

Intel Outlines High-End Roadmap

Willamette, Foster Extend IA-32 Line; New Details on Merced, McKinley



by Linley Gwennap

In an unprecedented outburst of openness, Intel has sketched out its plans for several new high-end processors that will roll out over the next four years. The roadmap features a mix of IA-64 and IA-32 (x86) processors and includes the first details of the next-generation Willamette core and a server version of that core, code-named Foster. The flow of new products should improve Intel's competitiveness in the high-end workstation and server markets until McKinley provides the *coup de grâce* in late 2001.

As the primary provider of PC processors, Intel has historically kept its long-range roadmap deeply hidden. But as the company has pushed into new high-end markets, with Xeon and ultimately Merced, it has found the price of entry to be additional disclosure of its plans. Vendors and even buyers of expensive systems plan two or three years ahead; they need assurances from Intel that its products will be competitive. The leading high-end system vendors, such as Sun (see MPR 10/5/98, p. 15), often publish their long-range plans.

At the Microprocessor Forum earlier this month, Intel vice president Steve Smith bore the latest news. He described Foster, a new IA-32 processor, as delivering integer performance similar to Merced's at about the same time. Although this information seems to undercut expectations of IA-64's superiority, he indicated that the second IA-64 processor, code-named McKinley, will far outperform any IA-32 processor when it ships in 2H01.

Cascades Moves L2 Cache on Die

After the new 450-MHz Pentium II Xeon (see MPR 10/26/98, p. 4), the next processors for workstations and servers will use the Katmai processor core. This core, which will be used mainly in PCs, appears to be little more than the current Pentium II with the Katmai New Instructions (see MPR 10/5/98, p. 1) added. Katmai's new memory block-move instructions will be useful in servers. Workstations, however, will see little benefit from KNI, which accelerates the single-precision floating-point math used in 3D games but not the double-precision FP used in high-end technical applications.

The Katmai processor will be added to the Xeon line in the form of a module code-named Tanner. We expect this product to appear in early 1999 at a clock speed of 500 MHz,

the same as the first Katmai processors. Tanner will be compatible with today's Slot 2 Xeon products, offering a small clock-speed upgrade. Like the current Xeon parts, Tanner will use an external full-speed level-two (L2) cache in sizes ranging from 512K to 2M.

In 2H99, Intel will deploy a new processor known as Cascades. This processor uses a 0.18-micron version of the Katmai core, which we expect to reach speeds of 700 MHz or more. Unlike Tanner, however, Cascades will use an on-die L2 cache. Moving the L2 cache onto the processor die reduces its latency and makes it easier to run at the full CPU speed. It also reduces manufacturing cost by eliminating the costly custom cache chips that Intel fabricates for today's Xeon. Intel will probably deploy multiple versions of Cascades,

with L2 cache sizes including 512K, 1M, and possibly larger. The larger cache sizes, however, may not appear until early 2000.

The challenge for Cascades is whether Intel can add enough L2 cache for large server applications without exceeding the maximum reticle size, which limits chips to about 500 mm². We estimate the die size of a Cascades with 1M of L2 cache to be about 250 mm², moderately large but quite manufacturable. A version with 2M of L2 cache, matching the largest size in the current Xeon line, would require nearly 400 mm², still well within the reticle limit and smaller than HP's PA-8500 (see MPR 10/26/98, p. 4).

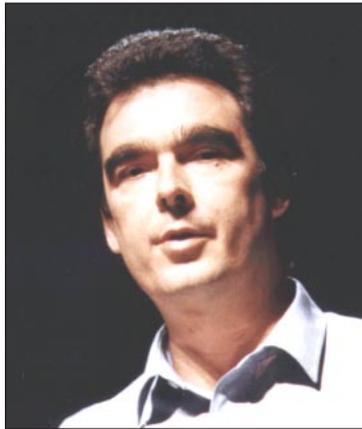
Although the enormous die size would reduce yield, this effect can be partially mitigated by incorporating redundancy into

the large cache array, a tactic Intel is already using in its Mendocino processor. The 2M Cascades would cost at least \$200 to build, but a 2M Xeon with four \$90 cache chips costs about \$450 to manufacture, according to the MDR Cost Model. Thus, the move to on-die L2 cache makes financial sense, even for large caches. Intel must hope, however, that its customers will not demand caches larger than 2M.

