VIEWPOINT

Room for Three Architectures in the 2000s

Intel/HP, PowerPC Look Good—Who Else Will Step Forward?

by Linley Gwennap

*Last issue (see **080801.PDF**), we discussed the ramifications of the Intel/HP partnership for the two partners; this story examines its effect on the competition.*

Intel's announcement of the successor to its x86 architecture places new light on the discussion of which instruction sets will prosper in the next decade. The partnership with HP is also a sign of consolidation: the six major instruction-set architectures will be reduced to at most five, and further consolidation is in the cards.

In the software industry, a rule of thumb is that there are two-and-a-half viable products in any given application area; in spreadsheets, for example, there are Excel and 1-2-3, with Quattro Pro hanging in there. As microprocessors move into the next decade, it looks like there will be an opportunity for two or three volume architectures for general-purpose computers, with everything else filling niches or moving to other markets.

Which will survive? Table 1 shows one possible scenario. As Intel throws its weight behind its post-RISC architecture, which we call P86, pure x86 processors will fall out of favor, ultimately disappearing from the market. With the exception of PowerPC, none of the current RISCs has either the software support or the financial backing to compete with the P86 team. This situation leaves many system vendors looking for a growth path.

	Current Status	Possibiity for 2000s
Compaq	x86 PC and server vendor	Deploy new architecture with AMD and SGI
IBM	x86/PowerPC chip and system vendor	PowerPC chip and system vendor
Dell, Gateway most PC vendors	x86 PC vendor	P86 PC vendor (some x86 PCs)
AMD	x86 chip maker	Chip vendor for new Compaq architecture
Sun/Fujitsu	SPARC system vendor	PowerPC system vendor
HP	x86 and PA-RISC system vendor	P86 system vendor (builds own chips)
Silicon Graphics	MIPS system vendor	System vendor for new Compaq architecture
Digital	Alpha chip and system vendor	Alpha system vendor (builds own chips)
NEC, Toshiba, IDT	MIPS chip vendor	Chip vendor for MIPS, new Compaq architecture
Hitachi, Oki, Winbond	Embedded PA-RISC vendor	Embedded PA-RISC vendor

Table 1. The Intel/HP announcement implies that many system and CPU vendors will need to change architectures later this decade; this chart shows one possible scenario.

Most PC vendors will follow Intel's lead and make the transition to P86 over time. Sun will need to abandon SPARC at some point, possibly for PowerPC. Other RISC system vendors may also look to P86 or PowerPC. But there is an opportunity for a third instruction set to compete over time, one that has the advantages of P86 (post-RISC design, x86 compatible) but not the disadvantage of a single source of processors (Intel).

This third architecture could be created by joining Compaq, which clearly dislikes having a single CPU source, with the MIPS team, bringing in NEC, IDT, and Toshiba as multiple chip sources. Silicon Graphics would provide CPU design expertise as well as system products complementary to Compaq's. NEC is the leading Japanese system vendor. The new architecture could offer compatibility with both MIPS and x86. To assist with the latter, AMD's design skills would be useful; AMD is also a CPU vendor that Compaq is comfortable with.

Although this scenario is just one possible outcome, pressures are building to cause such a radical industry realignment. To arrive at this conclusion, I make two assumptions. First, that Intel is truly committed to replacing x86 with P86; admittedly, the company has a poor record for teamwork, but its partnerships have always worked out best—for Intel. Second, I believe that, five or ten years from now, buyers will want chips that deliver thousands of MIPS. After all, software has always managed to use all available CPU performance in the past; future systems will probably need that performance for the 3D video user interface and, more important, the virtual-reality version of Rebel Assault.

Why x86 Will Fail

There are both technical and business reasons why x86 will not be a viable architecture into the next decade. From a technical standpoint, Intel has essentially admitted that x86 performance will not be competitive by the end of this decade, a point that others have been advocating for years: recently, even AMD's Mike Johnson admitted that x86 is inferior to RISC from a technical standpoint while making the business case that x86 is too entrenched to fail *(see **071605.PDF**)*.

That business case has changed now that Intel plans to throw its weight behind a new design. The crux of Johnson's argument was that "no company or consortium" could stand against the combined forces of Intel, AMD, Cyrix, and other x86 chip vendors. Now, however, it looks like the battle will feature Intel and HP on one

side against AMD, Cyrix, and other x86 chip vendors on the other. Against Intel's marketing muscle, can these vendors sustain buyers' interest in x86 for long? Suddenly, the x86 backers looks vulnerable.

