dataquest's own industry services. Secondary sources include various governmental and trade sources on sales, production, trade, and public spending. Semiconductor content assumptions are based on both surveys of producing OEMs and physical teardown evaluations by Dataquest analysts of representative systems.

The forecast methodology is based on various methods and assumptions, depending upon the area. To form a solid basis for projecting system demand, capital, government, and consumer spending assumptions are made for various regions of the world. For specific markets, saturation and displacement dynamics are considered as well. Key exogenous factors such as new software introductions, exchange rates changes, and government policies are also considered. Semiconductor content forecasts are based upon interviews of system marketers and designers (including makers of enabling semiconductor technology) along with an analysis of historical trends.

Project Analysts: Geoff Ballew and Dale Ford

Chapter 2
Personal Computer Markets

Market and Production Trends

Personal Computers

Market and production trends in PCs are as follows:

- Although notebook and subnotebook computers have enjoyed strong growth, desktop systems will continue to dominate the computing platform because of their lower manufacturing costs, the ability to handle robust configurations with large amounts of storage, and their role in offering the fastest and latest computing technologies.
- Notebooks have enjoyed solid unit and revenue growth because of their ability to replicate desktop environments with features such as Windows applications, the introduction of more powerful microprocessors such as the Pentium, and larger hard drive capacities.
- Ultraportable and notepad shipments continue to surge with popular offerings from Apple, Olivetti, Gateway 2000, Hewlett-Packard, and others. This market is expected to exceed notebook shipments worldwide by the end of 1998.
- Shipments of laptops and transportable units continue to decline as smaller notebook systems make them unattractive.
- In the world of mobile computing, handheld devices such as PDAs have commanded incredible attention as they have captured the imagination of consumers and manufacturers alike. Dataquest has defined two categories for these handheld devices:
  - Expandable organizers: These are computers that typically measure 3.0 x 6.0 x 0.75 inches and weigh less than a pound. They are distinguished by the capability to have the user add applications and memory, and the fact that expansion is proprietary to a particular device or family. The operating systems typically are proprietary. The market for these devices is expected to peak at more than 0.5 million unit shipments in 1995 and then slowly shrink to fewer than 0.4 million by 1999.
  - Standard handheld computers: These devices typically measure 4 x 7 x 1 inches and weigh about a pound. They are distinguished from expandable organizers by their adherence to hardware and software compatibility standards. The operating systems are open and licensed, and application development and memory expansion are open to third-party developers and may be distributed in a standard format (such as PCMCIA). PDAs fall into this category.
- Strong growth in the worldwide home PC market will continue across all regions. This sector is projected to account for nearly half of all PC shipments in 1998 in the United States, the most mature market.
- Microsoft's new mainstream operating system "Windows 95" was launched in the second half of 1995 with actual product sales beginning on August 24. Its multitasking nature will have an impact on hardware such as SCSI controllers. It also will raise the demand for memory, forcing users to upgrade to 12MB or more.

- Emerging markets and regions will account for a sharply increased proportion of shipments by 1998.
- The multimedia-ready PC market continued its torrid growth in 1994. Shipments of complete multimedia systems grew fourfold from 1993 to 1994, reaching 10.3 million shipments. Dataquest expects sales of complete multimedia systems to continue to grow during the next five years.
- Major PC manufacturers will continue to place a strong emphasis on the branding of PCs.
- Manufacturing strategies are dictated by economies of scale with strong influence from free trade zones, the availability of components, and the quality and price of labor. The trend is toward assembling PCs close to the end market or at least final configuration being executed on demand close to the end market.
- There is a clear trend toward using turnkey or contracted manufacturing sources such as Intel, SCI, and Solectron.

Figures 2-1 to 2-3 and Tables 2-1 to 2-7 present Dataquest's worldwide PC market forecasts. Figures 2-4 and 2-5 show the regional variation in markets and production from 1994 to 1999. Figure 2-6 illustrates the North American PC and workstation motherboard production forecast given in Table 2-8.

Figure 2-1
Worldwide Personal Computer Shipments by Microprocessor Type (Thousands of Units) (Includes Motherboard Upgrade Market; Excludes Standard Handhelds)

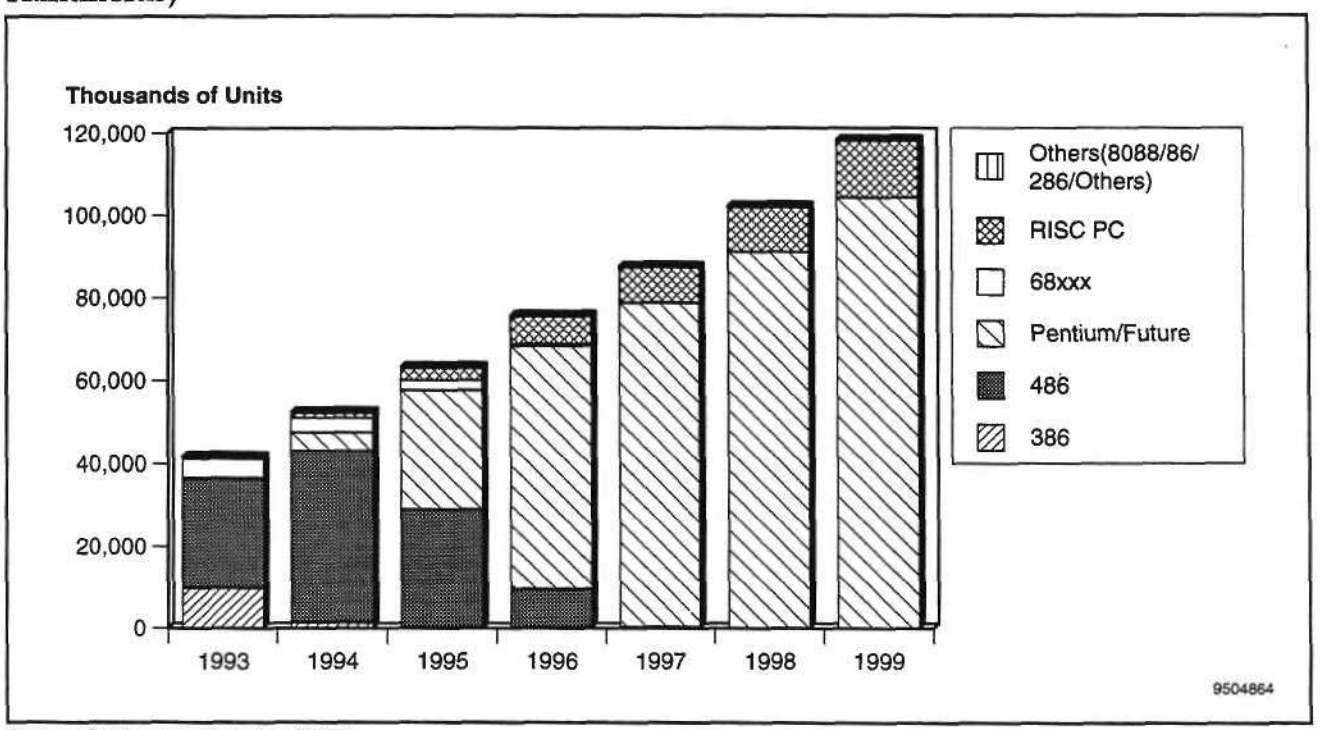

Source: Dataquest (October 1995)

Figure 2-2
Worldwide Desktop/Deskside Personal Shipments by Microprocessor Type (Includes Motherboard Upgrade Market)

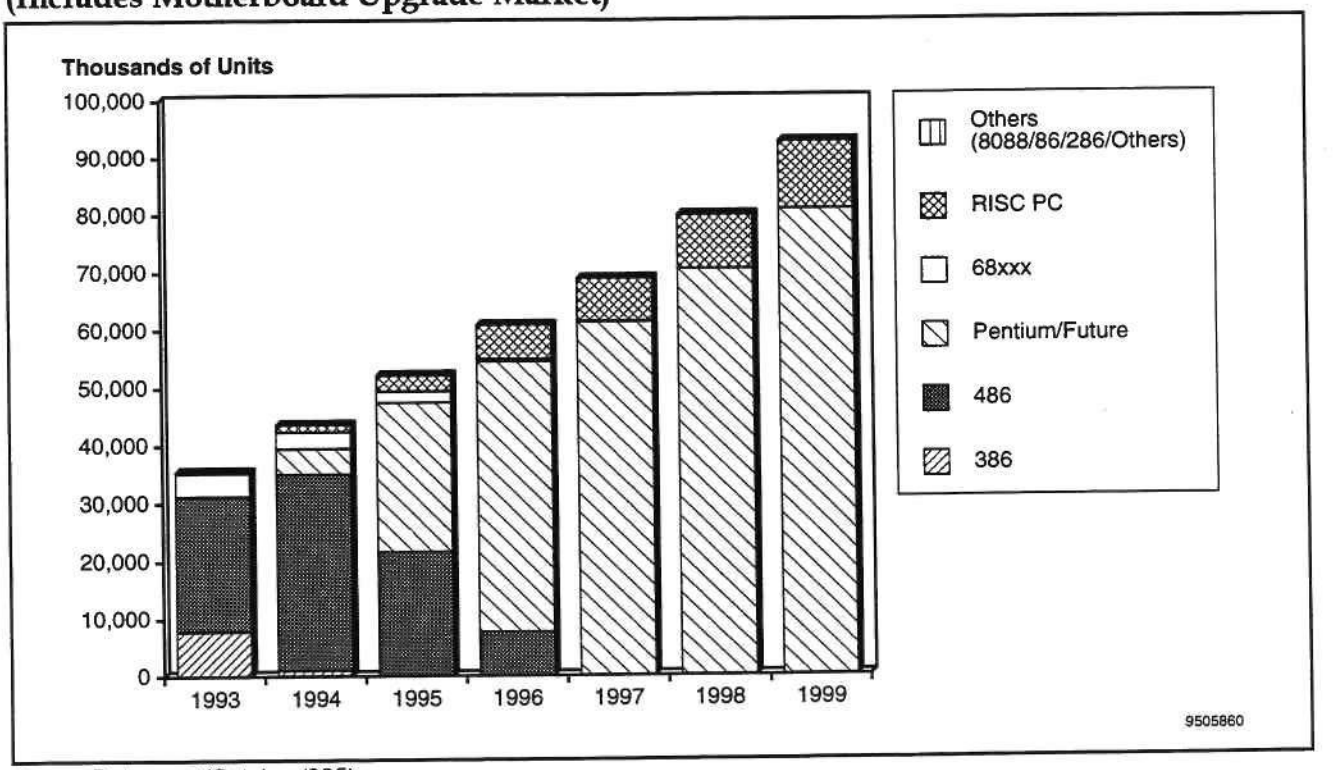

Source: Dataquest (October 1995)

Figure 2-3
Worldwide Mobile Personal Computer Shipments by Microprocessor Type

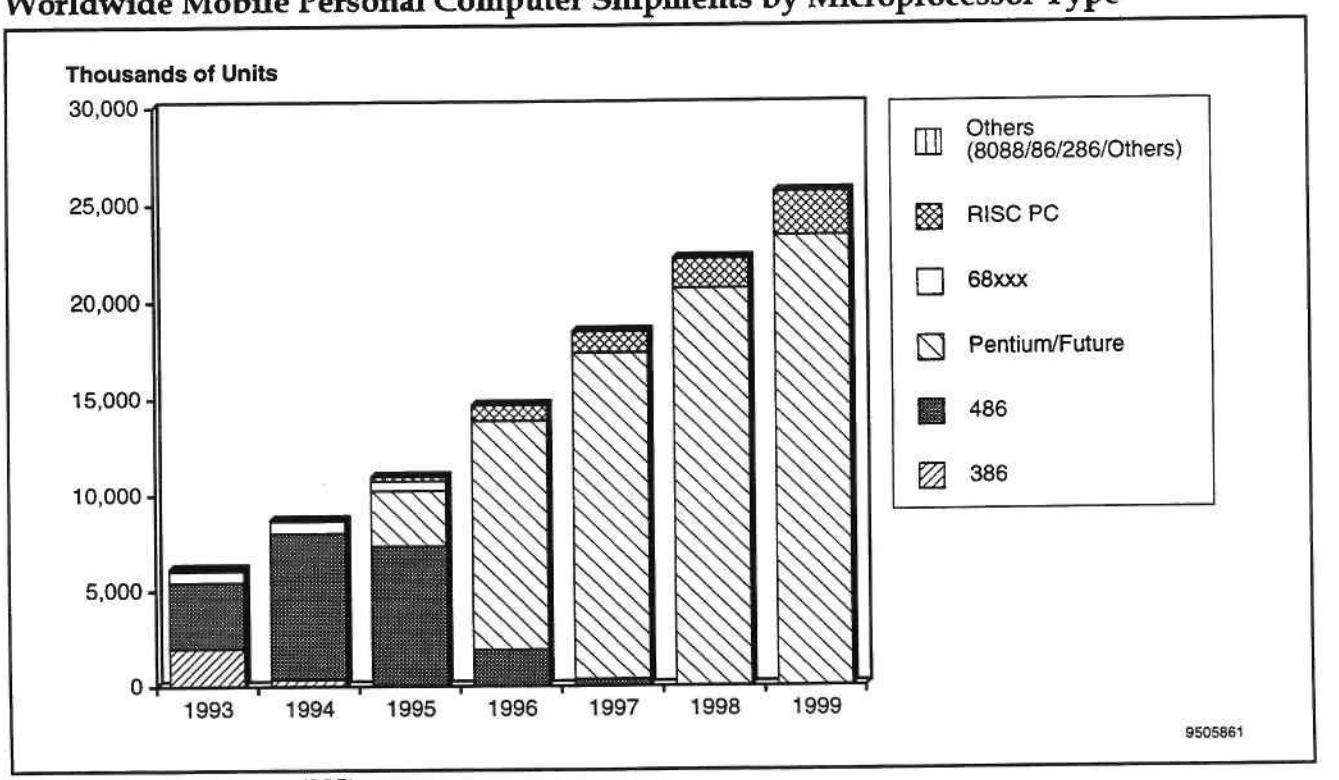

Source: Dataquest (October 1995)

Table 2-1
Worldwide Personal Computer Market by Product Type*

Product Type	1993	1994	1995	1996	1997	1998	1999	CAGR (%) 1994-1999
Desktop/Deskside								
Units (K)	32,712	39,219	46,623	54,056	60,865	70,221	81,419	15.7
ASP ($K)	1.80	1.87	1.95	1.90	1.84	1.79	1.74	-1.4
Factory Revenue ($M)	59,039	73,297	91,012	102,579	112,204	125,847	141,916	14.1
Transportable								
Units (K)	76	122	39	32	18	16	15	-34.2
ASP ($K)	4.11	6.72	6.69	6.07	5.38	4.92	4.51	-7.7
Factory Revenue ($M)	312	817	262	194	97	79	68	-39.2
Laptop								
Units (K)	170	47	13	0	0	0	0	-100.0
ASP ($K)	3.90	5.04	4.98	0	0	0	0	-100.0
Factory Revenue ($M)	662	237	63	0	0	0	0	-100.0
Notebook/Tablet								
Units (K)	5,477	7,489	9,414	11,559	12,379	13,077	12,780	11.3
ASP ($K)	2.35	2.49	2.52	2.36	2.32	2.18	2.06	-3.7
Factory Revenue ($M)	12,851	18,638	23,746	27,256	28,696	28,554	26,369	7.2
Ultraportable/Notepad								
Units (K)	417	1,007	1,479	3,065	6,051	9,041	12,786	66.3
ASP ($K)	1.67	1.83	1.78	1.79	1.77	1.70	1.64	-2.3
Factory Revenue ($M)	697	1,846	2,639	5,473	10,727	15,337	20,922	62.5
Standard Handheld (Palmtop)								
Units (K)	237	390	542	775	1,225	2,600	5,245	68.2
ASP ($K)	0.57	0.82	0.48	0.46	0.41	0.35	0.35	-15.7
Factory Revenue ($M)	135	318	258	357	502	913	1,842	42.1
Motherboard Upgrades								
Units (K)	2,759	4,404	5,476	6,738	7,992	9,582	11,190	20.5
ASP ($K)	1.25	1.32	1.40	1.39	1.33	1.28	1.24	-1.2
Factory Revenue ($M)	3,449	5,814	7,673	9,362	10,655	12,238	13,887	19.0
Total PC (Including Handheld and Motherboard Upgrade)								
Units (K)	41,847	52,677	63,586	76,225	88,530	104,537	123,435	18.6
ASP ($K)	1.9	2.0	2.0	2.0	1.9	1.8	1.8	-2.2
Factory Revenue ($M)	77,144	100,968	125,653	145,222	162,881	182,969	205,004	15.2

*ASP and factory revenue include standard peripherals and memory upgrades installed up to the point of sale.
Source: Dataquest (September 1995)

Table 2-2
Worldwide Personal Computer Shipments by Microprocessor Type (Thousands of Units) (Includes Motherboard Upgrade Market; Excludes Standard Handhelds)

	1993	1994	1995	1996	1997	1998	1999	CAGR (%) 1994-1998
386	9,873	1,432	47	0	0	0	0	-100.0
486	26,619	41,693	28,923	9,667	416	0	0	-100.0
Pentium/Future	227	4,451	28,746	58,681	78,342	91,183	104,279	87.9
68xxx	4,609	3,503	2,461	471	158	0	0	-100.0
RISC PC	76	1,161	2,863	6,631	8,389	10,754	13,911	64.3
Others (8088/86/286/Others)	205	40	0	0	0	0	0	-100.0
Total	41,609	52,280	63,040	75,450	87,305	101,937	118,190	17.7

Source: Dataquest (September 1995)

Table 2-3
Worldwide Desktop/Deskside Personal Computer Shipments by Microprocessor Type (Thousands of Units) (Includes Motherboard Upgrade Market)

	1993	1994	1995	1996	1997	1998	1999	CAGR (%) 1994-1999
386	7,875	1,075	35	0	0	0	0	-100.0
486	23,187	34,062	21,608	7,754	53	0	0	-100.0
Pentium/Future	227	4,429	25,832	46,745	61,333	70,524	80,929	78.8
68xxx	3,994	2,858	1,960	471	158	0	0	-100.0
RISC PC	75	1,161	2,662	5,825	7,313	9,279	11,680	58.7
Others (8088/86/286/ Others)	113	38	0	0	0	0	0	-100.0
Total	35,471	43,623	52,097	60,795	68,857	79,803	92,609	16.2

Source: Dataquest (October 1995)

Table 2-4
Worldwide Mobile Personal Computer Shipments by Microprocessor Type (Thousands of Units) (Excludes Standard Handhelds)

	1993	1994	1995	1996	1997	1998	1999	CAGR (%) 1994-1999
386	1,998	357	12	0	0	0	0	-100.0
486	3,432	7,631	7,315	1,913	363	0	0	-100.0
Pentium/Future	0	23	2,914	11,936	17,009	20,659	23,350	299.8
68xxx	615	645	501	0	0	0	0	-100.0
RISC PC	1	0	201	806	1,076	1,475	2,231	-
Others (8088/86/286/ Others)	93	2	0	0	0	0	0	-100.0
Total	6,139	8,657	10,943	14,655	18,448	22,134	25,581	24.2

Source: Dataquest (October 1995)

Table 2-5
Worldwide Personal Computer Revenue by Microprocessor Type (Millions of Dollars) (Includes Motherboard Upgrade Market; Excludes Standard Handhelds)*

	1993	1994	1995	1996	1997	1998	1999	CAGR (%) 1994-1999
386	13,617	1,453	40	0	0	0	0	-100.0
486	54,022	76,429	53,261	14,388	564	0	0	-100.0
Pentium/Future	1,075	13,913	62,747	118,429	148,015	164,845	181,489	67.1
68xxx	8,081	5,872	3,736	592	181	0	0	-100.0
RISC PC	74	2,938	5,604	11,458	13,620	17,210	21,671	49.1
Others (8088/86/286/Others)	140	22	0	0	0	0	0	-100.0
Total	77,009	100,627	125,388	144,866	162,380	182,055	203,160	15.1

*Factory revenue includes standard peripherals and memory upgrades installed up to the point of sale.
Source: Dataquest (September 1995)

Table 2-6
Worldwide Desktop/Deskside Personal Computer Revenue by Microprocessor Type* (Millions of Dollars) (Includes Motherboard Upgrade Market)

	1993	1994	1995	1996	1997	1998	1999	CAGR (%) 1994-1999
386	9,829	948	27	0	0	0	0	-100.0
486	44,844	57,074	36,285	10,964	67	0	0	-100.0
Pentium/Future	1,075	13,720	54,558	90,684	111,196	123,743	138,159	58.7
68xxx	6,590	4,410	2,725	592	181	0	0	-100.0
RISC PC	72	2,938	5,088	9,703	11,416	14,342	17,643	43.1
Others (8088/86/286/ Others)	78	21	0	0	0	0	0	-100.0
Total	62,488	79,111	98,683	111,942	122,860	138,085	155,802	14.5

*Factory revenue includes standard peripherals and memory upgrades installed up to the point of sale.
Source: Dataquest (September 1995)

Table 2-7
Worldwide Mobile Personal Computer Revenue by Microprocessor Type* (Millions of Dollars) (Includes Motherboard Upgrade Market; Excludes Standard Handhelds)

	1993	1994	1995	1996	1997	1998	1999	CAGR (%) 1994-1999
386	3,788	505	13	0	0	0	0	-100.0
486	9,179	19,355	16,976	3,424	497	0	0	-100.0
Pentium/Future	0	193	8,189	27,745	36,819	41,102	43,330	195.3
68xxx	1,491	1,462	1,011	0	0	0	0	-100.0
RISC PC	2	0	516	1,755	2,204	2,868	4,028	-
Others (8088/86/286/ Others)	62	1	0	0	0	0	0	-100.0
Total	14,521	21,516	26,705	32,924	39,520	43,970	47,358	17.1

*Factory revenue includes standard peripherals and memory upgrades installed up to the point of sale.
Source: Dataquest (September 1995)

Figure 2-4
Regional Personal Computer Markets (Unit Shipments)

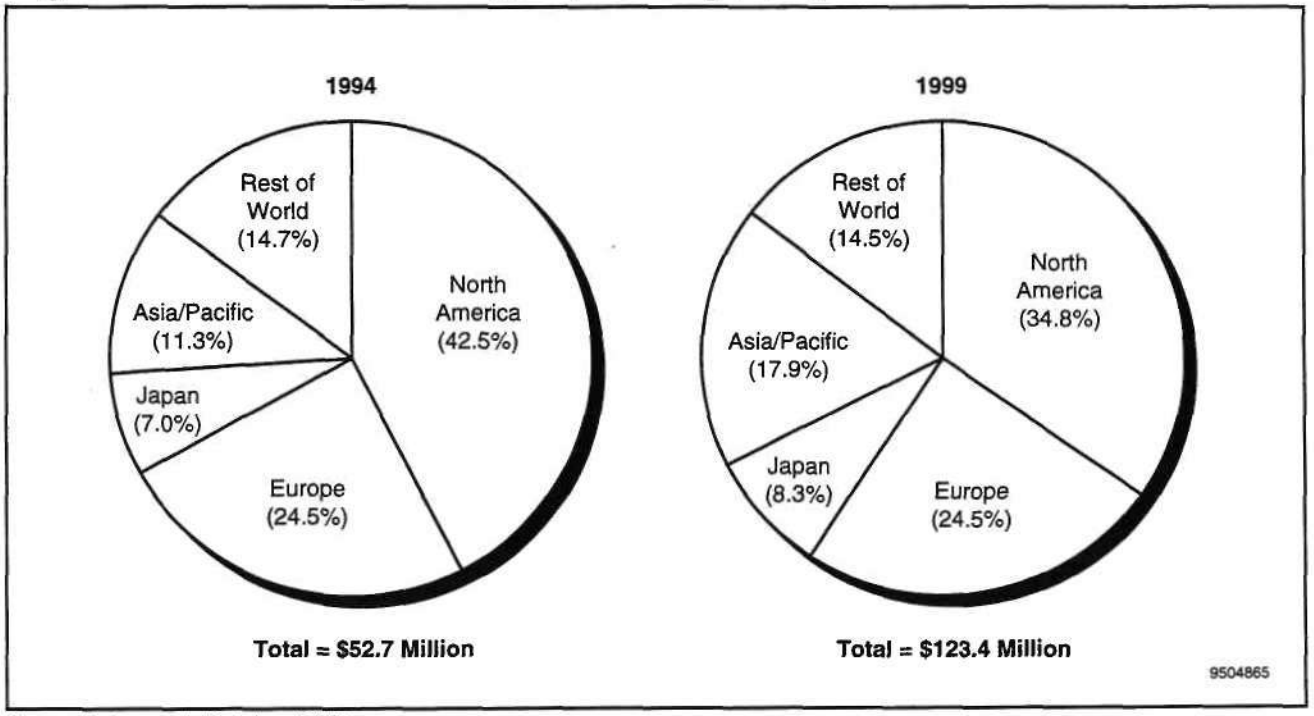

Source: Dataquest (October 1995)

Figure 2-5
Regional Personal Computer Motherboard Production (Unit Shipments)

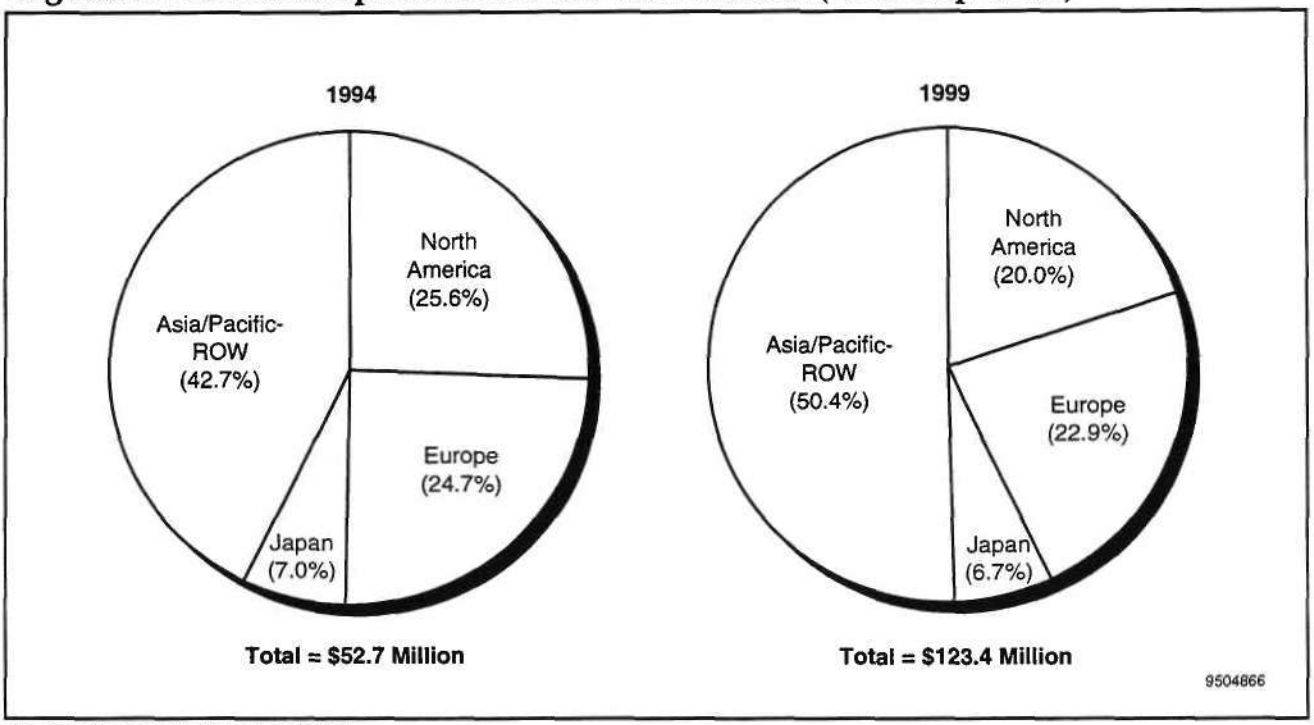

Source: Dataquest (October 1995)

Figure 2-6
North American Motherboard Production

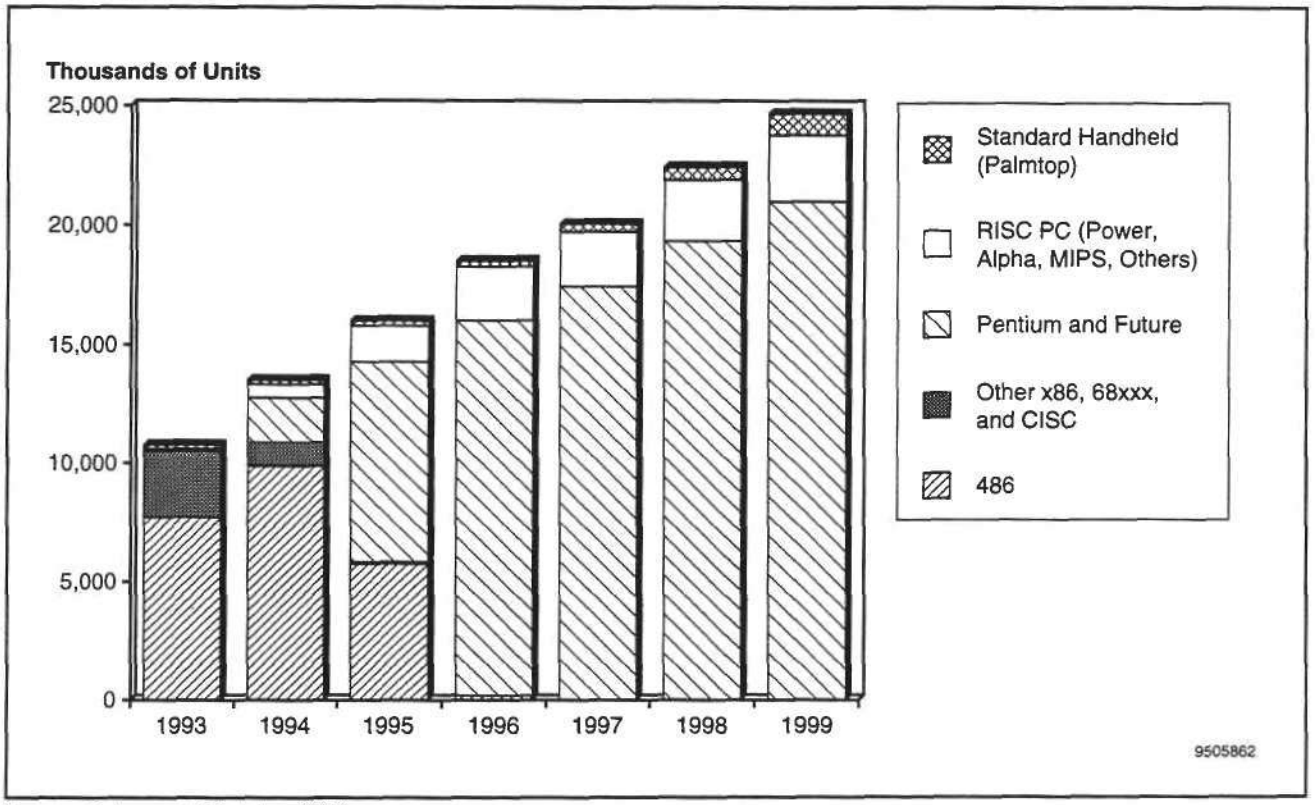

Source: Dataquest (October 1995)

Table 2-8
North American Motherboard Production by Microprocessor Type (Thousands of Units)

	1993	1994	1995	1996	1997	1998	1999	CAGR (%) 1994-1999
486	7,725	9,883	5,767	199	0	0	0	-100.0
Other x86, 68xxx, and CISC	2,765	984	74	0	0	0	0	-100.0
Pentium/Future	127	1,905	8,443	15,813	17,441	19,348	20,994	61.7
RISC PC (Power, Alpha, MIPS, Others)	0	568	1,510	2,261	2,300	2,568	2,788	37.5
Standard Handheld (Palmtop)	130	160	184	229	303	511	888	40.9
Total	10,747	13,500	15,979	18,503	20,043	22,427	24,669	12.8

Source: Dataquest (September 1995)

OEMs

Tables 2-9 and 2-10 list the key OEMs ranked by revenue and unit shipment market share. For the first time in more than a decade, a company other than IBM or Apple Computer leads the world in PC shipments in 1994. Compaq Computer not only leapfrogged the perennial PC shipment leaders, but did so in dramatic fashion by shipping 847,000 more PCs than any other manufacturer. Unable to continue grabbing market share away from the clone makers in 1994, the large PC suppliers were forced to turn their guns on each other, with Compaq and Packard Bell winning the battle. They were the only companies in the top five to increase their share of the market in 1994.

Table 2-9
Personal Computers Worldwide Revenue Market Share (Percentage)*

	1993	1994
Compaq	9.2	10.5
IBM	11.8	8.8
Apple	10.1	8.5
NEC	5.3	5.5
Packard Bell	1.7	3.9
Toshiba	2.3	3.7
Dell	3.6	3.3
Gateway 2000	2.2	2.7
Hewlett-Packard	1.9	2.7
Acer	1.2	2.3
Other Vendors	50.7	48.2

*Does not include motherboard upgrade revenue
Source: Dataquest (September 1995)

Table 2-10
Personal Computers Worldwide Unit Market Share (Percentage)*

	1993	1994
Compaq	8.05	10.03
IBM	10.84	8.26
Apple	9.43	8.26
Packard Bell	2.94	5.16
NEC	4.16	4.05
AST Research	2.68	2.72
Dell	3.02	2.68
Hewlett-Packard	1.75	2.66
Acer	1.41	2.63
Toshiba	1.73	2.53
Other Vendors	53.99	51.02

*Does not include motherboard upgrade shipments
Source: Dataquest (September 1995)

Systems Technology and Architecture Trends

The technology road maps of PCs and workstations principally are driven by the twin factors of economy and improved end-user features. PC makers typically try to give the buyer more features with each new product generation while maintaining price points. Workstation makers are known more for emphasizing performance and other features, although the low end of the workstation market is attempting to be competitive with the high end of the PC range. Some specific trends are listed in the following sections.

Personal Computers

Trends in PCs are as follows:

- Desktop power demand will come from servicing graphical user interface (GUI) and WYSIWYG-oriented applications including processing (compressing and decompressing); store, forward, and read multimedia (OLE enabled); and real-time multimedia such as desktop videoconferencing. Networking support increasingly will require faster servicing as transfer rates soar (for example, toward 100 Mbps over Ethernet) and the size and frequency of transferred files and e-mail soar as well.
- Expect continued proliferation of variations on the Pentium theme as Intel, Advanced Micro Devices, and Cyrix cover the entire price/performance/power consumption spectrum.
- The Power PC alliance will continue to be a major element in Apple, IBM, and other units.
- The backplane bus, long pointed out as one of the bottlenecks limiting PC performance, will be supplemented by the local bus with domination by the PCI standard. ISA, EISA, and other buses will continue to exist for some time because of the vast bulk of add-in cards based on these buses.
- Plug-and-play PCs, with support from major players such as Microsoft, Intel, Compaq, AMD, and others, will be a significant feature in the drive to make PCs more user-friendly.
- The "green" PC as stimulated by the Energy Star program and similar requirements coming out of Europe should become a viable selling point for future PCs.
- The SCSI interface has had a short-term upswing in use for accessing larger disk capacities, disk arrays, CD-ROM drives, and other peripherals such as scanners. Because SCSI is more expensive to implement, IDE and its enhanced version (EIDE) will remain the dominant RDD interface.
- Serial and parallel I/O have remained fairly constant, but now new technologies such as 1394/Firewire, Access.bus, and universal serial bus (USB) are vying to displace those trusted standbys.
- The USB standard was rolled out with strong support from Intel, Microsoft, and Compaq at the Windows Hardware Engineering Conference '95 (WINHEC '95). This local bus for the PC will provide a common standard for communicating among PCs, modems, scanners, joysticks, keyboards, mice, monitors, and printers.

- The worldwide PC Card market started to take off in 1994 with 3.4 million unit shipments. Customers and vendors have started to realize how PC Cards can help mobile workers be more productive by enabling them to simply and quickly add new capabilities to their mobile devices. While modem cards were by far the largest market in terms of unit shipments in 1994, vendors will need to incorporate multiple functions in PC Card technology in order to stay viable in the long term.
- A key trend in mobile technology will be the purchase of skeleton systems by users that are then enhanced with shrink-wrapped PCMCIA RDDs, memory cards, and communications features.
- Other key trends in mobile computing will include commercialized low-power-consuming components, standardization of operating voltage for circuits and components at 3V/3.3V, continued increases in sophistication of battery management, improved battery technology, and effective handwriting/voice recognition.

Table 2-11 highlights projected trends in PC technology. Figure 2-7 shows areas that offer opportunities for enhancement.

Tables 2-12 through 2-16 provide forecasts for key PC and workstation technologies.

Table 2-11
The Evolving PC

Standard Features	1993	1995	1997	Trend
Main Memory (MB)	4-6	8-12	16-20	To wider organizations, sub-70ns, and cache replacement. Toward EDO Architecture
CPU*	386 and 486 (SX/DX)/68K	486DXx/Pentium PowerPC	Pentium/ PowerPC/ Pentium Pro	Toward more power and higher price/ performance ratio
Cache (K)	64/128	128/256	256	Possible integration with MPU or Core Logic
Bus*	ISA et al.	ISA with VL/PCI	ISA with PCI	PCI dominant and higher-speed versions of PCI
Video	VGA/Mac (1/4MB buffer)	32/64-bit Accelerated GUI (1MB buffer)	64-bit Accelerated GUI (1MB buffer)	Addition of digital video and 3-D features
		Accelerated/Mac	Motion Video Acceleration	
Storage	IDE/SCSI	EIDE/SCSI	EIDE/SCSI	Enhanced IDE (EIDE) will dominate
Local I/O	RS-232/422	RS-232/422	RS-232/422/USB	Toward USB, possible entrance of 1394/ Access.bus

*And corresponding core logic
Source: Dataquest (September 1995)

Figure 2-7
PC Performance Enhancement Opportunities

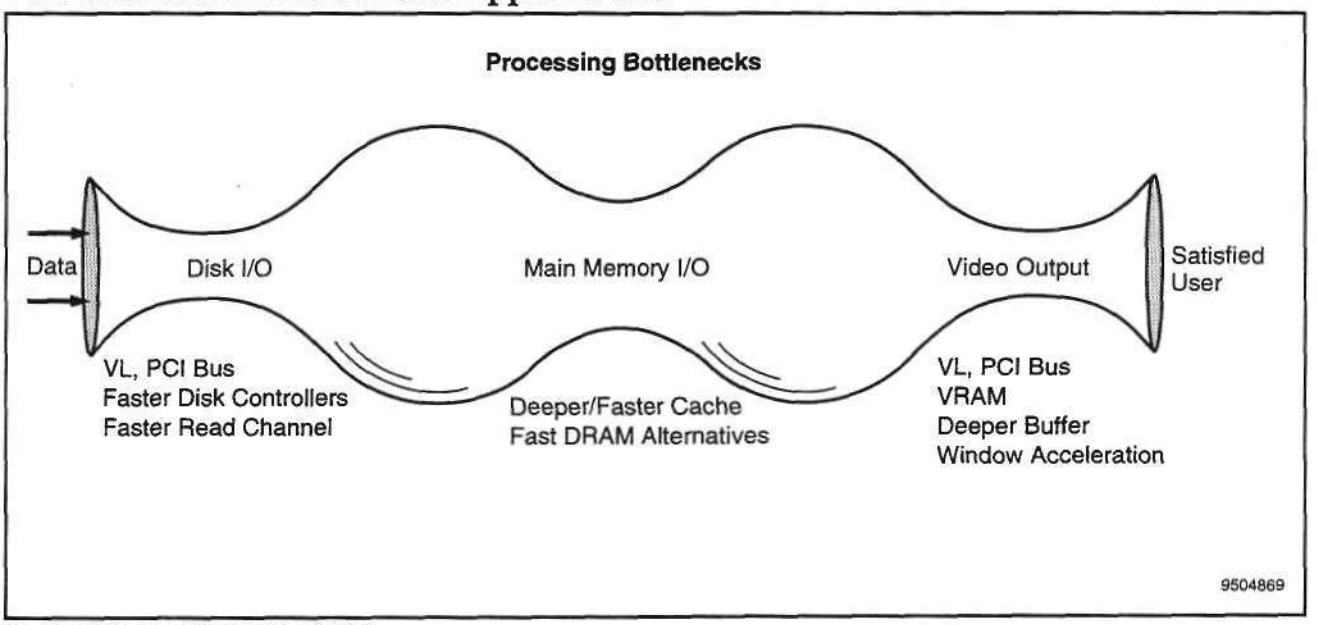

Source: Dataquest (September 1995)

Table 2-12
Worldwide PC Bus Requirements (Percentage of Total or Category)

	1993	1994	1995	1996	1997	1998	1999
AT/ISA Total	68	70	70	67	65	64	64
ISA Only	18	6	1	0	0	0	0
With VL	81	75	50	23	5	1	1
With PCI	1	19	49	77	95	99	99
EISA Total	3	3	3	3	3	2	2
EISA Only	79	23	8	4	2	0	0
With PCI	21	77	92	96	98	100	100
MCA Total	3	2	1	1	1	0	0
Proprietary/Specialty	10	8	7	8	8	9	10
Proprietary/Specialty Only	100	100	95	92	90	85	85
With PCI	0	0	5	8	10	15	15
No Bus (Notebook, among Others)	16	18	19	21	23	24	24
Total	100	100	100	100	100	100	100

Source: Dataquest (October 1995)

Table 2-13
Worldwide Penetration of PCMCIA/PC Card Slot Forecast (Thousands of Slots)

	1993	1994	1995	1996	1997	1998	1999
With One or More Type II Slots	4,869	10,389	14,545	16,751	19,173	21,571	23,813
With One or More Type III Slots	1,889	4,040	5,333	5,973	6,272	5,958	5,482

Source: Dataquest (October 1995)

Table 2-14
Worldwide PC (Host) Storage Interface Requirement (Percentage of Category)

	1993	1994	1995	1996	1997	1998	1999
x86/68xxx/Others							
IDE/AT	84	85	86	88	90	92	94
SCSI (I/II/III)	16	15	14	12	10	8	6
FC-AL/SSA/1394	0	0	0	0	0	0	0
Total	100	100	100	100	100	100	100
Pentium and Future PC							
IDE/AT	10	76	85	86	85	83	80
SCSI (I/II/III)	90	24	15	13	13	12	10
FC-AL/SSA/1394	0	0	0	1	2	5	10
Total	100	100	100	100	100	100	100
RISC PC							
IDE/AT	2	10	14	16	17	17	17
SCSI (I/II/III)	98	90	86	82	79	75	69
FC-AL/SSA/1394	0	0	0	2	4	8	14
Total	100	100	100	100	100	100	100

Source: Dataquest (October 1995)

Table 2-15
Worldwide PC Graphics Standard Penetration (Motherboard-Based or Add-In)

	1993	1994	1995	1996	1997	1998	1999
VGA (Percentage of Total)	56	28	5	1	0	0	0
Accelerated VGA (Percentage of Total)	41	66	77	45	14	2	0
Accelerated VGA + Motion Video and Future Features (Percentage of Total)	0	3	18	54	86	98	100
Others (XGA, pre-VGA)	3	3	0	0	0	0	0

Source: Dataquest (October 1995)

Table 2-16
Worldwide Embedded Communications and Multimedia in PCs (Percentage of Total)

	1993	1994	1995	1996	1997	1998	1999
Overall PC							
Modem	1.5	1.6	1.8	2.0	2.3	2.5	2.8
LAN	1.0	1.5	2.9	6.0	10.0	17.0	26.0
USB	-	-	-	8.0	35.0	94.0	100.0
Sound I/O	14.0	16.0	19.0	25.0	34.0	45.0	56.0
Graphics	44.0	46.0	50.0	55.0	57.0	60.0	63.0
MPEG Decoder (%)	-	-	-	4.3	3.5	2.8	2.0
General-Purpose DSP (%)	0.1	0.5	1.0	2.0	3.0	4.0	5.0
Desktop/Deskside PC							
Modem	1.4	1.5	1.7	2.0	2.4	2.7	3.0
LAN	1.0	2.0	4.0	7.0	13.0	22.0	33.0
USB	-	-	-	10.0	40.0	98.0	100.0
Sound I/O	12.0	13.0	15.0	20.0	27.0	36.0	46.0
Graphics	34.0	35.0	40.0	44.0	46.0	49.0	53.0
MPEG Decoder (%)	-	-	-	5.0	3.8	2.5	1.0
General-Purpose DSP (%)	0.1	0.4	0.8	1.6	2.3	3.2	4.0
Mobile PC							
Modem	2.0	2.0	2.0	2.0	2.0	2.0	2.0
LAN	0.1	0.1	0.1	0.1	0.1	0.1	0.1
USB	-	-	-	2.0	20.0	80.0	100.0
Sound I/O	23.0	29.0	36.0	46.0	60.0	76.0	92.0
Graphics	100.0	100.0	100.0	100.0	100.0	100.0	100.0
MPEG Decoder (%)	-	-	-	1.2	2.4	3.9	5.6
General-Purpose DSP (%)	0.2	0.9	1.8	3.6	5.5	7.0	8.5

Source: Dataquest (October 1995)

Chapter 3
Market Semiconductor Trends for PCs

Semiconductor Market Opportunities and Technology Trends

Semiconductor market opportunities and technology trends are as follows:

- New generations of Pentium (Pentium Pro, P7, and P8), PowerPC, Alpha, PA-RISC, SPARC, and MIPS processors will drive future PCs, workstations, and servers. Future trends include internal frequencies soaring past 100-MHz for PCs and toward 200-MHz for workstations. Multiprocessing features will also be available in many workstations and high-end PCs.
- Proprietary processors such as ARM will continue to be used in palmtop devices.
- PCs with 8 to 12MB average DRAM main memory in 1995 will move to 32MB DRAM in 1998.
- Increased opportunities will emerge for fast/synchronous DRAMs as MPU speeds outstrip conventional memory architectures.
- Bursting, MPU-specific SRAM will be used for cache design. SRAM modules are becoming more popular.
- New chipsets to support MPUs and the new buses (particularly PCI) will continue to emerge.
- Mixed-signal accelerated window graphics on a single chip (with RAMDAC) will be mainstream. High-resolution (1024 x 768, 16-color) 3-D graphics (Open GL) and digital video feature will become more prevalent.
- Mixed-signal I/O chips will support high-speed storage and peripherals communications (EIDE, SCSI, 1394/Firewire). Chips integrating LAN and SCSI control will complement existing super I/O functions.
- Sound and video codecs and I/O functions will be used in multimedia systems. Microsoft, IBM, and Apple are continuing to provide multimedia and telephony interfaces to their operating systems.

Figures 3-1 to 3-4 provide block diagrams for generic Pentium desktop, Power Macintosh, generic notebook, and Apple Newton systems.

Figure 3-1
Generic Pentium-Based Desktop PC Using PCI Bus

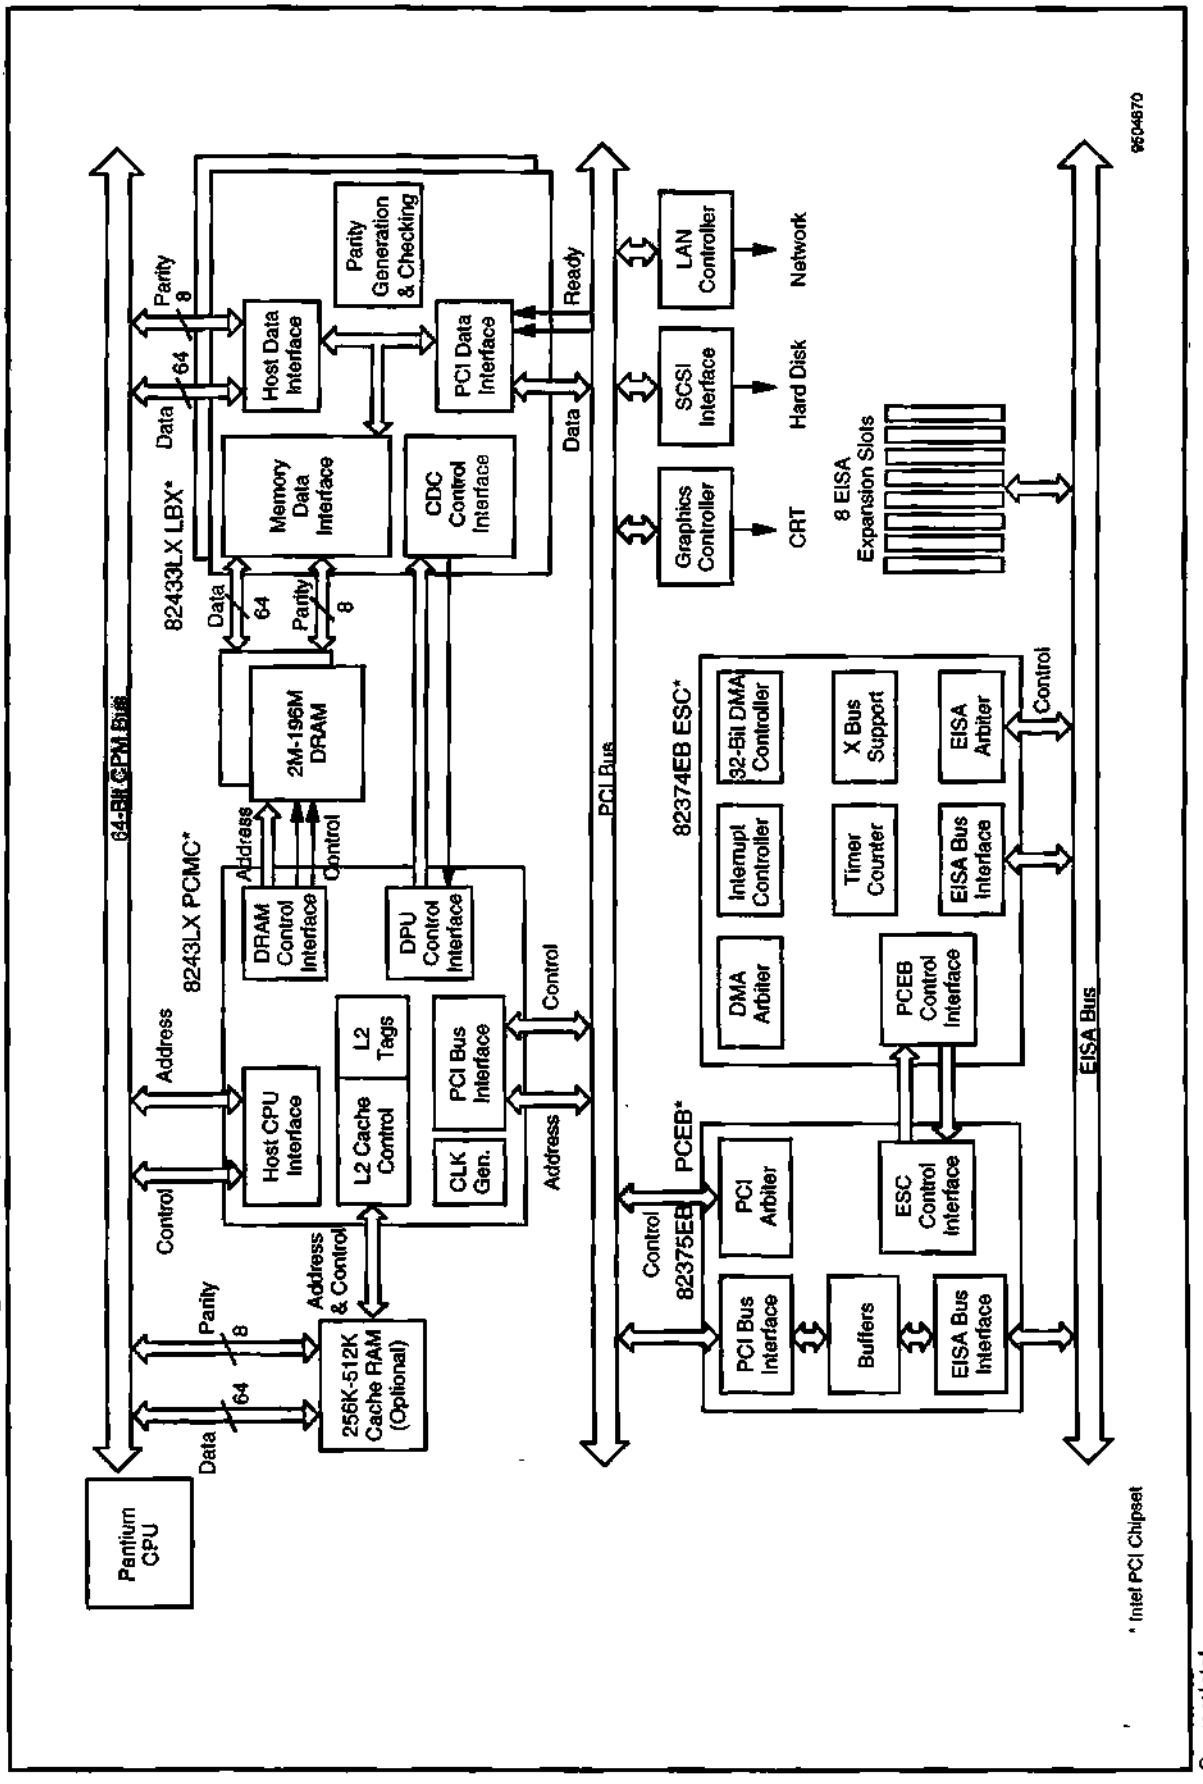

Source: Intel

Figure 3-2
Power Macintosh Hardware Architecture

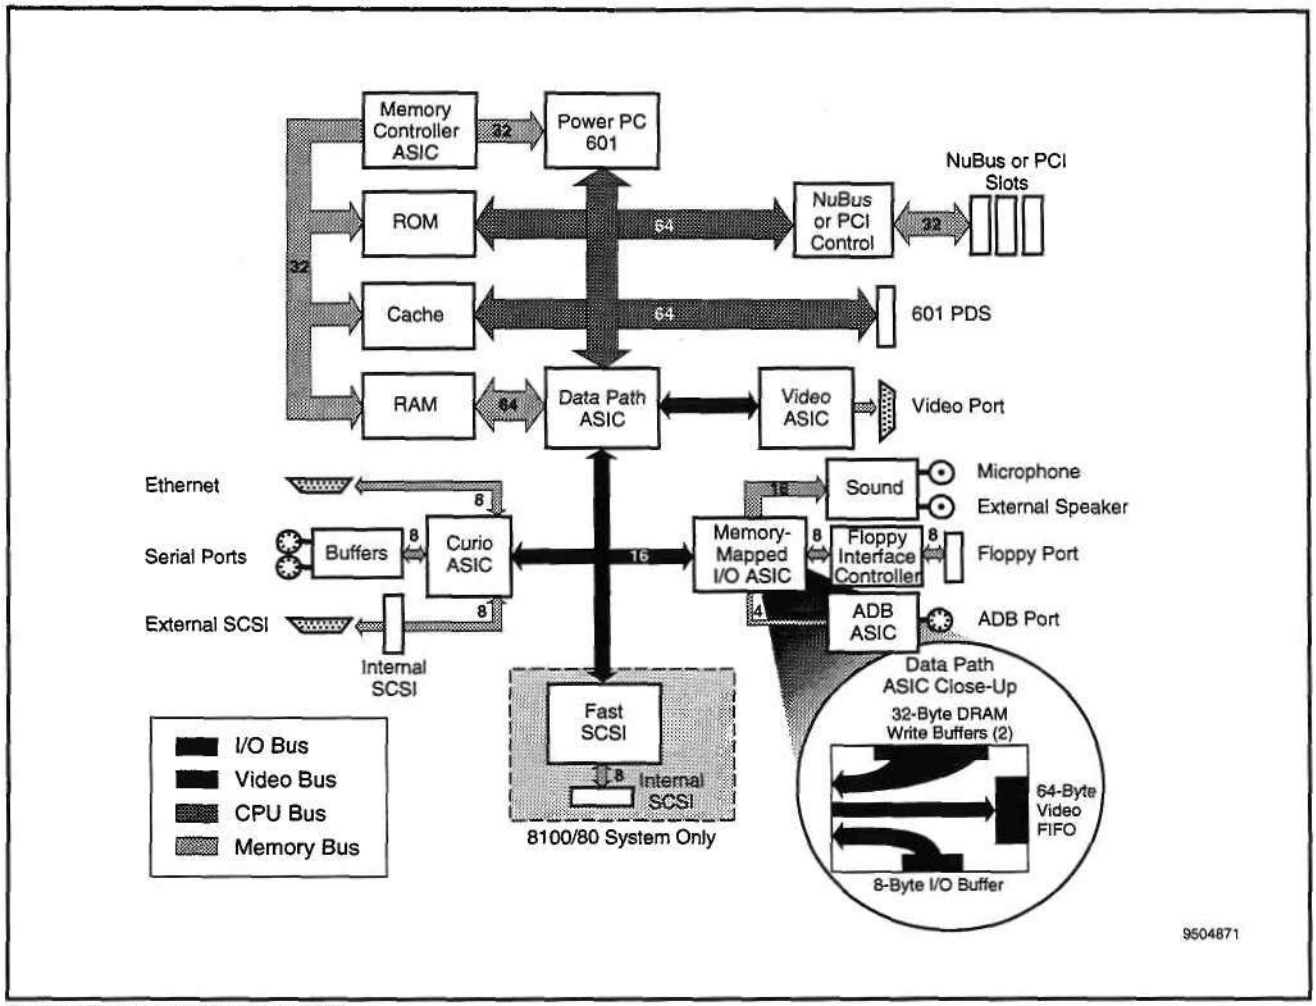

Source: Dataquest (October 1995)

Figure 3-3
Generic Notebook Block Diagram

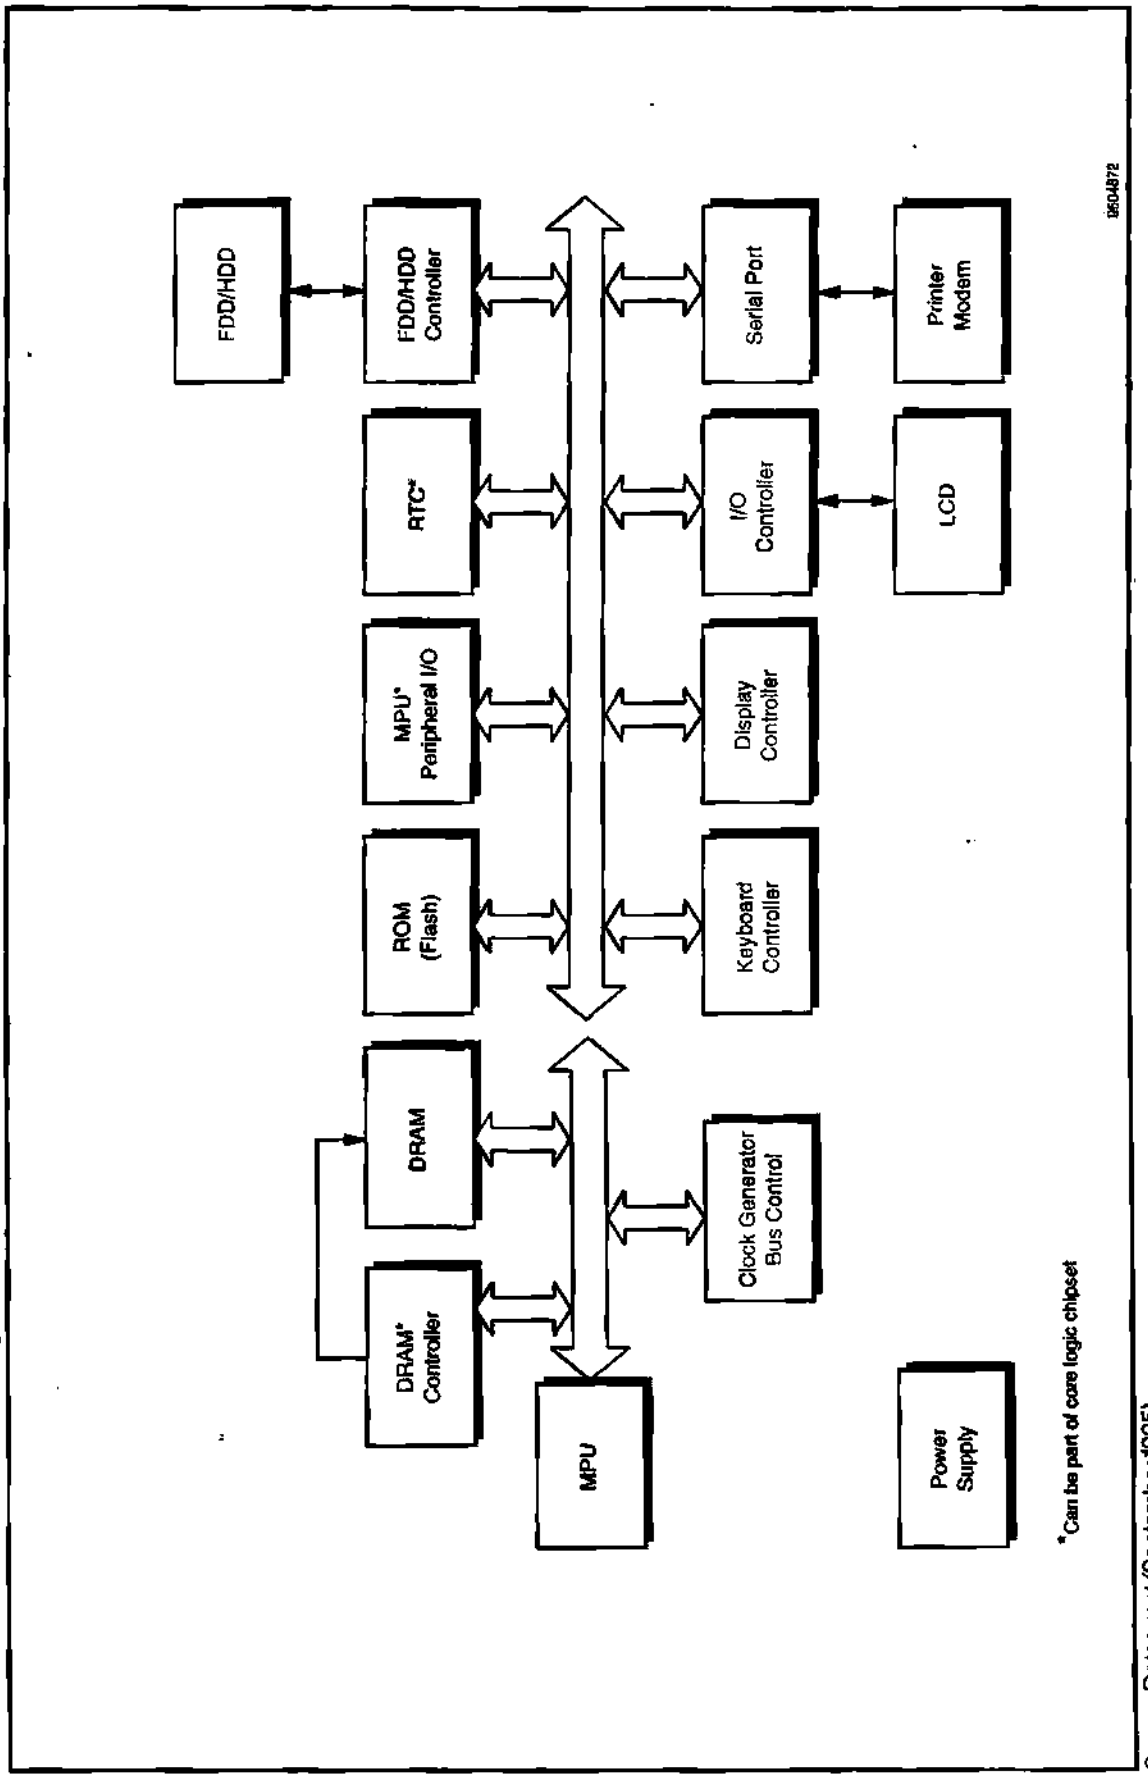

Source: Dataquest (September 1995)

Figure 3-4
Apple Newton MessagePad Block Diagram

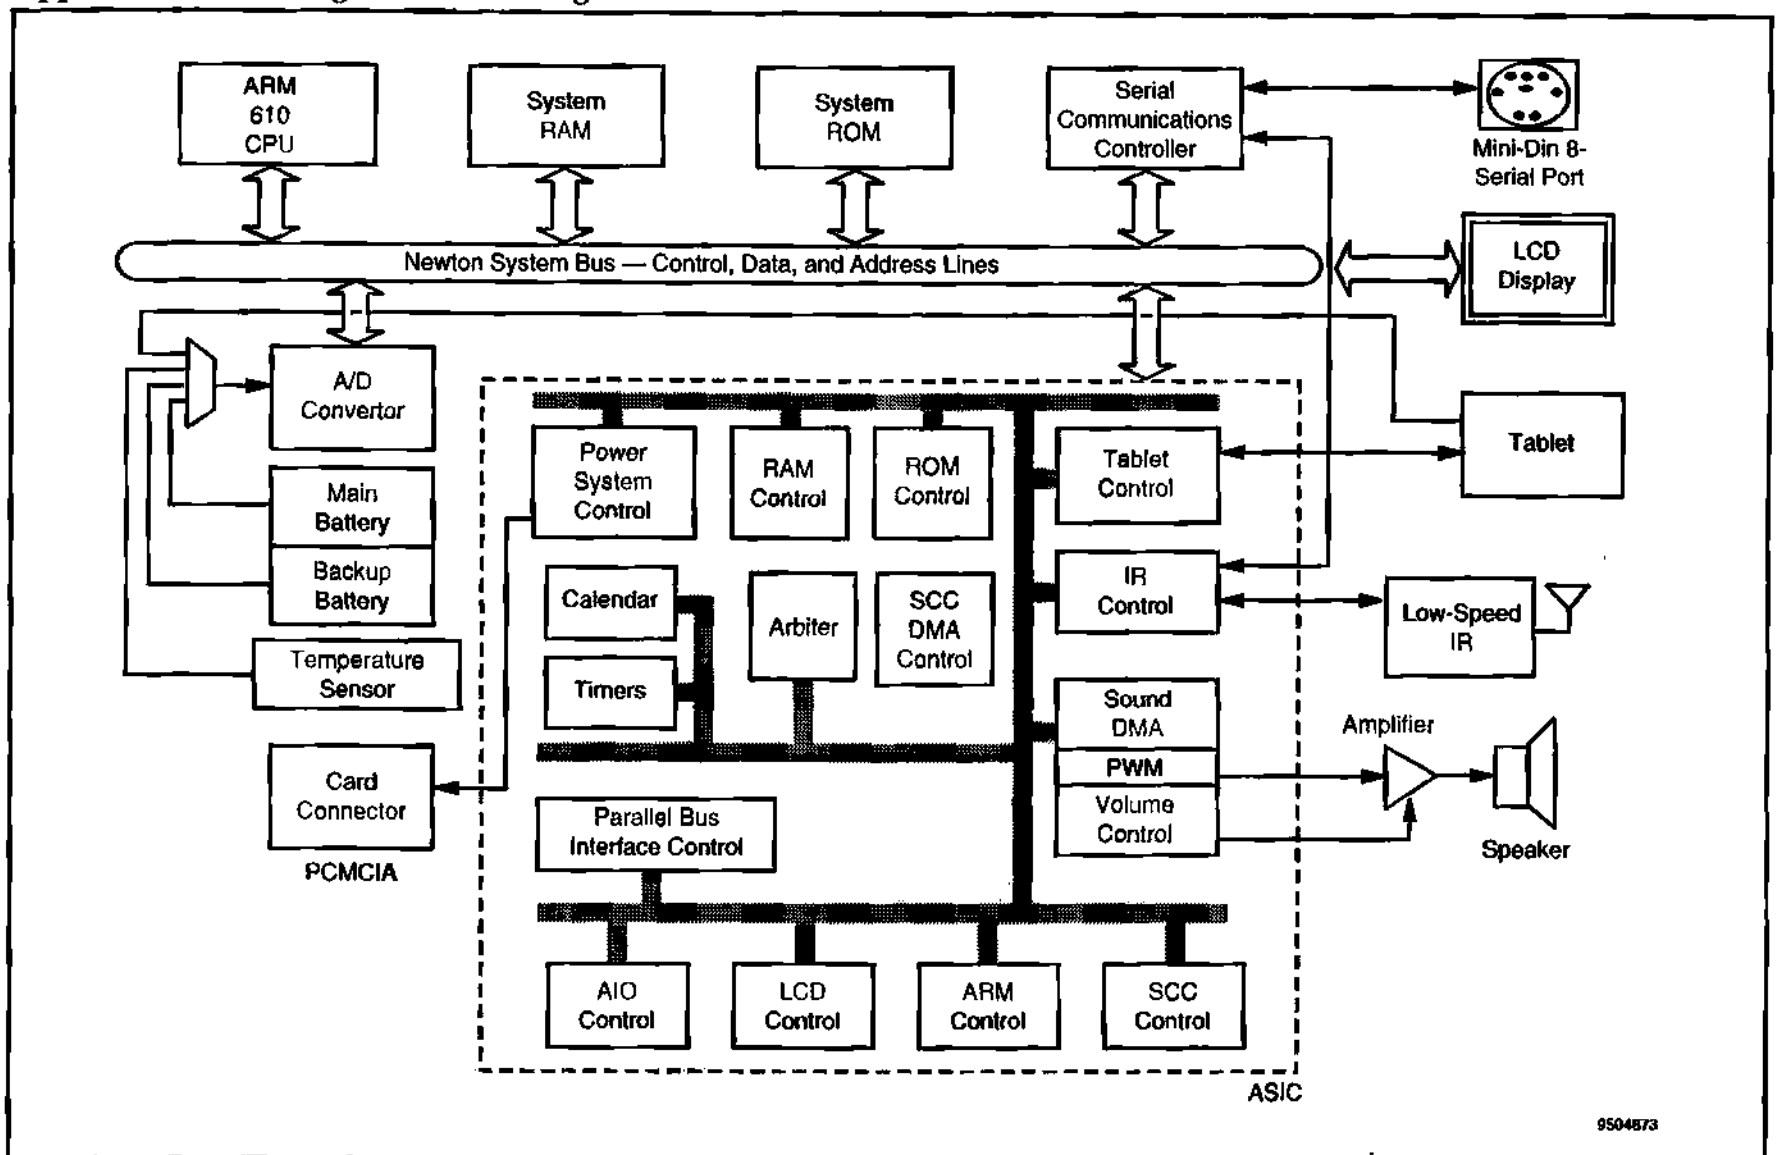

Source: Apple, PC Week

Chapter 4
Semiconductor Market Size and Forecast for PCs and Workstations

Figures 4-1 through 4-3 and Tables 4-1 through 4-2 provide forecasts and illustrations of semiconductor opportunities in the PC and workstation markets, including a focus on opportunities by region, system type, and semiconductor device type.

Figure 4-1
Worldwide PC Semiconductor Market

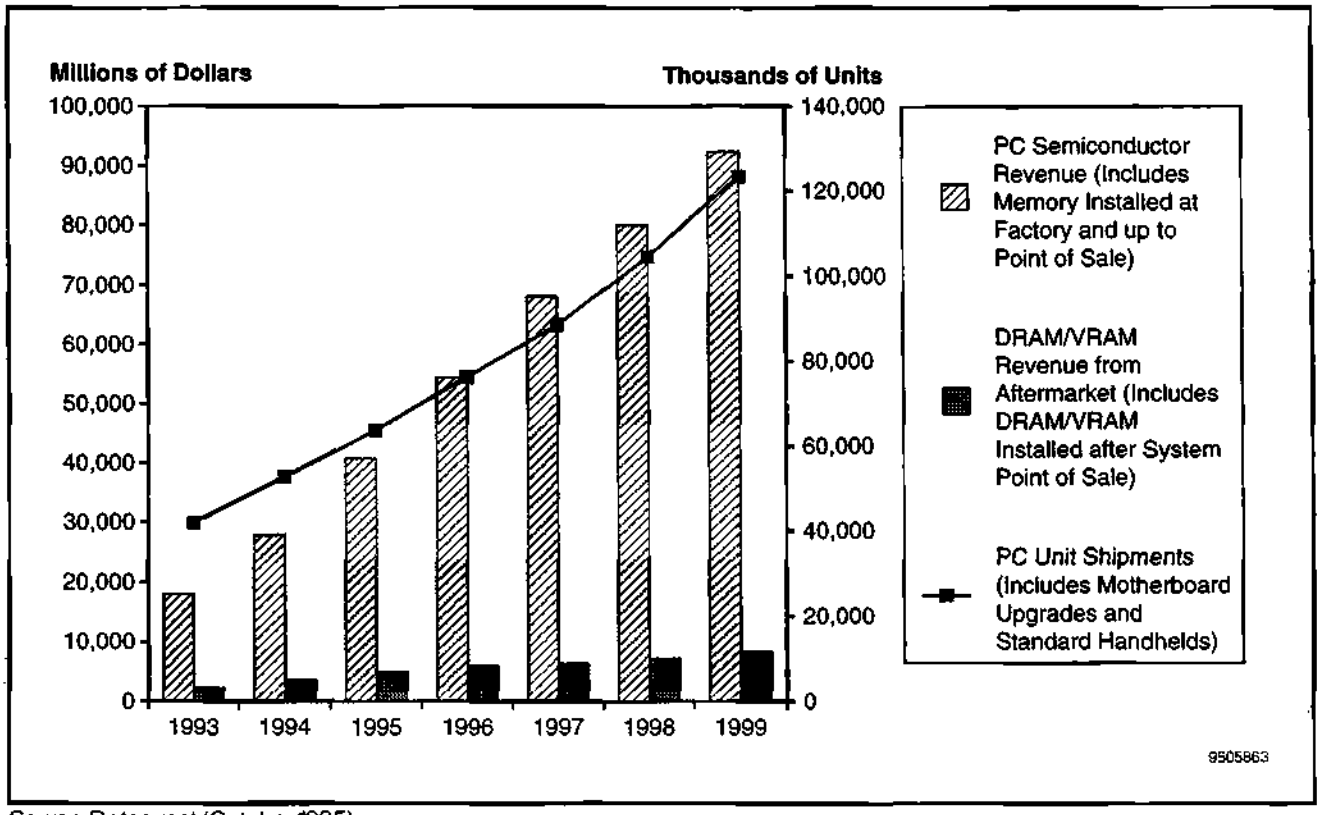

Source: Dataquest (October 1995)

Figure 4-2
PC Semiconductor Revenue by Region

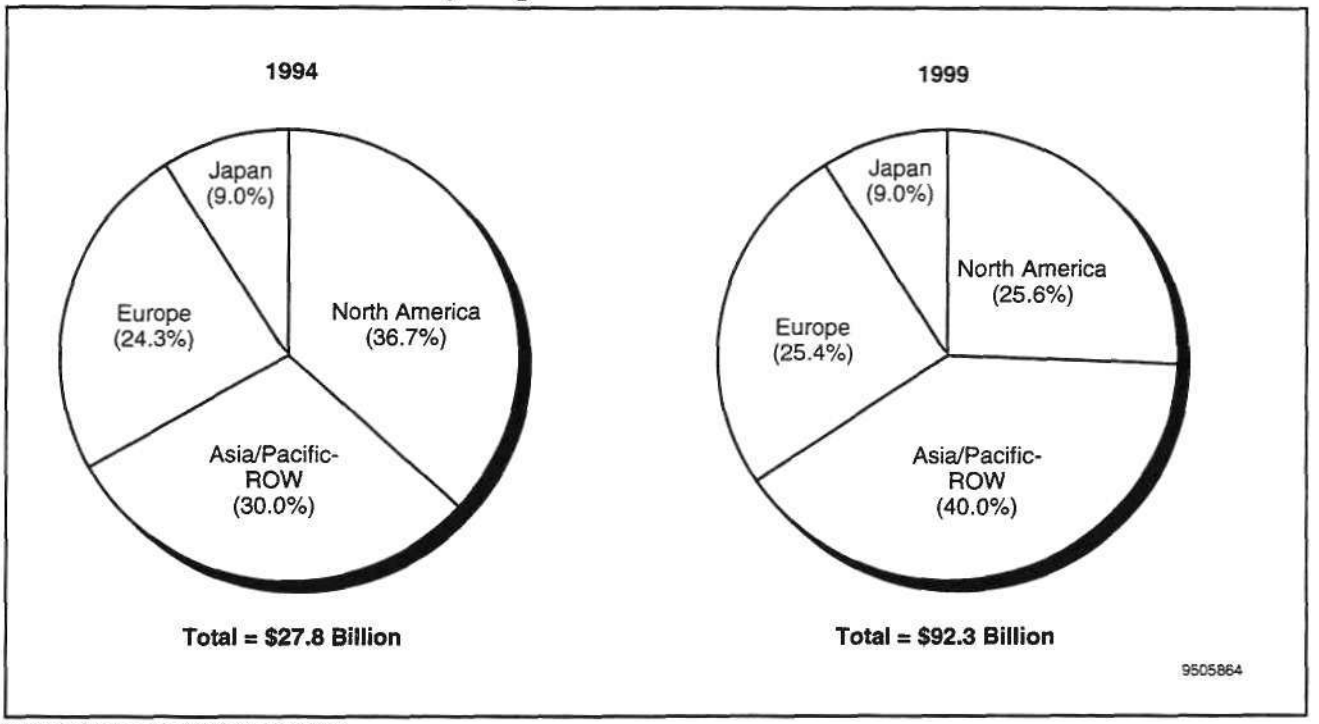

Source: Dataquest (October 1995)

Figure 4-3
Worldwide PC Semiconductor Consumption by Product Type

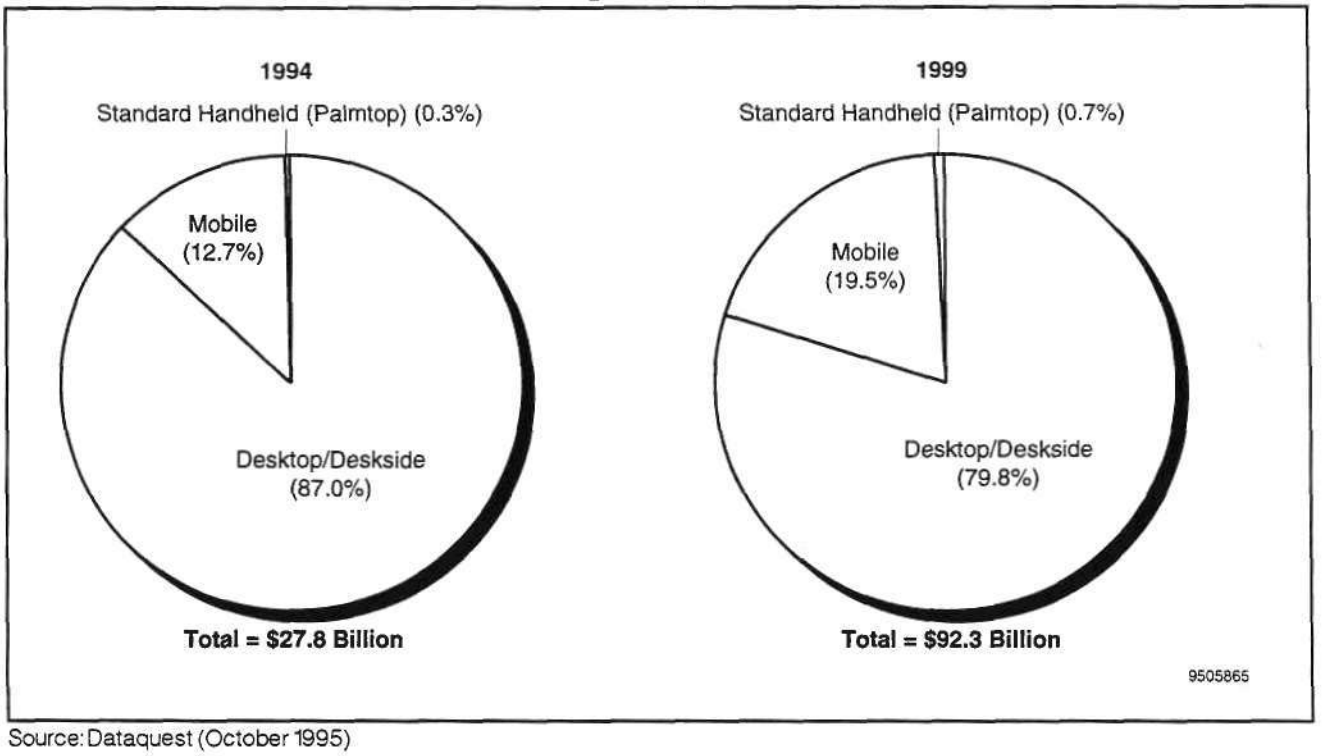

Source: Dataquest (October 1995)

Table 4-1
Worldwide PC Production and Semiconductor Market (Less Peripherals; Includes Standard Handhelds)*

	1993	1994	1995	1996	1997	1998	1999	CAGR (%) 1994-1999
Complete PC Systems								
Units (K)	39,087	48,266	58,106	69,487	80,538	94,955	112,245	18.4
Factory ASP ($)	1,441	1,541	1,605	1,573	1,515	1,427	1,348	-2.6
Factory Revenue ($M)	56,323	74,379	93,253	109,323	121,989	135,493	151,359	15.3
Semiconductor Content ($)	430	526	638	711	765	760	743	7.2
Semiconductor TAM ($M)	16,799	25,375	37,069	49,390	61,590	72,203	83,436	26.9
Desktop/Deskside PC Systems								
Units (K)	32,712	39,219	46,621	54,057	60,865	70,221	81,419	15.7
Factory ASP ($)	1.30	1.37	1.45	1.44	1.38	1.32	1.29	-1.2
Factory Revenue ($M)	42,456	53,719	67,740	77,829	84,118	93,023	104,806	14.3
Semiconductor Content ($)	444	555	669	731	806	807	796	7.5
Semiconductor TAM ($M)	14,527	21,776	31,212	39,490	49,069	56,666	64,806	24.4
Mobile PC Systems								
Units (K)	6,376	9,047	11,485	15,430	19,673	24,734	30,826	27.8
Factory ASP ($)	2.17	2.28	2.22	2.04	1.92	1.72	1.51	-7.9
Factory Revenue ($M)	13,867	20,659	25,512	31,491	37,869	42,468	46,553	17.6
Semiconductor Content ($)	356	398	510	642	636	628	604	8.7
Semiconductor TAM ($M)	2,272	3,598	5,856	9,899	12,521	15,536	18,630	38.9
Motherboard Upgrades								
Units (K)	2,759	4,404	5,476	6,738	7,992	9,582	11,190	20.5
Factory ASP ($)	1,250	1,320	1,401	1,389	1,333	1,277	1,241	-1.2
Factory Revenue ($M)	3,449	5,813	7,672	9,359	10,653	12,236	13,887	19.0
Semiconductor Content ($)	444	555	669	731	806	807	796	7.5
Semiconductor TAM ($M)	1,225	2,445	3,666	4,922	6,443	7,732	8,907	29.5
Total Production								
Units (K)	41,846	52,670	63,582	76,225	88,530	104,537	123,435	18.6
Factory ASP ($)	1,428	1,523	1,587	1,557	1,498	1,413	1,339	-2.5
Factory Revenue ($M)	59,772	80,192	100,925	118,682	132,642	147,729	165,246	15.6
Semiconductor Content ($)	431	528	641	713	768	765	748	7.2
Semiconductor TAM ($M)	18,024	27,820	40,735	54,312	68,033	79,935	92,343	27.1

*Includes graphics, serial/parallel I/O, storage adapters, and memory SIMMS installed up to point of sale
Source: Dataquest (October 1995)

Table 4-2
North American PC Production and Semiconductor Market (Less Peripherals; Includes Motherboard Upgrades and Standard Handhelds)*

	1993	1994	1995	1996	1997	1998	1999	CAGR (%) 1994-1999
Units (K)	10,747	13,500	15,979	18,503	20,043	22,427	24,669	12.8
Factory ASP ($)	1,522	1,547	1,622	1,580	1,537	1,436	1,378	-2.3
Factory Revenue ($M)	16,357	20,885	25,918	29,235	30,806	32,205	33,994	10.2
Semiconductor Content ($)	665	756	884	960	1,005	981	959	4.9
Semiconductor TAM ($M)	7,147	10,199	14,125	17,763	20,143	22,001	23,650	18.3

*Includes graphics, serial/parallel I/O, storage adapters, and memory SIMMS installed up to point of sale
Source: Dataquest (October 1995)

Semiconductor Opportunities by System

Figures 4-4 through 4-16 and Tables 4-3 through 4-5 detail semiconductor opportunities by system.

