



May 29, 2001

**Page One Feature**

## **Intel Gambles It Can Move Beyond The PC With New Microprocessor**

### **Chip Giant Places Huge Bet on Starting From Scratch With Model for Servers**

**By DAVID P. HAMILTON**

Staff Reporter of THE WALL STREET JOURNAL

**Intel** Corp. has long dominated the world of microprocessors, using its control of the engine of personal computers to build one of the most lucrative franchises in corporate history. Its chips supply the brains for more than half a billion PCs world-wide.

But a decade ago, Intel saw that it was nearing a crossroads. Its fingernail-size chips were based on an aging design, one that would require major revision in order to fulfill Intel's dream of extending its reach beyond PCs to high-end computer servers that power corporate networks and the Internet.

Like a homeowner torn between renovating an old house and building a new one, Intel realized it could spiff up the existing microprocessor line, known as the x86, with a technical facelift, or it could start on a whole new design. After heated internal discussions, Intel decided to plunge into the unknown, forging an ambitious partnership with **Hewlett-Packard** Co. to design an all-new chip.

Tuesday, Intel unveiled the Itanium, the first of a series of processors it hopes will extend its dominance to the entire computing universe and cement its position for decades. The stakes for Intel are high. With prices of its PC processors falling steadily, it badly needs to move to higher ground. The fate of Itanium could determine whether the company's dominance slips as the PC wanes in importance -- or whether Intel helps shoulder aside big servers from **Sun Microsystems** Inc. and **International Business Machines** Corp., the way Intel-based PCs once overran rivals such as **Apple Computer** Inc.

### **Operation Merced**

Developing Itanium, previously known by the code name Merced, has been an intense and unpredictable effort that sometimes teetered on the brink of disaster. Time and again, a project team of as many as 500 circuit engineers, chip architects and software wizards found it had underestimated the difficulty of its task, more than once sinking into a quagmire of complexity with no obvious way out.

Like carpenters forced to build new hammers and saws as they went along, Intel's engineers designed and tested new software tools at the same time that they were sketching out parts of the tiny chip. The team broke into separate groups, each working on one piece without knowing just how they would fit together.



John Crawford

"Everything was crazy," says John Crawford, the chip's chief architect. "We were taking risks everywhere. Everything was new. When you do that, you're going to stumble."

As a result of the many setbacks, the first Itanium chip is two years late, an eternity in the world of technology. What would have been a speedy processor if introduced on time in 1999 now will run only half as fast as Intel's next version of the Pentium 4, which still uses the x86 architecture.

For this first Itanium, expectations are low. Many in the industry think it will be used mainly for testing, with corporate customers waiting for a second-generation chip code-named McKinley that is due out next year. **Microsoft** Corp. has only just completed a limited edition of its Windows operating system tailored for Itanium and doesn't plan to release a final version until the end of the year. Since the Itanium is aimed at servers and

workstations, Intel will continue making Pentium and Celeron chips for PCs.

"This will end up being one of the world's worst investments, I'm afraid," predicts David House, a former Intel official who is now chief executive of Allegro Networks, a network-equipment start-up. Mr. House, Intel's chief of corporate strategy in the early 1990s, worries that Intel will never get back the \$1 billion to \$2 billion that analysts estimate Itanium has cost so far. Mr. House, who approved the project at the time despite his own reservations, says the scale of the Merced project "scared the everloving bejesus out of me." Intel says that his doomsaying is nonsense and that the Itanium family will make it plenty of money.

The Itanium story began in the early 1990s as Intel's designers started to chafe at limitations of the x86 design. The most serious: It processes data in chunks of 32 bits -- each a one or a zero. For technical reasons, that limits the memory of a PC or server to four gigabytes.

The limit is still distant for today's PCs, which average 128 megabytes of memory. Servers, however, use far more memory because they do so much heavy lifting, from storing Web pages to running giant company databases. Memory requirements also double every 18 months or so, making the memory limit a distinct threat to Intel's high-end computing ambitions.

The solution: Intel needed to develop a new processor architecture, or underlying design, that could handle data in chunks of 64 bits. That would make possible memory sizes four billion times as large as with the x86.

A new architecture is a complex ecosystem of devices and programs, so creating a new one involves far more than just designing a single chip. Engineers need to build, generally from scratch, software to help programs run better on the new processor, sets of "assistant" chips for the processor, and batteries of testing and verification programs. Little wonder that a group of Intel engineers argued for simply tweaking the x86 architecture to handle 64-bit data.

