Most Significant Bits

Intel 8-Mbit Flash Hit by One-Year Slip

Customers of Intel's 8-Mbit flash memory have been hit with a delay of up to one year in availability of this highly successful product, causing a potentially catastrophic slip for some customers. Intel will not accept new orders for the 8-Mbit flash chips, or for the 4-, 10-, and 20-Mbyte memory cards that use them, until 4Q93. Intel claims to have 85% of the flash memory market, so its production problems will be widely felt.

The lengthy production delay stems from high demand coupled with the failure of Intel's fabrication partner, NMBS in Japan, to achieve acceptable yields and write endurance. Intel began manufacturing the chips in its Santa Clara D2 development line, and about 350,000 chips were produced. Intel expected that this would be enough to meet demand while NMBS ramped up its production. NMBS produced its first 8-Mbit chips for Intel in late 1991, but serious problems cropped up that still are not completely solved. Meanwhile, the D2 line was torn apart and rebuilt to serve as Intel's 8" wafer development line, making it impractical to build any more parts on that line.

Intel also planned to shift its lower-density flash parts, currently built on a 1-micron process in Fab 7 in Albuquerque, to the NMBS line, but the difficulties there have prevented them from doing so. As a result, production is constrained on these chips as well. Fab 7 is now being converted to the 0.8-micron process so the 8-Mbit parts can be built there, but it will not come on-line until the second half of '93.

The flash problems are rippling into other areas as well. Intel has pulled out of the EPROM business to provide more capacity in Fab 7 for flash production. Under pressure from Intel, NMBS is pulling out of the DRAM market so it can focus all its resources on flash.

Intel is also working with Sharp on its next-generation, 0.6-micron flash process, which is intended to produce 16-Mbit flash chips. While this was originally expected to be well behind the NMBS project, delays at NMBS coupled with better than expected results at Sharp mean that the Sharp line may be production-ready at the same time or sooner than the NMBS line. Intel is now planning to build 8-Mbit parts on the Sharp line as well.

In summary, the bad news is that there simply won't be any additional supply of 8-Mbit flash chips for at least six months, and possibly as long as a year; all of the existing chips have been allocated. The good news is that three separate fabs are expected to be on-line by the end of '93, creating an enormous capacity.

Intel's troubles could present an opportunity for other vendors, but none have chips to offer at this density. Toshiba has recently announced a 16-Mbit device, but production is not scheduled to begin until March 1993.

AMD Puts Intel on Notice for 486 Microcode

On October 27, AMD sent a letter to Intel giving 30 days notice that it was planning to ship a 486 microprocessor with Intel's microcode. The 30-day notice is required by the agreement between the two companies to give Intel time to specify how the parts should be marked. Intel disputes AMD's claim that the agreement applies to microcode in microprocessors, claiming that the term "microcode" in the agreement really refers to firmware. While a jury ruled that the license did not apply to microcode in the 287 math coprocessor, Judge Ingram ruled that the jury verdict applied only to the 287, and the judge is expected to issue his own ruling soon on the broader issue of whether the agreement covers microcode in microprocessors. AMD apparently wants to be able to ship its 486 microprocessor with Intel microcode immediately if the judge hands down a positive ruling.

Cyrix Adds Extended 486SLC

Cyrix has begun shipping its Cx486SLC/e, which adds a system-management mode to the original 486SLC. The SMM functions were described when the part was originally introduced *(see **060501.PDF**)* and were originally promised to be available in July, but shipments are just now beginning. The 486SLC/e will replace the original 486SLC in early 1993. Pricing is $75 for the 25-MHz version and $99 for the 33-MHz part; the 25-MHz part is also available in a 3.3V version for $89.

Cyrix Reveals Its First 486-Pinout Processor

Cyrix also announced its first 486-pin-compatible chip, the Cx486S2/50, code-named M6. This chip is based on the same CPU core as the 486SLC and 486DLC, but with a 2K write-back cache (instead of the 1K write-through cache) and an on-chip clock doubler. It plugs into a 25-MHz socket but operates internally at 50 MHz. While the core is not quite as fast as Intel's 486 and the cache is smaller, the write-back cache design is more effective than Intel's write-through approach. The Cyrix chip also has an eight-level write buffer, while Intel's has only four levels. In addition, Cyrix's chip implements burst writes as well as burst reads. Cyrix claims that its 486S2/50 is "significantly faster" than Intel's 486SX-33, but no benchmark results are yet available.

