

# Intel, Merced, and the Fate of Microprocessor Forum

Michael Slater



[michael@webvanta.com](mailto:michael@webvanta.com)

# MICROPROCESSOR REPORT

VOLUME 11, NUMBER 14  
OCTOBER 27, 1997

THE INSIDERS' GUIDE TO MICROPROCESSOR HARDWARE

## Intel, HP Make EPIC Disclosure

*IA-64 Instruction Set Goes Beyond Traditional RISC, VLIW*



by Linley Gwennap

Breaking out of the 1980s RISC mind set, Intel and Hewlett-Packard have designed a new instruction set, IA-64, geared toward the highly parallel processors of the next decade. IA-64 goes beyond previous CISC, RISC, and VLIW instruction sets with a new set of features that its creators call EPIC (explicitly parallel instruction computing). This strategy should give Merced, the first IA-64 chip, a leg up on its old-fashioned competitors when it debuts in 1999.

### Compiler Provides Explicit Directions

Modern compilers analyze a program to exploit opportunities for parallel instruction execution, carefully arranging the instructions for optimum performance on today's super-scalar processors. Yet the serial programming model of RISC and CISC instruction sets forces the processor to dynamically re-evaluate the compiled code, instruction by instruction, and perform its own parallelization. Today's instruction sets provide no means for the compiler to communicate parallelism to the hardware, forcing the processor to duplicate this work using complicated out-of-order circuitry.

# MICROPROCESSOR REPORT

VOLUME 12, NUMBER 8  
JUNE 22, 1998

THE INSIDERS' GUIDE TO MICROPROCESSOR HARDWARE

## Merced Slips to Mid-2000

### *Delay Jeopardizes Attempt to Gain Performance Lead*

by Linley Gwennap

Reality has reached out and tossed sand into the gears of Intel's product-development machinery. The company regretfully reported that an eight-month schedule slip has pushed the first volume shipments of its Merced processor from late 1999 to mid-2000. This slip will delay IA-64's penetration of the workstation and server markets and make it more difficult for Merced to achieve the performance lead, as Figure 1 shows. In the long term, however, the delay will probably have little effect on Intel's success in these markets.

to 1999 shipments. At Microprocessor Forum, Intel's Fred Pollack promised that Merced would deliver "industry-leading performance" when it shipped. The design team at this point consisted of several hundred engineers, and the logical design was nearing completion.

Despite (or perhaps because of) this enormous staffing level, keeping the Merced program on schedule continued to be difficult. After a recent schedule review, senior management was shocked to discover that the chip was nowhere near tape out and in fact could not be expected to ship until the middle of 2000. Since Intel had publicly committed to 1999 shipments, it was forced to publicly acknowledge the

# Mercy, Mercy, Merced

*Schedule Slips? Performance Problems? It's the Same Old Story.*

*by Nick Tredennick & Brion Shimamoto*

Despite Merced's recent tapeout, we continue to hear that the project's performance is falling below target, and that the official mid-2000 schedule for system shipments is unlikely to be met. No one should be surprised, however. Unrealistic targets are endemic to a highly visible project managed by executives disconnected from the engineers doing the real work.

## **Leading-Edge Projects: the Theory**

Here's how leading-edge microprocessor projects are planned.

At month zero, executives and engineers meet to negotiate features, goals, and schedules. One input to this negotiation is the current competitive state of the industry. Another input is Moore's Law.

Moore's Law says that if they design a microprocessor on, say, a 36-month schedule, its performance has to be four times that of today's new microprocessors to be competitive. At their month-zero meeting, they plot their first performance target on a graph of performance versus time. Their

## **Leading-Edge Projects: the Practice**

The competition might deliver surprisingly good performance, better than Moore's Law. The executives raise the performance target to allow for that. Also, projects can slip. If a project slips too much, it pushes the performance target below Moore's curve. If this happens, there's no market for the part, so the executives raise the performance target to allow for that, too. In practice, a project's target starts out looking like that in Figure 1.

At the beginning of a project to design a leading-edge microprocessor, the engineers are enthusiastic and confident. They sign up for a lofty performance target and for an aggressive schedule. But then the engineers begin the refinement process that leads to the real design.

Almost from the beginning, the actual performance of the design diverges from what is being "pitched." This happens because, as engineers decide on implementation details, they compromise performance to solve problems affordably. Figure 2 illustrates this situation.

Over time, the gap between "PowerPoint" Performance (what executives are telling each other and their cus-

Unrealistic targets are endemic to a highly visible project managed by executives disconnected from the engineers doing the real work.





# MICROPROCESSOR REPORT

VOLUME 13, NUMBER 13  
OCTOBER 6, 1999

THE INSIDERS' GUIDE TO MICROPROCESSOR HARDWARE

## Merced Shows Innovative Design

*Static, Dynamic Elements Work in Synergy With Compiler*



By Linley Gwennap

At this week's Microprocessor Forum, Intel unwrapped the Merced microarchitecture, showing how IA-64's EPIC design results in hardware that is both simpler and more powerful than traditional RISC or CISC processors. Gone are the complex instruction reorder buffers and register alias tables found in modern superscalar processors. In their place are more registers, more function units, and more branch predictors. These trade-offs eliminate unneeded complexity while leaving some dynamic structures in the hardware to handle events the compiler can't easily predict.

### Six-Issue EPIC Processor

As Figure 1 shows, Merced can fetch and issue six instructions per cycle to a pool of function units that includes four integer units, two FPUs, and three branch units. Two of the integer units can also handle load/store instructions. Additional operations can be achieved using the SIMD integer and FP capabilities, the pointer post-increment feature of the load and store instructions, and the loop-counter update in special branch instructions (see MPR 5/31/99, p. 1).

The ability to handle up to three branches per cycle is unique among announced server processors; in fact, most can handle only one. In IA-64, branch-prediction instructions consume some of the branch slots, but there should be



THE WALL STREET JOURNAL.

MARKETPLACE





