EVSX Emerges From Exponential Ashes

When we last checked in on Exponential Technology (see MPR 6/2/97, p. 4), the company had shut down its operations except for a 30-person team in Austin (Texas) that was working without pay, reportedly on an x86 processor. Since that time, the now-defunct company has auctioned off its patent portfolio to pay off creditors and spun out the Austin team as a new company, now called EVSX.

The patent auction, conducted in August, raised several million dollars. In addition to some bipolar-circuitry patents, the company held several microarchitecture patents, including one for a dual-mode processor with CISC and RISC instruction sets (see MPR 6/23/97, p. 18). The identity of the new patent owner remains a secret, although the company will eventually be identified when the patents are reregistered with the U.S. Patent Office.

EVSX recently announced it has received a $14 million contract from an undisclosed corporate partner, rumored to be Texas Instruments, to develop "high-performance microprocessors for visual computing." The new company is led by Paul Nixon, who oversaw Apple's role in the Somerset design center before heading Exponential's Austin team. Jim Blomgren, lead architect for Exponential's failed PowerPC chip (and former x86 architect at Chips & Technologies), is also part of the new venture, but most of the other participants in the original PowerPC design did not transfer from San Jose to Austin.

Even before Exponential collapsed, the Austin team was rumored to be developing an x86 processor (CMOS, not bipolar). It isn't clear whether "visual computing" is merely a smokescreen for the x86 effort or whether the contract is keeping the company going while it seeks funding for the x86 effort. We suspect EVSX will join Rise, Transmeta, IMS, Metaflow, and others working on x86 processors that have not yet seen the light of day. —*L.G.*

DRAM Vendors Gear Up for 440BX Chip Set

While Intel has yet to announce its forthcoming 440BX system-logic chip set for Pentium II processors, the company has released specifications for a critical new aspect of the 440BX, its 100-MHz main-memory interface. Vendors have already begun sampling dual inline memory modules (DIMMs) designed to meet the new specifications. One vendor, Smart Modular Systems (*www.smartm.com*), has even volunteered the information that Intel has begun sampling the chip set itself.

The 440BX is expected to be similar to the current 440LX chip set except for the faster DRAM interface and the CPU bus, which is also expected to move to 100 MHz. Greater bandwidth between processor and main memory will help multimedia applications; in particular, host-based DVD decoding will finally be possible at a full 30 frames per second. While the 440BX's AGP interface will have the same features as that in the 440LX, bandwidth-hungry AGP texturing operations can take immediate advantage of the faster SDRAM.

Intel's PC SDRAM specifications may be found at *developer.intel.com/design/pcisets/memory*. Currently, they include version 1.51 of the PC SDRAM definition, intended to define the subset of the full JEDEC SDRAM standard that Intel's chip sets require; version 0.9 of the PC SDRAM unbuffered DIMM specification, which defines the electrical and mechanical requirements for 168-pin 3.3-V 100-MHz DIMMs with and without ECC support—also an Intel-specific version of a JEDEC standard; and revision 1.2 of Intel's serial-presence-detect (SPD) standard, which allows Intel chip sets to determine the size and configuration of installed DIMMs. —*P.N.G.*

Acer Offers AGP Chip Set for Socket 7 CPUs

Though no Socket 7-compatible processors are yet shipping with 100-MHz bus speeds, Acer Labs (*www.acerlabs.com*) has already introduced an AGP chip set that supports the faster clock rate. The Aladdin V chip set, available in versions for desktops and notebook systems, centers around the M1541 north bridge, which includes a 100-MHz processor interface, an L2 cache controller with integrated tag RAM, a 66-MHz main-memory controller, a 2×-mode AGP interface, and a 33-MHz PCI bridge.

Matching a 100-MHz bus interface with a 66-MHz SDRAM controller makes sense in the Socket 7 universe, where L2 caches can benefit from the higher speed without boosting the cost of the whole main-memory subsystem. This configuration won't be seen for P6-bus processors, however, as the only reason to increase that bus's clock is to improve bandwidth to main memory. AMD and Cyrix have committed to 100-MHz Socket 7 bus interfaces (see MPR 10/27/97 p. 20), while Centaur's C6+ may do the same; Intel, however, will not. This limits the Aladdin V to non-Intel Socket 7 CPUs, a large but shrinking market.

The desktop configuration comes with the M1543C south bridge, which includes an integrated super I/O block, while the mobile Aladdin V provides the M1533 south bridge. Both versions provide the normal complement of other interfaces, including IDE and USB, and comply with the advanced configuration and power-management interface (ACPI) standard, a requirement for 1998 systems. —*P.N.G.*

Intel Changes NC Into Lean Client

Trying to respond to market demand below the low end of its PC processor line, Intel is developing specifications for what it calls lean clients. The company offered few details for its vision, but the lean client will apparently be a stripped-down PC running Windows CE instead of Windows. The plan is

for these systems to include an embedded Pentium processor, of which Intel is currently offering two: a 100-MHz version and a 166-MHz version, neither with MMX. These processors carry a list price of $85, hardly a bargain.