By the early 1990s, however, new processors such as Digital Equipment Corp.'s Alpha chip were already showing greatly improved performance with new designs. Mr. Crawford led a faction that said Intel had to push forward with a new architecture or risk trailing competitors "from day one." Technical arguments raged until his group eventually convinced Albert Yu, then general manager for microprocessors, that an aggressive approach was best.

Across Silicon Valley, Hewlett-Packard also had some decisions to make. A power in high-end computing, H-P had long built its own microprocessors but reserved them for its own workstations and servers. Recently, engineers in its corporate labs had designed an advanced processor architecture called PA-WideWord, which promised blazing speeds by letting the chip perform several calculations at once in what is known as parallel processing.

But the costs of chip-making were soaring to the point that a single new manufacturing facility could run more

than \$1 billion. Despite its tradition of going it alone, H-P decided it needed a partner that could share the financial burden and help sell the chip to other computer makers, possibly making it an industry standard.

H-P approached Intel, which was intrigued. Not only did H-P's team include several chip-design luminaries, but Intel also noticed that H-P had a big head start. "When we saw WideWord, we saw a lot of things we had only been looking at doing, already in their full glory," Mr. Crawford says.

At a preliminary technical exchange, says WideWord architect Rajiv Gupta, "I looked Albert Yu in the eyes and showed him we could run circles around PowerPC [an IBM processor], that we could kill PowerPC, that we could kill the x86. Albert, he's like a big Buddha. He just smiles and nods."

Lawyers worked out some ground rules, and by early 1994, technical discussions had begun in earnest. They took place at an out-of-the-way H-P sales office, with no documents allowed out of the room. At each day's end, material was stored in a double-locked filing cabinet, with one key held by Intel's Mr. Crawford and one by H-P's Mr. Gupta. "The idea was that if we didn't do a deal, we would take the filing cabinet to the parking lot and blow it up," Mr. Crawford jokes.

In June 1994, the companies announced a partnership, with Intel leading the design of a 64-bit processor that would use many of H-P's ideas. An exuberant Mr. Yu incautiously declared: "If I were competitors, I'd be really worried. If you think you have a future, you don't."

## **Culture Clash**

But the task was just beginning. H-P officials, accustomed to consensus-driven management, were jarred by an Intel culture of "constructive confrontation" that embraces argument and interruptions. For their part, Intel officials wondered why H-P people wouldn't stand up for their ideas. After some meetings, says H-P designer John Wheeler, "the H-P folks were exhausted, while the Intel guys would be slapping us on the back saying, 'This is the best meeting we've ever had.'"

Managers eventually drew up a 75-page guide to cooperation that explained, among other things, how to interpret the behavior of both sides. Engineers quickly dubbed it the "owner's manual."

Some disputes about technical issues dragged on. H-P engineers argued that certain arithmetic functions known as floating-point operations could be handled outside the chip by software, saving chip space. Intel wanted the functions designed into the chip, so they'd be faster. As the two sides arm-wrestled, Intel suddenly found itself facing a flap over a flaw in its Pentium chip, one that just happened to involve a hardware fault in the floating-point unit. "That just sort of ended the discussion," says H-P mathematician Alan Karp. The team went with the software method.

Once H-P and Intel had hammered out the design basics, the development team expanded rapidly. But with the Internet boom just starting, they were hard-pressed to find experienced engineers. Intel ended up hiring many recent college graduates. "I was one person signed up to design this giant thing that had no idea what it was," says Nadeem Firasta, who joined Intel in 1995.

Managing such a diverse, intense and fast-growing team took a toll on managers. The first to head the project, Avtar Saini, lasted only a year before requesting a reassignment. He now runs Intel's operations in India.

## **Too Big**

His successor, a laconic Missourian named Gary Thomas, joined the project in 1995 -- just in time to see it hit a devastating roadblock.

The chip's architects had divided functions into separate modules, like letting teams of subcontractors design

the rooms of a house. In mid-1996, Mr. Thomas slotted the modules together for the first time in what the team called the floor plan. Bad news: The floor plan was larger than anyone had expected, far too big to fit on a die of silicon that Intel could manufacture economically.

Mr. Thomas maintains that his only emotion was a realization that "it was time to get to work." Mr. Crawford, the chip architect, is less restrained. "We had blown out the walls," he says. "This was a lot worse than anything I'd seen before."

