

VIEWPOINT

Intel and HP Shoot Each Other in the Foot

New Alliance Cripples Current Product Plans

by Nick Tredennick, PhD.

I used to get compliments from Intel employees who read my opinions. Others wondered if I secretly worked for Intel. This opinion should end the compliments and the speculation. Intel has done something so dumb even I noticed. Intel and HP have shot each other in the foot with the announcement that they will merge x86 and PA-RISC into a common architecture that *Microprocessor Report* calls P86 (see [080801.PDF](#)). P86, which Intel will apparently implement in its P7 processor, will probably be a VLIW derivative that will execute x86 code, PA-RISC code, and (presumably) its own native code with better performance than would have been achieved by any x86 or PA-RISC implementation alone. Faster, better, cheaper. Sure. Another free lunch.

Why It Happened

Why it happened is probably the most interesting aspect of the alliance. Here's my theory. Universities produce two basic kinds of engineers: those that work on engineering projects and those that work on their careers. (I think they are two fundamentally different personality types.) The engineers who work on their careers eventually run companies, and the ones who work on engineering projects work on engineering projects. I'll call the former the engineering executives and the latter the working nerds. As long as the company is run by the engineering executives, there is some technical relationship between management and the working nerds, since the managers have engineering backgrounds. Both time and company growth work against the strength and value of this relationship.

Intel has grown to an age and size that finally broke the technical relationship between management and the working nerds. Intel is large enough and old enough now that some of the managers of Intel's Microprocessor Products Group are business or economics majors with no engineering background. I've never liked the theory that a professional manager can manage any organization, but it seems to be an axiom of business today.

Once this technical relationship is broken, the company begins to rely on designated experts. Designated experts are created by credentials: you get a PhD or you get credit for some major technical accomplishment, then you work in a research organization, a university, or a staff position. Designated experts track the latest fads but can't afford too much time for detail. Designated

experts are the liaison between managers and the working nerds. Managers rely on designated experts for interpretation of the technical environment.

Surf's Up: The VLIW Wave Approaches

Designated experts tend to believe in truth, justice, and the free lunch. At least that's how it looks to me. We muddled through 13 years of the RISC fad and have now entered the "post-RISC" era. As the RISC wave is breaking on the beach, there's a VLIW wave on the horizon. VLIW is positioned well to become the next technofad. The claims being made for VLIW are identical to the claims that helped RISC become a major fad.

The claims say it's faster, cheaper, and smaller. Yet another free lunch. Intel's and HP's designated experts sold their managers on a VLIW derivative as a convergent path to glory. Now it is left to the working nerds to create reality from the dream. The alliance may have been born of apprehension about the perceived momentum of the PowerPC alliance. They should have waited: PowerPC panderers have oversold the technology and probably set expectations beyond what they can deliver.

What It Means

The alliance puts Intel and HP in a weaker position with respect to their competitors than before the alliance. The alliance also weakens the competitive positions of both Intel's and HP's customers.

HP's position is weakened as it changes development focus to a new architecture, which will once again cause a transition for its installed base. Before the alliance, HP was already planning a transition to a new architecture, supporting both PA-RISC and a new VLIW-like architecture. Now, it will be supporting a third architecture, the x86. Costs will rise, performance will decrease, and, despite added resources, product introductions will be delayed.

HP may have decided to align with Intel because HP can no longer afford to develop competitive IC process technology. The partnership ensures continued access to leading process technology without prohibitive process development cost. HP also may reduce processor development cost, since costs will be shared with Intel.

Intel's position is weakened. Whatever its original P7 plans, they are undoubtedly delayed substantially (and probably not for the last time). If its original P7 plan called for the introduction of a new instruction set in addition to the x86, at least it could have been defined

by the Intel working nerds with the transition in mind. The new P7 instruction set could have been to the x86 what the Alpha instruction set is to the VAX. No more. Now the new instruction set has to accommodate both the x86 and PA-RISC. That accommodation can't be free and probably isn't beneficial to Intel.