Figure 4-4
Worldwide PC Semiconductor Market by Microprocessor Type (Including Motherboard, Graphics, Storage, Serial/Parallel I/O Adapters, and Add-On Memory Installed up to Point of Sale)

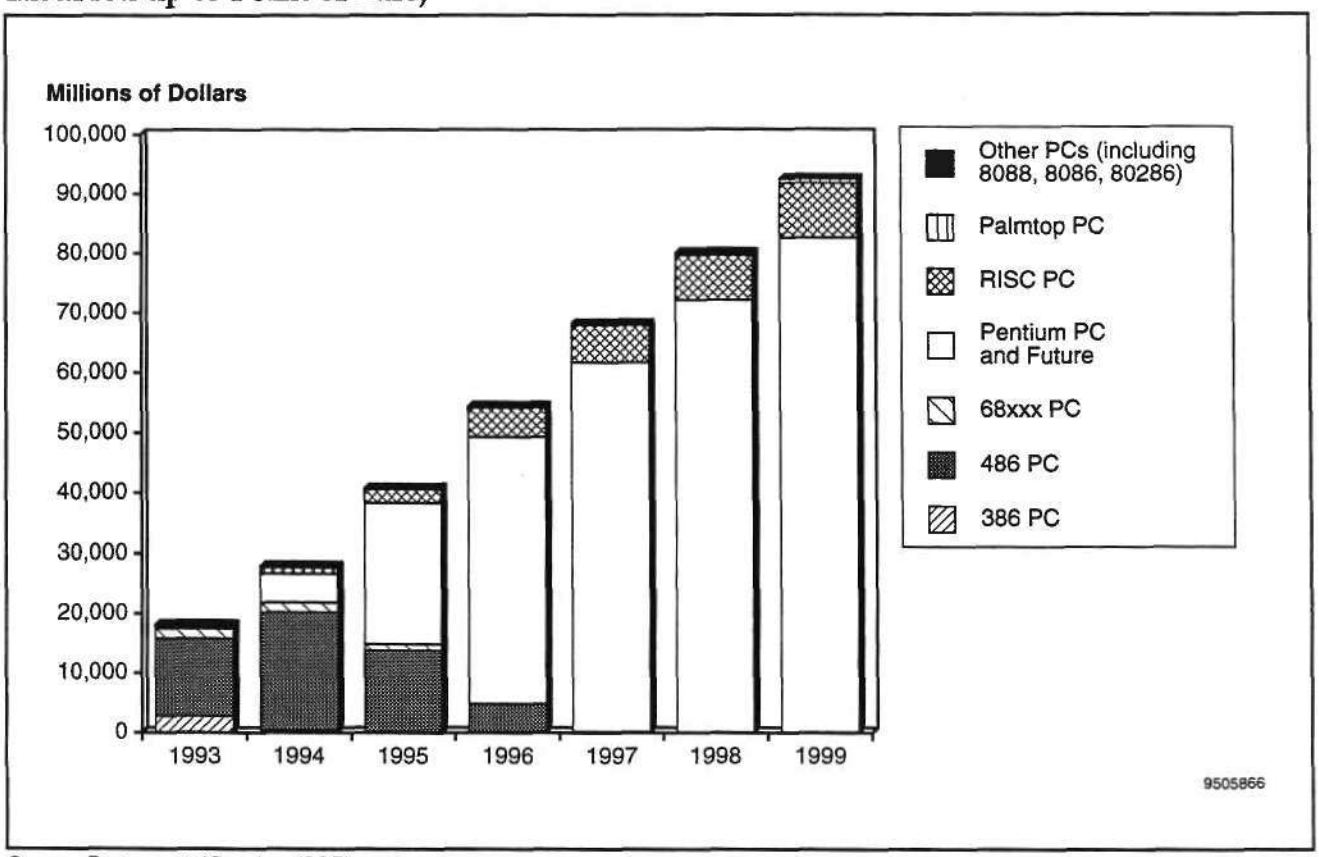

Source: Dataquest (October 1995)

Figure 4-5
Worldwide Desktop/Deskside PC Semiconductor Market by Microprocessor Type* (Including Motherboard, Graphics, Storage, Serial/Parallel I/O Adapters, and Add-On Memory Installed up to Point of Sale)

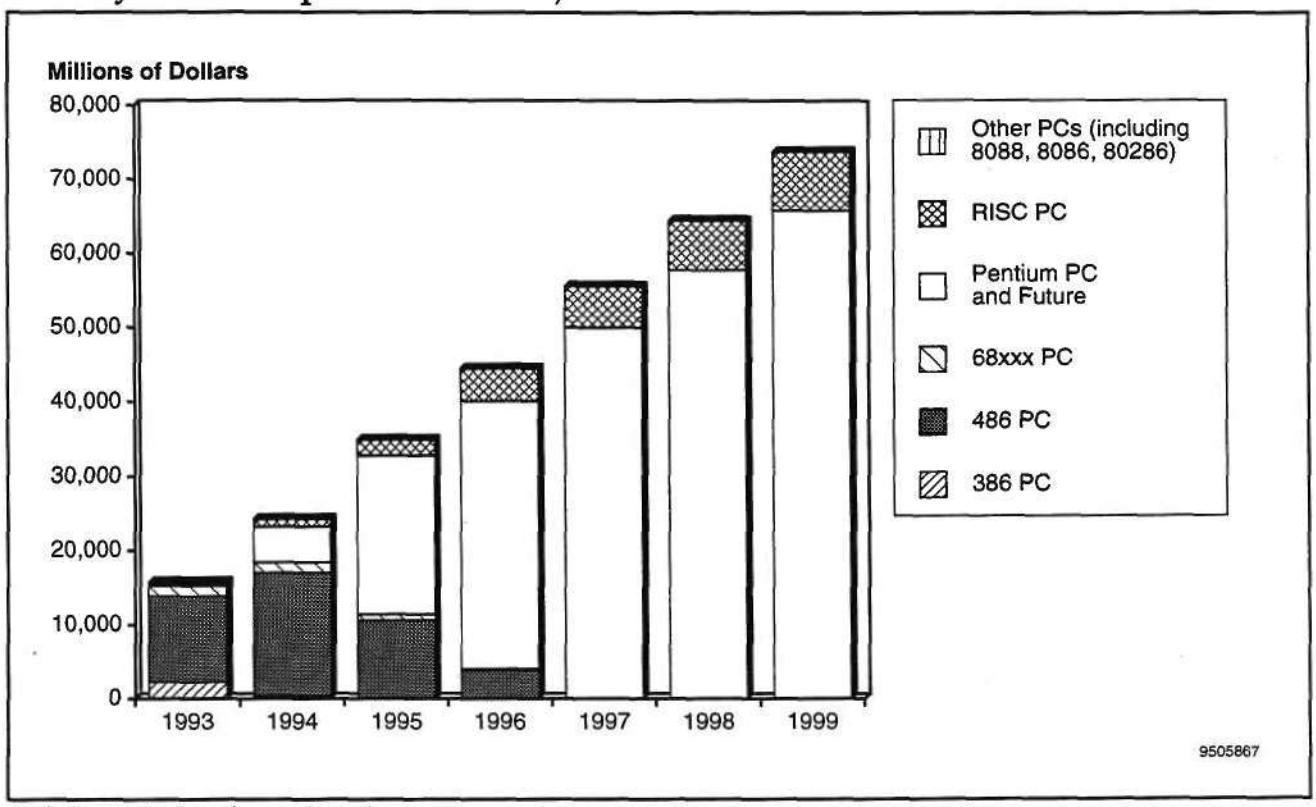

*Includes motherboard upgrade market
Source: Dataquest (October 1995)

Figure 4-6
Worldwide Mobile PC Semiconductor Market by Microprocessor Type (Including Motherboard, Graphics, Storage, Serial/Parallel I/O Adapters, and Add-On Memory Installed up to Point of Sale)

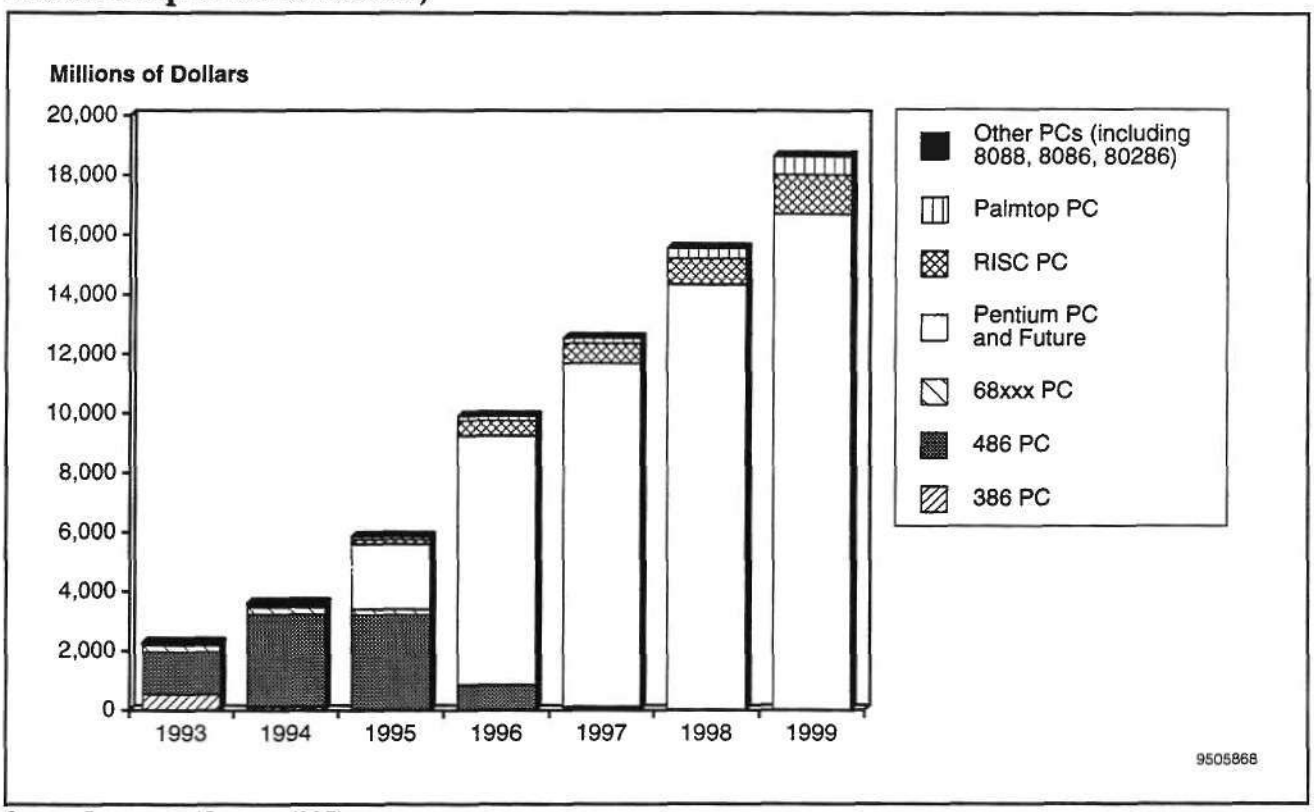

Source: Dataquest (October 1995)

Figure 4-7
Average Overall PC System Semiconductor Content (Including Motherboard, Graphics, Storage, Serial/Parallel I/O Adapters, and Add-On Memory Installed up to Point of Sale)

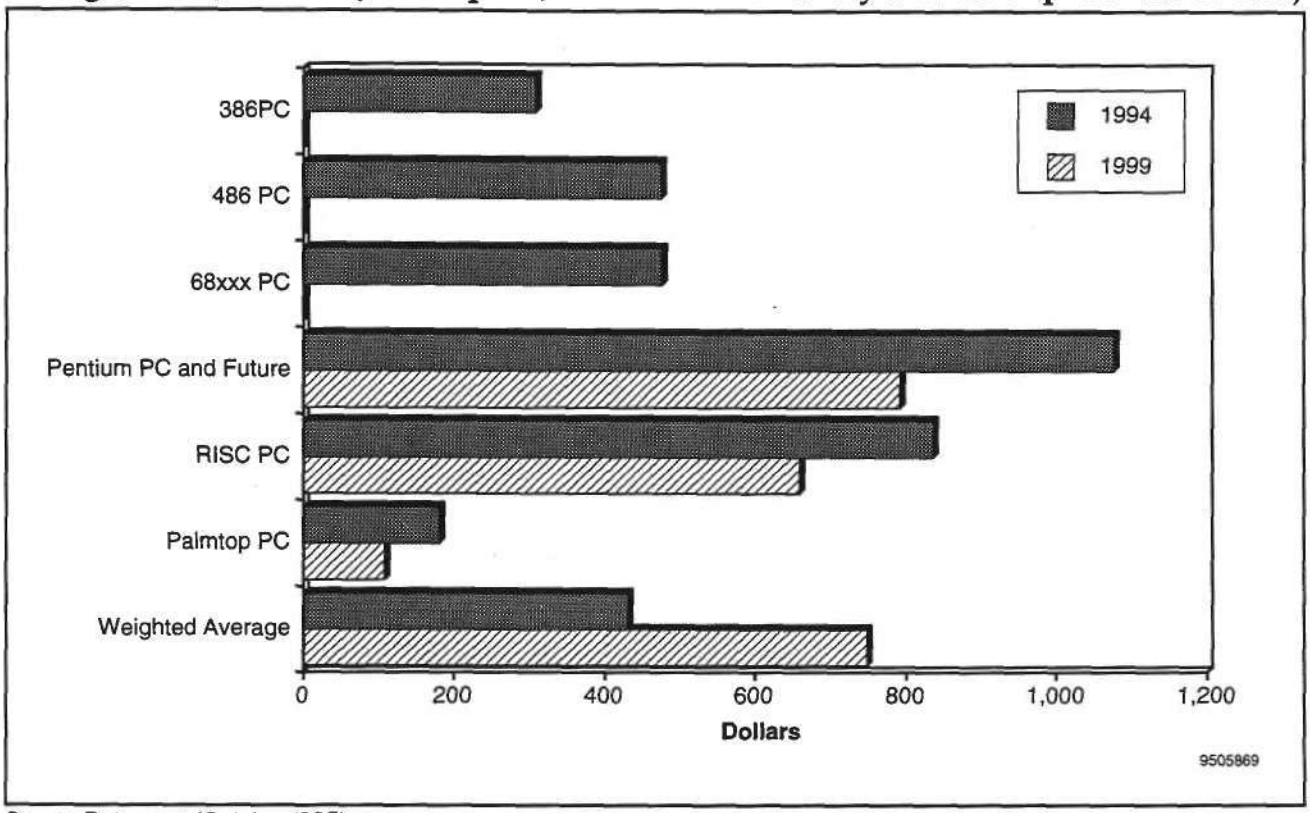

Source: Dataquest (October 1995)

Figure 4-8
Average Desktop/Deskside PC System Semiconductor Content (Including Motherboard, Graphics, Storage, Serial/Parallel I/O Adapters, and Add-On Memory Installed up to Point of Sale)

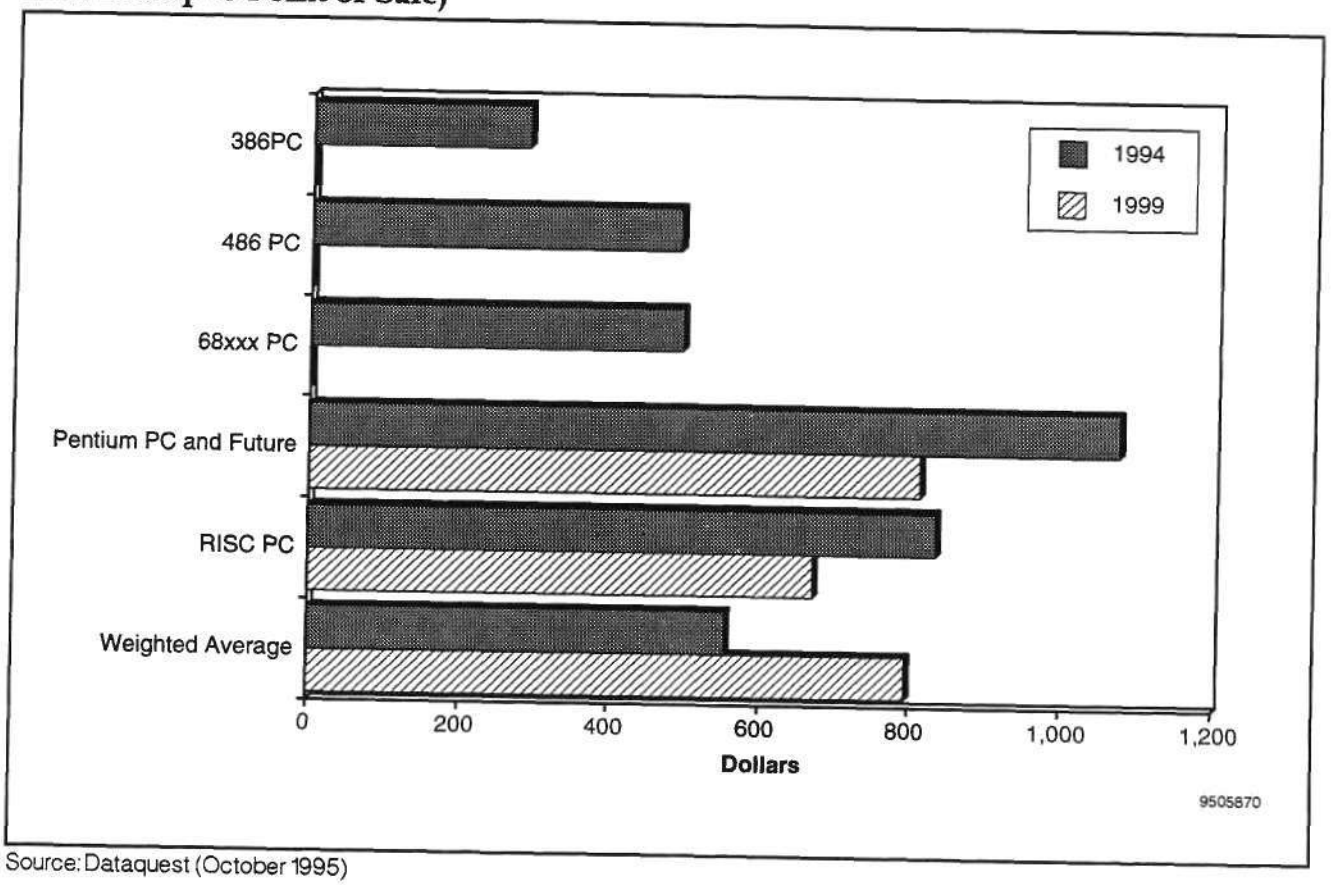

Source: Dataquest (October 1995)

Figure 4-9
Average Mobile PC System Semiconductor Content (Including Motherboard, Graphics, Storage, Serial/Parallel I/O Adapters, and Add-On Memory Installed up to Point of Sale)

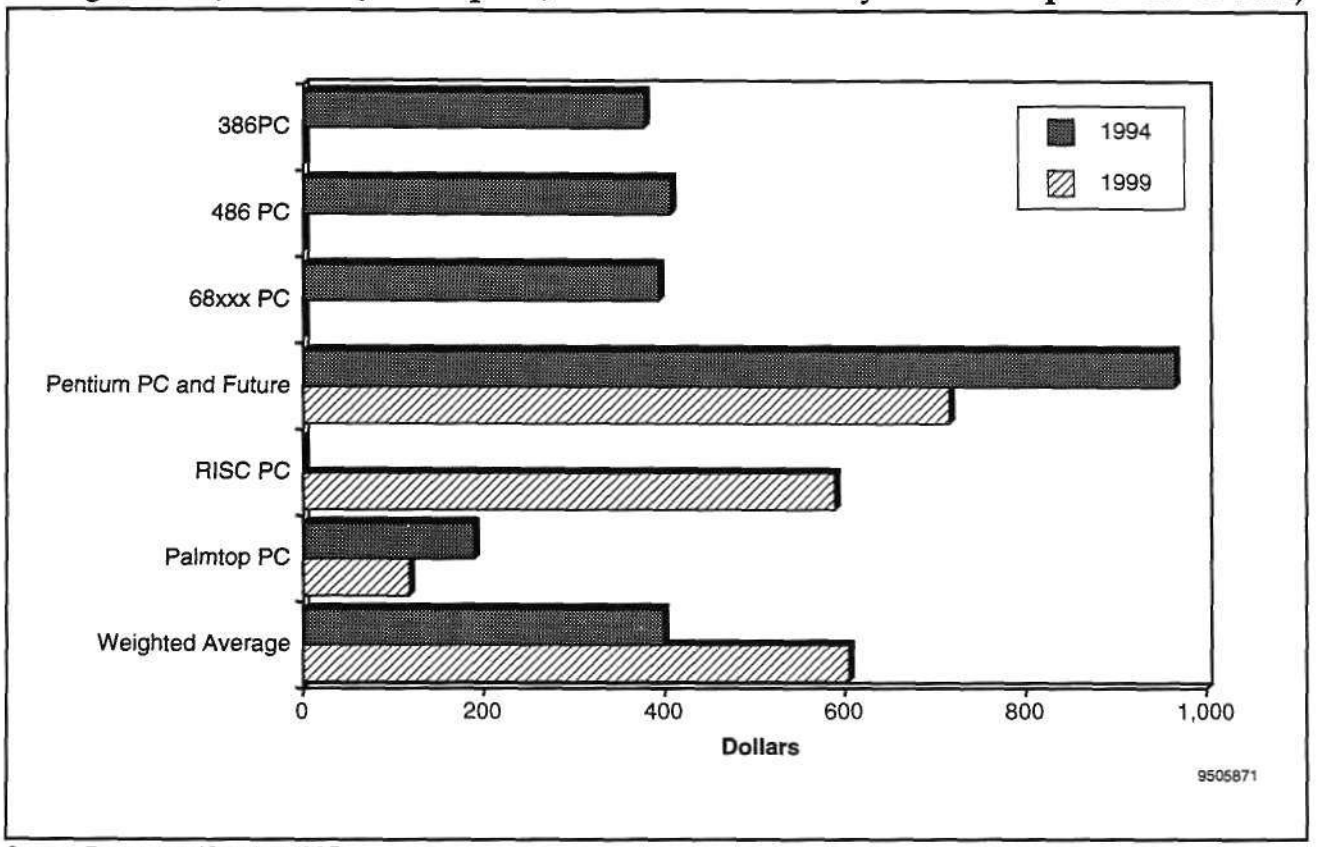

Source: Dataquest (October 1995)

Figure 4-10
Overall Pentium PC Semiconductor Content Trend (Including Motherboard, Graphics, Storage, Serial/Parallel I/O Adapters, and Add-On Memory Installed up to Point of Sale)

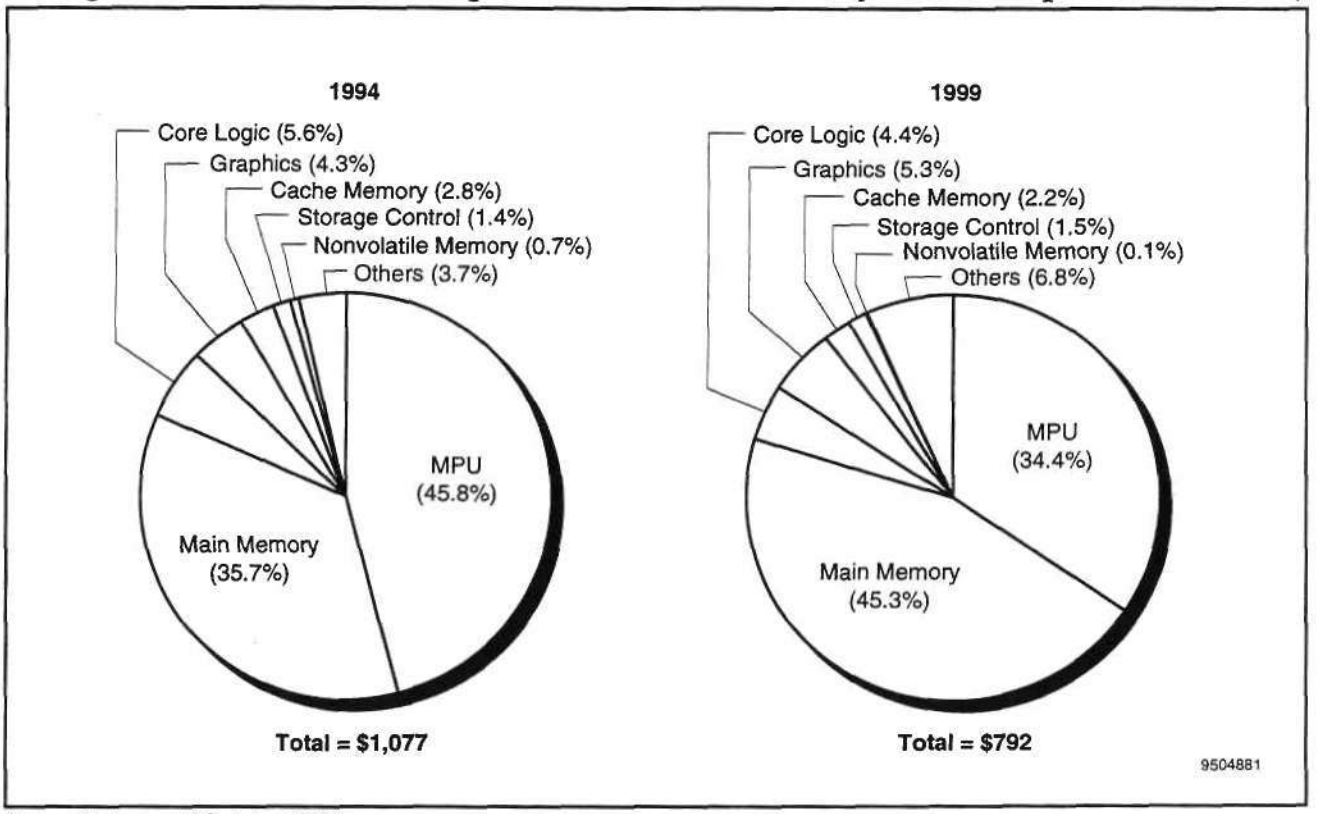

Source: Dataquest (October 1995)

Figure 4-11
Desktop/Deskside Pentium PC Semiconductor Content Trend (Including Motherboard, Graphics, Storage, Serial/Parallel I/O Adapters, and Add-On Memory Installed up to Point of Sale)

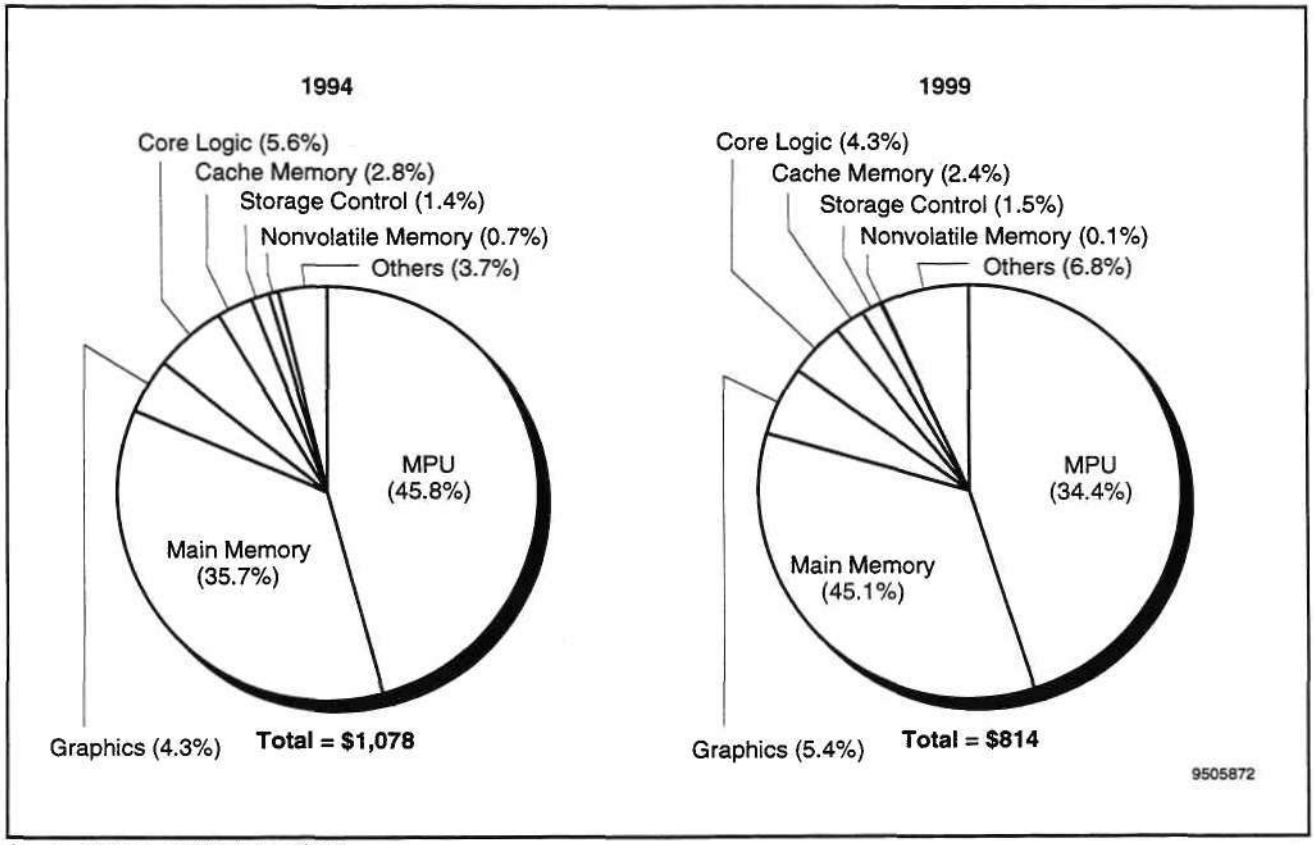

Source: Dataquest (October 1995)

Figure 4-12
Mobile Pentium PC Semiconductor Content Trend (Including Motherboard, Graphics, Storage, Serial/Parallel I/O Adapters, and Add-On Memory Installed up to Point of Sale)

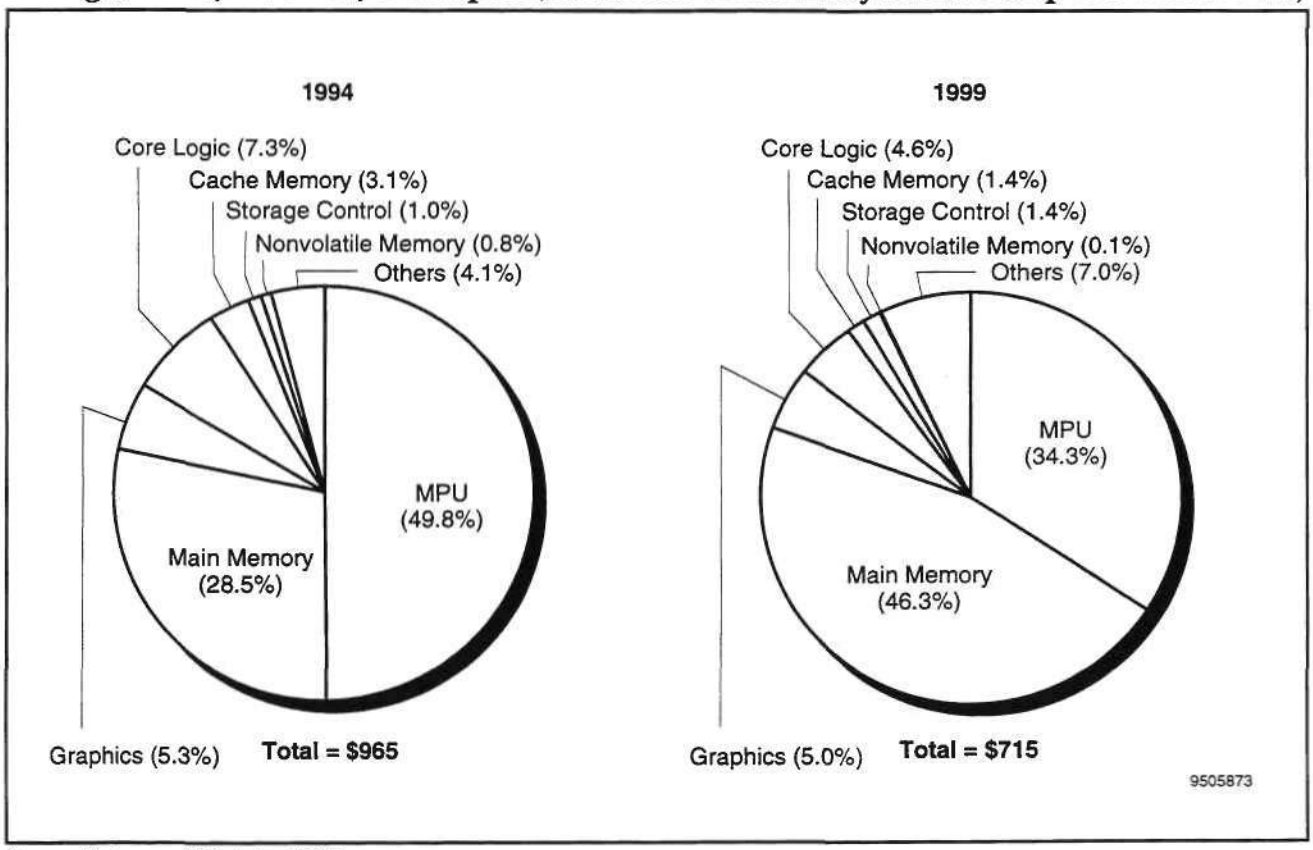

Source: Dataquest (October 1995)

Figure 4-13
Overall RISC PC Semiconductor Content Trend (Including Motherboard, Graphics, Storage, Serial/Parallel I/O Adapters, and Add-On Memory Installed up to Point of Sale)

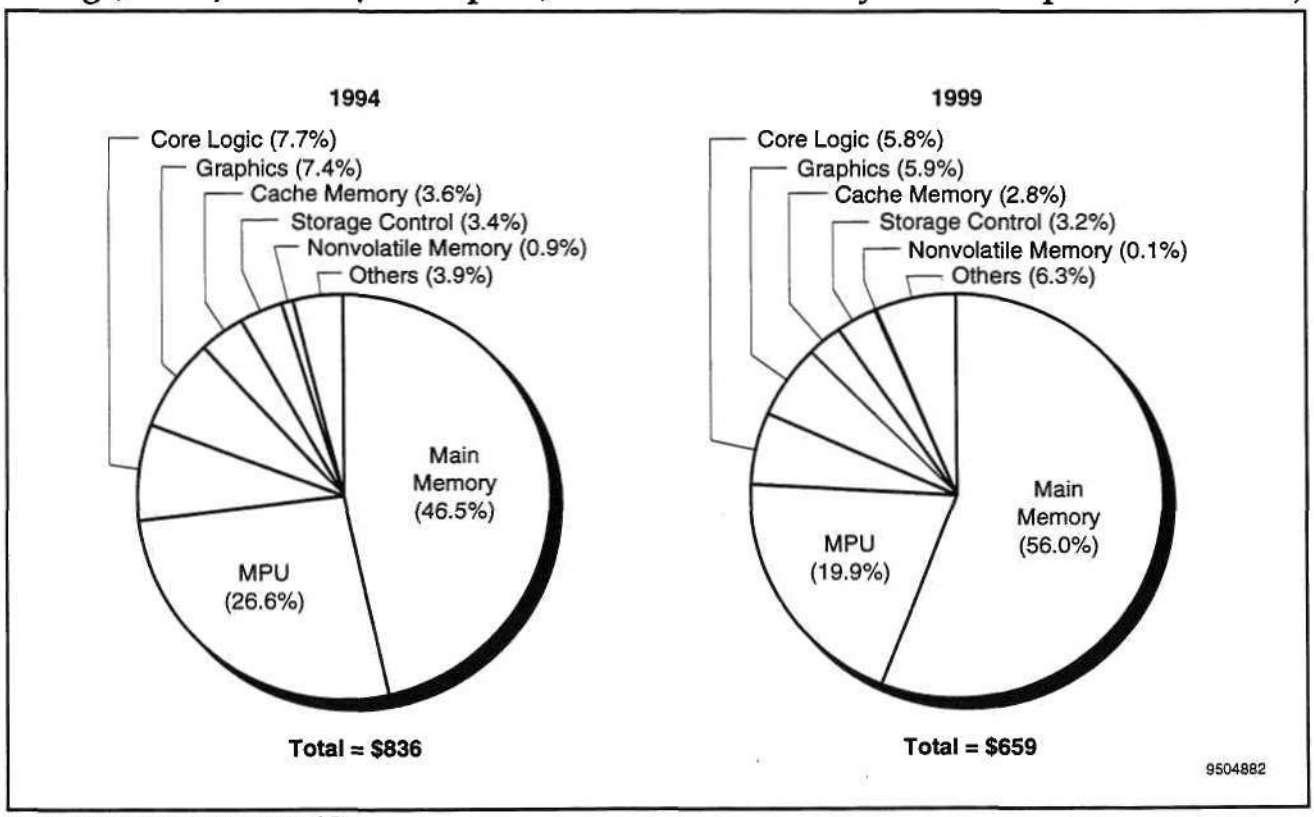

Source: Dataquest (October 1995)

Figure 4-14
Desktop/Deskside RISC PC Semiconductor Content Trend (Including Motherboard, Graphics, Storage, Serial/Parallel I/O Adapters, and Add-On Memory Installed up to Point of Sale)

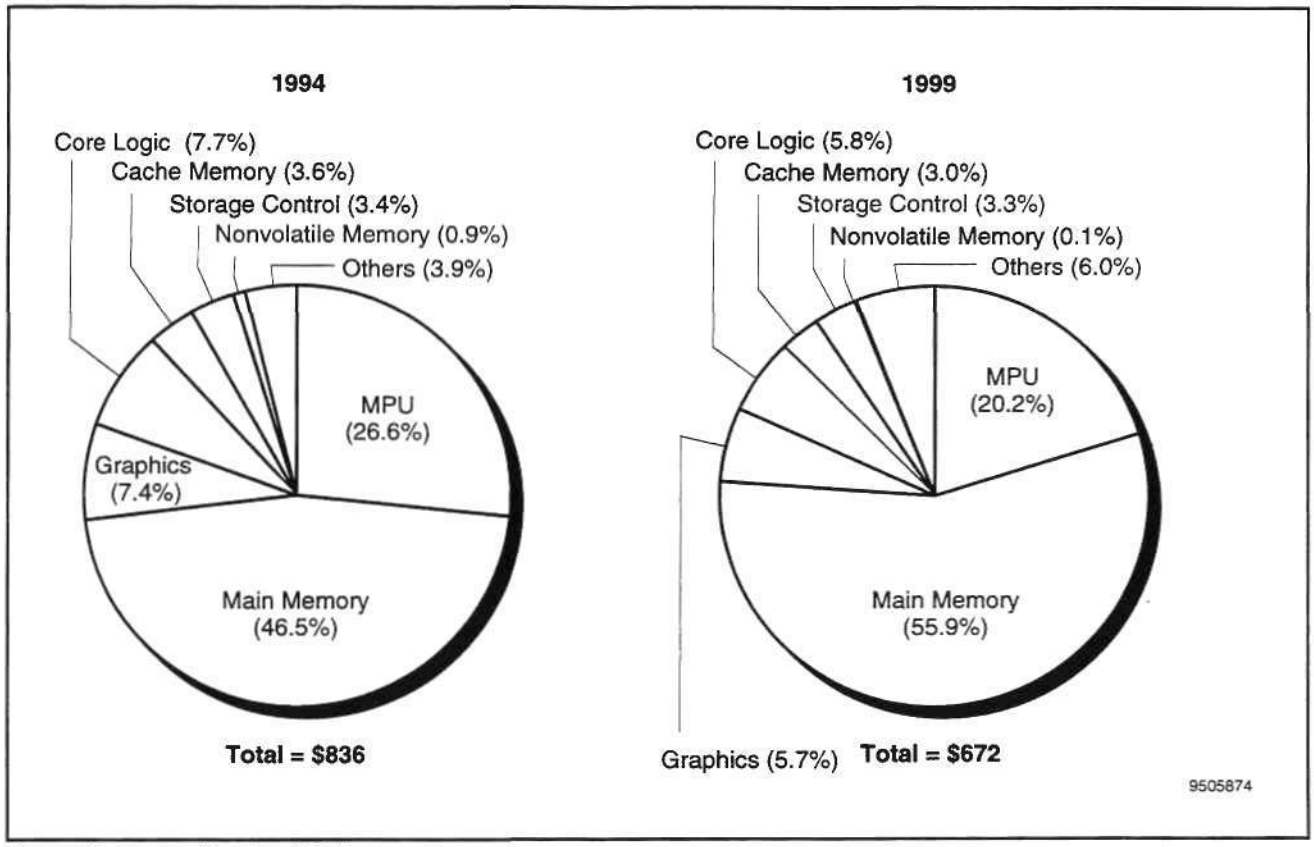

Source: Dataquest (October 1995)

Figure 4-15
Mobile RISC PC Semiconductor Content Trend (Including Motherboard, Graphics, Storage, Serial/Parallel I/O Adapters, and Add-On Memory Installed up to Point of Sale)

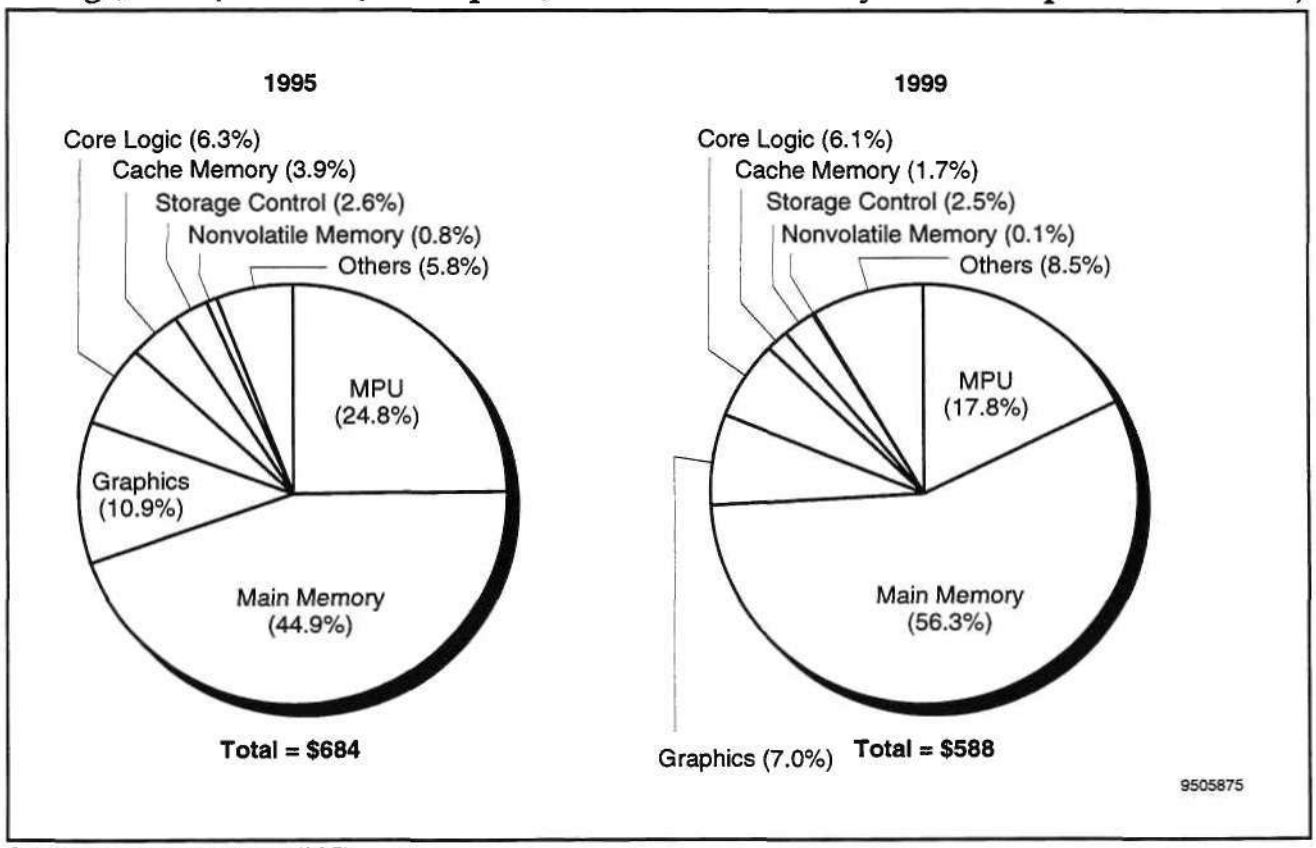

Source: Dataquest (October 1995)

Figure 4-16
Standard Handheld (Palmtop) Semiconductor Content Trend (Including Motherboard, Graphics, Storage, Serial/Parallel I/O Adapters, and Add-On Memory Installed up to Point of Sale)

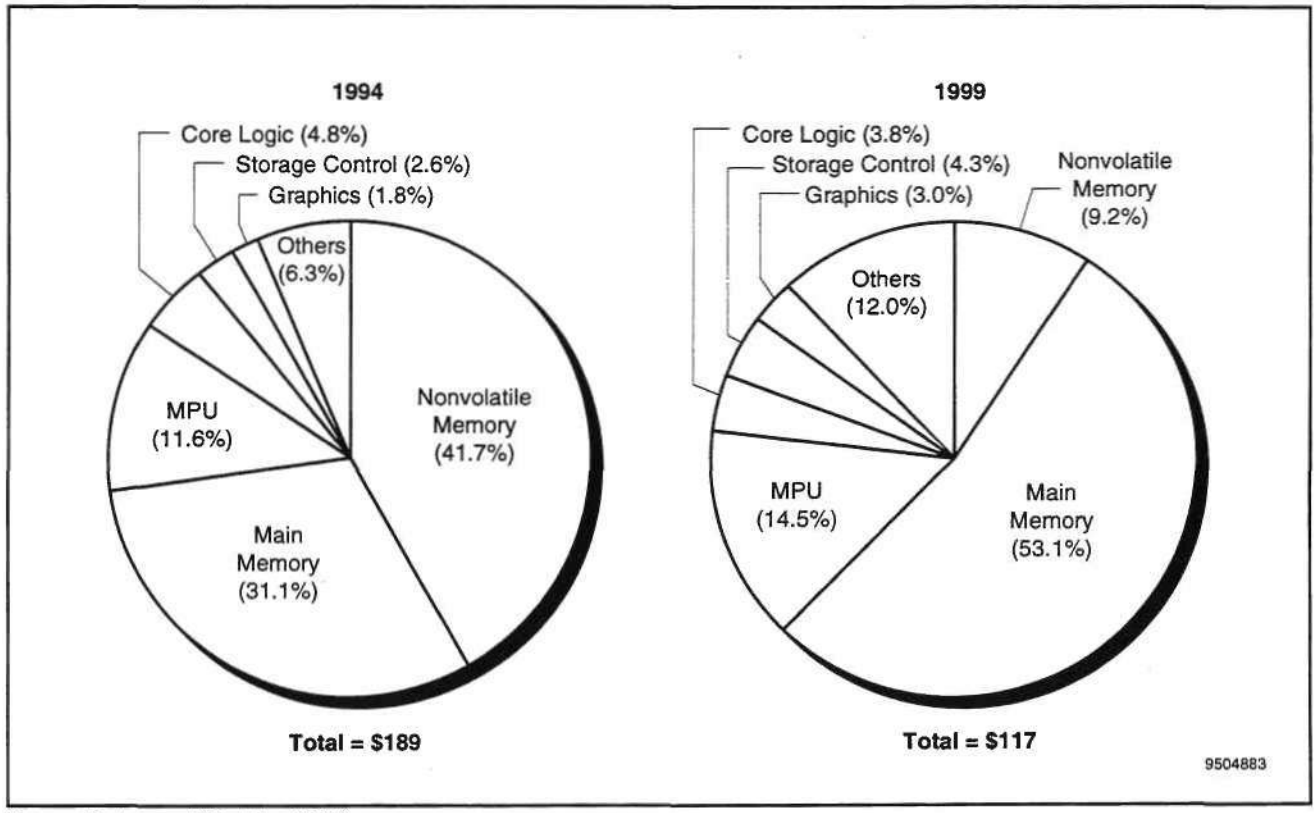

Source: Dataquest (October 1995)

Table 4-3
Worldwide PC Semiconductor Market by System* (Millions of Dollars) (Includes Motherboard, Graphics, Storage, Serial/Parallel I/O Adapters, and Add-on Memory Installed up to Point of Sale)

	1993	1994	1995	1996	1997	1998	1999	CAGR (%) 1994-1999
386 PC	2,832	442	13	0	0	0	0	-100.0
486 PC	13,006	19,845	13,820	4,698	153	0	0	-100.0
68xxx PC	1,689	1,672	1,046	208	59	0	0	-100.0
Pentium PC and Future	341	4,794	23,543	44,469	61,510	72,119	82,565	76.7
RISC PC	57	971	2,207	4,797	6,136	7,474	9,164	56.7
Palmtop PC	38	74	106	140	176	343	613	52.7
Other PCs (Including 8088, 8086, 80286)	61	22	0	0	0	0	0	-100.0
Total	18,024	27,820	40,735	54,312	68,033	79,935	92,343	27.1
Year-to-Year Growth (%)	51	54	46	33	25	17	16	

*Includes motherboard upgrade market
Source: Dataquest (October 1995)

Table 4-4
Worldwide Desktop/Deskside PC Semiconductor Market by System* (Millions of Dollars) (Includes Motherboard, Graphics, Storage, Serial/Parallel I/O Adapters, and Add-on Memory Installed up to Point of Sale)

	1993	1994	1995	1996	1997	1998	1999	CAGR (%) 1994-1999
386 PC	2,273	308	9	0	0	0	0	-100.0
486 PC	11,599	16,743	10,607	3,835	20	0	0	-100.0
68xxx PC	1,452	1,419	852	208	59	0	0	-100.0
Pentium PC and Future	341	4,772	21,340	36,079	49,953	57,792	65,862	69.0
RISC PC	56	971	2,069	4,291	5,481	6,606	7,851	51.9
Other PCs (Including 8088, 8086, 80286)	31	10	0	0	0	0	0	-100.0
Total	15,752	24,222	34,879	44,413	55,512	64,399	73,712	24.9
Year-to-Year Growth (%)	50	54	44	27	25	16	14	

*Includes motherboard upgrade market
Source: Dataquest (October 1995)

Table 4-5
Worldwide Mobile PC Semiconductor Market by System (Millions of Dollars) (Includes Motherboard, Graphics, Storage, Serial/Parallel I/O Adapters, and Add-on Memory Installed up to Point of Sale)

	1993	1994	1995	1996	1997	1998	1999	CAGR (%) 1994-1999
386 PC	559	135	4	0	0	0	0	-100.0
486 PC	1,408	3,102	3,213	863	134	0	0	-100.0
68xxx PC	237	253	194	0	0	0	0	-100.0
Pentium PC and Future	0	22	2,203	8,390	11,557	14,326	16,703	276.6
RISC PC	1	0	138	506	655	867	1,314	-
Palmtop PC	38	74	106	140	176	343	613	52.7
Other PCs (Including 8088, 8086, 80286)	30	12	0	0	0	0	0	-100.0
Total	2,272	3,598	5,856	9,899	12,521	15,536	18,630	38.9
Year-to-Year Growth (%)	58	58	63	69	26	24	20	

Source: Dataquest (October 1995)

Semiconductor Opportunities by Device Type

Figures 4-17 through 4-27 and Tables 4-6 through 4-22 detail semiconductor opportunities by device type.

Figure 4-17
Worldwide Total PC Semiconductor Market by Device Type (Including Motherboard, Graphics, Storage, Serial/Parallel I/O Adapters, and Add-On Memory Installed up to Point of Sale)

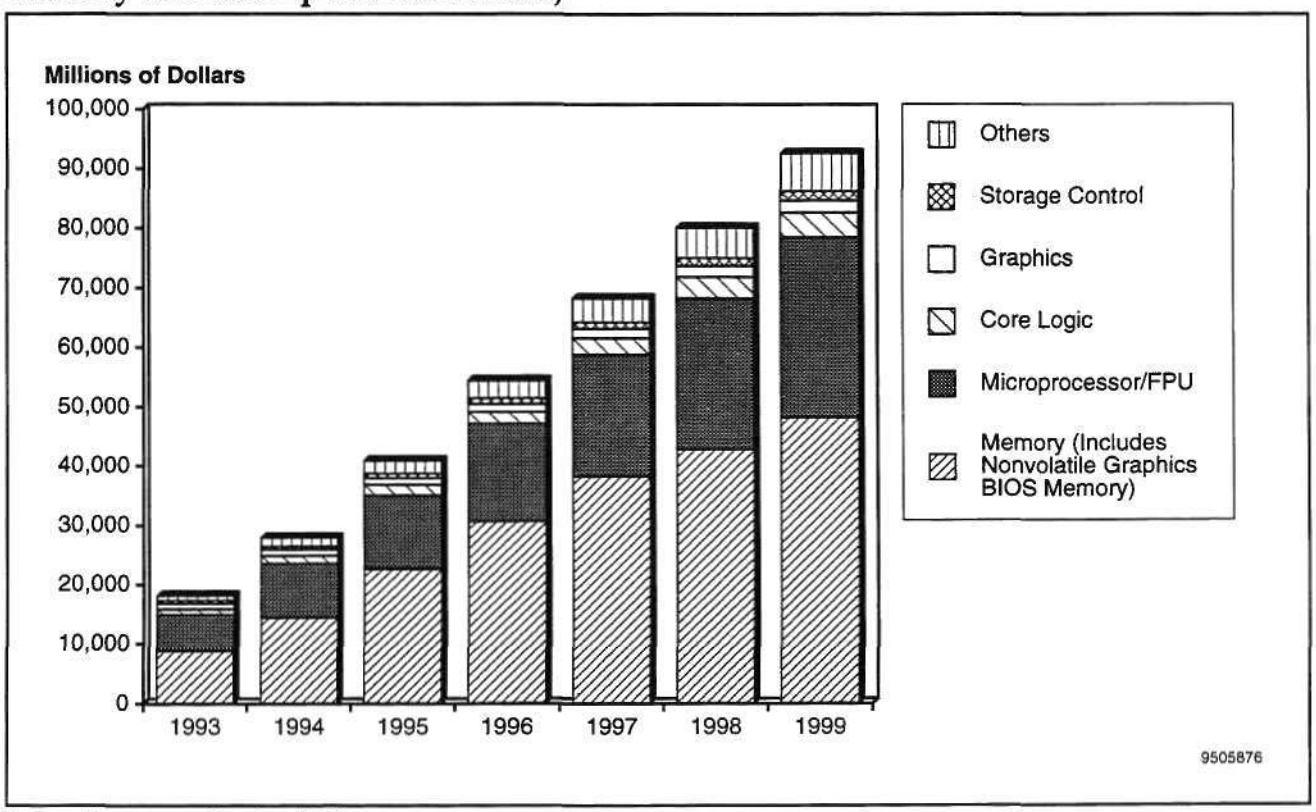

Source: Dataquest (October 1995)

Figure 4-18
Worldwide Desktop/Deskside PC Semiconductor Market by Device Type* (Including Motherboard, Graphics, Storage, Serial/Parallel I/O Adapters, and Add-On Memory Installed up to Point of Sale)

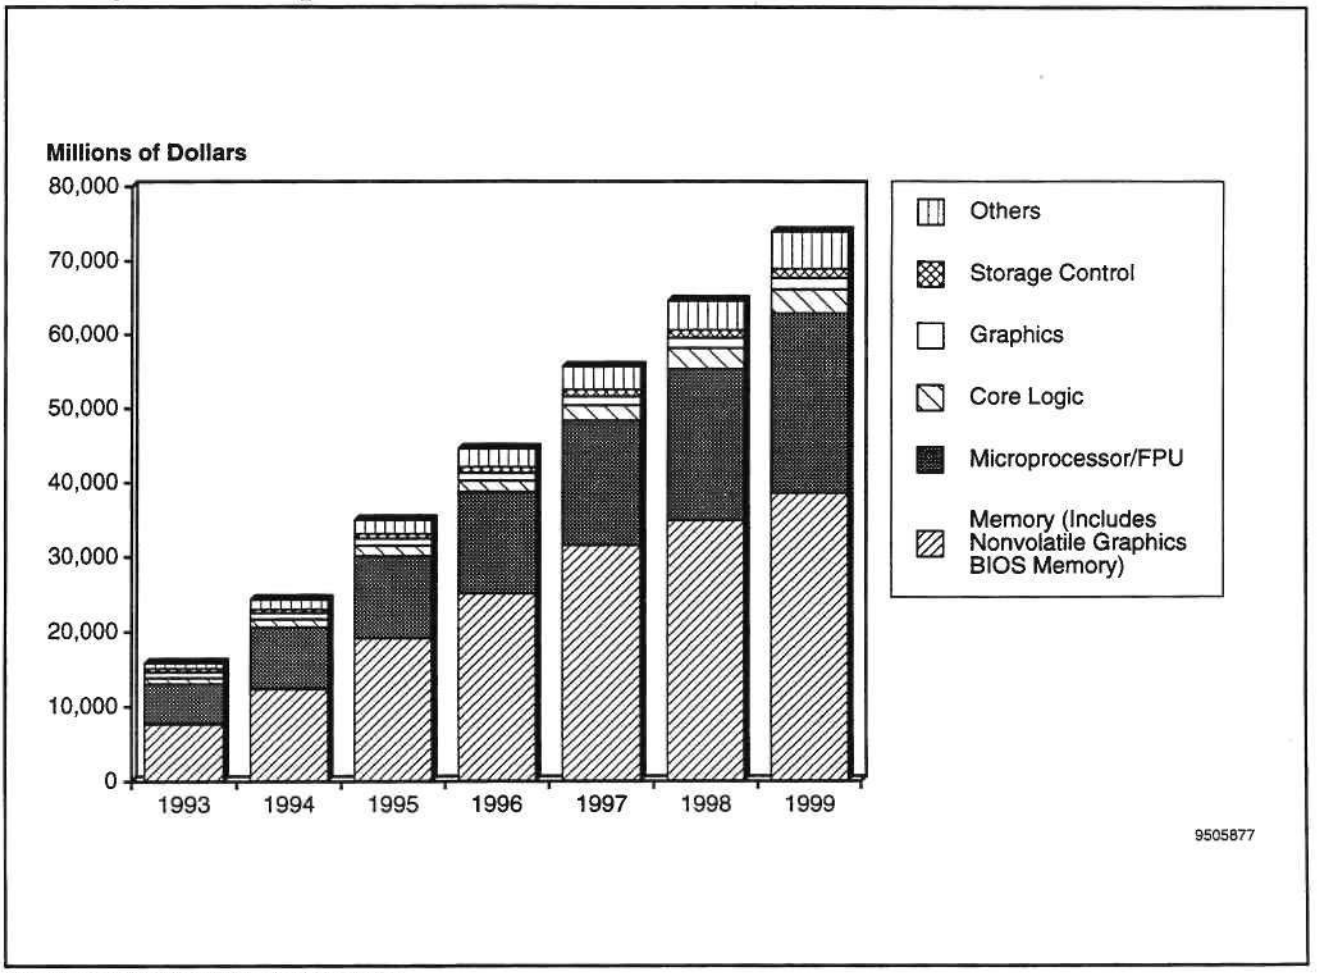

*Includes motherboard upgrade market
Source: Dataquest (October 1995)

Figure 4-19
Worldwide Mobile PC Semiconductor Market by Device Type (Including Motherboard, Graphics, Storage, Serial/Parallel I/O Adapters, and Add-On Memory Installed up to Point of Sale)

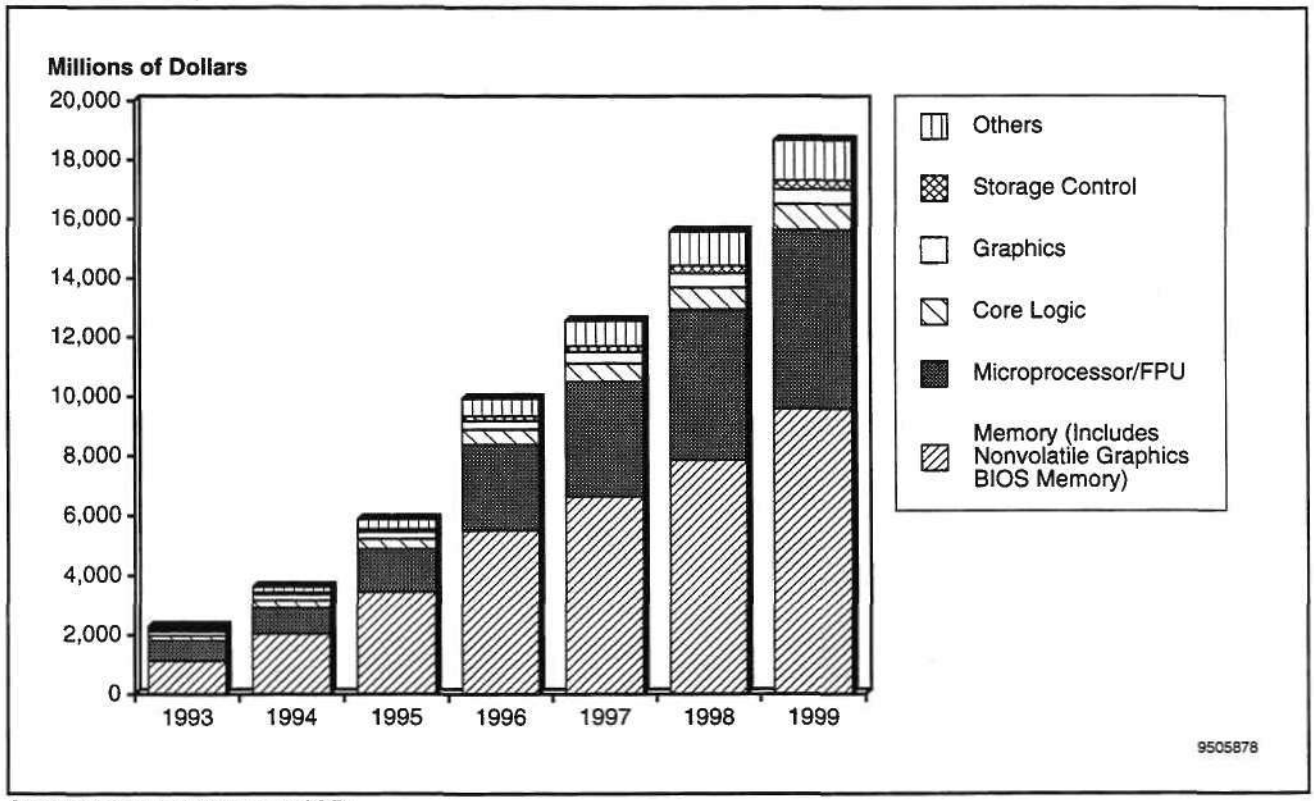

Source: Dataquest (October 1995)

Figure 4-20
Worldwide Microprocessor Opportunities in PCs

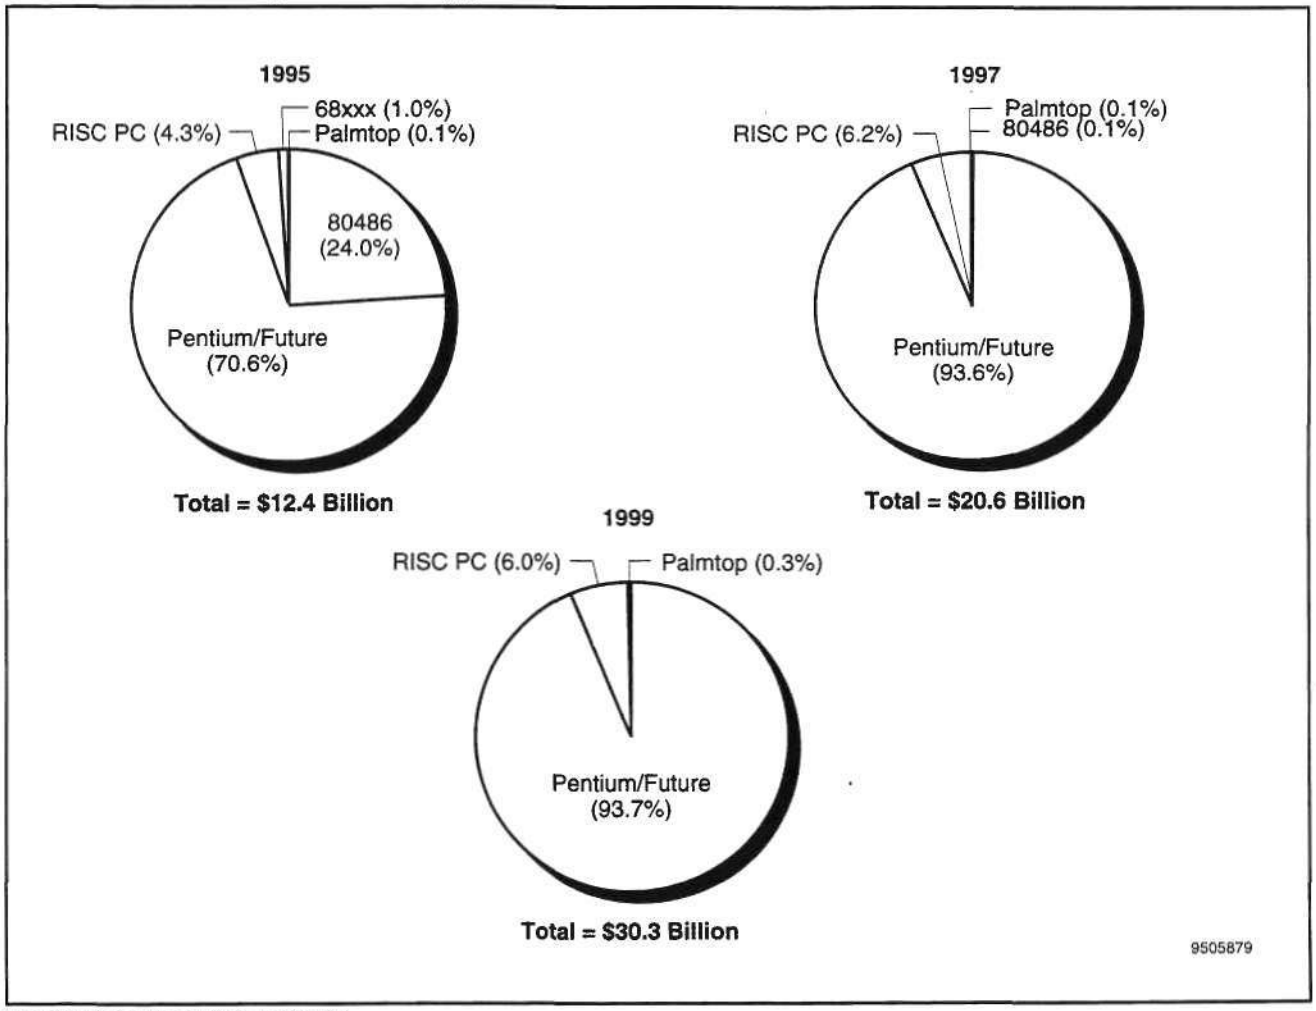

Source: Dataquest (October 1995)

Figure 4-21
Worldwide Microprocessor Opportunities in Desktop/Deskside PCs*

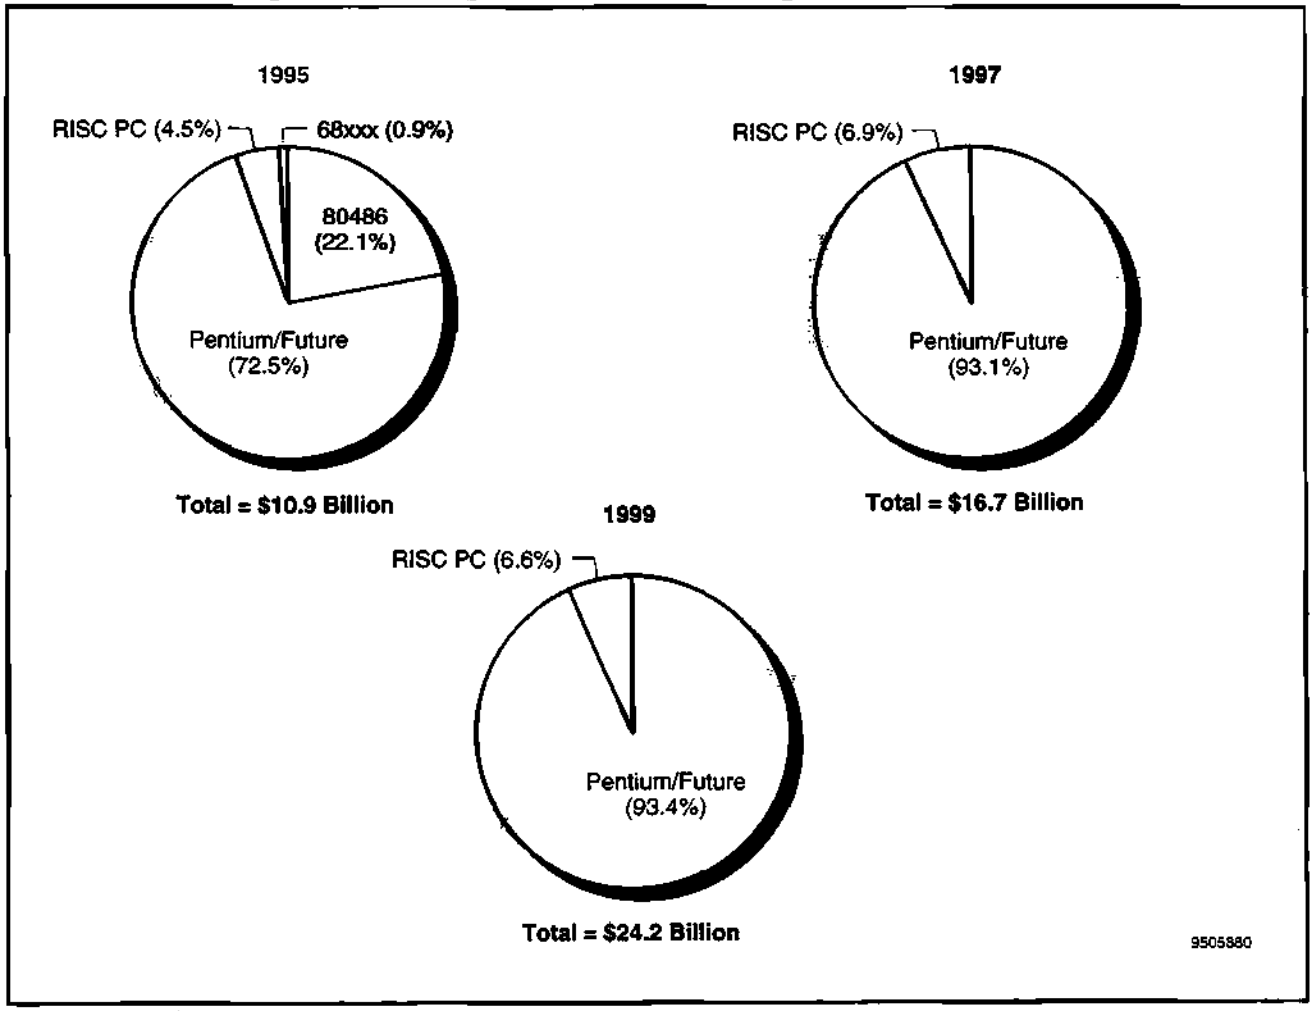

*Includes motherboard upgrade market
Source: Dataquest (October 1995)

Figure 4-22
Worldwide Microprocessor Opportunities in Mobile PCs

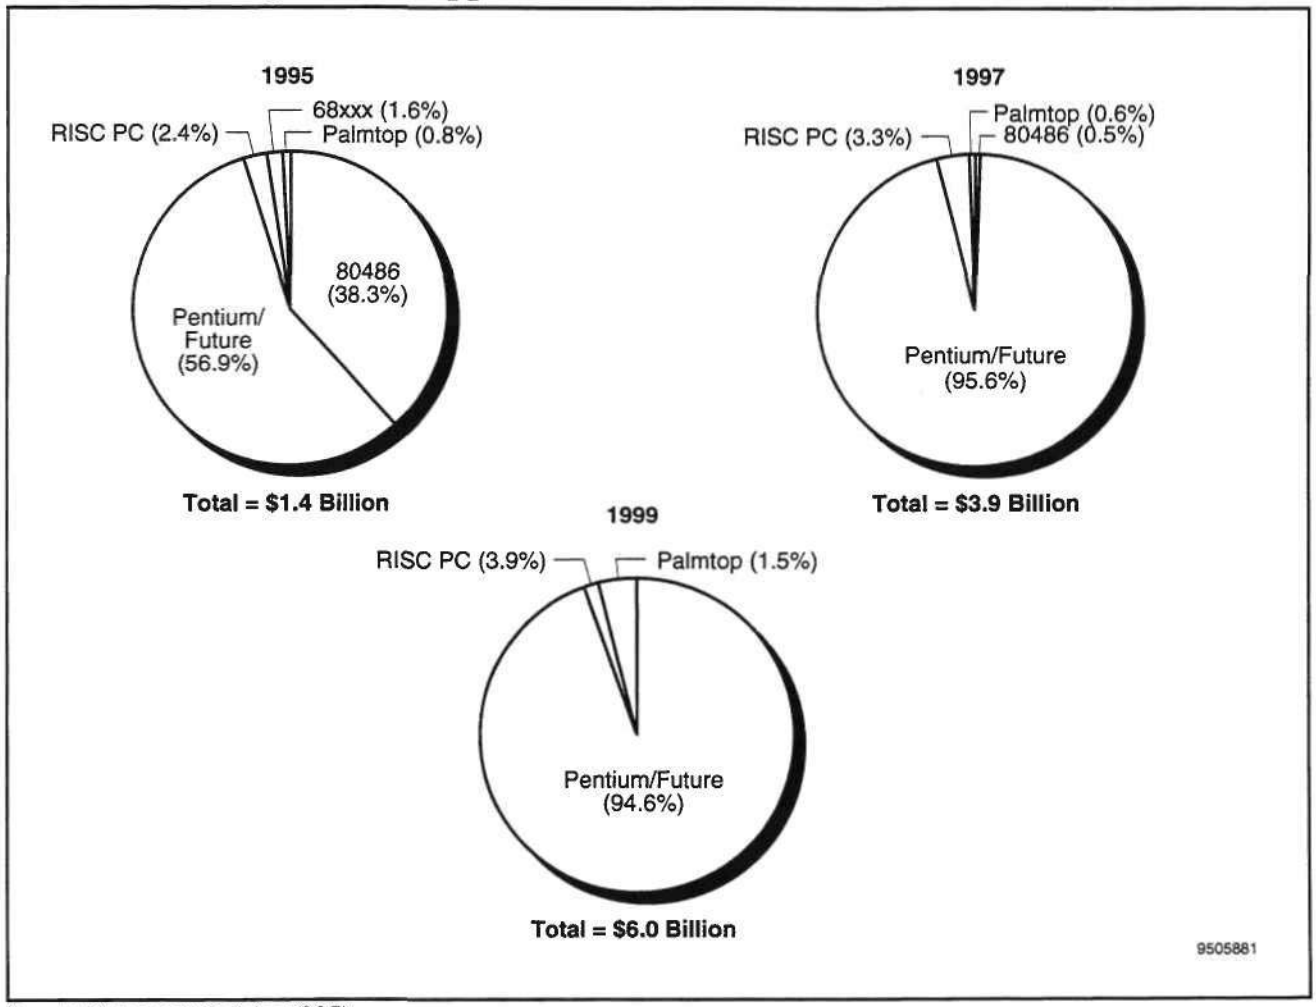

Source: Dataquest (October 1995)

Figure 4-23
Worldwide Core Logic Opportunities in PCs

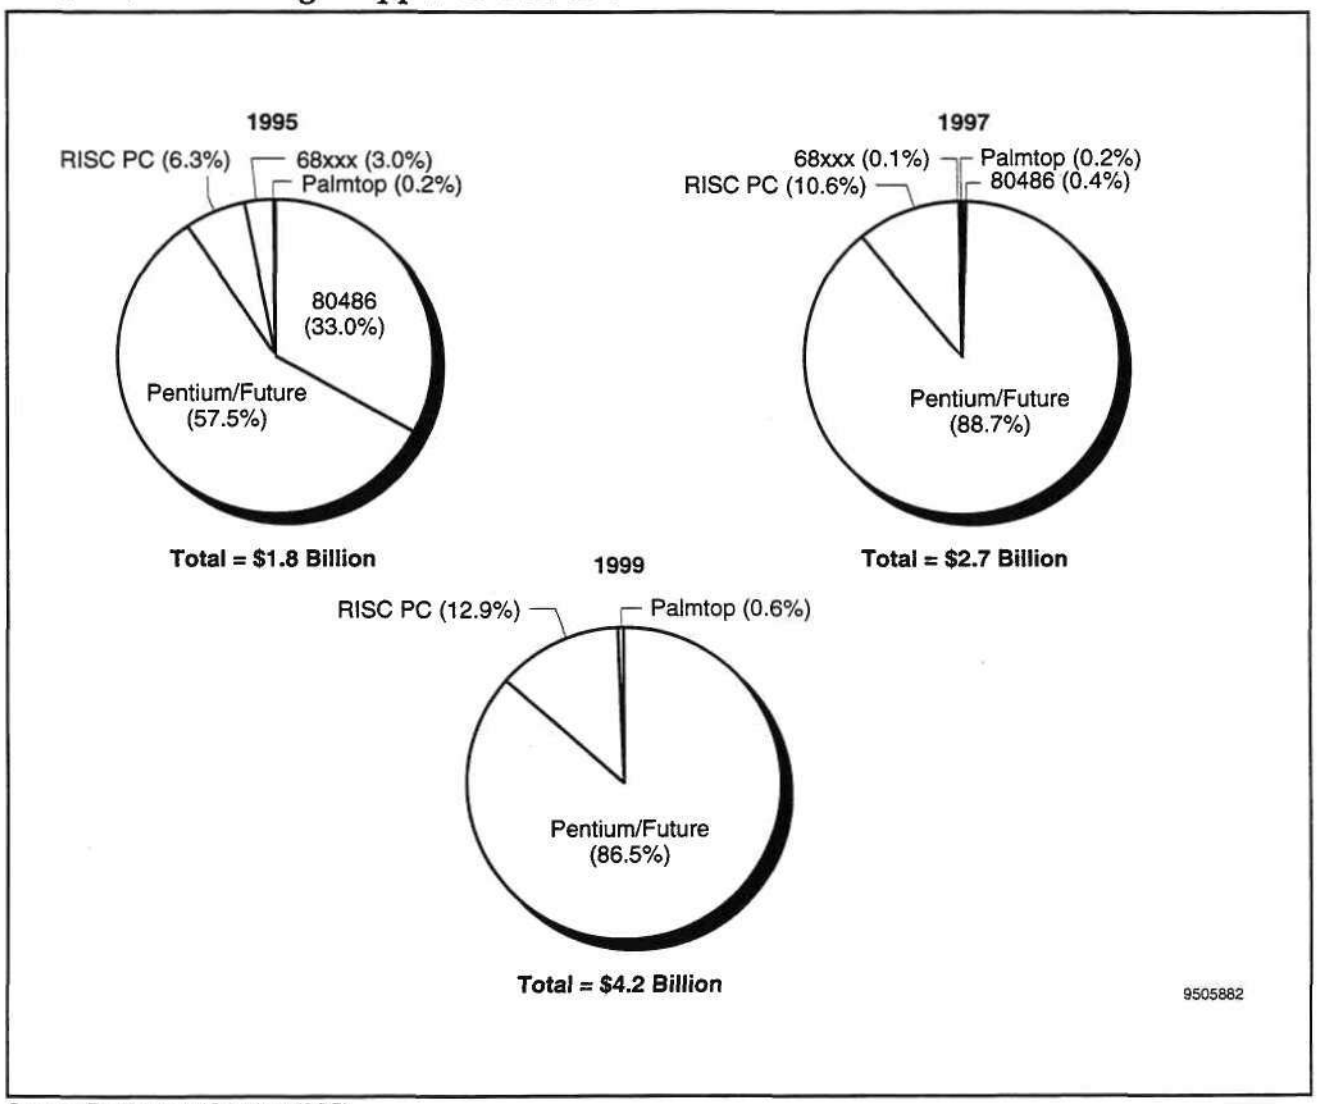

Source: Dataquest (October 1995)

Figure 4-24
Worldwide Core Logic Opportunities in Desktop/Deskside PCs*

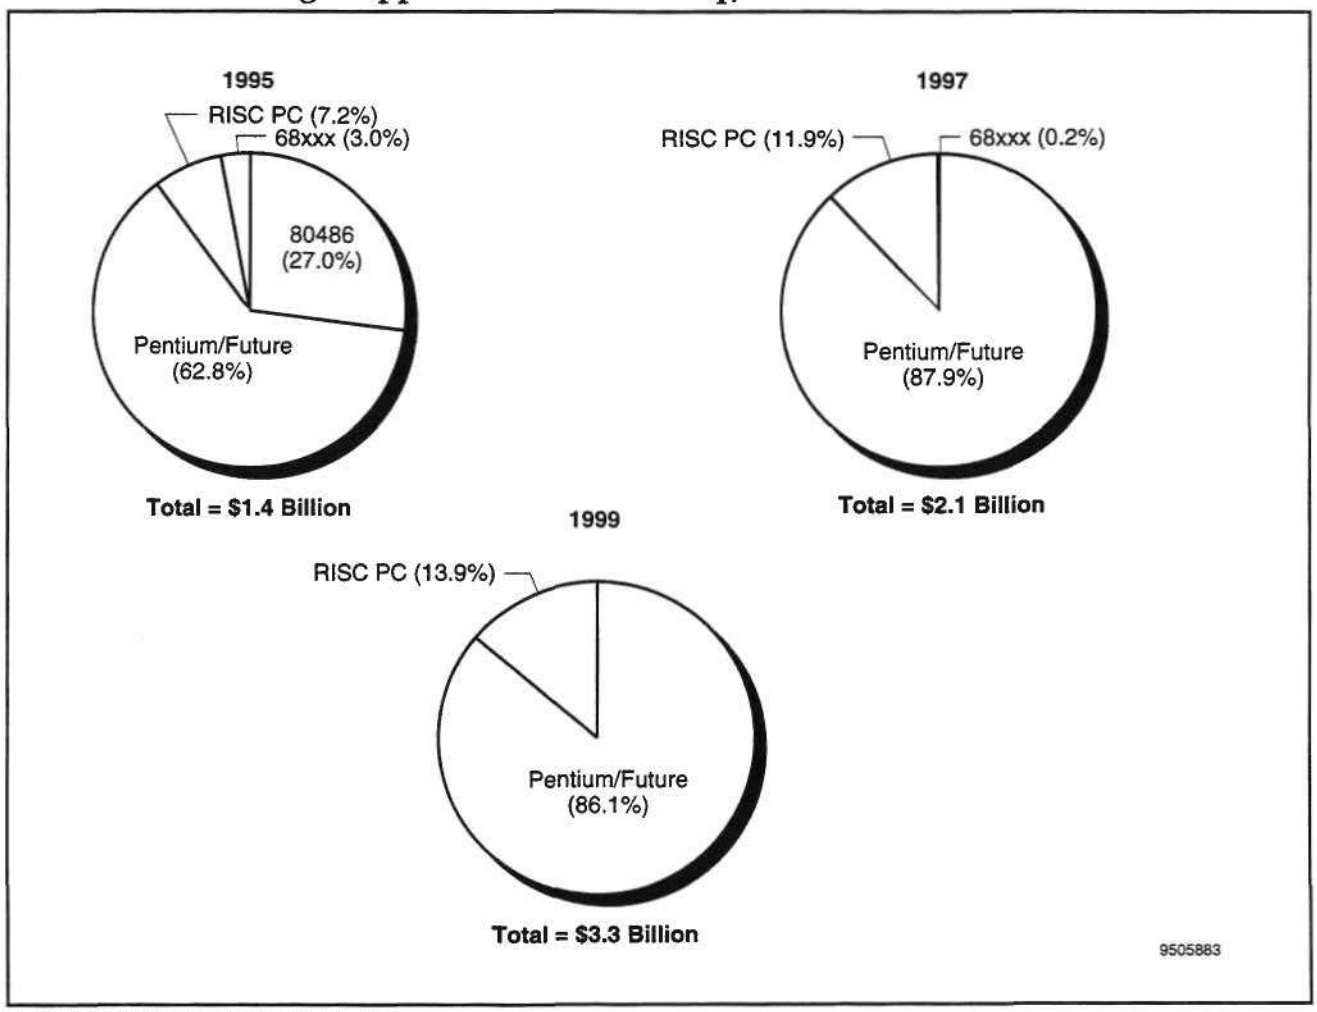

*Includes motherboard upgrade market
Source: Dataquest (October 1995)

Figure 4-25
Worldwide Core Logic Opportunities in Mobile PCs

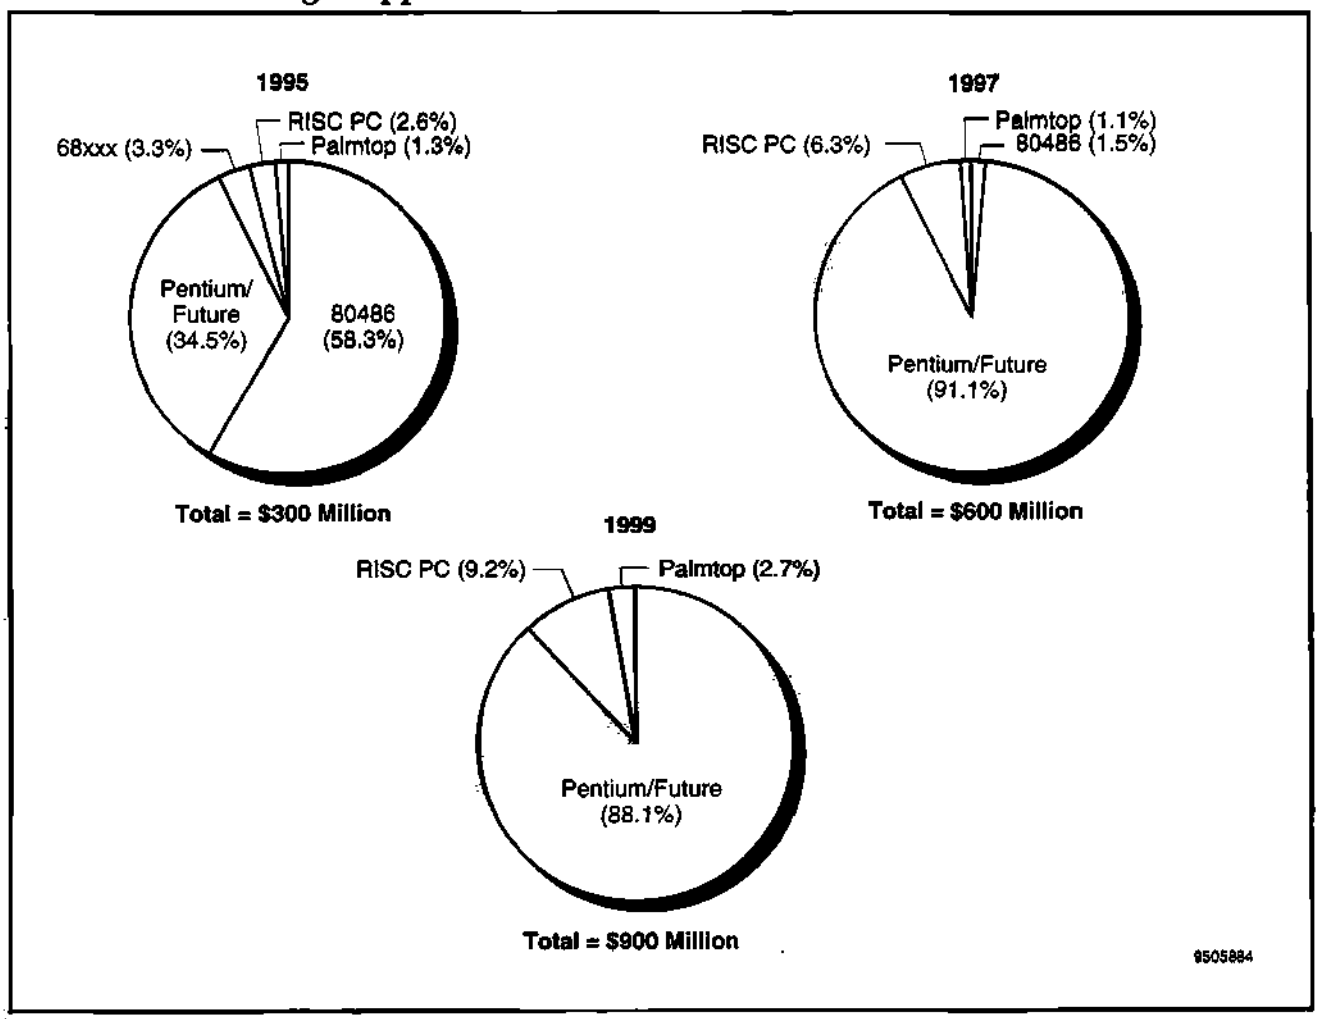

Source: Dataquest (October 1995)

Figure 4-26
Worldwide Aftermarket DRAM/VRAM Memory Demand in PCs (Includes Memory Purchased for Installation in PCs after PC System Point of Sale)

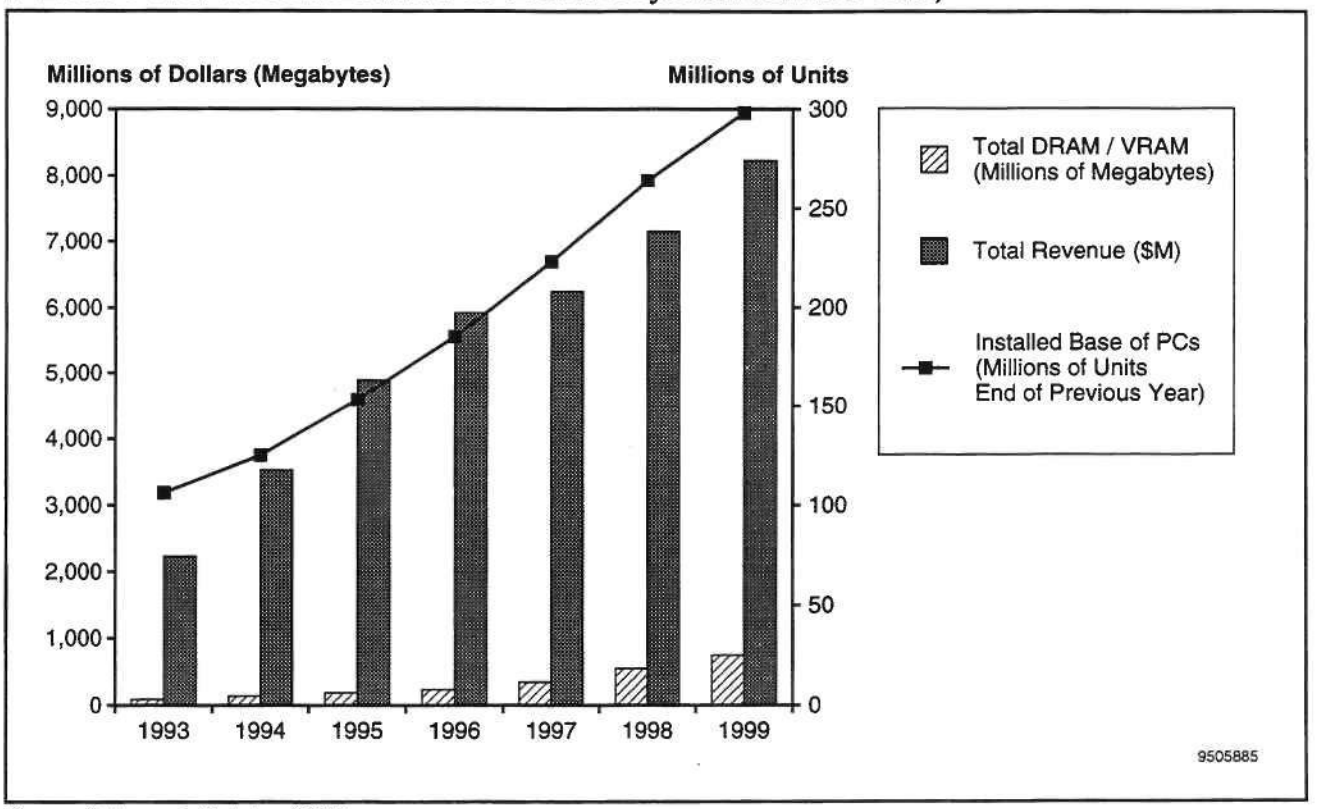

Source: Dataquest (October 1995)

