Fantastic Forecasts

Crystal Ball Sometimes on the Nose, Other Times on the Blink

OS/2 Has Dim Future

SEPTEMBER, 1987—It has become apparent that OS/2 will be slow and will require large amounts of memory (2 megabytes minimum if the compatibility box is used). All things considered, the personal computer world may well have been better off without OS/2. Many users may skip it entirely, and continue with DOS 3.x until they move up to a 386-based machine with a 386-based OS.

Microsoft foisted OS/2 onto IBM, where it languished and died.

Who Will Win the CPU Wars?

JUNE, 1988—Who will be the big winners? Here are our bets:

• Intel's 386 family, including the 386SX 16-bit-bus version and the future 486, will continue to dominate the IBM PC world for at least a few years. It will eventually have competition at the high end from 386-compatible chip sets, and from other architectures via software emulators and other compatibility bridges.

• Motorola's 88000 family will be the premier high-end RISC processor for general-purpose computing applications. Superior system-level integration is its key advantage.

• AMD's 29000 will capture the lion's share of the high-end embedded control market but will not have a significant impact on general-purpose computers.

• Motorola's 68000 family will be the Z80 of the 1990s—the natural choice for low-cost control applications. It will slowly decline in importance at the high end and may eventually lose much of the workstation market. The Mac will keep the 68000 family a key player in personal computers.

And the contenders:

• SPARC will be important because of Sun but will be eclipsed by the 88000 and 29000. Sun's own products will account for the majority of the volume. [SPARC] will have little success in embedded control.

• Intel's 80960 family will have considerable success in the embedded control market, particularly in military and aerospace applications, but will never achieve the popularity of Intel's other architectures.

• MIPS will have a hard time getting a significant share of the chip market, despite impressive processor performance, and [the company] will concentrate on their systems business.

Last and unfortunately least, the also-rans:

• Intergraph's Clipper will continue to languish, despite significant upgrades. Intergraph's workstations will be the dominant application. Chip users just don't trust Intergraph as a vendor.

• Inmos' T800 and other Transputers will be a favorite for networked multiprocessor experimentation but will never enter the mainstream. Few people have ever figured out how to program them.

• Zilog's Z8000 and Z80000 will never achieve significant success. Too little too late sums it up. Zilog has lost all credibility as a CPU provider, other than for the low-end Z80 and Z8 families.

A fairly accurate set of predictions dating back to our first year. We overestimated the impact of the 88000, and the 29000 did well for a time but eventually failed to pan out. Our other winners won, our contenders contended, and the also-rans quickly lost.

Heller Predicts Modules to Become Dominant

OCTOBER, 1989—[According to former IBM Fellow Andy Heller,] "The packaging and microprocessor design over the next five years will become indistinguishable—you will not buy the chip, you will buy the module."

Now that Intel is converting its entire product line to modules, this forecast has come true, but it will take ten years, not five.

Conflicts Surface at Microprocessor Forum

OCTOBER, 1989—[Roger Ross, of Ross Technology, said,] "I had the choice of joining any particular RISC camp, and we chose SPARC. ... There's one reason why SPARC is going to win: SPARC has now gone over $1 billion in R&D-basis installed software applications. That is a huge amount of software—there is nothing else that compares to it. The growth curve is greater than for all the other [RISCs] as well, which means they're not even catching up. To me, ladies and gentlemen, that means the war is over."

SPARC remains the best-selling workstation RISC but has been surpassed on the desktop by PowerPC.

Microprocessor Architectures in the 1990s

JANUARY 24, 1990—SPARC is a triumph of strategy over technology, as was Intel's 8086. SPARC's hardware implementations to date have been clumsy at best. ... SPARC will succeed not because of the technology, but because of the chip and system makers that are committed to it. Of key importance, of course, is

Sun, which is 100% committed to SPARC and is a driving force in the workstation business.

… Computing in the 1990s is likely to be dominated by architectures that already exist. At some point, however, a repeat of the RISC phenomenon is bound to occur. A new crop of architectures will appear that are better matched to the implementation technology of the time. Perhaps these will be 64-bit architectures, or VLIW designs.

In the case of IA-64, the answer is both.