Willamette Provides New Core

Intel's roadmap provided the first official confirmation of the long-rumored Willamette (see MPR 11/18/96, p. 4), the company's next-generation x86 core. This processor is a completely new core, scheduled to appear in late 2000. If Willamette meets this schedule, it will be the first new x86 core from Intel in five years, a surprisingly long duration.

The company provided few details about Willamette's design. It said the new core is "superpipelined" to achieve



MICHAEL MUSTACCHI

Intel VP Steve Smith provides the first disclosure of Merced's internal architecture.

high frequency. Since the P6 already sports a lengthy 12-stage pipeline, this statement implies the Willamette pipeline is long enough to carry Alaskan oil. The target frequency for Willamette is at least 1 GHz in a 0.18-micron process, about 40% faster than the P6 pipeline is likely to achieve in the same process.

One reason for Intel to revamp the pipeline is that clock speed is one of the most effective levers in increasing overall performance. Adding execution units helps, but only if programs have enough parallelism to utilize the extra units. A second reason to aim for higher clock speeds is that megahertz sells. Or in this case, gigahertz. Unsophisticated consumers frequently buy PCs on the basis of the processor's clock speed, not performance, and Intel doesn't want another CPU vendor to break the gigahertz barrier first.

To keep the pipeline flowing even in programs with many branches, Willamette will include a trace cache. This unit will store instructions in the order they are executed, regardless of branches. In contrast, the P6 instruction cache delivers 32 bytes per cycle, but if the fifth byte is a branch instruction, the remaining instructions in the line are discarded and the processor must fetch a new line.

The trace cache is ideal for x86 programs, where there are often only a few instructions between branches. Furthermore, because Willamette's trace cache will store decoded instructions, these instructions can bypass the initial stages of the pipeline and feed directly into the execution core, improving performance. This method is particularly helpful in a long pipeline; the P6 requires six stages to fetch and decode instructions, and Willamette presumably needs more.

To further reduce pipeline stalls, Willamette will have "advanced branch prediction." The P6's two-level prediction method was innovative when it was released but has since fallen behind. Willamette presumably takes into account more recent advances in the art.

Foster Brings Willamette to Servers

Intel did not disclose any details of the Willamette processor, which is designed to bring the new core into the high-end PC market sometime in 2H00. The company chose to focus on a second processor, code-named Foster, that will bring the Willamette core to the high-end workstation and server markets. Foster will succeed Xeon and Cascades, as Figure 1 shows, by combining the Willamette core with 1M or more of on-die L2 cache and a new 3.2-Gbyte/s system interface. This powerful interface will deliver four times the bandwidth of the current Slot 2.

Foster should reach the same 1-GHz clock speed as Willamette but will debut slightly later, around the end of 2000 or possibly the beginning of 2001. With its large L2 cache and high-bandwidth system interface, Foster should fare well in high-end servers. Intel is developing the Colusa chip set to support systems with up to four Foster processors; some system makers may deploy systems with eight or more Fosters using custom chip sets.

Merced Includes Fast FP MAC Units

Smith's presentation contained an eclectic collection of details about Merced without providing a complete picture of the chip, which is still several months from tapeout. For instance, he declined to provide any information about Merced's clock speed and pipeline or how the chip fetches, decodes, and dispatches instructions. He did not discuss the integer units or the load/store architecture.

He did disclose a few details about the floating-point units. One key goal of Merced and IA-64 is to fix the poor floating-point performance of Intel's x86 processors, which falls well behind that of high-end RISC chips. Merced will include two fully pipelined FP MAC (multiply-accumulate) units. In contrast, today's Pentium II has one FP add and one FP multiply unit, the latter not even fully pipelined. As a result, Merced will be able to execute up to four FP operations per cycle (two MACs) at any precision up to 80 bits.

