Memories Yield to Processors at ISSCC

IBM Reaches 1 GHz, Puts Multithreading on a Chip; Intel Tips Slot-2 Hand

by Brian Case

In early February, the forty-fifth International Solid State Circuits Conference (ISSCC) drew an estimated 3,000 attendees to San Francisco for its annual sampling of advances in integrated circuit technology. Although mainly aimed at circuit designers, ISSCC has evolved into a conference of more than passing interest to microprocessor designers. For example, attendees heard about the coming Slot 2 Pentium II processors, a 1-GHz PowerPC subset chip, and possible techniques for future generations of superscalar chips. For variety, there was also an evening discussion with a Buddhist monk as a panelist.

1-GHz PowerPC Wows Many

One of the star attractions of the show was the IBM paper "A 1.0-GHz Single-Issue 64-bit PowerPC Integer Processor." The chip is an experimental integer-only PowerPC processor implemented in IBM's 0.25-micron six-layer-metal CMOS process without copper interconnect. This is the same process (CMOS-6X) used to fab the current 604e PowerPC chips, which run at a top speed of 350 MHz. The test chip was built to help explore circuit, clocking, microarchitecture, and design methods for high-frequency processors.

There are several caveats to keep in mind when assessing the relevance of this chip's clock rate. On the downside, the chip does not implement the full PowerPC integer instruction set, leaving out multiply- and divide-related ops. The core is a single-issue design. The chip's instruction and data caches are small, at 4K each, and they are direct-mapped. There is no address translation and no external system interface. Load and store instructions execute in 1.15 ns, reducing the realistic maximum operating frequency to 870 MHz. And, as Stanford's Mark Horowitz points out, ISSCC is a "room temperature/typical" conference; this chip was evaluated at 25° C with a nominal 1.8-V supply. The designers estimate that characterization over full temperature and voltage ranges would reduce performance by another 20%.

On the upside, design compromises leave much room for improvement. Full-custom-logic building blocks were used, such as fast 8-to-1 multiplexers and hand-tuned 64-word × 64-bit ROMs, but the design is not as optimized as possible; the small design team estimates a 10% improvement could be had with a more careful design using a larger team. Copper interconnect, recently announced by IBM as production ready (see MPR 8/4/97, p. 14), was not used, but the designers expect the next-generation copper process would improve performance by 20%. The chip was tested only with a wafer-probe station; a normally packaged chip would have much better power delivery (C4 packaging vs. the pad-ring delivery of the probe station), which could also improve performance.

In the final analysis, it appears that by designing and packaging this chip as a product and using the most aggressive technology, a clock rate of 1 GHz is well within reach. Further, deepening the chip's four-stage pipeline, shown in Figure 1, and pipelining the caches should allow a clock frequency beyond 1 GHz. It is worth noting that the chip dissipates a modest 6.3 W at 1.8 V and 1 GHz.

The designers tapped some clever techniques to attain the high clock rate. When accessing the data cache, the chip does not scrimp on PowerPC addressing modes: the full register+register address computation is implemented. Thus, in one cycle, the chip performs an addition, a cache access, and the multiplexing needed for byte reversal and unaligned accesses. To fit all this in a single cycle, the brilliant trick of using carry-save instead of carry-propagate adders is used. Carry-save adders eliminate the slow carry propagation chain, and they can be merged with the SRAM decoders. Many people have recently discovered this trick—necessity is the mother of invention—but Sun filed a patent application a couple of years ago. Sun presented a paper describing the technique: "64K Sum-Addressed Memory Cache."

Another clever technique employs the wires needed for carry propagation in the integer adder to implement the shifter and rotator. This design, shown in part in Figure 2, has several advantages, including sharing transistors and wires, reducing fan-out, and eliminating a result multiplexer.

