Intel Demos Merced, Discloses Chip Set

At Intel Developer Forum, just two weeks after first silicon came out of the fab, Intel demonstrated two Merced systems—one running a 64-bit version of Windows 2000 and the other running Linux and the Apache Web server. The company declined to disclose the clock speed or give any indications of performance. Intel does seem to have raised its performance expectations from a year ago (see MPR 10/26/98, p. 16); a new roadmap projects Merced's performance to be somewhat above Foster's rather than about the same. Intel also claimed a four-way Merced system will deliver better transaction-processing performance than an Alpha box with four 1.1-GHz 21264 CPUs.

The demonstration did nothing to validate these claims but did show that the first silicon, as well as the systems and software to support it, are in good shape, given the early stage of the chip's evolution. Later this year, hardware and software developers will begin working with real silicon, which will help them prepare for the Merced processor launch, planned for the middle of next year.

The demonstration system used the 460GX chip set, which Intel is developing for servers and workstations with up to four IA-64 processors. This chip set will enable OEMs to easily deploy Merced systems. Several large system vendors, however, are developing their own Merced chip sets to support greater numbers of processors and other differentiating features.

The 460GX is similar in structure to Intel's current high-end chip set, the 450NX (see MPR 7/13/98, p. 11), but requires twice as many chips. For pin-count reasons, the 460GX north bridge is implemented as two chips, the SAC (system address controller) and the SDC (system data controller). These chips connect one or two memory cards, each of which contains eight DIMM slots. Each card requires two MAC chips and four MAD chips.

The new design supports up to 64G of memory, eight times the capacity of 450NX systems. The peak memory bandwidth is 4.2 GBytes/s, twice what we expect of the Merced bus and far more than delivered by the 450NX's EDO memory. Intel did not disclose the new chip set's memory configuration; we expect it to use four banks of PC133 SDRAM.

The 460GX enables high availability, with support for ECC on the memory and system buses; deconfiguration of failed memory chips ("chip-kill"); and error scrubbing. Although Intel is moving to Rambus DRAM for its PC chip sets, this memory is not appropriate for servers, due to its limited capacity and lack of cost-effective chip-kill support.

The chip set includes up to four PCI bridge chips, each driving one 64-bit PCI bus or two 32-bit buses at speeds up to 66 MHz. This design delivers a maximum of 2.0 GBytes/s of I/O bandwidth. For greater system uptime, the 460GX supports PCI parity checking and hot-plug card replacement.

The total number of chips in a full 460GX implementation is thus 18, versus 9 for the 450NX. Intel has not disclosed pricing for the new chip set, which will not ship until mid-2000, along with the first Merced systems; it is likely to cost more than the $400 Intel charges for a full 450NX set today. Intel also offers a stripped-down version of the 450NX, for about $200, that supports less memory and fewer I/O channels; a similar strategy is likely for the 460GX.

With its greater memory capacity, improved bandwidth, and greater I/O support, the 460GX offers a big step forward from Intel's existing four-processor chip set, making it well suited to the higher performance of Merced. Along with the 460GX, Intel is also developing a four-processor motherboard code-named Lion. With these off-the-shelf technologies, OEMs will be able to deploy powerful Merced workstations and servers with little upfront investment. We expect many system makers to take advantage of this opportunity, resulting in a plethora of relatively inexpensive Merced boxes by the end of next year. —*L.G.*

NGIO, Future I/O Merge

In a surprise conclusion to more than six months of contention, Intel's NGIO effort (see MPR 3/29/99, p. 14) and the competing Future I/O architecture have been merged to form a unified standard for next-generation I/O, initially on servers. The new System I/O initiative is supported by Intel, Microsoft, and five major system OEMs. Microsoft is perhaps the most important addition to the list of sponsors, since the System I/O architecture is likely to require significant OS support.

System I/O will use a physical layer much like that of NGIO, consisting of one or more pairs of 2.5-Gbit/s unidirectional links. Each pair yields 500 MBytes/s of peak data throughput, plus overhead. Initial implementations are expected to include one, four, and twelve pairs, for up to 6 GBytes/s of bandwidth. These links will be optimized to carry relatively large data packets and should be very efficient for communicating with networking and storage subsystems.