Cyrix plans to support two schemes for dealing with write-back cache coherency. One method, which can be enabled by software, causes all dirty data in the cache to be written to main memory whenever HOLD is asserted before HLDA is asserted in response. This doesn't require any system hardware changes, but it does add an unex-

pected latency in responding to hold requests. The preferred method is to modify the system logic to support new signals that use pins that are no-connects on the standard 486. These new signals implement an "abort and retry" protocol, in which the processor snoops bus reads by other masters. If a snoop hit occurs on dirty data, the processor asserts a signal telling the system logic to abort the read. The processor then writes the dirty data to memory, and the system logic retries the read. While this requires support in system-logic chip sets, Cyrix says that chip-set vendors are already implementing this logic because it is required by Intel's Pentium processor.

Samples of the 486S2/50 will be available by year-end, with production in 1Q93. Pricing is $249 in thousands, which is $60 more than Intel's 486SX. This introductory price is unrealistically high, and pricing is likely to drop to the $150 range as production ramps up. A separate FPU is available for only $35, for a combined price of $284—about $200 less than Intel's 486DX2-50. (Cyrix added signals to mimic the 386/387 coprocessor interface on the 486S2/50.) The floating-point performance of the Cyrix CPU/FPU pair will not match Intel's 486DX, however, because of the coprocessor communication overhead.

Another 486SX-pin-compatible product, code named M6n, is planned for the first quarter and will add power-management features similar to Cyrix's 486SLC/e. Cyrix also revealed plans for the M7, to be introduced next year, which will add an on-chip FPU and a larger cache. It will also operate at higher clock rates. This chip appears to be a head-to-head competitor for Intel's 486DX2-66.

Intel 960 Wins Key HP LaserJet Design

Just when it seemed that AMD's 29000 was capturing the lion's share of high-volume laser-printer design wins, Intel has scored the biggest one of all: HP's next-generation, mainstream laser printer. The new LaserJet 4 provides 600 dpi resolution for only $2199, using a new Canon engine. The LaserJet 4M, priced at $2999, is designed for Macintosh users and includes Postscript and a LocalTalk interface. Both printers are based on a 20-MHz Intel 960KA processor.

HP's LaserJet line has previously used 68000-family processors for all except the high-end model, which is based on an AMD 29000. HP's adoption of the 960 for its new mainstream line signals the end of one of the 68000 family's highest-volume applications, and it also indicates that HP will not use the 29000 throughout the family. HP also uses 960 processors in its PaintJet XL300 color printer, DesignJet plotter, and 700/RX× terminals.

Design wins continue to captured by the 29000 as well, however; Lexmark, formerly IBM's printer division, just announced the Color Jetprinter PS 4079, a $3495 color inkjet Postscript printer that uses AMD's 16-MHz 29200.

Intel Announces New Interrupt Controller

Finally replacing the ancient 8259 interrupt controller (originally designed for use with the 8080!), Intel has announced a new interrupt controller design—the APIC (Advanced Programmable Interrupt Controller) architecture—and its first implementation, the 82489DX. The APIC is a descendent of the MPIC (multiprocessor interrupt controller) that was originally intended to support the i860XP as well as the 486 but was never shipped. With the near-demise of the 860, however, the 860-specific features, including support for the PAX loop-level parallelism technique, have been dropped, and other features have evolved.

The APIC is the first off-the-shelf interrupt controller designed for multiprocessor systems. Previously, developers of multiprocessor x86 systems had to develop their own ASICs for interrupt control. The APIC is divided into two blocks: one that processes interrupt inputs, and another that takes processed interrupt requests and passes them to the processor. In a uniprocessor system, the two blocks are connected together directly within the chip. In a multiprocessor system, a dedicated interrupt control bus connects multiple APICs, allowing interrupts received by any APIC to be routed to any processor. Interrupts can be assigned to a fixed processor, or they can be dynamically routed to the processor that is currently running the lowest-priority task.

