IBM to Build Chips for AMD

In an effort to solve its long-running production problems, AMD has signed IBM as a foundry for the K6. Using only its own Fab 25, the company has not been able to meet its production targets for the past three quarters. The IBM agreement will not offer immediate relief, however, as it will take some time for the design to be brought up on the IBM production line; the first shipments of IBM-made K6 chips are not expected until the fourth quarter.

This transition would take even longer if not for the similarities between IBM's and AMD's processes. The K6 was originally designed by NexGen to be built on IBM's fab lines. When AMD bought NexGen, it made some modifications to the K6 design but for the most part revamped its own fab lines to match IBM's process technology. Indeed, these extensive process modifications probably play a large role in AMD's current production problems.

Perhaps the most frustrating part of AMD's problems has been the strong demand for the K6. If only the company could build more chips, it could surely sell them. Yield on the 0.35-micron parts has increased over time, but the new 0.25-micron K6 is not yet yielding well. In fact, AMD has not even announced availability of the 0.25-micron version, since all units are currently being consumed by two large customers, Compaq and IBM (see MPR 1/26/98, p. 1).

IBM will use some of the chips it builds for its own PCs; the foundry will also supply chips to AMD for sale to its other customers. IBM does not have the right to resell the K6 under its own label, which would be complicated by its role as a vendor of Cyrix's competing 6x86 line.

The move seems a tacit admission by AMD that it cannot guarantee it will resolve its fab difficulties in the foreseeable future. Last year, the company claimed that it would manufacture and sell 15 million K6 chips in 1998, but more recently it has backed down from this number. AMD expects Fab 25, and its sister facility under construction in Germany, to be world-class facilities competitive with any of Intel's fabs, but now the company must look outside to gain adequate production capability. This situation is not desirable; AMD is likely to pay more to IBM for K6 chips than it would to build its own chips.

AMD may still be hoping to get Fab 25 operating at peak efficiency this year; the IBM deal could be a backup plan in case the fab improvements don't occur. Customers nervous about AMD's fab situation may be reassured by the IBM arrangement.

The foundry deal could also portend a closer relationship between AMD and IBM in the future; the two could, for example, work more closely in the 0.18-micron generation, perhaps using the K7 design. In the near term, however, it appears AMD has merely found a more expensive, albeit hopefully more reliable, way to build K6 chips. —*L.G.*

Tanner Bridges to Slot M

Sources indicate that Intel is developing a processor, codenamed Tanner, that is designed to bridge the gap between the company's high-end x86 products and Merced. Tanner, due to ship in 1999, is said to incorporate an x86 processor core, probably Katmai, along with an interface to the so-called Slot M interface that will be used by Merced. If Tanner ships before Merced, which we now expect to appear in 2H99, it could act as a development vehicle for Slot M systems, allowing workstation and server vendors to prepare and market "Merced-ready" products.

Even after Merced ships, Tanner could be a useful product. System makers could market high-end workstations or servers based on Merced while using Tanner in the midrange, all with a single motherboard and system design.

Tanner should also provide better performance than Slot 2 products using the same x86 core. We expect Slot M to include a 128-bit bus, twice the width of Slot 2's bus. Slot M is also likely to run at 200 MHz, providing a peak bandwidth of 3.2 Gbytes/s, enough to match the bandwidth of a two-channel Direct RDRAM memory system.

In 1999, Slot 2 is likely to move to 200 MHz (or 100-MHz double data rate) as well, doubling the 100-MHz speed it will use at its introduction later this spring. Even at this speed, however, Slot 2 will max out at 1.6 Gbytes/s, the equivalent of one Direct RDRAM channel. While this performance is likely to satisfy the needs of most systems in that timeframe, high-end workstations and multiprocessor servers can make good use of the greater bandwidth of Slot M. —*L.G.*

Samsung May Design Alpha Chips

Further illustrating the passing of the Alpha torch, Samsung has acquired from Digital the right to design its own Alpha chips in the future. The Korean company's original license (see MPR 7/8/96, p. 4) gave it only the right to second-source Digital's own Alpha designs. But with Alpha's future at Digital in doubt, Samsung now has the right to carry on the Alpha legacy without Digital's help.

Digital's pending sales of its semiconductor group to Intel and the rest of its business to Compaq have cast doubt on the long-term viability of Alpha (see MPR 2/16/98, p. 4). The recently completed 21264 is likely to last Digital through at least two process shrinks, but Samsung wants to drive Alpha into low-cost markets using products like the 21164PC. The companies had discussed a future 21264PC product; it now appears likely that Samsung will develop this product, or at least fund its development.