Figure 4-27
Total DRAM/VRAM Consumption in PCs (Millions of Dollars)

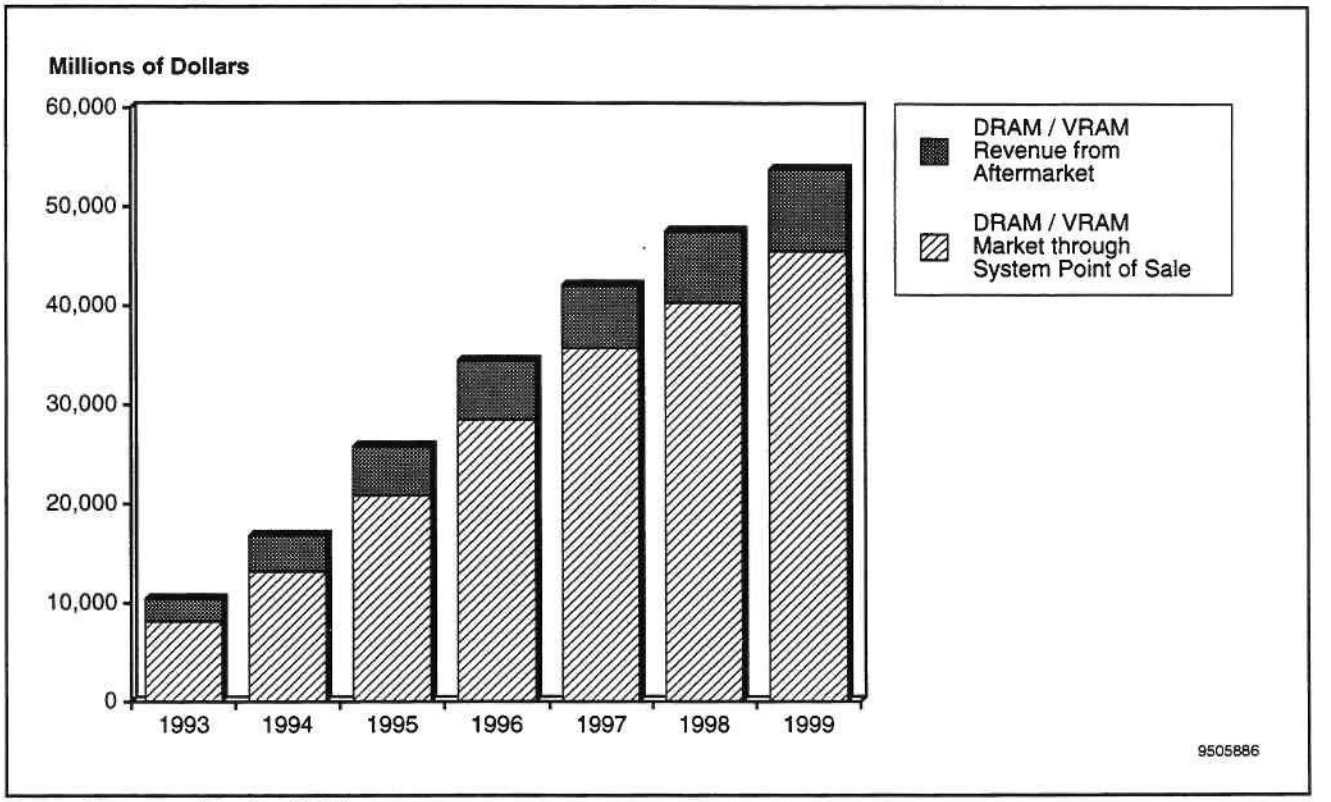

Source: Dataquest (October 1995)

Table 4-6
Worldwide Total PC Semiconductor Market by Device Type (Millions of Dollars) (Includes Motherboard, Graphics, Storage, Serial/Parallel I/O Adapters, and Add-On Memory Installed up to Point of Sale)

	1993	1994	1995	1996	1997	1998	1999	CAGR (%) 1994-1999
Memory*	8,879	14,471	22,684	30,695	38,287	42,844	48,252	27.2
Microprocessor/FPU	6,026	9,113	12,386	16,503	20,598	25,465	30,293	27.2
Core Logic	1,013	1,356	1,779	2,005	2,709	3,577	4,163	25.2
Graphics	925	1,044	1,167	1,320	1,578	1,873	2,029	14.2
Storage Control	490	510	713	956	1,109	1,292	1,521	24.4
Others	753	1,405	2,101	2,948	3,885	5,040	6,271	34.9
Total Semiconductor	18,087	27,899	40,830	54,426	68,166	80,092	92,528	27.1
Percentage Growth	51.2	54.2	46.4	33.3	25.2	17.5	15.5	

*Includes nonvolatile graphics BIOS memory
Source: Dataquest (October 1995)

Table 4-7
Worldwide Desktop/Deskside PC Semiconductor Market by Device Type[1] (Millions of Dollars) (Includes Motherboard, Graphics, Storage, Serial/Parallel I/O Adapters, and Add-On Memory Installed up to Point of Sale)

	1993	1994	1995	1996	1997	1998	1999	CAGR (%) 1994-1999
Memory[2]	7,744	12,417	19,231	25,172	31,638	34,963	38,644	25.5
Microprocessor/FPU	5,348	8,251	10,948	13,634	16,737	20,396	24,249	24.1
Core Logic	841	1,102	1,440	1,487	2,093	2,830	3,288	24.4
Graphics	784	843	924	1,025	1,190	1,381	1,524	12.6
Storage Control	431	447	628	810	915	1,050	1,228	22.4
Others	658	1,226	1,785	2,375	3,043	3,897	4,918	32.0
Total Semiconductor	15,806	24,287	34,957	44,504	55,615	64,518	73,851	24.9
Percentage Growth	50.4	53.7	43.9	27.3	25	16	14.5	

[1]Includes motherboard upgrade market
[2]Includes nonvolatile graphics BIOS memory
Source: Dataquest (October 1995)

Table 4-8
Worldwide Mobile PC Semiconductor Market by Device Type (Millions of Dollars) (Includes Motherboard, Graphics, Storage, Serial/Parallel I/O Adapters, and Add-On Memory Installed up to Point of Sale)

	1993	1994	1995	1996	1997	1998	1999	CAGR (%) 1994-1999
Memory*	1,135	2,054	3,453	5,523	6,649	7,881	9,608	36.1
Microprocessor/FPU	678	862	1,438	2,869	3,861	5,069	6,044	47.6
Core Logic	173	253	338	518	616	747	874	28.1
Graphics	141	201	243	295	388	492	504	20.2
Storage Control	59	63	85	146	194	242	293	36.2
Others	95	179	316	572	842	1,143	1,352	49.9
Total Semiconductor	2,281	3,611	5,874	9,922	12,551	15,573	18,676	38.9
Percentage Growth	57.6	58.3	62.6	68.9	26.5	24.1	19.9	

*Includes nonvolatile graphics BIOS memory
Source: Dataquest (October 1995)

Table 4-9
Worldwide Semiconductor Memory Demand in PCs by Device* (Millions of Dollars) (Includes Memory Installed in Factory and Add-On Memory Installed up to Point of Sale)

	1993	1994	1995	1996	1997	1998	1999	CAGR (%) 1994-1999
DRAM/VRAM	8,186	13,226	20,864	28,512	35,780	40,329	45,520	28.0
SRAM	395	834	1,327	1,819	2,201	2,210	2,419	23.7
EPROM/OTP/ROM	239	164	49	18	0	0	0	-100.0
Flash	60	246	444	346	306	306	313	4.9
Total	8,879	14,471	22,684	30,695	38,287	42,844	48,252	27.2

*Includes nonvolatile graphics BIOS memory
Source: Dataquest (October 1995)

Table 4-10
Worldwide Semiconductor Memory Demand in Desktop/Deskside PCs and Motherboard Upgrades by Device* (Millions of Dollars) (Includes Memory Installed in Factory and Add-On Memory Installed up to Point of Sale)

	1993	1994	1995	1996	1997	1998	1999	CAGR (%) 1994-1999
DRAM/VRAM	7,141	11,331	17,670	23,449	29,554	32,982	36,610	26.4
SRAM	371	793	1,189	1,453	1,870	1,781	1,840	18.3
EPROM/OTP/ROM	186	117	37	14	0	0	0	-100.0
Flash	46	176	335	257	214	200	194	2.0
Total	7,744	12,417	19,231	25,172	31,638	34,963	38,644	25.5

*Includes nonvolatile graphics BIOS memory
Source: Dataquest (October 1995)

Table 4-11
Worldwide Semiconductor Memory Demand in Mobile PCs by Device* (Millions of Dollars) (Includes Memory Installed in Factory and Add-On Memory Installed up to Point of Sale)

	1993	1994	1995	1996	1997	1998	1999	CAGR (%) 1994-1999
DRAM/VRAM	1,045	1,895	3,194	5,063	6,226	7,347	8,910	36.3
SRAM	24	42	138	366	331	428	580	69.1
EPROM/OTP/ROM	53	47	12	5	0	0	0	-100.0
Flash	13	70	109	89	92	106	118	11.1
Total	1,135	2,054	3,453	5,523	6,649	7,881	9,608	36.1

*Includes nonvolatile graphics BIOS memory
Source: Dataquest (October 1995)

Table 4-12
Worldwide Aftermarket DRAM/VRAM Memory Demand in PCs (Includes Memory Purchased for installation in PCs after PC System Point of Sale)

	1993	1994	1995	1996	1997	1998	1999	CAGR (%) 1994-1999
Total DRAM/VRAM (Millions of Megabytes)	84	137	186	235	347	550	748	40.4
Total Revenue ($M)	2,228	3,527	4,889	5,915	6,238	7,151	8,223	18.4
Installed Base of PCs (Millions of Units End of Previous Year)	106	125	153	185	223	264	298	18.9

Source: Dataquest (October 1995)

Table 4-13
Worldwide Semiconductor Memory Demand by Total PC System (Market Value by Function; Millions of Dollars)

	1993	1994	1995	1996	1997	1998	1999	CAGR (%) 1994-1999
386 PC								
Main OEM and POS Add-On	1,438	232	7	0	0	0	0	-100.0
Display Buffer	324	37	1	0	0	0	0	-100.0
Cache	50	7	0	0	0	0	0	-100.0
Nonvolatile	37	6	0	0	0	0	0	-100.0
Total	1,848	282	9	0	0	0	0	-100.0
486 PC								
Main OEM and POS Add-On	4,368	8,419	7,187	2,685	91	0	0	-100.0
Display Buffer	940	1,067	903	341	7	0	0	-100.0
Cache	319	633	366	110	0	0	0	-100.0
Nonvolatile	113	195	158	45	1	0	0	-100.0
Total	5,739	10,313	8,613	3,181	100	0	0	-100.0
68xxx PC								
Main OEM and POS Add-On	749	948	583	135	37	0	0	-100.0
Display Buffer	222	199	113	17	3	0	0	-100.0
Cache	4	3	3	1	0	0	0	-100.0
Nonvolatile	69	55	27	2	1	0	0	-100.0
Total	1,044	1,204	726	155	40	0	0	-100.0
Pentium and Future PC								
Main OEM and POS Add-On	96	1,711	10,136	20,593	30,002	33,657	37,470	85.4
Display Buffer	8	114	757	2,064	1,962	2,236	2,611	87.1
Cache	8	134	773	1,352	1,902	1,805	1,840	68.9
Nonvolatile	2	35	157	153	125	91	63	12.2
Total	113	1,994	11,822	24,162	33,992	37,789	41,983	83.9
RISC PC								
Main OEM and POS Add-On	32	451	1,026	2,343	3,375	4,156	5,133	62.6
Display Buffer	2	48	151	334	302	280	306	44.9
Cache	3	35	149	287	214	224	254	48.7
Nonvolatile	1	9	16	17	13	11	8	-1.9
Total	37	543	1,340	2,981	3,904	4,670	5,702	60.0
Palmtop PC								
Main OEM and POS Add-On	12	23	36	70	85	181	326	70.0
Display Buffer	0	0	0	0	0	0	0	-
Cache	0	0	0	0	0	0	0	-
Nonvolatile	14	31	41	32	33	47	57	13.0
Total	27	54	78	102	118	228	382	48.0

(Continued)

Table 4-13 (Continued)
Worldwide Semiconductor Memory Demand by Total PC System (Market Value by Function; Millions of Dollars)

	1993	1994	1995	1996	1997	1998	1999	CAGR (%) 1994-1999
Other PC (Including 8088, 8086, 80286)								
Main OEM and POS Add-On	6	1	0	0	0	0	0	-100.0
Display Buffer	1	0	0	0	0	0	0	-100.0
Cache	0	0	0	0	0	0	0	-
Nonvolatile	0	0	0	0	0	0	0	-100.0
Total	7	1	0	0	0	0	0	-100.0
Total PC								
Main OEM and POS Add-On	6,700	11,784	18,975	25,826	33,590	37,994	42,928	29.5
Display Buffer	1,498	1,465	1,925	2,756	2,274	2,516	2,917	14.8
Cache	382	812	1,290	1,749	2,116	2,029	2,094	20.9
Nonvolatile	236	331	398	250	173	149	128	-17.4
Total	8,817	14,392	22,589	30,581	38,154	42,688	48,067	27.3

Source: Dataquest (October 1995)

Table 4-14
Worldwide Semiconductor Memory Demand for Desktop/Deskside PC Systems (Market Value by Function; Millions of Dollars)

	1993	1994	1995	1996	1997	1998	1999	CAGR (%) 1994-1999
386 PC								
Main OEM and POS Add-On	1,185	189	6	0	0	0	0	-100.0
Display Buffer	287	28	1	0	0	0	0	-100.0
Cache	46	7	0	0	0	0	0	-100.0
Nonvolatile	30	4	0	0	0	0	0	-100.0
Total	1,548	228	7	0	0	0	0	-100.0
486 PC								
Main OEM and POS Add-On	3,839	6,929	5,346	2,146	11	0	0	-100.0
Display Buffer	846	872	711	293	1	0	0	-100.0
Cache	312	615	348	107	0	0	0	-100.0
Nonvolatile	87	135	118	40	0	0	0	-100.0
Total	5,083	8,551	6,523	2,586	13	0	0	-100.0
68xxx PC								
Main OEM and POS Add-On	649	833	485	135	37	0	0	-100.0
Display Buffer	193	162	90	17	3	0	0	-100.0
Cache	3	2	3	1	0	0	0	-100.0
Nonvolatile	60	45	21	2	1	0	0	-100.0
Total	905	1,042	599	155	40	0	0	-100.0
Pentium and Future PC								
Main OEM and POS Add-On	96	1,704	9,247	16,701	24,553	27,211	29,733	77.1
Display Buffer	8	113	680	1,764	1,656	1,834	2,226	81.4
Cache	8	133	694	1,077	1,671	1,574	1,608	64.5
Nonvolatile	2	35	141	122	98	71	49	6.8
Total	113	1,986	10,762	19,663	27,977	30,689	33,615	76.1
RISC PC								
Main OEM and POS Add-On	32	451	964	2,100	3,030	3,696	4,394	57.6
Display Buffer	2	48	140	293	263	241	257	39.9
Cache	3	35	143	268	199	207	232	46.0
Nonvolatile	1	9	15	15	12	9	7	-5.2
Total	37	543	1,262	2,677	3,504	4,154	4,890	55.2

(Continued)

Table 4-14 (Continued)
Worldwide Semiconductor Memory Demand for Desktop/Deskside PC Systems
(Market Value by Function; Millions of Dollars)

	1993	1994	1995	1996	1997	1998	1999	CAGR (%) 1994-1999
Other PC (Including 8088, 8086, 80286)								
Main OEM and POS Add-On	3	1	0	0	0	0	0	-100.0
Display Buffer	1	0	0	0	0	0	0	-100.0
Cache	0	0	0	0	0	0	0	-
Nonvolatile	0	0	0	0	0	0	0	-100.0
Total	4	1	0	0	0	0	0	-100.0
Total PC								
Main OEM and POS Add-On	5,803	10,108	16,048	21,082	27,631	30,907	34,127	27.6
Display Buffer	1,338	1,223	1,622	2,367	1,923	2,075	2,483	15.2
Cache	371	793	1,189	1,453	1,870	1,781	1,840	18.3
Nonvolatile	179	228	295	179	111	80	56	-24.6
Total	7,691	12,352	19,153	25,081	31,535	34,843	38,505	25.5

Source: Dataquest (October 1995)

Table 4-15
Worldwide Semiconductor Memory Demand for Mobile PC Systems (Market Value by Function; Millions of Dollars)

	1993	1994	1995	1996	1997	1998	1999	CAGR (%) 1994-1999
386 PC								
Main OEM and POS Add-On	253	43	1	0	0	0	0	-100.0
Display Buffer	36	9	0	0	0	0	0	-100.0
Cache	3	1	0	0	0	0	0	-100.0
Nonvolatile	7	1	0	0	0	0	0	-100.0
Total	300	54	2	0	0	0	0	-100.0
486 PC								
Main OEM and POS Add-On	529	1,489	1,841	539	80	0	0	-100.0
Display Buffer	94	195	193	48	7	0	0	-100.0
Cache	7	17	17	3	0	0	0	-100.0
Nonvolatile	26	60	40	5	1	0	0	-100.0
Total	656	1,762	2,090	595	87	0	0	-100.0
68xxx PC								
Main OEM and POS Add-On	100	115	98	0	0	0	0	-100.0
Display Buffer	30	37	23	0	0	0	0	-100.0
Cache	1	0	1	0	0	0	0	-100.0
Nonvolatile	9	10	5	0	0	0	0	-100.0
Total	139	162	127	0	0	0	0	-100.0
Pentium and Future PC								
Main OEM and POS Add-On	0	6	890	3,892	5,450	6,446	7,736	315.0
Display Buffer	0	1	77	300	306	403	385	266.2
Cache	0	1	78	275	232	231	232	220.3
Nonvolatile	0	0	16	31	27	21	14	138.8
Total	0	8	1,061	4,498	6,015	7,100	8,368	304.4
RISC PC								
Main OEM and POS Add-On	0	0	62	243	345	460	739	-
Display Buffer	0	0	11	41	39	38	49	-
Cache	0	0	5	19	15	16	22	-
Nonvolatile	0	0	1	2	2	1	1	-
Total	0	0	79	305	400	516	812	-

(Continued)

Table 4-15 (Continued)
Worldwide Semiconductor Memory Demand for Mobile PC Systems (Market Value by Function; Millions of Dollars)

	1993	1994	1995	1996	1997	1998	1999	CAGR (%) 1994-1999
Palmtop PC								
Main OEM and POS Add-On	12	23	36	70	85	181	326	70.0
Display Buffer	0	0	0	0	0	0	0	-
Cache	0	0	0	0	0	0	0	-
Nonvolatile	14	31	41	32	33	47	57	13.0
Total	27	54	78	102	118	228	382	48.0
Other PC (Including 8088, 8086, 80286)								
Main OEM and POS Add-On	3	0	0	0	0	0	0	-100.0
Display Buffer	1	0	0	0	0	0	0	-100.0
Cache	0	0	0	0	0	0	0	-
Nonvolatile	0	0	0	0	0	0	0	-100.0
Total	4	0	0	0	0	0	0	-100.0
Total PC								
Main OEM and POS Add-On	897	1,677	2,927	4,744	5,959	7,087	8,801	39.3
Display Buffer	161	242	303	389	351	441	434	12.4
Cache	11	19	102	296	246	247	254	68.1
Nonvolatile	57	103	104	70	63	69	72	-6.9
Total	1,126	2,040	3,436	5,500	6,620	7,844	9,562	36.2

Source: Dataquest (October 1995)

Table 4-16
Average Memory Configuration Assumptions by Overall PC System (Megabytes per System)

	1993	1994	1995	1996	1997	1998	1999
386 PC							
Main (OEM)	4.00	4.00	4.00	NM	NM	NM	NM
Main Add-On[1]	0.84	1.90	1.89	NM	NM	NM	NM
Display Buffer	0.90	1.00	1.00	NM	NM	NM	NM
Cache (KB)	64.00	64.00	64.00	NM	NM	NM	NM
Nonvolatile[2]	0.13	0.13	0.13	NM	NM	NM	NM
486PC							
Main (OEM)	4.44	5.91	7.75	8.80	9.13	NM	NM
Main Add-On[1]	1.02	1.44	1.69	2.24	3.05	NM	NM
Display Buffer	0.97	1.00	1.19	1.40	1.00	NM	NM
Cache (KB)	115.62	162.72	151.53	160.34	27.94	NM	NM
Nonvolatile[2]	0.14	0.15	0.25	0.45	0.28	NM	NM
68xxx PC							
Main (OEM)	4.87	7.26	8.00	8.00	12.00	NM	NM
Main Add-On[1]	0.53	2.58	0.99	3.37	1.00	NM	NM
Display Buffer	1.00	1.00	1.00	1.00	1.00	NM	NM
Cache (KB)	64.00	64.00	128.00	128.00	256.00	NM	NM
Nonvolatile[2]	0.50	0.50	0.50	0.50	0.50	NM	NM
Pentium and Future PC							
Main (OEM)	12.00	11.98	9.90	10.80	15.78	21.55	26.55
Main Add-On[1]	2.00	2.00	3.50	3.15	5.49	6.85	6.11
Display Buffer	1.00	1.00	1.00	1.40	1.39	1.89	2.28
Cache (KB)	256.00	256.00	256.00	256.00	456.42	454.00	454.68
Nonvolatile[2]	0.25	0.25	0.25	0.25	0.25	0.25	0.25
RISC PC							
Main (OEM)	11.95	12.00	9.93	10.88	15.87	21.73	26.68
Main Add-On[1]	2.00	2.14	3.68	3.17	6.48	8.00	6.87
Display Buffer	1.00	1.50	2.00	2.00	2.00	2.00	2.00
Cache (KB)	256.00	256.00	494.03	480.88	479.16	476.89	470.94
Nonvolatile[2]	0.25	0.25	0.25	0.25	0.25	0.25	0.25

(Continued)

Table 4-16 (Continued)
Average Memory Configuration Assumptions by Overall PC System
(Megabytes per System)

	1993	1994	1995	1996	1997	1998	1999
Palmtop PC							
Main (OEM)	1.00	1.25	1.60	2.50	3.25	4.00	4.00
Main Add-On[1]	0	0	0	0	0	0	0
Display Buffer	0	0	0	0	0	0	0
Cache (KB)	0	0	0	0	0	0	0
Nonvolatile[2]	2.00	2.50	3.50	4.00	4.25	4.50	4.50

[1]Includes memory added through system point of sale
[2]Excludes graphics BIOS typically using a 32Kx8 EPROM/OTP
NM = Not meaningful
Source: Dataquest (October 1995)

Table 4-17
Average Memory Configuration Assumptions for Desktop/Deskside PC Systems (Megabytes per System)

	1993	1994	1995	1996	1997	1998	1999
386 PC							
Main (OEM)	4.00	4.00	4.00	NM	NM	NM	NM
Main Add-On[1]	1.00	2.40	2.40	NM	NM	NM	NM
Display Buffer	1.00	1.00	1.00	NM	NM	NM	NM
Cache (KB)	64.00	64.00	64.00	NM	NM	NM	NM
Nonvolatile[2]	0.13	0.13	0.13	NM	NM	NM	NM
486PC							
Main (OEM)	4.50	6.00	8.00	9.00	10.00	NM	NM
Main Add-On[1]	1.00	1.40	1.40	2.00	2.00	NM	NM
Display Buffer	1.00	1.00	1.25	1.50	1.00	NM	NM
Cache (KB)	128.00	192.00	192.00	192.00	0.00	NM	NM
Nonvolatile[2]	0.13	0.13	0.25	0.50	0.50	NM	NM
68xxx PC							
Main (OEM)	5.00	8.00	8.00	8.00	12.00	NM	NM
Main Add-On1	0.40	2.60	1.40	3.37	1.00	NM	NM
Display Buffer	1.00	1.00	1.00	1.00	1.00	NM	NM
Cache (KB)	64.00	64.00	128.00	128.00	256.00	NM	NM
Nonvolatile[2]	0.50	0.50	0.50	0.50	0.50	NM	NM
Pentium and Future PC							
Main (OEM)	12.00	12.00	10.00	11.00	16.00	22.00	27.00
Main Add-On[1]	2.00	2.00	3.60	3.20	6.24	7.68	6.40
Display Buffer	1.00	1.00	1.00	1.50	1.50	2.00	2.50
Cache (KB)	256.00	256.00	256.00	256.00	512.00	512.00	512.00
Nonvolatile[2]	0.25	0.25	0.25	0.25	0.25	0.25	0.25
RISC PC							
Main (OEM)	12.00	12.00	10.00	11.00	16.00	22.00	27.00
Main Add-On[1]	2.00	2.14	3.76	3.33	7.02	8.64	7.20
Display Buffer	1.00	1.50	2.00	2.00	2.00	2.00	2.00
Cache (KB)	256.00	256.00	512.00	512.00	512.00	512.00	512.00
Nonvolatile[2]	0.25	0.25	0.25	0.25	0.25	0.25	0.25

[1]Includes memory added through system point of sale
[2]Excludes graphics BIOS typically using a 32Kx8 EPROM/OTP
NM = Not meaningful
Source: Dataquest (October 1995)

Table 4-18
Average Memory Configuration Assumptions for Mobile PC Systems (Megabytes per System)

	1993	1994	1995	1996	1997	1998	1999
386 PC							
Main (OEM)	4.00	4.00	4.00	NM	NM	NM	NM
Main Add-On[1]	0.20	0.12	0.03	NM	NM	NM	NM
Display Buffer	0.50	1.00	1.00	NM	NM	NM	NM
Cache (KB)	64.00	64.00	64.00	NM	NM	NM	NM
Nonvolatile[2]	0.13	0.13	0.13	NM	NM	NM	NM
486PC							
Main (OEM)	4.00	5.50	7.00	8.00	9.00	NM	NM
Main Add-On[1]	1.12	1.60	2.56	3.20	3.20	NM	NM
Display Buffer	0.75	1.00	1.00	1.00	1.00	NM	NM
Cache (KB)	32.00	32.00	32.00	32.00	32.00	NM	NM
Nonvolatile[2]	0.25	0.25	0.25	0.25	0.25	NM	NM
68xxx PC							
Main (OEM)	4.00	4.00	8.00	NM	NM	NM	NM
Main Add-On[1]	1.40	2.48	-0.60	NM	NM	NM	NM
Display Buffer	1.00	1.00	1.00	NM	NM	NM	NM
Cache (KB)	64.00	64.00	128.00	NM	NM	NM	NM
Nonvolatile[2]	0.50	0.50	0.50	NM	NM	NM	NM
Pentium and Future PC							
Main (OEM)	NM	8.00	9.00	10.00	15.00	20.00	25.00
Main Add-On[1]	NM	2.00	2.60	2.96	2.80	4.00	5.12
Display Buffer	NM	1.00	1.00	1.00	1.00	1.50	1.50
Cache (KB)	NM	256.00	256.00	256.00	256.00	256.00	256.00
Nonvolatile[2]	NM	0.25	0.25	0.25	0.25	0.25	0.25
RISC PC							
Main (OEM)	8.00	NM	9.00	10.00	15.00	20.00	25.00
Main Add-On[1]	-2.00	NM	-2.64	-2.00	-2.80	-4.00	-5.12
Display Buffer	1.00	NM	2.00	2.00	2.00	2.00	2.00
Cache (KB)	256.00	NM	256.00	256.00	256.00	256.00	256.00
Nonvolatile[2]	0.25	NM	0.25	0.25	0.25	0.25	0.25
Palmtop PC							
Main (OEM)	1.00	1.25	1.60	2.50	3.25	4.00	4.00
Main Add-On[1]	0	0	0	0	0	0	0
Display Buffer	0	0	0	0	0	0	0
Cache (KB)	0	0	0	0	0	0	0
Nonvolatile[2]	2.00	2.50	3.50	4.00	4.25	4.50	4.50

[1]Includes memory added through system point of sale
[2]Excludes graphics BIOS typically using a 32Kx8 EPROM/OTP
NM = Not meaningful
Source: Dataquest (October 1995)

Chapter 5
Input/Output and Dedicated Systems

Tables 5-1 through 5-8 detail system and semiconductor market data on selected computer I/O systems and high-volume dedicated systems such as copiers and expandable organizers.

Key Trends

Sound Boards

Key trends in sound boards are as follows:

- Market ramp-up continues through 1995, then saturation points are hit for the home market. PC motherboard embedded sound I/O becomes more prevalent in later years for "home" and mobile PC use.
- Microsoft has adopted DSP Group's TrueSpeech technology as a compression standard. Several vendors already support ADPCM for compression.
- Movement to 16 bit is nearly complete. Wave-table synthesis is gaining market share.
- Key chip functions include FM and wave-table synthesis (1MB to 4MB ROM-based), ASSP/ASICs (mixed-signal and digital CMOS), audio amplifiers, and mixers. SCSI host adapters are found on many boards. IDE (fast IDE) is becoming more prevalent.

Graphics Boards

Key trends in graphics boards are as follows:

- There is steady replacement by embedded motherboard chips.
- Boards are moving to greater than 1Kx1K-pixel resolutions, accelerated BitBLT-based, 64-bit data paths, and RAMDAC technology moving from 85 MHz to 135 MHz. Also, digital video, 3-D, and sound capability are appearing in the high-end boards.
- DRAM, and especially EDO DRAM, is becoming more common. Minimal buffers start at 1MB and move to 4MB with optional SIMMs. Most high-end add-in boards have separate RAMDAC, and in some case digital video functions are being integrated.

Digital Video Boards

Key trends in digital video boards are as follows:

- There will be continued penetration into the multimedia content creator market (software title development, market communications, and training). Playback board growth will be limited to full-screen 15- to 30-fps acceleration. Other opportunities exist for TV tuner, capture pass-through, and integrated audio and graphics boards.

- MPEG 1 decode boards became popular in 1994. Use of MPEG chips on PC motherboards will begin in late 1995 on some models.
- Key semiconductor opportunities include compression decoders for MPEG and JPEG, among others (and encoders for real-time algorithms), PAL/NTSC to CCIR decoders and encoders, ASICs (CMOS), digital video processors (scaling, among others), and DRAM/VRAM pixels buffers.

Leading digital video board OEMs worldwide are ATI, Creative Technologies, Diamond Multimedia, FutureTel, IBM, Intel, Matrox, Media Vision, Optibase, Optivision, Orchid, Sigma Designs, SuperMac, and Video Logic.

Monitors

Key trends in monitors are as follows:

- Color will grow to 98 percent of the market in 1999.
- 15 inches will become the predominant size in 1995; 16/17 inches will be predominant in 1997.
- There will be a chip content of about $5 for primarily video amplifier (moving to 135 MHz) and CRT controls.

Keyboards

The trend in keyboards is commodity items moving toward ergonomic and wireless versions.

Leading keyboard OEMs worldwide are Keytronic, Silitek, and MaxiSwitch.

Table 5-1
Worldwide Sound Board Application Market

	1993	1994	1995	1996	1997	1998	1999	CAGR (%) 1994-1999
Board Units (M)	4.7	13.7	14.4	15.0	15.7	16.4	17.9	5.5
Board ASP ($)	93.0	64.6	50.0	44.0	39.0	35.0	32.0	-13.1
Board Factory Revenue ($M)	440.4	883.1	718.5	661.1	612.2	575.3	573.1	-8.3
16-Bit Penetration (%)	65.0	97.0	99.0	100.0	100.0	100.0	100.0	
Semiconductor Content ($)	18.6	16.1	15.0	13.6	12.5	11.6	10.6	-8.1
Semiconductor Market ($M)	88.1	220.8	215.6	205.0	195.9	189.8	189.1	-3.0
ASSP/ASIC ($M)	37.0	94.9	92.7	84.0	74.4	72.1	71.9	-5.4
Synthesis ($M)	33.5	81.7	75.4	65.6	54.9	53.2	53.0	-8.3
Analog/Discrete ($M)	13.2	33.1	32.3	30.7	27.4	26.6	26.5	-4.4
Memory/Others ($M)	4.4	11.0	15.1	24.6	39.2	38.0	37.8	27.9

Source: Dataquest (August 1995)

Table 5-2
Worldwide Sound Board OEMs (1994 Unit Share)

Company	Percent
Creative Technologies/Labs	66
Aztech	19
Media Vision	4
Cardinal	3
Others	8

Source: Dataquest (August 1995)

Table 5-3
Worldwide Graphic Board Application Market

	1993	1994	1995	1996	1997	1998	1999	CAGR (%) 1994-1999
Board Units (K)	4,020	8,423	8,408	8,846	9,466	9,924	10,191	3.9
Board ASP ($)	218	160	141	125	112	104	97	-9.5
Board Factory Revenue ($M)	876	1,344	1,187	1,109	1,064	1,030	988	-6.0
Semiconductor Content ($)	76	57	51	45	40	37	35	-9.5
Semiconductor Market ($M)	307	484	427	399	383	371	356	-6.0

Source: Dataquest (August 1995)

Table 5-4
Worldwide Graphic Board OEMs (1994 Percentage)

Company	Percent
Diamond Computer	20
ATI	14
Matrox	14
RasterOps	13
Radius/SuperMac	10
Others	29

Source: Dataquest (August 1995)

Table 5-5
Worldwide Digital Video Board Application Market

	1993	1994	1995	1996	1997	1998	1999	CAGR (%) 1994-1999
Board Units (K)	232.0	605.0	1,139.0	1,464.0	1,713.0	1,815.0	1,967.0	11.5
With Compression and/ or Decompression	35.0	169.0	448.0	589.0	695.0	744.0	580.0	5.3
With TV Tuner	30.0	100.0	250.0	340.0	425.0	500.0	750.0	24.6
Board ASP ($)	387.5	310.0	279.0	251.1	226.0	203.4	183.1	-8.1
Board Factory Revenue ($M)	139.5	187.6	317.8	367.6	387.1	369.2	360.1	13.9
Semiconductor Content ($)	92.9	87.8	78.2	67.2	57.5	52.8	48.2	-11.3
Semiconductor Market ($M)	21.6	53.1	89.1	98.4	98.5	95.9	94.8	12.3
TV Signal Processsing ($M)	0.5	1.5	3.8	5.1	6.4	7.5	11.3	49.6
Video Scaling/Processing ($M)	9.3	20.0	28.5	29.3	27.4	23.6	23.6	3.4
Compression/ Decompression ($M)	1.6	5.9	11.2	11.8	11.1	9.7	7.5	5.0
Memory (DRAM/VRAM) ($M)	9.1	22.7	40.0	45.0	45.1	46.0	42.6	13.4
Others ($M)	1.2	3.0	5.7	7.3	8.6	9.1	9.8	26.6

Source: Dataquest (August 1995)

Table 5-6
Worldwide Monitor Application Market

	1993	1994	1995	1996	1997	1998	1999	CAGR (%) 1994-1999
System Units (K)	33,865	41,227	48,211	55,422	61,243	68,836	76,725	13.2
System ASP ($)	356	352	356	366	368	375	364	0.7
System Factory Revenue ($M)	12,054	14,494	17,140	20,300	22,550	25,820	27,900	14.0
Semiconductor Content ($)	5	5	5	5	5	5	5	2.2
Semiconductor Market ($M)	157	188	231	274	304	361	391	15.7

Source: Dataquest (August 1995)

Table 5-7
Worldwide Monitor Label/OEMs (1994 Percentage)

Company	Percent
Compaq	7
IBM	6
NEC	6
Apple	6
Samtron	4
Packard Bell	4
CTX International	4
Philips	3
Others	61

Source: Dataquest (August 1995)

Table 5-8
Worldwide Keyboard Application Market

	1993	1994	1995	1996	1997	1998	1999	CAGR (%) 1994-1999
System Units (K)	46,318.8	55,961.6	65,304.4	75,797.7	85,541.1	97,527.6	110,180.8	14.5
Semiconductor Content ($)	2.5	2.5	2.5	2.5	2.5	2.5	2.5	0.2
Semiconductor Market ($M)	113.5	137.1	160.0	185.7	209.6	238.9	273.2	14.8

Source: Dataquest (August 1995)

For More Information...

Geoff Ballew, Industry Analyst ..... (408) 468-8676
Internet address ..... gballew@dataquest.com
Via fax ..... (408) 954-1780

Dataquest

a company of The Dun & Bradstreet Corporation

DATAQUEST WORLDWIDE OFFICES

NORTH AMERICA

Worldwide Headquarters

Dataquest Incorporated
251 River Oaks Parkway
San Jose, California 95134-1913
United States
Phone: 1-408-468-8000
Facsimile: 1-408-954-1780

Dataquest Incorporated

Nine Technology Drive
P.O. Box 5093
Westborough, Massachusetts 01581-5093
United States
Phone: 1-508-871-5555
Facsimile: 1-508-871-6262

Dataquest Global Events

3990 Westerly Place, Suite 100
Newport Beach, California 92660
United States
Phone: 1-714-476-9117
Facsimile: 1-714-476-9969

Sales Offices:
- Washington, DC (Federal)
- New York, NY (Financial)
- Dallas, TX

LATIN AMERICA

Research Affiliates and Sales Offices:
- Buenos Aires, Argentina
- Sao Paulo, Brazil
- Santiago, Chile
- Mexico City, Mexico

EUROPE

European Headquarters

Dataquest Europe Limited
Holmers Farm Way
High Wycombe, Bucks HP12 4XH
United Kingdom
Phone: +44 1494 422 722
Facsimile: +44 1494 422 742

Dataquest Europe SA

Immeuble Défense Bergères
345, avenue Georges Clémenceau
TSA 40002
92882 - Nanterre CTC Cedex 9
France
Phone: +33 1 41 35 13 00
Facsimile: +33 1 41 35 13 13

Dataquest GmbH

Kronstadter Strasse 9
81677 München
Germany
Phone: +49 89 93 09 09 0
Facsimile: +49 89 93 03 27 7

Sales Offices:
- Brussels, Belgium
- Kfar Saba, Israel
- Milan, Italy
- Randburg, South Africa
- Madrid, Spain

JAPAN

Japan Headquarters

Dataquest Japan K.K.
Shinkawa Sanko Building
6th Floor
1-3-17, Shinkawa
Chuo-ku, Tokyo 104
Japan
Phone: 81-3-5566-0411
Facsimile: 81-3-5566-0425

ASIA/PACIFIC

Asia/Pacific Headquarters

7/F China Underwriters Centre
88 Gloucester Road
Wan Chai
Hong Kong
Phone: 852-2824-6168
Facsimile: 852-2824-6138

Dataquest Korea

Suite 2407, Trade Tower
159 Samsung-dong, Kangnam-gu
Seoul 135-729
Korea
Phone: 822-551-1331
Facsimile: 822-551-1330

Dataquest Taiwan

11F-2, No. 188, Section 5
Nan King East Road
Taipei
Taiwan, R.O.C.
Phone: 8862-756-0389
Facsimile: 8862-756-2663

Dataquest Singapore

105 Cecil Street
#06-01/02 The Octagon
Singapore 0106
Phone: 65-227-1213
Facsimile: 65-227-4607

Dataquest Thailand

12/F, Vanissa Building
29 Soi Chidlom
Ploenchit Road
Patumwan, Bangkok 10330
Thailand
Phone: 662-655-0577
Facsimile: 662-655-0576

Research Affiliates and Sales Offices:
- Melbourne, Australia
- Beijing, China

Dataquest

a company of
The Dun & Bradstreet Corporation

Market Analysis

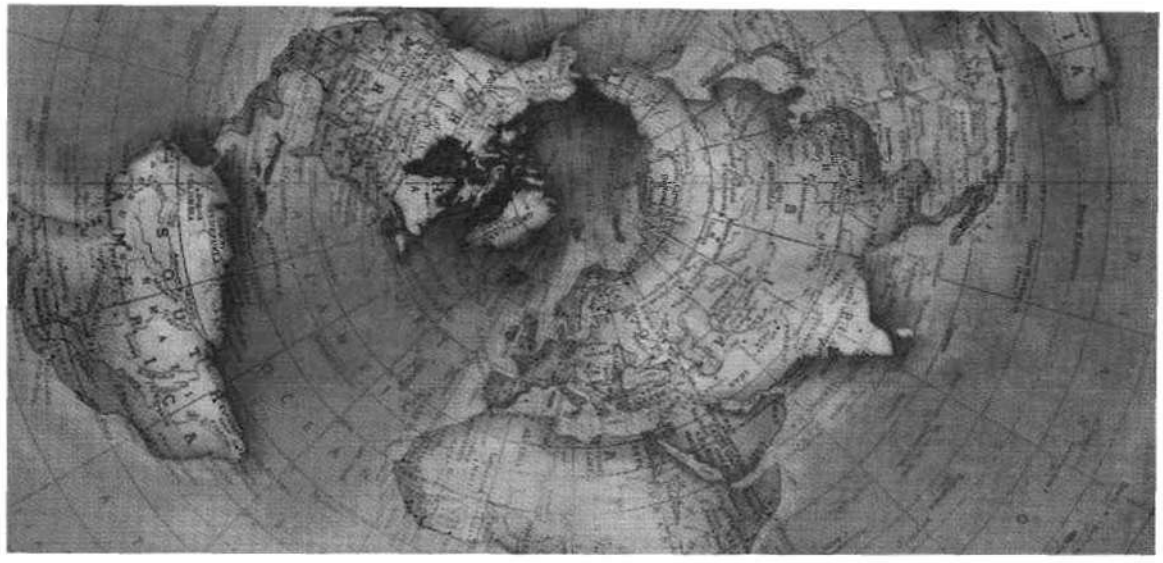

Semiconductor Directions in PCs

Market Analysis

Quarterly x86 Compute Microprocessor Forecast

Abstract: *This microprocessor forecast is an update to the previous forecast dated June 1995. The scope of this document is limited to compute applications for x86 microprocessors and does not include embedded applications. The forecast includes considerable changes in the market opportunity for Pentium and Pentium Pro microprocessors. There are two principal reasons for this. First, the PC market appears to be more robust than forecast, resulting in greater demand for MPUs. Second, new information about the Pentium Pro's performance when running 16-bit software has reduced the market opportunity for that MPU until the installed base of software shifts to 32-bit code or the performance issues of running 16-bit code on the Pentium Pro are resolved.*
By Nathan Brookwood

Higher Expectations for Unit Shipments

Dataquest increased its x86 unit forecast through the end of the decade, compared with the June forecast, as the robust PC market continues to exceed expectations, driving increased demand for CPUs. Table 1 and Figure 1 show the anticipated unit growth for x86 CPUs and x86-based PCs through 1999. CPU unit shipments exceed PC unit shipments because of the use of x86 CPUs in non-PC compute applications, along with inventory in the channel. Name-brand OEM x86 PCs correlate with complete systems sold by "traditional" PC vendors, as tracked by Dataquest. Other x86 PCs include PCs assembled by "no-name" resellers and assemblers, along with chip and motherboard upgrade purchases.

Dataquest

Program: Semiconductor Directions in PCs
Product Code: PSAM-WW-MA-9504
Publication Date: November 27, 1995
Filing: Market Analysis

Table 1
x86 Personal Computer and x86 Compute Microprocessor Shipment Comparison (Thousands of Units)

	1993	1994	1995	1996	1997	1998	1999	CAGR (%) 1994-1999
x86 MPU for Compute Applications	45,575	52,593	65,573	73,500	80,850	96,000	105,500	14.9
Name-Brand OEM x86 PCs	34,464	43,621	52,729	62,308	71,633	82,715	94,500	16.7
Other x86 PCs	2,441	3,998	4,990	6,040	7,125	8,468	9,779	19.6
Total x86 PCs	36,905	47,619	57,719	68,348	78,758	91,183	104,279	17.0
Difference between x86 MPU and x86 PCs	8,670	4,974	7,854	5,152	2,092	4,817	1,221	

Source: Dataquest (November 1995)

Figure 1
x86 Personal Computer and MPU Shipment Comparison

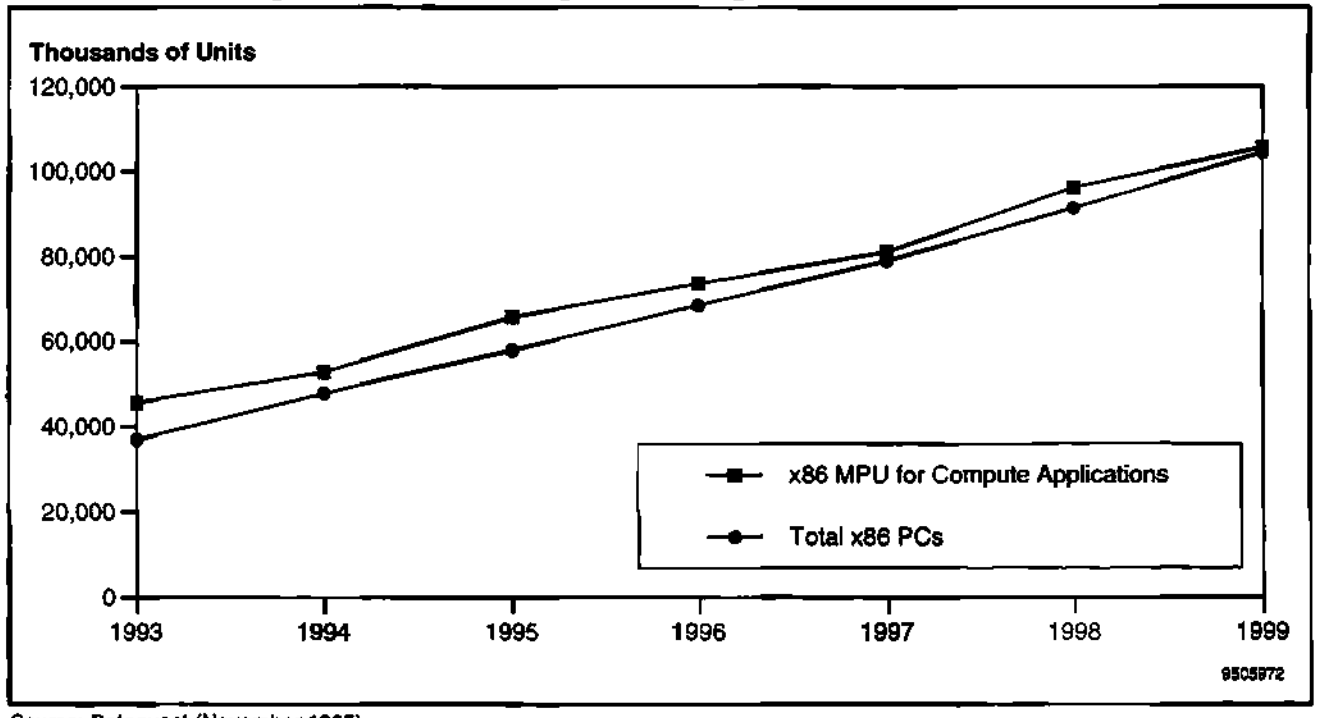

Source: Dataquest (November 1995)

Pentium Pro Will Ramp More Slowly than Anticipated

Dataquest reduced its Pentium Pro (formerly P6) forecast to 100,000 units for the last quarter of 1995 and to 2.5 million units for all of 1996. While we continue to have the highest expectations for the eventual success of this next-generation architecture, we believe demand for this product will be lower than previously expected, owing largely to the somewhat disappointing 16-bit performance of this CPU. In previous intergenerational transitions, desktop users could count on improved performance from the latest generation and demand was limited only by the high initial price of the newest products. The Pentium Pro delivers this performance boost only in

full 32-bit environments (such as Windows NT and UNIX), but in 16-bit environments (such as DOS, Windows 3.1, and even much of Windows 95), performance gains are marginal and sometimes nonexistent when compared to a far less costly system based on earlier Pentium designs. Thus, the high-value, performance-sensitive desktop applications that previously migrated first to the newest technology, regardless of cost, will not migrate to Pentium Pro in large numbers.

Servers have traditionally been a key market segment for early deployment of more powerful CPUs, and Intel clearly targeted the Pentium Pro at this market as well, with support for four-way multiprocessor configurations built into the basic Pentium Pro design. Technical concerns have forced Intel to constrain initial server configurations to single- and dual-processor systems only, and it appears that these technical issues will not be fully resolved before the second quarter of 1996. This will limit initial Pentium Pro demand within the enterprise server market.

In the near term, the Pentium Pro fits best in low-end enterprise server and technical workstation markets. It clearly offers compelling price/performance advantages in these segments when compared with typical RISC-based competitors. However, neither of these segments generates high unit demand for CPUs, compared with traditional desktop PC volumes. Dataquest anticipates that, even with these constraints, Intel will sell 2.5 million Pentium Pros in 1996, more than double the volume of all the projected RISC-based workstations and servers that year, but a far cry from the 5 million Pentium processors sold in the earlier generation's first year of existence.

Dataquest anticipates that demand for Pentium Pro will increase as the installed base of software applications migrates from its current 16-bit orientation to 32 bits. Even so, a few key components of Windows 95 system code continue to operate in 16-bit mode, and these components will drag down overall system performance. Future versions of Windows 95 will, we hope, take on an increasingly 32-bit character and will go a long way toward allowing Pentium Pro systems to demonstrate clear performance advantages for mass-market desktop applications, when compared with Pentium processors at similar clock speeds. We see such a scenario unfolding in 1997.

Table 2 and Figure 2 show Dataquest's unit forecast for Pentium and Pentium Pro CPUs by quarter through the end of 1996. Table 3 combines a look back at historic unit shipments for earlier x86 generations with a look ahead to shipments over the next four years. Figure 3 is a graph of x86 life cycles by generation.

Table 2
586 and 686 Compute MPU Forecast by Quarter (Thousands of Units)

	Q1/95	Q2/95	Q3/95	Q4/95	Q1/96	Q2/96	Q3/96	Q4/96
586-Class	4,807	6,649	9,150	12,000	14,000	15,000	16,000	17,000
686-Class	0	0	0	100	250	500	750	1,000

Source: Dataquest (November 1995)

Figure 2
586 and 686 Compute MPU Forecast by Quarter

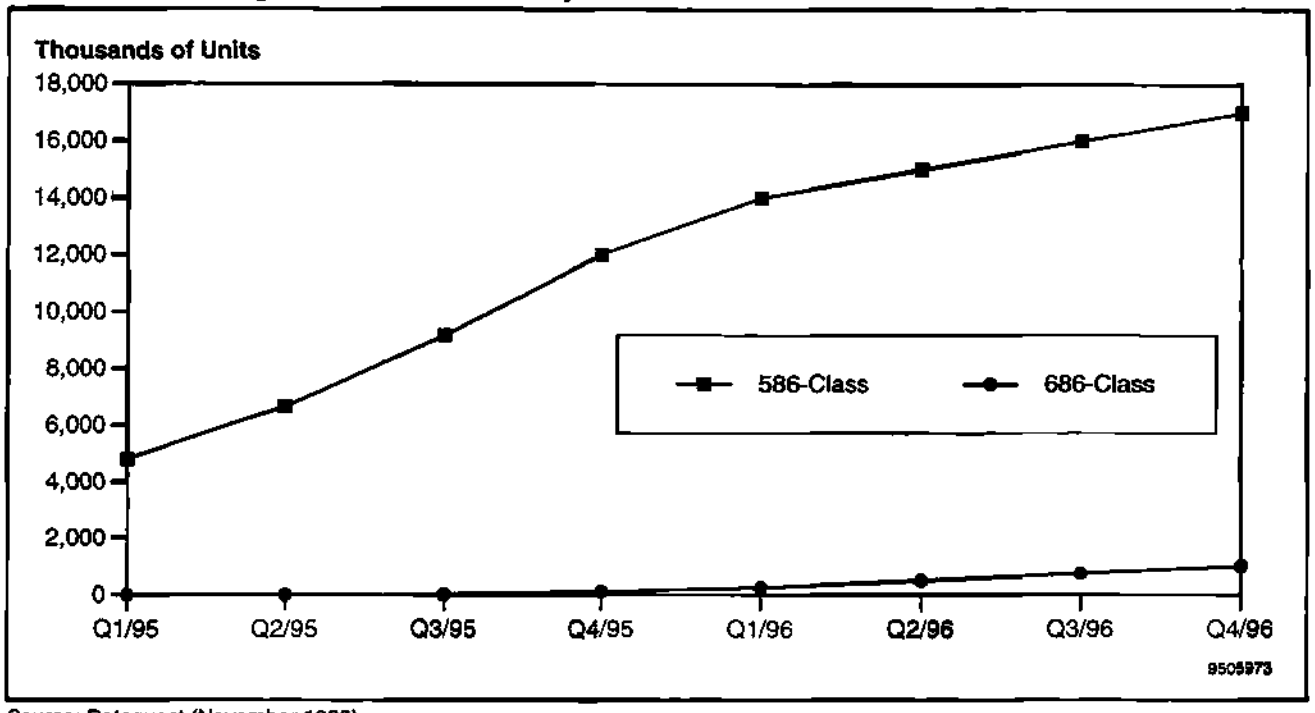

Source: Dataquest (November 1995)

Alternative x86 Vendors

Advanced Micro Devices (AMD) continues to hold the No. 2 position within the x86 CPU market, although 486 pricing is eroding rapidly and AMD's K5 program has encountered several highly visible delays. If AMD can stabilize its design and demonstrate price and or price/performance advantages in 1996, it remains well positioned from a manufacturing capacity standpoint to maintain its No. 2 position. Dataquest believes that the recently announced merger of AMD and NexGen Microsystems could allow AMD to become a serious player in the 686 generation at about the same time that Intel can begin to move the Pentium Pro into high-volume desktop market segments. Solid execution will be required on the part of AMD and NexGen to realize this potential.

Cyrix has now positioned its 5x86 M1sc processor against Intel's Pentium line and its 6x86 M1 CPU against the Pentium Pro. Shipments of both models have begun; Dataquest anticipates that Cyrix can ship about 500,000 5x86 units in 1995 and fewer than 100,000 6x86 during that period. Cyrix's ability to deliver competitive-performance 6x86 devices economically awaits the successful introduction of the die-shrunk version of the M1 in late 1995 or early 1996.

Table 3
x86 Product Forecast for Computing Applications

	1993	1994	1995	1996	1997	1998	1999	CAGR (%) 1994-1999
8086/186/286								
Units (K)	218	0	0	0	0	0	0	NM
ASP ($)	9	NM	NM	NM	NM	NM	NM	NM
Revenue ($K)	1,924	0	0	0	0	0	0	NM
80386-Class								
Units (K)	11,435	3,776	50	0	0	0	0	-100.0
ASP ($)	33	23	20	NM	NM	NM	NM	NM
Revenue ($K)	377,355	86,848	1,000	0	0	0	0	-100.0
80486-Class								
Units (K)	33,597	43,752	32,818	9,000	500	0	0	-100.0
ASP ($)	190	152	101	65	45	NM	NM	NM
Revenue ($K)	6,398,114	6,644,558	3,321,878	585,000	22,500	0	0	-100.0
80586-Class								
Units (K)	325	5,065	32,605	62,000	60,000	42,000	5,000	-0.3
ASP ($)	810	493	325	250	200	150	100	-27.3
Revenue ($K)	263,250	2,497,250	10,596,650	15,500,000	12,000,000	6,300,000	500,000	-27.5
80686-Class								
Units (K)	0	0	100	2,500	20,000	50,000	55,000	NM
ASP ($)	NM	NM	950	750	500	350	200	NM
Revenue ($K)	0	0	95,000	1,875,000	10,000,000	17,500,000	11,000,000	NM
80786-Class								
Units (K)	0	0	0	0	350	4,000	45,000	NM
ASP ($)	NM	NM	NM	NM	900	750	450	NM
Revenue ($K)	0	0	0	0	315,000	3,000,000	20,250,000	NM

(Continued)

Table 3 (Continued)
x86 Product Forecast for Computing Applications

	1993	1994	1995	1996	1997	1998	1999	CAGR (%) 1994-1999
80886-Class								
Units (K)	0	0	0	0	0	0	500	NM
ASP ($)	NM	NM	NM	NM	NM	NM	900	NM
Revenue ($K)	0	0	0	0	0	0	450,000	NM
Total Compute x86								
Units (K)	45,575	52,593	65,573	73,500	80,850	96,000	105,500	14.9
ASP ($)	154	175	214	244	276	279	305	11.7
Revenue ($K)	7,040,643	9,228,656	14,014,528	17,960,000	22,337,500	26,800,000	32,200,000	28.4
Name-Brand OEM x86 PCs (K Units)	34,464	43,621	52,729	62,308	71,633	82,715	94,500	16.7
Other x86 PCs (K Units)	2,441	3,998	4,990	6,040	7,125	8,468	9,779	19.6
Total x86 PCs (K Units)	36,905	47,619	57,719	68,348	78,758	91,183	104,279	
Difference between x86 MPU and x86 PCs	8,670	4,974	7,854	5,152	2,092	4,817	1,221	

NM = Not meaningful
Source: Dataquest (November 1995)

Figure 3
Life Cycles of x86 Products (Thousands of Units)

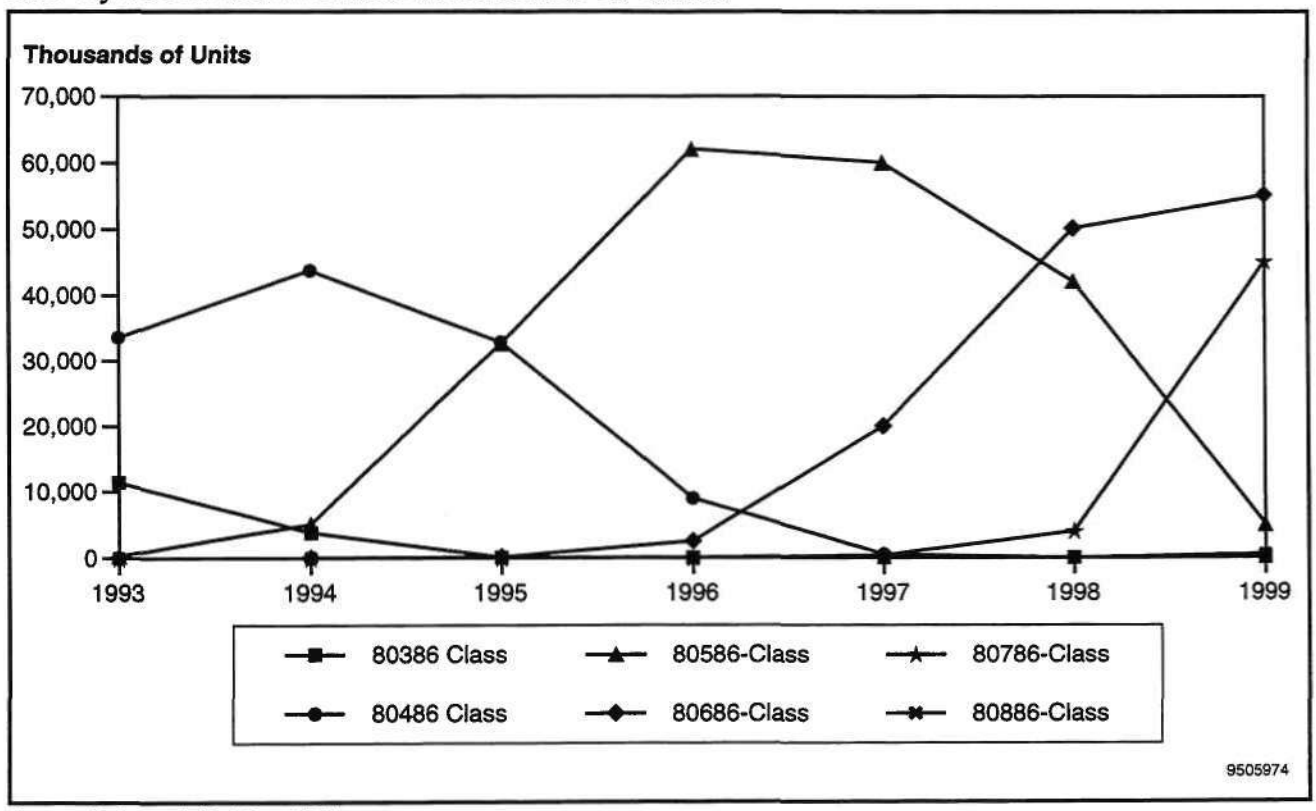

Source: Dataquest (October 1995)

For More Information...
Nathan Brookwood, Principal Analyst..... (408) 468-8537
Internet address.....nbrookwood@dataquest.com
Via fax ..... (408) 954-1780

Dataquest
a company of The Dun & Bradstreet Corporation

Market Analysis

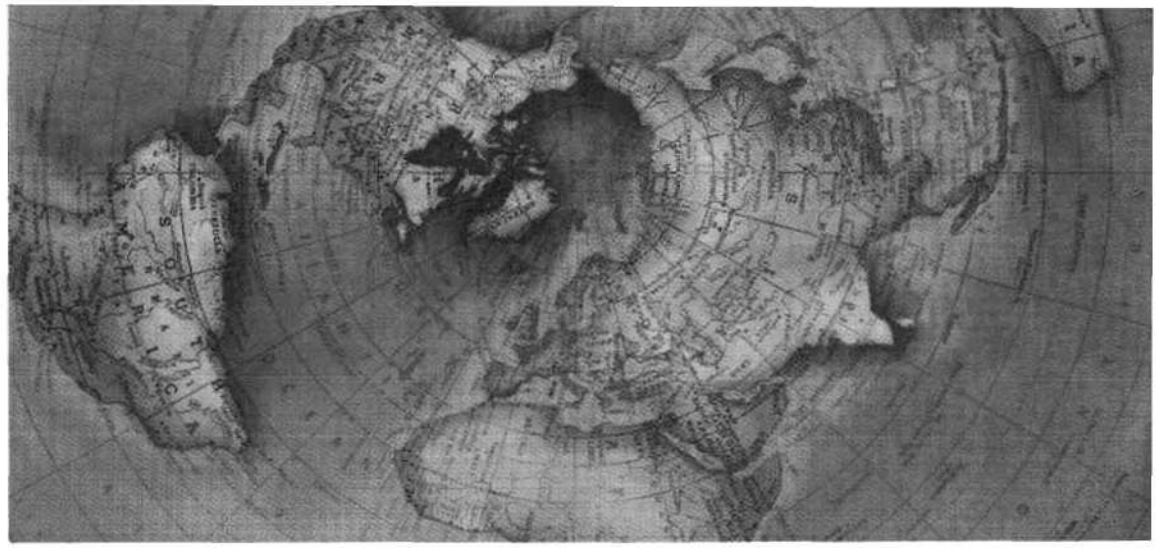

Semiconductor Directions in PCs

Market Analysis

PC OEMs Shift Cost of Memory to PC Buyers

Abstract: *The DRAM and PC markets are inherently intertwined: PCs consume more DRAM than any other application and DRAM is the largest single cost in a typical PC. The current dynamic between these two markets is increasingly interdependent because unit growth for PCs and their increasing memory requirements continue to create demand for DRAM. This high demand has kept prices up and has created a growing point-of-sale market for PC memory because PC OEMs are caught between consumer demands for more memory and relatively constant average selling prices for PCs. The average PC appears to have less memory than the average buyer demands, so PC buyers are adding more megabytes at the point of sale.*
By Geoff Ballew

PC and DRAM Markets Are Still Strongly Correlated

The high correlation between DRAM and personal computer revenue is tied to the fact that DRAM represents the largest single cost in a typical PC. When DRAM is in short supply, as it is now, PC OEMs are squeezed between demands for larger memory sizes and relatively constant (or even declining) PC prices. As process technologies improve, semiconductor devices are supposed to become smaller, faster, and cheaper, but the severe DRAM shortage has caused prices for many DRAM devices to increase from 1994 to 1995 rather than decrease. DRAM prices are expected to ease slightly from 1995 to 1996 and then follow a more traditional model where prices decline appreciably over time.

The higher DRAM prices affect the PC industry because most other components (notably microprocessors) are declining rapidly in price for a given performance. Larger microprocessors generally require larger amounts of memory, so computer buyers expect more memory to be offered in the basic models with relatively constant average selling prices. As long as

Dataquest

Program: Semiconductor Directions in PCs
Product Code: PSAM-WW-MA-9503
Publication Date: August 14, 1995
Filing: Market Analysis

DRAM prices were declining, PC OEMs could install more memory each year for about the same cost per system. Average selling prices for PCs have remained relatively constant over time because OEMs add as many features to new models as they can without increasing or compromising the historical price points. This strategy works well as long as component prices decline for a given capacity or level of performance. However, when the new models have a faster MPU, better graphics, and a larger rigid disk drive, users want more memory to complete the package. If memory prices do not decline at the rate that other components' prices decline, then OEMs are faced with either spending more per system on DRAM or configuring their models with less DRAM than buyers expect. Dataquest believes that the typical PC is underconfigured when it is shipped from the factory, and this creates increasing demand for add-in memory at the point of sale of the PC.

Tables 1 and 2 and Figures 1, 2, and 3 show the latest forecasts for worldwide PC unit shipments and revenue.

Table 1
Worldwide PC Market Shipment Forecast by Product Type (Thousands of Units)

	1994	1995	1996	1997	1998	1999	CAGR (%) 1994-1999
Desktop/Deskside	39,219	45,984	53,056	58,815	66,421	74,419	13.7
Mobile	8,664	11,278	14,656	18,448	22,134	25,581	24.2
Total	47,883	57,262	67,712	77,263	88,555	100,000	15.9
Percentage Growth	22.3	19.6	18.2	14.1	14.6	12.9	

Source: Dataquest (August 1995)

Figure 1
Worldwide PC Market Shipment Forecast by Product Type (Thousands of Units)

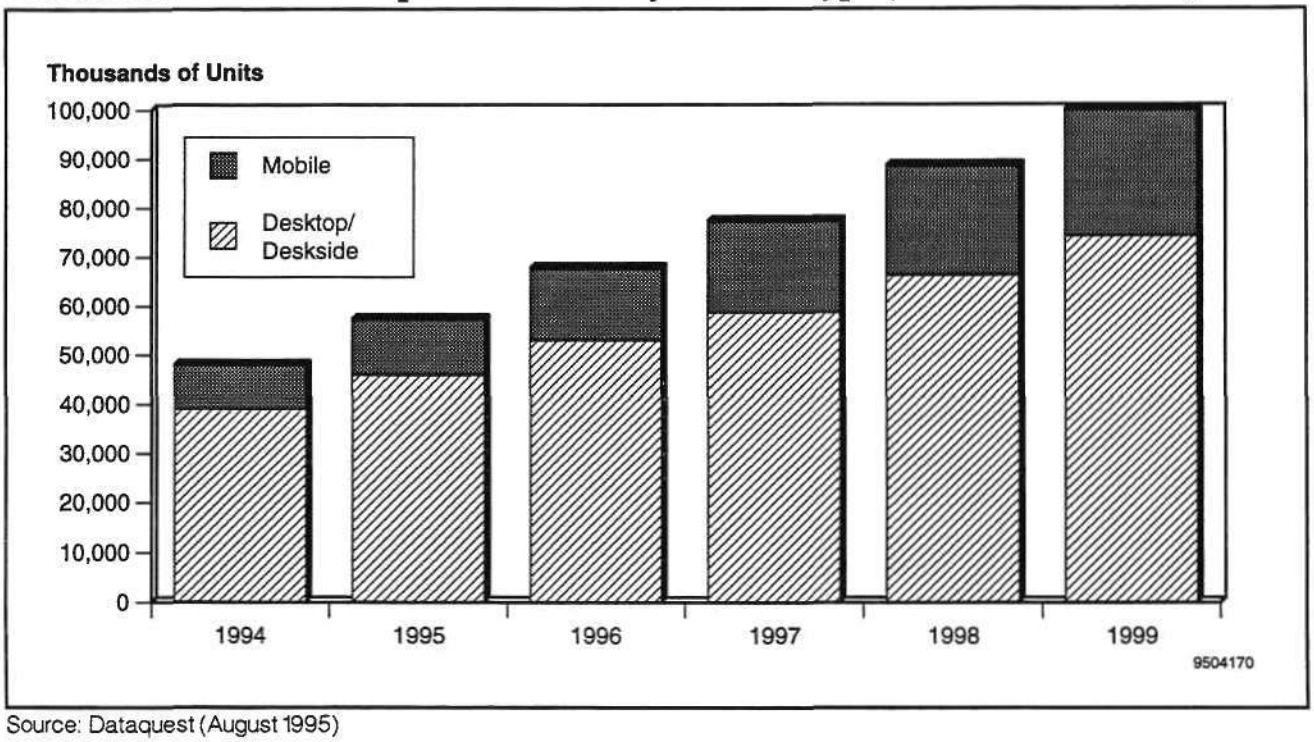

Source: Dataquest (August 1995)

Figure 2
Worldwide PC Market Revenue Forecast

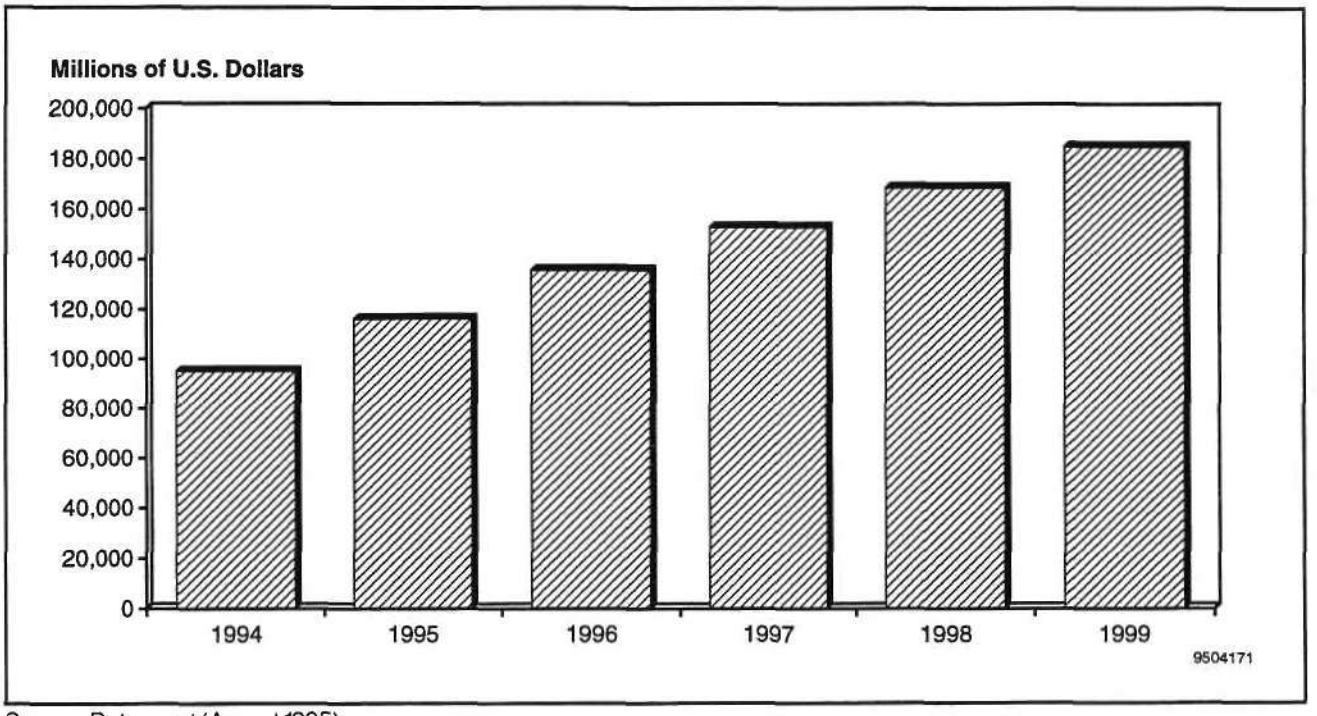

Source: Dataquest (August 1995)

Figure 3
Worldwide PC Shipment Forecast by Microprocessor Type

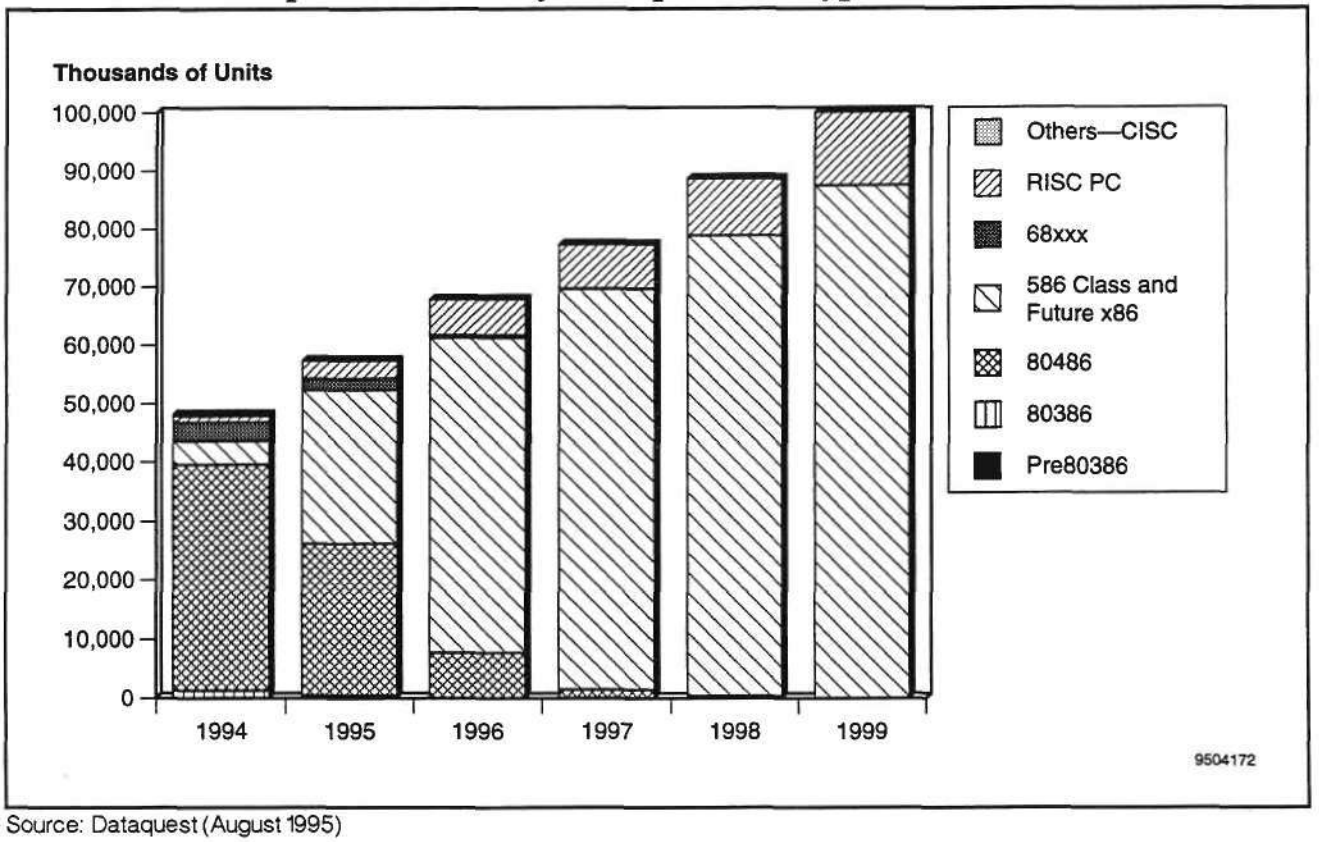

Source: Dataquest (August 1995)

Figure 3 shows PC unit shipments by microprocessor, with the numerical data in Table 2. The lion's share of PC units are forecast to have 586-class- and future x86 MPUs. Intel's rapid introduction of faster versions of the P54C (120-MHz and 133-MHz versions this spring alone) and aggressive pricing ensure that shipments of the Pentium MPU will increase at a record pace. The P6 is not broken out from the 586-class and future x86 category, but will have shipments of a couple of hundred thousand units by the end of this year and rapid growth in 1996 and 1997.

Dataquest Perspective

As a consequence of the short supply of DRAM, PC manufacturers are likely to produce PCs with less memory than the average user requires. This has the effect of passing on the cost of the memory to the consumer without significantly increasing the advertised price for the PC. Those buyers who demand additional memory will pay for it themselves, increasing the total market for point-of-sale memory revenue, which is not reflected in system unit prices. It is a "hidden" way of increasing the average selling price without impacting standard configuration prices. An additional benefit for the PC manufacturer is greater inventory flexibility. By reducing the total amount of DRAM their assembly operations consume, these manufacturers have reduced the probability that a DRAM inventory shortage will affect their production schedule.

Semiconductor vendors are advised to consider the point-of-sale upgrade market when planning their products. DRAM vendors should see a larger market outside of their large OEM accounts as more memory is added at the point of sale of PCs. Other semiconductor vendors may realize some benefit from this dynamic as PC OEMs search for ways to add features and value to their standard models without requiring more DRAM. Some examples might be a faster modem, better graphics controller, or audio features.

Figures 4 and 5 show a forecast for the average megabytes per system installed by the PC manufacturer (factory-installed megabytes) and the average megabytes per system purchased by the user at the point of sale (POS megabytes). The gap between the top line on both figures (total

Table 2
Worldwide PC Shipment Forecast by Microprocessor Type (Thousands of Units)

	1994	1995	1996	1997	1998	1999	CAGR (%) 1994-1999
Pre80386	32	0	0	0	0	0	-100.0
80386	1,323	530	0	0	0	0	-100.0
80486	38,261	25,748	7,801	1,490	494	0	-100.0
586-Class and Future x86	4,004	26,100	53,500	68,093	78,421	87,500	85.3
68xxx	3,214	1,950	426	140	0	0	-100.0
RISC PC	1,044	2,934	5,985	7,540	9,640	12,500	64.3
Others – CISC	3	0	0	0	0	0	-100.0
Total	47,883	57,262	67,712	77,263	88,555	100,000	

Source: Dataquest (August 1995)

Figure 4
Megabytes per System for Mobile PCs

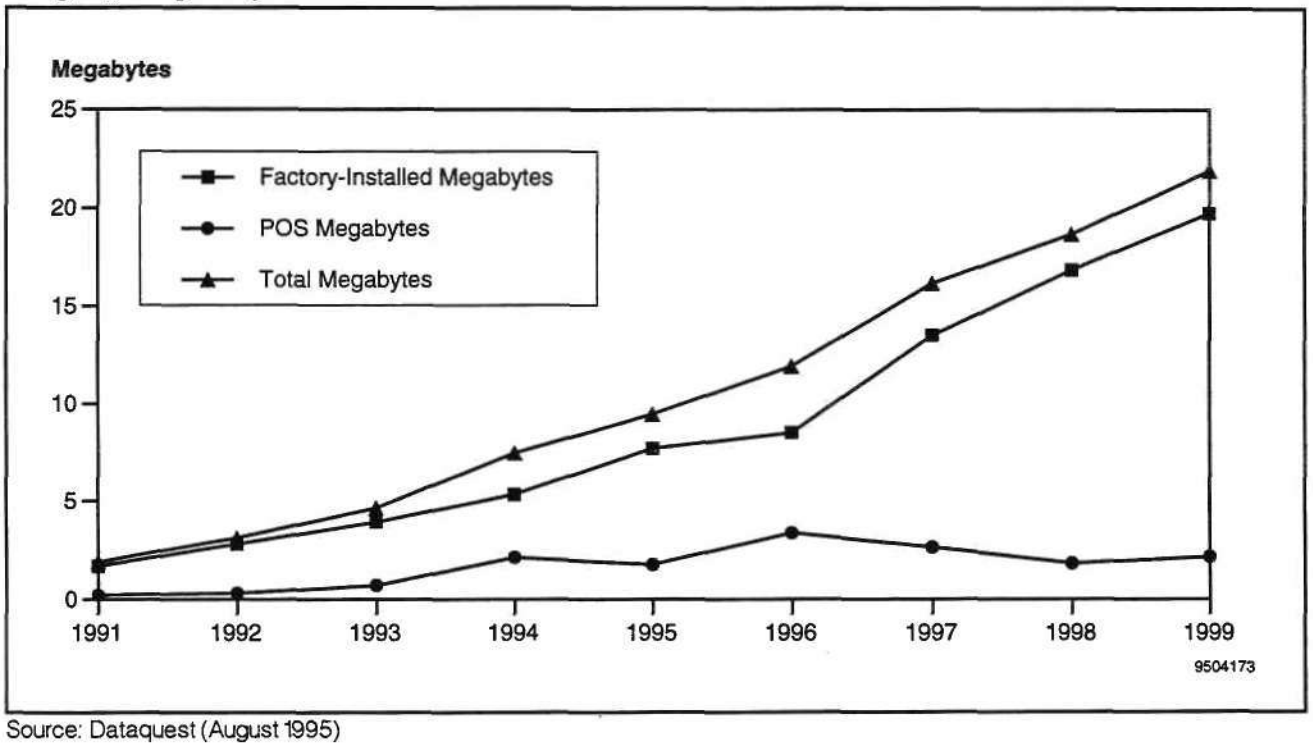

Source: Dataquest (August 1995)

Figure 5
Megabytes per System for Desktop PCs

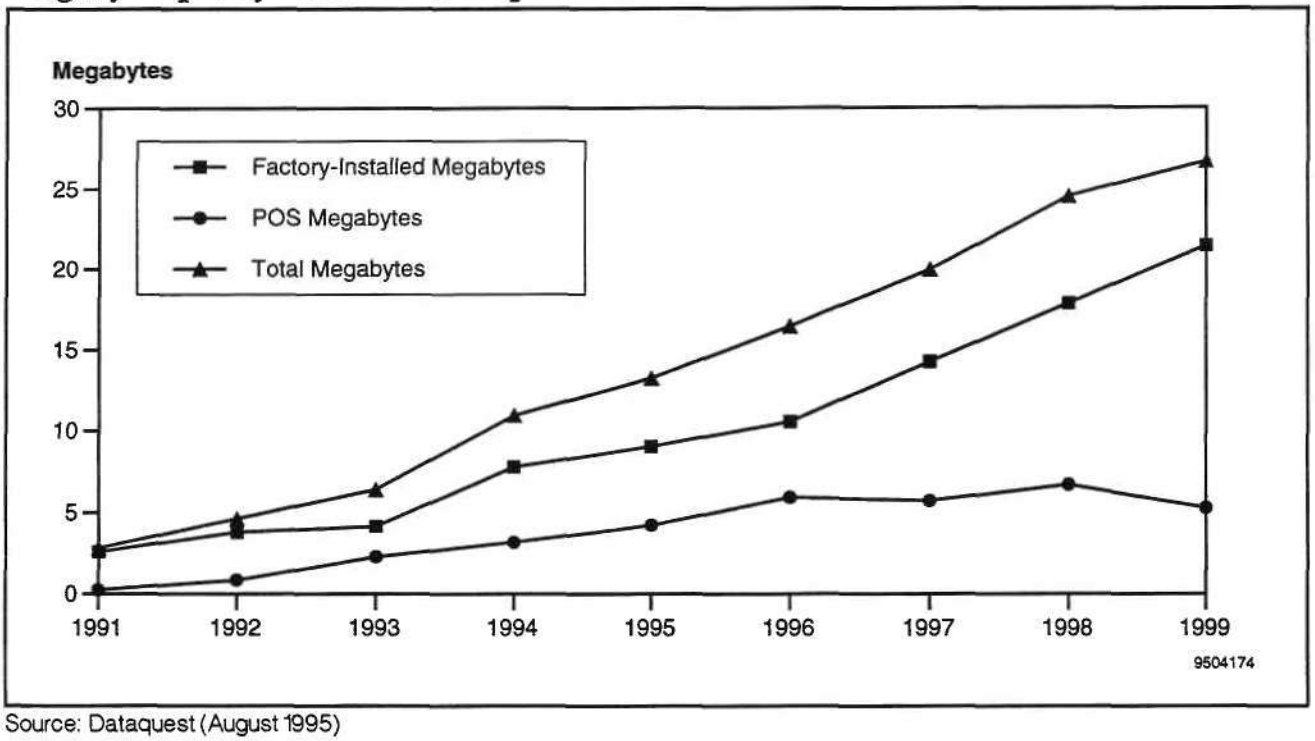

Source: Dataquest (August 1995)

megabytes) and the factory-installed megabytes line is equal to the POS megabytes line. Figure 5 shows this best—the gap between total megabytes and factory-installed megabytes starts to increase in 1993 but gets wider and wider through 1994, 1995, and 1996. This gap represents the additional memory that PC buyers must purchase to have the systems they want.

For More Information...

Geoff Ballew, Industry Analyst ..... (408) 437-8676
Internet address ..... gballew@dataquest.com
Via fax ..... (408) 437-0292

Market Analysis

Semiconductor Directions in PCs

Market Analysis

Graphics Controllers: Who Is at the Top?

Abstract: *This newsletter provides the latest market share and unit shipment statistics for graphics controllers in desktop and mobile PCs. These final statistics are a preview of the information in the upcoming Graphics Focus Report, which will provide market analysis in addition to detailed market statistics.*
By Geoff Ballew

1994's Rankings Look a Lot Like 1993's

The 1994 market share rankings do not show much change from 1993. The top four companies in each category (mobile units, desktop units, and total units) in 1993 remained in the top four in 1994, although there was some jockeying for position: S3 took the No. 3 spot from Tseng Labs in the desktop market, and Western Digital overwhelmed Cirrus Logic to take the No. 1 slot on the mobile side with over 100 percent growth.

Figure 1 and Table 1 show the company rankings for all desktop and mobile PC graphics controller unit shipments. Figure 2 and Table 2 show the company rankings for all desktop PC graphics controller unit shipments. Figure 3 and Table 3 show the company rankings for all mobile PC graphics controller unit shipments.

Look for detailed market analysis of this data in the upcoming Graphics Focus Report. These statistics are published in this newsletter to give our clients access to them as quickly as possible.