While it is most compelling in multiprocessor systems, the APIC will benefit uniprocessor systems as well. Because of the more flexible interrupt priority scheme and much faster register access, the overhead associated with processing interrupts is drastically reduced. Intel estimates that system throughput could increase by as much as 6% over a system using the traditional pair of 8259A interrupt controllers. Unfortunately, existing operating systems and even some DOS applications assume that an 8259A interrupt controller is present. Windows NT and various UNIX implementations will provide support for the APIC, however, and many high-end PC makers are likely to include the APIC along with an 8259A (actually, typically the equivalent logic integrated into a chip set) for compatibility with older software. Eventually, the APIC architecture will be implemented in system-logic chip sets, and it may be integrated on future processors as well.

The 82498DX is in production now and is priced at $26 in thousands. For more information, call 800/548-4725 and ask for literature packet #E9-P01.

Sun Extends Product Line at Top and Bottom

Sun announced two new systems that greatly extend the breadth of its SPARC offerings. At the high-end, the SPARCcenter 2000 will eventually allow up to 20 SuperSPARC processors to work together in a single system. At the low end, the SPARCclassic is a fully-config-

ured color workstation for $4,495 based on TI's microSPARC chip.

The highly-integrated microSPARC, combined with two peripheral chips from NCR *(see **061402.PDF**)*, greatly reduces the number of components needed to build a SPARC system, allowing the Classic to reach its low price. Sun rates the system at 26.4 SPECint92 and 21.0 SPECfp92, about 15% higher than the figures given by TI for its chip; Sun attributes the increase to compiler tuning. This performance falls well short of a high-end 486DX2-66 on integer code but beats it on floating-point applications.

The base price includes 16M of memory, a 207M hard disk, a 15" color monitor with 1024 × 768 × 8 graphics, two SBus expansion slots, and fast SCSI and Ethernet interfaces. Sun believes that the Classic is priced lower than comparably equipped 486DX systems from Dell, Compaq, and IBM, although Sun cautions that $4495 is the "street" price; few discounts will be available.

Filling out the high end, the SPARC 2000 uses new bus technology to considerably increase the number of processors that can be combined in a single shared-memory system. The backbone of the system is a pair of buses using the XDBus (previously called DynaBus) protocol. These two 64-bit buses run at 40 MHz, the same speed as the SuperSPARC processors. Unlike MBus, XDBus uses a packet-switched (split-transaction) protocol to reduce overhead cycles, achieving a combined sustainable bandwidth of over 500 Mbytes per second. XDBus uses Gunning Transceiver Logic (GTL), a CMOS technology with a 0.8V swing, to improve transmission characteristics across the relatively long backplane. XDBus was developed at Xerox PARC and can be licensed from Xerox, although the actual bus interface chips in the SPARC 2000 are jointly owned by Sun and Xerox.

The SPARC 2000 puts Sun in direct competition with HP, Pyramid, and Sequent to put large UNIX systems in the data center, the so-called "glass house." The highly-parallel system organization is well-suited to large OLTP (on-line transaction processing) applications. At just $95,000 for an entry-level two-processor system, Sun is aggressively pricing this box to quickly gain market share. No performance figures are available, but Sun hopes the system will offer equivalent performance to the most powerful HP and Sequent systems at a much lower price. Along with price and performance, Sun will also have to convince potential customers that it can offer the reliability and service required for such critical applications. The first SPARC 2000 systems will support up to 8 processors and will ship next April; support for the full 20 processors is expected by mid-1994.

Die Size Errata

Due to an editorial slip, on page 1 of our 10/28/92 issue we may have given the mistaken impression that the PowerPC 601 is a gargantuan chip: the die size is 10.95 *mm* square, not 10.95 cm square.

On page 4 of our 10/7/92 issue, the die size of the redesigned 486SX is shown as 1 × 1.5 cm, which is the size given by Intel at the time. Apparently, someone at Intel rounded up to the nearest 0.5 cm; the actual die size is 6.9 × 10.4 mm (270 × 410 mils). ♦