In addition, Merced will support SIMD FP operations similar to those in the Katmai New Instructions. Whereas KNI adds a set of eight 128-bit registers to the existing eight 80-bit FP registers in x86, we expect all of the 128 FP registers in IA-64 will be 128 bits wide. Thus, any IA-64 FP register can hold either a 32-bit single-precision (SP) value, a 64-bit double-precision (DP) value, an 80-bit extended-precision value, or a set of four SP values. Intel did not disclose a dual-DP mode.

Merced uses the standard FP MAC units to handle two of the four operands in a SIMD instruction. The IA-64 chip includes two additional single-precision MAC units to handle the other two elements in a SIMD operand. As a result, the SIMD instructions are supported by adding only two relatively small SP MAC units. Using the SIMD instructions, Merced can compute up to eight SP operations per cycle, twice as many as Tanner.

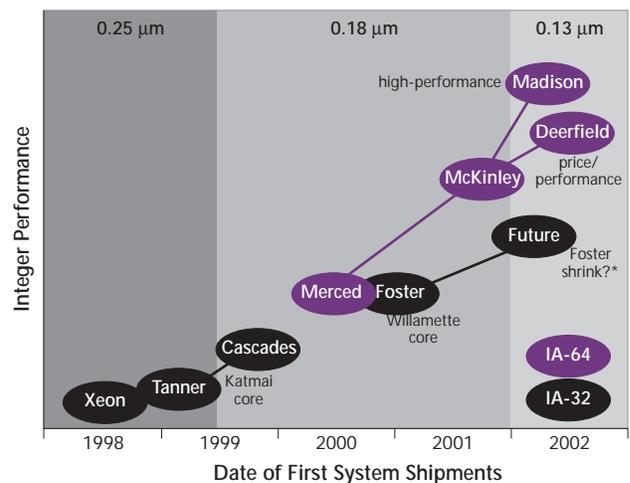


Figure 1. Intel's high-end workstation and server roadmap through 2002 shows the IA-64 line starting about even with IA-32 (x86) in performance but opening a lead once McKinley appears. (Source: Intel, except *MDR)

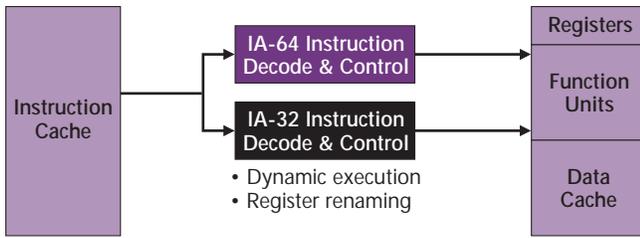


Figure 2. IA-64 and IA-32 (x86) instructions share the same cache but are routed to different instruction control units. Once the instructions are decoded, they are executed using a single pool of function units and a single data cache. (Source: Intel)

P6-Like Decoder Provides x86 Compatibility

Smith also clarified the method of x86 compatibility in Merced. IA-32 and IA-64 instructions will cohabit a single instruction cache. Depending on the state of an internal mode bit, instructions fetched from the cache are routed to either an IA-64 decode unit or an IA-32 decode unit, as Figure 2 shows. Once the instructions are appropriately decoded and converted to a standard internal format, they are executed by a single set of function units, using a common physical register file and a common data cache.

This arrangement, which is very similar to that disclosed in Intel’s U.S. patent 5,638,525 (see MPR 3/31/97, p. 16), avoids the need for separate caches and execution units for x86 code. In addition, x86 instructions can take advantage of the speed and flexibility of the IA-64 execution resources, particularly the fast floating-point units.

In this design, the chip’s x86 performance is determined by how many native operations per cycle the IA-32 decode unit can generate. A simple one-to-one translator

would provide poor performance. To reach Intel’s goal of “mainstream PC performance” on x86 code, the company has chosen a decode unit similar to the P6 front end. This unit can decode multiple x86 instructions per cycle, dynamically reordering them and renaming their registers. Although this is a significant piece of logic, it consumes less than 10% of the Merced die, as Figure 3 shows.

Merced Requires External L2 Cache

Contrary to expectations, Merced has only modest amounts of on-die cache. The on-die cache is divided into small instruction and data caches backed by a larger unified cache. Intel calls the primary caches “L0” and the secondary cache “L1.” Even the primary caches require a two-cycle access, indicating the chip has a fairly deep pipeline. Intel did not disclose the exact size of the caches. We expect the total amount of on-die cache to be less than 256K.