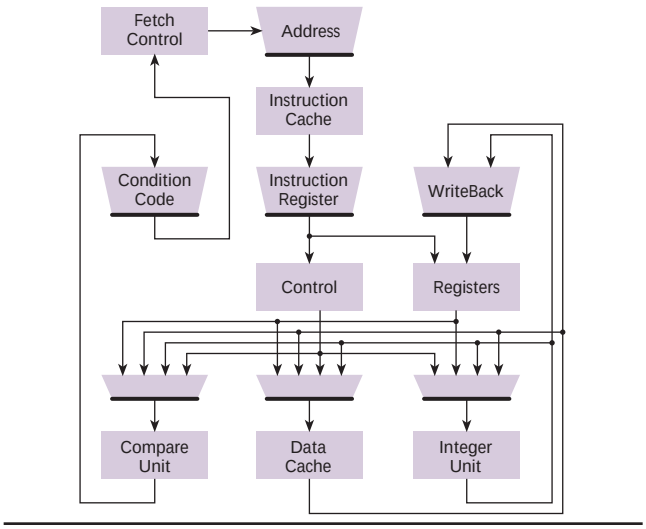

Figure 1. Simple four-stage pipeline of IBM's experimental 1-GHz PowerPC chip. (Thick lines represent pipeline registers.)

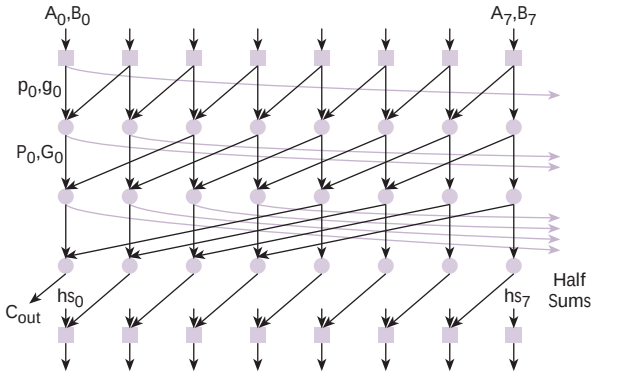

Figure 2. A slice of a single-network implementation of addition and rotation in IBM's experimental 1-GHz PowerPC chip. The black diagonal wires are used both for carry propagation and up shifting/rotation.

The right-to-left wires are needed for both propagating and generating signals for carry lookahead and for shifting/rotating. The left-to-right wires are needed only for rotating high-order bits to lower-order positions. The structure's last stage combines half-sums with carry-lookahead signals.

Copper Mainly Helps Clock Distribution

IBM presented another paper, "A 480-MHz RISC Microprocessor…With Copper Interconnects," that described a PowerPC 750 microprocessor in a 0.2-micron six-layer-metal CMOS process (CMOS-7S). The chip has separate 32K on-chip caches and an integrated controller for an off-chip L2 cache of 256K, 512K, or 1M. The die size is a minuscule 40 mm^2, down from 67 mm^2.

The main benefit of copper is in reducing the RC delay of long lines. Short lines do not improve much because their RC delay is already low. IBM compared the new implementation to the existing 0.25-micron chip and found that of the total 77% improvement in clock speed, only 12% is directly attributable to the use of copper. Perhaps the biggest impact was seen in the delay and skew of the clock-distribution tree. Copper reduced the clock latency from 170 ps to 85 ps and improved skew from 90 ps to 35 ps. Thus, while the benefit for this chip seems modest, the value of copper interconnect seems to be more pronounced at higher clock speeds, since the 1-GHz design team estimates a 20% benefit from copper. In relative terms, this 40-mm^2 PowerPC chip is tiny; the benefit may also be magnified for larger chips.

This PowerPC implementation allows either 3.3-V or 1.8-V external cache chips. With the 1.8-V option, the external L2 can operate at the full CPU clock speed.

IBM Tries Multithreading on a Chip

The third IBM presentation in the Microprocessors session described "A Commercial Multithreaded RISC Processor," which is a PowerPC chip for the AS/400 line of minicomputers. The processor implements duplicated register files (both integer and FP) and duplicated address-translation hardware (TLB and segment TLB) to support two threads in hardware.

The thread-switch penalty is three cycles, which are needed to drain the four-stage pipeline in preparation for the thread switch. The processor can switch on L1 cache misses, L2 cache misses, and misses in the TLBs. Since L1 misses have a penalty of five cycles, switching to a new thread reduces the penalty by 40%, to three cycles. Simulation predicts a performance increase of up to 35%, but less than 30% is typical under normal circumstances. These simulations used a commercial, transaction-processing workload, which can experience very high cache-miss ratios; a performance improvement for PC or workstation applications should also be possible for some types of code.

Mobile Deschutes Uses C4 Package

Intel gave two presentations relating to the forthcoming Slot 2 Pentium II (Deschutes) processors. "A 450-MHz IA-32 P6 Family Microprocessor" described the processor chip while "A 450-MHz 512K Second-Level Cache With a 3.6-Gbyte/s Data Bandwidth" detailed the custom L2 cache chip.