These sketchy specifications sound much like the network computers that other vendors have been pushing, none with significant success. Intel expects lean clients to be used as terminal replacements and in single-function applications such as for bank tellers and purchasing agents. Intel is working with several major software and PC makers on the specification, but none has committed to delivering systems based on the lean-client document.

We believe customers who want PC compatibility are more likely to buy a full-fledged PC selling for much less than $1,000; these systems often have non-Intel x86 chips today. Customers who just want to read e-mail and access the Web don't need PC compatibility; these users are candidates for a Windows CE box, but such a system could use a RISC processor that delivers more performance at a much lower price than an embedded Pentium. Thus, we doubt the lean client will fare much better than the ill-fated NetPC, Intel's previous foray into reducing cost of ownership. —*L.G.*

Intel, Sun Share Patents

In contrast to its litigious counterparts, Sun Microsystems agreed to share its patents with Intel without a lawsuit. The two companies announced a patent cross-license agreement that will allow both vendors to design future high-end processors without fear of each other's patents. Intel already has such agreements with HP, Digital, and IBM; as it moves further into the high-end space with Merced and other IA-64 chips, access to patents regarding high-performance CPUs and systems becomes more critical. With access to Intel's vast patent portfolio, Sun will also benefit.

As part of the same announcement, Sun confirmed that it will port its Solaris operating system to IA-64, following its current strategy of supporting both Intel and SPARC processors. NCR had revealed in October its plans to use Solaris on Merced, so the Sun disclosure was somewhat anticlimactic. Sun hopes Solaris will provide an alternative to Windows NT on IA-64 systems, but so far Solaris-on-x86 has offered little competition to NT, and we doubt Solaris-on-Merced will fare better. Most Unix system vendors are sticking with their own proprietary operating systems, and those vendors seeking a cross-platform OS are gravitating to Windows NT.

Sun denied any plans to offer IA-64 systems in the future and reiterated its commitment to SPARC. With an IA-64 version of Solaris in hand, however, Sun could easily change its mind in the future. —*L.G.*

NEC Virtually Improves the SDRAM

We didn't think there was room for yet another DRAM interface in a market already glutted with them. SDRAM, DDR SDRAM, CDRAM, EDRAM, MDRAM, VRAM, SGRAM, RDRAM, and Direct RDRAM all compete for our attention on the basis of various theoretical advantages in their core or interface designs. NEC, however, has proposed a new way to design and use DRAMs that provides benefits similar to those offered by each of the memory architectures listed above while adding only slightly to DRAM die size and testing cost.

NEC's Virtual Channel Memory (VCM) architecture, to be implemented first in an SDRAM-like device called the VC-SDRAM, breaks the strict physical relationship between DRAM-array banks and the external interface. Instead, VC-SDRAMs have 16 virtual banks, or channels, each of which can be assigned by the DRAM controller to a different—and possibly overlapping—region of the physical DRAM array, which may itself consist of multiple banks.

Like Mitsubishi's Cached DRAM (CDRAM) (see MPR 2/15/93, p. 18), VC-SDRAMs hold multiple rows of array data in registers with very low access latency. Unlike the CDRAM cache, however, the registers in the VC-SDRAM are explicitly managed by the controller, so VC-SDRAMs don't need the complex address comparators used by CDRAMs to detect cache hits. This complexity is instead shifted to the controller, where it can be centralized.

When used with a controller that can assign the virtual channels to sources of memory requests—like the individual applications in a multitasking PC, or screen refresh in a graphics controller—VC-SDRAMs should achieve channel hits much more frequently than conventional SDRAMs achieve page hits, decreasing average latency and increasing effective bandwidth. This technique is well suited to a 3D graphics controller, which typically accesses several unique streams (drawing operations, texture maps, screen refresh, etc.) that are easily identifiable. Tracking data streams from multiple software applications using PC main memory would be more challenging.

NEC says the die-size penalty for VC-SDRAMs compared with SDRAMs in the same process is only about 3%, with a similar increase in test time; these costs are low enough that the company expects to offer VC-SDRAM products at price-per-bit parity with SDRAM. NEC is offering licenses for the VCM technology at no cost, a practical decision considering the intense competition. No other DRAM vendors have announced support for VCM at this time.

First silicon of the VC-SDRAM is expected in April 1998, with second sources and chip-set announcements also due by that time. The new VCM enhancements make the most sense when combined with double-data-rate (DDR) SDRAMs, boosting the peak transfer rate as well as the bus efficiency; such products are expected in 2H98.

The VCM parts appear most likely to break into the graphics and embedded markets, which are smaller than the PC main-memory market but more flexible. Access to the main-memory market is unlikely unless Intel accepts the technology. NEC's investment in this new design has been fairly small, however, and could easily be justified without the high volumes of main memory. —*P.N.G.* M