This relatively paltry amount of on-die cache requires a large high-speed external cache. One advantage of the L0/L1 notation is that it allows Intel to call the external cache an L2, making it consistent with the current Xeon parts. Smith said Merced will achieve more than 10 Gbytes/s of bandwidth to the external cache, using an interface running at the full CPU speed. We expect the interface to be 128 bits wide, indicating CPU speeds in excess of 625 MHz. (We expect the initial shipments will operate at 700 to 800 MHz.)

Merced will be delivered in a cartridge that, like the current Xeon module, will contain the CPU chip along with one or more Intel-built SRAM chips. Assuming Intel converts the current Xeon SRAM chip to its 0.25-micron process to gain the speed needed for Merced, it would double in capacity to 1M per chip. With two SRAM chips per cartridge, Merced could support 2M of L2 cache, although Intel is likely to deploy a 1M version at a lower price.

Intel did not disclose details of Merced’s system bus but admitted it would have significantly lower bandwidth than Foster’s 3.2-Gbyte/s bus. The two buses will require different chip sets (the 460GX for Merced, Colusa for Foster) but are “protocol compatible.”

Despite Merced’s lower bus bandwidth, Merced systems will support larger numbers of processors than Foster will, according to Intel. The Merced bus will support more deferred transactions and other undisclosed performance features. Smith also stressed Merced’s reliability features for enterprise-class servers. For example, the L1 and L2 caches will be protected by ECC, as will the L2 cache bus and the system bus.

Although Intel is counting on features such as these to position Merced above the x86-based Foster, that task may be difficult. Intel says the two processors will have similar performance on many integer and floating-point applications, even if they are recompiled for Merced, and the x86 chip will have better system bandwidth. Merced will mainly have an edge on applications that take advantage of its 64-bit addressing. If Foster carries a lower list price than Merced, it could discourage Merced’s adoption in servers.

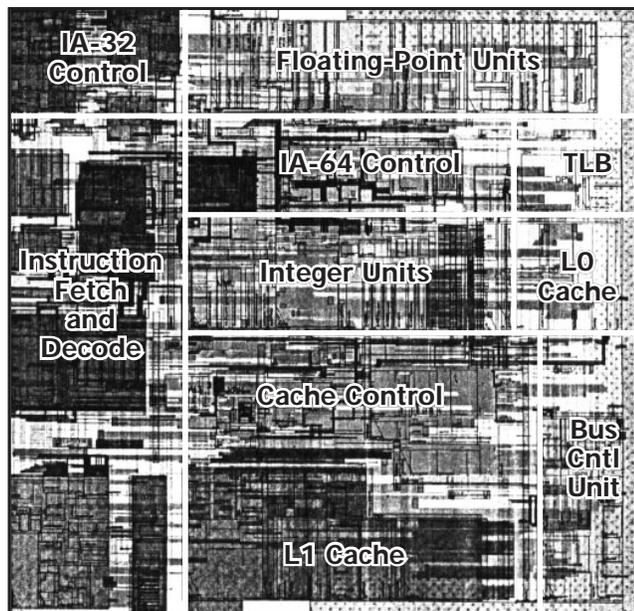


Figure 3. This die plot of Merced, with an overlay from Intel, shows the IA-32 compatibility unit is relatively small. The lack of a pad ring indicates C4 bonding. Intel did not disclose the die size.

McKinley Towers Over Merced

Intel also disclosed a few details about its second IA-64 processor, code-named McKinley. That chip is due to ship in 2H01, using the same 0.18-micron IC process as Merced, although it will quickly shrink to 0.13-micron. Although it is set to appear less than 18 months after the first IA-64 chip, McKinley is said to be everything Merced is not.

For example, Intel says McKinley will operate at 1 GHz or more, making it faster than Merced. Yet McKinley will achieve these faster clock rates using a shorter pipeline than its predecessor. McKinley includes more execution units than Merced and delivers far more performance, yet its core is actually smaller than Merced's in the same IC process.