Like many other high-performance microprocessors, Deschutes implements on-chip bypass capacitors to smooth noise in the power-supply grid. The fifth layer of metal allowed the designers to reduce the channel area between blocks of logic by more than 50%. Bypass capacitors are used to fill any unused space remaining between blocks. When space is available, bypass capacitors are also implemented within blocks. Finally, bypass capacitors for the output drivers are placed under the wire-bond pads. The result is a total on-chip bypass capacitance of about 117 nf, which is sufficient to eliminate the need for expensive package-mounted bypass capacitors.

The Deschutes die can be packaged with either C4 flip-chip technology or standard wire-bonding. The C4 package is small at 31 × 31 mm, while the plastic land grid array is 42.5 × 42.5 mm. The on-chip bypass capacitance makes C4 packaging more practical.

The small C4 package makes sense for use in laptop PCs. The power consumption is low enough: at 1.6 V and 266 MHz, the worst-case power dissipation is 7.15 W. Intel says power consumption drops to about 4 W when running typical office application software. This is an encouragingly low number, but it does not account for the L2 cache. Laptops will likely use commodity pipelined-burst SRAM (PBSRAM) running at one-half the processor clock rate, so their power dissipation should be just a watt or two.

Slot 2 Deschutes Stands Tall

In desktop PCs and servers, Deschutes will be running at its full speed. At 2.0 V and 450 MHz, power consumption is 18.9 W worst case. Assuming a relationship between worst-case and typical power similar to what Intel claims for 1.6-V operation, typical power dissipation running office applications at 450 MHz and 2.0 V should be around 11 W.

The 450-MHz custom L2 cache chip is fabricated in a 0.35-micron four-layer-metal process (P854) for a die size of 12.9 × 17.23 mm (222 mm^2). It will be built only by Intel, presumably in some of the 0.35-micron fabs that would otherwise be idle after Pentium and Pentium II production has shifted to 0.25-micron lines. Each cache chip has about 5 million transistors to implement a 512K, four-way set-associative cache. This is the second generation of the cache chip currently used in the Pentium Pro.

The chip operates at the processor core frequency and has a four-cycle latency. Thus, a full four-transfer burst of 36 bytes (32 plus parity) takes eight cycles total. A cache data request may be initiated every four cycles, yielding a sustained transfer rate of 3.6 Gbytes/s.

The cache chip's supply voltage is 2.5 V, but the I/O runs at 2.0 V for compatibility with Deschutes. Its maximum power dissipation is 4.5 W at 450 MHz, assuming back-to-back reads. Typical power should be significantly less.

Thus, the total power for a Deschutes Slot 2 processor with 1M of L2 cache should be about 28 W maximum and somewhere around 17 W typical. With 2M of cache—the maximum configuration—maximum power jumps to 37 W and perhaps 23 W typical. Even 37 W is better than the 42-W dissipation of the first-generation Pentium II at 300 MHz.

This cache chip may have a long life in Intel product line: the paper reports that the chip operates at 657 MHz under nominal conditions; the chip probably runs faster, but the tester cannot. Running with the chip's built-in self-test logic, 750-MHz operation was achieved. While the high speed of this cache chip makes it a natural complement to the next-generation Katmai processor, Intel could shrink it to 0.25-micron for Merced.

To support the new clock and bus speeds, Deschutes implements two new bus-to-core-clock ratios: 2:9 and 1:5. With a 100-MHz system bus, these ratios set a core speed of 450 MHz and 500 MHz, respectively. Given the speed of the cache chip and the fact that Deschutes can run at up to 2.2 V if needed, Intel may be able to introduce a 500-MHz Deschutes using its current process.

In both the processor-chip and cache-chip presentations, a picture of a bare Deschutes Slot 2 module PC board was shown. Figure 3 shows an approximation of this board design. Intel has acknowledged that Slot 2 modules will be wider than Slot 1 modules and twice as tall.

On a Slot 1 PC board (see MPR 5/12/97, p. 7), the cache chips are mounted next to the processor. These PBSRAM chips, however, are in fairly small packages, so there is room for them as well as the discretes next to the processor. The packaged custom cache chips, on the other hand, are in land-grid-array (LGA) packages as big as that of the processor, so there is insufficient space on a Slot 1 PC board.