Version 1.0 of the System I/O specification is due out by the end of the year, and products are expected in 2001. The new standard should eventually supplant Fibre Channel and similar interfaces. —*P.N.G.*

Compaq Dumps NT for Alpha

Finally admitting that Alpha has no chance on the desktop, Compaq pulled the plug on Windows NT for that platform. With new CEO Michael Capellas struggling to get costs in line, the 100 employees handling porting, validation, and support for NT-on-Alpha were an obvious target.

After Compaq's announcement, Microsoft quickly followed suit. Ironically, all development of the new 64-bit version of NT is being done on Alpha systems, but that version will be shipped only on IA-64 systems. Whereas NT was once supported on three RISC platforms, it is now available only on x86, just like its DOS-based predecessors.

Compaq continues to develop new Alpha chips and to support its Tru64 Unix and OpenVMS operating systems for those processors. The company needs Alpha for the high end of its product lines, at least until Intel ships McKinley in late 2001. With a new cost-cutting CEO, however, Compaq's long-term commitment to Alpha remains a question mark.

Also questionable is Samsung's Alpha Processor (API) subsidiary (see MPR 6/21/99, p. 19). Although the company continues to supply most of Compaq's Alpha processors, its hopes of gaining significant revenue from other customers will suffer without NT support. The only viable OS strategy for API is now Linux, which is growing in popularity, particularly for small servers, but remains a small part of the market. Already in a difficult position, API must rework its strategy just as Merced is beginning to make headway. —*L.G.*

Rapid Pentium III Price Cuts Continue

Like a machine gun, Intel is squeezing out price cuts at a blistering pace. The company today cut prices on its two fastest parts, the 550-MHz and 600-MHz Pentium III, for the third time in three months. At $615, the 600-MHz part now lists for 25% less than at its July introduction; the price of the 550 has fallen 36% during the same period. Although the price of the Pentium III-500 didn't change today, it dropped 41% just three weeks ago, as the table below shows.

	1,000-piece list price as of			Change from	
	7/18/99	8/22/99	9/12/99	7/18	8/22
Pentium III-600	$824	$669	$615	–25%	–16%
Pentium III-550	$658	$487	$423	–36%	–13%
Pentium III-500	$423	$251	$251	–41%	0%
Pentium III-450	$230	$183	$183	–20%	0%

While these dramatic cuts come just weeks after the introduction of AMD's Athlon (see MPR 8/23/99, p. 1) at the same clock speeds, Intel says this is just coincidence. Instead, the company is seeing excellent yields on its soon-to-be-announced Coppermine, the 0.18-micron Pentium III, and is dropping prices to prepare for a rapid transition to the new device. We expect the current versions to see another 10–15% price cut when Coppermine appears in October. According to Intel, Coppermine will achieve speeds of at least 700 MHz (presumably meaning 733 MHz) at this introduction.

	1,000-piece list price as of			Change from	
	6/6/99	8/1/99	9/12/99	6/6	8/1
Celeron-500	—	$167	$153	—	–8%
Celeron-466	$147	$114	$99	–33%	–13%
Celeron-433	$113	$93	$79	–30%	–15%
Celeron-400	$93	$73	$64	–31%	–12%

Having gone a full six weeks since the last price move, the Celeron line received a trim, with prices dipping 8–15%, as the table above shows. The fastest Celeron, a 500-MHz part, now lists for $153, while the 366 falls off the price list, replaced by the Celeron-400 at $64. Celeron prices should remain flat, perhaps with another small trim, for the rest of the year.

Intel adjusted its mobile prices last week, with most prices falling 30–40% from the previous quarter, as the table below shows. The top-of-the-line Mobile Pentium II-400 now carries a list price of just $358, leaving room for Mobile Coppermine processors as fast as 500 MHz to appear in October. Similarly, the top of the Mobile Celeron line is now just $106, well below its usual $180 mark, which leaves room for faster parts in that segment as well.