The Need for Speed

FEBRUARY 7, 1990—There are countless other reasons why additional processing power will be useful. Features such as automatic spell checking while you type become more practical. Programs are easier to write when they don't have to be carefully optimized for speed. Much more error checking can be performed, and programs could suggest intelligent alternatives instead of just responding with "unknown command." Emulation of foreign instruction sets and use of processor-independent intermediate languages become more practical.

In this editorial, Michael foresees the emergence of Word 7.0, FX!32, and Java.


Special Issue: Celebrating the 25th Anniversary of the Microprocessor

MICROPROCESSOR REPORT

THE INSIDERS' GUIDE TO MICROPROCESSOR HARDWARE

Microprocessors Changing the World

Proliferation Only Beginning—Biggest Changes Ahead

by Michael Slater

Inside: Bell • Patt • Ditzel • Dally • Sites • Case • Tredennick

The Decline of Macintosh

APRIL 4, 1990—In the next few years, most of the important Macintosh applications will appear under Windows, OS/2, and Unix. Most new software development will be targeted to multiple platforms, and while the Mac will be one of them, it will not be the only one. It's going to take two or three years for OS/2 and Windows to stabilize and develop a broad range of applications, but when they do, the Mac will be in serious trouble.

If Apple continues their existing proprietary strategy, they may well be able to survive as a successful computer company, but their overall importance in the world of computing will gradually decline.

A regrettably accurate forecast.

Should Architectures Be Protected?

AUGUST 22, 1990—Architectures should be as available as [programming] languages are. If I want to write a program, I pick the language most suitable for the task, without concern for who developed the language. If I want to build a microprocessor, I should be able to pick the best instruction set for the task. … The architectures that dominate in the long run will achieve this public domain status, whether by choice or by force.

By force was Intel's choice.

MIPS and Sunset (by Nick Tredennick)

JUNE 26, 1991—The workstation market is about to saturate. PCs are going to eat into workstation sales. As engineering software becomes available for the PC, the number of engineering seats will increase rapidly, because of the significantly lower cost of PC-based engineering development systems. At the same time, sales of "serious" engineering workstations will decrease as they give way to the PC. The current workstation companies will be forced into niche markets such as servers, special requirements platforms, and front-end platforms for test equipment and the like. Sunset for raw mips and the workstation companies.

A long slow sunset, but it seems to be happening. RISC workstation unit growth has stalled, while sales of "PC workstations" are growing rapidly.

A New World for Intel

MAY 6, 1992—It is inevitable that Intel's share of the microprocessor market, as well as its profit margins, will decline. This doesn't necessarily mean the decline of Intel, but it does mean that the company will have to accept lower profit margins and devote some resources to developing products its customers want, not just those that optimize its profits. Ultimately, Intel is dependent on the PC business returning to high growth rates. The only way Intel can continue to increase its PC microprocessor business is for the market to grow faster than Intel's share of it declines.

Intel's margins and share have declined since the heady days of 1992, but the PC market has grown much faster, in part due to Intel's platform efforts.

Multivendor 386/486 Market Burgeoning

JANUARY 25, 1993—Unless Intel stumbles badly with Pentium, it will retain its performance lead among x86 implementations. … Pricing of 486DX chips will surely follow the trend of the 386 and 486SX, eventually heading under $100. The big price cuts probably won't occur until 1994, however, when AMD, Cyrix, TI, and probably others as well will have ramped up production. By then, Intel will once again have established the high ground with a range of Pentium-family processors, and Intel may gradually give up the 486 market just as it has with the 386.

A fairly accurate forecast of the next two years.

The Power Play (by Nick Tredennick)

JANUARY 25, 1993—Microsoft and the IBM-compatible personal computer will capture a giant chunk of the Apple Macintosh segment of the personal computer

market—provided Microsoft has a great story for migrating from the Mac to the PC. It only has to be a great story, not a great reality.

Nick followed this up with more specific numbers: SEPTEMBER 12, 1994—I think the Power Macs will find a place. If Apple does everything right, I think it will get a 7–8% market share for the Power Macintosh. I don't think it will get back to the 14% share it had with the 680x0-based Mac.

