# What's Wrong With Merced

## *Implementation Woes May Overshadow IA-64 Advantages*

While Intel was remarkably forthright in disclosing that the scheduled delivery date of its first IA-64 processor, Merced, had slipped to mid-2000 (see MPR 6/22/98, p. 1), the company was not as forthcoming regarding the reasons for the slip. Sources indicate that the project, under way for more than four years, is facing problems that could jeopardize Merced's existence as a viable product. Even if that chip is compromised, however, IA-64 itself is likely to prosper.

Intel's claims that Merced would deliver "industry-leading performance" were based in part on a plan to deliver the chip before most other 0.18-micron processors. This IC process advantage, not an instruction-set advantage, might have boosted Merced beyond competing products in performance. But with the latest delay, Merced will be competing on a more level playing field.

Sources indicate that Merced will debut at a clock speed of about 800 MHz. This sounds impressive, but we expect that, in the same 0.18-micron process, Intel's x86 line will reach 700 MHz, and the 21264 Alpha processor will exceed 1 GHz. So, at least in its initial implementation, IA-64 doesn't seem to offer much of a clock-speed advantage—this despite claims that eliminating complex out-of-order logic would allow Merced to run faster than RISC processors.

Clock speed is not the only factor in performance, of course, but it is generally the most important. We expect Merced to execute up to 6–8 instructions per cycle, whereas the 21164 can execute up to four. In practice, however, processors rarely execute more than two instructions in any given cycle, due to the limited amount of instruction-level parallelism in typical applications. Thus, we expect Merced to get some small benefit from its peak execution rate, but not enough to make up for the 21264's faster clock rate.

Although Intel has not changed its public position, behind the scenes the company and its partners are now downplaying Merced's performance. In fact, some IA-64 system makers are quietly saying "Wait until McKinley," referring to a second IA-64 processor due to ship in 2001. McKinley, the story goes, will be twice as fast as Merced in the same IC process, showing off the true performance characteristics of the IA-64 instruction set.

What will make McKinley so fast? Perhaps we should go back to the question of what makes Merced so slow. My theory is that it is much easier to throw away performance than to improve it. Even if the IA-64 instruction set offers some inherent advantages over RISC, the performance gain of an ideal IA-64 processor over an ideal RISC chip is likely to be small. A weak implementation of IA-64, on the other hand, could easily be slower than a good RISC design, despite any instruction-set advantages.

Sources indicate that the Merced implementation is suffering from the sheer size of the design team. Intel's standard approach is to first have a few top architects define the implementation, then corral hundreds of low-level engineers to design the circuits. This approach may work well on an x86 chip, where long experience has identified most of the key design issues. With a new architecture, however, a top-down style makes it difficult to optimize for clock speed and overall performance, as problems surface only when the entire design is put together and simulated.

We suspect the recent Merced delays are due to just such an event: problems surfacing late in the design process that forced some redesign to meet speed and performance goals. If so, there is no guarantee that this redesign (and this slip) will be the last.

If Merced slips another couple of quarters and McKinley stays on track, the two designs will both appear in 2001. This overlap could lead Intel to not market Merced as a product but use it only as a development vehicle. HP is already talking about a PA-8700, and possibly a PA-8900, to fill the gap until it can switch over to IA-64, and SGI has started to look at developing a MIPS processor to follow the R14000. Other system makers would have to stick with the x86-based Xeon line for a longer time.

Intel must ensure that, unlike Merced, McKinley meets its lofty goals. The company has consistently stated that it is responsible for every transistor in the Merced design. If the processor fails to meet its schedule and performance goals, there is no other place to put the blame. Fortunately, Intel's IA-64 partner, HP, has a proven CPU design team that doesn't follow Intel's top-down approach. That team is expected to have a large role in the McKinley design.

With McKinley already in the on-deck circle, Intel's IA-64 effort is likely to succeed whether or not Merced meets its revised schedule. IA-64 processors will clearly outrun Intel's x86 line while offering a compatibility mode for older software. The workstation and server companies that are already supporting IA-64 have little alternative but to stay the course. Whether that course leads to Merced or McKinley remains to be seen. M