Dataquest

Program: Semiconductor Directions in PCs
Product Code: PSAM-WW-MA-9502
Publication Date: July 31, 1995
Filing: Market Analysis

Figure 1
Graphics Controller Unit Shipments Worldwide

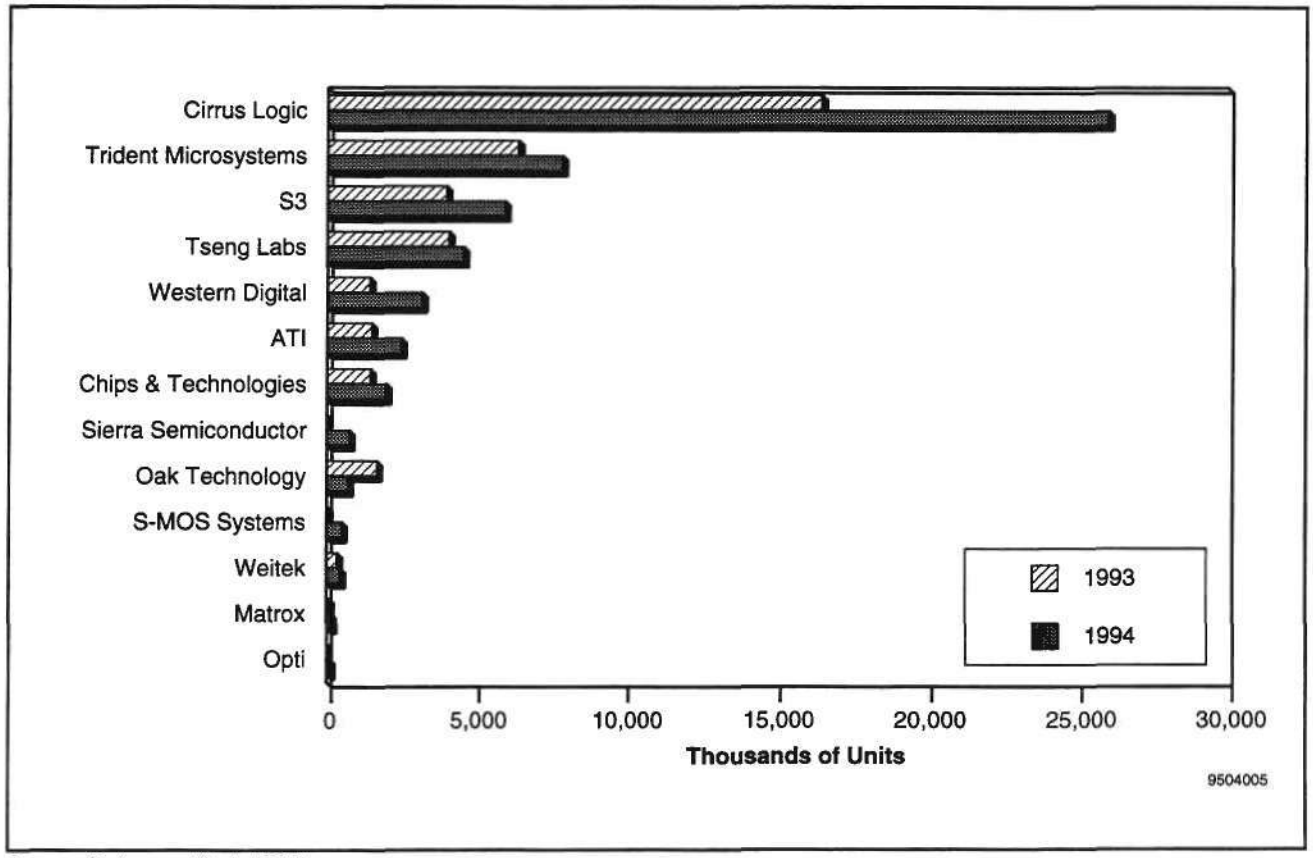

Source: Dataquest (July 1995)

Table 1
Graphics Controller Unit Shipments Worldwide (Thousands of Units)

1994 Rank	1993 Rank		1993 Units	1994 Units	Percentage Change	1994 Market Share (%)
1	1	Cirrus Logic	16,500	26,000	57.6	47.4
2	2	Trident Microsystems	6,436	7,903	22.8	14.4
3	4	S3	4,000	6,000	50.0	10.9
4	3	Tseng Labs	4,100	4,600	12.2	8.4
5	8	Western Digital	1,430	3,175	122.0	5.8
6	6	ATI	1,500	2,500	66.7	4.6
7	7	Chips & Technologies	1,461	2,013	37.8	3.7
8	13	Sierra Semiconductor	0	750	NA	1.4
9	5	Oak Technology	1,700	705	-58.5	1.3
10	11	S-MOS Systems	25	500	1,900.0	0.9
11	9	Weitek	350	452	29.1	0.8
12	10	Matrox	75	150	100.0	0.3
13	12	Opti	0	100	NA	0.2
		Total Market	37,577	54,848	46.0	100.0

NA = Not applicable
Source: Dataquest (July 1995)

Figure 2
Desktop Graphics Controller Unit Shipments Worldwide

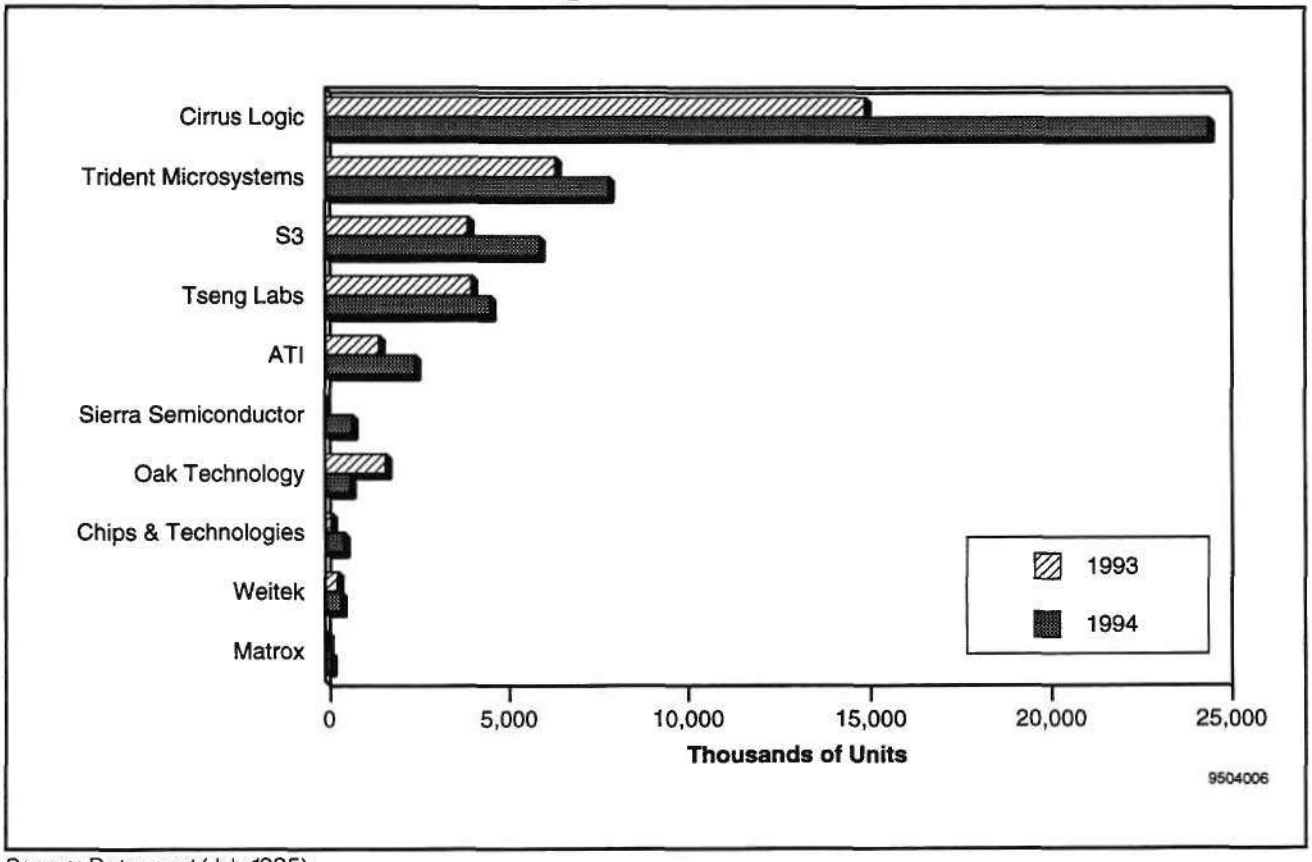

Source: Dataquest (July 1995)

Table 2
Desktop Graphics Controller Unit Shipments Worldwide (Thousands of Units)

1994 Rank	1993 Rank		1993 Units	1994 Units	Percentage Change	1994 Market Share (%)
1	1	Cirrus Logic	15,000	24,500	63.3	51.0
2	2	Trident Microsystems	6,436	7,903	22.8	16.4
3	4	S3	4,000	6,000	50.0	12.5
4	3	Tseng Labs	4,100	4,600	12.2	9.6
5	6	ATI	1,500	2,500	66.7	5.2
6	10	Sierra Semiconductor	0	750	NA	1.6
7	5	Oak Technology	1,695	689	-59.4	1.4
8	8	Chips & Technologies	161	513	218.6	1.1
9	7	Weitek	350	452	29.1	0.9
10	9	Matrox	75	150	100.0	0.3
		Total Market	33,317	48,057	44.2	100.0

NA = Not applicable
Source: Dataquest (July 1995)

Figure 3
Mobile Graphics Controller Unit Shipments Worldwide

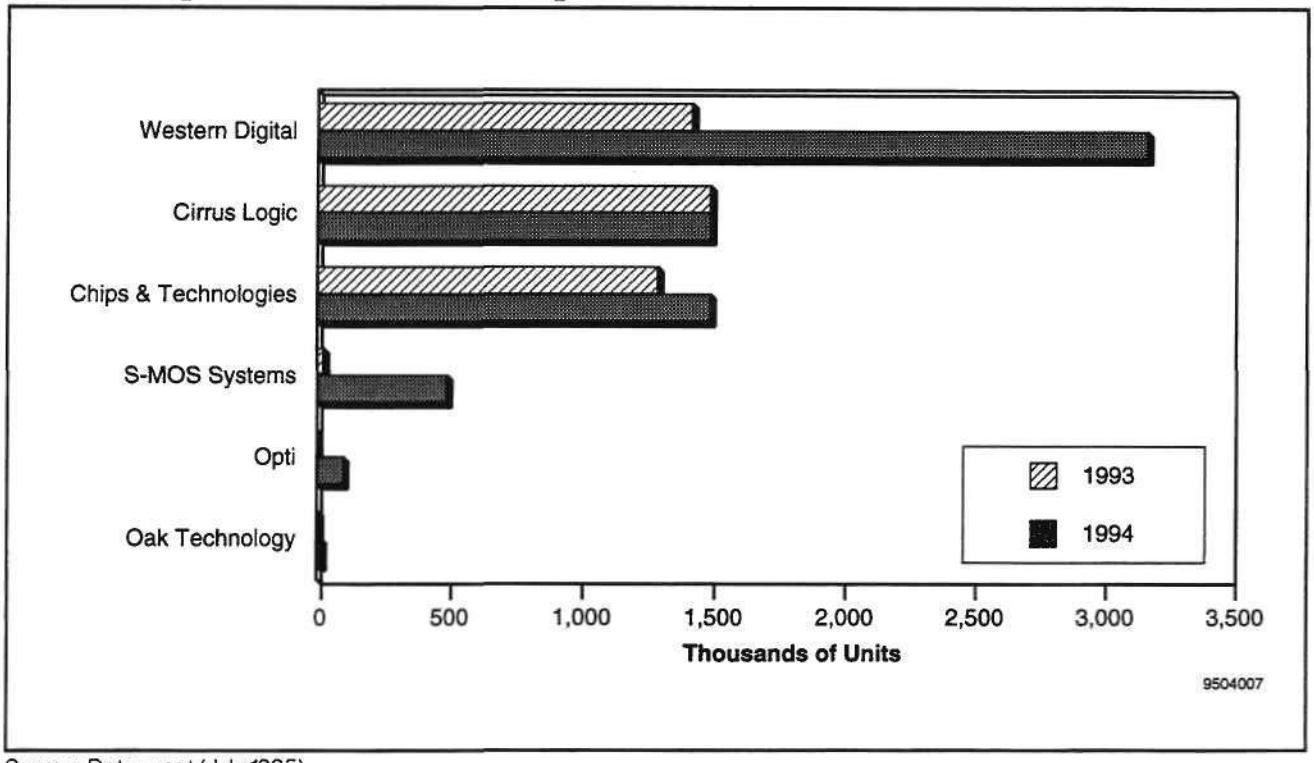

Source: Dataquest (July 1995)

Table 3
Mobile Graphics Controller Unit Shipments Worldwide (Thousands of Units)

1994 Rank	1993 Rank		1993 Units	1994 Units	Percentage Change	1994 Market Share (%)
1	2	Western Digital	1,430	3,175	122.0	46.8
2	1	Cirrus Logic	1,500	1,500	0	22.1
3	3	Chips & Technologies	1,300	1,500	15.4	22.1
4	4	S-MOS Systems	25	500	1,900.0	7.4
5	6	Opti	0	100	NA	1.5
6	5	Oak Technology	5	16	220.0	0.2
		Total Market	4,260	6,791	59.4	100.0

NA = Not applicable
Source: Dataquest (July 1995)

For More Information...
Geoff Ballew, Industry Analyst ..... (408) 437-8676
Internet address ..... gballew@dataquest.com
Via fax ..... (408) 437-0292

Dataquest
a company of The Dun & Bradstreet Corporation

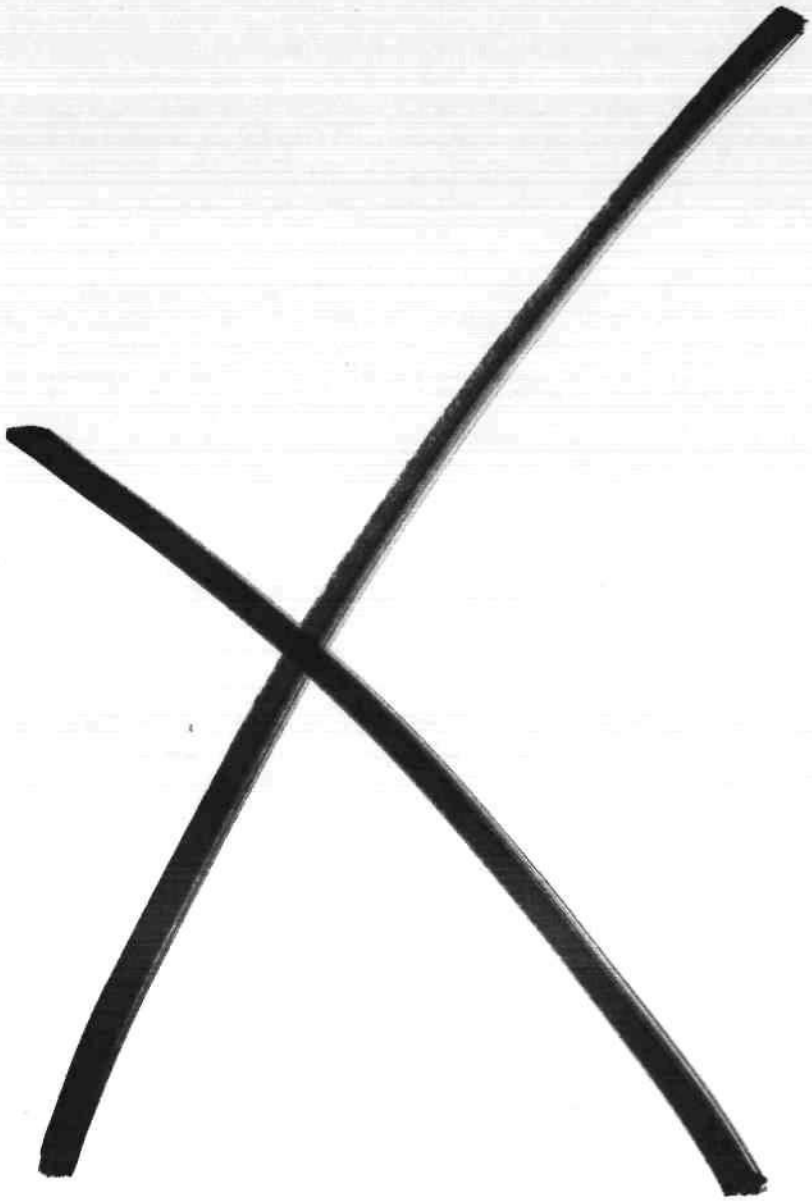

Dataquest

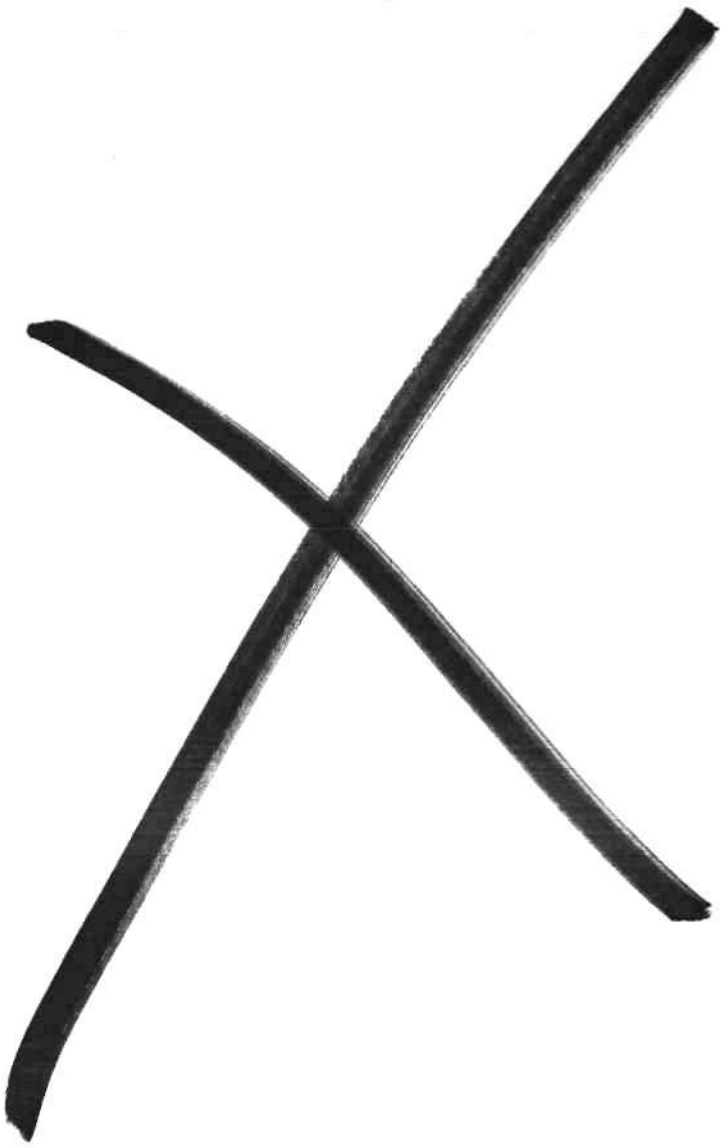

Dataquest

Perspective

PC Directions in Semiconductors

Technology Analysis

1 + 1 = 1? Integration Propels Modem and Audio Functions on a Collision Course

Abstract: *Modem and audio functions are becoming standard features on PCs. These features are checklist items for the consumer and small office/home office (SOHO) markets and are making inroads into the corporate PC market, as well. This document provides forecasts for both the modem and audio chipset markets and discusses the collision course that these two technologies are taking. Synergy between these two markets is driving integration at both the board and chip level, creating a variety of hardware options for adding these and other features to PC designs.*
By Geoff Ballew

Major Shift in Modem Chipset Market

PC user demand for connectivity is pushing modems closer to becoming a standard feature for PCs. The transition of home PCs from computing tools to communications tools, as well as changes in the business world, are driving this demand. Modems make new sources of information accessible by connecting personal computers to other computers over vast public and private networks. Figure 1 and Table 1 show Dataquest's forecast for worldwide shipments of PCs and modem chipsets. Modem chipsets are defined as one or more chips that provide the modem functions.

Note that the five-year compound annual growth rate (CAGR) for modem chipsets is lower than that for PCs. This does not mean that fewer PCs will have modem functions; in fact, quite the opposite is true. It does highlight a shift in how modem functions are sold to the end user, which in turn affects the market dynamics for modem chipsets.

Dataquest

Program: PC Directions in Semiconductors
Product Code: PSAM-WW-TA-9502
Publication Date: December 11, 1995
Filing: Competitive Dynamics

Figure 1
Worldwide Forecast for PC and Modem Chipset Markets

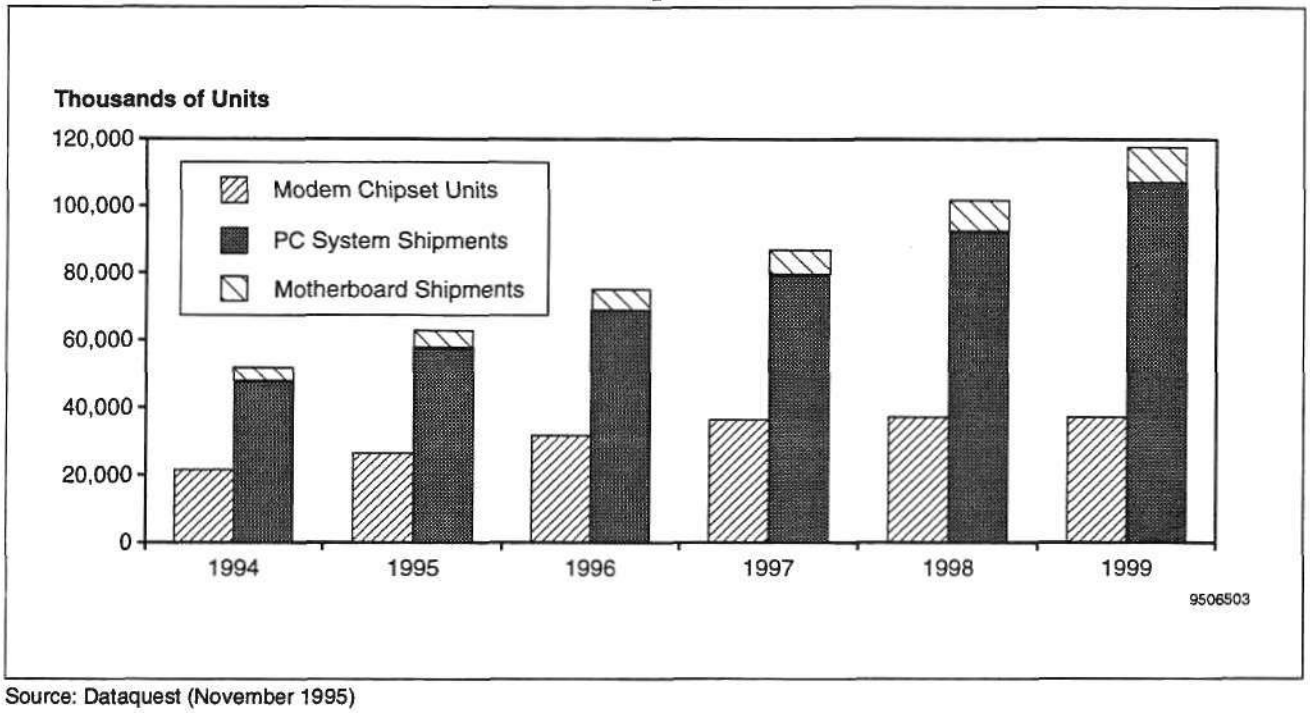

Source: Dataquest (November 1995)

Table 1
Worldwide Forecast for PC and Modem Chipset Markets

	1994	1995	1996	1997	1998	1999	CAGR (%) 1994-1999
Modem Chipset Units (K)	21,531	26,489	31,733	36,396	37,211	37,166	11.5
Modem Chipset ASP ($)	26.3	22.2	21.0	20.3	20.2	20.7	-4.7
Modem Chipset Revenue ($M)	566	587	667	738	751	770	6.3
PC System Shipments (K)	47,876	57,564	68,712	79,313	92,355	107,000	17.4
Motherboard Shipments (K)	4,404	5,476	6,738	7,992	9,582	11,190	20.5
Total PC Unit Shipments (K)	52,280	63,040	75,450	87,305	101,937	118,190	17.7
Ratio of Modem Chipsets to Total PCs (%)	41.2	42.0	42.1	41.7	36.5	31.4	

Source: Dataquest (November 1995)

Modem sales today are a combination of retail sales and sales to PC OEMs. Over the next four years, modem sales to PC OEMs will grow faster than retail sales and will be the primary driver for market growth in 1998 and 1999. For the purposes of this document, retail sale is defined as the sale of a modem directly to an end user, whether it is an internal or an external modem. The market for external modems is almost entirely a retail market; PC OEMs that provide modem functions in their PCs prefer add-in card modems.

Retail sales have been robust for two reasons. First, the increasing popularity of online services and access to the Internet has caused many new users to purchase modems for their existing PCs, which did not have modems installed. Second, many modem users upgrade their modems more frequently than their PCs because modem speeds have been increasing and the cost to upgrade is relatively small. Modem speeds are slow compared to other data transfer rates in even a low-end PC, and modem users are sensitive to that performance bottleneck.

This market is changing, however. Dataquest believes that the retail sales of modems directly to end users will decline, and that sales to PC OEMs will drive market growth at the end of this decade. There are two reasons for this. First, modems are now standard features for PCs targeted at the consumer and small office/home office (SOHO) markets, so it is unnecessary for the buyer to purchase a modem directly. Second, modem speeds are not increasing as fast as they have previously, which reduces the incentive for users to buy a new modem. Dataquest expects modem speeds to peak at 28.8 Kbps, the current standard known as V.34, for communication over standard telephone lines.

Without faster modems, the retail upgrade market will decline gradually as the installed base becomes saturated with V.34 modems. The addition of digital simultaneous voice and data (DSVD) to V.34 modems in 1996 will help postpone the decline until 1998. DSVD will allow modem users to transfer voice and data simultaneously over a single telephone line. Although some techniques for this exist today, they are less sophisticated than DSVD and involve switching between voice and data modes rather than processing both streams simultaneously. Not every modem user will have a V.34 modem, but those users willing to upgrade will not find new modems compelling if they do not have significantly enhanced features. Other users may be content with V.32bis modems, which operate at 14.4 Kbps.

Users who demand greater performance may upgrade to digital modems using integrated service digital network (ISDN) lines instead of standard analog telephone lines (plain old telephone service—POTS). ISDN modems and services are available in some locations today, but cost and availability issues will delay widespread adoption. Dataquest forecasts that 11 million ISDN lines will be in use by the year 2000, so ISDN modems will remain high-end products through that time. In other words, analog modems are expected to dominate the market at least through the end of this decade. Figure 2 shows the market positioning of modem standards through 1999.

Figure 2
Market Positioning of Modem Standards

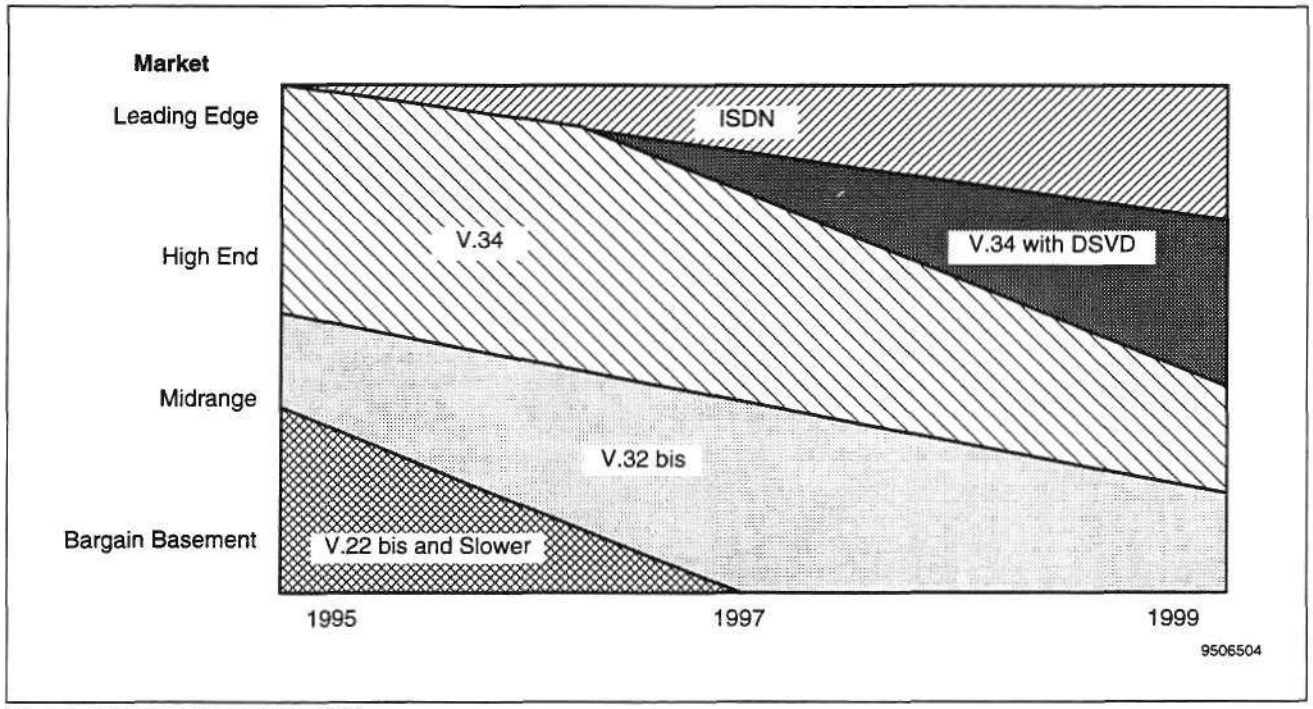

Source: Dataquest (November 1995)

Modem chipsets are rarely implemented on the motherboard because of Federal Communications Commission (FCC) certification issues. Modems are considered FCC Class A devices, while motherboards are generally FCC Class B devices. However, if a designer adds a modem chipset to a motherboard design, the motherboard becomes a Class A device and must meet the more stringent Class A specification. For this reason, modem chipsets are typically implemented on add-in cards or as external modems.

The Audio Market Marches On

Useful? Maybe. Fun? Yes. Either way, audio features are more common in PCs today than ever before. With capabilities ranging from "business" audio (that is, no Sound Blaster compatibility for games) to high-end wavetable solutions, PCs just want to be heard these days. Audio capabilities have already become standard features for PCs targeted at the consumer and home office markets. The corporate market is a tougher nut to crack because many information systems managers still question the need for audio features in the workplace. Notebook PCs are one exception to this rule; the target market is the corporate world, but audio features are becoming common. Of course, Apple Macintosh computers have always had built-in audio capabilities, whether they were selling into consumer or business accounts. Figure 3 and Table 2 show the forecast for worldwide shipments of PCs and PC audio chipsets. Audio chipsets are defined as one or more chips that provide the audio functions. Note that 1998 is a milestone; it is the first year that Dataquest believes more computers will be configured with audio features than without.

Figure 3
Forecast for PCs and PC Audio Chipsets

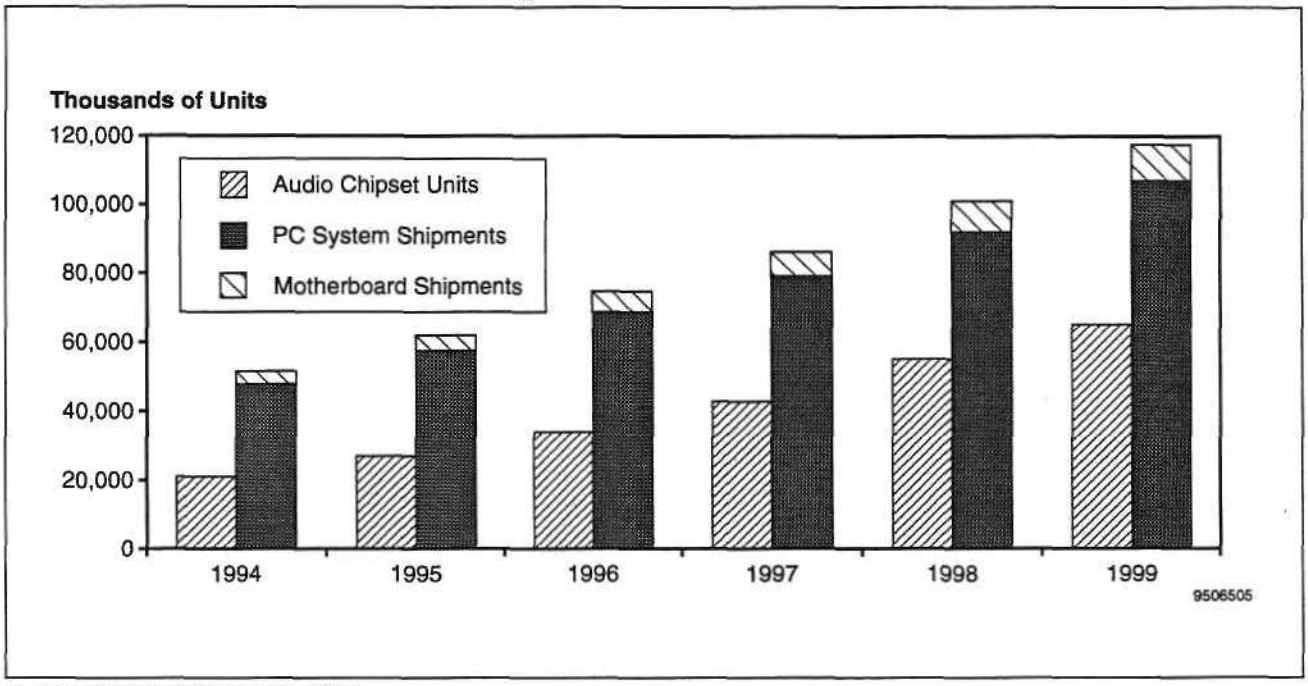

Source: Dataquest (November 1995)

Table 2
Forecast for PCs and PC Audio Chipsets

	1994	1995	1996	1997	1998	1999	CAGR (%) 1994-1999
Audio Chipset Units (K)	20,996	26,978	33,889	42,762	55,173	65,184	25.4
Audio Chipset ASP ($)	16.14	15.00	13.64	12.48	11.55	10.56	-8.1
Audio Chipset Revenue ($M)	338.9	404.7	462.2	533.7	637.2	688.3	15.2
PC System Shipments (K)	47,876	57,564	68,712	79,313	92,355	107,000	17.4
Motherboard Shipments (K)	4,404	5,476	6,738	7,992	9,582	11,190	20.5
Total PC Unit Shipments (K)	52,280	63,040	75,450	87,305	101,937	118,190	17.7
Ratio of Audio Chipsets to Total PCs (%)	40.2	42.8	44.9	49.0	54.1	55.2	

Source: Dataquest (November 1995)

The PC audio market is a mix of add-in cards and motherboard implementations. For 1995, the audio add-in card market represents about half of the total demand for audio chips, while motherboard implementations account for the other half. Motherboard implementations are heavily skewed toward Macintosh and notebook PCs, but that is changing. Multimedia PCs with x86 microprocessors will increasingly have audio functions implemented on the motherboard. Dataquest expects motherboard implementations to represent over 70 percent of the demand for audio chipsets by 1999, with add-in card implementations absorbing the remainder. In that year, 40 percent of PCs and motherboards sold will have audio functions integrated onto the motherboard, compared to only 20 percent in 1995.

The CAGR for audio chipset shipments is substantially higher than that for total PC unit shipments (25.4 percent versus 17.7 percent). This highlights the migration of audio features into a higher percentage of PCs sold.

Multifunction Solutions Will Dominate

As demand grows for audio and modem functions in PCs, these two feature sets are propelled on a collision course. The synergy between the modem and audio markets is tied to common signal-processing elements as well as competition for system resources. If a PC OEM buys an audio card and an add-in card modem, it is purchasing some semiconductor devices twice; both cards require their own codecs and amplifiers.

Competition for resources is also an issue because the modem chipset vendors have added call-processing features such as voice mailboxes and speakerphone capabilities to their chipsets. If a user wants to answer a telephone call using the modem as a speakerphone, then the modem needs control of the speakers as well as control of the microphone. But the speakers and microphone are plugged into the sound card, and nobody wants to pay for a second set of speakers and microphone for the modem. The situation is worse if the user happens to be playing a game or an audio CD and is actually using the speakers at the time the telephone rings.

Multifunction cards are the trend of the future. Board vendors have taken the first integration step by creating dual-function cards that use two chipsets—one audio chipset and one modem chipset. Figure 4 shows the layout of one of these combination cards.

This example is an Aztech audio card with a Rockwell chipset for the modem functions and other telephony features. All of these functions can be integrated into a single chipset, but there is little overlap between the audio chipset vendors and the modem chipset vendors. Rockwell has over 60 percent of the modem chipset market but does not participate in the audio chipset market. Aztech designs its own chipsets for the audio market but does not design modem chips. Two exceptions are Creative Laboratories and Cirrus Logic; both of these companies sell audio products as well as modem products. Cirrus' products are sold at the chip level, while Creative sells at both the chip level and the board level. Either one of those companies could conceivably produce a dual-function chipset that provides both feature sets by integrating its current products.

One company has already introduced a dual-function chipset. The Quartet chipset from Sierra Semiconductor provides audio, modem, fax, and speakerphone functions in a single chipset.

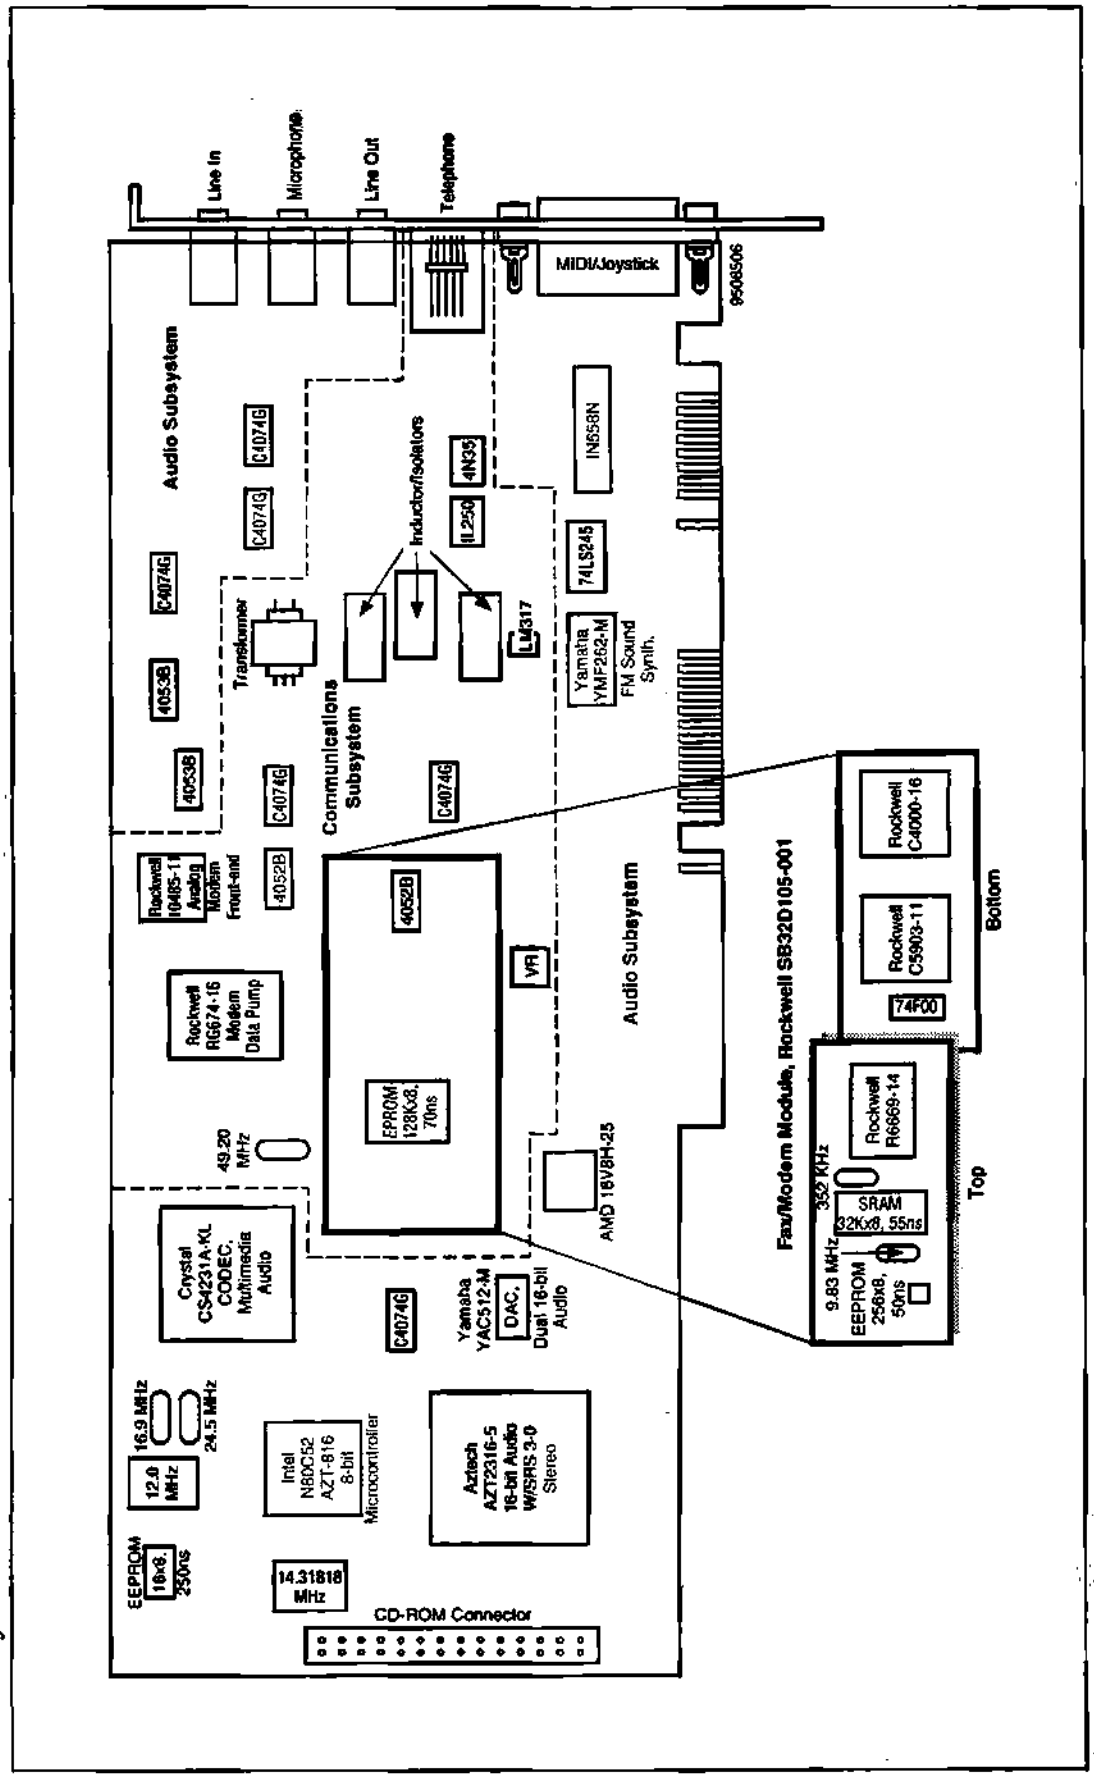

Figure 4
Board Layout for a Dual-Function Audio and Modem Card

Source: Dataquest (November 1995)

There are also alternatives to dedicated chipsets. Dedicated chipsets are defined as chipsets in which device compatibility and compliance to standards is designed at the hardware level. Several companies have stepped away from fixed-function designs and are using special processors that run software to provide specific functions and device compatibility. Two examples are known as digital-signal processors (DSPs) and media engines. Figure 5 shows the current application areas for media engines and DSPs.

Media engines and DSPs are more expensive than fixed-function chipsets, but they are also more flexible. PC OEMs and end users can add new functions simply by adding software rather than adding silicon. That level of flexibility could make them less expensive than buying several fixed-function chipsets. Media engines are generally more powerful than DSPs and can handle some graphics and motion-video processing as well as offload some tasks from the CPU. Three companies promoting media engine products are Philips Semiconductor, Chromatic Research, and MicroUnity. Philips and Chromatic are targeting the PC market for their products, which are named Trimedia and Mpact, respectively. MicroUnity appears to be targeting non-PC consumer devices such as digital set-top boxes. One example of a DSP-based solution is the Mwave chip from IBM. Mwave chips are shipping in some IBM PCs. Figure 6 shows a road map for hardware implementation trends.

Figure 5
Function Map for Alternative Hardware Solutions

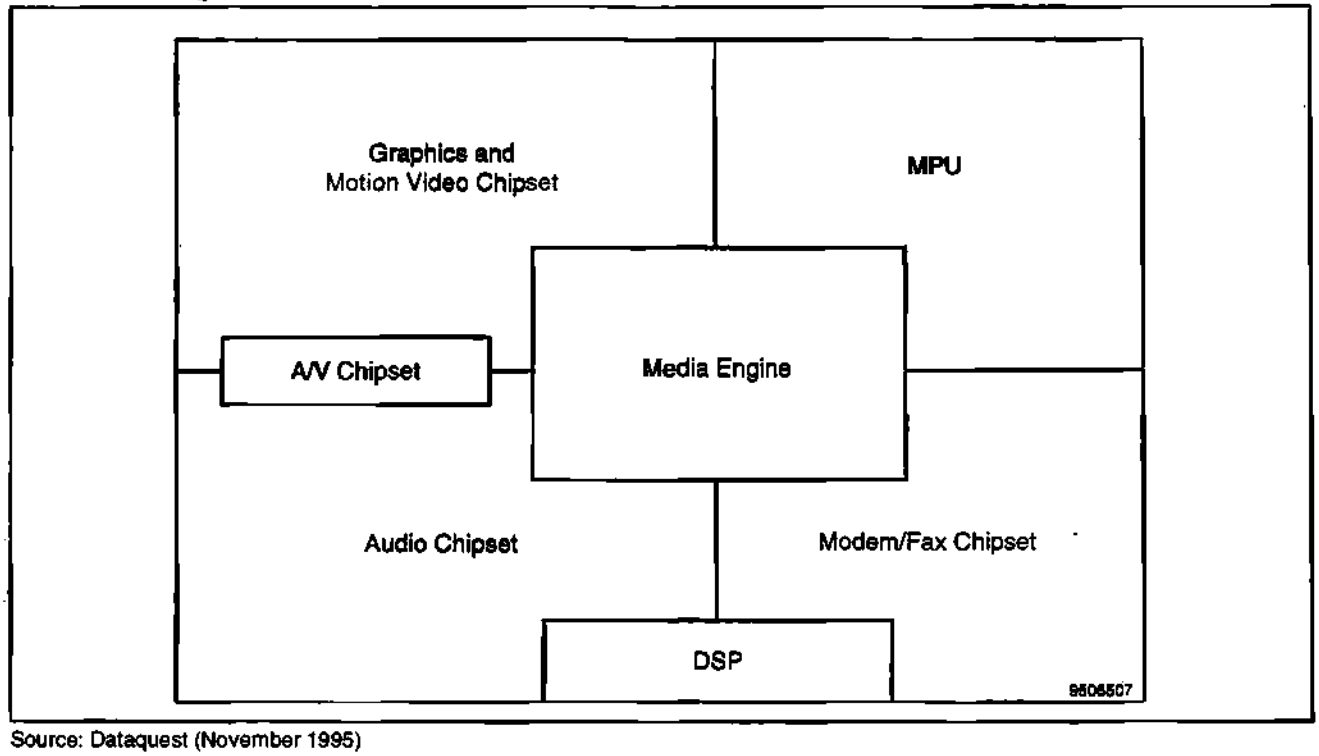

Source: Dataquest (November 1995)

Figure 6
Hardware Implementation Trends

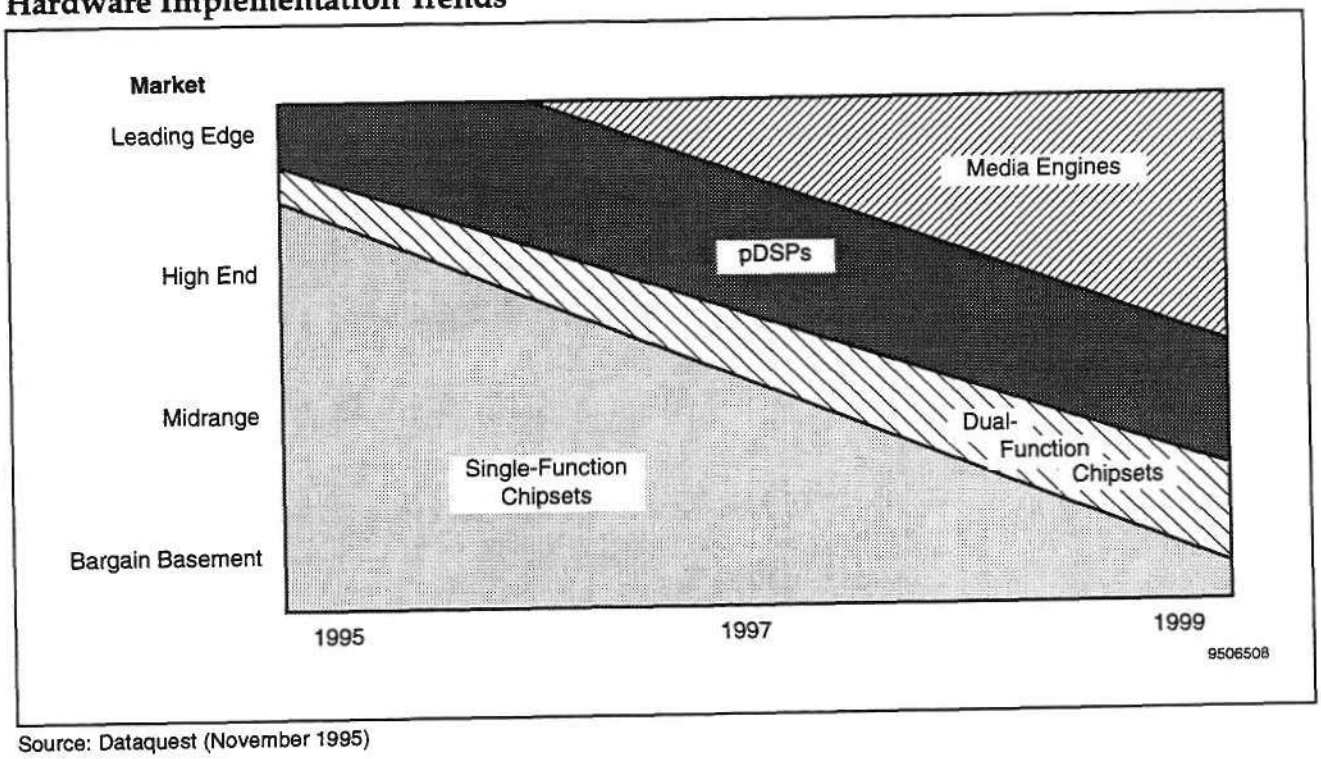

Source: Dataquest (November 1995)

Dedicated chipsets are the least expensive solutions today, so alternative solutions need to be cost-effective on a price-per-function basis to gain significant market share.

Dataquest Perspective

PCs are being used more frequently as communications tools, and this shift is driving the market for both audio and modem chipsets. Semiconductor vendors are adding value to their products by integrating multichip solutions for audio and modem functions onto a single chip as well as pushing performance levels to new heights. Audio and modem functions have already been integrated at the board level by a couple of vendors, and that trend will continue until those functions are integrated into a single dual-function chip. At the same time, those dedicated chipsets will be challenged by alternatives that use generic processing capability and provide functional compatibility through software.

Media engines could take significant market share from dedicated chipsets in 1997 or 1998, but cost of implementation will continue to be an issue and may delay broad-based acceptance. DSPs have a lower price point than media engines and could challenge dedicated chipsets in 1996, but cost remains an issue. If those solutions can handle digital modem functions at full ISDN speeds, they could dominate the top end of the modem market because of their additional functionality. Otherwise, the cost issue will determine how quickly the market shifts from dedicated chipsets to new alternatives.

For More Information...

Geoff Ballew, Industry Analyst ..... (408) 468-8676
Internet address ..... gballew@dataquest.com
Via fax ..... (408) 954-1780

Competitive Dynamics

Semiconductor Directions in PCs

Technology Analysis

Trends for Peripherals Controllers in PCs

Abstract: *Peripherals controllers are the key to moving information into or out of your computer. Standard PC peripherals such as monitors, disk drives, keyboards, mice, serial ports, and parallel ports require dedicated silicon inside the PC as well as in the peripheral itself. This document examines trends in the implementation of peripherals controllers in the host PC for graphics, keyboards, serial ports, and parallel ports. Specific examples of the I/O implementations of 16 recent PC designs are provided from Dataquest's PC Teardown program, which documents the semiconductor content of PCs shipping today.*
By Geoff Ballew

Integration Marches Onward

The trends for I/O controllers in PCs are classic examples of semiconductor integration. Those who haven't looked at a new PC design lately may be surprised at how few ICs are required to create a full-featured design. Fewer ICs are needed to implement all of the standard I/O functions, such as disk and keyboard controllers, as well as serial and parallel ports. Graphics solutions are also more highly integrated as new features are added without increasing chip count. Figure 1 shows a system card from a Packard Bell multimedia PC. This system card is an example of the trends identified in this document.

Graphics Controllers Get New Functions on a Single Chip

Graphics controllers are excellent examples of integration trends. Previously, these controllers required several support chips such as RAMDACs or clock generators for a functional graphics subsystem. That is simply no longer the case. Single-chip solutions are now available from all major vendors complete with the features that today's PC buyers demand. Graphics memory remains off chip, so these single-chip solutions do require memory components for a functional graphics subsystem. Many single-chip

Dataquest

Program: Semiconductor Directions in PCs
Product Code: PSAM-WW-TA-9501
Publication Date: November 13, 1995
Filing: Competitive Dynamics

Figure 1
System Board from a Packard Bell Legend 417CDT

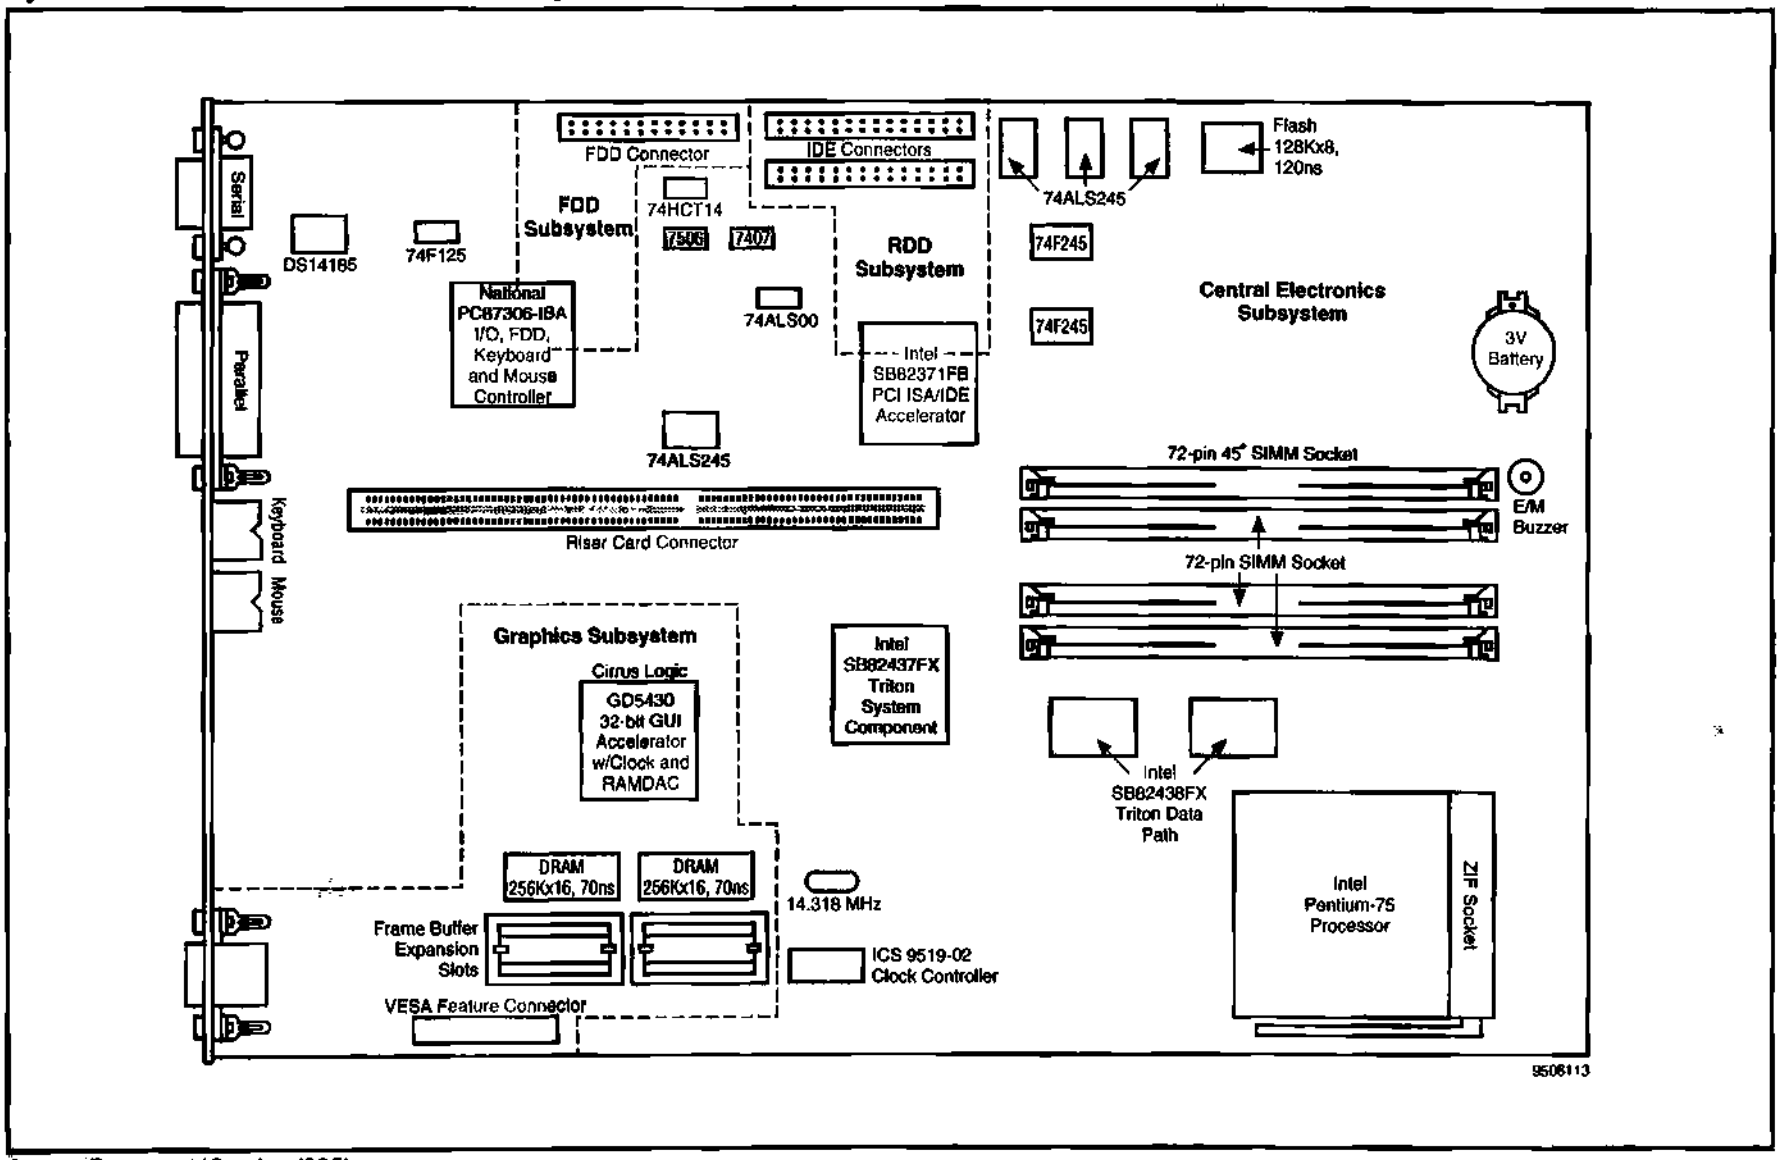

Source: Dataquest (October 1995)

solutions were available last year, but they did not include the advanced features, such as 64-bit processing and motion video acceleration, that are increasingly common today.

A typical graphics subsystem now consists of only three ICs (the graphics controller and two 4Mb DRAMs) with some connectors, a crystal, and some passive components. That low part count is a significant improvement over the previous designs, which required as many as eight or nine ICs. The additional ICs were RAMDACs, clock generators, and standard logic used as line buffers. Figure 2 shows a section of an older graphics subsystem layout compared to a newer design. This figure is not to scale, but the differences in design complexity and required board space are easily seen.

This year is the final chapter for external RAMDACs and clock chips for mainstream graphics applications. Table 1 lists the graphics controllers designed into 16 PCs examined by Dataquest over the past year. This table shows integrated solutions from Cirrus Logic, S3, Chips & Technologies, and Western Digital, as well as one less-integrated solution. Note that only four of the 16 systems have an external RAMDAC and that all four of those

Figure 2
Actual Board Layouts for Two Desktop PC Graphics Subsystems

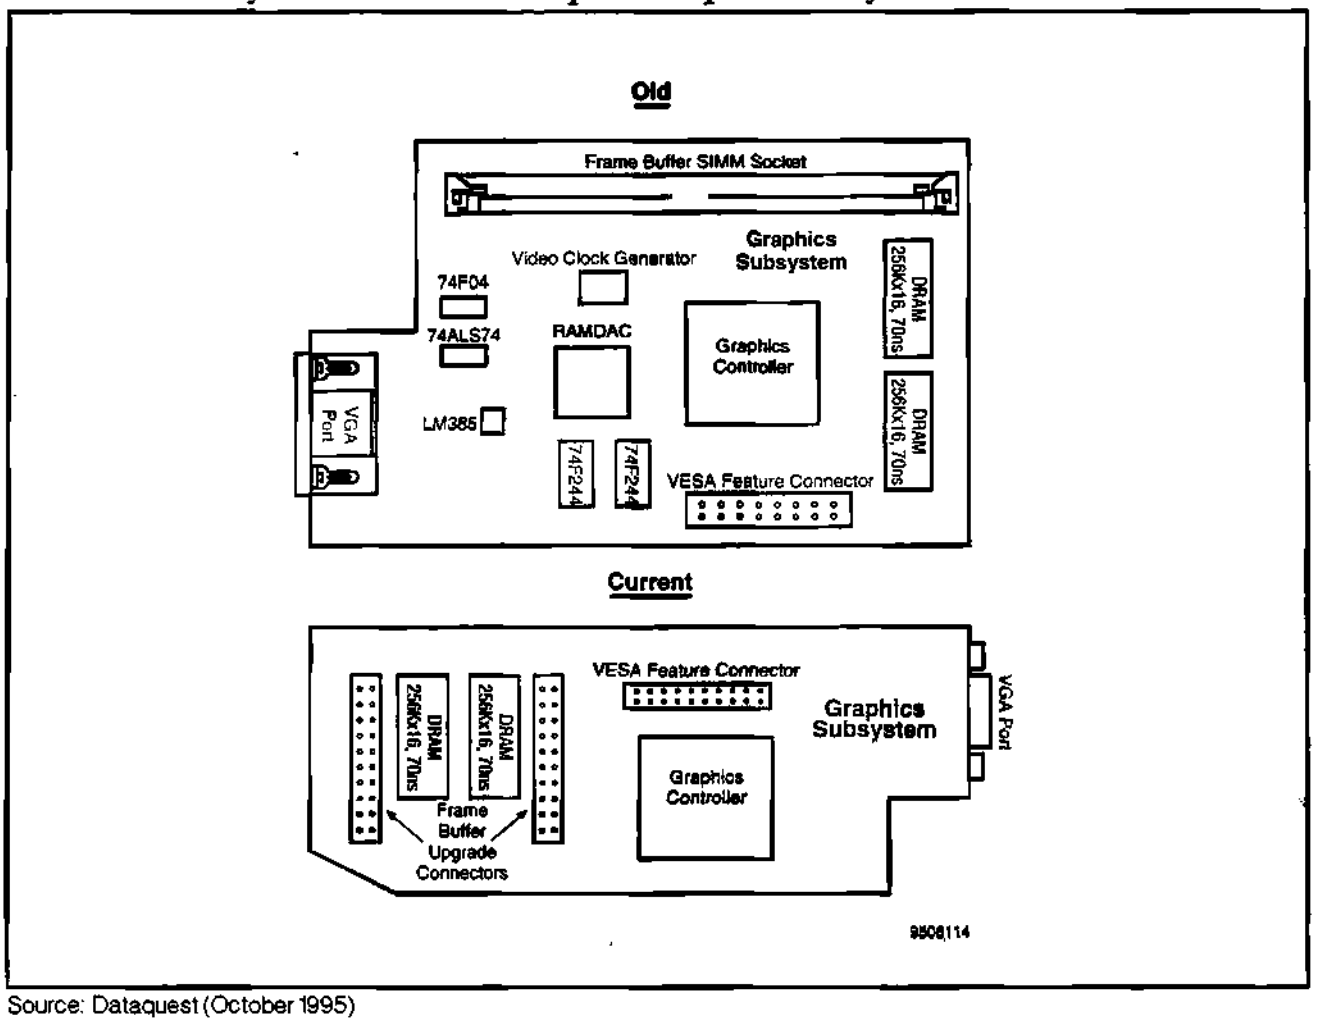

Source: Dataquest (October 1995)

Table 1
Graphics Implementations for 16 PCs

	Supplier	Part Number	Package	Bus Used	Bit Width
Acer Acros 486DX/2-66	Cirrus Logic	GD5428	PQFP-208	VL	32
IBM Aptiva 330	Cirrus Logic	GD5426	PQFP-208	VL	32
AST Bravo LC 4/100t	Cirrus Logic	GD5428	PQFP-208	VL	32
Packard Bell Legend 25CDT	Cirrus Logic	GD5428	PQFP-208	VL	32
HP Vectra XM 4/100i	S3, SGST	86C864, STG1702 RAMDAC	PQFP-208, PLCC-44	PCI	64
Compaq Presario 954 CDS	Cirrus Logic	GD5434	PQFP-208	VL	64
Digital HiNote Ultra CT450	Chips & Technologies	65545	PQFP-208	VL	32
HP Omnibook	Western Digital	90C24	PQFP-208	VL	32
DEC Celebris 590	S3	86C864, 86C716 RAMDAC	PQFP-208, PLCC-68	PCI	64
AST Bravo MS P90	Cirrus Logic	GD5434	PQFP-208	PCI	64
Dell Optiplex XM 590	S3	86C864, 86C716 RAMDAC	PQFP-208, PLCC-68	PCI	64
Packard Bell Legend 417CDT	Cirrus Logic	GD5430	PQFP-208	PCI	32
IBM PC 750	S3	86C864, 86C716 RAMDAC	PQFP-208, PLCC-68	PCI	64
Hewlett-Packard Vectra XE	Cirrus Logic	GD5430	PQFP-208	PCI	32
Micron P120 Millenia	S3	Trio64	PQFP-208	PCI	64
Toshiba Satellite Pro T2150CDT	Chips & Technologies	65545	PQFP-208	VL	32

Source: Dataquest (October 1995)

have the same graphics controller. S3's Vision864 64-bit graphics controller does not have an integrated RAMDAC. This product has been replaced by the Trio64, which does have an integrated RAMDAC and is designed into the Micron PC shown on Table 1.

Delays in bringing the Trio64 to market caused OEMs using the Vision864 to continue using external RAMDACs longer than they had planned. Now that all major graphics providers are shipping their volume products with integrated RAMDACs, external RAMDACs will be used primarily for niche markets willing to pay extra for higher RAMDAC performance. Integrated RAMDACs are implemented in CMOS and do not match the speed or analog accuracy that BiCMOS RAMDACs have, so the trade-off is a question of price versus performance.

The bus interface used for graphics controllers is largely tied to the microprocessor choice: PCI for Pentium systems and VL-bus for 486s. One exception to that general rule is the Vectra XM 4/100i, which uses a PCI bus despite the fact that it has a 486DX/4. To simplify their product lines during the transition from VL-bus to PCI, several graphics vendors designed their parts with support for both bus standards. One example shown in Table 1 is the Cirrus Logic GD5434. In the Compaq Presario 954 CDS this chip was wired to the VL-bus, but in the AST Bravo MS P90 the same part was used on a PCI bus. Only one bus is used at a time, of course, but support for multiple bus standards demonstrates a level of integration that simplifies production and inventory management. It reduces the likelihood that either the vendor or the OEM will be stuck with inventory because it supports the wrong bus.

Bit width is shifting to 64-bit designs. Only two of the seven Pentium systems used 32-bit graphics controllers. The other five all had 64-bit designs. The 486 systems shipped this year generally had 32-bit graphics controllers. Expect shipments of 32-bit graphics controllers to decline as the 486 weathers another year or two of low-end PC designs. The volume shipments of graphics controllers moving forward will be increasingly weighted toward the 64-bit chips.

A few companies are working to integrate the graphics memory onto the graphics controller chip. This is a difficult task because of the differences in the manufacturing processes used for memory as compared to logic products. It is unlikely that anyone will produce a price-competitive product that integrates logic and memory for production within the next 12 months. However, Dataquest believes that demand for an integrated product would be high from OEMs striving to deliver greater functionality at traditional price points.

The bottom line for trends in graphics subsystems is simply fewer ICs. Current product offerings from all of the major graphics vendors provide the functions that PC buyers demand integrated into a single controller chip. As buyers demand more functions, expect those functions also to be integrated onto the single controller chip. Graphics memory will continue to require additional chip count for the short term and probably even through the next two or three years.

Super I/O Chips Get New Features

Super I/O controllers are omnipresent in PC designs today. These chips provide an array of features, including, at a minimum, a floppy disk drive controller, a basic serial port, and a basic parallel port. Many Super I/O chips also have an IDE interface, support for an infrared port, and enhancements to the basic serial and parallel ports. All of these features are frequently packed into a single 100-pin PQFP package. Super I/O controllers were originally introduced in 5V versions but are now available in 3.3V designs. The 3.3V Super I/O chips bring OEMs one step closer to PC designs that operate entirely at 3.3V for system board logic.

The latest enhancement to Super I/O chips is the integration of additional features. These new chips are called Ultra I/O chips by Standard Microsystems (SMC) and include an 8042-compatible keyboard controller, ROM BIOS, and even a real-time clock. National Semiconductor has also updated its Super I/O product line to include keyboard controllers and real-time clocks. Ultra I/O chips from SMC and the updated Super I/O chips from National Semiconductor are typically packaged in 160-pin PQFPs.

Table 2 shows the chips used for general I/O functions in 16 different PCs. One system, the Packard Bell Legend 417CDT, stands out as a trendsetter. Figure 1 showed the board layout for this PC, including the PC87306 Super I/O chip from National Semiconductor. The PC87306 is one of the new breed of Super I/O chips with an integrated keyboard controller and real-time clock. Note that 13 of these PCs feature standard Super I/O controllers with a separate keyboard/mouse controller. These 13 systems represent the typical PC design today. The remaining two systems use custom solutions as well as custom core logic. The Super I/O and Ultra I/O product markets are dominated by National Semiconductor and SMC. Their representation here (13 of 16 systems) is indicative of their presence in the market, but this list is biased toward North American designs. United Microelectronics Corporation (UMC) scored the design win for the Aptiva 330 as well as for other designs not listed here.

Single-Chip RS232 Transceivers Replace Three Other Chips

New RS-232 transceiver ICs provide both the line driver and line receiver functions, which previously required three separate pieces of standard logic. Table 3 shows the number of standard logic parts used for the serial ports in 16 PCs. Typical designs today use two line driver ICs and three line receiver ICs for a dual serial port configuration or three pieces of standard logic (one driver and two receivers) for a single serial port. Note that three designs use only one piece of standard logic. All three of these PCs (two notebooks and one desktop) have a single serial port and use a transceiver chip to replace both the line driver and line receiver ICs: IC count just dropped from three chips to one for a single serial port. A dual port design would use two ICs instead of the five shown on most of the designs listed in Table 3.

Table 2
Super I/O and Keyboard Controller Implementations for 16 PCs

	Super I/O Part Number	Supplier	Keyboard Controller Part Number	Supplier
Acer Acros 486DX/2-66	Custom	ALI	Custom	ALI
IBM Aptiva 330	82C863	UMC	8042	Intel
AST Bravo LC 4/100t	37C651	SMC	82C114	VLSI
Packard Bell Legend 25CDT	37C665	SMC	82C42	Intel
HP Vectra XM 4/100i	37C665	SMC	Custom	NEC
Compaq Presario 954 CDS	87332VLJ	National	82C114	VLSI
Digital HiNote Ultra CT450	87334VLJ	National	Custom	Hitachi
HP Omnibook	37C665	SMC	Custom	Mitsubishi
DEC Celebris 590	37C665	SMC	82C42	Intel
AST Bravo MS P90	37C665	AST	Custom	Intel
Dell Optiplex XM 590	87332VLJ	National	PC9011	National
Packard Bell Legend 417CDT	87306	National	Integrated into Super I/O	-
IBM PC 750	87332VLJ	National	8042	Intel
Hewlett-Packard Vectra XE	87332VLJ	National	Custom	NEC
Micron P120 Millenia	37C665	SMC	8242	Intel
Toshiba Satellite Pro T2150CDT	Custom	Toshiba	Custom	Toshiba

Source: Dataquest (October 1995)

Table 3
Standard Logic Used for Serial Ports

	Number of Parts	Number of Serial Ports	Type of Parts
Acer Acros 486DX/2-66	5	2	2 drivers, 3 receivers
IBM Aptiva 330	5	2	2 drivers, 3 receivers
AST Bravo LC 4/100t	5	2	2 drivers, 3 receivers
Packard Bell Legend 25CDT	3	1	1 driver, 2 receivers
HP Vectra XM 4/100i	5	2	2 drivers, 3 receivers
Compaq Presario 954 CDS	3	1	1 driver, 2 receivers
Digital HiNote Ultra CT450	1	1	1 transceiver
HP Omnibook	1	1	1 transceiver
DEC Celebris 590	5	2	2 drivers, 3 receivers
AST Bravo MS P90	5	2	2 drivers, 3 receivers
Dell Optiplex XM 590	5	2	2 drivers, 3 receivers
Packard Bell Legend 417CDT	1	1	1 transceiver
IBM PC 750	2	1	2 drivers, 3 receivers
Hewlett-Packard Vectra XE	5	2	2 drivers, 3 receivers
Micron P120 Millenia	5	2	2 drivers, 3 receivers
Toshiba Satellite Pro T2150CDT	None	None	NA

NA = Not applicable
Source: Dataquest (October 1995)

Dataquest Perspective

The result of these trends is simply that it takes fewer ICs to provide the basic functions in a PC, even as those "basic" functions grow in capability. These integration trends impact virtually all PC designs because the functions discussed here are universal. Every standalone PC needs a graphics controller, a floppy disk drive controller, a keyboard controller, serial and parallel ports, and a real-time clock. These chips will face new competition from competing designs as other standards, such as Universal Serial Bus, vie to replace multiple ports (serial, parallel, keyboard, and mouse) with a single multifunction port. Until then, these new chips effectively reduce board space and chip count without compromising on features. Semiconductor vendors supplying real-time clocks, keyboard controllers, and standard logic beware—your piece of the PC just got smaller.

For More Information...
Geoff Ballew, Industry Analyst ..... (408) 468-8676
Internet address ..... gballew@dataquest.com
Via fax ..... (408) 954-1780

Dataquest
a company of The Dun & Bradstreet Corporation

Competitive Dynamics

Semiconductor Directions in PCs Worldwide

Product Analysis

PC Memory Trends Are Both DIMM-Witted and Savvy

Abstract: *The personal computer market is plagued by high competition on one side and increasing memory component costs on the other side. The worldwide DRAM shortage and strong yen have resulted in stable or increasing DRAM prices. What is a PC OEM to do in this market? The answer: Tailor memory configurations in PCs to specific target markets to keep costs low. This article highlights trends in main memory configuration for office desktop, consumer multimedia, and notebook PCs.*
By Geoff Ballew

Introduction

Gross margins for PC OEMs are squeezed between severe price competition for personal computers and increasing prices for main memory devices. Memory is the largest single expense in a typical PC. DRAM represents about one-third of the manufacturer's direct costs, so memory prices and availability have a profound impact on system prices. PC OEMs are striving to minimize this impact by tailoring their computer models specifically to the computing needs and budget of each target market. This article highlights trends in this area. All of the data for this analysis is taken from Dataquest's PC Teardown Reports service and is provided in Table 1.

4MB? I Need a Minimum of 8MB ...

The most significant change to factory-installed memory configurations over the past year is the growth of the minimum amount of memory for mainstream configurations. Most PCs today are shipping with 8MB instead of 4MB factory-installed memory, probably because of a combination of software needs and the 64-bit memory bus of Intel's Pentium microprocessor. Because the most common organization for DRAM is 1Mbx4, a typical configuration for a Pentium system is 16 of these devices. Each device is

Dataquest

Program: Semiconductor Directions in PCs Worldwide
Product Code: PSAM-WW-PA-9502
Publication Date: July 17, 1995
Filing: Competitive Dynamics

Table 1
Standard and Expanded Memory Configurations

Type	System	MPU	MPU Type	Total MB	Max. MB	Soldered	Modules	Organization	Voltage	Speed (ns)
Office Desktop	Acer Altos 900	486SX-33	486	4	36	4	SIMM	1Mx4	5	70
Office Desktop	Vobis Highscreen Colani	486DX2-66	486	4	32	0	4MB SIMM	1Mx4	5	70
Office Desktop	AST Bravo LC 4/100t	486DX4-100	486	8	64	0	4MB SIMM	1Mx4	5	60
Office Desktop	AST Bravo MS P90	Pentium-90	Pentium	8	128	0	4MB SIMM	1Mx4	5	70
Office Desktop	Dell Optiplex 466/Le	486DX2-66	486	8	64	0	8MB SIMM	1Mx4	5	70
Office Desktop	Digital Celebris 590	Pentium-90	Pentium	8	128	0	4MB SIMM	1Mx4	5	70
Office Desktop	HP Vectra M2 433/s	486SX-33	486	8	96	0	8MB SIMM	1Mx4	5	70
Office Desktop	HP Vectra XM2 4/100i	486DX4-100	486	8	96	0	8MB SIMM	1Mx4	5	70
Office Desktop	Dell Optiplex XM 590	Pentium-90	Pentium	16	128	0	8MB SIMM	1Mx4	5	70
Consumer Multimedia	AcerAcros 486DX2/66	486DX2-66	486	4	36	4	SIMM	1Mx4	5	70
Consumer Multimedia	IBM Aptiva 330	486SX2-50	486	4	64	0	4MB SIMM	1Mx4	5	70
Consumer Multimedia	Apple Performa 638	68LC040-66	68040	8	36	4	4MB SIMM	1Mx4	5	70
Consumer Multimedia	Compaq Presario CDS 860	486SX2-66	486	8	64	4	4MB SIMM	1Mx4	5	70
Consumer Multimedia	Compaq Presario CDS 954	486DX4-100	486	8	104	8	SIMM	1Mx4	5	70
Consumer Multimedia	Packard Bell Legend 25CDT	486DX2-66	486	8	64	4	4MB SIMM	1Mx4	5	70
Consumer Multimedia	Packard Bell Legend 417CDT	Pentium-75	Pentium	8	72	0	4MB SIMM	1Mx4	5	70
Notebook	Apple PowerBook 540C	68LC040-50	68040	4	36	4	Custom	1Mx4	5	80
Notebook	Compaq LTE Elite 4/40 CX	486DX2-40	486	4	24	4	Custom	1Mx4	3	70
Notebook	DEC HiNote Ultra CT450	486DX2-50	486	8	24	8	DIMM	1Mx16	3.3	70
Notebook	Dell Latitude XP 450	486DX2-50	486	8	36	4	4MB DIMM	1Mx4	5	80
Notebook	HP Omnibook 4000C	486DX4-100	486	8	32		Custom	2Mx8	5	70
Notebook	Toshiba Satellite Pro T2150CDT	486DX4-100	486	8	32	8	DIMM	1Mx16	3.3	70

Source: Dataquest (July 1995)

4 bits wide, so 16 devices create a memory bank that is 1Mb deep and 64 bits wide, which is the same width as the Pentium bus and is 8MB total. The bar was raised for 486 systems when Pentium systems started shipping with 8MB of memory as a minimum, and most manufacturers have upgraded their standard models to a minimum of 8MB of memory. This trend is observable across all categories of mainstream PCs: office desktops, consumer multimedia PCs, and notebooks.

SIMM, DIMM or None of the Above

The packaging for main memory devices demonstrates how PC OEMs are tailoring their products to specific markets. For desktop or deskside PCs, single in-line memory modules (SIMMs) are the most common packing for main memory. SIMMs are popular because they mount into standard sockets without soldering. Because they are standard modules, they can be purchased from multiple vendors and with different amounts of memory installed on each one. As shown in Table 1, all of the office desktop and multimedia PCs use standard SIMMs for memory expansion.

Office PCs and consumer PCs are different with regard to standard memory. In this case, standard memory is the memory installed on the system by default as a standard feature of that particular model. Almost all of the office desktop PCs examined here have the standard memory installed on SIMMs. This relationship is turned upside down for consumer multimedia PCs, however. All but two machines have some or all of their base memory soldered down on the system card rather than mounted on SIMMs. This is a trade-off between two different sets of costs.

Implementing memory on SIMMs requires additional parts such as SIMM sockets and bare SIMM printed circuit boards (PCBs), but does use less space on the system card. One major benefit is increased inventory flexibility because SIMMs can be installed on the system card during final assembly for the whole PC. If PC manufacturers manage their inventory on a just-in-time (JIT) model, then they can reduce their inventory holding costs for DRAM by reducing the length of time the DRAM is held in inventory before final assembly of the PC. This has the added advantage of decoupling the system card assembly from the DRAM availability. If no DRAM is soldered to the system card, then assembly can continue even if DRAM shipments are delayed. Of course, DRAM availability still affects a manufacturer's schedules for final assembly of the PC itself. This increased flexibility does come at a cost, however. A 4MB SIMM is more expensive than the cost of the memory chips alone. One rule of thumb for DRAM SIMM prices is 110 percent of the cost of the memory devices themselves. For example, a SIMM that sells for $110 will have about $100 worth of memory devices installed on it.

In spite of the logistics advantages enabled by using SIMMs, several OEMs for consumer multimedia PCs have chosen to solder some or all of the factory-installed memory to the system card. This policy requires the memory to be purchased and in inventory before the system card assembly can be done. It also reduces the flexibility for changing the memory configuration. To change the memory configuration, the PCB must be redesigned or one of the memory expansion sockets must be used for factory-installed memory, which is the case for three of the multimedia systems examined here.

Using a combination of soldered-down and SIMM-based memory appears to be a fence-straddling policy where the manufacturer would like to use only soldered-down memory but is compelled to hedge against DRAM shortages by reducing the per-system amount of soldered-down memory. The cost of this trade-off is the continued (but reduced) probability that production is interrupted by DRAM part shortages, in addition to the higher cost per megabyte of the SIMM-based memory. SIMM sockets are necessary for memory expansion regardless of whether the factory-installed memory is SIMM-based, but manufacturers that avoid soldered-down memory entirely reap benefits in terms of simplified board design (both space and complexity) in addition to the freedom from the logistics hassles associated with the short supply of DRAM.

The latest trend in main memory packaging for notebook computers is the increasing popularity of standard modules. Until recently, most notebook and ultraportable (subnotebook) designs used custom memory modules for memory expansion. Factory-installed memory was either soldered-down or provided on the same types of custom modules used for memory expansion. Packaging methods are changing in favor of more standard memory modules. These new modules are dual inline memory modules (DIMMs) and are very similar to SIMMs with two exceptions: DIMMs are half the size and require a double-sided socket.

DIMM sockets have two rows of 36 pins each; SIMM sockets have a single row of 72 pins. Some SIMMs do have memory mounted on both sides of the card, but DIMMs are truly two-sided in terms of electrical contact with the socket. SIMM sockets make physical contact with both sides of the SIMM, but contacts on the top and bottom are shorted together (connected electrically) instead of carrying unique signals. In short, a DIMM is a SIMM that has been folded in half and bent over onto itself.

Table 1 shows that half of the notebooks listed use DIMMs for memory expansion. What makes this statistic compelling is the fact that all three of those systems are recent designs. The two oldest notebook designs in Table 1 are the Apple PowerBook and the Compaq LTE Elite, both of which use custom modules for memory expansion. DIMMs are rapidly gaining acceptance as the preferred memory module because they offer to notebook manufacturers the same advantages that SIMMs offer to desktop PC manufacturers: a standard memory interface and configuration flexibility. End users benefit also because they are no longer required to purchase additional memory through a single source, as is the case with custom memory modules.

For More Information...

Geoff Ballew, Industry Analyst ..... (408) 437-8676
Internet address ..... gballew@dataquest.com
Via fax ..... (408) 437-0292

Competitive Dynamics

Semiconductor Directions in PCs Worldwide

Product Analysis

MPU Complex in 1994: 18 PCs Examined in Depth

Abstract: *The personal computer market is both competitive and price-sensitive. To stay in business, PC manufacturers must balance feature set versus production cost and tailor their product lines to specific markets. Design trends in the microprocessor complex show how designers are trimming costs on systems targeted at the home market and the office market without sacrificing the features these buyers demand. Two design trends discussed here involve secondary cache implementation and microprocessor upgradability.*
By Geoff Ballew

Overview

Although 1994 was a big year for Pentium, 486 machines were still the most popular in terms of total sales. The year 1995 will see the Pentium overtake the 486, so it is important to examine the changes in the design of the microprocessor complex as this shift continues. This article examines the secondary cache implementation and the microprocessor upgradability of 18 PCs. These PCs are representative of the typical PCs sold in 1994 and were examined as part of the PC Teardown Project at Dataquest. A complete list of these systems is in Table 1.

Cold, Hard Cache

Secondary cache is a requirement in today's desktop PCs, although experts disagree how much cache memory PCs should have. However, once a PC manufacturer has decided to include secondary cache, there are a number of implementation options. This section highlights trends in the amount of cache memory installed in PCs and the shift in both packaging methods and operating voltage as the mainstream PC grows from a 486 to a Pentium.

Dataquest

Program: Semiconductor Directions in PCs Worldwide
Product Code: PSAM-WW-PA-9501
Publication Date: April 10, 1995
Filing: Competitive Dynamics

Table 1
Cache Memory by Target Market

Target Market	PC	MPU	Factory-Installed Cache (KB)	Maximum Cache (KB)	Cache Voltage	MPU Voltage	MPU Upgrade Path
Home	IBM Aptiva 330	486SX2-50	0	0	NA	5V	DX4 OverDrive
Home	Compaq Presario CDS 860	486SX2-66	0	256	5V	5V	P24T OverDrive
Home	Acer Altos 900	486SX-33	0	256	5V	5V	DX4 OverDrive
Home	AcerAcros 486DX2-66	486DX2-66	0	256	5V	5V	P24T OverDrive
Home	Packard Bell Legend 25 CDT	486DX2-66	0	512	5V	5V	P24T OverDrive
Office	NCR 3227	486DX2-66	0	0	NA	5V	P24T OverDrive
Office	Hewlett-Packard Vectra M2 4/33s	486DX-33	0	256	5V	5V	P24T OverDrive
Office	Dell Optiplex 466/Le	486DX2-66	128	128	5V	5V	DX4 OverDrive
Office	Dell Dimension XPS 466V	486DX2-66	128	256	5V	5V	P24T OverDrive
Office	Gateway 2000 4DX2-66	486DX2-66	128	256	5V	5V	P24T OverDrive
Office	Vobis Highscreen Colani	486DX2-66	128	256	5V	5V	DX4 OverDrive
Office	Dell Dimension XPS P60	Pentium-60	256	256	5V	5V	P54C OverDrive
Office	Compaq Deskpro XL 450	486DX2-50	256	256	5V	5V	P24T OverDrive
Office	Dell Dimension XPS 4100V	486DX4-100	256	256	5V	3.3V	No upgrade
Office	AST Bravo LC 4/100t	486DX4-100	256	256	5V	3.3V	P24T OverDrive
Office	AST Bravo MS P90	Pentium-90	256	256	3.3V	3.3V	Future MPU
Office	Dell Optiplex XM 590	Pentium-90	256	512	5V	3.3V	Future MPU
Office	Digital Celebris 590	Pentium-90	256	512	3.3V	3.3V	Future MPU

NA = Not applicable
Source: Dataquest (April 1995)

Cache Memory Is Strongly Tied to Target Market

Cache memory in PCs follows several distinct trends, even though they may appear random on the surface. A few systems examined have no provision for cache memory but these are rare because most buyers demand installed cache memory or at least the opportunity to add some later. The upgradability of machines is important because cache represents a cost with less of an understood need than some other components. Home consumers demand varying amounts of cache memory, but business buyers are more predictable in their demands. Consequently, machines targeted to the extremely cost-sensitive home market leave the cache memory as an option for the retailer or the consumer, while machines targeted to businesses have cache memory when they leave the PC manufacturer. As the mainstream PC changes from a 486 to a Pentium the demand for cache memory will increase. The faster designs and higher clock speeds require larger secondary caches to avoid waiting on main memory for either instructions or data. The data in Table 1 highlights the amounts of cache memory and expandability for two PC categories: home PCs and office PCs.