This is all good news for Intel's competition—both x86 and others. If AMD, Cyrix, NexGen, and other x86 producers concentrate on x86, they can produce better performance on x86 object code than the P86. They can offer continuity and painless migration. Contrary to press pundits' predictions that the x86 is "out of gas," there's still room for improvement (*see 081102.PDF*).

Software application developers will probably not be thrilled with the Intel/HP alliance. They can look forward to supporting an increasing variety of VLIW-like implementations. (Compiler developers might like the prospect of selling a custom compiler for each implementation.) Developers aren't exactly falling over themselves to support everything Windows NT runs on, probably because the additional cost of supporting more than just x86 isn't returned by the fraction of a percent in increased market share. The same is likely to be true if Intel introduces a VLIW derivative in the x86 market while other x86 producers (including Intel, which owns most of the high-end x86 production capacity) continue to produce pure x86 implementations for the market. Developers will ignore fringe implementations and concentrate on serving the volume market.

Customers Are Weakened

The transition weakens both HP's and Intel's customers. VLIW implementations move dependency checking and flow control from hardware to software and dedicate more hardware resources to execution. Hardware provides computing resources, and software provides instruction scheduling and hazard detection. This may be fine for accelerating execution, but it also means the compiler is closely matched to the implementation. Each successive implementation will require recompilation of the applications or suffer the equivalent of an emulation or interpretation penalty. Binaries won't work from one member of the same architectural family to the next. Worse, each new hardware implementation requires new versions of the corresponding software.

Complex hardware is difficult to develop and debug, so high-end microprocessor implementations are always larger and more complex, and they tend to arrive later than expected. And hardware development has a much better track record than complex software development. No matter how good the story is, there is going to be a performance penalty for the transition. Intel's and HP's customers are likely to face this transition, while their competitors (i.e., customers of PowerPC, SPARC, MIPS, Alpha, and conventional x86) may not.

As if dilution of technical focus weren't enough, Intel customers that compete with HP will be further weakened by the alliance, since HP will have advance information about the compilers, operating systems, and chips coming from Intel. Also, HP will have access to P86 derivatives inaccessible to Intel's customers.

For Intel's customers, the best they might hope for is that the Intel/HP alliance has all the impact of the ACE initiative or the Intel/IBM agreement to cooperate in the Robert Noyce Design Center. Intel can certainly afford to spend a couple of hundred million dollars on a VLIW experiment, but it need not lose its focus on the x86. It can afford to do both at the same time. Current Intel customers should be very interested in the product plans of AMD, Cyrix, and NexGen. The x86 is sufficiently entrenched that even Intel cannot move its customer base to a new architecture.

P86 and the Illiac IV

Remember the Illiac IV? There must be some people at Intel and HP that don't. Illiac IV was the product of an overly ambitious computer design project, combining new architecture, new technology, a new design team, new implementation techniques, and new software techniques. It failed. Illiac should have taught the computer design industry a lesson in controlling risk in new projects. I asked John Wharton what he thought. "Iliac? Why bother with Iliac? Intel has examples of their own: the 432, BiiN, and the i860." Good point.

The P86 combines new architecture, new technology, a new design team, new implementation techniques, and new software techniques. Who would like to quote odds on its success? Surely, Intel isn't going to bet its multibillion-dollar cash cow on a single advanced design experiment. Even if it wastes the efforts of the P7 working nerds on the Illiac-like P86 effort, the working nerds from the P6 design can build the P8.

Intel is on the verge of the situation IBM faced once clones of the PC came out. A flood of IBM-compatible PCs entered the market and destroyed IBM's ability to control margins. IBM tried to lead the market away from the clones by introducing MicroChannel. It obviously didn't work. The market was too large and too well established.

Intel would like to lead the market away from the x86 compatibles. Instead of simply introducing a new CPU, it has begun with press releases and a campaign to convince its lead customers that the transition is a good thing. Despite endorsements and investments, it won't work. Intel was the leading light of the PC-compatible CPU market through the introduction of the 486. Intel can still dominate the x86 CPU market by virtue of investments in technology and capacity, but Intel can't control the architecture. The x86-compatible market is too large and too well-established. The x86 has become the light and Intel is just a moth. ♦