To date, however, the market for such low-end Alpha processors remains tiny. Despite unimpressive results, Samsung appears ready to spend more money to further its goal of entering the PC processor market. Mitsubishi, in contrast,

is standing pat with its Alpha license. That company has announced no new Alpha plans, but if Samsung's efforts succeed, the Japanese vendor could step back into the market. At this point, however, Samsung's Alpha quest appears quixotic at best. —*L.G*

Nvidia Regains 3D Speed Lead

Nvidia has unveiled an improved version of the RIVA 128 dubbed the RIVA 128ZX. The company says the 128ZX surpasses the performance of the Intel740 (see MPR 2/16/98, p. 1) and 3Dfx's Voodoo 2, achieving a score of 781 on Ziff-Davis's 3D WinBench, compared with 691 and 724, respectively, for these competitors.

The 128ZX, sampling now with production scheduled for 2Q98, incorporates various internal performance enhancements, resulting in higher sustained pixel-fill rates. The peak rate remains at 100 Mpixels/s. The new chip also increases the RAMDAC clock rate from 206 to 250 MHz, enabling support for 1600 × 1200 × 24 displays with an 85-Hz refresh rate.

The 128ZX supports up to 8M of local memory, breaking the 4M ceiling of the original RIVA 128, and includes a full 2× AGP interface that boosts performance compared with the 1× implementation found in the RIVA 128. While the current RIVA 128 is almost always sold with 4M of local memory, Nvidia's customers are likely to offer both 4M and 8M configurations with the 128ZX, providing additional product differentiation.

The RIVA 128ZX will arrive with new Direct3D and OpenGL drivers that also support the RIVA 128, yielding moderate performance improvements for the older chip. The 128ZX will be priced at $32 in 10,000-unit quantities. This is slightly below Intel's list price for the 740 and in line with other performance leaders in the mainstream market. The original RIVA 128 will remain available at a lower (but undisclosed) price.

Nvidia (*www.nvidia.com*) has inked an agreement with chip foundry TSMC to provide an alternate source for the RIVA 128ZX and other Nvidia products. Nvidia will continue its strategic relationship with primary vendor SGS-Thomson, an architectural licensee that comarkets and sells Nvidia-designed graphics chips. —*P.N.G.*

3Dlabs Signs SGS-Thomson as Second Source

SGS-Thomson has more 3D business than just Nvidia's. It is also a strategic partner on Chromatic's Mpact media processor and recently announced that it will be a second fab for the popular Permedia 2 graphics chip from 3DLabs (*www.3Dlabs.com*). The 3Dlabs relationship is simpler than the comarketing deal with Nvidia; SGS-Thomson will manufacture—but not market—3Dlabs graphics chips.

3Dlabs is currently engaged in a legal battle with its primary fab vendor, Texas Instruments, over alleged misappropriation of 3Dlabs trade secrets; finding an alternate source may have been a prerequisite for this action.

SGS-Thomson is an unexpected choice, as it has had difficulty making enough RIVA 128 chips to satisfy Nvidia. It's unclear whether 3Dlabs will actually benefit from the new agreement, or whether it is just trying to hang on to enough fab capacity to meet demand. In either case, the deal gives SGS-Thomson more options in the graphics industry. —*P.N.G.*

Next-Generation PowerVR Planned

VideoLogic clearly has no worries about partner NEC's ability to produce silicon. The two companies have begun to describe the next generation of PowerVR graphics chips; the new roadmap includes five derivatives to appear in 1998 alone. The new chips will target three different markets: PCs, game consoles and set-top boxes, and arcade games. PowerVR will be in the next-generation Sega game console (see MPR 6/2/97, p. 5) and has long been a serious player in the arcade-game market.

The new PowerVR chips should be more popular in PCs than NEC's previous PCX1 and PCX2 PowerVR products, which were 3D-only accelerators requiring a separate 2D graphics adapter. Integrated 2D acceleration and higher performance—rated at 2–5× the speed of the PCX2—will be combined with an easier-to-use programming interface that game developers are sure to appreciate.

Including 2D support will give PowerVR better access to the lucrative, high-volume world of motherboard and OEM designs. These customers have shown little interest in 3D-only accelerators apart from a few home-entertainment systems such as Gateway's Destination series PC.

PowerVR's tile-based architecture has been both a help and a hindrance to developer support for first-generation chips. Even though the supplied drivers include software routines that can hide most of PowerVR's unique features, peak performance can be achieved only by exposing and using these features. The new PowerVR chips accelerate these routines in hardware, making it easier to get good performance from applications that are not written specifically for PowerVR. This includes commercial DirectX and OpenGL titles, which represent the bulk of the game software market.

The new PowerVR architecture will also include advanced quality features such as 32-bit floating-point Z computations, bump and environment mapping, and full-scene antialiasing. These features will greatly improve visual quality even compared with other new 3D products, including those from Intel, 3Dfx, and Nvidia.

VideoLogic/NEC have not announced specific configurations, pricing, or availability for second-generation PowerVR accelerators. We expect the first such chips to appear later this year. More information will be available on the Web at *www.powervr.com*. —*P.N.G.* M