	1,000-piece list price as of		
	6/13	9/5	%Change
Mobile Pentium II-400	$530	$358	–32%
Mobile Pentium II-366	$316	$187	–41%
Mobile Pentium II-333	$187	$161	–14%
Mobile Celeron-400	$187	$106	–43%
Mobile Celeron-366	$144	$85	–41%
Mobile Celeron-333	$106	$74	–30%

Coppermine will reduce Intel's manufacturing cost for its Pentium III line by about 40% while enabling much higher clock speeds. Intel is determined to pass along the savings to its customers—and to AMD's as well. —*L.G.*

S3 Preps Savage2000 for New Millennium

With its announcement of the Savage2000, the first mainstream 3D chip with an integrated geometry engine, S3 has regained the position of technology leadership it lost in the industry transition to 3D graphics. The new chip, expected to enter production this month, is substantially faster than S3's popular Savage4 (see MPR 2/15/99, p. 14).

The company recognizes that geometry acceleration will not be of great importance to end users for the next six months or so, since there are currently very few 3D games that can benefit from it—chiefly OpenGL titles based on Id Software's Quake engine. S3 expects more developers to take advantage of geometry acceleration over the course of 2000 as more vendors offer chips with geometry processors.

The rendering engine on the Savage2000+ can draw two pixels per clock at 200 MHz, with one or two bilinear-filtered textures per pixel. The chip's 400-Mpixel/s dual-textured throughput is better than that of Nvidia's competing GeForce 256 (see next item), but the Nvidia chip is faster in other modes.

S3 is prepared for the competition: it plans to sell the Savage2000 for just $29, while the Savage2000+ will list for $35. These prices are meaningful for only a few board makers, such as Number Nine, that use these chips in niche markets. S3's purchase of mainstream board vendor Diamond Multimedia means that the Savage2000 will not be offered to direct competitors. —*P.N.G.*

Nvidia GeForce Offers Acceleration

Close on the heels of S3's Savage2000 comes Nvidia's own geometry-accelerated PC graphics chip, the GeForce 256. Fewer details are available about the Nvidia product, but it looks to have rendering performance comparable to that of the Savage2000. Nvidia says the slowest speed grade of the GeForce 256 will run at 120 MHz and issue up to four pixels per clock for a pixel-fill rate of 480 Mpixels/s. The company has yet to determine the maximum speed of the new design.

Nvidia estimates the new chip's polygon throughput to be at least 15 Mpolygons/s. S3 has made no specific claims regarding the Savage2000's geometry performance; either chip could be faster on real-world applications. Nvidia says the new chip provides less than a 2× gain on Quake 3 compared with its current TNT2, however, suggesting that substantial optimization remains to be done.

As its name implies, the GeForce 256 includes a 256-bit rendering engine. This engine is coupled to a 128-bit local-memory controller that supports up to 128M of DDR SDRAM (vs. 64M of SGRAM on the Savage2000). The 23-million-transistor chip is being built in a 0.22-micron five-layer-metal process. Nvidia has not announced pricing or availability for the GeForce 256, but we estimate the chip price is as much as $50 for the fastest versions. Boards should be available for the Christmas buying season. Nvidia also announced deals with six board vendors and five major PC OEMs, showing that the new chip is already a major force in the industry. —*P.N.G.*

SGI Reorganizes, Allies With Nvidia

In its latest restructuring, SGI is backing away from its recently introduced Windows NT workstation effort, separating its Cray supercomputer unit, and transferring a key 3D engineering group to Nvidia.

Though it has been just eight months since SGI introduced its NT-based Visual Workstation systems, the company now plans to create a joint venture with an as-yet-unnamed PC maker to develop and market these machines.

SGI intends to retain its NT-based servers, including the SGI 1400 family announced last month, as well as forthcoming servers based on Intel's Profusion chip set (see MPR 8/23/99, p. 22).

As part of the restructuring, SGI announced plans to transfer the engineering team responsible for the next generation of SGI's midrange workstation 3D hardware to Nvidia, once the current development project is completed. After the transfer, Nvidia will become responsible for developing the chips SGI will use in subsequent product generations.

Nvidia will use the same chips in its own family of professional PC CAD accelerators. Nvidia says it sees little risk in the plan, despite the current glut of pro-3D vendors in the PC market (see MPR 9/13/99, p. 18). We believe SGI will cover most of Nvidia's expenses during the development of the next generation of chips, greatly reducing Nvidia's cost of entry into this overcrowded market. —*P.N.G.*