These contradictions defy the traditional rules of microprocessor design. Either the McKinley implementation is extremely good, or the Merced implementation is extremely poor. We think the answer lies somewhere in the middle, but closer to the latter extreme. Intel designers stress "how much we've learned from Merced." Even an experienced cook will often burn the first pancake trying to judge the griddle's temperature, then get the second one right. Merced is starting to look crispy around the edges.

The fact that Foster's x86 core can deliver the same performance as Merced's IA-64 core indicates that either IA-64 has no advantage over x86 or that Merced's implementation weaknesses are larger than IA-64's advantages. Despite the supposedly simpler nature of IA-64, the Merced core will achieve lower clock speeds and is probably physically larger than the Foster/Willamette core, again indicating inherent inefficiencies in the Merced design. The IA-64 chip's small split-level on-die cache looks to be a replay of the Alpha 21164's weakest area. Finally, Merced's lack of system bandwidth is potentially embarrassing.

If McKinley meets its goals, however, it should display the true worth of the IA-64 architecture. If it can deliver twice the performance of Merced (and therefore Willamette) with the same or slightly more die area than Willamette's core, it would validate the superiority of IA-64 in both performance and cost/performance. McKinley solves Merced's lack of on-die cache and will offer three times the bus bandwidth of its IA-64 predecessor.

McKinley should thus drive a stake into Merced. Intel's next-generation CPUs are typically faster than their predecessors but also bigger. This imbalance results in an extended transition as the company trades off cost for performance. Once McKinley is available, however, there will be no room in the product line for the bigger, slower Merced core.

Madison and Deerfield Advance IA-64 Line

Although the McKinley core may be smaller than Merced's, the McKinley die may not be: Intel plans to pile at least 2M of on-die cache onto McKinley, probably boosting the die size beyond 400 mm². Within a few months after the new chip's release, however, Intel should have its 0.13-micron process available. In this process, McKinley could reach speeds of

For More Information

Intel has not announced price or availability for any future processors. For more information, check www.developer.intel.com/design/processor/future/ia_road3.

1.6 GHz while shrinking well below 300 mm², even with 2M of on-die cache. Intel will offer this part, called Madison, for high-end workstations and servers, and it could be tempted to build a 4M version to sell at even higher prices.

Intel's roadmap also includes a "price/performance" IA-64 processor, code-named Deerfield. This part could be a 0.13-micron McKinley core with a 1M cache and possibly a different bus interface to reduce cost. These changes could cut the die size to 200 mm² or so, about the same as Klamath, the original Pentium II. Such a part would be quite suitable for midrange workstations and servers as well as, perhaps coincidentally, high-end PCs. If Intel sells Deerfield at the same price as today's least-expensive Xeon, IA-64 could ease into the high-end PC market as early as 2002.

Ultimately, McKinley could be the core that succeeds Willamette in the PC market, offering a significant performance upgrade to those users willing to convert to IA-64. By 2004, McKinley will move to a 0.10-micron process, further reducing its die size and enabling high-volume production. Any new IA-32 core that Intel could deploy in this timeframe is likely to be larger than McKinley and offer weaker native performance. Whether Intel chooses to follow this course, however, remains to be seen.

Attacking the Server Space

Intel's extensive roadmap for its high-end processors makes it clear that the company is serious about taking over the workstation and server markets from its RISC competitors. Tanner and Cascades will offer moderate improvements for those system makers that have already adopted the Xeon line. These 1999 products will remain in Slot 2, providing a simple upgrade path.

In 2000, Foster will provide a large performance boost due to its new Willamette core but will require a new slot. In the same year, Merced will give system makers their first look at IA-64. Except for 64-bit addressing, however, Merced offers no apparent advantage over Foster, making the IA-64 chip mainly a software-development vehicle for McKinley. This second IA-64 processor is the one likely to ignite a full-fledged transition to the new instruction set in the workstation and server markets.

For the RISC makers to withstand this attack in the long run, they must offer better performance than IA-64. Given Merced's apparent shortcomings, this may be possible, at least at first. But McKinley looks to be a much tougher competitor that should give even the fastest Alpha processor a run for its money—and the money of its customers. \square