Table 1 shows that PCs targeted to the home have no cache memory installed. This highlights the need for PCs targeted to home consumers to be price-competitive based on newspaper ads and at the same time feature the "checklist" items that consumers may use to distinguish among systems. The importance of secondary cache is less widely understood than other features such as microprocessor, hard disk capacity, or multimedia hardware. A consumer that does not understand the need for cache memory may overlook a system because of a small price difference when compared to an identical, but cacheless, system.

PCs targeted to the business or general market must also be price-competitive, but the requirements are different. Office systems are configured for business productivity rather than home productivity and entertainment. The price/performance trade-off is not complicated for most businesses by issues of bundled software or multimedia features, and this is reflected in the levels of installed cache. Business buyers therefore demand installed cache memory because it boosts performance.

SIMM, Soldered-Down, or None of the Above

Another important characteristic of secondary cache is the method of packaging. System vendors are constantly searching for ways to minimize cost without sacrificing other features. In terms of secondary cache packaging, more designs are breaking the mold of "DIP-socketed, 5V SRAM," and this is reflected by the Pentium systems and the more aggressive 486 systems on our list. Figure 1 shows the relative frequency of new packaging methods.

Until recently, cache memory was almost always mounted in DIP sockets. Nine of the 12 486s with cache memory had DIP sockets for the secondary cache. However, DIP sockets were not used on any of the Pentium systems. Those systems were split evenly between soldered-down and SIMM designs. By using a SIMM for cache memory, the PC can be reconfigured quickly at any time before or after the sale and users get exactly the amount of cache they want. SIMM-based cache will become mainstream for the same reasons that SIMM-based DRAM is popular: low cost, efficient use of board space, and upgradability.

Figure 1
Cache Packaging by MPU Type

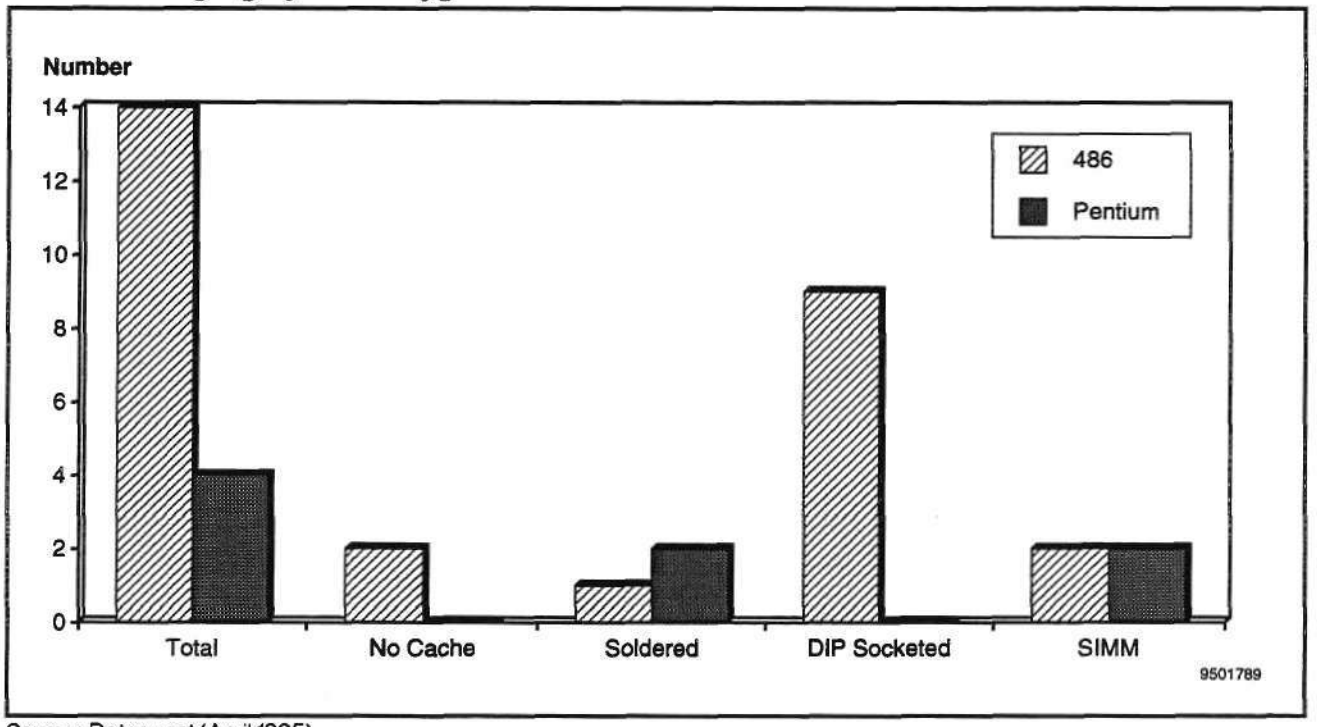

Source: Dataquest (April 1995)

3.3V Cache Appears on 3.3V Pentium PCs

Low-voltage cache memories are breaking ground in 3.3V Pentium PCs. This is a sharp contrast to the 486 PCs. Figure 2 shows that 3.3V cache is being designed into Pentium PCs but not 486 PCs. The reasons for this are simple: design complexity and cost.

Any voltage difference between the microprocessor and the cache memory increases the design complexity because these two components are so closely related. It is more important for the microprocessor and cache memory to operate at the same voltage than any other components in a PC. This became an issue when Intel began making 3.3V microprocessors for PCs in addition to the standard 5V microprocessors. Because the fastest processors were 3.3V instead of 5V, PC designers were forced to design a mixed-voltage interface (3.3V and 5.0V) or use 3.3V cache memory. Most designers used the more complex mixed-voltage interface because of the cost premium for 3.3V cache memory. SRAM is the type of memory used for cache on PCs, and the cost premium for 3.3V SRAM compared to 5.0V SRAM was prohibitive. However, prices on 3.3V SRAM are getting closer to those for 5.0V SRAM, and this cost premium is no longer prohibitive. Many PC designers are now using 3.3V cache memory when a 3.3V microprocessor is used. As Figure 2 shows, two of the three 3.3V Pentium PCs used 3.3V cache memory. New designs increasingly will match the cache voltage to the MPU voltage. This trend will continue until all 3.3V microprocessors are paired with 3.3V cache memory because it reduces design complexity without significant costs.

Figure 2
Cache Voltage by MPU Type for PCs with 3.3V MPUs

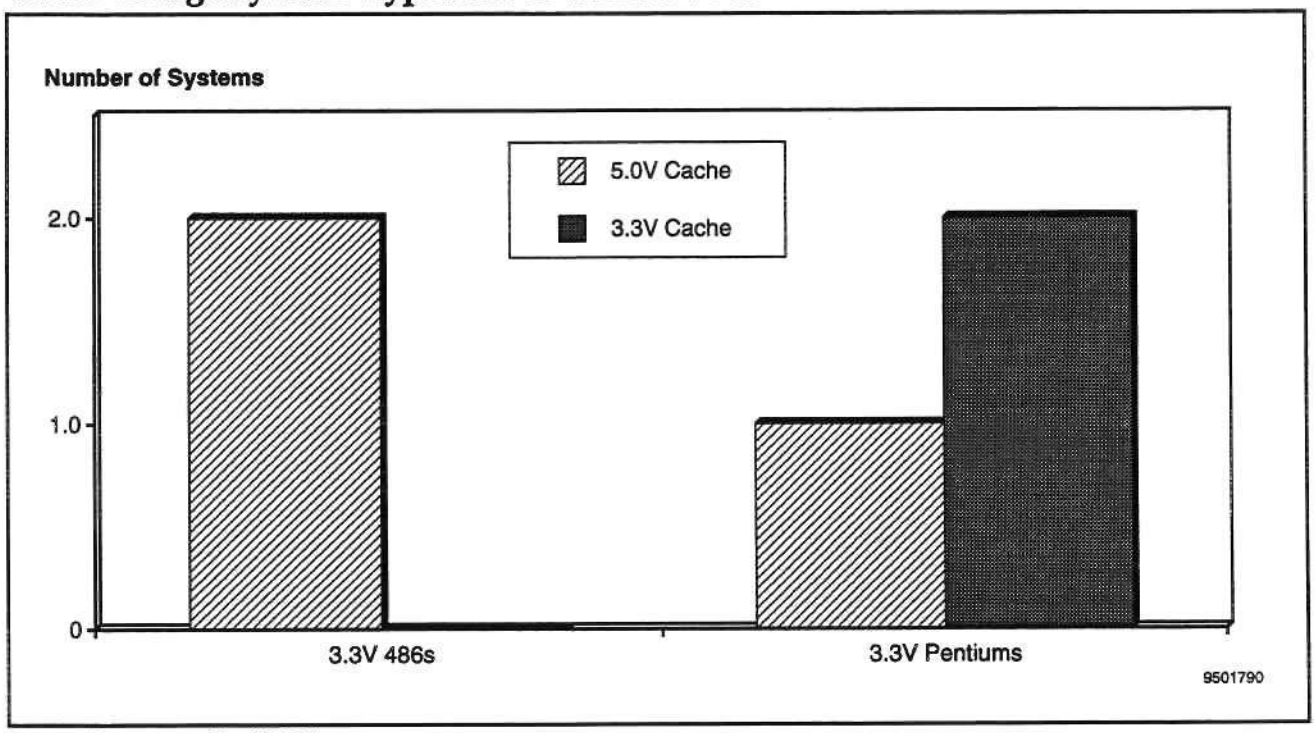

Source: Dataquest (April 1995)

Slipping into OverDrive: Obsolescence or Upgradability?

Performance expectations for PCs are rising as fast as hardware prices are falling. As corporations upgrade to new versions of their software suites and home users get more excited about multimedia features, more people are facing the need to replace or upgrade their PC hardware. Choices include upgrading the microprocessor, upgrading the system card, or replacing the PC entirely. Replace or upgrade, that is the question. And the question may be answered by one small, inexpensive part of a PC: the microprocessor socket. This section highlights trends in upgrade options and explores the trade-offs between upgrading a PC and completely replacing it.

Microprocessor Upgrade Options

The number of microprocessor upgrade options is overwhelming. However, the simplest way to determine a particular system's upgrade path is to count the number of pins in the MPU socket. Table 2 shows, for current microprocessors and sockets, what the microprocessor upgrades could be.

How 18 Systems Stack Up

The most common upgrade path among the 14 486 PCs is upgrading to a P24T OverDrive microprocessor, as shown in Table 3. The P24T OverDrive microprocessor is a Pentium microprocessor with a 32-bit external bus just like a 486 microprocessor, so it is targeted directly at users of high-end 486 PCs. The Pentium PCs were split three to one. The Dell Dimension XPS P60 is a candidate for a P54C OverDrive processor, but the other three

Table 2
Typical MPU Upgrade Path

Number of Pins	Current MPUs	Upgrade MPUs
168, 169	486 SX, DX, SX2, DX2	OverDrive 486 DX4
168, 169	486 DX4	No reasonable upgrades
235	486 DX4	P24T (Pentium with a 32-bit bus)
237, 238	486 SX, DX, SX2, DX2, DX4	P24T (Pentium with a 32-bit bus)
273	Pentium 60 or 66	P54C OverDrive
320	Pentium 90 or 100 (P54C)	Unknown (faster Pentiums?/P6 OverDrive?)

Source: Dataquest (April 1995)

Pentium PCs already have Pentium-90 microprocessors. No upgrade MPUs have been announced for those three systems, but faster Pentiums and next-generation OverDrive microprocessors are likely candidates.

The five 486 PCs that are not upgradable to P24T OverDrive MPUs highlight the delegation of 486 processors to the low-end econo-box end of the PC spectrum. These PCs are all niche-oriented. The Dell Dimension XPS4100V is already a DX4 machine, but does not have the proper socket for a P24T upgrade. It is not upgradable because faster microprocessors will not fit in its socket. The other four 486 PCs—the Acer Altos 900, the Vobis Highscreen Colani, the Dell Optiplex 466/Le, and the IBM Aptiva 330—were designed for extreme low cost and were not meant to serve a broad customer base. The Acer Altos 900 was actually a resurrection of an old board design to serve the low-end home market in Asia, while the IBM Aptiva 330 is a low-end product targeted to the U.S. home market.

Finding the Sweet Spot

The sweet spot of the upgrade market is 486 PCs that are upgradable to P24T OverDrive microprocessors. These PCs are likely to have accelerated, local-bus graphics and IDE controllers in addition to plenty of RAM and rigid disk capacity. In other words, all major parts of the PC are fast, but they could be improved with a faster microprocessor. Outside of this sweet spot are PCs with slow graphics and IDE controllers and PCs that are not upgradable to P24T OverDrive microprocessors.

So who will buy upgrade microprocessors? Anyone who needs a boost in performance and has a machine that is only a little out of date. This rule applies to owners of name-brand PCs, because no-name clone motherboards with MPUs do not cost much more than OverDrive processors. This cost/benefit relationship causes more owners of clone machines to upgrade their system card at a near-commodity price rather than simply upgrade the processor.

However, owners of name-brand PCs face a higher replacement cost for a system card that fits inside their case. This makes upgrading to either an entirely new machine or an OverDrive processor more attractive than just upgrading the system card. To illustrate this point, imagine that a new system costs $2,000, a new system card with a faster MPU costs $1,000 from a name-brand vendor or $600 from a no-name clone vendor, and an OverDrive MPU costs $500. Owners of no-name clones face a choice of

Table 3
Upgrade Options for 18 PCs

Upgrade Path for 486 PCs	Number
To 486DX4 OverDrive	1
To 486DX4 OverDrive	4
To P24T OverDrive	9
Total	14
Upgrade Path for Pentium PCs	
To P54C OverDrive	1
To future OverDrive MPUs	3
Total	4

Source: Dataquest (April 1995)

$500, $600, or $2,000 to upgrade the MPU, system card, or entire system, respectively. They probably will choose either the $600 system card or the $2,000 system because the OverDrive processor does not provide near as much value as the $600 system card. Owners of name-brand PCs face a choice of $500, $1,000 or $2,000. The choice in this case is less obvious because the system card cost has not been pushed down as close to the MPU cost. Indeed, any one the three options may be attractive, given the necessity of an upgrade and the opportunity cost of spending the extra money.

Overall, the biggest detriment to the microprocessor upgrade market is the sheer pace of development in all aspects of PCs. It is hard to justify spending several hundred dollars for incremental gains, when much faster and much bigger (capacitywise) devices are only slightly more expensive. Buying a new PC is attractive because you get a bigger rigid disk drive, faster graphics, quicker CD-ROM drive, and a faster MPU all at once. The upside for both PC manufacturers and semiconductor companies is simple: more people buying new PCs means greater sales.

Dataquest Perspective

PC manufacturers must tailor the features of their products to the target market. Home consumers demand a different feature set than do business buyers. The 18 PCs examined here show clear design differences in the MPU complex of PCs targeted at the home rather than the business community. Business buyers focus on the price/performance ratio while home consumers may trade performance for other features. Trends in both secondary cache implementation and microprocessor upgradability should be taken into account by PC manufacturers.

For More Information...

Geoff Ballew, Industry Analyst ..... (408) 437-8676
Internet address ..... gballew@dataquest.com
Via fax ..... (408) 437-0292

Focus Studies

Dataquest

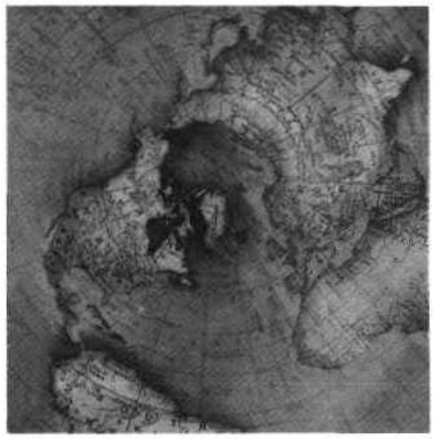

Dataquest

PC Graphics Controllers: A Focused Analysis

Focus Report

Program: Semiconductor Directions in PCs
Product Code: PSAM-WW-FR-9502
Publication Date: August 21, 1995
Filing: Focus Studies

PC Graphics Controllers: A Focused Analysis

Focus Report

Program: Semiconductor Directions in PCs
Product Code: PSAM-WW-FR-9502
Publication Date: August 21, 1995
Filing: Focus Studies

Table of Contents

List of Figures

List of Tables

Chapter 1
Executive Summary

The graphics controller market for PCs grew 48.5 percent from 1993 to 1994 because of strong PC sales and a trend away from captive graphics solutions. Unit shipments will continue to grow year to year, but the growth rate will follow PC unit growth more closely than it did last year. Revenue will also grow but will not keep pace with unit shipments because of price pressure in this competitive market.

Cirrus Logic remained the largest graphics provider by shipping roughly 26 million units representing 44 percent of the total graphics controller market. S3 showed renewed strength in desktop graphics controllers and moved into the No. 3 slot behind Trident Microsystems. Dataquest expects the 1995 rankings to look a lot like the 1994 rankings, with Cirrus Logic on top because of its overwhelming lead and strong product line. The competition for No. 2 promises to be an exciting race between Trident Microsystems and S3—Trident now leads, but S3 has greater momentum.

New features will play a larger role in product differentiation in the next few years as motion video acceleration becomes a standard feature and graphics vendors look for new ways to add value. Integrated RAMDACs and power management are more the rule than the exception in 1995, and advanced features such as 3-D graphics, advanced memory interfaces, and motion video acceleration will become standard over the next few years.

Tremendous opportunity exists for graphics companies as PCs fulfill growing requirements for communication and entertainment as well as personal productivity. Those companies who adapt best to this changing role will occupy the top lines of the market share rankings through the year 2000.

Chapter 2
Graphics Controller Market Share Statistics

Cirrus Logic retained the No. 1 spot for 1994 in the PC graphics controller market with a whopping 44 percent of the market, according to a recent Dataquest survey. No other company came close to matching Cirrus' sales volume for either 1993 or 1994. Much of Cirrus Logic's success can be attributed to its high level of integration. Cirrus was the first company to integrate RAMDACs onto graphics controllers, and that feature gave the company a significant edge against its competitors. Dataquest expects Cirrus Logic to remain No. 1 in unit shipments for 1995 also. Look for S3 to challenge Trident Microsystems for the No. 2 spot for 1995 because of S3's renewed strength with some impressive design-wins, including part of the Compaq Presario line. Figure 2-1 and Table 2-1 show worldwide unit shipments of graphics controllers.

Figure 2-1
Graphics Controller Unit Shipments Worldwide

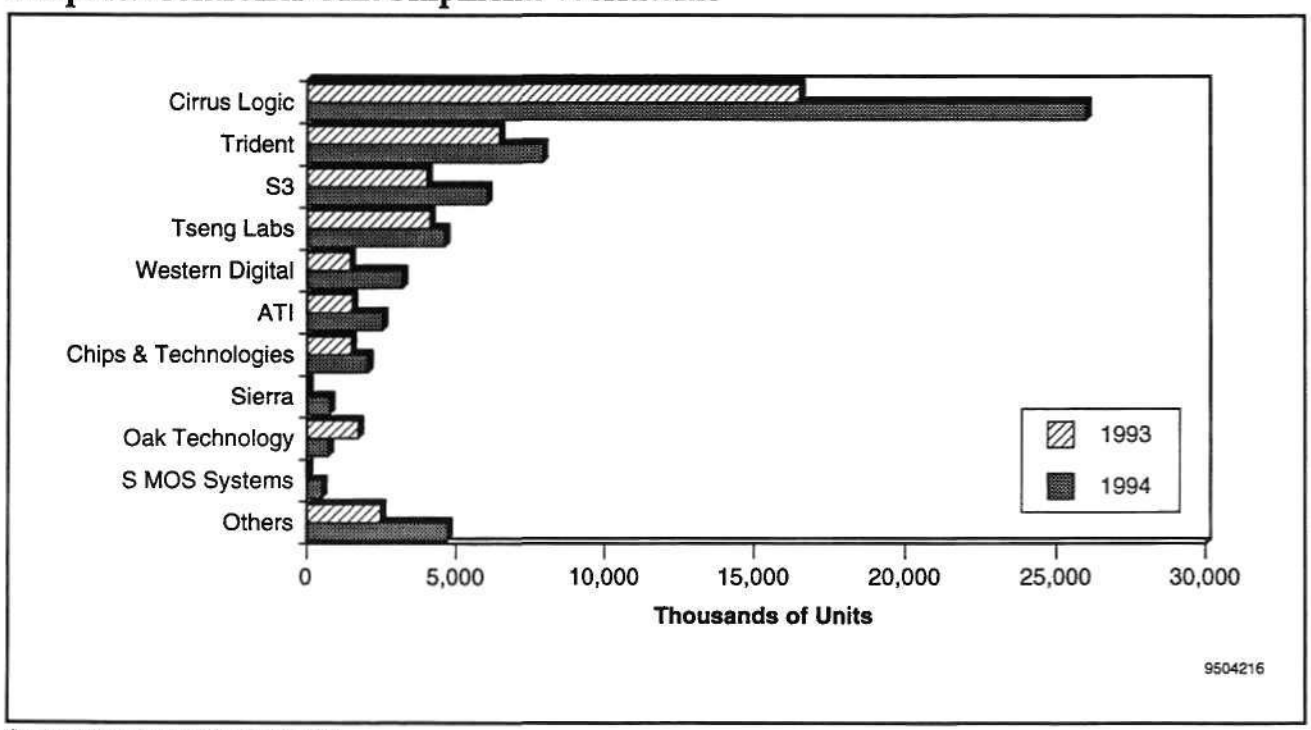

Source: Dataquest (August 1995)

Table 2-1
Graphics Controller Unit Shipments Worldwide (Thousands)

1994 Rank	1993 Rank		1993	1994	Percentage Change	1994 Market Share (%)
1	1	Cirrus Logic	16,500	26,000	57.6	44.2
2	2	Trident Microsystems	6,436	7,903	22.8	13.4
3	4	S3	4,000	6,000	50.0	10.2
4	3	Tseng Labs	4,100	4,600	12.2	7.8
5	8	Western Digital	1,430	3,175	122.0	5.4
6	6	ATI	1,500	2,500	66.7	4.2
7	7	Chips & Technologies	1,461	2,013	37.8	3.4
8	13	Sierra Semiconductor	0	750	NA	1.3
9	5	Oak Technology	1,700	705	-58.5	1.2
10	11	S-MOS Systems	25	500	1,900.0	0.8
		Other	2,465	4,694	90.4	8.0
		Total Market	39,617	58,840	48.5	100.0

NA = Not applicable
Source: Dataquest (August 1995)

Desktop Players Dominate the Market

The desktop market share rankings look a lot like the total market share rankings. No. 1 Cirrus Logic shipped almost as many graphics controllers into the desktop market as all other suppliers combined. Cirrus' strength in the desktop market is unmatched and should remain unchallenged for 1995. S3's momentum could propel it into second place, ahead of Trident Microsystems next year, with strong sales of S3's Vision and Trio graphics controllers. Chips & Technologies showed spectacular unit growth (219 percent) in desktop graphics controllers, but will need triple-digit growth rates for a couple of consecutive years to play in the top five. Figure 2-2 and Table 2-2 show worldwide unit shipments of desktop graphics controllers.

Mobile Graphics Has a New No. 1

Western Digital displaced Cirrus Logic to claim the No. 1 spot in the mobile PC graphics controller market. Its commanding presence (40 percent) in the mobile graphics arena is similar to Cirrus Logic's share of the desktop graphics controller market. Expect another year of good growth from Western Digital because of its design-wins in current notebook PCs, including the HP Omnibooks. Cirrus suffered flat growth (0 percent) from 1993 to 1994 and is challenged by Chips & Technologies for the No. 2 position. Chips & Technologies has some impressive design-wins in its pocket, including the DEC Hinote Ultra and the Toshiba Satellite Pro. The unit shipments for these two companies are too close to differentiate, but Dataquest has ranked Cirrus as No. 2 because of its former No. 1 status. Expect Chips & Technologies' current momentum to carry it into No. 2 next year. S-MOS was the percentage growth leader for 1994 and should gain market share again in 1995.

Figure 2-2
Desktop Graphics Controller Unit Shipments Worldwide

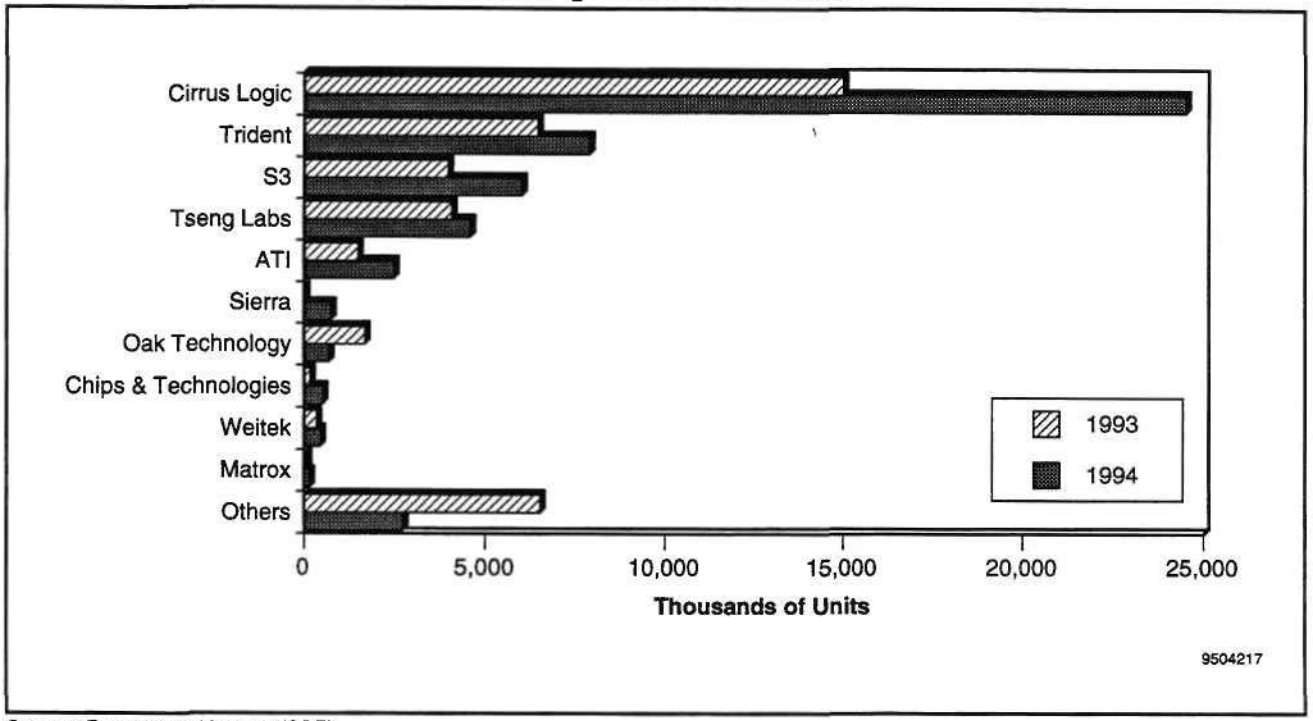

Source: Dataquest (August 1995)

Table 2-2
Desktop Graphics Controller Unit Shipments Worldwide (Thousands)

1994 Rank	1993 Rank		1993	1994	Percentage Change	1994 Market Share (%)
1	1	Cirrus Logic	15,000	24,500	63.3	48.2
2	2	Trident	6,436	7,903	22.8	15.6
3	4	S3	4,000	6,000	50.0	11.8
4	3	Tseng Labs	4,100	4,600	12.2	9.1
5	6	ATI	1,500	2,500	66.7	4.9
6	10	Sierra	0	750	NA	1.5
7	5	Oak Technology	1,695	689	-59.4	1.4
8	8	Chips & Technologies	161	513	218.6	1.0
9	7	Weitek	350	452	29.1	0.9
10	9	Matrox	75	150	100.0	0.3
		Others	6,500	2,754	-57.6	5.4
		Total Market	34,505	50,811	47.3	100.0

NA = Not applicable
Source: Dataquest (August 1995)

The roster of mobile PC graphics companies varies significantly from the desktop graphics companies. The total market is much smaller and requires selling directly to the PC OEMs because the third-party, add-in board market does not exist. Graphics controllers in mobile PCs are soldered to the system cards because of space requirements, and any cards purchased for a docking station should be considered part of the desktop PC graphics controller market.

The 57.1 percent unit growth for mobile PC graphics controllers exceeds the unit growth for mobile PC shipments from 1993 to 1994 because of a shift from captive solutions to merchant graphics solutions. Figure 2-3 and Table 2-3 show worldwide unit shipments of mobile graphics controllers.

Figure 2-3
Mobile Graphics Controller Unit Shipments Worldwide

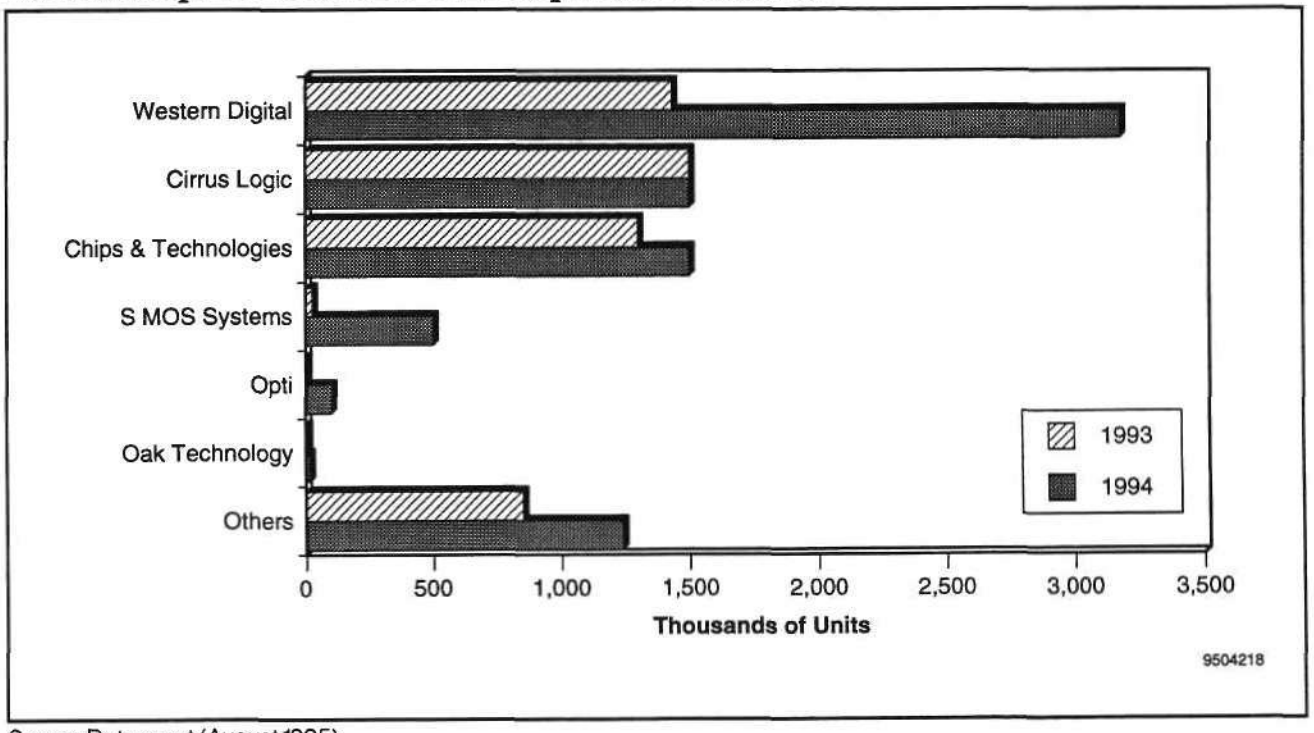

Source: Dataquest (August 1995)

Table 2-3
Mobile Graphics Controller Unit Shipments Worldwide (Thousands)

1994 Rank	1993 Rank		1993	1994	Percentage Change	1994 Market Share (%)
1	2	Western Digital	1,430	3,175	122.0	39.5
2	1	Cirrus Logic	1,500	1,500	0	18.7
3	3	Chips & Technologies	1,300	1,500	15.4	18.7
4	4	S-MOS Systems	25	500	1,900.0	6.2
5	6	Opti	0	100	NA	1.2
6	5	Oak Technology	5	16	220.0	0.2
		Others	852	1,238	45.3	15.4
		Total Market	5,112	8,029	57.1	100.0

NA = Not applicable
Source: Dataquest (August 1995)

Chapter 3
Graphics Controller Market Forecast

The graphics controller market will continue to show strong growth as the boom in PC sales continues. Unit shipments as a percentage (greater the 100 percent) of PC sales will slowly decline as mobile computers command a larger share of PC shipments. Table 3-1 shows the graphics controller unit forecast. Figures 3-1, 3-2, and 3-3 present the data from Table 3-1 graphically. Table 3-2 shows the average selling price (ASP) assumptions. Table 3-3 shows the graphics controller revenue forecast. Figure 3-4 shows graphically the information from Table 3-3.

Assumptions

- Unit growth will continue at a strong pace, tied to the expanding PC market.
- Additional unit demand for desktop graphics controllers is created by the retail add-in card market for upgrading existing systems.
- Graphics controllers for mobile computers are not upgradable.
- Graphic controller shipments precede PC shipments by roughly two months.
- ASPs will remain relatively stable through the forecast period because new features and performance increases will offset downward price pressure.
- 3-D graphics acceleration is the next new feature after motion video acceleration that semiconductor vendors will integrate into graphics controllers to add value.

Table 3-1
PC Graphics Controller Unit Shipment Forecast (Thousands)

	1994	1995	1996	1997	1998	1999	CAGR (%) 1994-1999
Desktop							
Standard VGA	13,649	3,000	1,000	0	0	0	-100.0
2-D Acceleration	35,512	46,324	27,293	6,863	0	0	-100.0
2-D and Motion Video Acceleration	1,650	13,066	40,940	61,770	64,475	51,636	99.1
2-D, 3-D, and Motion Video Acceleration	0	50	2,750	11,750	26,050	50,697	NA
Total Desktop	50,811	62,390	71,983	80,384	90,525	102,333	15.0
Mobile							
Standard VGA	3,349	250	0	0	0	0	-100.0
2-D Acceleration	4,680	10,662	11,587	6,995	1,890	0	-100.0
2-D and Motion Video Acceleration	0	366	3,069	10,492	17,009	18,220	NA
2-D, 3-D, and Motion Video Acceleration	0	0	0	961	5,125	7,361	NA
Total Mobile	8,029	11,278	14,656	18,448	22,134	25,581	26.1
Total Market	58,840	73,668	86,639	98,832	112,659	127,914	16.8

NA = Not applicable
Source: Dataquest (August 1995)

Figure 3-1
Desktop PC Graphics Controller Unit Shipment Forecast

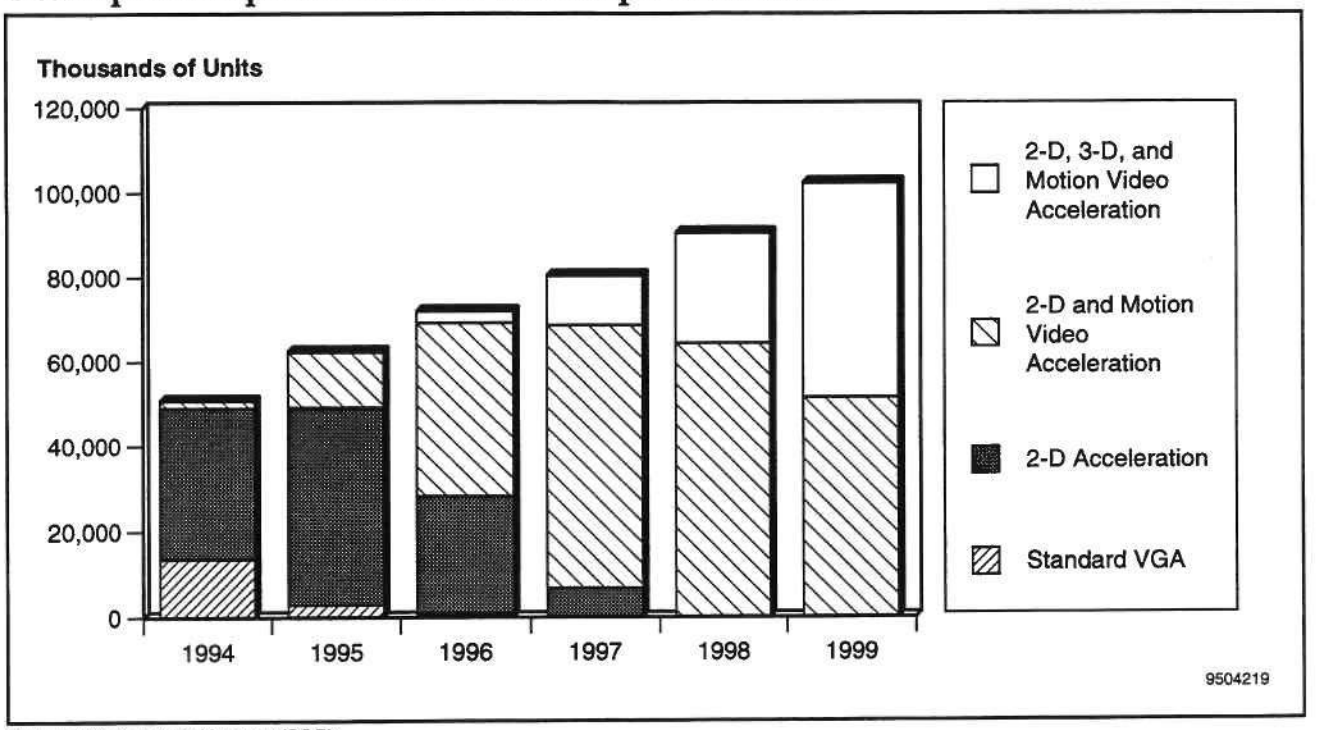

Source: Dataquest (August 1995)

Figure 3-2
Mobile PC Graphics Controller Unit Shipment Forecast

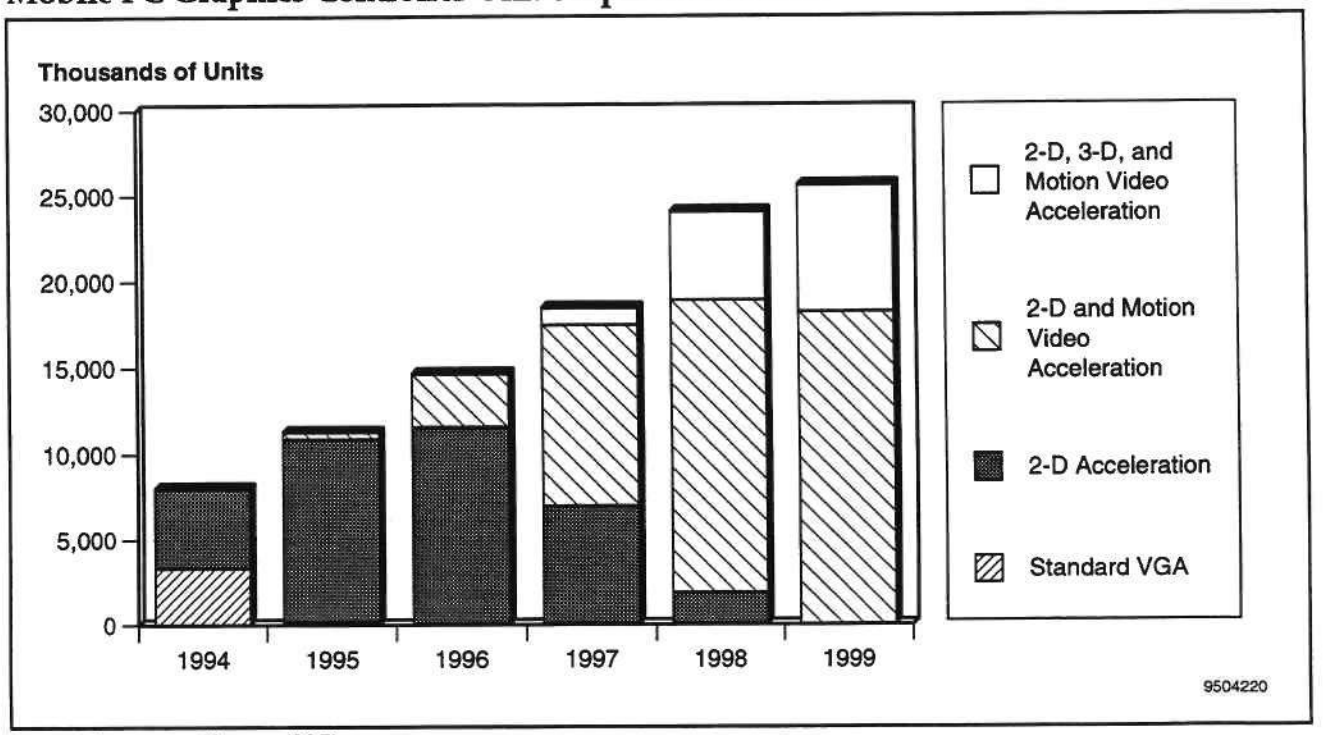

Source: Dataquest (August 1995)

Table 3-2
PC Graphics Controller ASP Forecast (U.S. Dollars)

	1994	1995	1996	1997	1998	1999	CAGR (%) 1994-1999
Desktop							
Standard VGA	10.00	9.00	7.00	7.00	7.00	7.00	-6.9
2-D Acceleration	19.50	17.00	15.00	14.50	14.00	12.00	-9.3
2-D and Motion Video Acceleration	30.00	21.00	17.00	15.00	14.00	12.00	-16.7
2-D, 3-D, and Motion Video Acceleration	45.00	45.00	45.00	31.00	25.50	21.00	-14.1
Weighted ASP	17.29	17.49	17.17	17.30	17.31	16.46	-1.0
Mobile							
Standard VGA	23.00	17.00	16.00	15.00	14.00	14.00	-9.5
2-D Acceleration	26.50	22.00	19.00	18.00	17.00	16.00	-9.6
2-D and Motion Video Acceleration	30.00	27.00	24.00	21.00	18.00	17.00	-10.7
2-D, 3-D, and Motion Video Acceleration	50.00	50.00	50.00	40.00	29.00	25.00	-12.9
Weighted ASP	25.04	22.05	20.05	20.85	22.00	19.30	-5.1

Source: Dataquest (August 1995)

Table 3-3
PC Graphics Controller Total Market Forecast (Millions of U.S. Dollars)

	1994	1995	1996	1997	1998	1999	CAGR (%) 1994-1999
Desktop							
Standard VGA	136.5	27.0	7.0	0	0	0	-100.0
2-D Acceleration	692.5	787.5	409.4	99.5	0	0	-100.0
2-D and Motion Video Acceleration	49.5	274.4	696.0	926.6	902.6	619.6	65.8
2-D, 3-D, and Motion Video Acceleration	0	2.3	123.8	364.3	664.3	1,064.6	NA
Total Desktop	878.5	1,091.1	1,236.1	1,390.3	1,566.9	1,684.3	13.9
Mobile							
Standard VGA	77.0	4.3	0	0	0	0	-100.0
2-D Acceleration	124.0	234.6	220.1	125.9	32.1	0	-100.0
2-D and Motion Video Acceleration	0	9.9	73.7	220.3	306.2	309.7	NA
2-D, 3-D, and Motion Video Acceleration	0	0	0	38.4	148.6	184.0	NA
Total Mobile	201.1	248.7	293.8	384.7	486.9	493.8	19.7
Total Market	1,079.5	1,339.8	1,529.9	1,775.0	2,053.8	2,178.0	15.1

NA = Not applicable
Source: Dataquest (August 1995)

Figure 3-3
PC Graphics Controller Unit Shipment Forecast

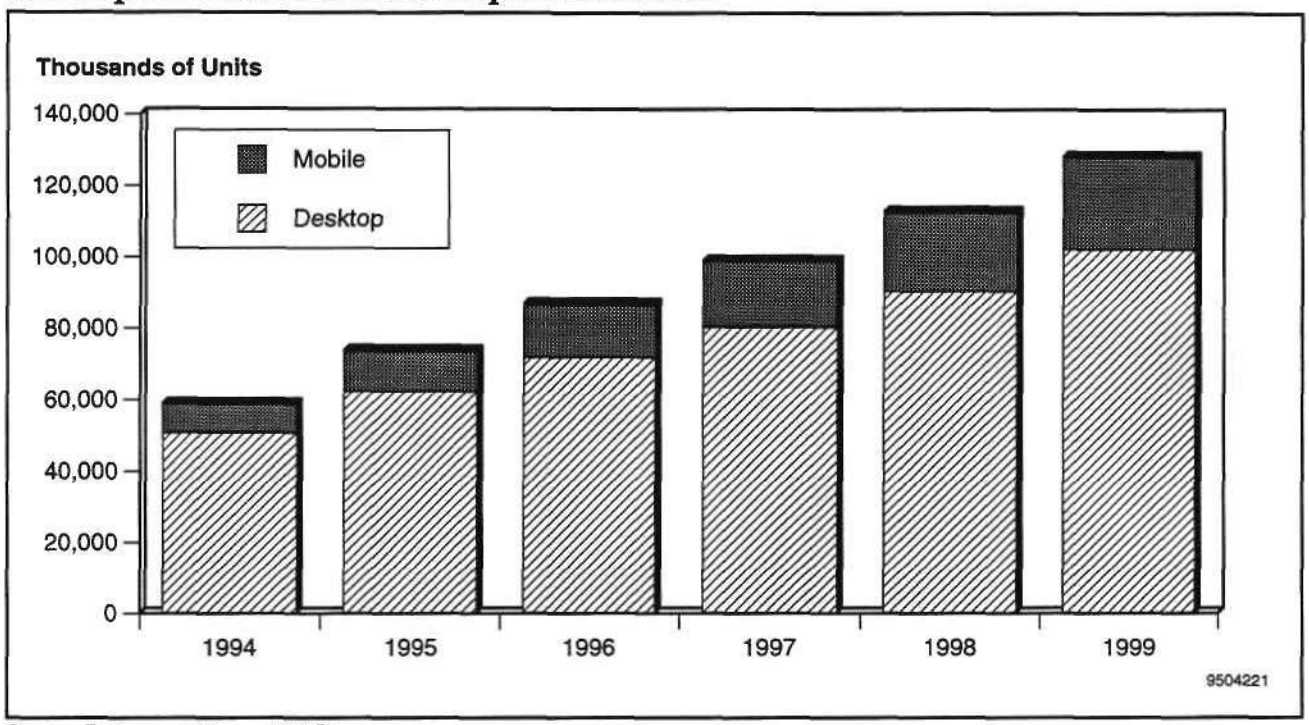

Source: Dataquest (August 1995)

Figure 3-4
PC Graphics Controller Total Market Forecast

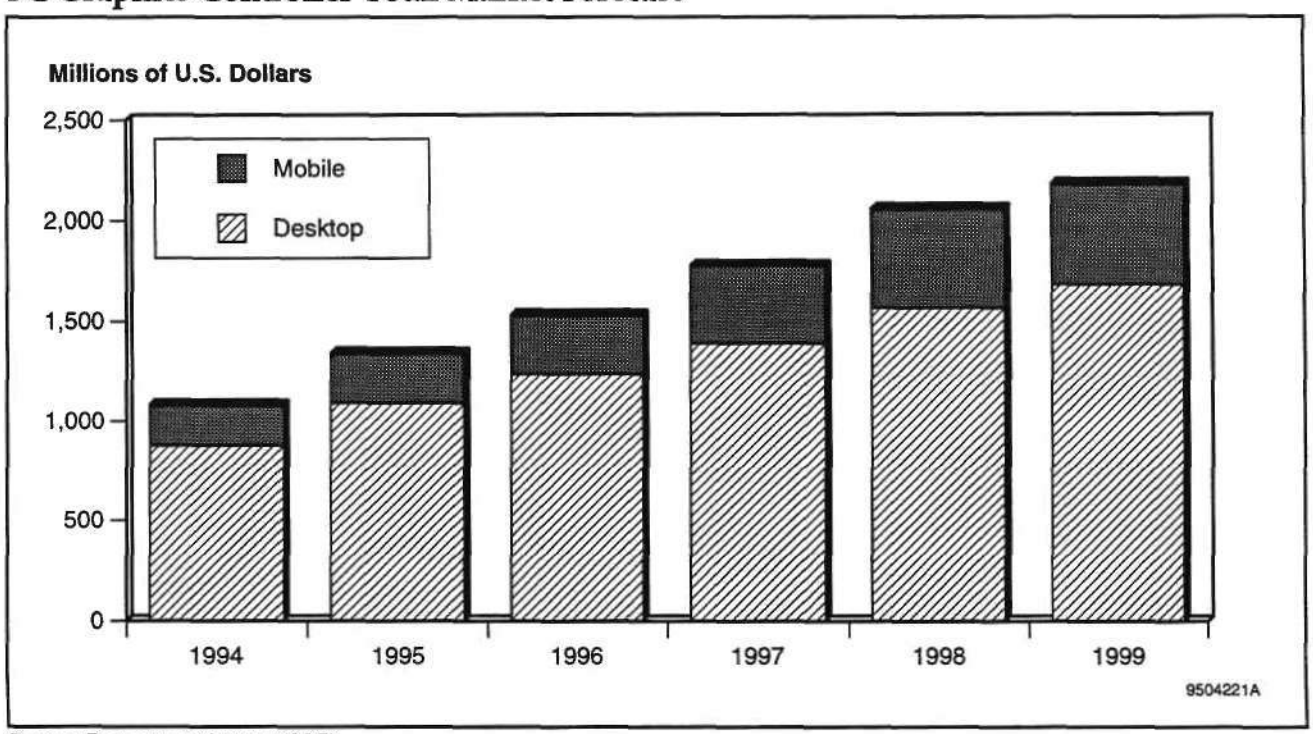

Source: Dataquest (August 1995)

Chapter 4
Feature Trends for Graphics Controllers

Higher Levels of Integration

A functional graphics subsystem requires fewer parts today than in previous years. Features such as on-chip RAMDAC, clock synthesizer, glueless bus interface, and glueless memory interface are becoming standard for mainstream PC graphics. This has the effect of eliminating the miscellaneous passive components and standard logic, as well as off-chip RAMDAC, that once were required for a complete solution. A direct quote from Cirrus Logic's Product Overview booklet is "VGA solution with one or more DRAMs," which applies to several of Cirrus' graphics controllers. Other graphics companies offer similar levels of integration. The bottom line here is that fewer parts are required for greater functionality, which means less board space, fewer components to mount, and cost savings for the OEM. Figures 4-1 and 4-2 show block diagrams of graphics controllers for both desktop and mobile applications. These diagrams show the high level of integration already designed into products shipping today.

Figure 4-1
Block Diagram for a Desktop Graphics Controller with Motion Video Acceleration

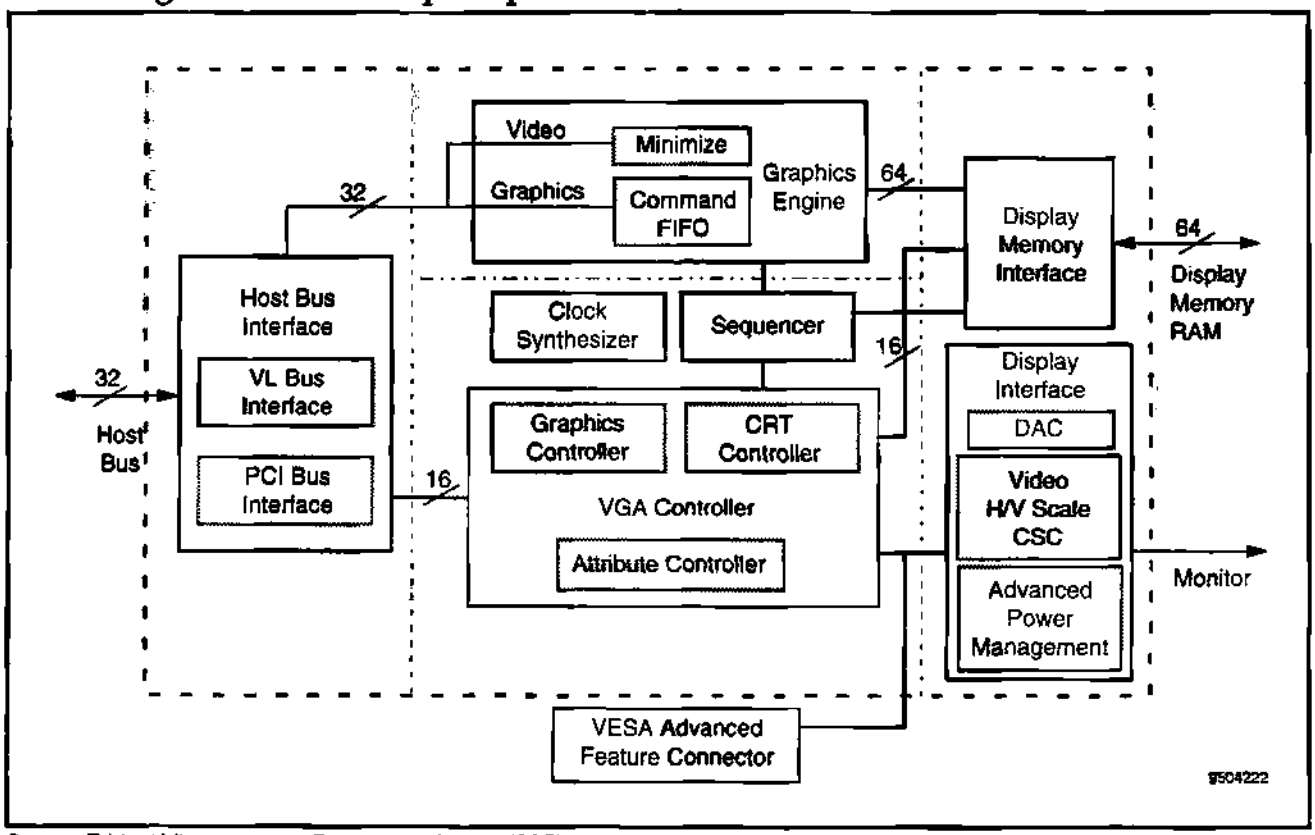

Source: Trident Microsystems, Dataquest (August 1995)

Figure 4-2
Block Diagram for a Notebook Graphics Controller

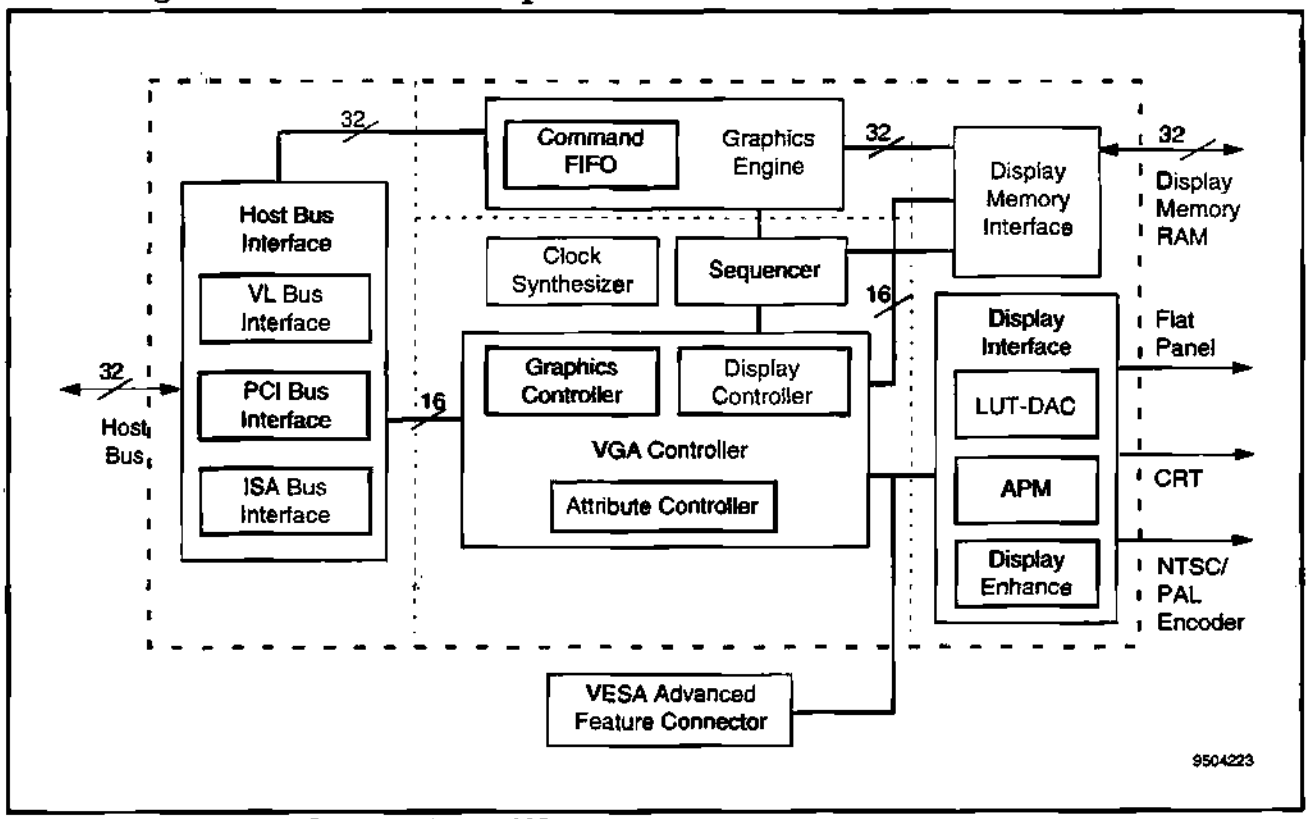

Source: Trident Microsystems, Dataquest (August 1995)

Memory Trends

The memory interface is a critical element of the graphics controller, because it has a strong impact on both performance and cost for the graphics subsystem. The two most popular types of memory for use with graphics controllers are standard DRAM and VRAM. VRAM is faster, but costs more than standard DRAM. Over the last few years, DRAM has replaced VRAM in almost all of the mainstream graphics applications because of the price/performance trade-off. Dataquest surveys show that 2 percent of desktop graphics controllers used VRAM in 1994, down from 3 percent in 1993. Mobile PC graphics controllers are 100 percent DRAM-based, according to the same vendor survey. Cost savings from using DRAM instead of VRAM can be as high as $18 to $20 per system, assuming a typical buffer size of 1MB. VRAM is still in demand for high-performance graphics applications, such as CAD and desktop publishing.

Other types of memory are being considered for frame buffer applications. The goal is better performance than standard DRAM at a lower price than VRAM. Leading candidates are EDO DRAM and Rambus memory (RDRAM). Roughly 1 percent of the desktop graphics controllers sold in 1994 supported EDO DRAM. Look for more variety in 1995 because Cirrus Logic has already announced a graphics controller that uses Rambus memory (RDRAM). RDRAM and synchronous DRAM (SDRAM) have both had design-wins in home video game units, which may lead to greater acceptance in personal computers.

Expansion Bus Interface

PCI made great gains against VL-bus during 1994. Expect this trend to continue for 1995 as Pentium shipments continue to climb. VL-bus graphics controller shipments will follow the 486 MPU shipments because the VL-bus standard was designed for use specifically with 486 microprocessors. 486 MPU unit shipments and VL-based core logic unit shipments are forecast to decline from 1994 to 1995; expect VL-bus graphics controllers to follow that pattern with a chance that some add-in graphics board sales will stretch that peak into 1995.

The outlook for PCI graphics is strong because PCI graphics controllers represent a growing share of a growing market. The PCI bus is processor-independent and is being designed into RISC PCs as well as x86 PCs. Where the PCI bus is designed in, PCI graphics controllers will follow. Look for PCI graphics chips in most Pentium PCs as well as in many of the new RISC PCs based upon the PowerPC architecture. The growth opportunities all point to PCI. Figures 4-3 and 4-4 show the bus interface forecast for desktop and mobile graphics controllers.

Figure 4-3
Bus Interface Forecast for Mobile Graphics Controllers

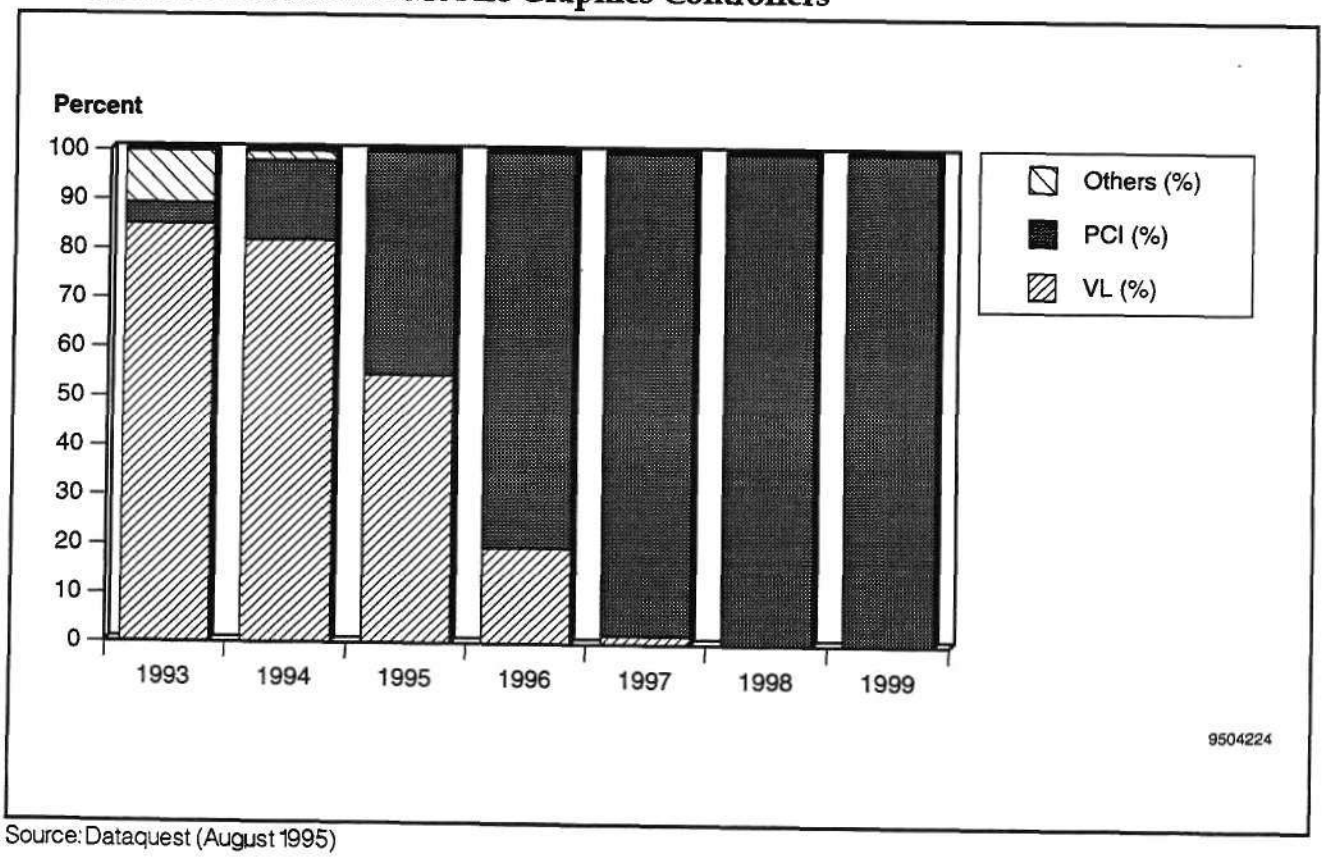

Source: Dataquest (August 1995)

Figure 4-4
Bus Interface Forecast for Mobile Graphics Controllers

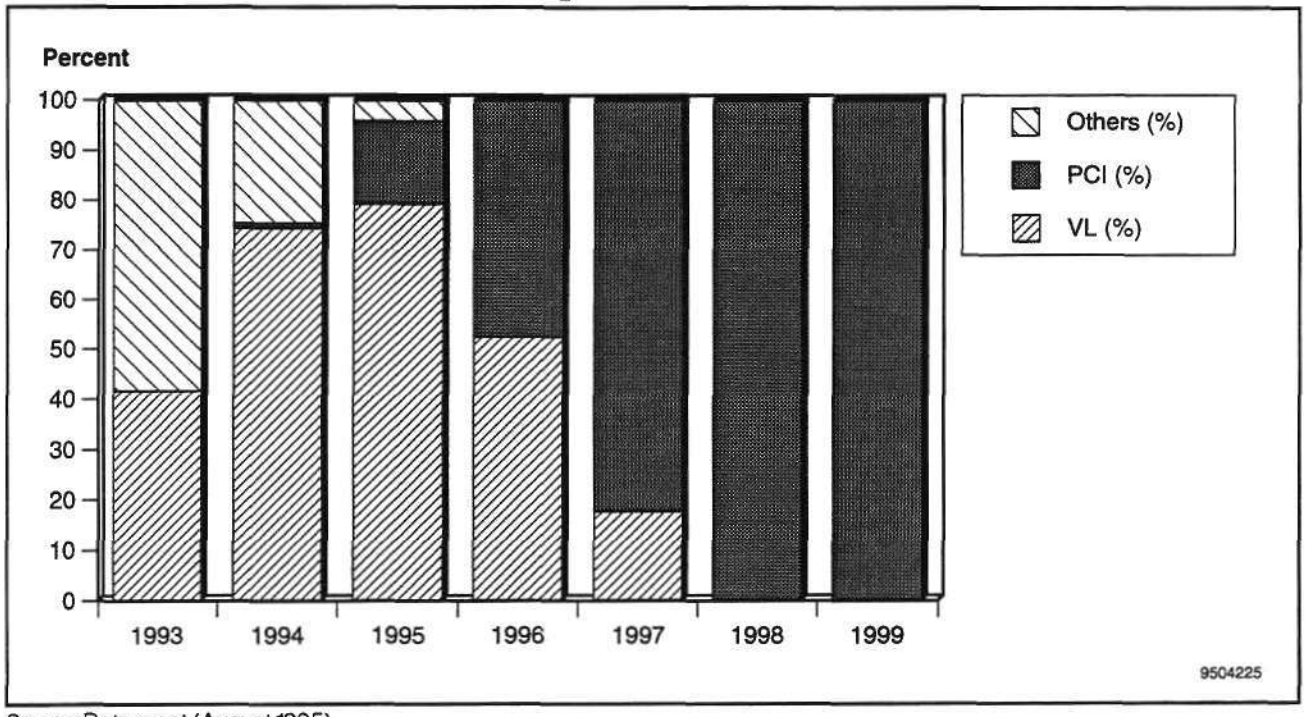

Source: Dataquest (August 1995)

Memory Bus Width

Demand for graphics performance is pushing graphics companies to increase the memory bus width for their graphics controllers. A wider memory bus increases memory bandwidth without requiring faster, more expensive DRAM. The transition to 64-bit buses has introduced a new problem, however; the wider bus requires more total memory, which increases the total cost for the graphics subsystem.

The most popular memory bus width today is 32 bits. This is a convenient width because 256Kbx16 DRAM devices are commonly used for graphics memory and two of these devices provide 1MB of total memory with a 32-bit width. A 64-bit memory bus requires 2MB of memory (four devices, each 16 bits wide). If OEMs do not want to increase the graphics memory to 2MB from 1MB (doubling the total cost of graphics memory), then they cannot take advantage of a 64-bit bus. For this reason, many graphics controllers have a flexible bus that tolerates both 32-bit and 64-bit implementations.

The growth in memory bus width has reached its logical end for mainstream applications because of the higher memory requirement. It is unlikely that PC OEMs will increase standard graphics memory configurations from 1MB to 2MB just to improve graphics performance because the $30 cost is too high. Graphics controller companies need to introduce products that support new DRAM architectures before this bottleneck problem can be solved. Synchronous architectures such as Rambus DRAM (RDRAM) and synchronous DRAM (SDRAM) are the

most likely long-term solutions for graphics memory. Semiconductor vendors might use the burst characteristics of synchronous memory to emulate a 64-bit (or larger) bus on a narrower, multiplexed bus. This technique would provide much greater bandwidth without requiring larger amounts of memory.

Motion Video Support

Support for motion video is one of the newest features for graphics controllers. About 3 percent of the graphics controllers shipped in 1994 provide some support for motion video, and that number is up from roughly 0.1 percent in 1993. Video features such as color space conversion, image scaling, shared frame buffers, and pixel interpolation are in demand for multimedia PCs. Decompression is also in demand, but hardware decompression will cost an OEM $20 or more to add to a PC. Software codecs (coder/decoder) are more popular because they can be changed easily and use the CPU instead of requiring special hardware. The video features popping up in graphics controllers now complement these soft codecs (as well as hardware codecs) by handling the video stream efficiently after the decompression step. Intel Corporation's Native Signal Processing (NSP) embodies the idea of running software codecs on the CPU instead of on fixed-function hardware for modem, audio, and even motion video codec functions. Also, Windows 95 offers new software interfaces to simplify the use of software codecs. A Pentium-based PC running Windows 95 should have motion video performance acceptable for the mass-market multimedia crowd.

The trade-off between software codecs and hardware codecs focuses on flexibility, performance, and price. Using a software codec running on the CPU allows the users to change standards and avoid the cost of dedicated hardware, but does involve loading the system resources (CPU and PCI bus) much more heavily. Intel and other MPU vendors will provide faster and faster MPUs to handle that data load, but application software vendors have always seized the opportunity to add more features to their own products, leaving fewer spare mips than one could otherwise expect.

Some graphics companies have taken a strong position by introducing hardware codecs for video data. This strategy may conflict with Intel's Native Signal Processing (NSP) at the low- and midrange performance levels. Semiconductor vendors should consider designing their low- to midrange products to complement the use of software codecs. Graphics controllers must handle the graphics windows anyway, so they are a natural location for handling video windows. Color space conversion and scaling are also easily handled by the graphics controller. This strategy leaves the codec functions completely to the CPU. High-end applications will probably require dedicated hardware, so there will continue to be a market for hardware codecs.

One other trend for graphics vendors to consider when planning their graphics strategy is the role of Intel's NSP as a market enabler. NSP may allow users to see good video performance and convince enough of them that a hardware upgrade is worth the price. In that case, the hardware codec market will develop as an aftermarket upgrade rather than a standard feature. Semiconductor vendors are advised to develop modular solutions where the hardware codec could be added in the retail upgrade channel.

3-D Graphics Acceleration

3-D graphics features such as polygon drawing and shading techniques are currently the domain of specialty chips like the GLINT chip by 3D Labs. These features have been designed into the latest generation of home video games, but have not yet penetrated the PC market. Demand for these features will increase when Windows 95 becomes more popular as a computer gaming platform. 3-D games are already popular in the arcades and on home video game systems, so it is likely that these titles will be available on the PC. Semiconductor vendors may want to add 3-D acceleration to their graphics controllers for 1997 to protect against lower ASPs and to add value to differentiate their products.

Power Management

Power management on the desktop has finally made it to the graphics controller. Large monitors can use as much energy as the rest of the computer and offer substantial energy savings if powered down (or off) when not in use. Many graphics controllers are now supporting VESA's display power management signaling (DPMS) as part of the "green" PC wave. Graphics controllers with DPMS actively place the monitor into sleep modes or even turn the electron gun off if the PC is idle. Software utilities allow the end user to change these time-out intervals and even turn this feature off. The ability to turn off the DPMS features is important because DPMS graphics controllers may damage non-DPMS monitors.

DPMS is essentially a software-controlled feature that requires support from the graphics hardware. Some screen savers complicate this issue by making the PC appear to be in use to the power management software. Many software companies are now including DPMS support into their screen savers so the power management features are actually used. As consumers and businesses continue to demand "green" PCs, power management features like DPMS are becoming necessities.

Chapter 5
Players and Groundbreaking Products

Players

Cirrus Logic

- Strong growth for both revenue and unit shipments in the graphics controller market, with an increasing share of a growing market.
- Diverse product offerings including core logic (Pico Power), audio IC's (Crystal Semiconductor), communications ICs, and mass storage ICs.
- Revenue for fiscal 1995, which ended in March 1995, was $889 million, up 60 percent (from $557.3 million) for fiscal 1994.
- Unit shipments for graphics controllers increased 57.6 percent from calendar 1993 to calendar 1994.
- Net income increased 35.2 percent to $61.4 million for fiscal 1995.
- Has a joint manufacturing venture with IBM known as MiCRUS to ensure growing access to manufacturing capacity. Recently accelerated investment in MiCRUS with target of significant production in second half of fiscal year 1996.
- Headquartered in Fremont, California.

Trident

- Very strong earnings growth, with $0.42 per share earnings for the first three quarters of fiscal 1995 versus $0.06 per share for same three quarters in fiscal 1994.
- Class-action lawsuit against company and directors has been settled, subject to court approval. If approved, Trident will pay $1.4 million of the $3.15 million settlement.
- Headquartered in Mountain View, California.

S3

- Revenue and earnings grew faster than unit shipments.
- Second-quarter 1995 revenue of $70.6 million compared to second-quarter 1995 revenue of only $27.5 million.
- Recently acquired Floreat Incorporated, for communications software expertise.
- Added NEC and LG Semicon to foundry list that includes IBM and TSMC.

Tseng Labs

- Faces a challenging time of sharply falling revenue.
- 1994 revenue was $0.49 per share.
- 1995 revenue and shipments will depend greatly on the ET6000 chipset meeting the schedule of sampling in the third quarter and volume shipments in the fourth quarter.
- Headquartered in Newton, Pennsylvania.

Western Digital

- Microcomputer Products Group represents only 7 percent of sales.
- Fourth-quarter fiscal 1995 revenue was $39.5 million for the Microcomputer Products Group, down 18 percent from $48.3 million in the third quarter, but only down 12 percent from the fourth quarter of fiscal 1994.
- 100 percent of units shipped to the mobile PC market.
- Headquartered in Irvine, California.

ATI

- Third-quarter fiscal 1995 revenue was $94.6 million versus 51.6 million for the same quarter last year.
- Sells graphics controllers as well as complete graphics cards.
- Recent design-win in Apple Power Macintosh line with the mach64 product.
- Headquartered in Thornhill, Ontario.

Chips & Technologies

- Fiscal 1995 revenue was $104.7 million, up 42.6 percent from $73.4 million in fiscal 1994.
- Net income was up 246 percent to $9.4 million for fiscal 1995.
- 75 percent of units were shipped to the mobile PC market.
- Headquartered in San Jose, California.

Sierra Semiconductor

- Second-quarter fiscal 1995 revenue was $46.1 million, up 107 percent from the same quarter last year, with a 184 percent increase in earnings per share for the same period.
- Diverse product portfolio with a combination of graphics controllers, audio ICs, and communications ICs.
- Headquartered in San Jose, California.

Oak Technology

- Revenue for the three quarters ending March 31, 1995, was $68.5 million, up 141 percent over the same period last year.
- Diverse product offerings, including IDE/ATAPI interface chips in addition to graphics controllers.
- Completed an IPO in February 1995.
- Headquartered in Sunnyvale, California.

S-MOS Systems

- Privately held, so financial information is not disclosed.
- Estimated to have $150 million annual revenue.
- Diverse product offerings.
- Estimated 230 employees.
- Headquartered in San Jose, California.

Groundbreaking Products

S3 Puts MPEG on the Motherboard

S3 became the first company to introduce an MPEG hardware solution for the mass market when it announced its Cooperative Accelerator Architecture on June 13, 1995. Other companies have MPEG-decoder chips, but S3 designed its chip as part of a three-chip multimedia chipset that includes a graphics controller, MPEG decoder, and audio DAC, all tied together with S3's proprietary Scenic Highway interface. This bold move runs against the current trend, embodied by Intel's Native Signal Processing (NSP), of using software to decode MPEG data.

S3's MPEG strategy appears to have a performance advantage over competing hardware codecs because of the integration with the graphics controller via the Scenic Highway bus. S3 has introduced its MPEG chip at $35 in 10K quantities, almost twice as much as some competing MPEG chips, which trade in the $18 to $20 price range. Dataquest expects the S3 MPEG solution to add $90 to $100 to the retail price of the PC. Compaq has endorsed S3's solution and will include it on some PCs for Christmas 1995. This will be the first big test for MPEG hardware bundled into mainstream PCs. Dataquest believes demand for hardware codecs is very elastic and that the current price is limiting market demand. The challenge for S3 and Compaq is to convince computer buyers that this feature is worth the additional cost when less expensive competing solutions have realized only marginal success.

Cirrus Logic Uses Rambus DRAM

Cirrus Logic raised the bar on graphics memory performance when it announced product support for Rambus DRAM. Other graphics vendors are waiting to see which of the alternative DRAM architectures will prevail, but Cirrus is betting on RDRAM. RDRAM offers a 500MB per second transfer rate with only a 10 percent price premium over standard DRAM, which makes it attractive on a price/performance basis. The looming question is product availability. Several memory manufacturers have announced they will make RDRAM, but mass market acceptance of the architecture is uncertain. The VisualMedia accelerator announced by Cirrus will work only with RDRAM, so there is additional market risk for this product associated with memory availability and pricing. Expect Cirrus' endorsement of RDRAM to boost its legitimacy as a product.

VLSI Ties Graphics to Core Logic

VLSI Technology broke the mold for graphics controllers by providing a special interface to its new Lynx chipset. A special two-wire interface between the Lynx chipset and a select group of VLSI's graphics controllers can boost performance by up to 15 percent. Status-checking information can be sent over this interface instead of the PCI bus, which should improve performance when the PCI bus is heavily loaded by disk or network activity. The graphics controllers will work with other vendors' PCI chipsets, but the extra performance is lost.

The pairing of these chips allows VLSI to jump-start sales of its graphics controllers by leveraging its design-wins for core logic. VLSI has shipped more core logic chipsets than any other vendor for the past two years, according to Dataquest surveys. This tactic is a departure from the trend

toward more standard products and fewer custom solutions because there is vendor-specific synergy between the graphics and the core logic. The main interface complies with the PCI specification, but the additional interface is a proprietary standard rather than an open one.

NeoMagic Integrates the Frame Buffer into the Graphics Controller

NeoMagic announced a notebook graphics controller with 1MB of integrated memory for the frame buffer. This one-chip graphics solution breaks ground by combining logic and DRAM on a single chip. DRAM manufacturing processes compromise transistor speed and levels of interconnect in favor of high reliability, so these processes are not optimized for advanced logic functions. NeoMagic claims it has addressed this issue and is making the chip with a process similar to that of 16Mb DRAM and a slightly larger die. The logic speed is slower because of this process, but the level of integration appears to compensate for the difference.

NeoMagic's challenge now is to manufacture the device at a competitive price. OEMs know the benefits of integration and low-power consumption, but may not pay much of a premium over competing solutions. If NeoMagic is successful, look for competing products for the notebook market. Similar products for the desktop are unlikely at least until 64Mb DRAM processes become common. The opportunity at that point will depend mostly upon high performance at a low price, because space and power requirements are less restrictive on desktop systems.

Chapter 6
Dataquest Perspective

The role of graphics controllers is changing just as rapidly as the role of the PC. A higher percentage of PC shipments now goes to the home, where the PC supports personal productivity, communication, and entertainment. Digital video is currently more important for consumer entertainment than business desktop productivity and will most likely continue to be so. The trend for graphics controllers is adding new functionality to address the changing role of the personal computer.

This year, the new feature is motion video acceleration. The next feature to be integrated into the graphics controller will probably be 3-D graphics acceleration. It is unlikely that hardware codecs will be integrated soon because of the additional gate count required and because of the increasing ability of faster MPUs to provide that function. Figure 6-1 depicts this graphically.

Another critical issue for semiconductor vendors is support for alternative DRAM architectures. Graphics performance is currently limited by memory bandwidth issues, and this will not change significantly as long as standard DRAM and VRAM are used for the graphics memory. Semiconductor vendors have an opportunity to differentiate their products by supporting new types of DRAM that solve the memory bandwidth issue without requiring larger frame buffers.

Figure 6-1
Technology Trends for PC Graphics Controllers

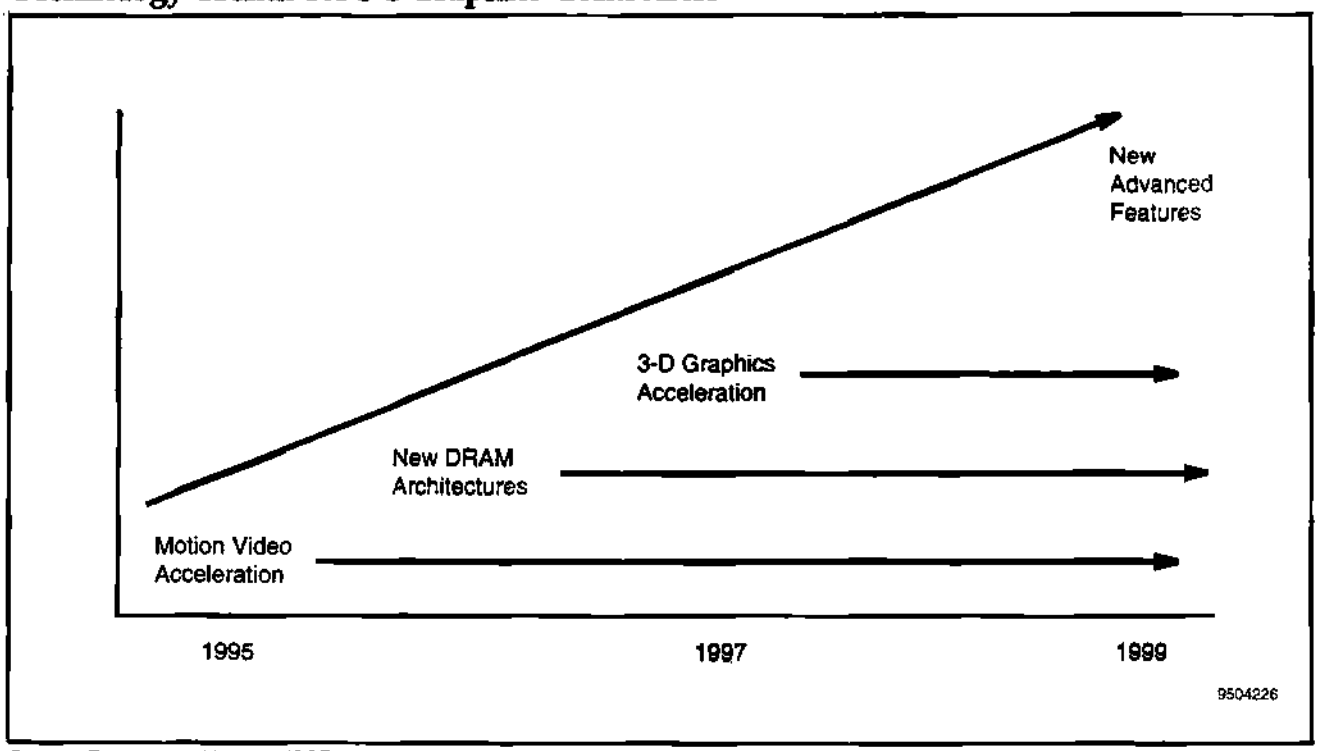

Source: Dataquest (August 1995)

One last issue that semiconductor vendors must consider is the large number of competitors in the graphics controller market. It is unlikely that the market can support all the players now on the roster, even with continued growth in PC shipments. PC OEMs and end users are sure to benefit as new entrants to the graphics market cut prices to gain market share and establish their presence in the industry. New features will play a critical role in differentiating products from competing semiconductor vendors. Those same features will also offer the greatest opportunity for semiconductor vendors to charge a relatively stable price for a graphics controller from one year to the next. Without new features, graphics controller companies will be faced with sharply lower ASPs and lost accounts.

For More Information...

Geoff Ballew, Industry Analyst ..... (408) 437-8676
Internet address .....gballew@dataquest.com@dataquest.com
Via fax ..... (408) 437-0292

Dataquest
a company of
The Dun & Bradstreet Corporation

DATAQUEST WORLDWIDE OFFICES

NORTH AMERICA

Worldwide Headquarters
Dataquest Incorporated
251 River Oaks Parkway
San Jose, California 95134-1913
United States
Phone: 1-408-468-8000
Facsimile: 1-408-954-1780

Dataquest Incorporated
Nine Technology Drive
P.O. Box 5093
Westborough, Massachusetts 01581-5093
United States
Phone: 1-508-871-5555
Facsimile: 1-508-871-6262

Dataquest Global Events
3990 Westerly Place, Suite 100
Newport Beach, California 92660
United States
Phone: 1-714-476-9117
Facsimile: 1-714-476-9969

Sales Offices:
Washington, DC (Federal)
New York, NY (Financial)
Dallas, TX

LATIN AMERICA

Research Affiliates and Sales Offices:
Buenos Aires, Argentina
Sao Paulo, Brazil
Santiago, Chile
Mexico City, Mexico

EUROPE

European Headquarters
Dataquest Europe Limited
Holmers Farm Way
High Wycombe, Bucks HP12 4XH
United Kingdom
Phone: +44 1494 422 722
Facsimile: +44 1494 422 742

Dataquest Europe SA
Immeuble Défense Bergères
345, avenue Georges Clémenceau
TSA 40002
92882 - Nanterre CTC Cedex 9
France
Phone: +33 1 41 35 13 00
Facsimile: +33 1 41 35 13 13

Dataquest GmbH
Kronstadter Strasse 9
81677 München
Germany
Phone: +49 89 93 09 09 0
Facsimile: +49 89 93 03 27 7

Sales Offices:
Brussels, Belgium
Kfar Saba, Israel
Milan, Italy
Randburg, South Africa
Madrid, Spain

JAPAN

Japan Headquarters
Dataquest Japan K.K.
Shinkawa Sanko Building
6th Floor
1-3-17, Shinkawa
Chuo-ku, Tokyo 104
Japan
Phone: 81-3-5566-0411
Facsimile: 81-3-5566-0425

ASIA/PACIFIC

Asia/Pacific Headquarters
7/F China Underwriters Centre
88 Gloucester Road
Wan Chai
Hong Kong
Phone: 852-2824-6168
Facsimile: 852-2824-6138

Dataquest Korea
Suite 2407, Trade Tower
159 Samsung-dong, Kangnam-gu
Seoul 135-729
Korea
Phone: 822-551-1331
Facsimile: 822-551-1330

Dataquest Taiwan
11F-2, No. 188, Section 5
Nan King East Road
Taipei
Taiwan, R.O.C.
Phone: 8862-756-0389
Facsimile: 8862-756-2663

Dataquest Singapore
105 Cecil Street
#06-01/02 The Octagon
Singapore 0106
Phone: 65-227-1213
Facsimile: 65-227-4607

Dataquest Thailand
12/F, Vanissa Building
29 Soi Chidlom
Ploenchit Road
Patumwan, Bangkok 10330
Thailand
Phone: 662-655-0577
Facsimile: 662-655-0576

Research Affiliates and Sales Offices:
Melbourne, Australia
Beijing, China

Dataquest

a company of
The Dun & Bradstreet Corporation

PSAM-WW-FR-9502 Qty: 1
Ms. Maria Valenzuela
Dataquest Incorporated
1-1100
.,
--INTERNAL DIST.--

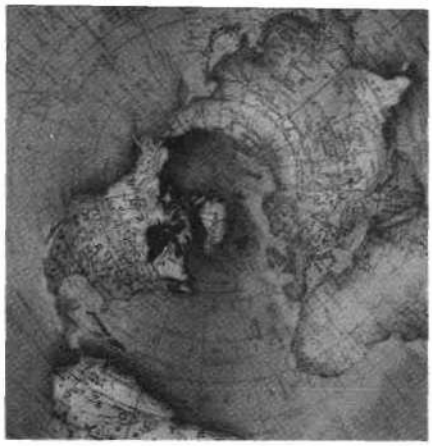

Dataquest

PC Core Logic: A Focused Analysis

Focus Report

Program: Semiconductor Directions in PCs
Product Code: PSAM-WW-FR-9501
Publication Date: June 12, 1995
Filing: Focus Studies

PC Core Logic: A Focused Analysis

Focus Report

Program: Semiconductor Directions in PCs
Product Code: PSAM-WW-FR-9501
Publication Date: June 12, 1995
Filing: Focus Studies

Table of Contents

List of Figures

List of Tables

Chapter 1
Executive Summary

PC core logic chipsets are as critical to the overall performance of the personal computer as any other component in the system. The PC core logic chips are the interface between the MPU complex (microprocessor and cache memory) and the rest of the PC. These chips manage the system resources that make up the PC. Many of the critical performance bottlenecks, such as the cache algorithms, DRAM refresh, memory management, power management, expansion bus control, local bus control, and DMA control, are handled directly by the PC core logic chips. A poorly designed chipset architecture will severely cripple a PC's performance, completely independent of the microprocessor used in the system.