Others argue that it is impossible to establish a new architecture because of the installed base of x86 software. P86 solves this problem by offering full compatibility with this software base. At the same time, Intel's backing will ensure that key ISVs port their applications to run natively on P86, offering superior performance.

There is every reason to expect that P86, in native mode, will deliver better performance than current architectures. History shows that creating a new instruction set to take advantage of the most recent research, as well as current manufacturing capabilities, can significantly boost performance: the i860 in 1989, POWER in 1990, and Alpha in 1992 all delivered leadership performance with their initial hardware implementations. If P86 is faster than RISC, and today's RISC chips are ahead of current x86 processors, then P86 should be far faster than any contemporary x86 devices.

Architectures such as the i860, POWER, and Alpha failed to dent the leadership of x86 despite offering better performance. But none of them offered compatibility with the existing software base. This compatibility will be the key to the success of P86.

x86 Vendors Have Limited Window

The x86 market will continue to be important for several years, even after the debut of the first P86 chips. There may even be a small window for pure x86 chips to compete against these early P86 chips, depending on the emulation performance of the Intel devices. Once P86 appears, however, the market for x86 will inexorably decline around the end of the decade and vanish by 2002, assuming there is demand for the higher performance of these P86 chips. Figure 1 shows this progression.

In the interim, x86 chip vendors can continue business as usual. They can even point out that P86 offers a growth path from their current x86 chips. Ultimately, however, these vendors must consider how to stay in business after the x86 market begins to decline.

The obvious option is to build processors compatible with P86, but Intel will make it very difficult for other companies to do so, both by keeping portions of the new architecture under wraps and by wielding chip- and system-level patents as much as possible. Over time, Intel's patent cross-licensing agreements with other foundries will expire, and Intel may convince these companies to accept more restrictive terms that prevent them from building chips for Intel's unlicensed competitors. These and other tactics will probably prevent a multivendor P86 chip market from emerging.

Another choice for x86 chip makers is to adopt PowerPC or another existing architecture. If the market for PowerPC continues to develop, IBM and Motorola may consider licensing other processor vendors. But both of these companies have significant manufacturing capacity and are unlikely to desire additional competition. If the PowerPC market becomes large enough, other vendors may build unlicensed PowerPC chips, although there would be some potential legal difficulties with this approach.

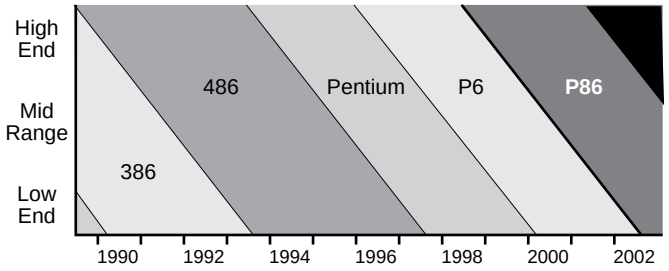

Figure 1. If Intel handles P86 as it has past generations of processors, pure x86 devices (such as Pentium and P6) will disappear from the market around the year 2002.

A third option is to develop a new architecture, an expensive and risky option. Partnering with other companies, as in the Compaq/SGI/AMD scenario, would allow the expenses to be shared and the risks to be reduced. Exiting the market for general-purpose CPUs would, of course, be a final option.

Why Most RISCs Will Fail

In RISC vs. CISC debates, the RISC guys usually demonstrate their technical superiority, but the CISC guys always win by pointing out that RISC volumes are far too small: a niche x86 vendor like Cyrix, for example, outships the leading RISC makers. Despite years of effort, most RISCs have been unable to establish even a toehold in the high-volume desktop market.

The sole exception has been PowerPC. With the backing of IBM, Motorola, and especially Apple, that architecture will be the first RISC to sell one million systems per year, and it has a clear path to several times that number. PowerPC's reasons for success are twofold: it has the unconditional backing of a high-volume system vendor, and it offers compatibility with a large base of existing (Macintosh) software applications.

None of the other RISC architectures has been able to achieve either of these goals. As a result, all are stuck in relatively low-volume markets with little hope of improving their lot. Without volume, it has been difficult for the RISC vendors to attract adequate ISV support *(see **0809ED.PDF**)*. Without applications, there is no clear reason for a major system vendor to adopt RISC.

