1993 COMPUTER PRESS ASSOCIATION WINNER ♦ BEST COMPUTER NEWSLETTER

MICROPROCESSOR REPORT

THE INSIDERS' GUIDE TO MICROPROCESSOR HARDWARE

VOLUME 8 NUMBER 8 JUNE 20, 1994

Intel, HP Ally on New Processor Architecture

Will Develop Compatible Replacement for x86, PA-RISC by 1998

by Linley Gwennap

Setting the stage for new architecture battles at the end of the decade, Intel and Hewlett-Packard have announced a partnership to develop a next-generation CPU architecture that will eventually replace the current x86 and PA-RISC architectures. The companies stress that processors implementing the new design will be fully compatible with current software. These processors are not expected to ship before 1998, although neither company is willing to discuss product details at this time.

With this announcement, Intel for the first time has admitted that the x86 architecture will run out of steam around the end of the decade. At the same time, the company revealed its plan to solve this problem by switching to a new post-RISC architecture that could deliver performance competitive with or superior to other processor designs in that timeframe. The companies have not yet named this new 64-bit architecture; for the sake of discussion, we'll refer to it as P86 (PA-RISC plus x86).

VLIW, or not VLIW?

Intel and HP are keeping the details of P86 close to the vest and probably have not resolved many of the key issues, as technical teamwork has only recently begun. Previously, HP had discussed using a VLIW (Very Long Instruction Word) instruction set in its PA-9000 processor, due in 1997 or '98; the company had already begun to develop the architecture for this device. Because Intel could have developed a next-generation architecture on its own, we believe that the x86 vendor teamed with HP specifically to gain access to this VLIW technology and accelerate its own development plans.

A VLIW design *(see **080205.PDF**)* would represent a step beyond current RISC and CISC architectures. VLIW is suited to microprocessors with tens of millions of transistors and would thus be appropriate for chips built late this decade. The basic concept of VLIW is to allow the compiler to directly schedule instructions for a large set of parallel function units, removing this burden from the hardware. This change simplifies the hardware design of complex superscalar processors, in theory allowing better performance at a lower cost.

There has been considerable debate about the effectiveness of VLIW, which requires highly sophisticated compiler technology; other alternatives have been proposed *(see **080605.PDF**)*. At HP, Josh Fisher, Bob Rau, and others have been working on VLIW designs for a couple of years after pioneering this technology in the 1980s at startups such as Multiflow and Cydrome. If Intel chooses this technology for P86, it would provide a strong endorsement for VLIW in general and HP's technologists in particular.

By moving directly to VLIW (or another advanced design), Intel would leapfrog the current RISC architectures. If executed correctly, this move could be a stunning coup: for the first time since the debut of RISC, Intel could have an x86-compatible processor with leadership performance. Combined with Intel's legendary manufacturing and marketing ability and its enormous bankroll, such a processor should deeply worry the competition. For its part, HP would get to be on a winning team instead of lagging far behind.

Intel has apparently learned from the fiasco of the 860, its first foray into RISC (see MPR 3/89, p. 1). That chip delivered leadership performance but had little software and system support; after propping it up for a few years, Intel has left the 860 to die. For its next try, the company will ensure that the complete portfolio of x86 software runs on any new processors.

Maintaining Compatibility Is a Challenge

The biggest challenge the partners face is retaining compatibility with not one but two software bases. Intel continues to assert that its processors will be compatible with x86 for years to come and that it won't let this transition stand in the way. Intel points to the 286-to-386 transition as an example of maintaining compatibility while moving from a 16-bit to a 32-bit architecture. The switch to P86 will be more difficult, however.

Most RISC vendors have moved to a 64-bit architecture simply by adding a mode bit that allows instructions to take advantage of extended 64-bit registers and addressing. Although Intel could extend x86 in the same way (as it did in moving to the 386), it still would be stuck with a CISC architecture and its two decades of accumulated baggage. P86 offers the opportunity for a clean break with the past, a blank piece of paper, and a chance to significantly improve CPU performance. This means creating a whole new instruction set optimized for the IC technology of 1998, not 1978.