One key point that will be supported in this report is that traditional users of custom PC core logic chipsets are rapidly moving away from the custom approach to a standard product approach. This transition took a step function change in 1994 when such companies as IBM and Compaq began a rapid move from custom solutions developed in-house to standard solutions. For IBM this involved a move to chipsets manufactured by OPTI, while Compaq made the switch from in-house-developed custom chipsets to standard product chipsets manufactured by VLSI Technology.

Also, the local bus of choice in 1993 and 1994 was the VESA or VL bus, but nearly 100 percent of new development work is focusing on the Peripheral Component Interface (PCI).

Intel is becoming a dominant force in PC core logic chipsets, driven primarily by its position in the Pentium chipset market. We expect Intel to continue its chipset focus in order to facilitate the early adoption of new generations of microprocessors. However, Intel needs to temper its position in PC core logic products. It appears to have modified its original intent of enabling the market to becoming a dominant supplier of PC core logic. Although it certainly has the resources to force this occurrence, following through on this path may create a situation that can hamper the time-to-market plans of its own next-generation microprocessors.

Several non-Intel-compatible microprocessor vendors are trying to enter the market for desktop PCs. Choosing which of these new entrants to support requires all of the market intelligence a company can muster. Advanced Micro Devices (AMD) now is the only truly Intel-compatible provider of a desktop solution for the desktop PC market. Choosing to support any other MPU vendor will require new R&D development. This R&D effort can be minimized, depending on the company in question. After AMD, Cyrix requires minimal modifications to support its data bursting scheme. Nexgen requires still further modifications, and the RISC vendors all require extensive R&D efforts. However, they should be quite helpful to any third-party PC core logic vendor seeking to provide support for its particular architecture.

The following chapters provide a more detailed look at these and other issues.

Chapter 2
The Players

According to a recent Dataquest survey, a familiar name is found at the top of the chart of chipset producers in 1994 (see Table 2-1 and Figure 2-1). VLSI Technology once again is the market leader. VLSI led all challengers vying for the No. 1 position in total PC core logic chips. Although the company's market share fell from 21 percent in 1993 to 17.9 percent in 1994, it was still able to hold off the hard charge of Taiwan-based Silicon Integrated Systems (SIS), which grew at a pace nearly double that of the PC core logic industry and leaped over both United Microelectronics (UMC) and OPTI. At its current pace, SIS should overtake VLSI in 1995. Both UMC and OPTI grew in the 20 percent range and were able to hold off the challenge of Intel, whose meteoric rise may propel it above SIS and VLSI into the No. 1 position before 1995 is finished. The only company to outperform Intel in terms of growth was the Cirrus Logic subsidiary, Pico Power, whose 253 percent growth rate propelled it from the No. 10 position to No. 8.

Growth Adjusted

This 43 percent market growth is skewed somewhat by the rapid transition of PC OEMs, such as Compaq and IBM, switching from the use of custom PC core logic to standard off-the-shelf PC core logic. As was estimated in an earlier report from Dataquest's Semiconductor Directions in PCs program, the use of custom PC core logic in desktop PCs fell from 15 percent in 1993 to 9 percent in 1994 and is expected to fall to less than 5 percent in 1995, while the use of custom PC core logic in notebook PCs fell from 65 percent in 1993 to 32 percent in 1994 and is expected to approach 15 percent in 1995. We expect this trend to continue because

Table 2-1
Top 10 Core Logic Chipset Producers, Worldwide (Thousands of Units)

1994 Rank	1993 Rank	Company	1993 Total	1994 Total	Percentage Change	Market Share (%)
1	1	VLSI	7,392	9,047	22.4	17.9
2	4	Silicon Integrated Systems	4,652	8,600	84.9	17.0
3	2	UMC	6,800	8,200	20.6	16.2
4	3	OPTI	6,344	7,775	22.6	15.4
5	7	Intel	1,500	5,300	253.3	10.5
6	5	ACER	3,900	4,200	7.7	8.3
7	6	ACC Microelectronics	1,983	2,563	29.2	5.1
8	10	Cirrus/Pico Power	325	1,850	469.2	3.7
9	8	Chips & Technologies	1,400	1,700	21.4	3.4
10	9	Symphony	410	860	109.8	1.7
		Others	800	535	-33.1	1.1
		Total Shipments	35,506	50,630	42.6	

Source: Dataquest (June 1995)

Figure 2-1
Top 1994 Core Logic Chipset Producers, Worldwide

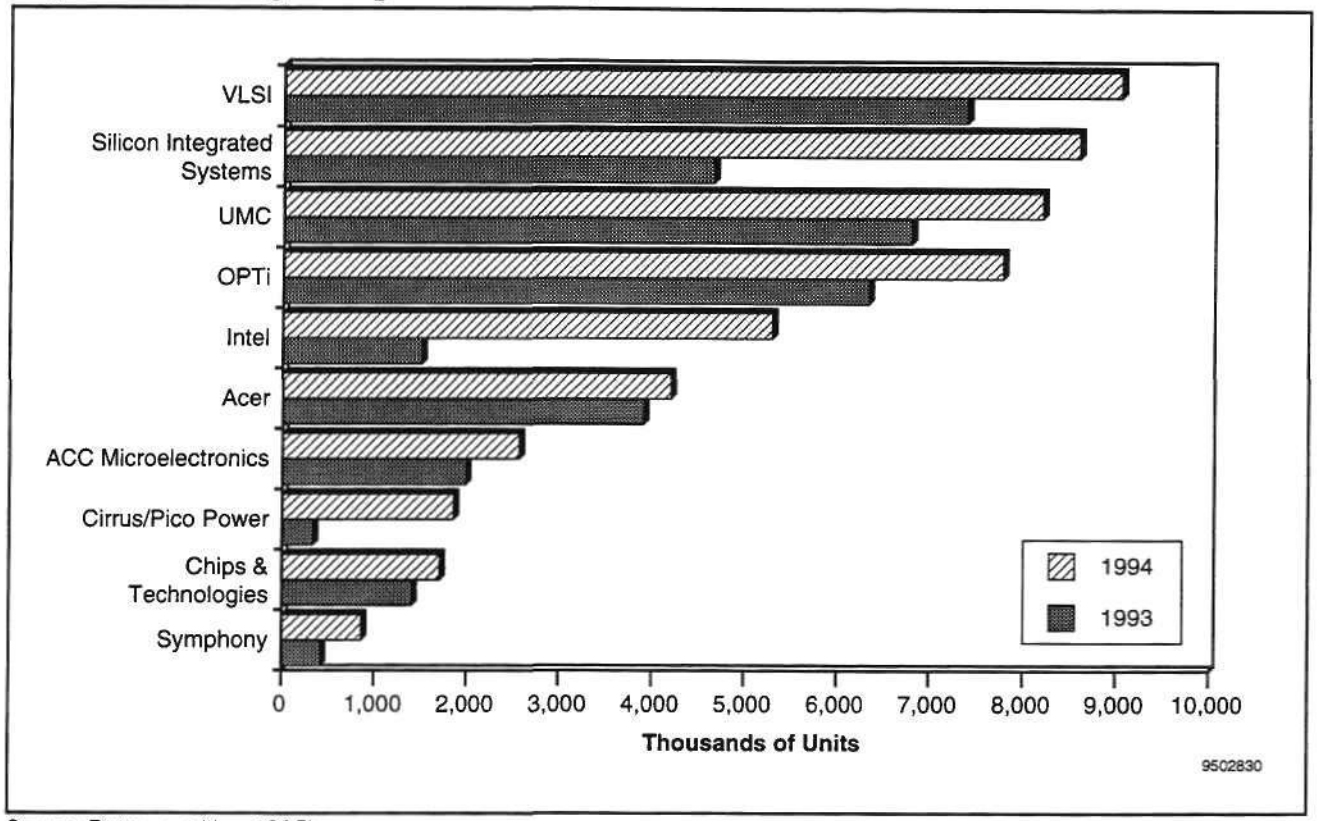

Source: Dataquest (June 1995)

major OEMs are rapidly disbanding, or significantly reducing, internal PC core logic R&D efforts. Because Dataquest does not include the captive PC core logic shipments in its market share calculations, the growth of the overall PC core logic market is overstated by about 6.3 percent for 1994. Adjusting for this, growth for the overall market for PC core logic becomes 37 percent. The 43 percent growth figure reflects the average growth experienced by the merchant market suppliers able to replace the custom products previously used, in addition to the growth of the PC and motherboard upgrade markets.

The Desktop: Still a Rising Star

In the high-volume world of desktop PCs, SIS vaulted from the No. 4 position to the No. 1 position, surpassing last year's desktop leader VLSI as well as the No. 2 and No. 3 players, UMC and OPTI (see Table 2-2 and Figure 2-2). Based on a 1 percent growth, VLSI fell to No. 4 in 1994 and is in danger of falling further in 1995. Intel is well positioned to challenge SIS for supremacy in the desktop category.

Only three companies experienced growth greater than the market average of 34 percent in 1994: SIS, Intel, and Symphony.

Intel's position in Pentium chipsets and the company's expanding motherboard manufacturing capability are two critical factors leading to our belief that Intel will seriously challenge for the No. 1 position in the

Table 2-2
Top Core Logic Desktop Chipset Producers, Worldwide (Thousands of Units)

1994 Rank	1993 Rank	Company	1993 Total	1994 Total	Percentage Change	Market Share (%)
1	4	Silicon Integrated Systems	4,652	8,510	82.9	19.0
2	2	UMC	6,800	8,200	20.6	18.3
3	3	OPTI	6,241	7,396	18.5	16.5
4	1	VLSI	7,080	7,152	1.0	15.9
5	6	Intel	1,500	5,300	253.3	11.8
6	5	ACER	3,900	4,000	2.6	8.9
7	7	ACC Microelectronics	1,446	1,774	22.7	4.0
8	8	Chips & Technologies	1,400	1,700	21.4	3.8
9	9	Symphony	400	850	112.5	1.9
		Others	0	0	0	0
		Total Shipments	33,419	44,882	34.3	

Source: Dataquest (June 1995)

Figure 2-2
Top 1994 Core Logic Desktop Chipset Producers, Worldwide

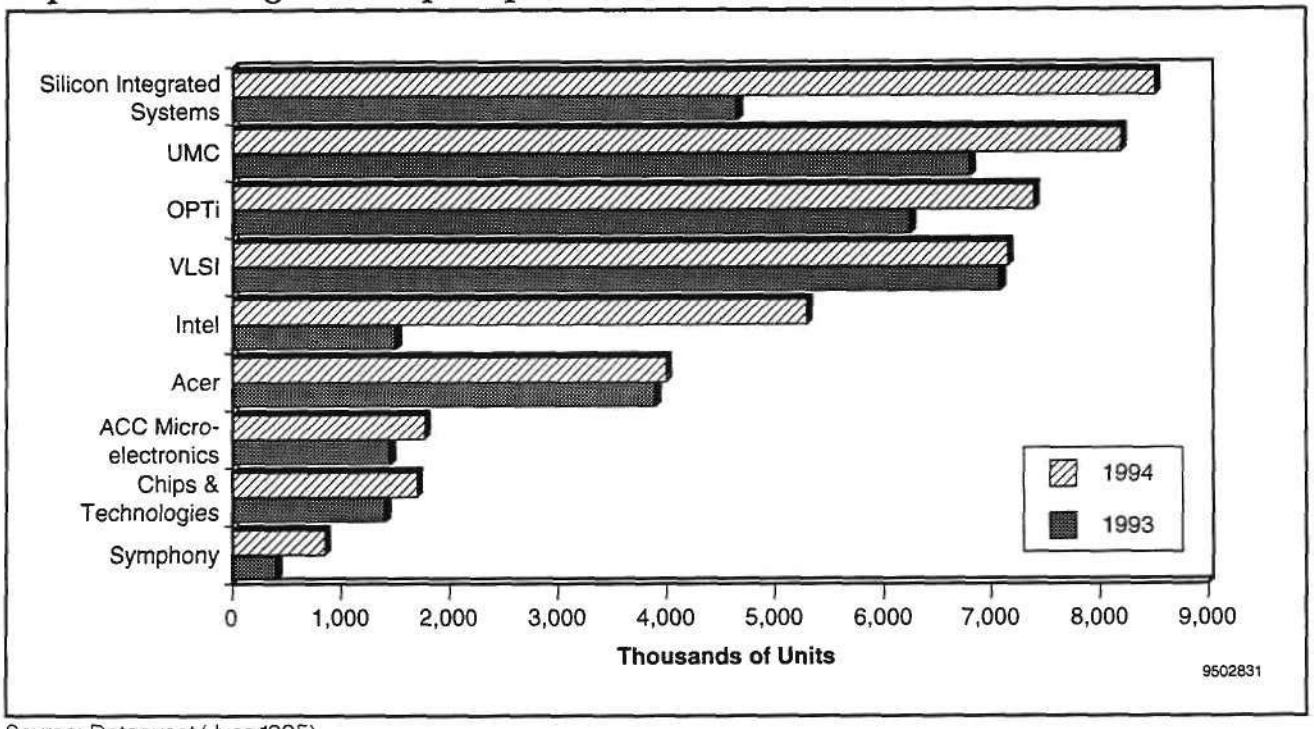

Source: Dataquest (June 1995)

desktop market. We expect Intel to manufacture about 10 million Pentium-based motherboards. Combining this with Intel's latest-generation Pentium PC core logic chipset, Triton, Intel is well positioned to ship well in excess of 10 million units in 1995.

The 1994 top four vendors – SIS, UMC, OPTI, and VLSI – will be hard-pressed to hold off the Intel challenge. A dark horse challenge may come from Acer Laboratories Incorporated (ALI), which has a significant internal market as well as a selling arm actively marketing its 486 and Pentium PC core logic chipsets. The market dynamics are in place to make 1995 a wild year for market share dominance on the desktop.

Notebooks: Amazing Growth?

The 175 percent growth in PC core logic chipsets shown in Table 2-3 for notebooks applies solely to the growth of standard product chipsets and does not include the rapid decline of the use of custom PC core logic. We expect this trend of switching to standard products to continue through 1995. This trend in notebook PCs has lagged a similar trend in desktop PCs because the notebook feature set is sufficiently different for notebook PCs that the PC core logic developed for the desktop PC market is inadequate for the notebook PC market. Third-party PC core logic suppliers rightly focused their efforts on the much larger desktop market because the desktop market is more mature and the feature sets for this market are much more straightforward. PC core logic suppliers supplying to this market could have had a higher level of confidence that their R&D efforts would bear fruit. The burden of developing PC core logic for the notebook market rested initially on the shoulders of the notebook OEM. Now that the notebook feature set is becoming more predictable and third-party PC core logic vendors are beginning to understand how to implement these features, we are starting to see notebook OEMs begin to disband their internal development efforts in favor of using third-party-developed standard products.

Table 2-3
Top Seven Core Logic Notebook Chipset Producers, Worldwide (Thousands of Units)

1994 Rank	1993 Rank	Company	1993 Total	1994 Total	Percentage Change	Market Share (%)
1	4	VLSI	312	1,895	507.4	33.0
2	3	Cirrus/Pico Power	325	1,850	469.2	32.2
3	2	ACC Microelectronics	537	789	46.9	13.7
4	1	Western Digital	800	500	-37.5	8.7
5	5	OPTI	103	379	268.0	6.6
6	NM	ACER	0	200	NM	3.5
7	NM	Silicon Integrated Systems	0	90	NM	1.6
		Others	10	45	350.0	0.8
		Total Shipments	2,087	5,748	175.0	

NM = Not meaningful
Source: Dataquest (June 1995)

Somewhat surprisingly, VLSI attained the No. 1 position in this fast-growth category (see Figure 2-3) because of a key design-win in one of Compaq's portable products, which is not a true indicator of VLSI's present position in the notebook market for PC core logic chips. However, a version of the company's new Eagle chipset is targeted at the Pentium notebook market. Cirrus Logic, via its acquisition of Pico Power, has jumped into the No. 2 position in this category. Cirrus' position is based upon a much broader portfolio of notebook OEM design-wins than that of VLSI, and we expect Cirrus to vault into the No. 1 position in 1995 as its design-wins evolve into high-volume production. ACC Microelectronics finds itself in the No. 3 position, but with less than half the market share of either VLSI or Cirrus. It is running at a tenth of these companies' growth rates, and we do not expect ACC to threaten for the No. 1 position in 1995. The leader in 1993, Western Digital, has begun pulling out of this market; it is not surprising to see it fall to the No. 4 position. Because we expect no change in its focus on this market, we believe that it will fall into the "others" category in 1995. Western Digital will continue to focus in product areas where it believes it has a stronger competitive position. OPTI is beginning to make a move in the portable arena and, along with Acer and SIS, should improve its market share position in 1995.

Figure 2-3
Top Seven Core Logic Notebook Chipset Producers, Worldwide

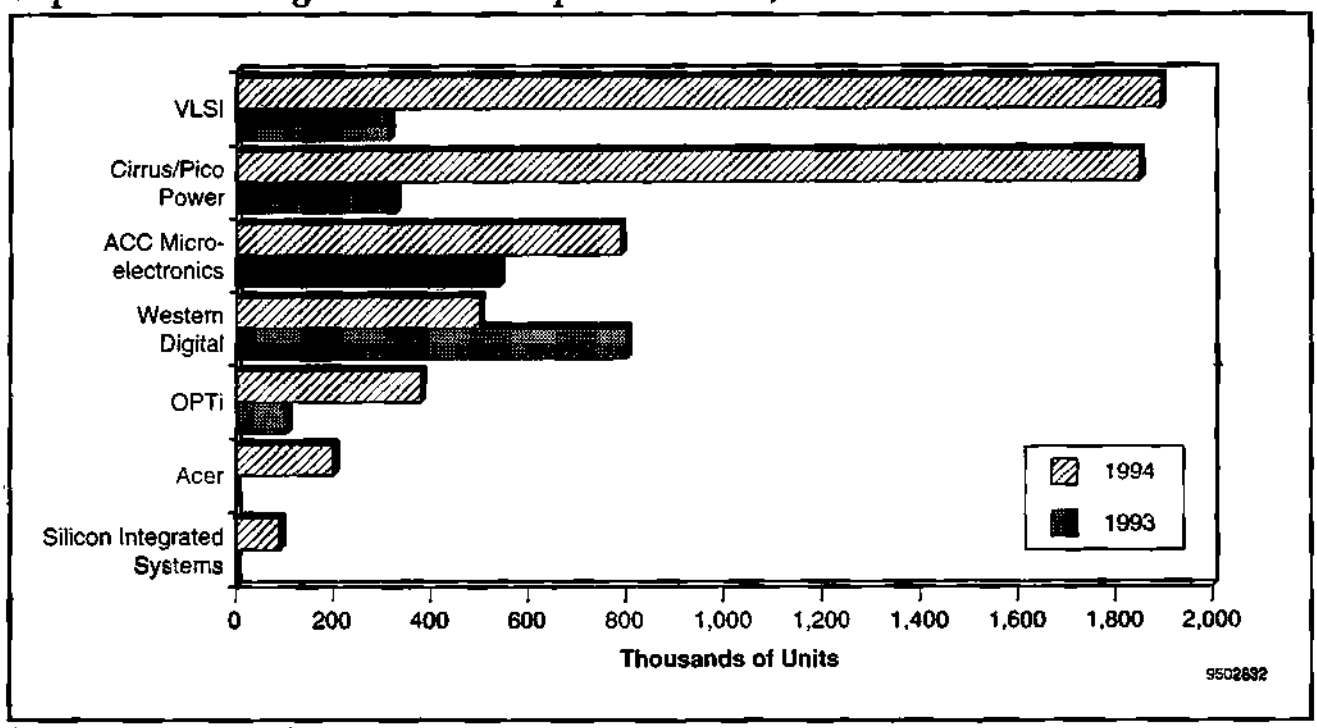

Source: Dataquest (June 1995)

Chapter 3
Where the Industry Has Been and Where It Is Going

The rate of change in the realm of PC core logic is on an accelerating path. Ten years ago, the industry was making the transition from the 286 microprocessor to the 386, with the only technical challenge being how to cope with the 32-bit-wide I/O bus of the 386DX microprocessor. Today we still have some lingering 386 designs, a vast array of 486 product offerings, the Pentium microprocessor with its 66-MHz, 64-bit I/O bus, Nexgen's 586-class processor, the soon-to-be-released clones from AMD and Cyrix, and Intel about to release its next-generation x86 microprocessor still known as the P6. The challenges for the PC core logic vendors have increased with each successive generation of x86 microprocessor.

Cache Is King

The transition from the 286 to 386 involved moving from a 16-bit to a 32-bit bus I/O. Moving from the 386 to the 486 required the addition of a secondary cache controller because PC performance was becoming increasingly important. This feature made the secondary cache controller a key product differentiator. An efficient cache controller could greatly improve the performance of a given PC. Inefficient controllers would often decrease performance because of a very low "hit rate" in cache. This lack of understanding by the PC core logic industry forced many PC OEMs to design their own PC core logic rather than purchase standard product. By the end of the 486 life cycle, everyone in the PC core logic industry seemed to have a good understanding of cache controller algorithms.

The Expansion Bus Wars

The only expansion bus 10 years ago was the Industry Standard Architecture (ISA) bus. Today we still have ISA, some remnants of the Extended Industry Standard Architecture (EISA) bus, the VESA local (VL) bus, and now the PCI bus. During the life cycle of the 386, the PC industry made a few attempts at replacing the incredibly slow ISA bus. The ISA bus, with its 16-bit width and 8-MHz clock rate, was a major performance bottleneck for the PC (see Figure 3-1). The first attempt at replacing this industry standard came from IBM with its Micro Channel Architecture (MCA) bus. IBM's own arrogance during the development and marketing of this bus doomed its chance of success at the outset. The industry had rapidly grown beyond a group of small companies that would happily follow IBM's lead to a group of very successful and fast-growing companies that believed they had a better handle of what the industry needed than did IBM. IBM's launch of the MCA architecture caused the leading companies in the PC industry to come together and develop the alternative expansion bus known as the Extended Industry Standard Architecture (EISA) bus. Although the EISA bus did have a broad number of backers, it had extremely limited success because of a number of factors, including its high cost, very limited add-in card support, and the fact that virtually none of the PC core logic vendors believed the return on investment warranted a new product design. As a result, the industry is still limping along with the ISA bus as the primary expansion bus in today's desktop PCs.

Figure 3-1
ISA Bus Implementation

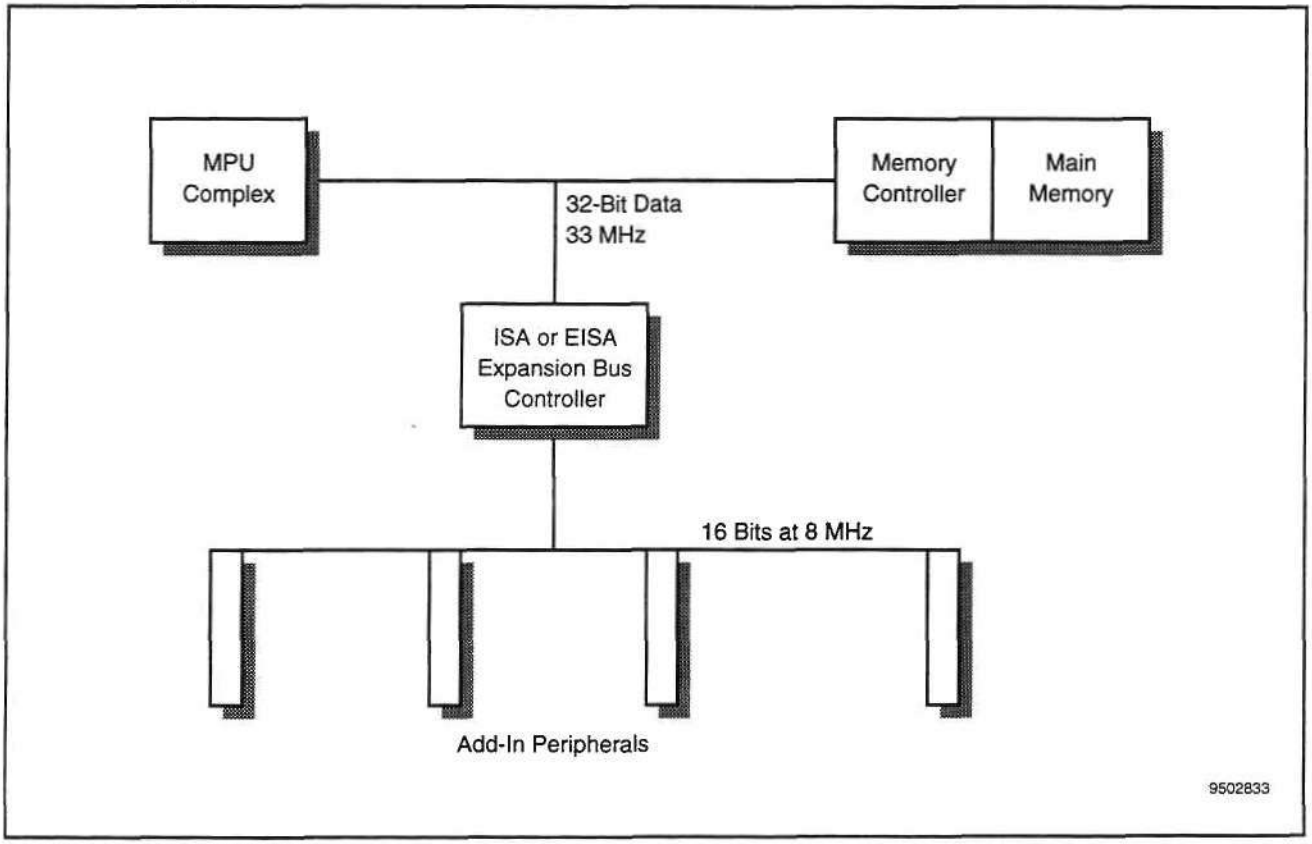

Source: Dataquest (June 1995)

Local Bus to the Rescue

The PC performance bottleneck was so severe that the industry kept working on alternative solutions that would prove acceptable to the industry and provide an adequate road map for future generations of PCs. This effort gave birth to another potential functional differentiator known as the local bus. The first industry-standard local bus to become implemented is the VESA local bus, or more commonly, the VL bus. Another local bus, the PCI, was in a standards committee when the VL concept was first launched, but because the PCI committee was attempting to develop a specification for a local bus that would easily extend into other microprocessor and/or PC architectures, the PCI standard would be years in the making. The VESA committee set its sights on a local bus that would be used in IBM PCs only, and version 1 was targeted specifically at the 486 microprocessor (see Figure 3-2). This less-aggressive specification allowed the VL bus to capture an early market share advantage over the PCI bus that peaked in 1994 with shipments of 33.9 million units giving it two-thirds of the market (see Table 3-1 and Figure 3-3).

Figure 3-2
VL Bus Implementation

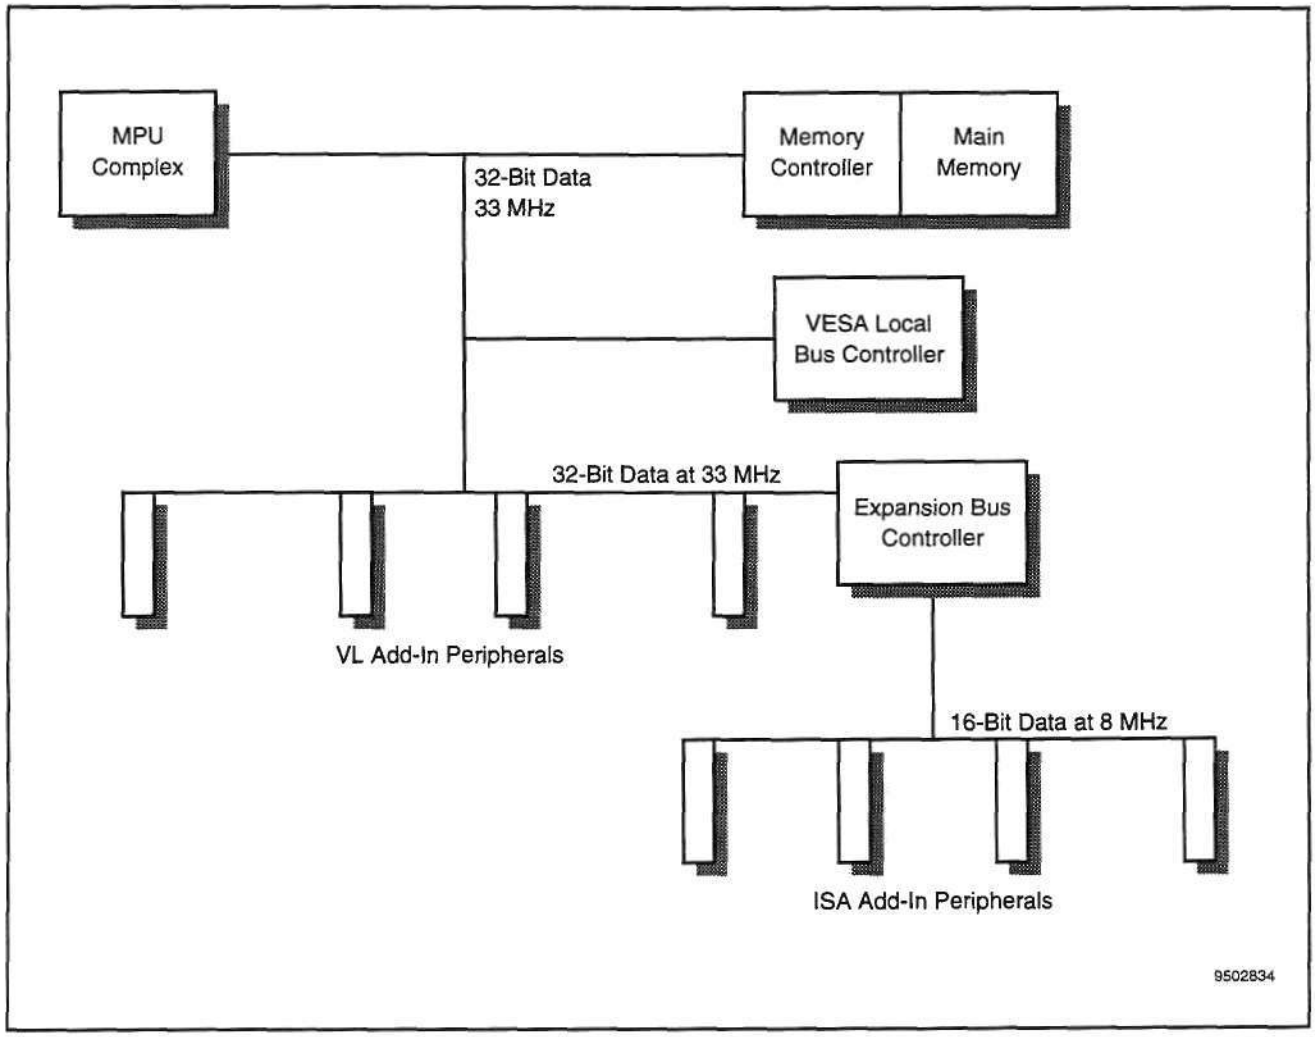

Source: Dataquest (June 1995)

Table 3-1
PC Local Bus Chipset Forecast (Thousands of Units)

	1993	1994	1995	1996	1997	1998	1999
VL	18,844	33,869	30,500	10,900	450	0	0
PCI	1,490	8,385	27,300	57,000	79,500	91,750	103,000
None	15,172	8,377	4,200	2,100	1,050	250	0

Source: Dataquest (June 1995)

Figure 3-3
PC Local Bus Chipset Forecast

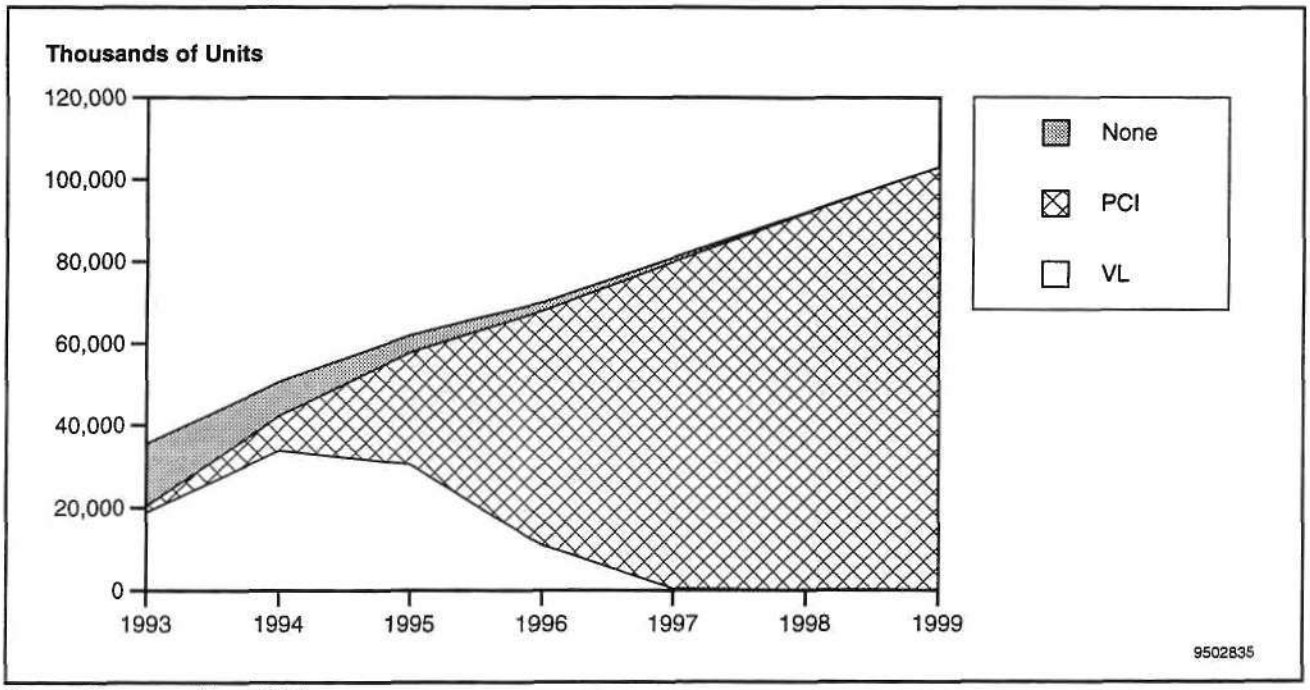

Source: Dataquest (June 1995)

The PCI bus has many advantages over the VL local bus:

- It is buffered from the microprocessor memory bus interface (see Figure 3-4), which allows the microprocessor to operate concurrently while the PCI bus is in use.
- It allows for bus mastering. An example of this would be a bus master IDE controller that could transfer data from a mass storage device directly to main memory with minimal microprocessor intervention.
- It is extensible to 64 bits using the same connector.
- It has a smaller form factor than VL (VL uses 80 pins and also requires use of the ISA bus slot; PCI requires a single 50-pin connector).
- It has much higher bandwidth at 132 MB/sec. This increases to 264 MB/sec at 64 bits and 528 MB/sec with the coming 66-MHz revision. Such high bandwidth is essential in multimedia and multitasking environments.
- It has broad industry backing.

The PCI local bus is rapidly gathering momentum, driven by the industry switch from the 486 to the Pentium microprocessor. The PCI specification was spearheaded by Intel, and it is Intel that is leading the transition from the VL bus to PCI. Intel is the No. 1 supplier of Pentiums, PC core logic supporting the Pentium, and Pentium-based PC motherboards. It should be no surprise that a company in such a position can cause such a rapid transition from one standard to another. Our analysis indicates that the PCI local bus that had only a 4 percent market share in 1993 will become the dominant local bus in 1996, achieving shipments of 57 million units, which translates into market share of 81 percent (see Table 3-1 and Figure 3-3).

Figure 3-4
PCI Bus Implementation

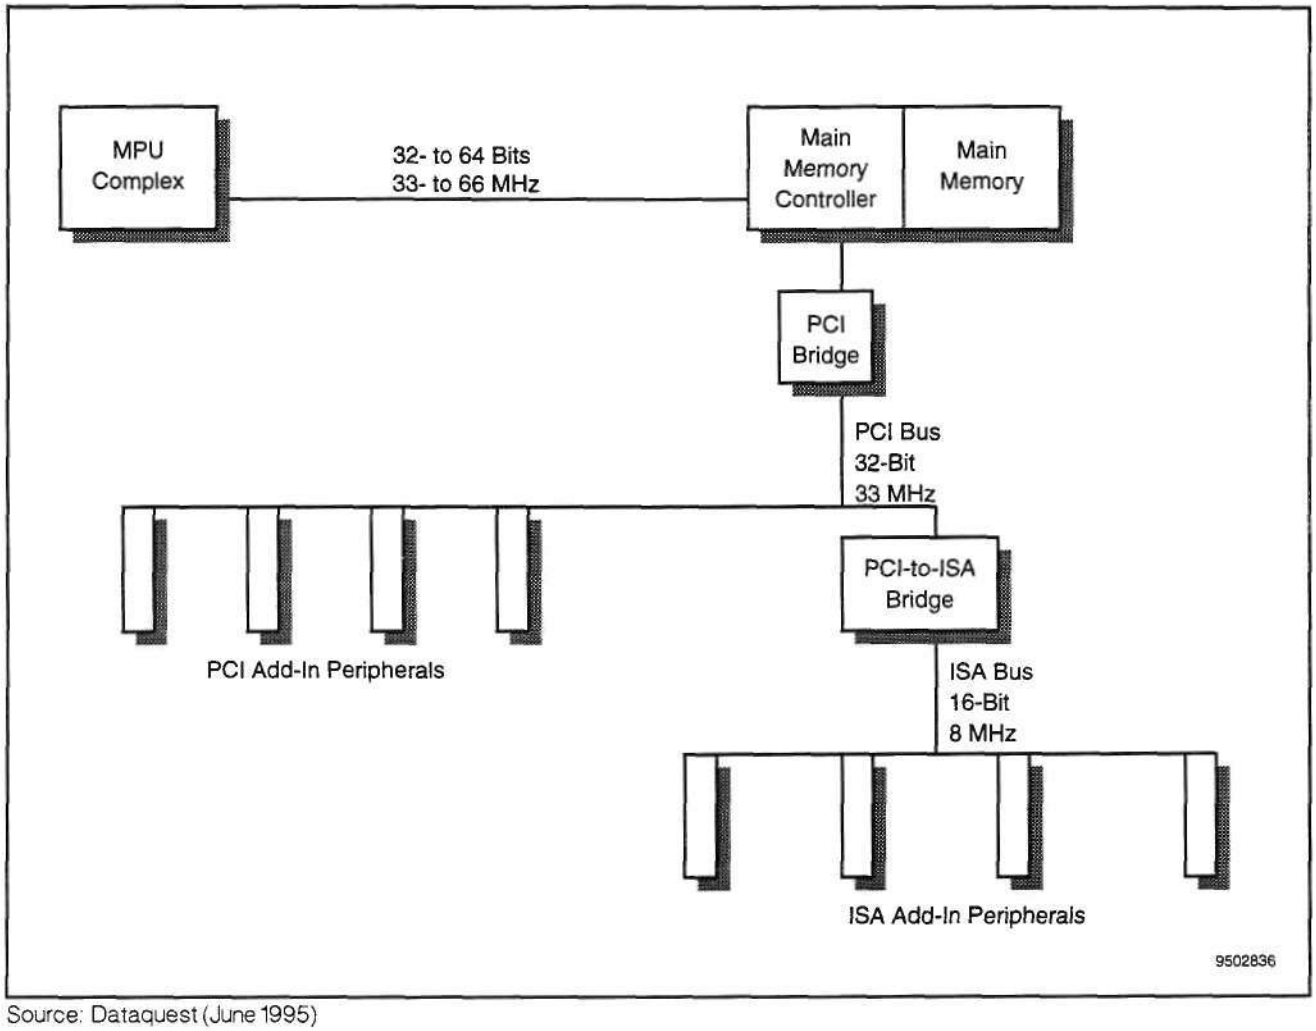

Source: Dataquest (June 1995)

NSP Takes a Ride on PCI

Another initiative being spearheaded by Intel is the Native Signal Processing (NSP) initiative. When NSP was first introduced at COMDEX/Fall in 1994, the message received by the world was that NSP would replace the function of multiple specialty chips that were trying to find a place in the PC. Since that first introduction, Intel has gone to great pains to tell the world that the message had been misinterpreted. At the WINHEC conference held in San Francisco in 1995, Intel began its re-education process in earnest. The message it is now clearly telling us is that NSP is an enabling platform for multimedia. It is trying to get the industry to agree upon a standard for implementing multimedia in the PC. Today there is a hodgepodge, with every card manufacturer developing its own unique feature set to implement some portion of multimedia in the PC. This approach often is wasteful because it can duplicate circuitry that already exists in the PC and quite often it is difficult to integrate these functions into today's PC. The NSP approach provides a baseline target or standard at which everyone can take aim with their particular solution. This should allow companies to focus their efforts on specific areas of competence and expertise and have the assurance that, if they follow the standard, their product will operate in an NSP-ready system.

Central to the operation of a multimedia PC is the need to transfer lots of data at very high rates to and from the PC, totally transparent to the user of the PC. For example, while a user is watching a training video on how to use a spreadsheet and simultaneously operating the spreadsheet and printing a document, the user should be able to receive a fax or other communication completely in the background without interrupting or noticeably impacting the session in progress. Such an example requires bandwidth and parallel operation not possible with either the ISA or VL bus architectures.

Power Management and the "Green PC"

Power management began with the need to extend the useful life of a battery in notebook computers. Notebooks could operate for only 1 to 2 hours, and this was viewed as unacceptable by the purchasing public and was a key element that slowed the acceptance of notebooks into the market. Battery technology moves at a relatively slow pace, so it became the task of the IC community to come up with a solution. The first and most obvious solution was to switch everything from 5V to 3.3V. This task was not as straightforward as at first thought, and the power savings, while significant, was not adequate. The task of finding a long-term solution was left primarily on the shoulders of the PC core logic vendors.

After a government study of the power consumed by personal electronics revealed that PCs were fast becoming the largest consumer of the world's electrical power plants, the U.S. Department of Energy developed a standard known as Energy Star. This environmentally safe standard soon was given the unofficial moniker of the "Green PC." The need for power management had now grown beyond the need of portable computers and was rapidly becoming a required feature in desktops as well.

The first versions of power management simply monitored which peripherals were not being used and shut them down. Although this technique does provide some benefit, it also created some problems. Some of the peripherals, once shut down, took time to power back up. A notable case is the hard disk drive, which can take up to tens of seconds to get back up to speed.

Later attempts involved a finer layer of granularity of shutdown. These varied from standby, to doze, to sleep, to off. This flexibility gave a better trade-off of power for performance, but the power savings was still not of breakthrough proportions.

More recently we have seen more sophisticated power management solutions that take a more proactive stance regarding power management. Foremost of these is the Cirrus Logic subsidiary Pico Power, which uses an active power management approach. Basically the PC core logic starts out in a standby or low power state and only allocates power as the need arises. This is the reverse of most power management approaches that start out at full power and only power down components or peripherals if they remain dormant for a set period. This active power management approach, although more complex to implement, in theory should be more efficient and effective than the traditional passive approach to power management. The passive approach is waiting for a reason to reduce power, while the active approach is waiting for a reason to apply power.

A new feature in PCs that is a direct fallout of the power management issue is referred to as "Instant On." This feature will allow a PC to be turned on and left on. In theory, a PC can be placed in a deep sleep, consuming less than 10W, and then be awakened at a keystroke or telephone ring. This would eliminate the need to physically turn a machine off and then reboot when the PC was needed at some later time. Of course, this will not yet completely eliminate the need to reboot. We still need to eliminate software crashes, and such hardware issues as internal modems that cannot be initialized will still require the system to be powered on and off. But once these bugs are fixed, Instant On will virtually eliminate the need to have easy access to the on/off power switch.

The whole area of power management is one that will be undergoing constant evolution and will involve close cooperation between IC companies and software developers, as well as peripherals vendors. Power management will continue to be a major challenge through the end of the decade and will be a key selling feature in all notebook computers and most desktop PCs.

Chapter 4
Time to Volume

Ten years ago, an x86 microprocessor took several years to reach volume production after the product was released. Today, the time to high-volume production (a 10 million-unit-per-year run rate) takes less than two years; it will take less than 18 months for the P6, and less than that for the P7 (see Figure 4-1).

Too Much Paranoia Can Hurt

As the time to volume for new microprocessors continues to accelerate, the pressure on the PC core logic vendors increases. When major microprocessor vendors announce a new generation of processor, the PC core logic vendors must race to develop new chipsets to support this new architecture. If they are late, they risk losing the entire market window. A prime example of this involves Intel's recent announcement of the P6 microprocessor. The P6 was first discussed with the public in February 1995. Conversations with Intel indicate that third-party suppliers of PC core logic did not receive their first look at the new architecture and the specifications they will need to support it until about the February time frame. Because the P6 uses a new bus, the PC core logic vendors must do a lot of ground-up engineering to support the product.

Figure 4-1
x86 Product Life Cycles

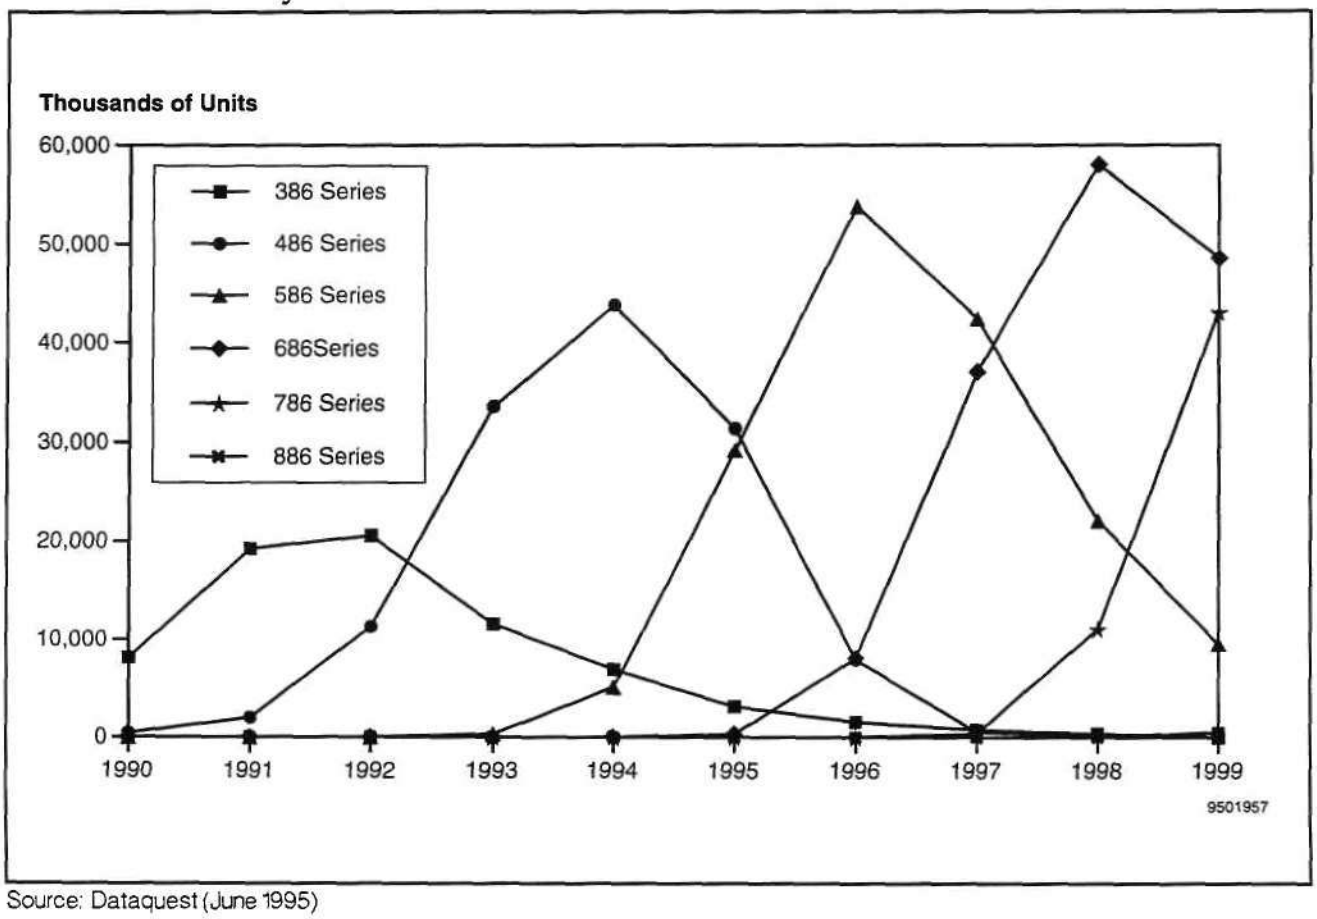

Source: Dataquest (June 1995)

To further place pressure on the third-party suppliers, Intel's own PC core logic group had seen the new bus requirements long before the public disclosure and had been given a tremendous jump-start in the race for market share. Also, with Intel manufacturing more than 10 million motherboards per year, the third-party PC core logic suppliers will have to develop a technologically superior product to the Intel chipset if they are to have a chance of capturing some of the high-volume business from Intel's motherboard business. If they develop a comparable product, they will probably be forced to eliminate 10 million units from their served available market.

Relying solely upon internal talent to develop such a critical set of components as the PC core logic chipset is not the most efficient way of getting the best product to market in a timely fashion. It may in fact delay the time to volume of a microprocessor if the internally developed PC core logic chipset design contains some undiscovered bugs or is missing a key feature that was overlooked or, for some unforeseen reason, is deemed unacceptable by the PC industry OEMs at large. Any of the aforementioned occurrences may slow the acceptance of high-margin microprocessors into the market.

A Better Solution

Intel and other microprocessor vendors must find a way to get third-party chipset suppliers into the architecture loop much sooner than happened on the P6. This may involve creative programs with third-party vendors that guarantee the safety of the microprocessor vendor's intellectual property. But, over the long run, this will be beneficial to all parties. The third-party PC core logic suppliers will have more opportunity to develop a good solution in a timely manner and the microprocessor vendor will have more, and possibly better, PC core logic chipset support when the microprocessor is ready to ship, thus eliminating a potential bottleneck in the critical "time to volume" issue.

Chapter 5
Whom to Support?

Ten years ago, the PC industry needed only to deal with an x86 microprocessor running DOS. Today, multiple processor architectures are vying for the desktop based upon new cross-architecture operating systems:

- x86
  - Intel
  - AMD
  - Cyrix
  - Nexgen
  - Texas Instruments
  - UMC
- RISC
  - IBM
  - Motorola
  - Digital Equipment
  - Sun Microsystems
  - Hewlett-Packard
  - MIPS

Today there is only one given for PC core logic chipset vendors: support Intel's implementation of the x86 architecture. Up through today, supporting Intel also meant supporting AMD, and this will probably continue down the road. Because AMD has been cleared to use Intel intellectual property, one could presume it will continue to capitalize on the installed infrastructure that supports the Intel architecture. It probably will benefit AMD to chart its own course if it wants to become truly positioned as an innovator in the microprocessor market, although this is further down the road. From this point on decisions on whom to support have to be well-thought-out. The others in the x86 camp are vying for the No. 3 slot. Of these, Cyrix and Nexgen are forging off into 586-class product, while Texas Instruments and UMC seem content to remain low-cost suppliers of 486 technology.

Nexgen is the only alternative source to Intel for a 586-class product. Its product uses a unique architecture to implement the 586 instruction set and requires a unique chipset to support it in a PC. Nexgen and VLSI technology have reached an agreement whereby VLSI will produce and market a chipset being designed primarily by Nexgen. Combining this third-party support with the announcement by Compaq to manufacture a computer based on the Nexgen product gives Nexgen a considerable boost in status. This boost may be enough to cause other first- and second-tier PC OEMs to consider using the Nexgen product. Additional third-party support for the Nexgen product has not yet been announced but may be a way for third-party PC core logic vendors to leverage themselves into Compaq.

Mixed Signals

Cyrix and IBM seem to be sending different signals on the production status of Cyrix's 586-class product, known as the M1. When Cyrix first introduced the product, it was more of a proof of concept than a real attempt to develop a manufacturable microprocessor. It had a prohibitively large die size of about 400 sq. mm. Cyrix is performing an optical shrink to get the die size down and expects to ship this product in the third quarter of 1995. IBM, on the other hand, is waiting for a complete redesign and relayout based upon its 0.5-micron CMOS process before taking the product to market. This is not expected to occur until the first quarter of 1996. On the compatibility front, the Cyrix product does target the Pentium pinout but, because of intellectual property reasons, it is forced to use a different data bursting scheme than does the Pentium. This unique bursting scheme will require specific support from PC core logic vendors. Assuming the M1 does not reach high-volume production status until the redesigned product is available in the first quarter of 1996, the 586 class will still have two years of good volume. And, as Intel begins to exit the 586 market in late 1996 or early 1997, opportunity exists for Cyrix and the other 586 vendors to ship millions of units.

The Super 486

In the shorter term, it appears that the M1-SC from Cyrix will have a more immediate impact on the PC market. This product is not a superscalar microprocessor like its big brother M1. The M1-SC, or scalar M1, is meant to provide 75-MHz Pentium performance in a 486DX4 socket. This product, and some of the higher-performance 486 products that operate in the 100-MHz and up range, will provide very attractive price/performance points that may compel some notebook OEMs to perform an upgrade of their existing 486 notebook products.

The Grand Experiment

At the low end is the grand experiment: the attempt to tap into the blue collar homes and greatly expand the PC base. For years we have been subjected to various versions of the pyramid shown in Figure 5-1.

The top of the pyramid represents the early adopters and power users, the middle represents the mainstream, and the bottom layer represents the entry level. The width of each segment is meant to indicate the relative size of each market. There has been much discussion recently about what a lower price point with the right feature set might mean in terms of volume. The basic thought, as represented in Figure 5-2, is that a multimedia-capable machine that is as easy to install as a home appliance with the complexity of a toaster, and sells for $795 to $1,000, will be that solution.

Maybe, just maybe, that time is approaching. However, there is an extreme difference between a $795 price point and a $1,000 price point in the mass merchant channel. If the industry is ready to make this grand experiment, it had better be prepared to offer this new box at the right price point. To that end, we are beginning to see some semiconductor manufacturers address this issue. Texas Instruments, which knows a thing or two about the mass merchant electronics market, is touting a 486SX2-66 for $33, and Pico Power has introduced a PC core logic chipset targeted to provide this market the multimedia features it needs while simultaneously facilitating a lower cost point. This will be an interesting

Figure 5-1
Traditional PC Installed Base Pyramid

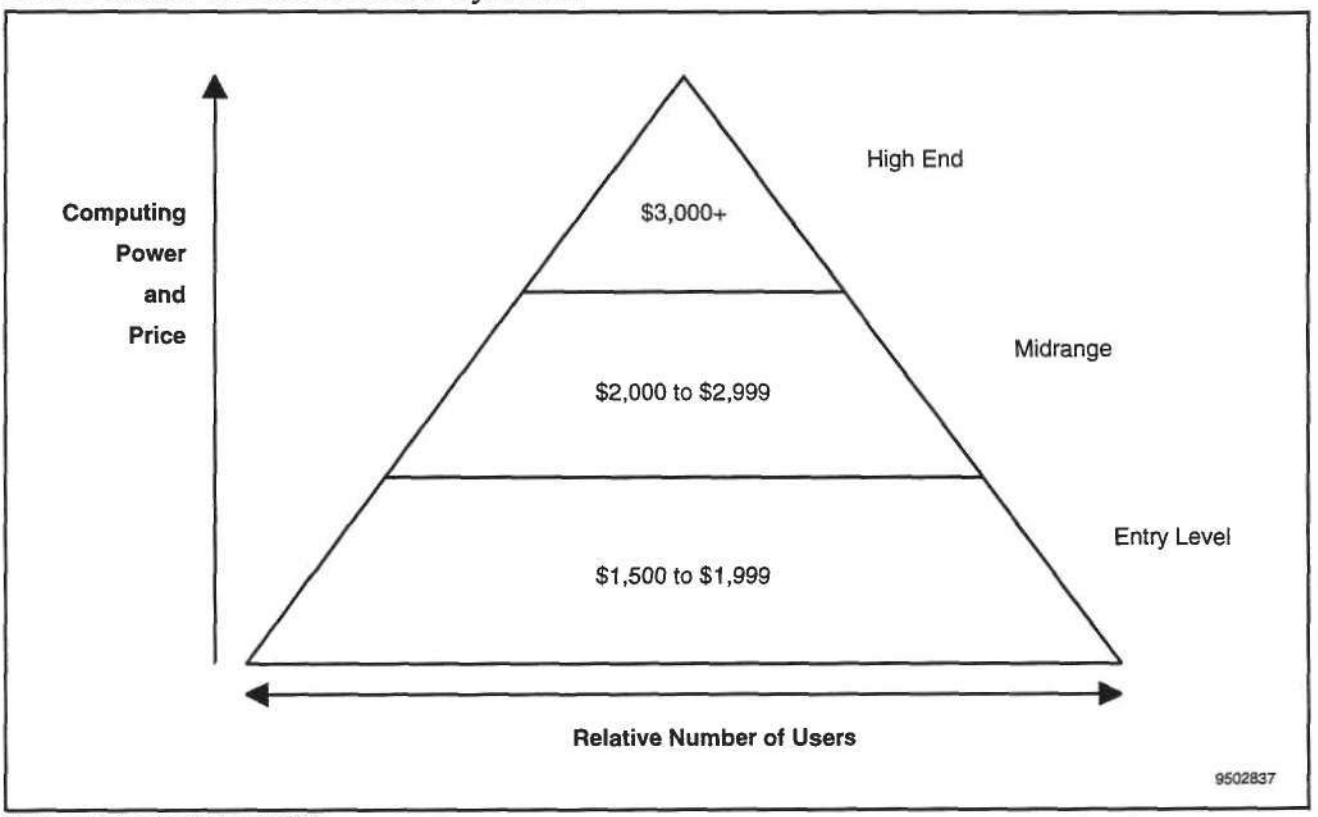

Source: Dataquest (June 1995)

experiment and, if well-thought-out, could greatly expand the total available market for PCs.

UMC's position in the 486 market seems a little tenuous, as UMC's own 486 advertisements in the international airport in Taipei, Taiwan, state clearly that its 486 is "not for export."

To RISC or Not to RISC

The vendors of the various RISC microprocessors are as fragmented in their individual approaches to gaining market share as is the open UNIX platform. (The "open UNIX platform" has at least as many versions as there are computers supporting the "standard".) Sun Microsystems' approach with the SPARC architecture appears to be to push a UNIX box into the mainstream PC desktop with "lower" prices. Hewlett-Packard is making no overt rush to the desktop with its PA-RISC architecture, but the relationship with Intel has not yet borne its fruit. MIPS et al. (referring to the various silicon partners) is still counting on Windows NT to carry it into the desktop. Digital Equipment is following a similar strategy with the Alpha microprocessor. Digital's difference is that it seems to have an architecture built for speed and has consistently demonstrated that it can lead the market in this area. We are not suggesting that the other RISC architectures are not high-performance, but recent history suggests that the Alpha is just faster. Over time, this speed advantage may come into play as a cross-platform operating system such as Windows NT or a derivative becomes mainstream.

Figure 5-2
PC Installed Base Pyramid with New Entry Point

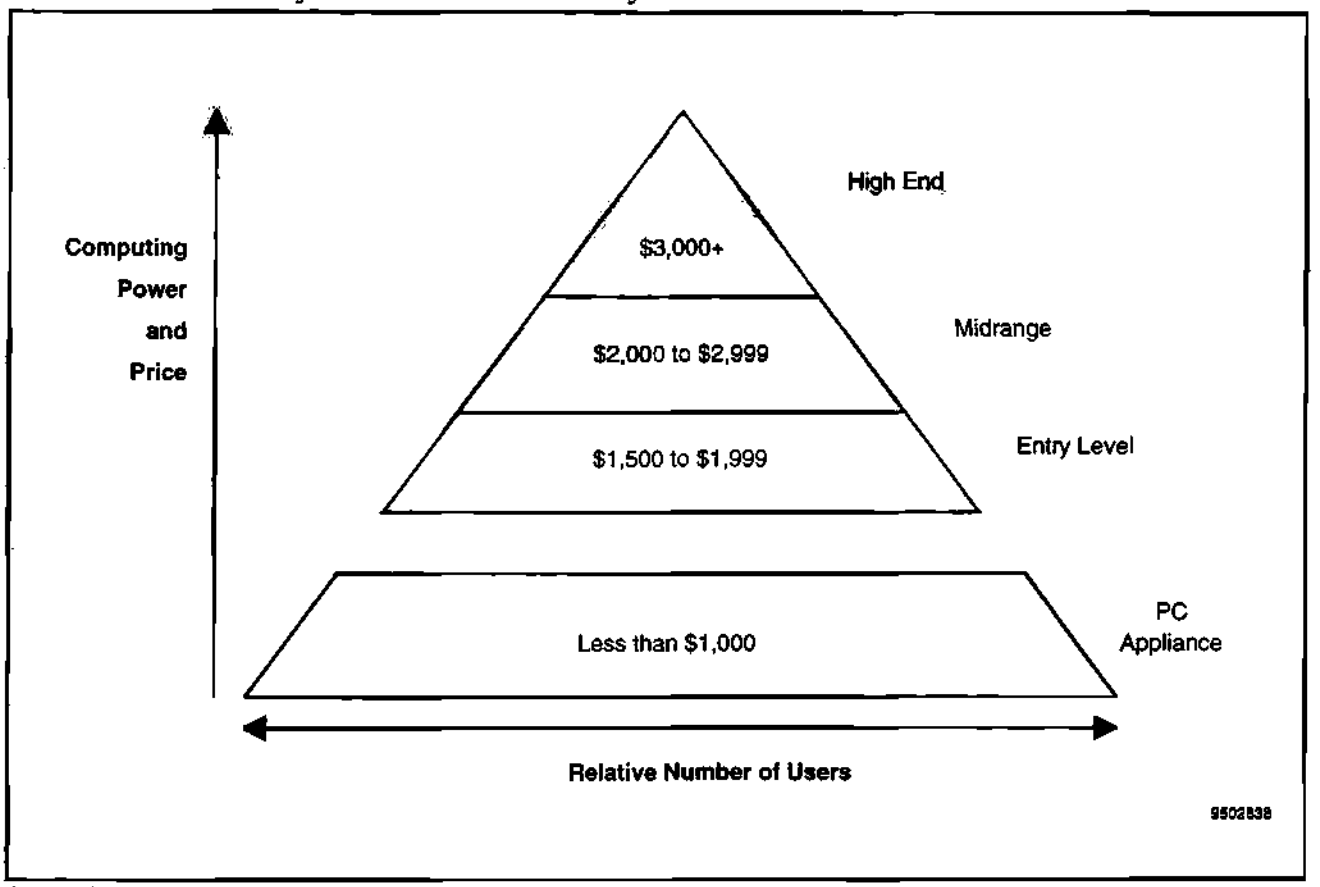

Source: Dataquest (June 1995)

Last but not least in the RISC processor camp is the PowerPC. The alliance of IBM, Apple, and Motorola cannot be ignored. Each of these companies brings formidable skills to the alliance. To name but a few: IBM's technology, Apple's experience and knowledge in developing user-friendly systems, and Motorola's highly reliable manufacturing muscle. Each of these companies also has other strengths, but we believe the point is proven that this alliance is potentially formidable. Unfortunately there exist some sizable negatives that have slowed the acceptance of the PowerPC into the market. Although Apple is making a massive shift to the PowerPC, it remains the only company to use the PowerPC in volume. The attempts to make the Macintosh an open standard have yet to yield any tangible results. Further confusing the issue is the operating system focus, or what appears to be a lack of focus. At this point it would seem that the only real chance for the PowerPC to compete head-on with the x86 architecture for a dominant position in the PC market is to attempt to expedite the acceptance of Windows NT on the desktop.

Windows NT's current lack of support for legacy 16-bit software and its larger system resource requirements has relegated it to file servers and other high-end applications. The upcoming release of Windows 95 will also serve to delay the need to shift to an NT-like operating system. As microprocessors continue on their high-performance migration path and

as Windows 95 is perceived to be running out of steam, the industry will finally be ready to move to an NT-like operating system. Unfortunately for the RISC vendors, such an event is not likely in the short term and will probably not occur until the end of the decade. At that point it is quite likely that the x86 vendors will be at, or very near, the performance of mainstream RISC microprocessors (see Figure 5-3). There is little argument that, with a given transistor budget, a RISC architecture will give more performance than a CISC architecture. However, if the CISC architecture is able to pay a transistor penalty, there is no reason why a CISC architecture cannot match the performance of a mainstream RISC architecture.

In Figure 5-3, the RISC and CISC (Old) lines represent historical views of the performance paths of the two basic architectural types. Of course, this assumption assumed that architectural techniques such as branch prediction, speculative execution, out-of-order execution, superpipelining, and superscalar architectures are the domain of RISC microprocessors only. However, it has recently been demonstrated by a number of people that these same techniques also can be applied to CISC microprocessors. In fact, the CISC microprocessors began to deviate from the traditional performance forecast with the Pentium that included two execution units. Announcements by AMD, Cyrix, Intel, and Nexgen indicate that the CISC vendors have every intention of making use of these same time-proven architectural techniques. The line labeled CISC (New) works under the assumption that the CISC suppliers can afford the transistor penalty to attain performance parity with their RISC competitors and that this will occur by the end of the decade.

Figure 5-3
SPECint92 Performance Forecast

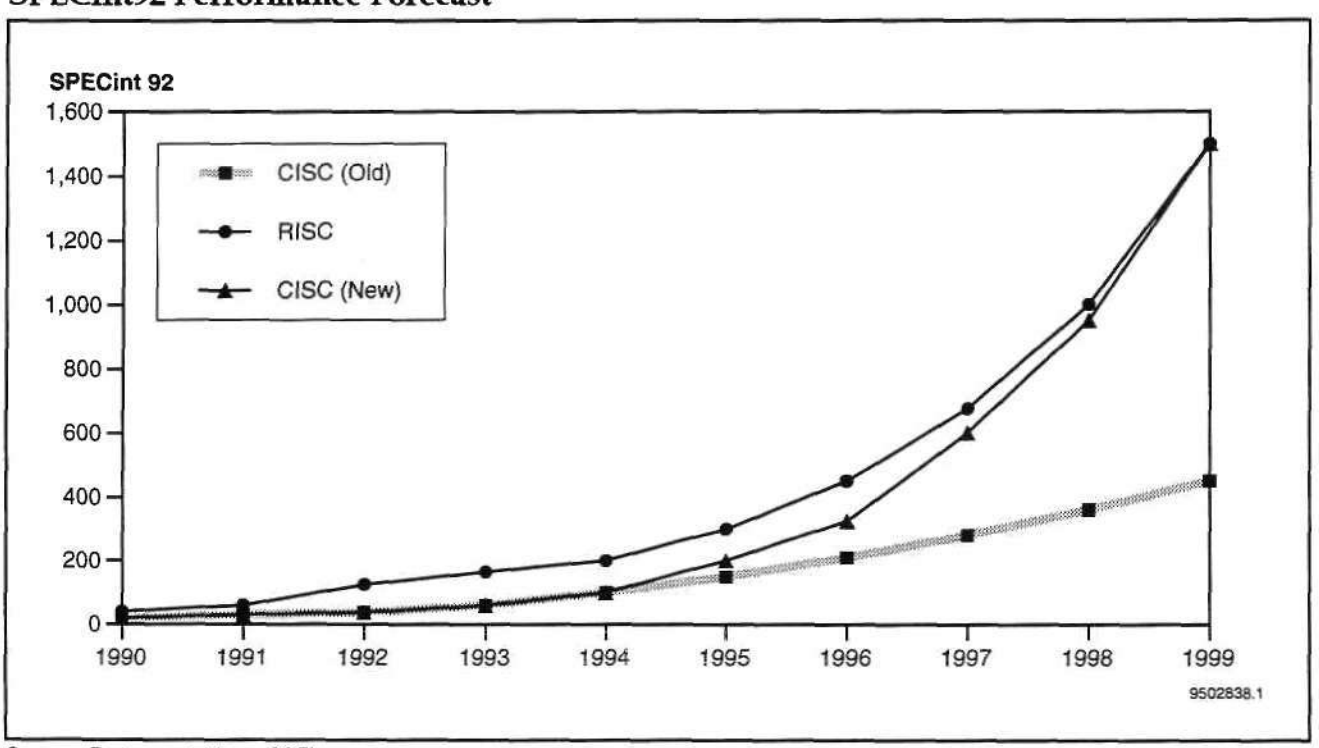

Source: Dataquest (June 1995)

If the x86 industry infrastructure is still in place when the move to an NT-like operating system occurs, if the x86 microprocessors have reached performance parity with mainstream RISC, if the various RISC vendors are unable to match the x86 industry's ability to ramp to volume (see Figure 3-1), what will motivate the requisite PC support infrastructure to develop the ability to support another microprocessor architecture (let alone several)? There are a lot of "ifs" in this scenario, but none of them is beyond the realm of reason. It will be an interesting story as it unfolds. Dataquest will be there to advise on its progress.

Chapter 6
Dataquest Perspective

A trend seems to be developing in the PC core logic market that is not good for the PC industry at large. Intel appears to be relying more heavily upon its own resources to develop and manufacture PC core logic chipsets. Now the world has been shown that Intel has no peer at developing high-performance microprocessors in high volume. Intel has reaped substantial financial rewards for developing this capability and for properly maintaining focus on this objective. However, is it possible that it is beginning to defocus its attention just a bit? Intel has a history of developing new technology and standards to enable the PC industry to build better, faster, and cheaper computers. The leadership provided by the Intel Architecture Labs has provided the entire industry. Of course, it is no coincidence that this direction also has been financially beneficial to Intel.

As of late, however, Intel's positioning in the PC core logic arena has taken on more of a role of a company trying to gain dominant market share rather than a company trying to lead the industry in a specific direction. Maintaining a PC core logic development effort certainly seems to be a critical program for Intel. One need not look back very far in history to see the importance of this program for Intel. During the ISA, MCA, and EISA bus wars, Intel was one of the few companies that developed an EISA chipset that was critical in selling high-end x86 microprocessors into server applications. Although this market never developed into significant volume, Intel's leadership caused other core logic vendors to develop EISA bus PC core logic. When the third-party PC core logic vendors were busy making money selling core logic based on the VESA local bus, Intel developed its own support for the PCI bus that it believed to be a better bus for the future. As the PCI standard solidified, the advantages of PCI over the VESA local bus were compelling. Now that the industry is switching to new motherboards based on the Pentium microprocessor, it is a natural time to switch from the VESA local bus to PCI. This general support of PCI is directly benefiting Intel because the PCI bus allows the power of the Pentium and next-generation x86 microprocessors to be more efficiently used.

The third-party PC core logic vendors are positioned with some very solid products to compete in this market that implement the PCI function and bring into play additional value-add that they have developed, such as a wider variety of main memory support, higher levels of integration, and better power management techniques, to mention but a few. However, Intel seems to be in an "I can do it better" mentality rather than a "here's how I would do it" mentality. The latter approach builds on the various levels of expertise developed by the third-party PC core logic suppliers and results in a product that implements the concepts Intel would like to see as well as the new innovations and value-add that the third-party suppliers can supply. The former approach seems a little risky and uncharacteristically shortsighted for Intel.

Today's leaders in developing PC core logic chipsets have come a long way in developing the requisite skills to support this critical portion of the PC market. Evidence of their growing skills is the fact that major PC OEMs

such as Compaq and IBM have almost entirely retreated from the development of custom solutions and are now purchasing standard products for the bulk of their requirements. Only a few systems still rely upon custom PC core logic, and we expect these to disappear in short order.

Given the tremendous strides that the PC core logic community has made in the development and design of PC core logic chipsets, it seems quite clear from a third-party observer's perspective that Intel would be well-served to bring the PC core logic vendors into the disclosure loop of the chipset requirements of a next-generation microprocessor much sooner than occurred on the P6. The PC core logic industry only received the necessary specifications to begin such a design nearly simultaneous with the announcement of the product itself. This allows just a few short months to develop P6-capable product. Such a time constraint serves no one's interests because many vendors will be fortunate to develop a bare-bones chipset in time, let alone develop one that will help Intel's ambitious goal of shipping several hundred thousand P6 microprocessors in 1995 and several million in 1996 (8 million, according to Dataquest's estimates).

The industry is growing at a rapid rate, and the time-to-volume pressures Intel has placed on itself do not allow for any stumbles along the way. Intel should be wary of becoming too enamored with its own development capabilities. The company would be better served enabling the industry rather than relying solely on its own resources to develop the support products for its highly lucrative microprocessor line. Such an enablement would better ensure that the right products were available at the right time and at the right price in order for Intel to achieve its ambitious goals of getting newer generations of microprocessors to market faster than prior generations.

For More Information...

Jerry J. Banks, Director/Principal Analyst ..... (408) 437-8677
Internet address ..... jbanks@dataquest.com
Via fax ..... (408) 437-0292

Dataquest
a company of
The Dun & Bradstreet Corporation

Dataquest

a company of
The Dun & Bradstreet Corporation

Corporate Headquarters
Dataquest Incorporated
1290 Ridder Park Drive
San Jose, California 95131-2398
United States
Phone: 1-408-437-8000
Facsimile: 1-408-437-0292

Dataquest Incorporated
Nine Technology Drive
P.O. Box 5093
Westborough, Massachusetts 01581-5093
United States
Phone: 1-508-871-5555
Facsimile: 1-508-871-6180

Dataquest Global Events
3990 Westerly Place, Suite 100
Newport Beach, California 92660
United States
Phone: 1-714-476-9117
Facsimile: 1-714-476-9969

European Headquarters
Dataquest Europe Limited
Holmers Farm Way
High Wycombe, Buckinghamshire
HP12 4XH
United Kingdom
Phone: 44-1494-422722
Facsimile: 44-1494-422742

Dataquest GmbH
Kronstadter Strasse 9
81677 München
Germany
Phone: 49-89-930-9090
Facsimile: 49-89-930-3277

Dataquest Europe SA
Immeuble Défense Bergères
345, avenue Georges Clémenceau
TSA 40002
92882 - Nanterre CTC Cedex 9
France
Phone: 33-1-41-35-13-00
Facsimile: 33-1-41-35-13-13

Japan Headquarters
Dataquest Japan K.K.
Shinkawa Sanko Building, 6th Floor
1-3-17, Shinkawa, Chuo-ku
Tokyo 104
Japan
Phone: 81-3-5566-0411
Facsimile: 81-3-5566-0425

Asia/Pacific Headquarters
7/F China Underwriters Centre
88 Gloucester Road
Wan Chai
Hong Kong
Phone: 852-2824-6168
Facsimile: 852-2824-6138

Dataquest Korea
Suite 2407, Trade Tower
159 Samsung-dong, Kangnam-gu
Seoul 135-729
Korea
Phone: 822-551-1331
Facsimile: 822-551-1330

Dataquest Taiwan
11F-2, No. 188, Section 5
Nan King East Road
Taipei, 105
Taiwan, R.O.C.
Phone: 8862-756-0389
Facsimile: 8862-756-2663

Other research offices in Beijing, Singapore, and Thailand

Sales agents in Australia, Belgium, Germany (Kronberg), Israel, Italy, South Africa, and Spain

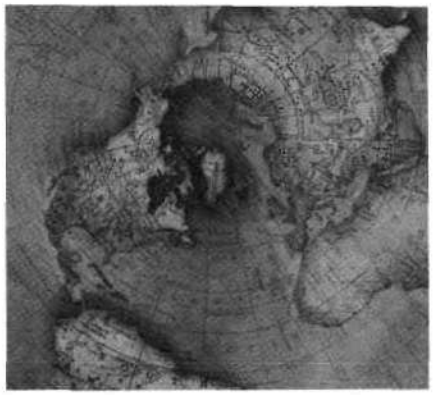

Dataquest

The Intel P6 Processor: A Summary for System Manufacturers

Focus Report

Program: Next-Generation PCs
Product Code: NGPC-WW-FR-9501
Publication Date: April 10, 1995
Filing: Focus Studies

The Intel P6 Processor: A Summary for System Manufacturers

Focus Report

Program: Next-Generation PCs
Product Code: NGPC-WW-FR-9501
Publication Date: April 10, 1995
Filing: Focus Studies

Table of Contents

Table of Contents (Continued)

Table of Contents (Continued)

List of Figures

List of Tables

Foreword

The 286, 386, 486; the P5, P6, and P7. Single-pipeline, dual-pipeline, super-pipelined cores, dispatching one IPC, two IPC, three—cha cha cha.

It is tempting to look at various x86-architecture processors as a logical progression, as though each new generation were merely a simple evolutionary enhancement over its forebears. Each new part may seem destined to follow a similar life cycle as the rest and be used in similar products.

In the case of the Intel P6, this assumption would be wrong. The device is clearly the logical successor to the 386, 486, and Pentium CPUs, and it will indeed further leverage the huge installed base of "legacy" x86 software, but I believe the part is much more than that. The P6 design has already established new paradigms within such disparate fields as design-team deployment, chip-level partitioning, architectural philosophy, and system connectivity. It represents a host of "firsts" for Intel:

- The P6 is the first processor in the x86 family to be designed at Intel's Hillsboro, Oregon, facility, away from the Santa Clara, California, headquarters. By effectively pipelining its design teams, Intel can now introduce major new product generations more frequently and ride more closely along the optimum technology curve.
- The P6 is the first microprocessor from any vendor to combine separate CPU and cache chips within a single common package. By splitting the design into two die, the P6 has broken Moore's Law, letting Intel fabricate the part for which years before it would have been practically using a conventional monolithic design.
- The P6 is the first microprocessor of any type to graft a sequential instruction-set architecture onto an internal data-flow engine. By rescheduling instruction sequences as they execute, the P6 is able to extract far greater parallelism from an instruction stream than earlier designs, delivering near-optimum performance from existing, unrecompiled software.
- The P6 is the first x86 processor to depart significantly from the basic system interface of the original 8086. By defining bus transactions at a higher level and decoupling transfer requests from their completion, the P6 makes far more efficient use of its local system bus, allowing four-way multiprocessing with zero added cost.

These factors each play a key role in defining P6 capabilities and potential. Together they will likely cause the part to be adopted more quickly and in higher volume than any x86 processor introduced to date. I believe the P6 will soon find its way into a broad spectrum of system classes, from handheld portables to multiple-processor mainframes. Product managers, system designers, and software developers that fail to consider the impact of the P6 on their corporate strategies will do so, I fear, at their own risk.

The significance of the P6 clearly cuts across many disciplines. Within this report, Dataquest does an exceptionally thorough and insightful job of detailing the P6 capabilities, discusses how its design will change the system engineering process, and describes the net results from a system

user's perspective. I commend Dataquest for having assembled a report with remarkable breadth and depth, which you will undoubtedly find as interesting to read as it is useful in forming company plans.

John H. Wharton
Palo Alto, California

(John Wharton is an independent consultant specializing in computer architectures and a contributing editor to Microprocessor Report. He is the editor of The Complete x86—the definitive guide to 386, 486, and Pentium-class microprocessors, available from its publisher MicroDesign Resources.)

Chapter 1
Overview

Background

Intel disclosed key features of the P6 architecture at the ISSCC conference on February 16.

This disclosure does not constitute a product announcement; that will happen at an unspecified future date that we believe to be in the third quarter of this year.

This report provides an in-depth understanding of the P6, covering its architecture, its implementation, the unique features associated with the two-chip design, the likely availability and pricing of the part, and our projections as to its effect on the various components of the system market. It also covers some key system implementation issues.

DQ Hotlist

- This is a real part; we have seen it running Windows—fast. It is not a paper tiger. The P6 is designed for speed; Intel may be understating the initial performance claims, practically a first. We would not be surprised to see 267-MHz parts in the second half of 1996.
- Intel is serious about volumes and performance on this product; system manufacturers need to develop implementation plans for all classes of system immediately. This is a P54C type ramp; be prepared.
- Intel is keeping its rate of progress of design and manufacturing technology high to keep competitors (AMD, Cyrix, NexGen, and so on) at bay; the P6 is born of that competition, and the users win.
- Notebooks will be in the market sooner than most people currently think.
- P6 system design is hard, a consequence of the 66-MHz enhanced Gunning Transceiver Logic (GTL+) bus and split-transaction protocol. Intel will be the sole supplier of core logic for some time: R&D dollars must be spent very carefully to develop a return when competing with Intel in this area, a caveat that applies to chipset and system vendors.
- Few vendors will be able to develop independent P6 technology: For many, Intel motherboards will be the only choice. Intel will grows its share—and its control—of this major segment of the industry, bringing Intel one step closer to the system business.
- The integrated cache will set a trend for high-performance processors and will reduce the opportunity and need for external cache devices. Note that this approach permits a semiconductor manufacturer to break with Moore's law.
- The second-tier RISC vendors could find themselves under threat from the P6, especially if it achieves the level of performance that we believe it is capable of; conversely, it is an opportunity for x86 system vendors to eat into the lucrative workstation market.

- Workstation vendors should evaluate the P6. At a minimum, workstation designs must leverage low-cost PC features including PCI and PC video accelerators.
- The CMOV (conditional move) instruction is a hidden jewel; software developers should start using it as much as possible, and Pentium cloners should look at building it into their designs for a future iteration.
- The opportunity exists for clever coders to take advantage of the architecture to implement algorithms for Motion Picture Experts Group (MPEG) compression and other looping software very efficiently; this characteristic may well bring MPEG to the desktop as a standard in 1997 or 1998.
- Cache memory manufacturers need to plan for the migration of cache memory from the system board into the processor package in 1997. Vendors that do not prepare for this transition will be surprised by a shrinking market. We do not believe that there is *any* opportunity for external cache in a P6 design.
- P6 will drive the system market to massive consumption of EDO DRAM in 1997. Smart memory vendors, and system vendors, will follow the standards selected by Intel.
- Core logic manufacturers face a stiff challenge in competing with Intel's chipsets. The GTL+ bus and split-transaction protocol complicates design, and differentiation through performance will be hard. There is perhaps an opportunity in notebook chips and the use of Rambus DRAM.
- The P6 could enable speech-to-text technology in late 1996 or 1997, another major market growth engine.

Overview

The P6 is a very advanced design, the technological equal or better of all microprocessors announced to date with the possible exception of the underperforming PowerPC 620. However, the architectural complications imposed by the x86 instruction set still seems to prevent Intel from getting ahead of the RISC vendors in terms of absolute performance, although it is this level of complexity that allows the P6 to run a broad range of existing applications at significantly higher speeds than previously possible.

With this level of technology comes a daunting manufacturing challenge; however, Intel is addressing this challenge with tremendous investments in manufacturing capacity. As the new plants come on line, Intel is in a position to produce massive quantities of the P6. We believe that it will follow the P54C ramp, with hundreds of thousands in 1995, millions in 1996, and tens of millions in 1997; this part follows the 486DX2 or P54C, not the 486/50 or original Pentium.

This report examines the multiple issues surrounding the device, from architecture and design through manufacturing and system implementation issues to its effect on the four major computing markets (desktop/deskside, mobile, server and workstation).

Project Analyst: Martin Reynolds

Chapter 2
What Could Go Wrong?

We believe that the P6 will be as influential on the market as the P54C Pentium, a part that has gone from nearly nothing to a major market force in less than a year. Intel's investment plans and design efforts are such that even this watermark introduction could be bested. This chapter summarizes the factors that could cause the P6 to fail to meet our expectations.

We also identify warning signs; note that Intel may make statements indirectly about any one of these factors as it talks more about expectations of the P6.

Packaging Failure

The dual-cavity pin grid array (PGA) is a tough package to build. Although not as sophisticated as a multichip module, there may prove to be technical problems that limit production capacity or yields from the package suppliers. If this happens, Intel will be forced to switch to a single-die design with a dramatic effect on manufacturing cost, availability, and perhaps performance.

Risk Assessment

We would judge the risk in this area to be very low. Although this is a new approach, it seems that the problems should be surmountable in the near term. Note that, once the problem is solved, it is solved for all potential takers of the dual-cavity technology.

Warning Signs

Warning signs include limited availability of the P6 and rumors from the industry about packaging problems.

Clock Rate Problems

The P6 is designed to go fast, and we have seen working systems. However, it must scale the clock rate rapidly to become a major market force as a 133-MHz P6 is not much more compelling than a 150-MHz Pentium, echoing the 486DX2/66-to-Pentium/60 step. Worse, a 150-MHz K5 (although we judge this to be unlikely until well into 1996) could severely blunt its edge.

Risk Assessment

This is truly the great unknown. Although 133 MHz will be interesting, 166 MHz and then 200 MHz are essential to the dramatic success of the part. We would judge the risk to be less than 50 percent but quite significant.

Warning Signs

Warning signs include rumors from the industry about thermal sensitivity of the part, recommendations from Intel suggesting the use of excessively large cooling fans, and absence of technology demonstrations of faster chips in 1995.

Functional Failure

A processor of this complexity runs the risk of a major functional bug. Such a bug may be caught and fixed, necessitating a product revision cycle, or it may leak out into the field and wreak later havoc along the lines of the Pentium FDIV bug.

Risk Assessment

After the Pentium $475 million education, we suspect that the risk of a functional bug surviving in the part is minimal.

Warning Signs

Product shipment date is out into 1996. Confinement to key markets such as servers where functional bugs are less likely to manifest themselves and are more easily controlled.

Failure of the 0.35-Micron Process

The 0.35-micron process that Intel is planning to produce the P6 on in 1996 is relatively advanced. There are many technical challenges involved in implementing such a process, but other companies have successfully demonstrated 0.35-micron parts for other applications. Intel has been working on the process in Oregon for some time and is implanting a full-size fab for manufacturing this year. Failure on this fab could be very costly. The capital investment for a 0.35-micron process line running in 1995 is well over $1 billion, and the fab must deliver working parts quickly to amortize its capital cost at a reasonable rate per part. Intel is also dependent on this process for its next generation of Pentium parts.

Risk Assessment

We believe that the risk here is low; 0.35-micron technology, although expensive to implement because of the capital cost, does work. Intel's 120-MHz Pentium part, initially delivered in this geometry, is earlier than might be expected from such a new process; this is a strong sign that the process will be extremely successful.

Warning Signs

The 0.35-micron fab will produce both faster Pentium parts and the 512KB L2 cache for the P6. If both of these fail to materialize in 1996, or the Pentium does not ship at an increased clock rate when it appears, there could be trouble ahead.

Failure of the Market—No Desire for Performance

The tremendous success of the PC market has been fueled by a desire or a need on the part of the end-user for greater performance. Many machines today are replacement machines—witness the 386 boxes of the early 1990s stacked in company closets to live out their depreciation cycle. The extra cost of a faster machine is outweighed by the gains in productivity, essential now more than ever in an economy that has turned to automation for leverage of its expensive desk workers.

If this devaluation of performance takes place, replacement purchases of systems will drop dramatically and Intel's capital investment will turn into a liability, forcing Intel to crash the price of the P6. Intel will show

weak financial results and will find a move into the system business appealing. This shift will affect both other processor manufacturers and many system manufacturers.

Risk Assessment

We judge the risk of the market failing in this fashion to be low.

Warning Signs

Pentium growth stops, and 486 processors maintain their hold on the market. No new technology that depends on processor performance introduced as mainstream in any one year (1995 will be on-board MPEG, 1996 telephone application program interface (TAPI), 1997 videoconferencing).