Given the expanded dimensions, it may be impossible to fit a Slot 2 motherboard and processor into a standard PC case. Since Intel is targeting Slot 2 for workstations and servers, this may be intentional. In any case, the Slot 2 edge connector has two key notches vs. the single key notch of Slot 1, so there is no danger of inserting a module into an incompatible motherboard.

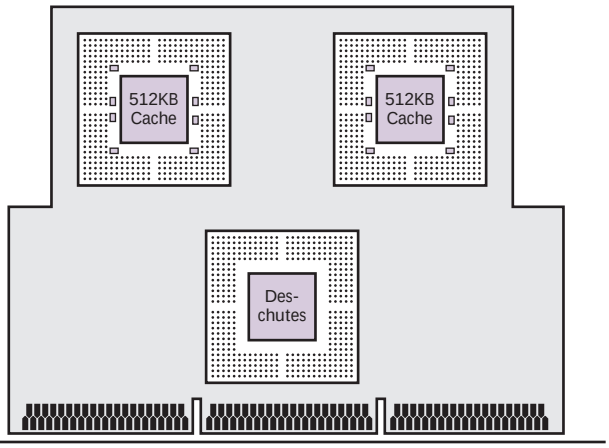

Figure 3. Possible Slot 2 module PC board form factor, as pictured in two ISSCC presentations. Note that the caches have package-mounted bypass capacitors.

AMD Reveals Some Details of 3D Instructions

At the Microprocessor Forum in October, AMD, Cyrix, and Centaur all revealed plans for philosophically similar but mutually incompatible x86 instruction-set extensions to speed geometry calculations for 3D applications. Later, the trio announced that they had agreed to support and implement AMD's new instructions (see MPR 12/8/97, p. 4). At ISSCC, a few details of these 3D extensions were disclosed.

Like the existing MMX subset, the new instructions use the FP registers, but they operate on two pairs of single-precision floating-point values. The format is IEEE-754 compatible, but, for most operations, only round-to-nearest mode is supported, and conversion results are always truncated. Overflows are always clamped to the maximum representable value, while underflows are always clamped to zero. The implementation recognizes neither NaNs (not a number) nor infinities, and instructions cannot generate exceptions and cannot set status flags. Simply put, the implementation ignores those parts of the IEEE standard that either are not relevant to graphics applications or stand in the way of high performance and simple implementation.

The instructions have a latency of two cycles and are fully pipelined for a throughput of one cycle per instruction. Thus, as we speculated, performance can reach a maximum of four operations per cycle. This throughput can be achieved by issuing two 3D instructions per cycle, but only one can use the 3D multiply pipeline; the other must use the 3D add/shift pipeline.

Also, AMD has quantified the value of the second MMX integer unit: clocks-per-instruction on the Intel Multimedia Benchmark has been improved by more than 20%. This could bring the K6 up to parity with Pentium/MMX but probably keep it behind Pentium II (see MPR 9/15/97, p. 18).

Three refinements have been added to the chip to improve performance on external accesses, which for the Socket 7–compatible K6 chips includes accesses to the L2 cache. The new chip's ability to pipeline data reads and writes on the bus, fill the data cache speculatively, and complete a store while a cache fill is in progress improve overall CPI in simulations by 2% to 10%.

StrongArm Gets Denser, Faster

Digital described two chips derived from its StrongArm core. The first, described in the Microprocessors session paper "A 200-MHz 32-bit 5W CMOS RISC Microprocessor," is a StrongArm core surrounded by the support circuitry needed for a PDA or HPC application. This chip is the SA-1100 for Windows CE devices (see MPR 9/15/97, p. 1).

The voltage has been reduced from 1.65 V in the previous StrongArm to 1.5 V. When idle, the chip dissipates less than 65 mW, and sleep current was simulated to be less than 50 µA but measured at a battery-saving 20 µA. Both data caches are 32-way set-associative and virtually addressed; previously, they were reported to be only two-way associative.

The second StrongArm-based chip, described in the Multimedia session, is the SA-1500, also described at last October's Microprocessor Forum (see MPR 12/8/97, p. 12). Power consumption at 2 V and 300 MHz has been measured at 1.57 W. At a lower voltage, the chip can reach 150 MHz with less than 0.5 W dissipation, and Digital was able to run the chip as high as 383 MHz in the lab.