The sheer size of its installed software and hardware base, however, forces Intel to maintain compatibility with the current x86 instruction set. HP also needs to support its PA-RISC customers. There are several ways to resolve this conundrum, and the partners have not yet revealed their plans. One method would be to support the x86, PA-RISC, and P86 instruction sets in a single chip. This strategy optimizes performance on older software, but the cost of supporting three instruction sets in hardware could overwhelm any advantages of the new architecture and may even reduce P86 speed.

Another tactic would be to use software emulation to execute old x86 and PA-RISC binaries. Apple is bridging its existing software base to the Power Macintosh using this tactic, and Digital is doing the same to help convert its VAX customers to Alpha. But both of these companies are having problems due to the poor performance of pure software emulation, resulting in a slow rate of conversion and the loss of some customers.

Unless Intel and HP develop a magic emulator with exceptional performance, their best choice would be to add some features to P86 that would accelerate the emulation of existing binaries without bloating the hardware design. There are a variety of techniques that can accomplish this feat *(see **080704.PDF**)* and, as we noted in that story, Intel is in the best position to implement them, given its complete understanding of the x86 architecture and its ownership of the needed patents.

Benefits for Both Sides

Although both parties in this new alliance clearly benefit, HP looks like the bigger winner. The company has spent the past few years waffling between offering its PA-RISC chips to other companies and focusing on in-house needs *(see **0710ED.PDF**)*, as well as between setting its processor investment to a reasonable level and trying to match the huge investments of its competitors. Now, HP can have it both ways: the company can spend a reasonable amount of money relative to its unit volume while Intel invests to match its competitors. HP can build chips for its own needs while Intel sells the same chips to many others, creating economies of scale.

HP has generated more profits than any computer system vendor over the past two years, and one key to its prosperity is a single set of PA-RISC platforms sold as both business and technical servers using either HP's legacy MPE software or its version of UNIX. When P86 processors become available, HP can use the same strategy to build a common desktop platform and configure it with different software for the PC and workstation markets. As a bonus, HP workstations will be able to run PC applications with excellent performance. This strategy is ideal as the line between PCs and workstations continues to blur *(see **0804ED.PDF**)*.

Intel's gains are less clear, and some may wonder why the company needs a partner at all, or why it chose HP. HP's RISC system revenues are greater than any other vendor's, but if that company builds its own P86 chips (as it intends to do), Intel won't see any direct revenue from converting HP to its architecture. Intel says that it wants to share HP's demonstrated expertise in designing high-performance CPUs; this sharing includes access to HP's VLIW program.

Intel also gains a strong partner in the workstation and server markets, where HP ranks in the top two in most market segments. Intel has found it difficult to penetrate these markets, but HP's backing should establish P86 as a major player in nearly all general-purpose computer markets, leaving few niches for other CPU vendors to exploit.

Of the major RISC vendors, HP seems most compatible with Intel. The company is already a big Intel customer through its PC business and is the largest buyer of Intel's 960 processors, which are used in HP's LaserJet printers. HP is the only CPU vendor that has not launched an assault on Intel's PC market using Windows NT. Finally, HP had already announced its plans to move to a new architecture in the next few years.

Opportunities for Competitors

For both companies, the hardest part of the allied effort will be to manage the transition to the new architecture. Any discontinuity gives customers an opportunity to defect to other vendors. Many Apple customers today are thinking, "Should I buy a Power Macintosh and get new PowerPC software, or should I get a Pentium box with new Windows software?" HP and Intel must plan carefully to avoid this trap; the new chips must execute old binaries as least as fast as previous-generation parts, and new native software must be plentiful and not carry a price premium.