From a technical standpoint, it is unclear whether RISC vendors will be able to keep up with the post-RISC P86 design. Given the freedom of a new instruction set, Intel and HP should be able to deliver superior perfor-

mance. The RISC vendors may be able to stay close by putting multiple processors on a chip *(see **080605.PDF**)* or by using other techniques.

It is also possible that other RISC vendors are quietly developing their own post-RISC alternatives and have simply chosen to conceal their efforts. These vendors may even be developing the emulation technology needed to offer compatibility with current architectures. Intel and HP would not have gone public at such an early stage had not the spectacle of large numbers of HP and Intel employees spending time in each others' facilities been likely to alert the press to the joint effort.

Any such efforts, though, will require extensive investments, perhaps more than a low-volume RISC vendor can afford. A new architecture will also need backing from major system vendors (preferably more than one) to become established. The IBM/Motorola/Apple triad has the resources to support PowerPC and any follow-ons for some time. Other RISC vendors may, like HP, need partners to help extend their current architecture or create a new one.

Building a Post-RISC Architecture

Note that other RISC vendors can use the same tactic as Intel/HP and offer compatibility with x86 along with their current instruction set. This tactic has the potential advantage of delivering access to the huge base of PC applications, but it greatly increases the difficulty of the design. Such a claim could also strain credibility unless an experienced x86 design team is participating.

As Apple is discovering with its Power Macintosh, merely emulating existing software is not enough to establish a new system; some native software must exist to allow users to take advantage of the full performance of the new processor. ISVs look for a commitment from major system vendors before delivering native applications. System vendors have been reluctant to back new architectures without a strategic reason. Thus, the best way for a processor vendor to establish a new architecture is to sign up one or more major system vendors as partners in the development of the design.

This need for partners leaves system vendors in the driver's seat. Assuming that HP and most x86 PC vendors move to P86 and that IBM and Apple stick with PowerPC, the companies that can make a big difference to a fledgling architecture are Compaq, Digital, and Sun. NEC is the strongest non-U.S. system vendor.

These companies can stay with their current processors and risk becoming uncompetitive in the long run. They can adopt P86 and lock themselves into a single chip vendor. They can move to PowerPC, but that architecture may not provide compatibility with these vendors' current software base.

Partnering on a new architecture is the riskiest solution for a major system vendor but potentially the best. In such a partnership, a system vendor can play a key role in defining the instruction set and implementations for its own needs, giving it an advantage over companies that buy off-the-shelf processors. With backing from one or two major system vendors, a new architecture could attract enough ISV support to prosper.

Digital, in the midst of a money-losing Alpha transition, seems unlikely to lead the development of a new architecture. Sun and Compaq could form a powerful third axis, but Sun seems to be heading toward PowerPC; these two companies also have overlapping product lines and pricing philosophies. If Sun does choose PowerPC, Compaq could still create a multisourced architecture by joining with Silicon Graphics and the MIPS vendors. With Compaq as the lead system vendor and with multiple chip sources, other leading PC vendors might adopt this x86/MIPS architecture rather than P86.

Processor Vendors May Seek Other Markets

At least 16 companies are selling processors for PCs, workstations, and servers, using at least seven architectures. This may be too many vendors, particularly as the number of processor architectures decreases. Some of these companies may choose to focus their efforts—and their architectures—on other markets.

Fortunately, several emerging markets should require large numbers of powerful CPUs in the future. These markets include game machines like Sony's PlayStation *(see **080902.PDF**)*, set-top boxes, and PDAs. Other applications for high-performance embedded chips include printers, networking, and cellular phones. There is little leverage between these devices and general-purpose systems, though, and the dominant embedded architectures may differ from the dominant general-purpose processors. Some RISC architectures, such as MIPS, may abandon the desktop for the embedded arena.

The Intel/HP partnership is an unmistakable signal of coming changes for microprocessor vendors and their customers, most of which face some sort of architecture transition to stay competitive. Many vendors will ignore these signals, hoping to make a change, if necessary, at a much later date. Given the long development cycle of new processors and particularly new instruction-set architectures, chip vendors that wait too long may find their options diminished as software vendors commit to a few top architectures.

System makers can change direction more easily, but those that simply buy processors may find themselves at a disadvantage against system vendors—such as HP, Apple, and IBM—that retain influence over the design and even the instruction set of their CPUs. Forming a third architecture is a viable strategy for a major system vendor, but the first one to act could prevent others from gaining enough software support. For both chip and system makers, change will come all too soon. ♦