Intel enters the system market.

Failure of the Market—No Volume

Although this slowdown could be driven by lack of interest in performance, the market could be affected by a major economic downturn, war, or other overriding socioeconomic factors.

Risk Assessment

We judge the risk of the market failing in this fashion to be low.

Warning Signs

Warning signs include typical economic indicators and unrest in emerging economies that will support volume growth at the low end of the market.

Failure of Support Chipsets

Support logic is a critical feature for the P6. Intel has not promoted the bus information to system logic manufacturers in sufficient advance of product availability in time to enable them to deliver products for 1995 and much of 1996. In part, this decision protects the P6 pinout against the cloners, but it could stifle the market. On the other hand, Intel has plenty of production capacity and may be able to support the market on its own for the critical first year.

Risk Assessment

We judge the risk of support chipsets failing to be low. Note that the risk is more severe in notebooks; there is no indication of a plan to develop P6 notebook logic from Intel, which could delay introduction of P6 notebooks beyond the optimum point of the fourth quarter of 1996. We note further that chipset manufacturers are usually reluctant to move quickly on a new generation of product; Intel's adoption of the cause certainly indicates that there will be a ready supply of parts.

Warning Signs

Intel fails to deliver Orion samples. Orion is priced too high. Turnkey motherboards from Intel is not available.

Chapter 3
Systems

P6 and the Windows Desktop

- This processor is for desktop machines.
- Systems must be ready as P6 is introduced.
- System manufacturers should start developing their P6 desktop products now.

Although the P6 is targeted at servers initially, the performance advantages of the P6 will be valuable to a small percentage of desktop users. Accordingly, we would expect to see desktop P6 systems available at introduction (no doubt built on Intel motherboards) at the traditional $5,000-to-$10,000 price range that used to grace new Intel-based systems. Again, savvy companies stand to do very well here by being early to market. We believe that the Pentium will continue to lead users to aspire to higher and higher performance (helped, of course, by a torrent of Intel education), and the P6 will enter a performance-hungry market.

We would recommend that system manufacturers work closely with Intel to develop their first P6 desktop products as soon as possible and put aggressive cost-reduction plans in place to compete. The initial high prices will drop rapidly as supply grows, and products with too high a cost structure will be driven out of the market.

Differentiating around the P6 will be very hard, a consequence of the integrated cache and the Intel chipset being the only game in town for at least the first year. The cache is a problem because it is very hard to improve on the performance of a Level 2 cache that, once primed, is just one clock cycle (6ns and shrinking) away from the processor; an external cache, which would require an advanced nonblocking cache controller, would probably have to consist of several megabytes of synchronous RAM to make a difference. Perhaps the only thing that will make a difference is the use of a high-bandwidth DRAM architecture, although Intel advises us that the fast L2 cache tends to mask improvement even from this alternative.

The P6 will also lead to another puckering of the business market in 1997 when it comes onstream in full force. As with the Pentium in 1995, the increased machine cost will suck budget dollars, reducing the total number of machine purchases—in other words, more revenue but less unit opportunity; so some manufacturers have to lose. Therefore, manufacturers should establish their position in the P6 desktop market early. Do not be fooled by all the server noises; this is a desktop processor.

On the other hand, the growth of the market expected in 1996 because of TAPI availability may carry on into 1997. However, we believe that users will save money by purchasing TAPI upgrade cards for existing systems, rather than increasing the tightly managed budgets typical of the 1990s.

We expect PC growth in emerging economies to remain strong and be essentially unaffected by the P6 introduction.

Note that, if the market contracts, Intel will divert production to the P6 to decrease its total unit shipments and to optimize revenue and profit. The price of the processor will drop, and all x86 participants will feel their profits squeezed. We believe that this will not happen in the foreseeable future, but it will be a body blow to the industry and could lead to a damping of the performance cycle that drives the industry today.

P6 and the Notebook Market

- Notebooks may happen sooner than people think (demo in PC EXPO 1996).
- Power, heat problems are manageable.
- Manufacturers should start developing their P6 notebooks now.

The power and packaging constraints that P6 imposes on the system designer, coupled with the need for a totally new set of core logic, would suggest that P6 notebooks are probably some way off in the future. We would estimate that such a product will not appear from a credible manufacturer until 1997 when the smaller geometry P6 is readily available and Intel and system designers have had the opportunity to pull some power-saving tricks.

However, the market for powerful personal computing seems to be all but insatiable, as Figure 3-1 illustrates with the projected penetration of Pentium into the notebook market. Thermal management techniques designed for Pentium will inevitably find their way into the first P6 products. Coupled with clock throttling and closed-loop thermal feedback, these techniques may well enable a P6 notebook to be built early on in the P6 life cycle with the ceramic package. Therefore, we recommend that notebook manufacturers design today with the P6 in mind, and that they press the logic suppliers for a P6 solution as soon as possible. Note that a notebook P6 chipset does not carry the overhead that a desktop system chipset does because it does not need multiprocessing capability.

The alternative view is that Intel could produce a product with integrated cache specifically for the notebook market. This device could be sampling in the middle of 1996, suggesting P6 for COMDEX/fall and systems in volume in 1997. Although the integrated cache part will cost more to build, the notebook market is strategic for Intel, and the company may well elect to invest resources in this strategy.

We have also been questioned about the double-digit power requirements of the P6. The solution to this is aggressive power management. When the user is typing, looking at the screen, or mousing around, very little processor performance is required. The processor can be clocked at a minimal frequency, reducing its power consumption to almost nothing. In this mode, most power goes to the other electronics and supports the backlight. The value of the P6 comes in when the hourglass appears. The full force of the processor will wipe out the hourglass twice as fast as the Pentium, a major advantage to the user. Now, this process also takes twice the power, but as it is half as long, processor energy consumption is a net even. However, the backlight was only "wasted" for half the time. To put it another way, the P6 will let the user get more work done from a single charge because he takes less time watching the hourglass. In the case of the notebook, the hourglass sucks not only productivity but also power.

Figure 3-1
Projected Notebook Shipments by Processor

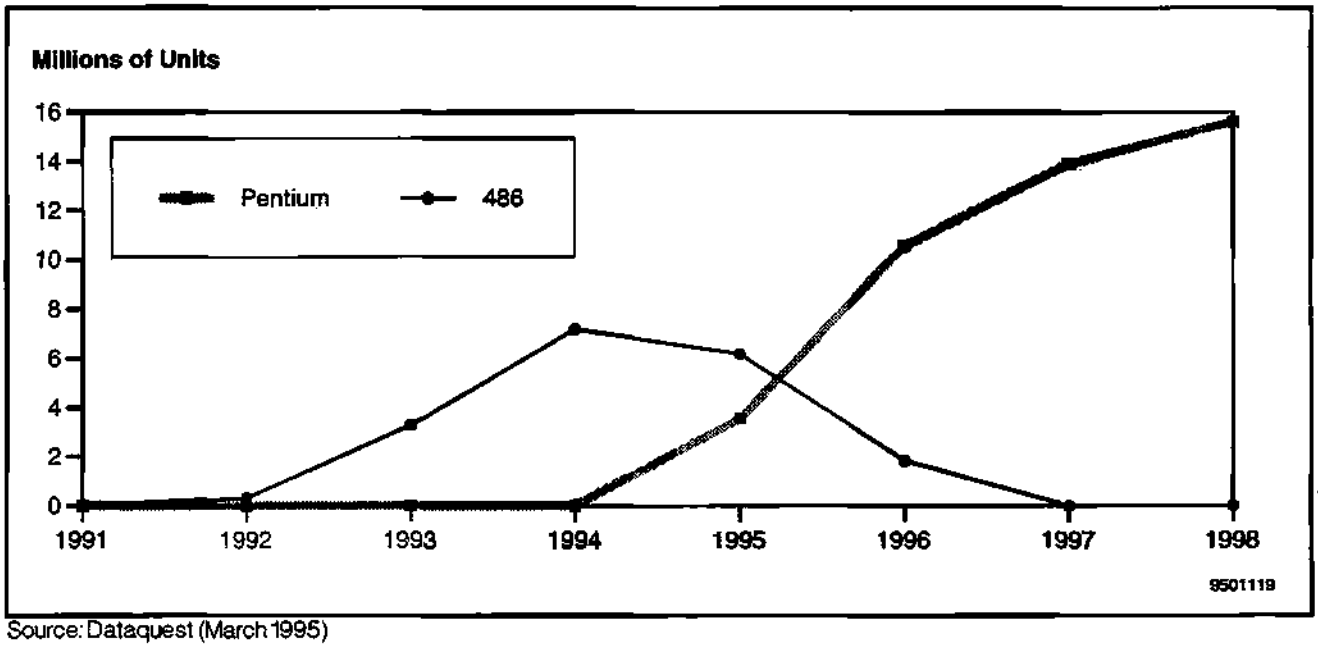

Source: Dataquest (March 1995)

P6 and the Workstation Market

- Digital's Alpha-based products are likely to keep the performance lead.
- Possible threat to second-tier (performance-wise) workstation suppliers exists.
- Sun's market share could be up for grabs.

We believe that workstation processors will probably continue to maintain an edge over the P6 at the high end. Even if P6 breaks the 400 SPECint barrier in 1996 (unlikely in our judgment), Hewlett-Packard, Digital Equipment, and Silicon Graphics Inc. (SGI) are all likely to have positions there already. Faced with an entrenched market base, Intel-based products may have a hard time dislodging these tenacious RISC players, although growth could be very tough to come by in the workstation business.

Digital Leads Workstations through Performance

Digital's model 900 workstation with the 275-MHz 21064 chip delivers a performance rating of 189 SPECint92, the fastest in the market today. Its 264 SPECfp92 rating was only recently overtaken by the IBM model 3CT at 267 SPECfp92, but this is nine months after Digital made initial deliveries of this product. Digital will have 300-MHz 21164 parts in the second quarter of 1995, and workstations based on this speed demon promise to eclipse all comers for another year or more. For the high end of the workstation market where absolute best performance is a key buying criteria, the Alpha-based workstations are growing share faster than any other. In a market where performance is king and the x86 water is rising over low-end players, this is at least a satisfying position. Continued growth will turn satisfaction to reward.

Sun's Workstation Share Is Declining

SuperSPARC has been the perennial bad boy of the RISC class, lagging not only its own brethren in performance but yielding to the Intel Pentium, with itself being a joke in the RISC business. Vendors of workstations using other RISC processors have stolen the lead technical applications that demand raw performance. In this process, Sun's price/performance value—but not price—in the mainstream workstation market, where bang-for-the-buck is an important buying criteria, has been eroded. SuperSPARC has dominated Sun's product line for the past few years, causing Sun to steadily lose market share since the third quarter of 1993. HP's PA-RISC 7100-based workstations started the rot, joined by Digital's Alpha (enabled by application support) in 1994. To compound the issue, IBM has attacked price/performance through aggressive pricing and new low-end models.

Table 3-1 shows the workstation market shares, with Sun declining and IBM and Digital gaining. Digital's workstation volume has shifted to higher-end, higher-priced models—so its revenue has grown more than its unit shipments. The other vendors all grew more at the low end of their model range (but SGI's average sales price still went up because its high-end models sold at much higher prices, loaded with expensive graphics options).

Table 3-1
Worldwide Workstation Market

	1993	1994	1993-1994 Growth (%)	1993 Market Share (%)	1994 Market Share (%)
Unit Shipments					
Digital Equipment	67,402	83,360	23.7	10.8	10.7
Hewlett-Packard	118,625	154,222	30.0	19.1	19.8
IBM	68,531	100,917	47.3	11.0	12.9
Other Vendors	94,970	112,639	18.6	15.3	14.5
Silicon Graphics	35,591	46,370	30.3	5.7	5.9
Sun Microsystems	236,800	281,877	19.0	38.1	36.2
Total	621,919	779,385	25.3	100.0	100.0
Growth over Previous Year (%)	4.1	25.3			
Factory Revenue ($M)					
Digital Equipment	976	1,219	24.9	9.5	9.8
Hewlett-Packard	2,330	2,596	11.4	22.7	20.9
IBM	1,582	2,315	46.3	15.4	18.7
Other Vendors	1,278	1,614	26.3	12.5	13.0
Silicon Graphics	925	1,290	39.4	9.0	10.4
Sun Microsystems	3,153	3,359	6.5	30.8	27.1
Total	10,247	12,395	21.0	100.0	100.0
Growth over Previous Year (%)	13.1	21.0			

Source: Dataquest (March 1995)

Workstation Market Impact

It is important to realize at this point that Intel has revealed the P6 architecture and has yet to announce the chip availability and pricing, let alone have a system vendor announce availability of a system product. Nevertheless, this architecture announcement will quickly come to market in systems and will have a dramatic impact on the traditional UNIX RISC-based workstations market. Intel is revealing a road map that increases the overlap between its and the RISC vendors' microprocessor performance. The P6 provides Intel four important improvements in relation to workstations:

- Improved data throughput by means of wider buses and data handling schemes, important to the more complex and data-intensive problems that workstations are called on to solve
- Faster and larger cache support, again important to more complex problem solving
- Four-way symmetric multiprocessing (SMP) without additional glue logic, fast becoming a means to tackle high-end desktop niches
- A jump in floating-point performance compared to the weak capability of the Pentium chips, important to technical applications and graphics

These improvements will effectively create a PC standard capable of capturing most of the opportunity in the low price points in technical and graphics-intensive computing, much more so than any Intel-based systems have done in the past.

The RISC architectures are popular in the traditional workstation market since they replaced the Motorola 68000 series and have captured applications that are not available for Intel-based systems. However, this situation will change as applications are ported to Windows NT, which provides more of UNIX's features needed by workstation users, particularly multitasking and OpenGL graphics support. The RISC workstations will still offer next-generation floating-point features and pipelines geared to graphics and multimedia beyond the capabilities of the P6.

The P6 makes it essential that workstation vendors take a PC design approach to their low-end RISC platforms. That is, they must use PC standards, especially the PCI bus, as much as possible to compete with the PC industry cost structure, or they will not be able to survive the PC value proposition.

HP, IBM, and Digital have PC product lines to complement their workstation business and should position products synergistically. Sun can position Solaris/x86 on PC platforms, but it will lose hardware revenue in that scenario. Sun's SPARC microprocessors suffer in the market from relatively low performance. The PowerPC remains the best alternative to Intel chips because it already has Apple driving unit volume, and thereby production and R&D sustenance. Alpha will remain the high-performance choice and the premium CPU for Windows NT. Intergraph has completed a transition to Intel-based workstations and Windows NT.

In analyzing the market changes and overlap of PCs and workstations, it is clear that there is abundant opportunity for the large PC vendors to capture a share of the workstation and advanced desktop growth. Key

enablers for this to happen, besides Intel delivering suitable high-performance chips like the P6, are the continued application wins for Windows NT and expected availability of advanced 3-D graphics for the PCI bus.

P6 and the Server Market

- P6 servers threatens high-end Pentium server products.
- It is still a market for low-end Pentium servers.
- Server manufacturers should get product ready for launch at P6 introduction.
- It is not embarrassing to use Xpress products.

The P6 processor has the potential to profoundly affect the midrange systems and server market. Certainly, the P6 will quickly seize the leading performance position in the PC and entry-level midrange server segments that are today dominated by the x86 architecture. Intel has been developing boards and chipsets for the P6 for some time and is no doubt ready to deliver P6-equipped Xpress platforms to its OEMs as soon as the product is introduced. The added performance of the P6 also signals the end of life for high-end multiprocessor Pentium systems. This is not to say that Pentium servers will instantly disappear from the market; these systems will continue to have a place in the lower-priced segment of the market, especially in the cost-effective dual-processor configuration. However, it needs to be emphasized that the introduction of the P6 has occurred earlier in the Pentium's life cycle than the introduction of the Pentium occurred in the 486 life cycle, meaning that Pentium will have a relatively short life as king of the server hill in comparison with the 486's reign.

Midrange and PC server vendors selling into performance-sensitive markets would do well to either manufacture an Intel product on an OEM basis or work with Intel's design groups to get a product to market quickly. HP, for example, already ships an Intel-built server and will open up devastating inroads into the market if it has free rein with a P6-based server early in the life cycle. With products from top to bottom based on Intel x86 and Pentium technology, AT&T GIS could also make significant gains assuming timely product introductions. This is not a market for the faint-of-heart. Compaq and IBM will find themselves under threat from competitors using Intel's products if they persist with a totally proprietary approach to their initial P6 products, unless they are ready with systems as the P6 debuts in the market. NetFRAME and Tricord, already under siege from the major PC manufacturers, also face serious challenges in this area. A sophisticated architecture is almost always trumped by a faster processor. Companies that stand to benefit include Dell, which could inflict significant pressure on Compaq and IBM at the departmental level, and Zenith Data Systems, which already uses Intel's Xpress box in its Z-Server line.

We believe that the P6 will make its first foray into the PC server and entry-level midrange server segments of the market. We expect the initial volumes to be in the one- to four-way SMP server systems. In the server market, it is becoming more crucial to have a multiprocessor-capable, multithreaded operating environment (Microsoft's NTAS or IBM's OS/2, for example) to take full advantage of the performance P6 offers, particularly in the SMP space. Intel's design, with its direct cache-to-cache

transfer over a 528-MB/sec bus and up to four CPUs, makes it ideal for the rapidly expanding low-end sweet spot in the server market. Intel has also equipped the P6 with some headroom and is quietly predicting that within two quarters of initial release it will bump the cache up from 256K to 512K. While doubling the cache does not double performance, the 512K cache option is yet another component that makes the P6 attractive to server system vendors because of the expected improvement in a multiprocessing environment.

We believe that the first P6 server systems could be announced as early as the third or fourth quarter of 1995. There are several Intel-based system vendors that will have either board and or system products ready as soon as Intel makes the P6 available. As with prior x86-based systems, a certain level of volumes needs to be reached to drive down the initial price and achieve the low competitive prices typical of the PC and entry-level midrange markets. This volume will come from the desktop market. As the P6 standardizes performance in the low-end segment, the added value for server vendors moves further into the area of service and support. Therefore, IBM, Compaq, and other global vendors have a key market advantage that must be fully exploited to fend off the invasion of high-performance low-cost clone servers.

To move the P6 beyond the PC and midrange entry-level server segments, some improvements are still needed. The relatively slow 528 MB/sec data transfer rate becomes an issue when the P6 competes against the 1GB+ server implementations that advanced MIPS-, PowerPC 620-, and HP PA-RISC-based servers will deliver; note that the ES/9000 mainframe delivers 2400 MB/sec between the main memory subsystem and the secondary cache In the medium-size and enterprise server segments of the market, I/O bandwidth, not raw CPU performance, is the constraining factor. Surrounding components need to be optimized to take advantage of—and complement—the type of performance provided by this (or any) new processor. This includes the operating environment and applications. With its relatively low-end maximum data transfer rate the P6 may not be ready to dislodge high-end RISC processor-based systems from their sizable midrange server market share. We suspect that the P6 will not move beyond the one- to four-way server segment before the second half of 1996 or later.

The P6 and the Mainframe

- The P6 is one generation away from mainframe performance functionality.
- The P6 shares many mainframe attributes.
- Difference is one of scale, not architecture.
- P6 is potentially faster than a mainframe.
- P6 lacks total integrity and I/O bandwidth.

The PC has evolved from a toy to a general-purpose computer that will become the backbone of office computing in the near future. Here, we are going to take a look at the driving forces behind the transition and show where the P6 fits in this progression.

Our approach is to look at the features that differentiate a mainframe from a PC product, and how the PC crosses the barriers that these features represent.

There are six fundamental barriers that protect the mainframe computing structure. Table 3-2 shows these barriers and Dataquest estimate of when PC-type technology will overrun them.

Table 3-2
Fundamental Barriers between the Mainframe and the PC

Feature	PC Delivers	Key Factor
Enterprisewide Applications	1992	NetWare
Individual CPU Performance	1995	P6; PowerPC; MIPS; PA-RISC
Multiple CPU Performance	1996	P6 and others
I/O Bandwidth	1998	Need 2 - 5 GB/sec bandwidth
Data Security and Integrity	1998	Not well understood
Application Code Base	1998	Emulation is key

Source: Dataquest (March 1995)

Enterprisewide Applications

In 1989, PC networking was in its infancy. The simple task of sharing a printer represented a major systems integration challenge with the potential for less than satisfactory results. By 1992, networking systems had improved to the point where centralized file sharing and management were reliable, and applications could be developed that relied on network resources and distributed computing.

Today, enterprisewide applications are being deployed wholesale on PC-based networks, and wide area networking allows them to operate on a national or even global basis.

Individual CPU Performance

As our review of the IBM 700 CPU architecture in Chapter 10 shows, the budget dollars associated with mainframe R&D allows mainframe manufacturers to invest tremendous amounts of money in the advancement of performance. However, the laws of physics are fast catching up with mainframe designers, and single-silicon processors look set to draw ahead over the next few years. The fastest RISC processors announced today have drawn level or past the fastest IBM mainframes and, although another generation of mainframes is due, will pass mainframe performance permanently with their next iteration.

Multiple CPU Performance

IBM's newest ES/9000 processors couple up to 10 processors together in a single system. These processors all run the same operating system code and share main memory.

Figure 3-2, taken from IBM marketing information, shows how performance scales as extra processors are added to the complex. Although there is some flattening of the curve, it is quite gentle and shows that IBM is able to get significant incremental performance in up to a 10-way multiprocessor system. This scalability reflects both the tremendous bandwidth of the memory buses in the machine and the dual 10-port L2 caches that feed the two memory banks.

Figure 3-2
IBM ES/9021 Scalability by Application Class

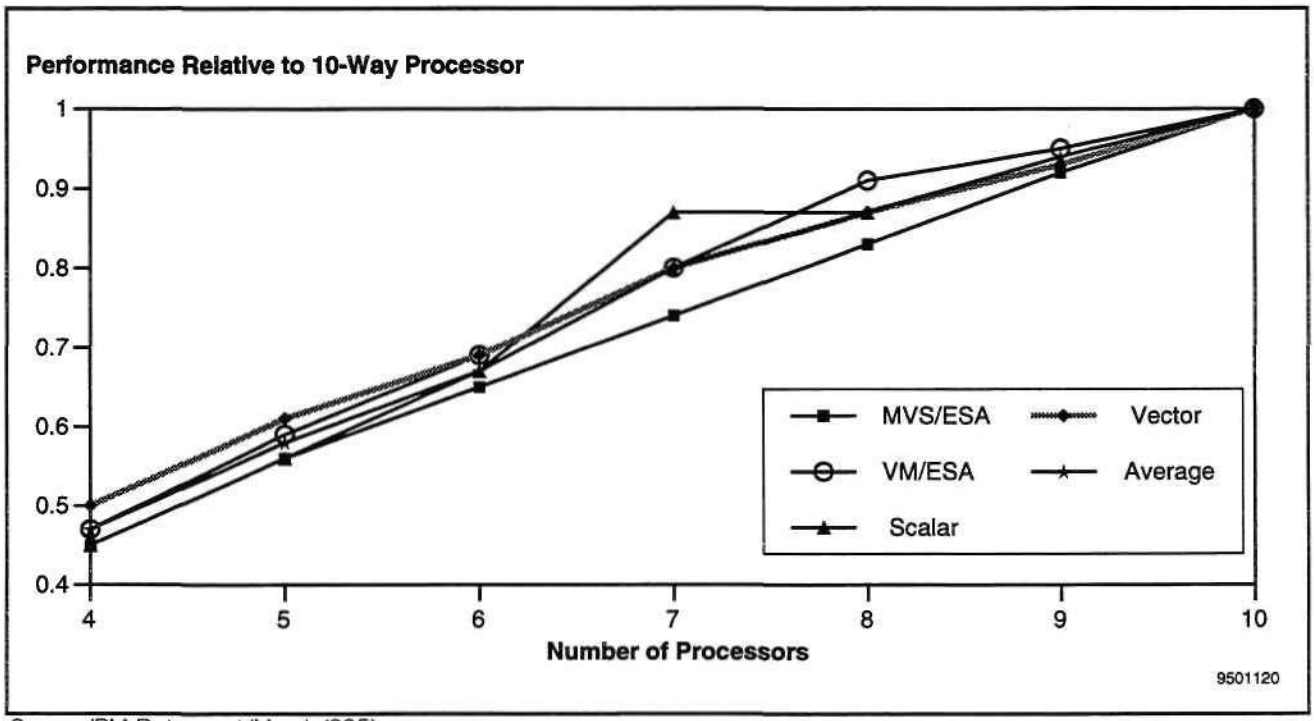

Source: IBM, Dataquest (March 1995)

The next generation of microprocessors from HP (PA-8000), MIPS (R10000), Sun (UltraSPARC), Intel (P6), and IBM (PowerPC 620) all incorporate some form of high-speed bus to provide improved scalability in multiprocessor applications. Whether they will achieve the level of scalability displayed by IBM remains to be seen, but they will certainly do well in four-processor clusters.

I/O Bandwidth

The buses in the IBM mainframe will support data transfer rates of approximately 2.4 GB/sec, well over the 1 GB/sec of the next-generation RISC processors and five times that of the P6. (Note, though, that all of these machines offer these data rates to the main memory subsystem; they all have decoupled cache architectures that reduce bus traffic, increasing the value of the available bandwidth.) This bandwidth is a function of internal architecture and of the width of the I/O buses and their clock frequencies. The mainframe has the benefit of 128-bit wide buses, a feature that is physically very difficult for today's microprocessors. The enabling technology here is the high-density ball-grid-array package, which offers both outstanding electrical properties (for high-frequency operation), and

sufficient interconnect for mainframe-class bandwidth. Our judgment is that the advanced RISC processors will hit these bandwidths in 1997, and that the P7 will match today's mainframe in 1998.

Data Security and Integrity

Many years ago, companies ran their mission-critical accounting applications through teams of accountants. Double-entry bookkeeping and careful, daily reconciliation guaranteed integrity of the data, without which the company would fail. To succeed in the business market in the 1960s, the mainframe had to offer not only functionality and performance but an iron-clad guarantee that data would not be lost or corrupted without notice. Note that loss or corruption, while undesirable, can be tolerated to a very limited extent as long as the error is flagged and can be corrected; the cardinal sin is to propagate the error through the information infrastructure.

Evolution fitted the mainframe with ECC protection on its buses and main memory, with parity on the caches and sometimes even through the execution units.

The P6 does not quite match this standard of performance, lacking address bus ECC. For these reasons, it cannot quite achieve the guarantee of integrity required to displace a mainframe. However, the next generation of product will have the luxury of far more I/O pins (better living through ball grid arrays) and will no doubt rectify this situation through the addition of the extra control signals.

Application Code Base

The application code base remains as the last challenge. The life cycle of mainframe-based software can span decades; the hardware is upgraded, but the software is dragged across with as little modification as possible. This approach reflects the enormous investment that it takes to build a system that a major corporation can trust with its critical records.

On the other hand, many of these applications were developed to run on mainframes that had significantly less processing capability than even today's Pentium-based PCs. Therefore, as microprocessor performance advances over the next few years, instruction emulation rather than direct execution of these applications becomes a real possibility. The P6 has sufficient power to emulate many of these applications today; progress will advance the art further, and by the end of the decade, we expect today's mainframe applications to be running effectively on microprocessors.

In summary, we conclude that the next generation of mainframe will be able to compete with the next generation of processors from Intel and others, but that this will change at the turn of the decade and mainframes will face their true watershed. A large system of either microprocessor or mainframe architecture will have an aggregate memory bandwidth of 9,600 MB/sec or more (double that of today's high-end systems), spread across multiple processors and memory banks.

Chapter 4
System Design Issues

Power

- A P6 system requires 5V, 3.3V, 2.9V, and 1.5V.
- Power-managed motherboards are required.
- There will be a new market for local switching regulator manufacturers (Micro Linear, Maxim).

At 2.9V and 2W, the P6 requires 6.9 amps of peak current. A linear regulator will have to dissipate 15W to provide this voltage from a 5V supply, and the power bill for the system will be over 35W because of the processor alone. This latter factor will make it difficult to achieve the EPA "Energy Star" standard. Therefore, we believe that the system will have to use several switching-type regulators to provide power to the P6; linear regulators will be too hot and too big.

P6 systems will incorporate four tightly regulated voltage supplies: 5V for the I/O bus and 5V logic parts, 3.3V for the system memory and 3.3V logic parts, 2.9V for the P6 itself, and 1.5V for the GTL+ bus. Each of these supplies – including the GTL+ bus supply – will need to provide several amps of current.

We expect system designers to move to the use of local regulation for these supply voltages, rather than try to build power supplies that support them. There are several reasons for this.

Regulation

5V logic supplies typically require tight regulation and are therefore the master voltage in most PC power supply designs today. These supplies are quite inexpensive (perhaps $25 for a 200W supply), the result of intense competition. In these designs, the other voltages are allowed to drift over a somewhat larger range as the application of these voltages is typically noncritical. In a P6 system, however, the 5V, 2.9V, 1.5V, and 3.3V supplies must all be tightly regulated, adding significant cost to the power supply.

Power Planes

To deliver these lower voltages successfully to the components, it is often necessary to have a power plane inside the board. Otherwise, the resistive drop becomes too large to guarantee regulation. A P6 design could end up with two extra power planes inside the board to distribute the 3.3V and 2.9V supplies, for an incremental cost that will be unacceptable as the P6 goes mainstream. Local regulation will make it simple to split the power planes without compromising noise immunity or introducing significant resistive losses, enabling designers to plan low-cost, four-layer boards.

Flexibility

The voltage drop hasn't finished yet. The P6 pinout reportedly incorporates a simple mechanism to signal the required operating voltage. Intel is quite likely to respecify the P6 operating voltage in a future generation, and local voltage regulation provides the flexibility to handle this change without a system redesign.

The P6 and Voltage Regulator Manufacturers

The P6 has an interesting side effect on the business of companies manufacturing advanced integrated voltage regulators, including Linear Technology, Maxim, and Micro Linear. At three regulators per board and 40 million boards a year, there is the opportunity to make hundreds of millions of dollars from P6-targeted regulators. We advise that system manufacturers enter discussions with these suppliers to evaluate their solutions.

Cooling in Fixed Systems

- The P6 needs decent airflow.
- Improved design will not need "supercooling" like original Pentium.

At 20W peak and 15W average, the P6 will get quite hot; this level of power dissipation is similar to that of the original Pentium. It harks back to days when system manufacturers were installing multiple, huge fans in their systems to cool the Pentium, and the Pentium's heat problem was the focus of many articles.

Those articles missed the true focus of the problem by suggesting that a machine might take fire, melt down, or otherwise self-destruct. The simple truth was that critical timing paths in the original Pentium became too long at relatively low die temperatures, leading to a temporary failure of the part if it became too warm. To keep the processor die at a relatively low temperature, the early Pentium systems needed significant airflow over the part leading to the substantial cooling arrangements mentioned earlier. In other words, the Pentium's problem was not excessive heat generation but excessive sensitivity to heat.

Intel managed to quietly sweep the whole issue under the carpet by tweaking the critical paths. Pentium processors still need good cooling, but it is no longer a consequence of design shortcomings.

The P6 is in line with many other next-generation processors that dissipate tens of watts of power. It will no doubt need a big cooling fan, but nothing that is beyond those shipping in Pentium systems today. We believe that Intel's fine-grained pipeline, coupled with the very careful clock distribution, will allow the P6 to meet performance expectations without excessive cooling requirements in well-designed systems.

Intel has been working with fan manufacturers to develop fan-assisted heatsinks ("fansinks") for the P6. The larger package area of the P6 permits the use of a larger fan in these devices, which in turn allows the fan to incorporate a more robust bearing than the small fans seen to date. This larger bearing delivers higher reliability than that of the smaller fans used to date, devices that have reportedly been subject to a small but significant failure rate. Intel has already developed multiple vendors of P6-specific fansinks.

On a final note, the fansinks for the P6 will draw their power from a disk power connector. The 2.9V supply does not work well with a brushless fan motor; the fansink manufacturers prefer the 12 volts available from the disk drive power connector.

Clearly, one side effect of this trend toward faster, hotter processors is that all the devices need some form of assisted cooling. Although fans and heat sinks for the desktop are not particularly advanced technology, there are products and manufacturers that offer innovative mounting and application solutions. Given the size of the market, even these manufacturers stand to make a gain from the introduction of the P6.

Cooling Advanced Processors in Notebook Computers

- Power consumption only seems impossible.
- Pentium techniques will help.
- Pico Power (Cirrus Logic) is likely to have the answer.

In the notebook market, however, cooling is a critical issue. Even though the notebook itself may suffer no ill effects for raised temperatures and hot spots arising from the use of a high-powered processor, the user may be disturbed by their presence.

There are three major problems—and hence three major opportunities—to be solved in the design of notebooks that use high-performance processors.

Reducing the Power Dissipation

The heart of the problem is to reduce the power that the processor itself consumes. Intel's System Management Mode (SMM) features help to some extent but lose value when the processor is required to perform at full speed. Lowering the processor's supply voltage decreases power consumption concordant with a square law, but transistor counts and clock frequency conspire together to increase power consumption faster than the reduced supply voltage decreases it. The problem is that the incremental thermal energy that must be dissipated may exceed the ability of the dissipation system to keep the processor running at an acceptable temperature; in the Intel solution, the processor must shut down when it exceeds its operating temperature limit. We believe that clock throttling solutions such as those offered by Pico Power, in conjunction with closed-loop thermal feedback under microcontroller control, can address this issue.

Getting the Heat out of the Unit

The heat has to go somewhere; there are six ways to unload it:

- A small fan—As demonstrated in the Compaq LTE Elite, it is possible to install a small fan in the unit to air-cool internal components. This approach has the advantage of reducing or eliminating hot spots but carries the stigma of a fan and can lead to trouble if the vent holes are blocked.
- The back of the keyboard—Typically a large metal plate, the back of the keyboard provides an opportunity to both distribute the heat and conduct it out of the unit.
- A metallic case component—The keyboard frame and cover assembly provides a good opportunity for the use of a magnesium casting: It benefits from the extra rigidity afforded by the metal, and the increased weight of the metal part over a plastic version is a minimal

issue because of the largely open construction of the part. The rectangular area that represents area not used by the keyboard presents an opportunity to both spread the heat from the processor and conduct it out of the unit. The high thermal conductivity of the metal provides a good heat sink path. The downsides are that magnesium tooling costs perhaps five times as much as the equivalent plastic tooling; the cost of the part is somewhat higher because it requires extra finishing operation; and the user may object to the noticeable heating of the part, particularly if it is effectively used as a palmrest in the design.

- A metallic subframe—Some Apple machines use a magnesium subframe to add rigidity and to provide mounting points inside the machine; this subframe also represents an opportunity to distribute heat inside the product, but the opportunity to move it outside is limited by the fact that it is enclosed inside the case.
- A metallic spreader inside the case—The heat can be conducted out through the plastic case, but the low thermal conductivity of the plastic requires that the conduction area be relatively large. A metal plate that spreads the heat over a large area of the back of the case can be used to dissipate heat. The low thermal conductivity of the plastic becomes an advantage in that it will feel cooler to the user than a metal component might; the user's hand is able to lower the temperature of the plastic that it contacts very quickly. We believe that Toshiba has implemented a solution like this in the Pentium-based 4900 notebook; NEC has vent holes in the base of its Pentium-based Versa P/75 unit that allow a metal plate to conduct heat to the atmosphere directly.

Transferring Heat Inside the Unit

Getting the heat out of the unit is one problem, but sometimes the design of the unit is such that the processor cannot contact a local heat sink. The IBM 700 series ThinkPads face this issue because the electronics are constructed from a multilayer stack that occupies only the rear of the cabinet; the front of the cabinet is reserved for the floppy disk drive, battery pack, and hard drive.

Heat sink manufacturers provide heat pipe solutions that can help solve this problem; the processor is cooled by a reservoir containing a fluorocarbon; the reservoir has a tube and a wick that connect to a dissipation plate located where heat can be dissipated out of the unit. In operation, the coolant evaporates in the reservoir, drawing heat from the processor. The vapor travels down the pipe to the dissipation plate, where it cools off (giving the excess heat up to its environment) and condenses. Then it travels back up the wick to the reservoir, and the cycle begins again.

The most extreme example of this class of solution was a unit that incorporated a liquid circulation system that transferred heated fluid through the hinge to a large radiator behind the noetbook's LCD panel; the large area of the radiator allowed the system to dissipate 15W without an unreasonable increase in temperature. This solution addressed both internal and external heat transfer issues.

Again, the introduction of the P6 bodes well for manufacturers of thermal management products. Notebook manufacturers will migrate all the faster to the Pentium as desktop machines open up the performance gap through the use of the P6, and the P6 itself will demand ever more advanced cooling systems for maximum performance.

Board Design at 66 MHz

- 66-MHz design is tough.
- Managed impedance approach will enable low-cost, four-layer boards.

The P6 bus is a GTL variant, termed GTL+ by Intel, that runs at 66 MHz and 1.5V. Intel claims that the move to GTL+ was necessary to support the multiple loads represented by up to four processors and the bus/memory control units. The bus requires termination at both ends, suggesting that it will require special attention at the design stage to ensure that its impedance falls within acceptable limits. One key benefit of this approach is that the GTL+ bus supply voltage is completely decoupled from the other supply voltages, preventing logic noise from affecting the bus supply and bus noise from affecting the logic supply.

If the board is not designed properly, the system can be unreliable in either of two ways: Systems may work well for some period of time, then fail in a variety of ways as they warm up, cool down, or run certain instruction sequences; alternatively, the system may not be manufacturable in volume as a percentage fails totally or exhibits symptoms as described earlier.

Traditionally, this class of products would be designed with a controlled-impedance motherboard, a process that requires specialized design and special manufacturing processes. We expect to see these techniques used in the initial P6 systems, but the manufacturing processes involved are too costly for mainstream P6 systems that we expect to see in 1996.

Therefore, we expect to see more system manufacturers switch to managed impedance designs. In this technique, the printed circuit board (PCB) manufacturing data is tuned to the target manufacturer's process, accounting for the parametric variations that may manifest themselves on the board through the manufacturing process. This tuning has the effect of centering the yielded trace impedance in the manufacturer's range, rather than arbitrarily off to one side. Although controlled impedance designs can offer a narrower distribution of impedance, this is probably unnecessary as long as the impedance is centered. The only requirement is that the manufacturer remains within his specified process parameters.

Figure 4-1 shows how these three approaches work. Assuming that the design requirement is 50 ohms +/- 10 ohms, it is clear that both the managed impedance and controlled impedance designs will result in very high yields. There is room for some process drift in either approach. The uncontrolled design, however, is clearly out of bounds in terms of its impedance range. Note that the yield is quite well controlled; it is just in the wrong place.

Systems built using boards that have not been designed for some kind of impedance control may well work most of the time but are likely to fail when the bus is fully loaded and other conditions are at the extremes.

Figure 4-1
PCB Trace Impedance Illustration

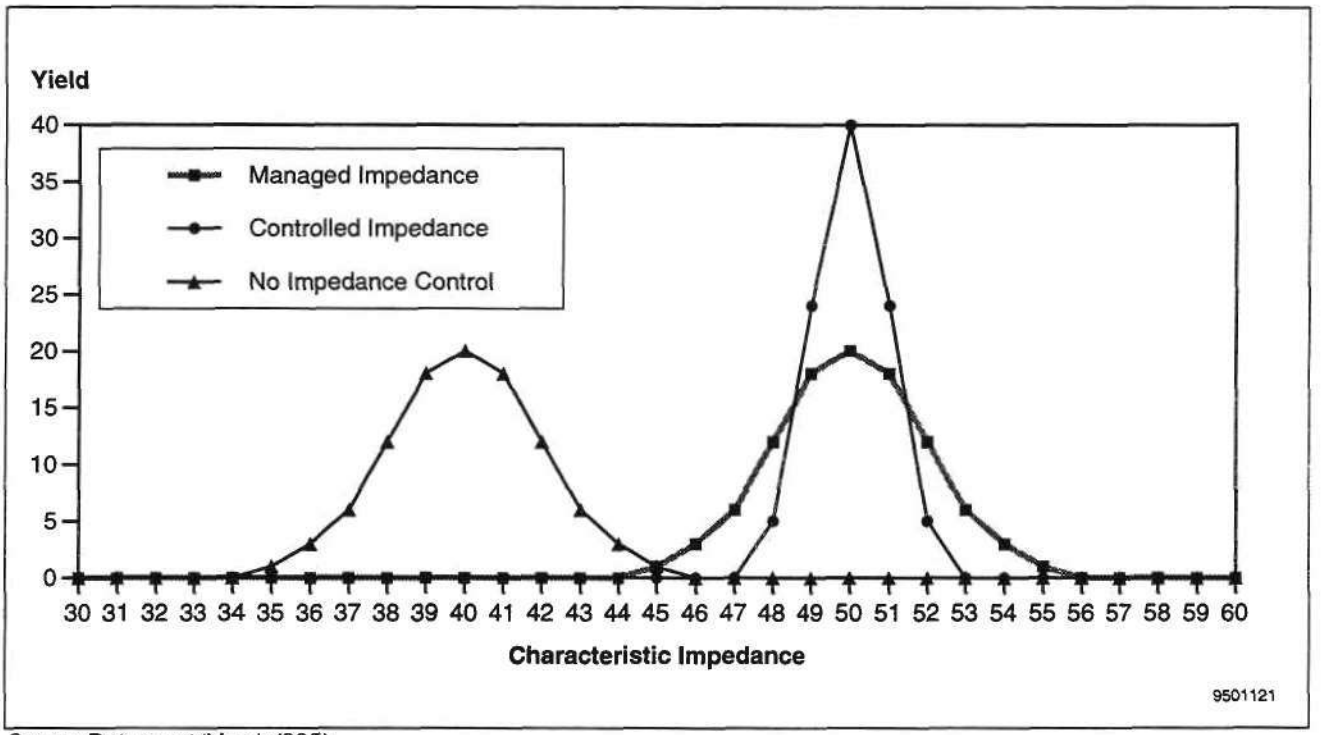

Source: Dataquest (March 1995)

P6 Technology Support

- Intel will be the only choice in 1995 and most of 1996.
- Designing for the P6 is very different.
- Core logic differentiation is difficult.
- Therefore, system differentiation is difficult.

Chipset Manufacturers

Intel has only recently disclosed P6 specifications to commodity core logic manufacturers. The P6 bus incorporates complex multiprocessing protocols and uses an unconventional logic interface, suggesting that the core logic companies cannot simply tweak existing designs to work with the P6. Therefore, these companies will be faced with a substantial lead time before they can introduce products that can compete with Intel's chipsets. Although this approach bolsters Intel's chipset division, it is a two-edged sword in that the lack of suppliers could dampen the market for P6 in its first 18 months.

There is, however, an opportunity to develop less sophisticated chipsets that support only a single or dual processor system, eschewing the complexity of the split-transaction bus for lower cost and perhaps better suitability to notebook or simple desktop designs.

Chipset Support

A multiprocessor split-transaction bus is a very different approach for a chipset manufacturer to deal with. Although it can be pared down to a fairly simple single-chip bus, chipsets that do this will face a challenge in the market when competing with Intel's variants. Coupled with the modified GTL interface and a requirement for a PCI bridge, it may be some time before a non-Intel variant becomes available.

We also note that Intel's P6 chipset was developed in conjunction with the P6 itself; the pair has been the subject of extensive cosimulation, an approach that greatly improves the functional reliability of the chipset in this complex transaction-based environment. We would postulate that the simulation models for the P6 bus are not available outside of Intel, presenting another challenge for chipset manufacturers.

Corollary and LSI Logic

We have always favored Corollary's C-BUS II multiprocessing logic, and we believe that much of that technology can be evolved into a P6 multiprocessor set that would enable P6 quads to be grouped into yet larger multiprocessors. LSI Logic has key cells and I/O technologies and could also be a contributor to a system manufacturer looking to differentiate its products. We recommend that system manufacturers consult these vendors to evaluate multiprocessing strategies beyond Intel's basic four-way systems.

Regulators, Heat Sinks, and Sockets

The P6 will engender a whole new scale of production of these components. With three discrete regulators on the system board and a guaranteed need for heat sinks for the foreseeable future, the market for these parts is set to grow as it follows the ramp of the P6. The P6 will reach a run rate of 40 million units a year by 1998, driving these previously sleepy markets to new heights.

Chapter 5
Architecture and Performance

Conclusion

After evaluating the architecture of the P6, we come to the conclusion that this design is built for speeds far in excess of the 133-MHz rate promised at introduction. There are four features that lead us to this conclusion:

- The careful clock distribution, which could run at 400 MHz if a 10 percent skew were allowed
- The fine-grained pipelining, which eliminates complex single-clock cache accesses and instruction decodes
- The 256KB cache, promised to go to 512KB late in 1995
- The ability of the decoders to outpace the execution engine

Furthermore, all of these features will scale with finer-geometry process technology. Time will tell, but we would not be surprised to see the part reach 200 MHz sooner rather than later; the lucrative server market may permit Intel to hand-pick such parts from the current 0.6-micron process lines.

Architectural Overview

- Superscalar; multiple execution units; like PA-8000, MIPS R10000
- Translates instructions internally, like NexGen and AMD's K5

Architecturally, the P6 follows the trend set by the latest round of next-generation processors that decouple execution resources from the instruction decoder. Although this class of machine is generally described as superscalar, we feel that it deserves a description that differentiates it from the dual-pipeline approach typified by the Intel Pentium and Cyrix M1; Intel chooses to describe this configuration as a "dynamic execution" machine; we prefer the phrase "decoupled resource," although MIPS has already used the term "Dynamic Execution" (it is the DE in ANDES, the R10000). This approach first appeared in the mainframes of the 1960s, notably the Control Data 6000 series and the IBM 360/91. It seems that microprocessors are evolving along the mainframe timeline—both classes took about 20 years to reach this level of complexity—although the evolutionary clock has been running fast for the past few years.

The machine can be described as four major subsystems: the in-order fetch and decode subsystem; the out-of-order execution core; the reorder logic that sorts out the instructions and registers after execution; and the memory interface. Intel is very coy about the memory interface, claiming that it incorporates trade secrets that are critical to achieving maximum performance.

The P6 decomposes x86 instructions into internal micro-ops (microcode operations), an approach previously disclosed in the AMD K5 and NexGen 586 products. Intel started translating x86 instructions to a regular format in the 486 and then the Pentium, a change that made the implementation of these pipelined processors easier to manage. This fact has not

been widely published, but Intel elected to call the new translated instructions microcode operations, or micro-ops, or uops because they are designed to work in the P6's out-of-order scheduling environment. These uops are 118 bits wide and are of regular format, an approach that makes their decoding and management in the out-of-order execution core simpler. The trade-off to this approach is that the P6 has to route the many wide instruction buses inside the machine, increasing the complexity of the design. By comparison, the HP PA-8000, which uses a similar architecture, has a 32-bit internal instruction bus.

The machine uses different instruction widths through different paths in the design (see Table 5-1).

Table 5-1
Microcode Widths in the P6

Pathway	Width	Comment
Microcode Sequencer Output	72	
Decoder Output	123	
Register Alias Table Output	118	
Reservation Station to Advanced Arithmetic Units	86	Data bits; instruction bits not disclosed (but known to be few)
Other Reservation Station Outputs	32	Data bits; instruction bits not disclosed (but known to be few)
Reservation Station Registers	254	Holds up to two floating-point numbers

Source: Intel, Dataquest (March 1995)

As a note on superscalar architectures, the P6 most closely reflects the design of the original Control Data 6600 machines. This machine queued instructions centrally, an approach echoed in the advanced processors from MIPS and HP. AMD's K5 and IBM's PowerPC architectures reflect the IBM 360/91, a machine that dispatched instructions to reservation stations located in execution units. This latter approach is followed in the modern IBM ES/9000 mainframes, an example that is briefly described in the mainframe section of this report.

The centralized reservation station is intuitively more flexible than the distributed approach, as instructions may be scheduled to any free unit. With distributed reservation stations, once the instruction is gone it cannot be recalled; this restriction could lead to imbalanced use of the execution resources under some circumstances. The spread of implementations between the two techniques, however, indicates that there is probably little real difference one way or the other.

P6 Architecture Summary Diagram

Figure 5-1 shows the overall architecture of the P6, identifying the in-order and out-of-order components.

Figure 5-1
The P6 Architecture

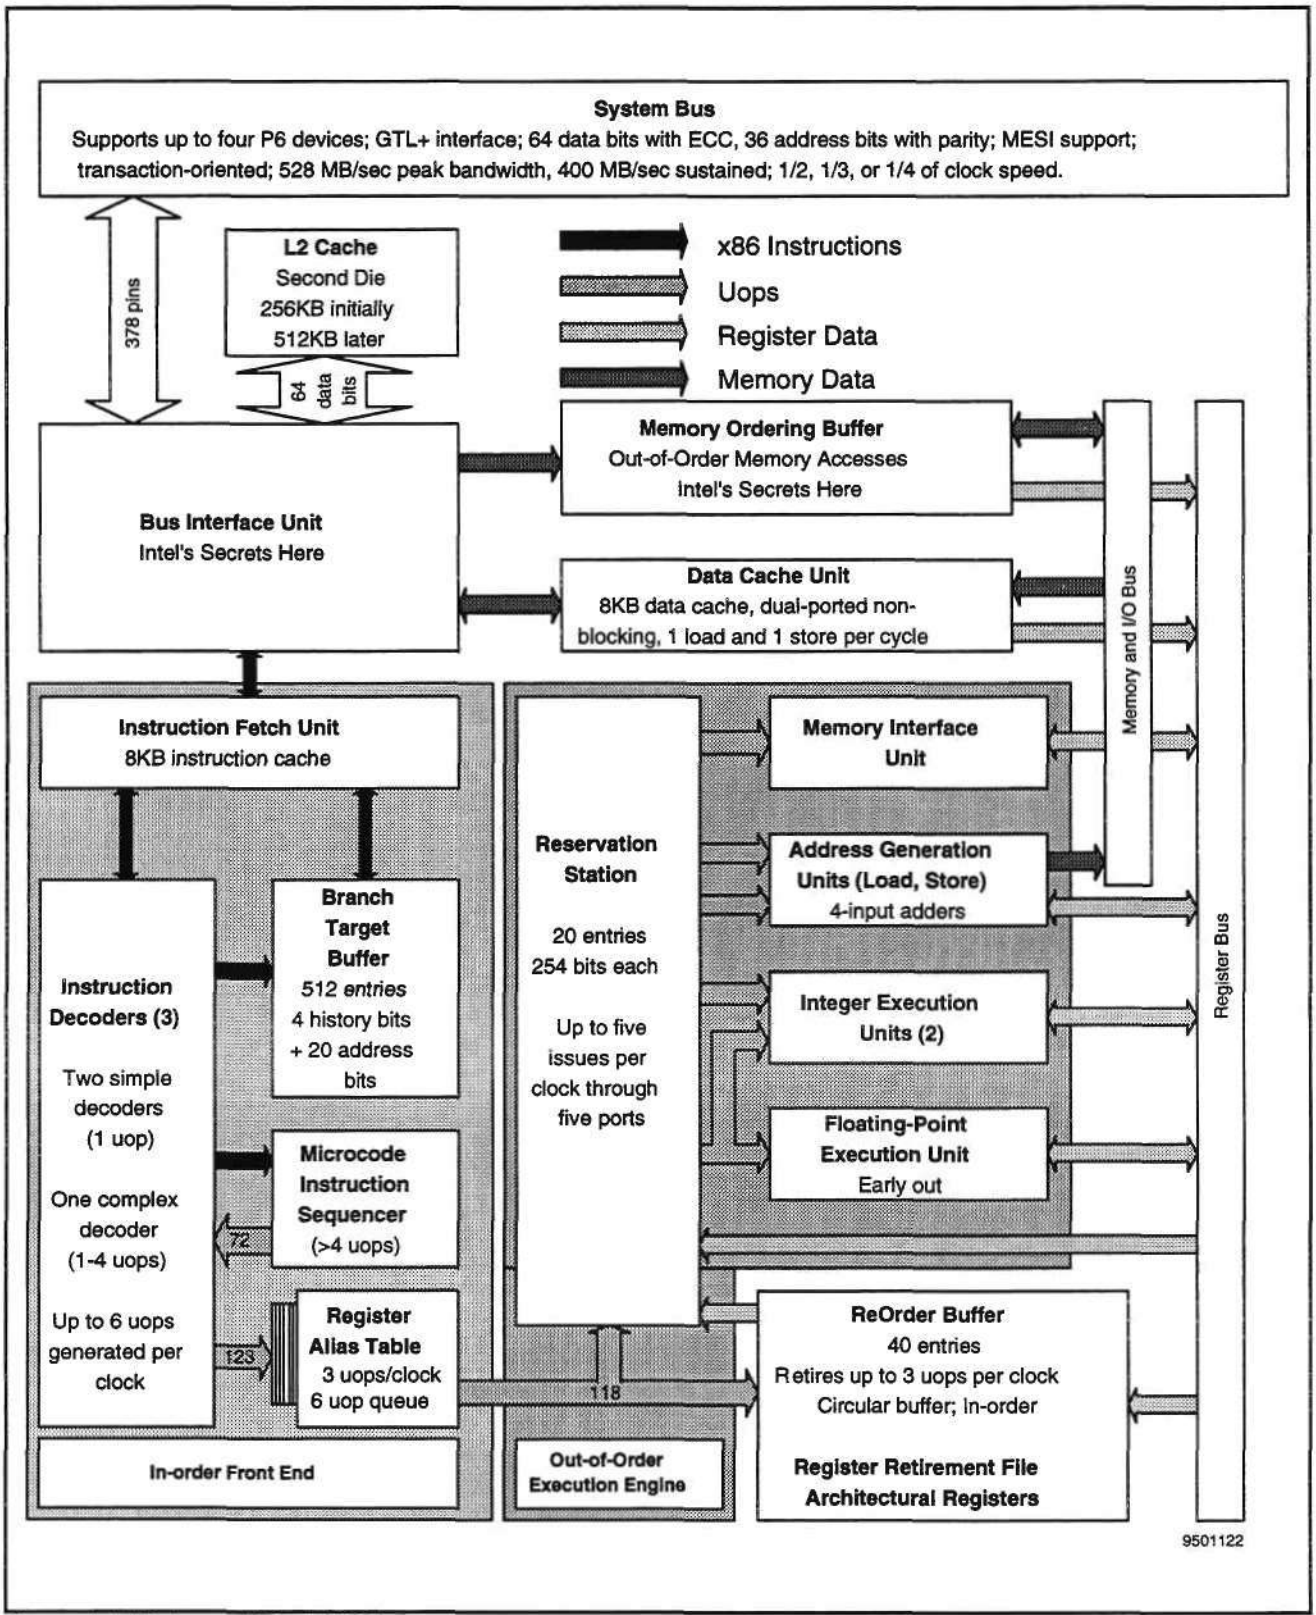

Source: Intel, Dataquest (March 1995)

Branch Prediction

The P6 incorporates an advanced Branch Target Buffer (BTB), which includes both four branch history bits and a Branch Target Address Cache (BTAC). BTAC contains a further (we estimate) 16 bits that record the address that the branch jumps to if taken. This sophisticated prediction and caching mechanism helps the P6 keep up the flow of instructions to the execution unit, a key design goal of the P6. The deep pipeline makes this sophistication essential—A missed prediction can result in a penalty of up to 15 cycles before the processor can start performing useful work again. Clearly, this does not have to happen very often to slow the processor significantly.

The BTAC in the P6 is very sophisticated compared to that of other microprocessors; Table 5-2 shows how the P6 compares with other advanced processors.

Table 5-2
BTAC Sizes in Advanced Processors

Processor	Branch Target Address Cache Entries
HP PA-8000	32
Intel P6	512
Cyrix M1	256
NexGen NX586	96
PowerPC 620	256
IBM 700 Mainframe	4,096

Source: Dataquest (March 1995)

The IBM 700 processor is included here for reference. IBM's mainframe designers have spent many years optimizing architectural features; its gigantic 4,096 entry BTAC is a harbinger of things to come, and we can expect the next round of processor disclosures to identify BTACs approaching this number of entries as designers gain access to more and more transistors.

Instruction Fetch and Decode

- Decodes up to three x86 instructions per cycle
- Can outpace execution engine
- Includes microcode sequencer for complex instructions
- The potential clock speed bottleneck

The instruction fetch and decode unit together take care of the task of detecting instruction boundaries and splitting out the individual x86 instructions. There are three decoders, of which two deal with the simpler x86 instructions and one deals with the complex instructions that require more than one uop; Intel claims that most instructions can be handled by the simple decoders. Assignment of instructions to the decoders is simple—they fall where they may. If a simple decoder cannot handle the instruction assigned to it, it passes it to the complex decoder in the next clock cycle (or whenever the complex decoder is available—all three instructions being decoded could be memory stores, a combination that would take three cycles to decode).

Simple instructions include register-to-register, immediate-to-register, and memory-to-register operations. A register-to-memory operation can generate up to four uops, two of which are required to write to memory.

The x86 instruction set includes instructions and implied loops that can execute repeatedly, including string searches and moves. The P6 accommodates these and other complex instructions by means of a microcode instruction sequencer that can generate strings of uops as required.

Again, Intel has designed the instruction decoder to be able to decode instructions and generate uops at a faster rate than the execution engine can typically handle. This approach guarantees that the core of the processor is kept as busy as possible, even when it may have to flush a few instructions that were executed speculatively.

The decoder can generate up to six uops per cycle, one from each of the two simple decoders and four from the complex decoder.

Once the instructions have been converted into uops, they are cast into the out-of-order execution core. Each uop is assigned a set of temporary registers on entry into the execution core through the Register Alias Table. The Register Alias Table is only three uops wide; if the decoders generated six uops, they are transferred in two cycles.

It is notable that the personality of the P6 architecture is totally defined in the in-order instruction decode stage; although there are perhaps some differences in the operation of the execution units, the out-of-order engine could be part of virtually any architectural implementation. Therefore, changing the programmable logic arrays (PLAs) and ROMs that control operation of the decoders could result in a totally new architecture, or the emulation of a current alternative; we do wonder just how fast a variant of this machine could run the PA-RISC instruction set. Loadable microcode, implemented through static memory instead of ROMs in the decoder, brings some interesting possibilities to this architecture.

We understand that decoding the instructions is the toughest timing challenge in the machine. The three candidate instructions must be aligned, loaded into the three decoders, and turned into uops once each cycle. The chip's designers conquered this problem by splitting the process across 2.5 clocks and overlapping it with the instruction cache read. This path is important because it represents the potential limit to clock speed. Clearly, the designers have paid a lot of attention in its subdivision, and the fact that they are well aware of it augurs favorably for prompt introduction of faster parts.

Instruction Execution

- The P6 has out-of-order core.
- It has central control of execution (as opposed to distributed).

The uops are loaded into a 20-slot reservation station, from which they are dispatched to the various execution units. The reservation station tracks operands (in temporary registers) for each instruction and dispatches an instruction and its associated data to an execution unit when it knows that the operands will be available. It has five instruction issue ports, which

combine to allow the P6 to execute a maximum of two arithmetic, one load, and one store instruction per cycle; note that a store requires two uops, one to perform address calculation and one to transfer the data.

The 20 slots in the reservation station each hold 254 bits, permitting any slot to support the two 86-bit values needed for floating-point operations. Each slot can connect to any port, leading to a very dense routing of interconnect.

The uops are also entered into the reorder buffer at this point. This is a circular buffer structure with head and tail pointers. New instructions increment the head: As instructions are retired, the tail increments; a flush of one or more instructions adds the appropriate count to the tail pointer.

Dataflow Aspects of the P6 Architecture

The P6 uses elements of a dataflow architecture to control the way that the reservation station operates. In a dataflow architecture, processes occur as a result of a data element being in a certain location in the system. In other words, execution events are triggered not by program flow but by data presence, and operations are completed by sending data through a target execution unit. The reservation station follows this concept: Data is dispatched to an execution unit when it is ready to go. The program sequence of the original instructions is nearly irrelevant to this process.

The Execution Units

The P6 has six execution units, three for arithmetic and three for load/store operations to memory. Each has its own data path, and all six are connected by a bypass bus (omitted from the diagram for clarity) that allows the transfer of results from one to the other without passing through the reservation station.

Integer Unit No. 1

This unit is fed from the same reservation station port as the floating-point unit. It includes basic integer arithmetic, integer multiply and divide, and a shift unit. The integer divide unit is faster than that of the Pentium, completing two bits per cycle.

The P6 concentrates all of its advanced arithmetic in this location for two reasons: Traces show that this allocation matches the proportional usage; and these units require a common 86-bit data path to the reservation station. The execution units on the other ports require only 32-bit data paths.

Floating-Point Unit

- No divide bugs
- Similar to Pentium unit
- Early-out speeds floating point

The P6 floating-point unit is (loosely) based on the unit used in the Pentium (minus small divide errors) and is quite identical in the add and multiply sections. The floating-point divide unit is similar to that used in the Pentium, but has been modified to produce early-out results. This is an

enhancement of little value in the sequential Pentium design, but of great use in the P6's free-for-all execution strategy where it lets the P6 turn floating-point divides up-to twice as fast as the Pentium can. The P6 includes an enhanced square-root unit that is twice as fast as that used in the Pentium, perhaps aimed at supporting native signal processing.

Integer Unit No. 2

The second integer unit includes only basic integer arithmetic (no multiplies, divides, or shifts) and is used for these operations or branch target address calculation. It has its own reservation station port; Intel's analysis showed that basic integer arithmetic is more important than all of the floating-point calculations, shifts, and integer multiplies/divides together.

Address Generation Units

The P6 incorporates separate address generation units for load and store operations to memory, enabling it to perform one of each per cycle. These units include sophisticated four-input adders to combine the base address, segment address, shifted index register value and displacement value, and a parallel set of four-input adders to perform segment limit checks. Each unit has its own reservation station port, allowing the P6 to issue one load and one store to the data cache per clock cycle.

Memory Interface Unit

The memory interface unit works in conjunction with the store address generation unit to provide the data that is to be stored at the target address. It also has its own port on the reservation station, indicating the high priority that Intel has assigned to completing loads and stores in a single cycle. The function of this unit is somewhat unclear from Intel's descriptions, and its proportion of the die area suggests that some fairly sophisticated operations are taking place in this component.

Instruction Retirement

As the instructions are completed, they are held in the reorder buffer until they can be retired or flushed. The circular buffer guarantees that retirement occurs in strict program order, protecting the original code sequence. If they are retired, the architectural registers are updated from the instruction's temporary registers and the reorder buffer entry deleted. If flushed, the reorder buffer entry is deleted, and the temporary registers freed for reuse. The reorder buffer is also the court of last resort for instructions that were speculatively executed and later rejected; this rejection can also occur earlier in the machine (where it would conserve execution and temporary register resources). The reorder buffer can hold up to 40 uops pending retirement.

The New Instruction

- CMOV can improve performance by eliminating branches.
- CMOV may also result in improved code density.
- Programmers take note.
- Competitors take note.

Most of the working P6 instructions are direct equivalents of those found in the 386, although Intel has added a number of specialized control instructions over the last few years to accommodate new features in the 486 and Pentium, including Intel's still-mysterious "Appendix H." These

instructions would not normally be used by application programs, partly because of compatibility issues with non-Intel x86 processors; in the P6, however, Intel has introduced the CMOV instruction. This instruction permits a move to be based on a conditional result from a previous operation and eliminates the possibility of a branch misprediction by simply eliminating the branch. It uses the same set of condition codes that are available to a conditional branch instruction (including those based on floating-point and integer calculations) and can therefore replace many branch-based decision structures. This instruction is very valuable in the P6 environment where a branch misprediction can be expensive because of the long pipeline. We can expect compilers to start incorporating support for this instruction in the near future as Intel disseminates its new compiler technology.

The Pipeline and the Clock

- Pipeline is no longer clearly defined.
- Fine-grained 14-stage pipeline eliminates long delay paths, enabling a higher clock speed.
- Long pipeline implies big mispredicted branch penalty.
- Advanced branch prediction and saturated decoders compensate.

In a processor as complex as the P6, the short, inflexible pipeline of previous generations of x86 has vanished. Now, pipelining is more of a design issue than an architectural feature (see Figure 5-2). Although Intel describes a pipeline in the P6, instructions may lurk about in one slot for many cycles, share a pipeline time slot, leapfrog each other, and even disappear from the middle of the pipeline. We also suspect that store uops may result in two dispatches from the reservation station (one for the address generation unit and one for the memory interface unit), a kind of spontaneous generation of instructions. The only thing that doesn't happen is for instructions to move backward.

Intel also describes the processor as superpipelined, reflecting the 14 distinct time slots that exist in the machine (see Figure 5-2). Of these, the typical instruction will only pass through 11 slots as some sections overlap, but it may take many, many more than 11 clock cycles to do it in extreme cases.

We believe that the pipeline has become less of an architectural feature and represents the aggressive reclocking strategy that Intel has adopted to increase performance.

The P6 master clock includes a phase-locked loop (PLL) that synchronizes the clock outputs. This clock is distributed to 80 subunits across the die, and the clocks are then buffered for distribution within the subunits. Intel claims that this approach, combined with careful layout and buffer design, has resulted in a clock skew of no more than 250 picoseconds across the chip.

Coupled with the fine subdivision of tasks through the pipeline, this approach portends a design that has a far higher clock headroom than previous Intel designs. In fact, early versions of both the 486 and the 5V Pentium were plagued by execution failures if the processors were not aggressively cooled.

Figure 5-2
Intel's P6 Pipeline

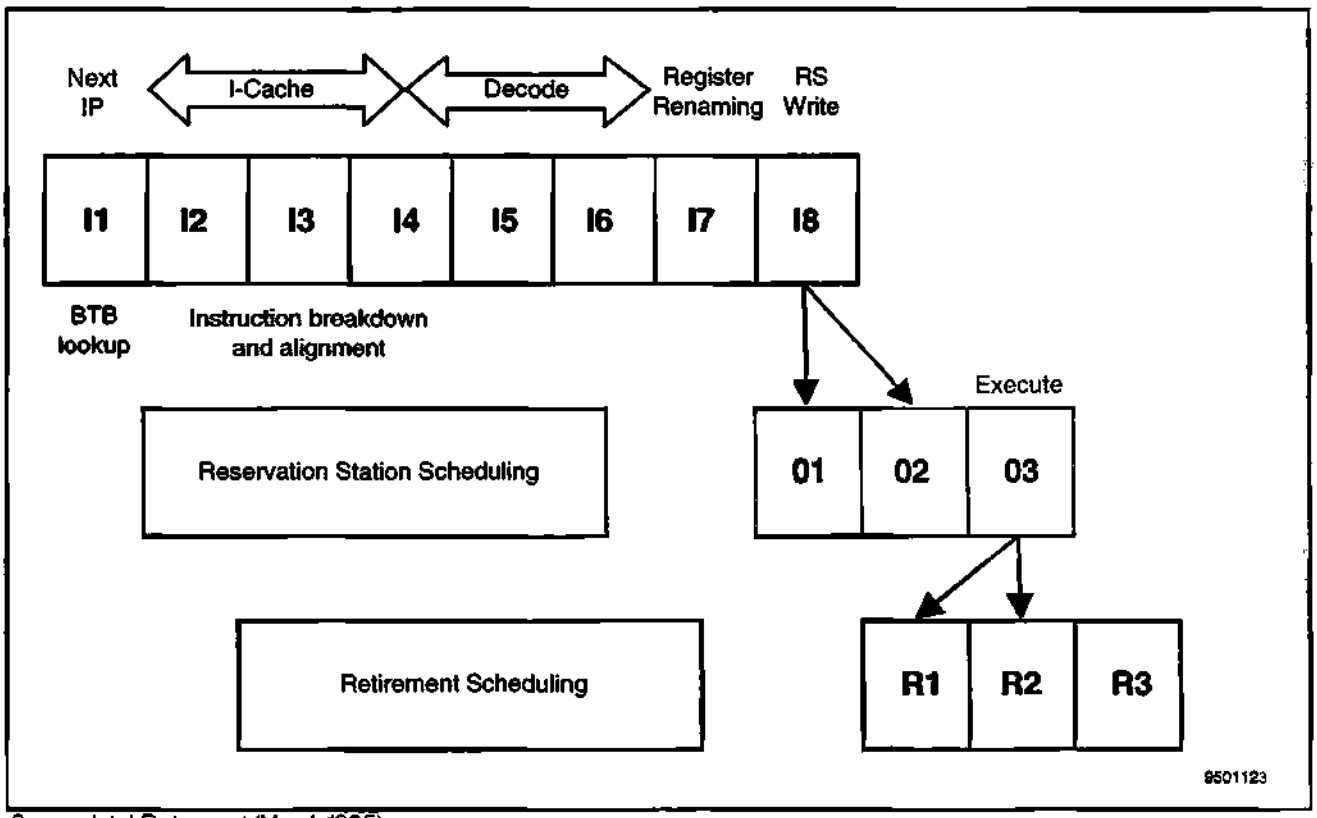

Source: Intel, Dataquest (March 1995)

With this fine-grained subdivision, however, comes a problem. If we assume that a jump is mispredicted, the processor could fill with instructions that should not be executed. In fact, the reorder buffer, reservation station, and pipeline slots together can hold as many as 49 instructions at any one time (46 uops and three x86 instructions in the decoder, more counting the instruction fetch unit). If these are derived from a mispredicted jump, it will be 11 cycles before another instruction can propagate through the pipeline and complete representing 33 pipeline bubbles. Keeping this lengthy pipeline fed with useful work becomes a major challenge that has driven Intel to the sophisticated BTAC and branch prediction techniques.

Opportunities for the Clever Coder

The deep pipeline and complex out-of-order execution engine creates an opportunity for the clever coder. It is possible to insert a nonexecuted branch in the instruction stream that results in a prediction followed by a code load from main memory. If this trick is employed well in advance of the code being needed, the code will be in the L2 cache in time to be loaded and executed. A similar opportunity applies to data; a blind load early in the instruction stream will result in the L2 cache being precharged with the data well in advance of the processor knowing that it will need it.

Writing code that performs these tricks is generally nonproductive: Few applications benefit from investment in this area, and the P6 automatically optimizes many code constructs. However, applications that run relatively small algorithms repeatedly could benefit from these techniques. Specifically, MPEG decoding and possibly MPEG-1 encoding (at higher clock rates) could be performed entirely in software and "hint" instructions inserted into the code stream could bring useful performance enhancements.

The Bus and Memory Interface Units

- These units contain Intel innovations—No details available.
- Check out the July 1992 issue of the *IBM Journal of Research and Development.*

Although the execution architecture of the machine is interesting, there is little that is new and exciting in the light of disclosures from AMD and NexGen and studies of the mainframe architectures of the 1960s. The bus and memory interface is interesting simply because of Intel's reticence to divulge any meaningful information about its operation beyond two simple statements: It holds the keys to wringing maximum performance out of the machine; and that even broad hints about these secrets would be valuable information to Intel's competition. We would direct readers to the July 1992 issue of the *IBM Journal of Research and Development* for an excellent description of a hierarchical memory interface structure in an out-of-order machine with a complex instruction set.

Intel has told us that the bus interface unit (BIU) registers occupy almost as much die area as the L1 cache; this parameter indicates that it is extremely sophisticated by virtue of being able to juggle many pending operations. The die photo and overlay (see Figure 6-2 in Chapter 6) shows this area of the chip in some detail.

Translation Lookaside Buffers (TLBs) and Caches

- L1 and L2 caches are nonblocking.
- L1 cache has three-cycle latency.
- L2 cache has four-cycle latency.

The P6 incorporates a 64-entry data TLB and a 32-entry instruction TLB. The instruction and data caches are both 8K with a 32-byte line size and are parity protected. The data cache is split into four banks and can support two accesses per cycle (provided that the accesses are to separate banks). The instruction cache is single-bank, single-port and does not support writes (as in self-modifying code)—a write causes a flush. The L1 caches take three cycles to access but can return one result per cycle thereafter. The L2 cache is 256K and also has a 32-byte line size, and incorporates ECC on a 64-bit word. It is four clocks away from the processor and can also return one result per cycle once the data is ready. Its data path is 64 bits wide, so a line fill takes four cycles. The L2 cache returns the segment of the line that includes the requested data first. Note that a miss in the L1 cache implies that the data takes six cycles to return to the processor as the accesses are managed sequentially (a miss in the L1 cache is detected before the 3-cycle sequence completes).

L1 and L2 cycles can happen in parallel, but for different fetches. The design only starts an access to the L2 cache if there is a miss in the L1 cache, so the effective data access time is seven cycles if the data is present in the L2 cache only. Intel advises us that nonblocking caches were almost as useful as simultaneous access and had the added benefit of not consuming all-important L2 cache bandwidth. The spare bandwidth is critical in multiprocessing systems, for the purpose of performing memory snoops; with simultaneous accesses, the L2 cache would be almost permanently busy.

The L1 and L2 caches are nonblocking, so the pipelining can be used to advantage—the next address can be entered while waiting for the first one to return its data.

Can CISC Catch RISC?

- The P6 is a monster chip with more core transistors than any other architecture.
- It reflects complexity of x86 on top of out-of-order execution.
- Maybe a clean-sheet CISC design could catch RISC.

The P6 is the most complicated processor design disclosed to date, reflected in its core transistor count (excluding cache) of about 4.5 million transistors. By comparison, the MIPS R10000 (with a claimed 50 percent or greater better performance) has a core of only 2.8 million transistors, and the PowerPC 604 has (after accounting for 32K of cache) an estimated 1.8 million transistors. Intel tends to quote transistor counts that run about 30 percent high by other vendors' standards, but even so, the transistor count for the performance return is very high.

As discussed above, it is possible that Intel could deliver 400 SPECint92 performance in 1996. It is likely that Intel would then be right in with the second-tier RISC players. Even so, the resources consumed to achieve this target represent an unfair advantage and would suggest that the similar resources expended on a RISC design could yield a better performance return yet.

This complexity is a consequence of the need for compatibility with code written 15 years ago, the traditional burden that mainframes must shoulder and, like the x86 architecture, the reason for their continued existence. RISC processors have the advantage of being able to shed their instruction set skins as they evolve, knowing that the system providers will recompile to take advantage of new instructions (either because they want to or, more often, own the architecture anyway).

Therefore, it is likely that the x86 architecture will continue to lag RISC performance. However, one could make the case that a brand-new CISC architecture could conceivably outperform RISC if it were designed for easy execution with: maximum instruction bandwidth through short instructions; full set of instructions designed to eliminate branching as much as possible; compiler-generated branch hints; a larger register set; relaxed-precision exceptions; no support for self-modifying code; and explicit control over program status word settings This is a laundry list of all of the architectural features known to make execution of code streams flow easily and, coincidentally, of most of the things that the x86 architecture does badly.

Why Is CISC Harder?

The size of the uops is intimately tied to the rich addressing modes of the x86 architecture and the limited register set. Intel, and also NexGen-and AMD, have to drag an instruction bus up to 118 bits wide all through the execution unit of these translation architectures. The advanced RISC architectures get away with 32 bits, plus perhaps a few assist bits. This instruction bus is large because the x86 architecture provides for rich addressing modes, permitting even memory-to-memory operations; the uops must in turn be large to incorporate all of these complex address modes. This architecture stems back to the dawn of computing when registers were invented to bypass the bottleneck of memory access, memory being where the early computers did all of their work. Until the concepts of RISC emerged (almost 30 years later), richness and complexity of instruction sets were seen as an advantage, allowing better leverage of the expensive human resource of programmers (although, in the early days, programmers were inexpensive relative to the cost of the hardware, and programming techniques reflected that).

Register availability is also an issue. The x86 architecture provides only eight registers, forcing many programs to transfer register data to main memory for temporary storage. Therefore, x86 instructions provide the ability to do memory-to-memory and memory-to register instructions, increasing the length of the microcode word because the source and destination pointers must address 32-bit memory addresses rather than perhaps 6-bit register addresses. Worse, the intermediate values have to be written to main memory. The hardware has no way of identifying these writes as to scratchpad registers; they could just as easily be intersystem communications. Therefore, these values wind through the whole chain of buffers, two levels of cache, and a full 32-byte line write to get to main memory. Paradoxically, the machine will reread them from buffers even before they hit the L1 cache, so the memory write and associated L1 and L2 cache transfers are pure overhead.

These complex instructions with their variable lengths lead to difficult decoding processes. This adds both time and silicon overhead to the decode and dispatch end of the processor.

These factors together combine to make the x86 architecture difficult to implement and costly in terms of total transistor counts.

On the plus side, the x86 architecture is far more productive for retro-type assembly programmers as fewer instructions are needed to complete a program; productivity is therefore increased. This is pretty much irrelevant in today's compiler-driven environment, although the Zoomer with its miserable 8086 core produces amazing results in conjunction with GeoWorks' assembly-coded operating system. A bigger advantage comes from the code space side, as x86 object code typically consumes less memory than an equivalent RISC program does. This compactness not only saves memory space, but increases the effective instruction bandwidth across the processor-memory interface and also increases the effective number of instructions that are stored in the instruction cache.

Test and Qualification

- Intel invested 270 man-years in design verification.
- Verification was half of total design budget.

The increased complexity of these processors brings many test and qualification issues to the market. The recent Pentium floating-point issue is a fine argument for increased product verification test, while manufacturing test time can be prohibitive without assistance from the part itself.

To address the product verification issue, Intel claims to have invested 270 engineer-years in test, a fifteenfold increase over that expended on the Pentium; in fact, Intel also claims that half of its design resources were invested in test, leading us to the conclude that the P6 took 540 engineer-years to design and suggesting an electronic design budget of around $100 million. Coupled with more rigid design techniques and advanced simulator tools that automatically detect specified exceptions, Intel believes that it has achieved a higher level of verification for the P6 than for any previous part.

The fact that the part is already running Windows indicates that Intel has been quite successful with the premanufacture verification process; the Pentium floating-point affair guarantees a substantial investment in the next phase of testing before the part is released for full production.

Performance

- We estimate 200 SPECint92 at 133 MHz.
- We estimate 190 SPECfp92 at 133 MHz
- We expect 200 MHz in the first quarter of 1996 and 266 MHz by the fourth quarter of 1996.

Intel claims 200 SPECint92 at 133 MHz. With a clock leverage of 1.5 (SPECint92 divided by clock speed), this is just at the low end of the range for a next-generation processor, but respectable nonetheless. The real challenge now is cranking up the clock speed, but the extensive pipelining should make that less of a roadblock than it has been in the past. The bus speed multipliers would suggest that Intel plans to drag this part up to 266 MHz in a future generation, perhaps with the 0.35-micron process due to enter limited production later this year. At that speed it should return 400 SPECint92, a performance level that we would not expect to see until the next generation of RISC processors gets established in 1996. Although this part can't catch Digital's Alpha, the nonshipping RISC architectures may find this device to be another embarrassing challenge to the RISC performance claim.

Per Intel, the Pentium could realize up to a 30 percent performance gain through compiler optimizations (along with 50 percent code bloat); the P6 is expected to be closer to a 15 percent gain, a consequence of the extensive dependency removals through the architecture.

We note at this time that the performance running existing x86 applications may be greater than the SPECint92 numbers suggest. The P6 architecture is apparently much better than the Pentium at extracting parallelism from the code, a function of its sophisticated decoupled execution. We suspect that Intel's Pentium SPECmarks use compiler technology

not available to users of standard packages, a suspicion borne out by *BYTE* magazine tests in which an advanced Pentium system returned only 66 SPECint92, significantly less than the 100 SPECint92 that Intel reports. Intel's demonstration of its new compiler technology took 79 bytes to complete its function as opposed to 54 bytes in the "old" way. This code inflation is undesirable for general use, so most software packages available today don't use it. If the P6 can speed up this code, it could offer three times Pentium-100 performance at its introductory clock speed.

Floating-point performance has not been addressed by Intel at this stage. The unit will be faster because of the early-out divide, and it will run benchmarks faster because of its increased parallelism, but it still only incorporates a two-bit SRT (Sweeney, Robinson, Tocher) algorithm for divide; we believe that some of the more advanced RISC processors are moving to a four-bit SRT algorithm, nearly doubling performance.

Based on the early-out divide and extra instruction resources, we would expect the processor to return between 160 and 200 SPECfp92 at its initial 133-MHz operating frequency.

Chapter 6
Manufacturing

- Intel is implementing a megafab a year.
- Capacity is coming on line at 0.35 microns.
- P6 shipments will be in hundreds of thousands in 1995, millions in 1996, and tens of millions in 1997.

Our models indicate that one of Intel's 0.6-micron megafabs (5,000 8-inch wafers per week, one every two minutes) can kick out about 4 million P6 cache and processor pairs a year, compared to a little under 18 million of the current Pentium-90 parts. When the devices switch to the 0.35-micron process (due to start its ramp in the second half of this year), the counts go to about 20 million P6 pairs or 58 million Pentiums. The gap between the two designs closes because the reduction in die size benefits the P6 more than the Pentium; as defect densities drop, the gap will close further, and if Intel brings on a second source for the cache, volumes jump even higher. Clearly, the P6 will be coming wholesale in the latter half of 1996 and is set to displace the Pentium somewhere in 1997. The initial P6 implementation is unlikely to be a major player, just as the original Pentium never really exploded into the market. However, unlike the original Pentium, the second P6 will be socket-compatible with the first-generation part so the transition will be seamless. Furthermore, the lines for the 0.35 micron part are being commissioned now for the next-generation Pentium, again pointing to relatively early adoption of the new design.

More Haste, Less Speed

Intel's previous processor generations have been a major stretch for the corporation. The first production parts have typically been at the limit of production technology, and ramps have been slow. The P6 breaks this tradition, primarily by means of a smaller relative increase in transistor count. Pentium had three times the transistor count of the 486, and the 486 represented a similar step over the 386; the P6 processor chip has less than double the transistor count of the Pentium. When the P6 starts shipping, it may only offer about 30 percent more performance than the anticipated 150-MHz Pentium; Intel has typically offered a doubling in performance as a new generation comes on line, however, and we suspect that the P6 may move along the fast track in this respect.

This conservative approach is no doubt a result of the desire to get P6 to market in 1995. When the P6 was first discussed several years ago, it was described as having 10 million transistors. Such an approach could have delivered higher levels of performance but would have set Intel's first delivery well into 1996. The 10 million transistor product would have echoed the slow start of the original Pentium processor, a situation that Intel cannot afford with its expensive new fabs to fill.

Physical Implementation

- L2 cache is a separate, large die.
- L2 cache is incorporated in dual-cavity package.
- CPU is 0.6-micron BiCMOS device, 17.5mm x 17.5mm die size.

The P6 itself sports a minuscule 16K cache, split as 8KB of data cache and 8KB of instruction cache, that returns data in two processor clocks. Intel chose to build an L2 cache on a separate die and package it with the P6 in a dual-cavity ceramic package. The L2 cache is a 256KB 6T-cell SRAM that can run at processor speeds. It is pipelined to permit one transfer per clock although it takes four processor cycles to return the first data, implying that data in the L2 cache is just one cycle further away than that in the L1 cache (although, for practical purposes, it starts its access one or two clocks later when the L1 cache has identified a miss).

The part is built in the same 0.6-micron BiCMOS process that the P54C uses (Intel process 852) and includes 5.5 million transistors on a die 17.5mm x 17.5mm. The L2 cache part consists of 16 million transistors, on a die 17.5mm x 12.2mm. The next version of the P6 will be built in Intel's new 0.35-micron process, scheduled to start production in mid-1995. Although speculation is that it will not be BiCMOS, Intel's use of the BiCMOS process at 0.6 microns was a surprise, and this could be repeated.

The Package

- Package is a dual-cavity ceramic staggered pin grid array with 387 pins.
- We believe that it can be manufactured in high volume at minimal premium.

The P6 ships in a 387-pin ceramic pin grid array that is unique in that it has dual cavities for the processor and the cache memory. Intel has included local capacitance on the processor chip itself, buried under the extensive routing channels, to improve the decoupling. As with the Pentium, the package includes a special high-dielectric constant layers, constructed from a ceramic with metallic particles dispersed through it, to provide extra decoupling within the package itself. These high-dielectric constant layers reduce overall yield on the package as they can result in electrical shorts. Although these packages never ship to the processor manufacturer, the lower yield increases cost and reduces availability. Therefore, we would expect Intel to devote some time to eliminating this feature in future revisions of the part. Figure 6-1 shows the P6 and L2 cache memory in its package. As a final note, we would not be surprised if Intel increases the pin count at some future date. Pentium systems are being shipped today with a 320-pin socket, 24 more than the Pentium needs. The signal for this will be the first P6 socket with more than 387 pins.

The P6 has provision for external decoupling capacitors, mounted on the top of the package. Intel indicated that these would not be needed, although they may become useful as the P6 clock speed cranks up.

As with the Pentium, the package incorporates the exotic copper-tungsten heat spreader to assist in cooling the chip itself. We believe that this may turn out to be an unnecessary precaution, and that it could disappear early into the product life cycle.

Figure 6-1
P6 Package Photograph

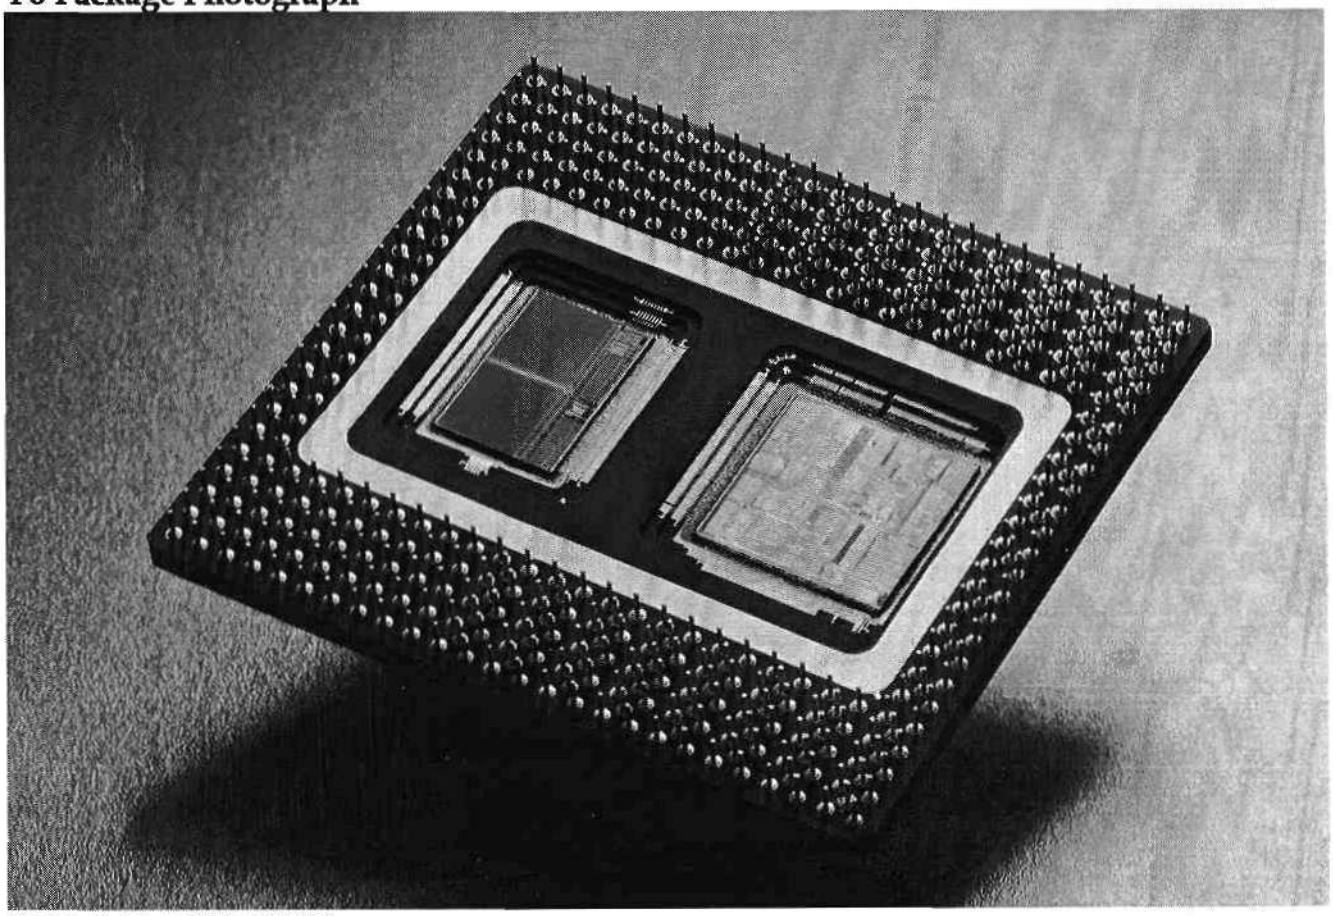

Source: Dataquest (March 1995)

Although the package has been described as a multichip module, we support the view that it is not. The opportunity to lay out the cache RAM to match up with the processor removes the need for complex routing layers within the part itself and, assuming that the package itself does not suffer from any problems as a result of the dual-cavity design, should not be significantly more expensive than a standard PGA to build.