In addition, the partners must avoid saddling P86 with too much baggage from previous generations. Otherwise, when competitors introduce their own next-generation architectures, Intel will again find itself to be the performance laggard, and will have dragged HP down as well. Thus, the P86 designers must walk a fine line between improving emulation performance and opti-

mizing native execution. One possibility is to make the emulation acceleration features optional so they can later be dropped in favor of full software emulation.

The biggest danger is that Intel and HP could make poor decisions in defining P86. Perhaps VLIW (or whatever is chosen) is the wrong way to go. Perhaps, in stressing compatibility with two existing instruction sets, native performance is hampered and cost is increased. Perhaps emulation performance is so poor that current customers switch to another vendor. Given these risks, HP and Intel have chosen to combine their best minds to reduce the odds of a fatal slip.

Ten Years and Counting

We expect that, in about 10 years, Intel will stop making pure x86 chips in favor of P86 chips. Intel will continue to milk the x86 cash cow as long as it can, and it certainly has the resources to develop both x86 and P86 chips simultaneously, particularly with HP's help on the latter. Intel's P6, due in late 1995, probably will be the last pure x86 core that Intel develops. The company may carry the P6 through two IC process shrinks, boosting the clock rate and reducing its cost. Although the initial P6 is rumored to be a two-chip set, process shrinks eventually will allow it to be combined into a single chip for low-cost versions.

Intel's P7 has been expected to debut in late 1997. We now believe that the P7 will be the first P86 chip, delaying its appearance until 1998. Even as the P7 rolls out, Intel will continue proliferating versions of the P6; it may take five or six years for P86 products to dominate Intel's line. On HP's side, the PA-9000 will probably be merged with the P7 effort. Given HP's limited resources and greater control over its customer base, it will probably move completely to P86 within a few years of that architecture's debut instead of continuing to develop its own PA-RISC processors.

Although Intel's motives for switching to a new architecture are probably rooted in technology issues, the switch also creates a new market dynamic: it will return Intel to the position of being the sole supplier of its CPU architecture, at least for some time. Intel appears uninterested in licensing P86 and probably will take steps to design the architecture to be virtually unclonable, having learned valuable lessons in the past decade. In fact, P86 may be the first architecture created with a lawyer on the design team.

A New Competitive Landscape

With HP giving up on RISC by the end of the decade and Intel continuing to rail against that technology, other RISC vendors must demonstrate a plan to compete with the forthcoming Intel/HP architecture. Some may choose simply to adopt P86 or to seek shelter behind the big boys in the PowerPC camp. Others may attempt to sustain their RISC architectures but now must explain how they will compete against a post-RISC design; otherwise, they may be forced into niche markets.

Intel's x86 competitors also have some difficult choices in the long run. Proliferating x86 without Intel's backing and marketing muscle won't be easy, particularly if Intel is correct in its belief that x86 performance will not remain competitive. Alternatively, cloning P86 will be difficult and expensive at best.

Other vendors may define post-RISC architectures to compete with P86. Unless these companies team up, however, it will be hard for any new architecture to gain critical mass against the Intel/HP partnership and the IBM/Motorola/Apple troika. Digital, still in the midst of its CISC-to-RISC transition, will find it most difficult to change architectures again so soon.

Of all the companies affected by this announcement, surely the PowerPC alliance must be the most distraught. IBM and Motorola dreamed that their processors would continue to soar in performance as Intel's chips ran into a wall at the end of the decade. Even if Intel switched to RISC, they reasoned, it would have to make a nasty transition merely to catch up to PowerPC. Now, it appears that Intel will leap to a new architecture, leaving PowerPC holding the RISC bag. HP wisely chose to join Intel's side. Other processor vendors, busy concentrating on next quarter's results, must now show a five- or ten-year plan to remain competitive. ♦

*Next issue, (see **0809VP.PDF**) we will further address the questions of how this new partnership affects Intel's competitors and what they can do to combat it.*