The team found itself sweating through a "die diet" as it worked feverishly to slim down bloated functions and subsystems. Mr. Crawford, who called himself the "chief liposuctionist," searched for cuts that wouldn't hurt performance too much. Internal memory was scaled back, as was a module that maintained compatibility with x86 chips.

## **Downsizing**

But individual modules, initially only rough designs, kept growing larger as they were refined. After months of struggle, senior Intel managers realized that they could solve the size problem only with a radical step: a new manufacturing process that would let engineers shrink every wire and transistor. The change would drop the chip's tiniest dimensions to 0.18 microns -- or millionths of a meter -- from 0.25 microns, making each module smaller.

The switch to a new fabrication process appeared likely to solve most of the Merced project's problems, at the cost of a few months of delay. But the project team soon found itself in a fresh predicament as they worked to tune up the movement of signals across the chip.

In a well-functioning chip, signals flit from module to module in a precisely timed choreography, with the speed of the chip as a whole determined by the slowest signals. Merced engineers started looking for those slowpokes and found ways to speed them up via slight changes to the chip design. Soon, however, it became clear that many of these changes were disrupting the chip's delicate signal ballet, forcing engineers of other modules to rework their designs as well.

The team found itself in a nightmarish world where a change to one module would ripple through the work of several hundred other people, leaving more problems in its wake. If engineers couldn't balance their signals, the only solution would be to slow down the entire chip -- unacceptable for what was supposed to be a groundbreaking design. By mid-1998, the problem had grown so serious that Intel announced the chip would be delayed at least six months beyond its planned late-1999 launch.

Not long after, Mr. Thomas decided to pack it in. "I was tired, and probably close to burnout," he says. He moved to another post within Intel.

## **New Boss**

By coincidence, an Intel engineering manager from Israel named Gadi Singer was visiting the U.S., planning to spend a week at headquarters in Santa Clara and then take his family for a trip to the Rockies. Mr. Yu caught up with him the night he landed and implored him to take over Merced. Mr. Singer made a quick check with his family, agreed, and started at once, not even taking time to pack up his apartment in Israel.

An intense, thickly bearded man given to sketching diagrams on any nearby whiteboard, Mr. Singer was a natural choice, having previously overseen the final stages of the Pentium design. His first priority was team morale, which had slumped badly as the signal-timing problem kept the design from "converging." He stepped up training programs for engineers to give them a break from their obsessive focus on design issues and started posting a list of "Merced babies" born to team members. One night he ran out to Toys "R" Us, returning with bags of Nerf toys to help people blow off steam.

Mr. Singer also shuffled managers' duties and demanded new information systems that helped engineers see more quickly the ripple effects of design changes. Slowly, Merced began to converge.

The design was finally completed on July 4, 1999. The next month Intel produced the first physical chip, which worked right off the assembly line. Up to that point, engineers had to run test computations on a simulator, but

### AMD's Plans

Itanium's challenges aren't over. While Intel had made sure the new chip would run older applications made for the x86 line, it is likely to do so more slowly than equivalent x86 chips. Intel's remedy has been to spend hundreds of millions of dollars encouraging software developers to tailor applications specifically to the new architecture, which they're doing.

In addition to H-P, **Compaq Computer Corp.**, **Dell Computer Corp.** and IBM have thrown their support behind Itanium, but customers will have to be convinced that Itanium systems can demonstrate the sort of reliability and security now available on high-end servers. That will take time, so few analysts expect an Itanium rush.

Meanwhile, rival **Advanced Micro Devices Inc.** is designing its own 64-bit x86 chip, although that project has

Now that the agonies of the Merced development are behind them, Intel officials say it's been worth the effort, and they've learned valuable lessons. Mr. Yu, for instance, says Intel has learned to limit the initial euphoria of its chip designers and to avoid packing a new development project with too many untested engineers.

Mr. Crawford, the architect, has felt the stings of Itanium's many setbacks more personally than most. "We're headed toward the thrill of victory, but we suffered a number of the agonies of defeat," he says.

**Write to** David P. Hamilton at [david.hamilton@wsj.com](mailto:david.hamilton@wsj.com)<sup>1</sup>

---

#### URL for this Article:

<http://interactive.wsj.com/archive/retrieve.cgi?id=SB991082952420872249.djm>

#### Hyperlinks in this Article:

(1) <mailto:david.hamilton@wsj.com>

---

**Copyright © 2001 Dow Jones & Company, Inc. All Rights Reserved.**

**Printing, distribution, and use of this material is governed by your Subscription Agreement and copyright laws.**

**For information about subscribing, go to <http://wsj.com>**

Close Window