The pending introduction of the P6 brings an interesting problem to the fore. The PC market dwarfs most other applications of ceramic packaging, and the newer processors have electrical and thermal properties that make ceramic packages almost mandatory. Furthermore, ceramic packages provide system manufacturers with a level of flexibility that results in reduced costs as the system can be configured close to the end user, eliminating the high carrying cost of an advanced processor. With the PC market growing, the percentage of processors packaged in ceramic growing, and the area of the package itself growing the demand for ceramic packages could put considerable pressure on the supply side of the equation. This supply-side pressure could make it difficult for alternative processor vendors to obtain packages, leading to yet another competitive advantage for Intel through its virtual buying oligopoly, a mirror image of the near-monopoly it enjoys on the processor supply side.

Die Layout and Resource Allocation

Figure 6-2 shows the die photograph as provided by Intel, along with an overlay identifying the functional units of the processor.

Figure 6-2
P6 Die Photograph and Functional Overlay

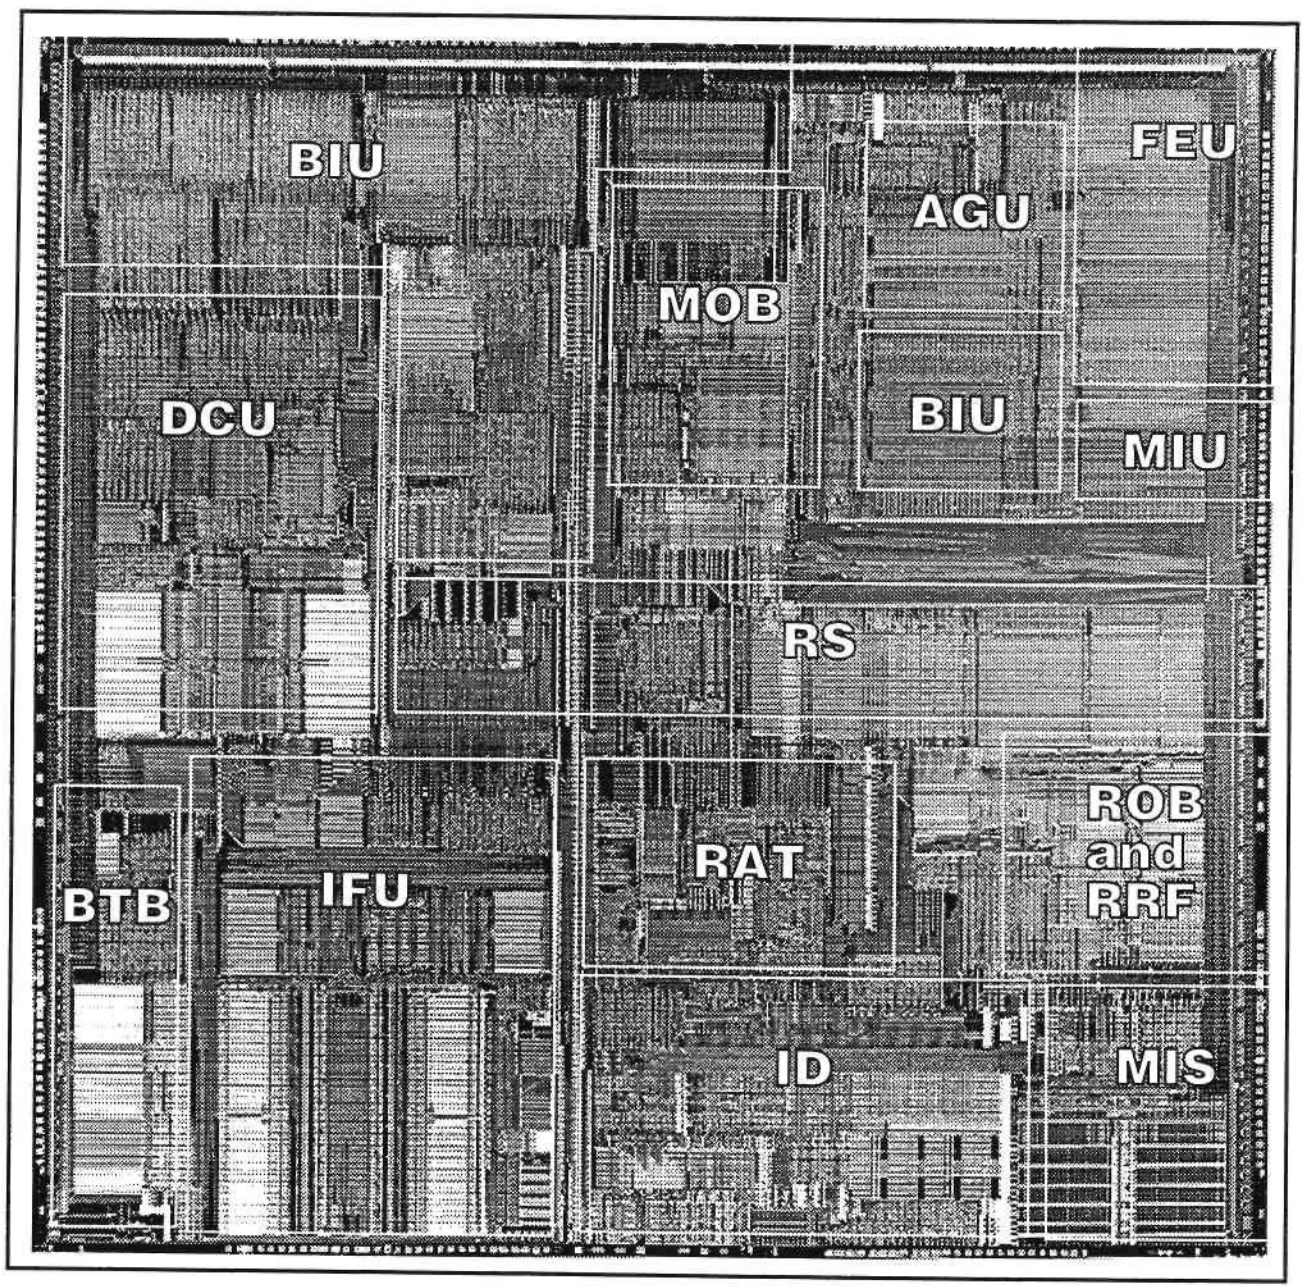

Source: Intel, Dataquest (March 1995)

Consumption of Die Area by Functional Units

We have used this information to identify the area associated with each functional unit, shown in Table 6-1. Intel's overlay is approximate, and the four metal layers make it difficult to clearly identify structures, but we have made appropriate assumptions where possible.

Table 6-1
Die Area Analysis of the P6

Unit ID	Description	Die Area (%)
MIS	Microcode Instruction Sequencer	4
IFU	Instruction Fetch Unit	12
RAT	Register Alias Table	4
ID	Instruction Decoder	9
ROB	ReOrder Buffer	5
BTB	Branch Target Buffer	4
RS	Reservation Station	8
DCU	Data Cache Unit	9
BIU	Bus Interface Unit	16
MOB	Memory Ordering Buffer	5
AGU	Address Generation Unit	3
IEU	Integer Execution Unit	2
FEU	Floating-point Execution Unit	5
MIU	Memory Interface Unit	1
Others	Routing, clock distribution etc.	11

Source: Dataquest (March 1995)

The first thing to note is that the execution units (FEU, AGU, IEU, and MIU) only occupy 11 percent of the total die area, underlining the complexity of this out-of-order architecture. The Instruction Fetch Unit and Data Cache Unit both incorporate 8K of cache. Based on expected cache transistor densities (in turn derived from the L2 cache parameters), we estimate that the caches occupy approximately 2.5 percent of the total die area apiece, barely more than the 4 percent of the die that Intel has allocated for test. Clearly, this is an area that Intel could enhance to maintain relative performance as the clock rate increases. The largest component of the device is the BIU. As mentioned earlier, this is where Intel believes that it has some of the most important innovations in the P6 architecture, and the engineers aren't talking.

The Cost of the P6

- The cost information is based on Dataquest processor model.
- It provides reference costs for P6 and its competitors.

The cost of a microprocessor has several major components. The cost of the silicon wafer, its processing, test, and packaging all contribute to the incremental cost. However, depreciation of the capital plant and equipment is a major contributor to the overall cost of the product.

Modeling Desktop Processor

To predict the likely cost, performance, and availability of future systems, we need to make projections about the processors that go in them. The model described here incorporates tools to predict cost, line yields, power consumption, and die size. Although not exposed in detail here, we can discuss specific questions on a case-by-case basis.

The data in Table 6-2 comes from several sources: manufacturer data, where available; Dataquest estimates; and formula-based models. In particular, defect densities represent judgments from the program analysts based on discussion with manufacturers and, in the case of Intel's parts, reconciliation of production to fab capacity.

We make the following assumptions:

- Cost of an 8-inch 0.6-micron wafer—About $2,000
- Cost of an 8-inch 0.35-micron wafer—About $3,000
- Cost of a Pentium package—About $25
- Cost of a P6 dual-cavity package—About $50

We use a Murphy model to predict overall yields. Die sizes are as published by the manufacturer or our estimates based on geometries and transistor densities.

The Power Model

The power model is under development. There are many factors that affect the power consumption of a device, many of which cannot be accounted for by simply looking at the characteristics of the die. The basic formula for power dissipation is:

$$\text{Power} \propto 1/2\, CV^2f$$

where C is the effective capacitance of the device, including junction and wiring capacitance, V is the supply voltage, and f is the operating frequency. C is difficult to determine because large areas of the die may remain dormant if not in use; processor designers have discovered that it is necessary to shut down the clocks to subsystems whenever possible to reduce power consumption, as lower power consumption results in less heat generated, which, in turn, leads to a higher operating frequency. We have found empirically that C appears to relate to the area of the die less the area of the cache, and the power model numbers reflect this approach. We will be refining this model over time as we accumulate more data. Note that the initial results are fitted to published power specifications, so they probably follow the trend of overstating actual power consumption by up to 100 percent.

Table 6-2 shows a variety of calculations that are of value in understanding the P6 and its market position.

Table 6-2
P6 and Its Competition

		Die Dimensions	Performance				Defects and Yields	8" Wafer	
Name	Comment	Area (sq. mm)	Clock Speed (MHz)	SPECint per MHz	Disclosed Power (W)	Power Model (W)	Defect Density	Annual Yield 5000 Start Line	Packaged Cost ($)
PowerPC 604		196	100	1.6	13	8.6	0.6	10.3	80
R10000		298	200	1.5	30	32.2	0.6	4.1	190
Alpha 21164		314	300	1.1	50	42.6	0.6	3.6	205
UltraSPARC		315	167	1.6	30	34.7	0.5	4.5	183
PM1	SPARC V9	296	153	1.7	60	37.3	0.5	4.8	256
P6		306	133	1.5	20	17.6	0.4	5.9	163
P6 L2 Cache	No package	214		0		0	0.3	13.7	44
P6S	DQ estimate	120	266	1.5		22.0	0.5	28.4	88
P6S L2 Cache	No package; DQ estimate	83		0		0	0.3	54.9	19
P6S L2 Cache 512K	DQ estimate	167		0		0	0.3	20.5	49
Pentium		249	66	1.0	16	15.5	0.4	9.4	75
Pentium (P54C)		164	100	1.0	10	7.4	0.5	17.3	71
Shrunk P54C	DQ estimate	148	100	1.0		6.6	0.4	23.7	61
Pentium-35	DQ estimate	81	200	1.0		6.8	0.4	58.1	51
NX586		196	87	0		11.3	0.6	10.4	82
NX586	DQ estimate	144	100	0		9.5	0.6	19.0	61
486DX4		77	100	0		3.6	0.45	60.5	28
AMD 486-100		44	100	0		2.9	0.5	121.0	23
M1		380	100	1.3	10	15.7	0.6	2.2	200
K5	DQ estimate	201	100	1.3		13.2	0.4	14.0	74
K5S	DQ estimate	109	150	1.3		8.8	0.4	38.1	56
P6 .35 Integrated	DQ estimate	202	267	1.5		16.3	0.4	13.9	124
P6 .35 Integrated 512K	DQ estimate	286	267	1.5		13.2	0.4	7.1	185

Source : Dataquest (March 1995)

Integrating the Cache

- The separate cache makes sense for Intel.
- The case is stronger for other manufacturers.
- Notebook parts may be the exception.

The use of a separate cache chip in the same package as the processor, closely coupled to the processor itself, is a radical change for a mainstream processor. However, once the packaging issues are dealt with, there are some clear advantages to this strategy. The cost of the P6 is the sum of the processor and the cache costs. In the model, the processor cost includes the cost of test and packaging; some overhead is also applied against the cache part, which inflates the cost somewhat.

From the model, we see the following costs:

- Initial P6 with 256K L2 cache—$206
- Initial P6 with 512K L2 cache—$212
- Shrunk P6 with shrunk 256K L2 cache—$111
- Shrunk P6 with 512K cache—$140
- Integrated P6 with 256K cache—$125
- Integrated P6 with 512K cache—$184

The manufacturing cost of the part drops dramatically with the shrink; this is a consequence of a relatively small die size, in turn a consequence of the split cache. Integrating the cache, using the numbers in this model, looks like a push; the cost goes from $111 to $125. Improved yields on the dual-cavity package (we assume a 4 percent loss over a single-cavity package here), or a better defect density could push the equation either way. The single chip is unlikely to offer better performance as the cache will probably still require three cycles to access its data, a consequence of the physical size of the cache. However, the 256K version would provide a viable P6 solution for notebooks. The 512K version remains best implemented as two die. Therefore, we may see a dual-pronged strategy to support both low-end and high-end products.

Note, though, that the larger integrated die reduces production capacity by about 30 percent if Intel makes the SRAM, 50 percent if it is procured from outside. Therefore, if Intel could acquire a cache device from an outside source, its production capacity would effectively be increased. This approach could be very feasible if the supplier were able to meet Intel's requirements with a 4 T-cell memory as the cost would be low enough to make it profitable for both parties. This is another reason why we may see a special notebook part.

We believe that the split cache will set a trend that all processor manufacturers will follow. Once the dual-cavity package bugs are worked out, this approach will be an effective one for delivering local, large, fast L2 caches. The advantage is even greater for smaller manufacturers that do not have Intel's volumes to drive yields up.

This technique could also impact manufacturers of fast cache memory; as P6 ramps up, there is little place for another cache in a P6 system as the 256K L2 cache is so large and tightly integrated into the system. Improving system performance with an external cache is a challenge that probably cannot be met in a mass-market desktop system, so the P6 cache memory will squeeze the fast SRAM manufacturers out of the desktop; other processor manufacturers will probably follow Intel's lead, so the situation will worsen. However, the opportunity exists to provide large cache memories to Intel and other processor manufacturers.

In summary, Intel's innovation here is a subject worthy of intensive study. There are many strategic advantages that accrue from this approach, including cost, yield, production capacity, configuration flexibility, and reduced design cycles.

A Separate Cache Chip?

As this report is going to press, we hear rumors around the industry of a possible version of the P6 that supports a custom-designed external L2 cache memory. Such a device would have to interface at speeds of up to 266 MHz if it were to meet the needs of all expected versions of the P6, and the P6 would have to add up to 100 pins to support the 64-bit data bus, the ECC bits, the address bus, and the control bits. Furthermore, the physical implementation will have to be carefully controlled to ensure that the final product is reliable. However, the strategy is credible although our analysis indicates that the dual-cavity solution is a reasonable one; perhaps the biggest obstacle to acceptance is the package pinout change.

- We estimate $999 at introduction for 256K part and $1,499 for 512K part.
- Price will decline rapidly (30 percent to 40 percent a year).
- Price will decline asymptotically to $350 to $400, $100 more than Pentium and $200 more than 486.

Intel declined to discuss price on the product at this stage. However, there are a couple of factors that differentiate the P6 from previous introductions. Unlike the Pentium-486 transition, the per-wafer revenue of the P6 is likely to compare favorably to that of the Pentium, especially in the 0.35-micron process. Secondly, the clean multiprocessing interface lends itself to industrial-strength x86 computing in a Windows NT or other network server environment. Accordingly, P6-based servers are likely to be extremely attractive to Windows NT-based environments because they will offer a significant performance boost over existing Pentium-based systems. As these servers can increase the productivity of many users, they command prices of up to the low hundreds of thousands of dollars; a $1,500 processor can easily be lost in the overall economics of such a system. Historically, the gray market has also bid up the prices of Intel processors at introduction to over Intel's book price. Therefore, we would not be surprised to see Intel introduce the part at up to $1,500 list price, although we would also predict a very rapid price decline curve. It may well be that the 512KB cache part occupies this premium spot, while the 256KB part is priced at the traditional $999 for desktop use. We also note that, although the price decline may be swift, it is likely to level off in the high $400 range. Intel's strategy has to be to continue to increase the average selling price of its processors, as market growth and market share growth do not provide sufficient incremental revenue alone.

Chapter 8
The System Interface

- The 64-bit burst mode bus peaks at 528 MB/sec.
- Split-transaction protocol supports up to eight outstanding operations.
- Bus design is optimized for glueless four-way multiprocessor support.
- Electrical protocol is a modified GTL (GTL+) interface.

The bus is still a 64-bit bus, but now includes support for full multiprocessing. The architecture supports split transactions and supports up to four parallel processors without additional glue logic.

Intel chose to switch to a GTL variant running at 1.5V (rather than the standard GTL 1.2V interface) at 66 MHz, described as GTL+. The switch was inspired by the desire to put up to four processors and several bus and memory controllers on the processor bus, a requirement that dramatically increases the bus loading. The bus is tightly specified and requires termination at each end.

The CPU core itself runs at 2x, 3x, or 4x the bus clock. Therefore, the initial devices have a 66-MHz bus clock and a 133-MHz core clock; we do not believe that the device is intended to run with a bus clock of less than 66 MHz. This ratio set implies that we will see 200-MHz and 267-MHz core clock versions as the product matures. The bus, with its tightly specified GTL interface, could run faster in later revisions of the product; this change would improve the maximum I/O bandwidth of the part, bringing it closer to the 1 GB/sec rates of the advanced RISC competitors.

The bus supports a MESI protocol, as well as the ability to stream data between processors without first writing to memory. Bandwidth is extremely high; the maximum burst rate is 528 MB/sec and could reasonably be achieved for 80 percent of the time—the P6 transfers data as one address cycle followed by four data transfers of eight bytes each. The data transfer portion can exceed four cycles if the processor transfers several consecutive 32-byte lines from the cache, and processors in a multiprocessor configuration can transfer data directly from cache to cache. In this last case, the memory controller is responsible for "snarfing" the data as it goes past and can slow the transaction down if it has to. Snooping and synchronization is taken care of in parallel with address output, so there is no dead time on the bus while other units complete their snoops (a problem faced by NexGen).

A single-byte write to memory or I/O, on the other hand, is awful; it still requires that the processor go through the whole five-cycle access. Fortunately, most of these accesses will be bundled up with other data before the write occurs, but, for example, programs that write directly to video memory (which frequently have a byte-wise pattern) may not extract the full potential of the P6.

The bus also supports split transactions: Up to eight memory read or write requests can be outstanding on the bus at any one time.

This bus is very advanced, but the gigabyte rates of the PowerPC 620 and MIPS R10000 products still have the edge. The P6 may therefore fall just short of being fast enough to penetrate the minicomputer world, but it is going to tear up the high-end server market.

The 64 bus data bits are protected with full ECC; the 36 address bits and the control bus are parity-protected. Intel negates many of the arguments for ECC on the address bus by recycling if the bus signals an address parity error, guaranteeing that the system will not shut down for a transient fault. However, the risk of a dual-bit error is still measurable and this exposure may well keep the P6 out of the higher end of the server market; witness NEC's recent decision to select PA-RISC systems over its own MIPS designs for mainframe-class servers.

The P6 and Rambus

Rambus remains one of our favorite next-generation memory interfaces. We are attracted by many of its key features.

- Broad sourcing including Toshiba, NEC, Goldstar, Samsung, Fujitsu, and OKI (and Hitachi for the ASIC interface)
- Cookbook design solves 250-MHz problem.
- Few signal pins
- 500 MB/sec transfer rate in a burst environment
- Excellent expansion granularity

These characteristics are exceptionally valuable in the cacheless P6 environment where burst transfer rate is critical to performance. The on-chip cache covers up the DRAM subsystem latency, allowing the system to make the most of the 500-MB/sec burst transfer rate. It is easy to envision a P6 memory controller with four Rambus control ports, permitting the implementation of a four-bank DRAM subsystem that can be expanded 8MB at a time. We recommend that core logic and system designers evaluate the use of Rambus parts in this application for implementation when Rambus DRAMs provide an economically viable solution.

Rambus will have its first high profile wins in 1995, with the Nintendo Ultra 64 video game and the anticipated video controllers from Cirrus Logic. SGI, through its Nintendo relationship, could conceivably end up using the Rambus design in its workstation video subsystems.

However, true high volume will only come from a success in the PC industry, which will require either a major system manufacturer settling on Rambus as the main memory interface or for Intel to integrate the Rambus interface directly onto the CPU. Although this is unlikely in a desktop environment, it may well work in the notebook environment.

Chapter 9
The Competition

The x86 Competitors

In the x86 section of the market, competing with Intel is more of a function of manufacturing capacity and die cost than of raw performance. Although the big money is made at the middle and high end of the market, the huge base market is very lucrative as long as the supplier can deliver. We note, though, that Intel has always managed to avoid the silver bullets from AMD and Cyrix in the past, a situation that we expect to continue.

The 486

- The 486 has low cost, good performance.
- It is still a contender in emerging economies.
- It will remain strong through 1995 at 100 MHz.

Despite the massive incursion of Pentium and Pentium-alikes, we expect the 486 market to remain strong and last longer than the 386 market did. This situation arises because the 486 is vastly more useful than the 386 was in the last year of its life, a function of its relatively high performance. Coupled with plenty of manufacturing capacity from Intel and the cloners, the outlook is positive for this part and the low end of the market. The signs are particularly strong for emerging economies and the European home market, although we would caution manufacturers that the 486/100—which may dry up a little as AMD shrinks its K5 part—is really the place to be late this year.

The Intel Pentium

- The Pentium foils competitors.
- The Pentium is the volume part of the profit equation.
- Intel balances Pentium against P6 for profitability.

The Pentium is a formidable competitor for the P6, simply because it will fill fab space at excellent margins. However, the shrunk P6 parts and Intel's monster fab investments will combine to relegate the Pentium downmarket; somewhere in the second half of 1997, P6 system production will exceed that of Pentium systems. The Pentium becomes more of a foil for the products of the competition, positioning them downmarket and hollowing out the margin structure at the low end of the market as the profit equation moves from technology to volume.

The Cyrix M1

- The M1 is a brilliant design but late on delivery.
- Window of opportunity is closing.
- It must show 100-MHz+ clock speeds to succeed.
- Possible opportunity exists for technology in notebooks.

The M1 architecture now looks somewhat jaded. Its dual-pipeline design, despite the sophisticated register renaming and interlocking, does not reflect the new architectural approach embodied in the latest RISC machines and the competing superscalar x86 architectures. Furthermore, it retains the dubious title of the largest x86 processor ever built, by our estimates yielding less than two million devices a year from a megafab. In desperate need of a shrink and with an undisclosed clock rate, M1 has to move fast to catch up with the Pentium. To this end, the relationship with IBM should enable Cyrix to crack its implementation problem and deliver Pentium-class product in time to ride the coattails of the market to success. The window of opportunity for glory, however, has closed, and Cyrix must now demonstrate its ability to participate in a mature x86 market. We would still like to see a 100-MHz 486 from this company, but clever exploitation of buses and cores (of which the "Chili" product is a hopeful sign) may enable Cyrix to repeat its earlier success in notebook products by providing notebook system manufacturers with a no-risk upgrade path.

The model reflects the huge die size of the M1 in its $195 manufacturing cost estimate. Curiously, our power estimate suggests that the M1 is closer to 15W than Cyrix's published 10W (which would be consistent with a 66-MHz clock speed). This could signal thermal problems ahead for Cyrix.

The AMD K5

- The K5 will not be a player until the die shrink in early 1996.
- It is a fine Pentium challenger.
- It may face clock speed limitations.
- It can be extended to challenge P6 in 1996/1997.
- There are excellent revenue prospects for AMD, but they are not a threat to Intel.

The similarities between the P6 and the K5 are obvious, although the performance claims underline the fact that K5 has a somewhat less sophisticated design. Mostly, the P6 has the advantage of scale over AMD's device: There are far more places to keep instructions, and the force-fed front end coupled with the monster cache permits a much higher clock speed without running the instruction queue dry. AMD's short pipeline requires that the processor complete many tasks in some cycles, a likely limited factor to the maximum clock speed of the device. The Tomasulo-style distributed reservation stations, with their small, sequential queues, impose a little more order on the execution engine. Again, this could translate to a slightly reduced relative performance, one of those rare examples where less order loses to relative chaos. These deficiencies, however, will undoubtedly be remedied in the K6, which AMD will claim to be the true P6 competitor.

Ultimately, AMD's greatest challenge remains Intel's capital spending plan: We estimate that at any time Intel has approximately five times AMD's manufacturing capacity. Coupled with Intel's ability to always creep past the next x86 killer chip's performance before it ships in meaningful quantities, AMD seems to be doomed to taking its share out of the low end of the market. Of course, we should all be so lucky as to face such a fate. AMD will be able to count its revenue in the billions of dollars even with these crumbs from Intel's table.

AMD's initial K5 part looks to have a reasonable cost ($74) and yield (14 million units a year from a megafab). However, capacity-constrained AMD can get nine times as many 486/100 parts (121 million units a year at $23 each) from the same fab. If the K5 sells for $500, it makes an estimated per-unit gross profit of $426, a splendid gross margin of 85 percent. The 486/100 is selling at $75, with a smaller gross margin of about $52 per unit or 69 percent. However, AMD can make nine times as many 486/100s as K5s, so the $426 profit from one K5 has to compete with the $468 profit from nine 486s. Therefore, it does not make economic sense for AMD to sell K5s over 486s unless the 486 price drops below $75. This is not likely to happen in 1995; once the K5 is shrunk, however, the cost turnover point jumps to $120 with the K5 selling at $300. Clearly, AMD will abandon the 486/100 as soon as the shrink is completed. The shrink brings another benefit: Power consumption should drop from about 13W to 9W, even at a higher clock speed.

The NexGen NX586

- The NX 586 has great performance and die size.
- It is a shame about the bus.
- Separate cache die and P54C pinout could advance to "Go."

NexGen pioneered the architectural approach and multiprocessor bus structure that we see implemented on the P6. Time pressure, however, reduced the multiprocessor bus functions to a vestige and limited the machine to a single x86 instruction dispatch per cycle. Even so, it achieves performance that matches that of today's dual-issue Pentium, and its decoupled architecture offers the promise of a faster second-generation part that could compete with the P6. Again, IBM's manufacturing skills will pay off for NexGen. Of all the pretenders, NexGen alone has the die size that looks like it belongs to a high-volume product. NexGen faces dual challenges: One is to implement a bus that system manufacturers can use; the other is to achieve a product cycle closer to nine months rather than nine years.

Of course, we would also like to see some level of integrated floating-point capability; this feature is rumored to be en route to the market.

The cost model shows NexGen sitting right at Pentium cost, power, and performance levels; a proper accounting would have to include IBM's cut of the deal, but NexGen is clearly a contender.

The RISC Competitors

- Competition is defined by straight performance; manufacturing capacity is less important.
- The second-tier RISC players are under threat from Intel.

Unlike the x86 market, Intel is not a leader in the desktop RISC market, and competing is more of a matter of keeping Intel's camel out of the relatively small RISC tent. The camel's nose is always under the flap, but the RISC manufacturers generally outperform Intel's parts by a decent margin and deliver faster processors earlier than Intel can. Competition, therefore, is not defined by cost and capacity but by straight performance.

Dividing the SPECint92 performance figure by the clock rate gives a measure of how effectively the processor extracts parallelism from the code that it is executing. Figure 9-1 shows the clock leverage for the P6 and other key processors. It is important to remember that clock leverage without the ability to clock fast is worthless, resulting in a slow, expensive processor rather than a slow, cheap processor. The two also tend to go hand-in-hand – a complex processor may well result in a drastically reduced clock rate. This is the problem that has cursed the SuperSPARC for the last couple of years.

Clearly, the P6 leads the x86 pack by a comfortable margin. It is also very competitive with the next generation of RISC processors, trailing slightly but within the target zone.

However, to keep up with these machines the clock speed must also be competitive.

The clock speed of a processor is a moving target. Figure 9-2 shows the clock speed of the processors, still sorted by clock leverage, as targeted when volume shipments begin.

Figure 9-1
Clock Leverage of Advanced Processors

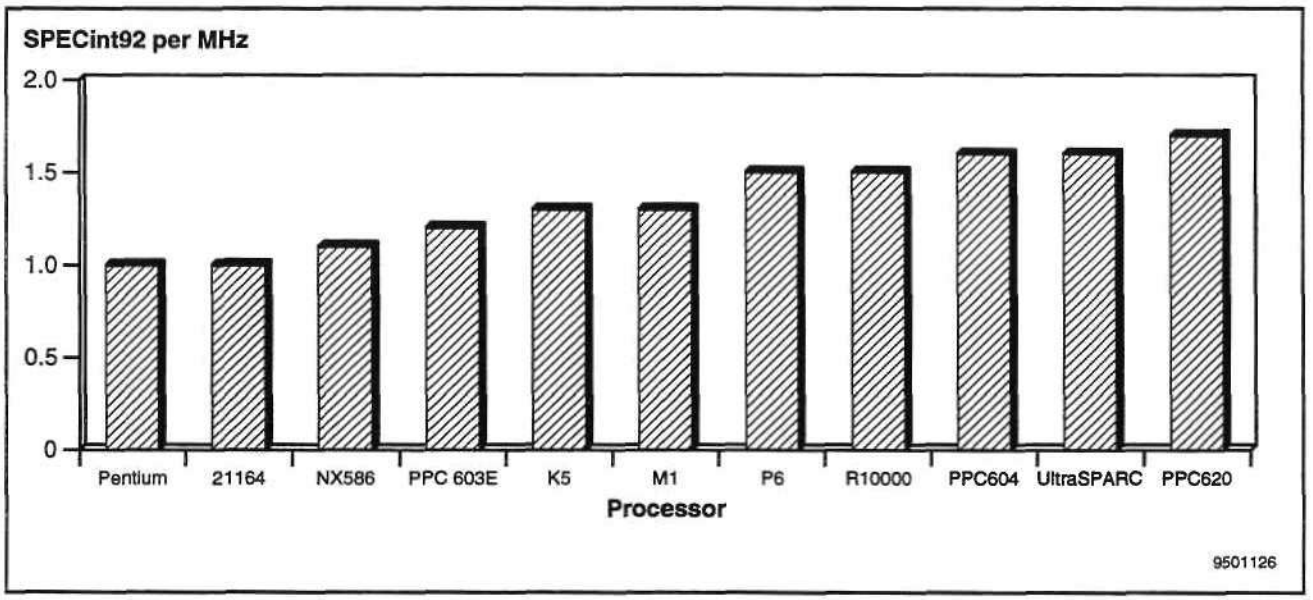

Source: Dataquest (March 1995)

Figure 9-2
Target Clock Speeds

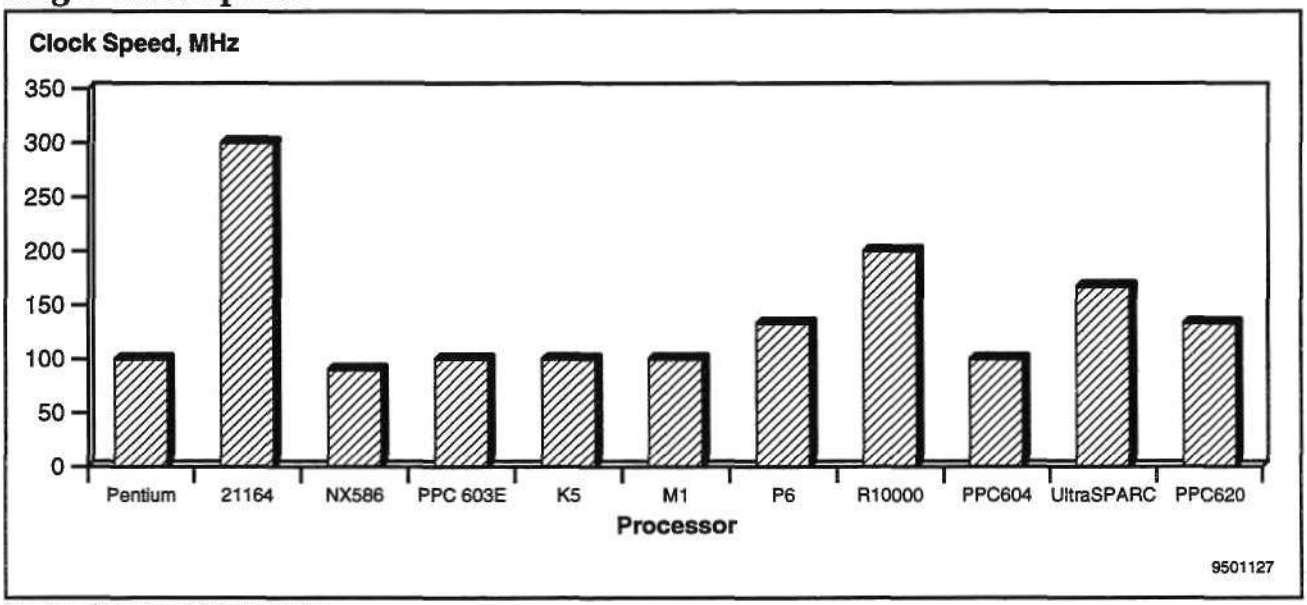

Source: Dataquest (March 1995)

Alpha

- Alpha is a brilliant architecture, and implementation is fast, fast, fast.
- It could maintain its performance lead.
- It is a possible basis for third-architecture partnership.

The Alpha 21164 stands alone. Already shipping at 300 MHz, this processor is likely to hit 500 MHz by the end of 1995. If Digital can tweak the architecture to extract just a little more parallelism, Alpha is likely to maintain its 1.5x-to-2x performance advantage over all comers. The wild card is Hewlett-Packard. HP has been able to take advantage of its compiler skills to increase parallelism in even relatively simple architectures and tends to show the second-highest clock rates. The PA-8000 could approach Alpha performance more closely than any other competitor's product, although its high operating frequency and external cache will make PA-8000 systems expensive to build. For the performance at any cost market typified by HP's workstation and systems products, however, this is not an issue. It is interesting to observe that the highest-performance products come from companies that consume almost all of their own production, a result of the ability to tune the silicon to the system design. Putting it another way, the more OEM system manufacturers there are for an architecture, the easier the system design has to be, resulting in compromises in system performance. Witness Intel's problems with just getting system designers to blow air over the processor.

We also note that Digital would make a fine partner for AMD and Compaq in establishing a third architecture. Alpha's performance, coupled with Compaq's system volumes and AMD's marketing, design, and manufacturing strengths, would be a formidable team.

PowerPC

- PowerPC has great performance.
- Early silicon, late systems weaken advantage.
- Where is the marketing?

The PowerPC parts are also sleepers; the PowerPC 620 is going to be troubled by competition from its less-expensive sibling, the 604, which will no doubt progress to 133 MHz before any significant number of systems ship with it. There are a couple of caveats in the announced PowerPC 620 numbers, however. The numbers assume unrecompiled code, which does not take advantage of the 64-bit features added throughout the 620; nor does it attempt to take advantage of the increased parallelism in the machine. The IBM/Motorola partnership is also conservative with its initial performance numbers; the P6 announcement prompted a hurried disclosure that the PowerPC 604 would see 150 MHz late this year, a speed boost that puts it on top of the P6 just as Intel-based systems start shipping. Unfortunately, this speed boost also helps it to outperform the PowerPC 620, a grotesquely expensive part with a performance disadvantage. The PowerPC 620 gets a boost for the fall 1996 models that will put it up to 300 SPECint, but this number would fall behind the expected performance of a 267-MHz P6. At these speeds, the huge cache of the P6 may prove to be an insurmountable advantage for Intel over the PowerPC 620 in any kind of volume system.

PowerPC's Lesson in Marketing

There is an interesting lesson in marketing here. The PowerPC family, when first disclosed in 1993, was positioned as always being one step ahead of Intel. The PowerPC 601 was going to be a Pentium-class box in a 486 market. Then, Apple introduced the Power Macintosh 6100 as the 486 competitor. Between emulation and Apple's cheapskate product managers (buried video, no L2 cache), the best-selling PowerPC system was not much better than the comparably priced 486 products shipping at the time. The lesson is learnt; L2 cache is now standard, and Apple's desire to differentiate pricing through performance appears to be waning somewhat. However, it may be too late. PowerPC 604 systems have not appeared yet, and the PowerPC 604 is likely to take much longer to hit the market in force than the P6 will. Worse, the P6 will probably ship with a performance advantage.

The end result of this is that the PowerPC family has been quietly repositioned by market forces so that it drops through the cracks in the x86 marketplace and competes with itself rather than x86.

The following is a brief summary of the positioning problems that the PowerPC partners have:

- The 603e outperforms today's Pentium but isn't shipping yet. It will be the mainstay of Apple's desktop product line late in the year but will by then face the challenge of faster Pentium chips.
- The 604 is faster than the Pentium but lags the Pentium by at least a year in terms of cost and availability. It should beat the P6 to market but is likely to not perform as well.

- The hot, new PowerPC 604 will match P6 but will face competition from faster P6 parts in mid-1996. Worse, the faster 604 displaces the 620 from any desktop opportunity.

These are not insurmountable problems. The root causes are delay in shipping systems to market, underperforming hardware (differentiation through performance), and insufficient attention to marketing. We believe that these problems can be overcome (licensing helps), and that the PowerPC will become a major force in the market.

UltraSPARC

- This one had better meet Sun's promises.

The UltraSPARC is ahead of the P6; it has a similar clock leverage but is going to clock about 25 percent faster than the P6 for a corresponding performance advantage. Sun has two good reasons to stay ahead of the P6: it has supported the notion that CISC can be faster than RISC for the last couple of years with the atrocious performance of the SuperSPARC (the only major RISC processor slower than Pentium); and the market will not tolerate another cycle of laggard performance, particularly if the P6 catches up in clock speed.

As indicated elsewhere, the fine-grained timing and the large 256K cache (and a 512K option coming!) in the P6 indicates that Intel has designed this part to run much faster than initial information indicates. On a final note, the UltraSPARC includes floating-point performance a generation ahead of the P6 with its multiply-accumulate and multiple execution units; this is a requirement for workstations, less necessary for mainstream desktop computing.

R10000

- R10000 has elegant design and strong implementation.
- Open RISC strategy seems to be working at the moment.
- Risks are typecast in just a few shows.

The R10000 targets a fast clock speed of 200 MHz and 1.5 SPECints per megahertz, jockeying for a leadership position in the second tier behind Digital. Like the UltraSPARC, it also has very advanced floating-point capabilities. Like the P6 and the PowerPC 620, it is also targeted at mainstream transaction-processing applications. Given MIPS' track record of performance and the support of its skilled fab partners, MIPS is likely to make good on its promises and deliver this high-performance engine. However, with Digital's Alpha becoming the preferred Windows NT performance engine and PowerPC set to be the volume Windows NT alternative, MIPS may be confined to its core clientele of workstation manufacturers, refugees from the vanished minicomputer market, embedded controllers, and Nintendo.

Chapter 10
The IBM 700 Mainframe

Figure 10-1 is a representation of the architecture of IBM's current high-end processor, the model 700 processor. Introduced in 1993 and based on a design originally introduced in 1990 (the 110-MHz model 520), the model 700 processor is the core of mainframes that supports up to 10 processors in a multiprocessor configuration.

Hardware Implementation

With a $20 million price tag and minuscule volumes by microprocessor standards, mainframes trade off higher unit costs against lower design investment. Therefore, these processors are implemented using multichip modules and highly customized gate arrays produced using direct-write e-beam technology. The processor PCB carries four multichip modules approximately five inches square, each incorporating a major processor subsystem. Each of these multichip modules (IBM calls them thermal conduction modules, or TCMs) carries about 100 80,000-transistor bipolar gate arrays that represent the functional logic. The chips are bonded face-down using IBM's C-4 connection process. Therefore, the processor core includes about 30 million transistors; even though the sea-of-gates construction renders many of them useless, the end result is a machine somewhat more complex than the P6 in terms of its raw transistor count. The cache memories are implemented using 2.5ns 8KB ECL SRAMs, packed 128K at a time onto two of these modules.

The operating characteristics of this machine make the P6 appear as no more than a punctuation mark; the four modules that make up the processor, operating at a little less than 2V, look for a supply current approaching 4,000 amps. This is more than 500 times the current drawn by the P6.

The TCMs used in this machine have 27 layers of interconnect, and performance is limited by intermodule propagation delays. The board design uses a half-clock phase shift to synchronize signals that cross between modules. Mainframe designers have fought clock speed limitations and the processor-to-memory time warp for nearly 40 years, and it is this richness of interconnect that permits the complex architectural features that we see here. Note that the enabling technology for the new generation single-chip microprocessors is extra levels of interconnect, an interesting parallel.

The TCMs also bring us a lesson in economics—there is no way that the TCM technology could be commercially viable as a product for sale outside of IBM. However, without it, the mainframe could not be manufactured. The lesson is that a captive manufacturing technology may not have to succeed in the free market to be valuable, and IBM and Digital are both overcoming this issue in their semiconductor manufacturing operations.

Figure 10-1
IBM 700 Mainframe Processor Architecture

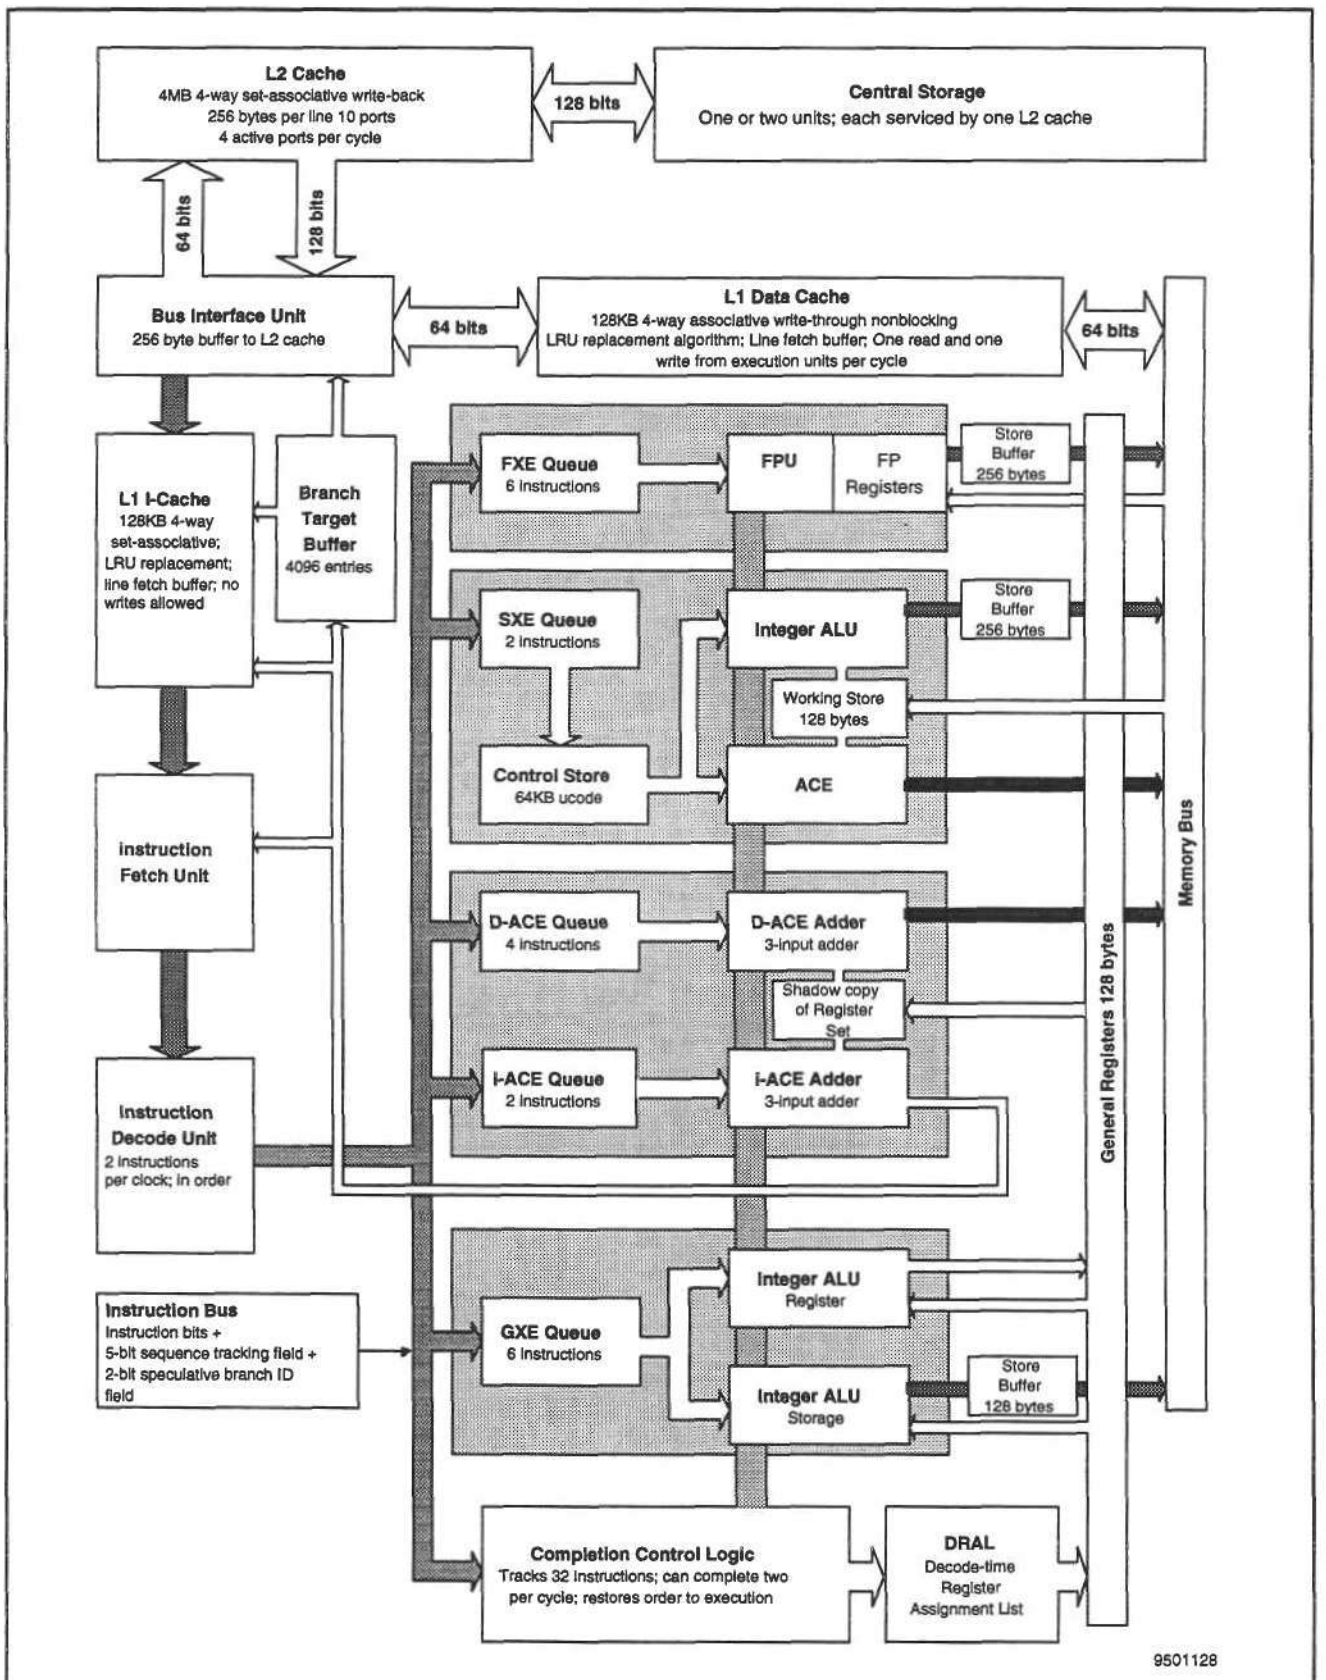

Source: Dataquest (March 1995)

Architecture

- It is superscalar, decoupled like K5, P6.
- It executes native code—doesn't translate.
- Sophisticated design components will migrate to micros.

The processor design is superscalar with decoupled execution resources. It can issue two instructions per cycle and incorporates register renaming and out-of-order execution and completion, and two levels of speculative execution along unresolved branch paths (that is, up to four alternative paths, one of which is correct). The instruction buffer manages up to 32 active instructions and can retire up to two instructions per cycle.

The instruction set of these machines has its roots in the 1960s, when machine code programming still existed and assembler code was king. As such, the instruction set was designed to facilitate programming and sports a rich range of CISC-like features including memory-to-memory arithmetic, instruction lengths ranging from 2 to 6 bytes, and the opportunity to design and build a new instruction on the fly for particularly time-consuming instructions that could be executed entirely in the microcoded processor core.

The design incorporates three arithmetic logic units (ALUs), aligned to each of the three addressing modes: register-to-register, register-to-storage, and storage-to-storage. The two simpler register-based ALUs are located in the General Execution Element (GXE), one of the four TCMs that compose the processor. This module incorporates its own 6-byte instruction queue, has a four-stage pipeline, and is controlled by hard-wired logic. It is capable of executing two instructions per cycle, one in each ALU. We see here the elements of a RISC-like architecture, using the benefits of deterministic instruction lengths and properties to execute at speed.

The memory-to-memory instructions are more complex. These are handled by the System Execution Unit (SXE), a microcode-driven module that handles the complex instructions including system control functions, for example, those that work with the timers, the program status word, and the control registers. Architecturally, this element represents the CISC components of the machine. Its monstrous 64KB writable microcode store (Writable Control Store, or WCS in IBM speak) enables it to fulfill the dreams of assembly-coding system programmers, while the rest of the machine clocks the more predictable instructions through hardwired logic units.

The Address Computation Element, or ACE, includes two address generators, one for data and one for instructions. It has a shadow copy of the register set and provides address calculations for the GXE and floating-point unit.

The registers in this machine consist of a pool that can be allocated to any instruction. The decode-time register assignment list (DRAL) tracks these registers in conjunction with the completion control unit and sets up pointers that identify which ones actually carry the values in the architected registers.

The machine also incorporates a complex memory hierarchy of queues, TLBs and caches. These features allow it to maximize its use of the available memory bandwidth. A full description of the multilevel structure is beyond the scope of this article but no doubt discloses some of the techniques that Intel has adopted in its still-cloaked memory interface units in the P6.

A full analysis of this architecture is beyond the scope of this document; however, the interested reader is referred to the July 1992 issue of IBM's *Journal of Research and Development* for a thorough treatment.

Key Architectural Differences between the Model 700 and the P6

Native Execution Core

This is a critical architectural feature. The Model 700 uses a 48-bit instruction bus with seven auxiliary control bits, far fewer than the 118 bits used in the P6. These instructions are essentially of the native IBM architecture; the machine feeds them to units that essentially represent little computers within the processor, rather than centrally-controlled execution units. This approach, with each processor having its own queue, looks more like the PowerPC or AMD K5. Intel's approach echoes the PA-8000 or MIPS R10000. Both approaches have merit; the IBM design probably results in more complexity in the subprocessors, in exchange for a simpler front end.

Branch Prediction

The branch prediction is much less sophisticated in this machine; newly-encountered backward branches are assumed taken, forward branches are assumed not taken. Any taken branch is stored in the history table, untaken branches are flushed. This is much less sophisticated than the four-bit Yeh algorithm used in the P6. To some extent, the difference is made up in scale: The 4,096-entry history table increases the overall accuracy by having fewer branches fall outside of the short-term memory of the machine.

Speculative Execution

The model 700 stores two tag bits with each instruction, identifying the branch path with which it is associated. Once a branch is verified, all of the other paths in process (up to three) are flushed from the machine. Therefore, the machine can only speculate to two levels deep. The P6 does not appear to have this restriction; we would speculate that the instruction time stamps also include branch path ID bits in order to accommodate the 40 or so instructions that can be sitting in the execution unit at any one time. The model 700 is only capable of holding 35 instructions at any one time.

How Fast Is This Machine?

Execution Speed

We have not been able to find SPECmarks for these machines, making comparisons difficult. However, cross-matching some older benchmarks and scaling for new machines lead us to conclude that these machines offer performance in the 300 SPECint range when running at speeds of about 150 MHz. Looking at the architecture, we can speculate that it could achieve one to two completions per cycle. The richness of execution

resources would suggest that it is going to achieve a better execution rate per clock than the P6; its less sophisticated branch prediction is compensated for by the enormous history table (4096 entries as opposed to 512), tremendous memory bandwidth and relatively low clock rate. However, the physical implementation of this machine imprisons it at clock rates below 200 MHz, giving the P6 room to draw ahead.

I/O Bandwidth

The last bastion of the mainframe is multiprocessing and I/O bandwidth. The sophisticated multiport L2 cache and banked central storage combine to permit the implementation of a 10-way multiprocessor that scales reasonably well. The P6 should scale well to four-way multiprocessing but may be restricted when a second system comes into play.

The P6 supports maximum sustained transfer rates in the 400-MB/sec range; the Model 700 can move 8 bytes to the L2 cache and 16 bytes back per cycle, for a sustained bandwidth of 3,600 MB/sec. The L2 cache communicates with main memory at only 2,400 MB/sec. Clearly, the Model 700 has headroom to spare and probably does not use more than a fraction of the available bandwidth; it is this bandwidth that supports the 10-way scalable multiprocessing so effectively.

Chapter 11
Speculating on the P7 and the HP-Intel (HPI) Architecture

- HPI will happen in 1998 for HP, in 2000 for Intel.
- The P7 will be another x86 derivative.
- Very long instruction word (VLIW) features in the compiler only.

The P6 offers us little in the way of clues about the P7 architecture. Intel's strategy would lead us to believe that it will be disclosed late in 1996 for introduction and first shipments in 1997 and real volume in 1998. This is just in time to catch the last tidal wave of the x86 market as Microsoft rolls to Windows 97, probably its implementation of a dedicated and hand-tuned operating system. For this reason, the HP-Intel architecture may not be incorporated in the P7 because it will add no value.

Therefore, we would speculate that the P7 will be a faster variant that builds on the P6 architecture with larger primary and secondary caches, more execution units including enhancements for digital signal processing (DSP)-like functions including motion video codecs and audio processing, a strategically shortened pipeline (implemented through more transistors in the decoder) to reduce branch miss penalties, larger TLBs and BTACs to improve prediction, and an even higher bandwidth instruction fetch and decode unit. It could also include a minimal level of programmable micro-code storage, perhaps to introduce flexibility and performance for key software algorithms.

The HPI Architecture

This strategy would mean that HP would introduce the first member of the joint family, a product faced with the formidable challenge of matching the position of Digital's current Alpha (that is 1.5 to 2x faster than the current tier-2 RISC processors) in the performance stakes. First application would be in HP mainframe replacements followed by high-end workstations and then, through the progeny of Windows NT, to the portable-operating-system-based desktop market in the year 2000.

The performance issue is critical: As portable software enables software vendors to hop among architectures with some facility, the HPI architecture must offer a major performance advantage over the PowerPC, MIPS, and x86 cloners that will cause manufacturers to settle on HPI as the architecture of choice. Digital may continue to occupy its wild card position by increasing parallelism in the Alpha architecture, taking it away from the pack through both clock speed and architectural performance.

An Opportunity for Apple

- Migrating to portable software and taking a slice of the x86 operating system market

There is a key opportunity for Apple here. If Copland is truly portable, Apple has the opportunity to port it to the HPI architecture to gain market share from Microsoft. At the same time, Apple has the opportunity to explore coding design and implementation methods that can deal with the

processor to main memory time warp by increasing the native parallelism of object code. There may even be a case for an x86 port of the Apple architecture, but that is beyond the scope of this report.

A Note on VLIW

- Hardware VLIW is no longer relevant.
- Software techniques derived from VLIW design will play a central role.

VLIW is often suspected of being the core of the HP-Intel architecture. Although we believe that it will implement techniques developed from VLIW research, we also believe that the pure VLIW concept is outdated in the modern desktop environment. The HPI hardware architecture is not likely to look like a classic VLIW processor with multiple pipelines; it is more likely to echo the decoupled-resource architectures of today's advanced microprocessors (and mainframes). The original concept of VLIW was to extract parallelism from the code at the compiler, delivering preparallelized object code to a simple hardware set with rich execution resources. In this approach, the hardware could be relatively simple as execution conflicts were all eliminated at the object code level. In other words, VLIW object code has a much higher degree of parallelism—better living through compilers.

However, the P6 and other next-generation processors show us that the advance of hardware technology will continue to deliver more and more execution resources with full interlocks, an environment in which many of the features of VLIW hardware are irrelevant. The true problem is that current object code does not have sufficient parallelism to take advantage of these resources; the gain from extra resources rapidly drops off as hardware is added.

VLIW compiler technologies, adapted to these resource-rich architectures, will produce object code with far more parallelism than we have today. In particular, compiler-generated hint bits—produced using VLIW compiler trace technology—will significantly increase the accuracy of branch prediction, eliminating a major cause of execution stalls.

VLIW, however, may have a future in low-cost applications where the effort required to develop highly parallel code is rewarded by performance in a high-volume application where unit cost is key. In this case, the expensive hardware required to resolve resource conflicts is eliminated in exchange for a larger up-front investment.

Chapter 12
Multiprocessing and the P6

Although we have touched on it elsewhere, the multiprocessing features of the P6 deserve further expansion. Intel introduced the concept of glueless multiprocessing years ago with the exotic i432 product, tested it with the i960K and then went mainstream with the P54C version of the Pentium processor; we expect Pentium MP-enabled systems to proliferate into the midrange and premium market through 1995, even though there is almost no software that takes advantage of the feature. The reason that it will appear is because the incremental cost of the second socket needed to add the MP feature is only a couple of dollars, and the current version of Windows NT supports it today. There are a few specialized applications available—for example, Adobe Photoshop—that can use the second processor when running under Windows NT.

The P6 goes even further toward glueless MP. The integrated cache solves the major performance bottleneck in dual processor Pentium systems, that of the bandwidth to the cache memory: Both processors must arbitrate for access to the L2 cache. In the P6, the L2 cache is located behind the processor's memory controller, and the memory controller can manage the cache coherency across the P6 system interface. This approach adds 256K of cache for each additional processor, greatly improving the scalability of the architecture over the Pentium design. The cache controllers support a full MESI protocol and include provision for direct cache-to-cache transfer, without any external electronics whatsoever. There is a price to be paid in the memory controller, which must be able to recognize these transactions as they pass across the bus.

The hardware implementation of an MP system, other than following the design guidelines for laying out a 66-MHz GTL+ bus, becomes very simple. The level of integration and performance possible with total control of the silicon eliminates much of the value add that server vendors have enjoyed in the past. Also, the MP capability adds absolutely minimal cost, and an MP-enabled system can be upgraded by the user for the cost of a processor. This further devalues engineering investment as the return on sale for upgrade processors cannot be realized.

In summary, a P6 MP system simply has up to four pin-to-pin linked sockets. This is a simple implementation issue for any board manufacturer. Furthermore, we can figure that second sockets will be a standard feature on all Intel-built motherboards.

Making MP Happen on the Desktop

The Pentium MP system is adequate and useful in certain applications running under Windows NT. To be successful, however, there must be some draw for MP support under Windows 95. We believe that some enterprising developers will come up with support dynamic link libraries (DLLs) that use the second processor to perform specific functions. An example of this would be an image transform under Photoshop, where the user could realize significant performance benefits without the overhead of having to run the Windows NT version.

As time goes on, we are likely to see standardized DLLs that provide asymmetric MP functionality to many Windows applications, and it is possible that the final revision of x86 Windows will support some limited amount of symmetric multiprocessing. The advent of the P6 makes the development of these support features all the more valuable.

P6 MP and the Technical Workstation Market

Many technical workstations are sold into compute-intensive applications including electronic design, mechanical design, and simulation. These are applications that have traditionally demanded maximum performance but are also capable of being divided and conquered on a parallel-processor machine. Windows NT is becoming a preferred operating system for technical applications by virtue of its absolutely standard nature. The combination of Windows NT and the P6 is a natural one for this market, offering the best combination of standards and performance; through the multiprocessor upgrade, a P6/Windows NT system becomes a serious competitor for most of the workstation market today because the incremental cost of the second processor is relatively low. The P6 thus presents a further challenge to all products based on better-than-Pentium performance, a segment of the market that is continually shrinking.

Chapter 13
Wrap-Up

To wrap up our position on the P6, there are a few things that bear repeating, as follows:

- Plan for desktop and server P6 systems to be introduced in the fall of 1995.
- Work on P6 notebooks right away; pester Intel for a single-chip TAB package P6 as soon as possible; demand chip sets without MP support for notebooks.
- Give serious consideration to using Intel-built motherboards for the first generation of product. It is not worth missing the market to add value.
- Include MP support in all products and prepare the market with dual-socket Pentium machines. The incremental cost is minimal.
- Do not count on the 486 to work in the U.S. market in 1996, except in basic notebooks and low-end value-class products.
- P6 is the big product for Christmas 1996.

For More Information...

Martin Reynolds, Director/Principal Analyst ..... (408) 437-8212
Internet address ..... mreynolds@dataquest.com
Via fax ..... (408) 437-0292

Dataquest

a company of
The Dun & Bradstreet Corporation

NGPC-WW-FR-9501 Qty: 1
Ms. Maria Valenzuela
Dataquest Incorporated
1-1100
.,
--INTERNAL DIST.--

Dataquest

a company of
The Dun & Bradstreet Corporation

Corporate Headquarters
Dataquest Incorporated
1290 Ridder Park Drive
San Jose, California 95131-2398
United States
Phone: 1-408-437-8000
Facsimile: 1-408-437-0292

Dataquest Incorporated
Nine Technology Drive
Westborough, Massachusetts 01581-5093
United States
Phone: 1-508-871-5555
Facsimile: 1-508-871-6180

European Headquarters
Dataquest Europe Limited
Holmers Farm Way
High Wycombe, Buckinghamshire
HP12 4XH
United Kingdom
Phone: 44-1494-422722
Facsimile: 44-1494-422742

Dataquest GmbH
Kronstadter Strasse 9
81677 München
Germany
Phone: 49-89-930-9090
Facsimile: 49-89-930-3277

Dataquest Europe SA
Immeuble Défense Bergères
345, avenue Georges Clémenceau
TSA 40002
92882 - Nanterre CTC Cedex 9
France
Phone: 33-1-41-35-13-00
Facsimile: 33-1-41-35-13-13

Japan Headquarters
Dataquest Japan K.K.
Shinkawa Sanko Building, 6th Floor
1-3-17, Shinkawa, Chuo-ku
Tokyo 104
Japan
Phone: 81-3-5566-0411
Facsimile: 81-3-5566-0425

Asia/Pacific Headquarters
7/F China Underwriters Centre
88 Gloucester Road
Wan Chai
Hong Kong
Phone: 852-2824-6168
Facsimile: 852-2824-6138

Dataquest Korea
2506 Trade Tower
159 Samsung-dong
Kangnam-gu, Seoul 135-729
Korea
Phone: 822-551-1331
Facsimile: 822-551-1330

Dataquest Taiwan
3/F, No. 87 Sung Chiang Road
Taipei 10428
Taiwan, R.O.C.
Phone: 8862-509-5390
Facsimile: 8862-509-4234

Dataquest ICC (conferences)
3990 Westerly Place, Suite 100
Newport Beach, California 92660
United States
Phone: 1-714-476-9117
Facsimile: 1-714-476-9969

Other research offices in Beijing, Singapore, and Thailand

Sales agents in Australia, Belgium, Germany (Kronberg), Israel, Italy, South Africa, and Spain

Focus Studies

PC Semiconductor Applications Worldwide

Focus Analysis

PC-Specific Standard Products

Abstract: *This article presents a unit-weighted look at trends in core logic, graphics, and super I/O regarding proprietary versus standard product, movement to the motherboard, and supplier shifts. The article is based on teardown analysis of personal computers over the past two years. The core of this information comes from Dataquest's own ongoing Personal Computer Teardown program. Supplemental data was also pulled from various trade journals. The information from these sources was combined with Dataquest's PC-by-model shipment statistics to obtain results weighted by unit shipments.*
By Jerry J. Banks

PC Core Logic: Standard versus Custom

The age-old decision of make versus buy is as critical in the PC marketplace as in any other market. Before Chips & Technologies invented the chipset business, PCs were built primarily from discrete logic elements. Every company attempting to play in this fledgling market had to come up with its own discrete design to manufacture a PC that not only was compatible with the machines offered by IBM, but also had to offer some differentiating features that would compel buyers to purchase it. Unfortunately, in the early days of the PC, IBM compatibility was a difficult task to master. Adding the complexity of feature differentiation made compatibility an even more difficult proposition. For a short time, various vendors would submit that a machine's ability to run a software game entitled Flight Simulator was tantamount to a guarantee of IBM compatibility.

Once Chips & Technologies established the concept of a standard core logic LSI chipset, IBM compatibility became a nonissue. PC OEMs quickly saw the benefits and rapidly moved to either a standard chipset solution or a custom implementation of the standard. As the market started to shake out, there was a clear delineation between the use of custom versus standard

Dataquest

Program: PC Semiconductor Applications Worldwide
Product Code: PSAM-WW-FA-9501
Publication Date: February 27, 1995
Filing: Focus Studies

products between those at the top of the market share ladder and those trying to move up the ladder. IBM and Compaq, which clearly were dominating the market, chose to use the custom logic approach. This decision was based on several factors: each had the requisite system expertise in house; each knew the customer base it was trying to penetrate; each commanded high enough margins to afford such an internal development effort; and smaller companies using standard core logic chipsets could only differentiate themselves on price and delivery while IBM and Compaq Computer could incorporate feature differentiation into their products, which would allow them to command a price premium.

The Impact of the "PC Clone"

As third-party core logic chipset suppliers gained experience in developing integrated LSI solutions for the PC, the gap between the feature set of proprietary products and standard products began to diminish. This dealt a significant blow to IBM and Compaq and other "name-brand" suppliers because the "PC clone" makers in the Asia/Pacific region were able to offer PCs with equivalent feature sets to the name-brand manufacturers at a fraction of the price. In fact, the clone-making label was not limited to the Asia/Pacific region. several North American companies, such as Gateway 2000 and Packard-Bell, adopted a standard product low-cost posture that was causing a severe erosion in market share for the higher-priced name-brand companies.

Name Brands Fight Back

This rapid market share erosion is what caused Compaq to make an abrupt turnaround in mid-1992 with the announcement of massive price cuts in an all-out effort to stop the erosion. Compaq was soon followed by the other name-brand manufacturers, and a full-scale PC price and market share war began that continues today. Beyond the issue of price, another tactic in this market share battle was to accelerate the introduction of new models in the hopes that competitors that could not keep up on the technology treadmill would drop out of the race.

New Market Factors Force the Switch

The new lower prices and the corresponding lower margins eliminated much of the R&D budget available to the OEMs, thus modifying a key variable in the make-or-buy equation. Another variable that was modified was time. The faster introductions of new models and the shorter life cycles for production products brought to bear much greater time-to-market pressures. PC OEMs now had to develop new products much more quickly, at lower cost.

To maintain this accelerated technology pace in conjunction with the reality of lower margins, many manufacturers were forced to forgo their preference for custom core logic and switched much of their product lines to standard product. Only the higher-margin machines still use custom core logic chipsets, and we expect these machines to give way soon to standard product offerings.

PC Core Logic Vendors Step Up

A factor making this switch from custom to standard logic more palatable is the fact that the core logic vendors are shipping very good product. The leading core logic vendors now have as much PC system expertise in-house as do most of the PC OEMs. It simply does not make any sense to maintain the R&D effort necessary to generate a next-generation chipset when it may at best match the price and features of product available off-the-shelf.

In 1993, 15 percent of the desktops PCs shipped included custom core logic chipsets (see Figure 1). By 1994, this had dropped to only 9 percent.

The Portable Difference

The issues in the portable marketplace were somewhat different from those in the desktop, and because the desktop market drives such a large portion of the total available market, the desktop PCs received primary attention from the PC core logic vendors. The result was that notebook OEMs that relied on standard core logic product offerings often were unable to market a competitive product in terms of such a critical feature as battery life. This left a larger challenge for PC notebook OEMs. Battery life and, as a result, power management are critical issues for notebook vendors and are a side issue for the desktop vendors. Notebook OEMs that relied on standard core logic product offerings were at a disadvantage to those with enough of an R&D budget and requisite systems expertise to develop their own power-managed solutions. Brute force techniques such as larger (and heavier) batteries or multiple battery solutions are often implemented to bandage an inferior power-managed solution. Such solutions not only increase the cost significantly, but also add significant weight to a weight-sensitive market.

Figure 1
Core Logic Trends in Desktop PCs

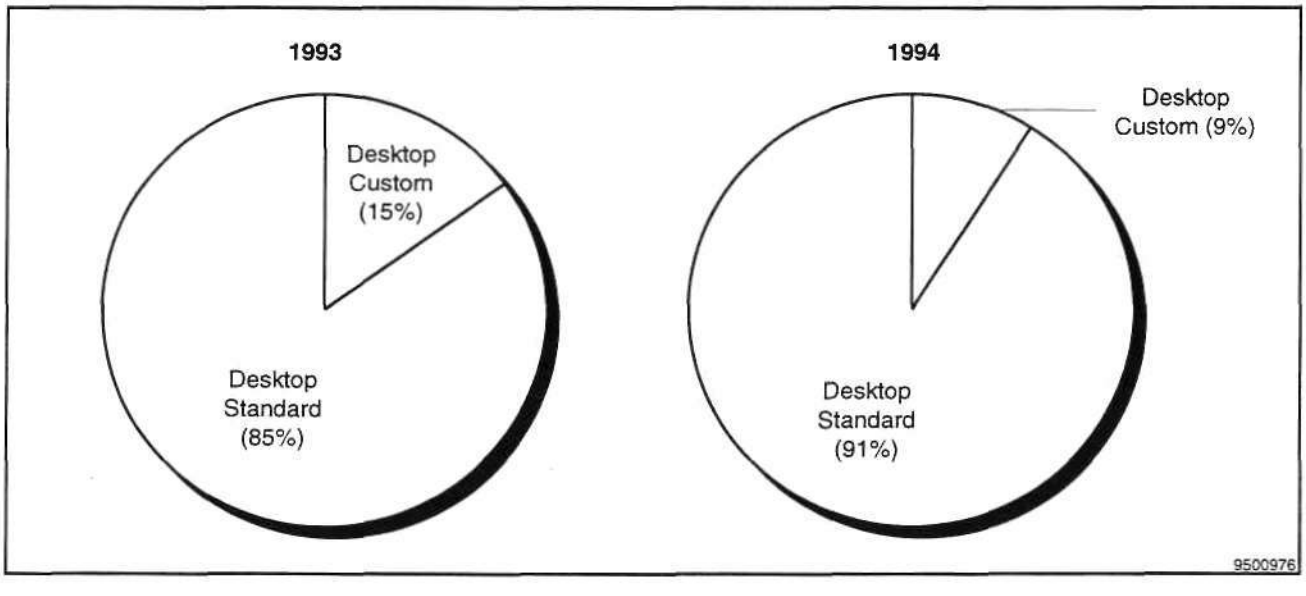

Source: Dataquest (February 1995)

Given these issues, it is no wonder that manufacturers of portable PCs have moved much more slowly to the use of standard core logic product offerings. Figure 2 demonstrates this fact. However, because of the recent attention given to the needs of the notebook market by PC core logic suppliers, the capabilities of the products offered have improved significantly over the past year. In fact, Figure 2 shows that a significant number of portable PC OEMs are beginning the move to standard core logic products. We expect to see this trend continue to the point that the use of standard core logic products in portable PCs will approximate that of desktop PCs in a three- to four-year time frame.

Whither Goest the Motherboard?

The desktop PC motherboard has long been the dream home for manufacturers of microperipherals products. An IC manufacturer that gains a design win with a motherboard manufacturer has a much greater chance of seeing a product reach high-volume production than one dealing with an add-in board manufacturer. Also, an IC manufacturer that sells to the motherboard has a greater amount of control in the IC of choice selection process. For example, if an IC manufacturer sells to an add-in card manufacturer, the IC vendor must hold its breath and hope that the add-in card vendor has the marketing savvy to sell its card solution to the end market. Only then will the IC vendor see any revenue. If the IC manufacturer sells directly to the motherboard manufacturer, particularly if the motherboard vendor is also the PC OEM, it can rest assured that appropriate effort is put into the design-win process and that its efforts will be rewarded. An additional benefit of such a close relationship with a motherboard or PC OEM is that the IC vendor can rest assured that customer input for next-generation products is on track and should result in products with the right features and price at the right time.

Figure 2
Core Logic Trends in Portable PCs

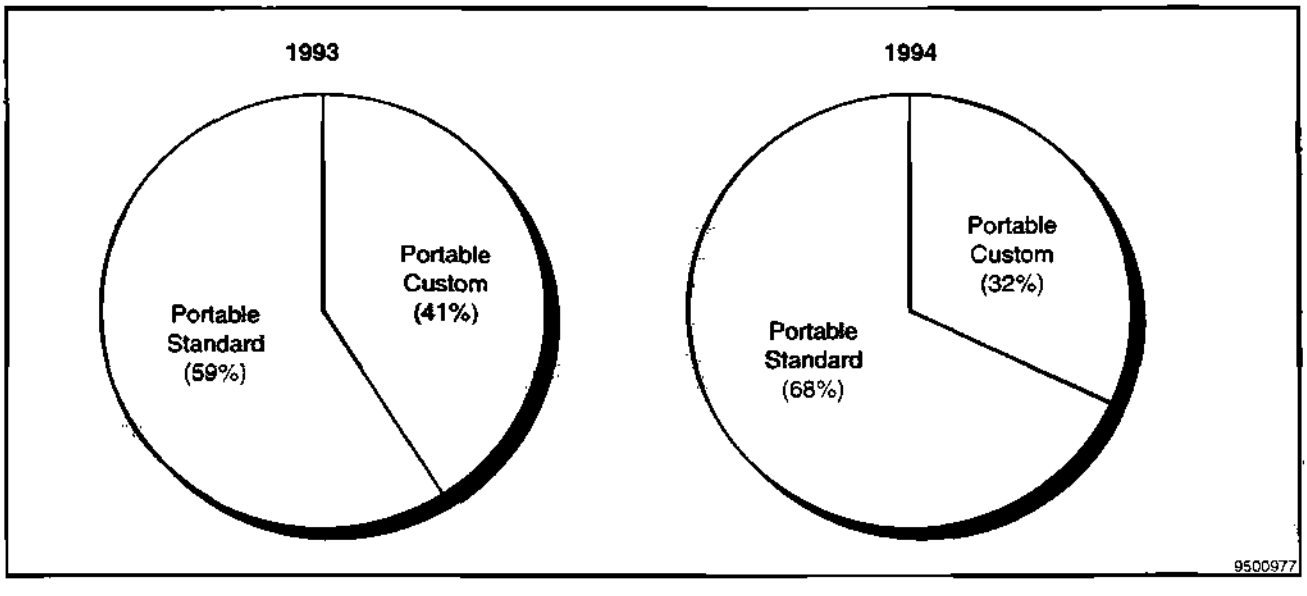

Source: Dataquest (February 1995)

The problem is that the motherboard manufacturers are reluctant to add components to the motherboard. If a particular component adds a feature required by a small number of end users, it is just a source of added cost with minimal or no value. In the low-margin business of PC motherboards, cost without value is the kiss of death. Conversely, if a motherboard manufacturer is not adding sufficient value or the right feature set, then its product again is noncompetitive. As a consequence, motherboard manufacturers are constantly evaluating the feature sets purchased by the end user. When a particular feature becomes a de facto, or official, standard and reaches the right price point, chances are it will find its way onto the motherboard.

A Graphic Discussion

The only microperipherals products to consistently find a spot on the motherboard have been PC core logic chips. Bringing a graphics controller onto the motherboard has been limited only to those with very high volume. Adding connectors also adds cost to the motherboard and typically requires the development of a custom cabinet, which also adds significantly to the cost of developing a new PC. If the volume is high enough, this cost can be amortized over many units. Another factor limiting movement to the motherboard is standardization. Graphics controllers are a classic case. It is the rare PC that does not require a graphics controllers. So, why don't all motherboards include a graphics controller? The issue here is standardization. Not all users demand a Windows accelerator. Some are quite happy with a basic, cheap standard SVGA controller.

Risk/Reward

If one were to assume that a certain number of systems would require Windows acceleration, wouldn't it be a simple matter to perform some market research to determine what percentage of systems required this feature and build the correct percentage of motherboards that contained a Windows accelerator chip? Well, now the decision comes down to which vendor, at what price, and at what performance? No vendor at this time second-sources any other vendor's product offering. As a consequence, committing to a particular vendor's Windows accelerator solution means a lock-in to a specific price performance point that may or may not meet the needs of the end user. This is particularly risky for pure motherboard vendors. A PC OEM may have a different idea as to the appropriate graphics price performance point as well as the choice of graphics controller vendor. Switching to a different Windows accelerator solution is not a simple matter of plugging in a new chip. All vendors have a different pinout and offer slightly different features. Also, at any point in time a new entrant is likely to appear on the scene with a superior feature set and lower price than the existing solutions. Such a change requires a new design and board layout on the part of the motherboard manufacturer. The wrong selection of a graphics solution by the motherboard manufacturer may put it out of the running.

PC OEMs that design their own motherboards are not totally without risk themselves, should they choose to integrate the graphics solution onto the motherboard. Although it is a sure bet that the board will find a home in the PC, it is not a sure bet that the PC itself will have the right price and feature set to succeed in the marketplace. An add-in card approach that allows end

users to select their own graphics solution is a much safer bet. However, such an approach does limit the OEM's ability to differentiate in other areas such as system packaging. The more flexible the approach, the more rigid the form factor of the PC itself.

We still continue to find that only the large name-brand PC OEMs are integrating graphics onto the motherboard. This allows them more freedom in package design, which can be a key differentiator in a large portion of the PC marketplace. The midrange and small PC vendors still rely on the add-in card strategy for graphics control.

Standard or Roll Your Own?

The choice of custom versus standard product graphics controllers is even more dramatic than is the trend for PC core logic. On a unit weighted basis, nearly 100 percent of PC manufacturers are shipping systems with a standard graphics controllers. This weighted average leans heavily to the standard products side, as depicted in Figure 3.

SIO Evolves

The SIO or super I/O function historically has consisted of serial and parallel communications ports, a game port, a floppy disk drive controller and interface, an IDE mass storage interface, and sometimes a mouse port. More recently, with the introduction of multiple CD-ROMs that use an IDE interface, an IDE CD-ROM control function has been added to the SIO (see Table 1). Also, the performance limitations of the standard IDE rigid disk drive interface have caused the advent of the enhanced IDE (EIDE) interface that can handle data transfer rates approaching 16 MB/sec. We are now seeing this interface being integrated into the SIO chip/chipset.

Figure 3
Graphics Controller Standard Product Trends

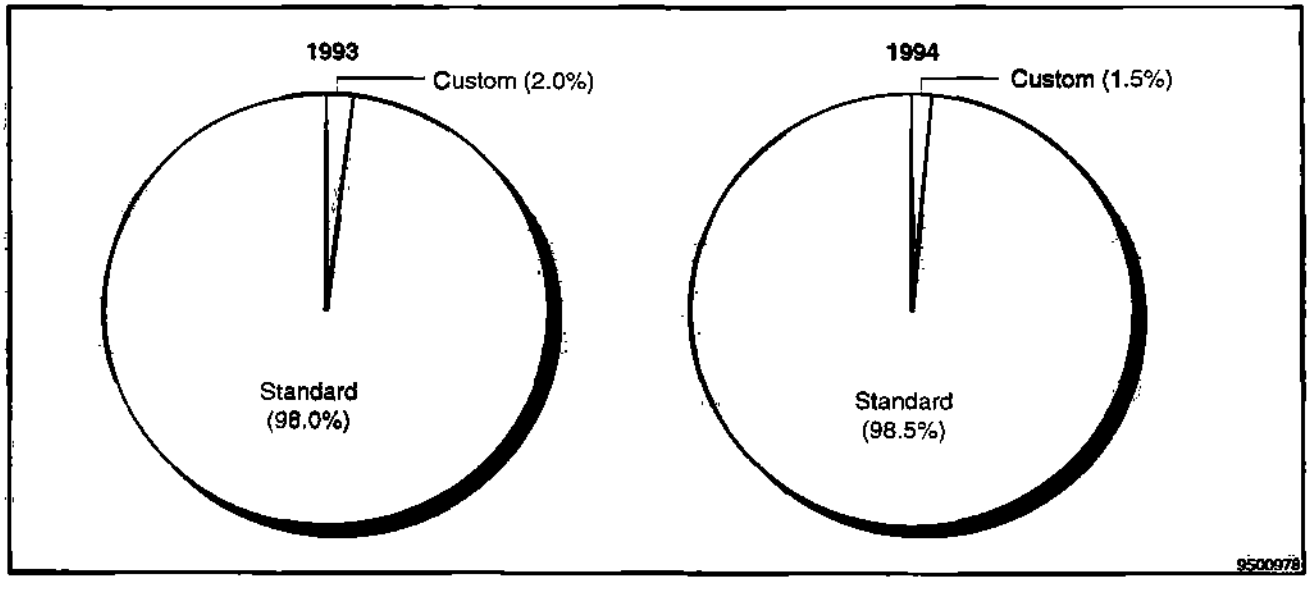

Source: Dataquest (February 1995)

Table 1
SIO Feature Evolution

Up to 1994	1995
Serial port	Serial port
Parallel port	Parallel port
Game port	Game port
Mouse port (optional)	Mouse port (optional)
Floppy disk controller	Floppy disk controller
IDE mass storage interface	Enhanced IDE
	Higher data transfer rate
	CD-ROM support

Source: Dataquest (February 1995)

This new super I/O functionality will dramatically increase the performance of the mass storage interface and provide the opportunity to get rid of the mishmash of CD-ROM interfaces that included proprietary solutions as well as a SCSI solution. This will also open up an expansion card slot of the PC motherboard, making room for features to be added to the PC.

The migration of the super I/O function to the motherboard has closely matched that of graphics control. The small- and medium-volume OEMs rely on the add-in card solution, while the high-volume OEMs will more quickly move functions onto the motherboard.

Dataquest Perspective

Three main factors are driving PC OEMs to use standard products in lieu of developing their own custom solutions:

- Lower margins
  - When Compaq Computer decided to halt the erosion of its market share to the PC clone manufacturers and started the PC price wars in the second half of 1992, the profit margin for PCs dropped dramatically. As a consequence, PC OEMs have fewer R&D dollars to spend on custom programs.
- Time-to-market pressures
  - One of the tactics in the market share wars is to accelerate the pace of new product introductions. This shortened time allowed for R&D hampers a PC OEM's ability to develop custom product.
- Better standard product availability
  - PC core logic vendors have greatly enhanced their internal capabilities and are producing high-performance, feature-rich, quality products at an affordable price.
  - Graphics logic vendors long ago provided this level of product and have almost completely eliminated the need for a PC OEM to develop a custom graphics controller.

With such compelling factors, it is no wonder that standard logic is becoming the choice of more PC OEMs. The "clone vendors" made this decision early in the game because price was their primary weapon against the name-brand vendors. Now that the PC-specific standard products have improved so dramatically and time and financial pressures on the PC OEMs have increased, we expect the move to standard products to continue on the desktop until virtually 100 percent of PC core logic and graphics controllers are standard products.

The notebook vendors, which until recently have not been able to find the ideal standard product solutions for their systems, are now beginning to make the move to standard products. Several companies have begun to address the unique needs of core logic in the notebook and will soon obviate the need for notebook OEMs to develop their own core logic products.

The add-in card market is still a good target market for those wishing to offer high value-added. This is the market in which high-performance solutions will find the quickest acceptance. Although a large PC OEM may be reluctant to add a high-cost, high-performance product to the motherboard, an add-in card manufacturer is much more willing to offer a product that offers significant feature differentiation with the expectation that higher margins will also result from such a product line.

Today's PC-specific standard product suppliers must have a broad strategy that covers not only the product line but also a broad spectrum of players in the PC marketplace, from PC OEMs to motherboard manufacturers to the add-in card market. This market is so dynamic that concentrating in just the low-cost, high-volume part of the market may cause such a player to be blindsided by new technologies or by a new competitor that heretofore was not visible to the narrow market player.

For More Information...

Jerry J. Banks, Director/Principal Analyst ..... (408) 437-8677
Internet address ..... jbanks@dataquest.com
Via fax ..... (408) 437-0292

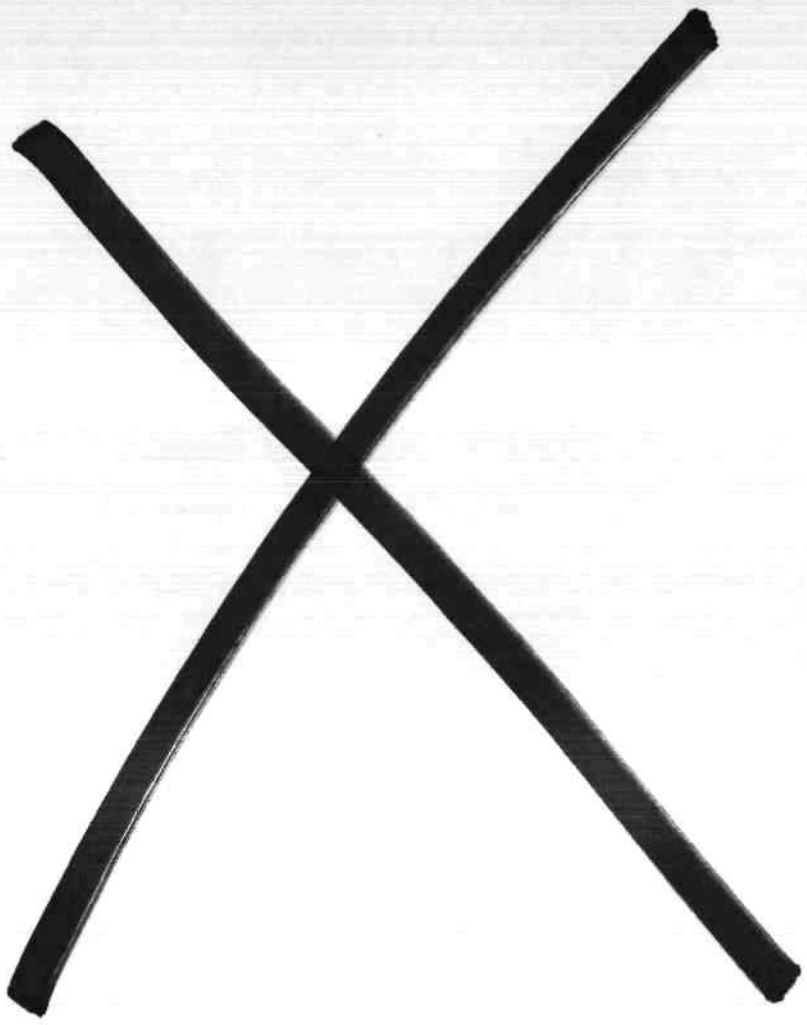

Dataquest