21164PC Gets Much Faster

One of the highest clock-rate chips was described by Digital in the Multimedia session in the paper "A 667-MHz RISC Microprocessor Containing a 6.0-ns 64-bit Integer Multiplier." This chip is a new version of the 21164PC in a 0.28-micron four-layer-metal process; with 5.7 million transistors, the die measures 6.7 mm × 15 mm. The chip supports a 2.5-V interface but requires a 2.0-V supply for the core. At 667 MHz, the chip dissipates a respectable 22 W. The instruction and data caches are doubled to 32K and 16K, respectively, and the paper describes a set-prediction scheme to speed operation of the instruction cache and a refined branch-prediction mechanism.

If this 21164PC could run at a lower voltage, the resulting reduced power consumption might make this chip a realistic candidate for use in an Alpha-based laptop. For example, at 1.5 V and around 500 MHz, the power dissipation might be reduced to less than 10 W while maintaining respectable performance.

DRAM Developments Reflect Reality

In years past, huge DRAM chips have been one of the main attractions at ISSCC. Three years ago, two 1-Gbit DRAMs were described at the conference, with a 4-Gbit chip following last year. NEC's 1-Gbit device in 1995 had an almost unbelievable die size of 936 mm^2. Last year, NEC's 4-Gbit chip—which is actually a 2-billion-cell chip that stores two bits per cell—measured 986 mm^2, essentially a factor of two improvement in die size.

This year, only Fujitsu described a gigascale DRAM, and it was "only" a 1-Gbit chip. Presumably, NEC has tired of building these expensive concept-car chips. The die size of the Fujitsu chip is 505 mm^2 in a 0.18-micron process, so the area per bit and thus the manufacturability of gigabit DRAMs has not improved. For comparison, DRAM expert Steve Przybylski estimates that some modern 16-Mbit chips are well below 50 mm^2 in 0.35-micron technology, and 0.25-micron 64-Mbit devices are just crossing the magic 100-mm^2 volume-manufacturability threshold.

IBM and Hitachi joined the 256-Mbit generation with chips that approach manufacturability. The Hitachi chip measures 205 mm^2, with the IBM chip only slightly larger at 220 mm^2. Both of these SDRAMs operate at more than 150 MHz, and the Hitachi device can operate in double-data-rate mode to achieve 300 Mbits/s per pin.

It is interesting to note that if Fujitsu used one-quarter of its 1-Gbit DRAM to build a 256-Mbit part, the resulting device would measure substantially less than the IBM or Hitachi chips. Fujitsu's chip, however, is fabricated in a 0.18-micron process, denser than the cheaper 0.22-micron processes of the IBM and Hitachi chips.

Another notable development in memory design was reported by Hitachi in the paper "A 128-Mbit Early Prototype for Gigascale Single-Electron Memories." Currently, this design has little to offer over conventional memories because cell size is limited by word- and bit-line pitches. Also, latency is very slow, measured in microseconds, not nanoseconds, but the burst rate is very fast, 10 ns in this device. Chip measurements indicate that a 256-Mbit device could be built now with a die size of around 120 mm^2, half the size of the other reported 256-Mbit chips. Thus, in the future, this technology promises very dense solid-state storage.

Specialty Memories Flourish

Some notable memory developments were presented in the Nonvolatile and Embedded session. First, there seems to be a trend toward development of multibit-per-cell storage in flash memories. Intel, for one, is already shipping 2-bit-per-cell "Strataflash" parts (see MPR 10/6/97, p. 5).

SGS-Thomson detailed its design of a 6-bit-per-cell flash memory that yields a 257-Mbit/cm^2 effective bit density. Thus, a 256-Mbit flash chip would theoretically be only 100 mm^2 in a 0.5-micron process. In a more modern process, the die size would probably be small enough to allow volume production. The text of the paper describes storing JPEG images and compressed speech, so the chip seems designed for consumer digital cameras and answering machines. Writing to the memory is slow, at about 100 microseconds per byte.

A group from the Korea Advanced Institute of Science and Technology described an 8-bit-per-cell memory. This

design is based on EEPROM technology, and the test chip contained only 1,024 cells in a 0.8-micron process.

A second notable development was Mosaid Technologies' graphics accelerator with an integrated 13.4-Mbit frame buffer. The chip operates at 66 MHz and uses a mixed 0.35-micron/0.5-micron three-layer-metal process. While not the first single-chip graphics solution, this chip is notable for integrating parts of the graphics processor within the DRAM structure.

