400-MHz K6-2 Nears Top of Intel's Line

Coming within one speed grade of Intel's Pentium II line, AMD has rolled out a 400-MHz version of the K6-2. By beginning volume shipments at this speed, the company fulfills the goal set when it announced a 350-MHz part last quarter. Compaq has already announced new Presario PCs using the fast K6-2.

The 400-MHz part is designed for the "Super 7" socket, a 100-MHz version of Socket 7. AMD is also offering a 380-MHz version that runs with a 95-MHz bus. Sources indicate yield at 400 MHz is limited, so some vendors may have to take the slightly slower version. The K6-2/400 carries a list price of $283, while the 380 lists for $213, providing a further incentive for PC makers to take the slower part. (All prices are in 1,000-unit lots.) We do not expect both the 380 and 400 to appear in the same PC product line.

Creating further potential for confusion, AMD also announced a 366-MHz K6-2 that plugs into the old 66-MHz Socket 7. This part will deliver less performance than the 350-MHz K6-2 on most benchmarks, since the 350-MHz part uses a 100-MHz bus. But the K6-2/366 lists for $187, slightly more than the K6-2/350 at $160. Again, PC makers are likely to choose only one of these parts for their product lines, depending on the bus speed they want to support.

Intel's Pentium II-400 has a list price of $375, so AMD is maintaining its 25% discount off Intel's price. AMD has not released performance data on the new part, which may lag the Pentium II slightly due to its reliance on a 100-MHz L2 cache, half the speed of the Pentium II's cache. But applications that take advantage of AMD's 3DNow instructions will fare better than on the Pentium II.

AMD plans to close the cache-speed gap with its Sharptooth part, which integrates a full-speed 256K L2 cache within the confines of the Super 7 socket. This part is slated to appear in 1Q99 at speeds up to 450 MHz. The K7 (see MPR 10/26/98, p. 1) should follow shortly thereafter, in 2Q99. That part will use the new Slot A, which enables external L2 caches at the same speeds as in Pentium II.

One challenge in establishing Slot A is getting chip-set support. At Comdex, both Acer Labs (*www.ali.com.tw*) and VIA (*www.via.com.tw*) announced plans to develop Slot A chip sets, although neither stated a timeframe for deploying them. Initial K7 systems are likely to use AMD's own Slot A chip set, which will be available simultaneously with the release of the CPU. —*L.G.*

Cyrix Alters Socket Strategy

Cyrix has shifted its strategy, which earlier this year called for a complete focus on integrated solutions for its next-generation processors, and now plans to develop both integrated and standard-pinout versions. The Cayenne core, a third-generation enhancement of the 6x86 design, was planned to debut only in the MXi (see MPR 12/8/97, p. 16), which combines the CPU with a 3D graphics unit, PCI bridge, and SDRAM controller. The MXi, originally due by the end of this year, has slipped to April 1999.

The MXi will be followed later in the second quarter by a new addition to the roadmap, code-named Jedi, which is a 100-MHz Socket 7 version of the Cayenne core. Cyrix has previously stated that it expects Cayenne to reach a clock frequency of 350 MHz in National's 0.25-micron CMOS 8 technology. Unlike AMD's Sharptooth, Jedi will not incorporate any L2 cache on the die.

Cyrix's next-generation CPU core, Jalapeno (see MPR 11/16/98, p. 24), is set to debut in the M3 in 1H00. Like the MXi, the M3 is an integrated part but with an on-chip 256K L2 cache and dual Direct RDRAM interfaces. Cyrix won't disclose any other specific plans, and it remains unclear whether Jalapeno, like Cayenne, will be offered in a standard pinout.

Last year, Cyrix hinted that a Slot 1 version of Cayenne was in the works, but after being absorbed into National, Cyrix's strategy shifted to a complete focus on integrated processors. Recognizing that this approach would abandon too much of the market unnecessarily, Cyrix now says it will offer its processors with the interfaces its customers demand. Given Intel's massive push to the P6 bus, and the debut of a new low-cost PGA (Socket 370) for Celeron, this interface is likely to be Cyrix's next target. The Cayenne core, if combined with on-chip L2 cache, could make a viable low-end Socket 370 offering in 2H99; in 2000, Cyrix will have the Jalapeno core to keep that product line competitive. —*M.S.*

Fujitsu Shipping Hal's Sparc64-III

Beating its original speed and performance targets, Hal's Sparc64-III is now shipping in Fujitsu Solaris-based servers. The new GP7000F models, which will begin shipping in 1Q99, support up to eight processors at clock speeds up to 275 MHz. An entry-level system with a single 225-MHz Sparc64-III is priced at ¥2,300,000 (about $20,000). The systems take advantage of the error-checking features built into the Hal chip to deliver high reliability.

When introducing the chip at the 1997 Microprocessor Forum (see MPR 12/8/97, p. 19), Hal's Hisashige Ando projected the Sparc64-III would reach 250 MHz and deliver 14 SPECint95 and 18 SPECfp95 (base). Although Fujitsu has not yet published SPEC results for the GP7000F, we estimate the 275-MHz CPU will reach 15 int/25 fp, based on currently available data. Hal expects its chip to reach 300 MHz by mid-1999. Big increases are slated for the next two years: 500 MHz in 2000 and 1 GHz in 2001, although the latter may require a modified version of the current CPU core.