Since this was published, the PC has taken roughly two-thirds of the Macintosh market share, putting it below even Nick's prediction. Nick notes that Apple has certainly not done everything right.

IC Manufacturing Drives CPU Performance

MAY 31, 1993—By the end of the decade, we expect to see gate lengths under 0.2 microns, processor clock rates over 500 MHz, and single-chip microprocessors with 50 million transistors. Because chip-to-chip communication will not improve at the same rate, the processor designer's challenge will be to select the right mix of high-speed memory and function blocks to use these transistors most effectively.

The clock-speed prediction was met in 1997, the transistor-count goal will be achieved in 1998, and the gate length will occur on schedule in 1999.

HP Twiddles While PA-RISC Burns

AUGUST 2, 1993—HP's lack of openness dooms PA-RISC to low volume, small R&D budgets, and niche markets. HP knows well what can happen to boutique RISCs; after buying Apollo a few years ago, management quickly eliminated all support for Apollo's PRISM architecture. With a larger installed base, PA-RISC will certainly survive for quite a while, but it is unlikely to thrive without a significant change in market strategy.

PA-RISC failed to thrive and will ultimately be absorbed into IA-64.

Digital Bets Company on Weak Hand

AUGUST 22, 1994—It is cruelly ironic that Digital has done nearly everything right to establish its new architecture in the market: merchant processor and chip set availability, adoption of standards such as NT and PCI, establishing a separate business unit for chip sales. But Alpha has been undone by the poor timing of its market entry. … Alpha may continue to exist as a means of serving [Digital's] customer base, but the architecture will probably become irrelevant to the rest of the market.

Digital remains sold on Alpha, but no other major vendor is producing Alpha systems.

Apple Plans Open Licensing of PPC Platform

DECEMBER 5, 1994—Apple won't catch up [to Windows 95] in its underlying kernel technology until it ships the OS code-named Copland (likely to be officially named Mac OS 8.0). Once planned for mid-95 delivery, the projected date for this OS has slipped into late 1995 or early 1996. … If too much time passes between the delivery of Windows 95 and Copland, Apple risks a significant erosion of its customer base.

Copland never shipped, contributing to Apple's decline. Mac OS 8.0, a subset of the original Copland, finally shipped in mid-1997.

New Processor Families Join Embedded Fray

DECEMBER 26, 1994—For Motorola, ColdFire may be too little, too late. Any self-aware embedded designer has known for some time that the 68060 would likely be the end of 68000 dynasty, and years of easy upgrades would be over. Given the typically long design cycles for embedded applications, these designers have been looking for alternatives for many months, all the while deafened by Motorola's silence about alternatives. … By the time ColdFire devices are shipping, many will have already chosen another path.

Coldfire has yet to catch on, failing to stem defections from the 68000 camp.

Is Intel Sandbagging on Speed? (John Wharton)

AUGUST 21, 1995—It seems early [P6] devices won't be limited to 133 MHz after all, but will actually do up to 150! … I can't say I'm surprised. I've been telling my clients for months that the P6 was clearly designed to run much, much faster than Intel initially let on. … But I think even the 12% boost that Intel just discovered is grossly conservative. Indications are that the current P6 design, using the current [0.6-micron] process, could someday run at 266 MHz or more.

The P6 eventually achieved 266 MHz in a 0.28-micron process.

Intel Boosts Pentium Pro to 200 MHz

NOVEMBER 13, 1995—By the end of 1996, the existing Pentium Pro design should be boosted to clock rates of 233 MHz, or possibly even 266 MHz. To make its move into the mainstream, however, we expect Intel to shift away from the two-chip package with the integrated L2 cache in 1997, introducing an enhanced version of Pentium Pro with a dedicated bus for an external L2 cache. This version will also presumably incorporate the multimedia extensions that will debut next year in the P55C and will probably include enhancements to boost performance on 16-bit code. It is likely to be fabricated in pure CMOS, rather than BiCMOS, enabling lower supply voltages and reduced power consumption. By 1998, as this version reaches high-volume production, Pentium Pro technology will move into the PC mainstream.

Although Pentium Pro never made it above 200 MHz, this was an accurate description of Pentium II. ♦