The result is that the graphics processor effectively has a 4,096-bit bus between the DRAM and the pixel logic for some operations. For example, screen clears can be performed with 4,096-bit writes, yielding a throughput of 33 Gbytes/s. Fortuitously aligned block copies can also be performed at half this rate, and arbitrary copies can be performed 256 bits at a time for a bandwidth of 2 Gbytes/s, still at least four times faster than a graphics chip plus separate 64-bit memory. In contrast, NeoMagic's successful integrated graphics chip (see MPR 6/23/97, p. 5) attains a 400-Mbyte/s on-chip bandwidth.

A third notable development is a compression/decompression DRAM intended to boost performance in unified memory architecture (UMA) systems. The idea is that by moving raw frame-buffer information over the memory bus in compressed form, more memory-bus bandwidth is available for the system processor, thus reducing the negative impact of the UMA organization on system performance. The claim is that the memory-bus bandwidth available to the system processor improves from 220 Mbytes/s to 595 Mbytes/s, using typical assumptions about frame buffer contents (mostly text characters; graphics and pictures would increase the compression ratio and therefore the benefit).

The downside to this CompressDRAM, as NEC calls it, is that a 16-Mbit chip is about 16 mm × 9 mm (144 mm^2) in a 0.35-micron process, at least three times the size of a commodity DRAM of similar capacity. Also, UMA is looking less attractive as PC prices fall.

Fractal Capacitor a Thing of Beauty

As microprocessors and system logic get faster and consume more instantaneous power, supply noise becomes a problem that needs a chip-level solution. Some commercial microprocessors now have integrated bypass capacitors, such as the aforementioned Deschutes.

One novel and beautiful way to increase capacitance per area of two separated metal plates is by exploiting lateral fringing fields. The effect of fringing fields is proportional to periphery, and an ideal fractal has infinite periphery within a finite area. Of course, a real fractal drawn using real metal lines has finite periphery, but this technique can increase capacitance per area by up to a factor of six.

Technology Directions Beyond Superscalar

David Luick from IBM presented his view of future processor architecture in a presentation "Beyond Superscalar RISC, What's Next? An Almost Unbiased View." He advocated some of the ideas described earlier here (see MPR 1/26/98, p. 21), including clustering of execution units and distribution of resources. Luick further predicted that future processors will use ever simpler instruction sets, much greater reliance on compilers to allow hardware to be simplified, a combination of hardware and software to learn about programs at run time, fewer logic gates per pipeline stage, and many levels of on-chip cache. He also advocated multithreading to reduce apparent pipeline depth.

In one of the evening discussion sessions, titled "Will Power Limit Microprocessor Performance?", architects, microarchitects, circuit designers, and process designers faced off for some finger pointing. Fundamentally, all parties agreed that power either is or will be a limiting factor in performance today, but that it need not be. Most agreed that reducing power consumption should be a design issue from the earliest stages of a project. An interesting point was made that not all signal paths on microprocessors are time-critical, but since most are designed to be as fast as possible and speed is bought with power, significant power is wasted.

Too Much Information, Too Little Time

In addition to microprocessor- and memory-related talks, the program also featured papers on MPEG-2 and MPEG-4 codecs, single-chip radio-frequency receivers, image sensors (including a 1,400-mm^2 single-chip dental X-ray imager), high-speed DRAM and chip-to-chip interfaces including talks from Rambus, and analog technology such as disk-drive electronics. Conspicuously absent from this year's program were talks on gate-array developments, a prominent topic in years past. Unequivocally, though, ISSCC presents an impressively broad sampling of the present and future of integrated electronics. M

ISSCC Modernizes Products

Tradition rules at ISSCC. Virtually unchanged is the weighty *Digest of Technical Papers*, with its familiar format, color scheme, and heft. This year the *Digest* is a mere 504 pages (up from 336 in 1993), each one a thick piece of coated, high-gloss paper, making it extremely heavy to carry and capable of causing injury if carelessly dropped. An equally hefty, but still useful, 480-page slide supplement volume was shipped in late February.

Much easier to wield is the conference CD-ROM. This valuable resource contains the entire *Digest* and the slides from the paper presentations. The 1997 conference CD-ROM—the first one available—also contains the *Digest of Technical Papers* from the 1996 VLSI Circuits Symposium. To help the weary, wired engineer keep up to date between conferences, the ISSCC now has a Web site at *www.isscc.org*.