To date, Fujitsu has seen little return from its large investment in Hal (*www.hal.com*), because only a small

fraction of its systems contain Hal chips. The new SPARC servers represent a great step forward: Fujitsu expects to ship 30,000 systems with more than 100,000 processors. The company says it will use Hal parts in more future systems as well. With Fujitsu's other SPARC house, Ross Technology, now shuttered (see MPR 6/22/98, p. 5), Hal takes on a bigger role in its parent's plans. —*L.G.*

Intel Reveals Next-Generation I/O Plans

Intel has finally published the specifics of its next-generation I/O (NGIO) initiative. Contrary to previous rumors (see MPR 10/5/98, p. 4), NGIO is not intended as a replacement for PCI, nor is it functionally similar to the Direct RDRAM interface. Instead, NGIO is meant to spark a radical change in server architecture, away from monolithic systems and toward multiple chassis connected by a fabric of serial links.

Each NGIO link provides one serial interface in each direction and is strictly point to point, not bused. These links carry packetized requests and responses. Packets consist of multiple cells and can be as long as 2^{32}–1 bytes. The cell format is reminiscent of a cross between TCP/IP and ATM, with complex header fields and a variable-length payload (up to 256 bytes). A complete transaction includes a request packet, a data packet, and a reply packet.

The NGIO physical layer is essentially that of today's Fibre Channel at 1.25 or 2.5 Gbits/s. Like Fibre Channel, NGIO uses 8B/10B coding, so the effective data rates are 125 Mbytes/s and 250 Mbytes/s each way. These peak rates are comparable to today's server-oriented PCI implementations, but NGIO will be more efficient on long data transfers due to a lack of bus-arbitration overhead. The new interface can span a distance of about 20 meters with copper links or about 100 meters with optical-fiber links. NGIO also supports link bundling, where multiple physical links are used to create a single virtual connection.

The best analogy to an existing I/O interface technology is the high-performance parallel interconnect (HIPPI), which runs today at 800 Mbits/s and is being upgraded to 6.4 Gbits/s. Both NGIO and HIPPI sit between the host and remote peripheral devices. Intel envisions NGIO as a way to connect servers to nearby cabinets of disk drives, LAN concentrators, and other high-performance I/O subsystems. NGIO supports switching controllers that would sit between multiple hosts and multiple I/O subsystems to support clustering and peripheral sharing.

NGIO is meant to eliminate conventional backplane buses, but it has characteristics that may prevent it from doing so. NGIO's packet and cell structure imposes high overhead and latency for short interactive communications, making it a poor match for many common PC I/O devices.

Intel is developing an NGIO bridge chip code-named VxB for its server chip sets as well as a separate PCI-to-NGIO bridge known as PVxB. Both of these chips will begin sampling in late '99, and systems with NGIO are scheduled for 2H00, about the same time as Merced.

NGIO clearly doesn't conflict with the 1-Gbyte/s PCI-X proposal working its way through the PCI SIG. Nor will it conflict with the 2-Gbyte/s post-PCI expansion interface that the PCI-X advocates are rumored to be working on, which is also supposed to have a switched architecture—but aimed at backplanes, not chassis-to-chassis communication. It appears that NGIO is designed for a very limited set of applications, but they may be sufficient to justify Intel's efforts. —*P.N.G.*

ATI, Silicon Motion Introduce Laptop 3D Chips

In the year since S3 introduced the first laptop graphics chip with 3D acceleration and integrated DRAM, no systems have shipped with the ViRGE MXi. However, embedded DRAM has been very successful for the 2D-only notebook graphics chips from NeoMagic. Now the second wave of 3D+DRAM devices has arrived, with new chips from ATI and Silicon Motion, and these chips are likely to give NeoMagic its first real competition.

Silicon Motion's Lynx3D is a straightforward extension of the company's earlier 2D-only chips, which have achieved a few design wins but little market share. The chip includes 2.5M of integrated 128-bit-wide SDRAM and can connect to 2–4M of additional discrete SDRAM through a 64-bit external interface. Though low-resolution 3D displays can be generated from the on-chip memory, most laptops today have enough screen resolution to require the use of external memory for best results with 3D rendering. The $49.50 Lynx3D can drive LCD, CRT, and TV monitors through three independent outputs. Silicon Motion (*www.siliconmotion.com*) offers the Lynx3D in 272-contact and 316-contact BGA packages. The company says the chip consumes just 300 mW (typical).

The ATI Rage Mobility-M takes a new approach to graphics-memory integration. Although ATI (*www.atitech.com*) has not revealed details of the $39 Mobility-M's internal configuration, company representatives have implied that the chip's 4M of integrated SDRAM is not on the same die with the graphics controller. Instead, sources say, the memory is implemented as two discrete 1M × 16 SDRAMs within the 328-contact BGA package. Another 4M of SDRAM may be connected externally. With 60% more memory than the Lynx3D, the Mobility-M consumes 600 mW, twice the power drain of the Silicon Motion chip.

ATI also announced the $32 Rage Mobility-P, which appears to use exactly the same graphics chip, but without integrated memory. Both chips include RGB, video, and low-voltage differential signaling (LVDS) LCD drivers and are pin compatible with ATI's current Rage Pro LT.

The Lynx3D and the Mobility-M both address the most critical problems facing notebook designers: board space and power consumption. Both Silicon Motion and ATI should be able to compete effectively with NeoMagic, especially with the addition of 3D to attract OEMs and end users. —*P.N.G.*