Microprocessors Displace DRAMs at Forefront of ISSCC

By Brian Case

Each year in late winter, circuit designers from around the world convene at the International Solid State Circuits Conference (ISSCC) to check the pulse of the IC industry. This year's conference, held in San Francisco on February 19–21, saw a lull in DRAM developments but a surge in important microprocessors and areas related to consumer electronics. While last year's conference had five papers on 64-Mb DRAMs, this year's had only one. Perhaps the DRAM makers' attention has shifted toward getting last year's designs into production.

Active areas at this conference were high-speed A/D converters, neural network chips, HDTV image sensors, and dense SRAMs. Two interesting medical application chips were described: a long, skinny (7.2 × 0.7 mm) CMOS blood pressure/temperature sensor designed to be inserted directly into the blood stream through a catheter, and a 16-channel neural stimulator designed to be jabbed directly into neural tissue. The neural stimulator has four long prongs with four stimulation sites each and is part of research toward devices suitable for use as a neural prosthesis for deafness and other disorders.

The opening plenary session included talks on three interesting topics: micromachines, automotive electronics, and personal communications. Some fascinating pictures of micromotors and gear trains appear in the proceedings. The various moving parts were fabricated with an E-beam process.

A BMW engineer presented a paper on the evolution of automotive electronics. Currently, about 10% of the cost of high-end autos is electronics, and another 10% is wiring. Of the electronics component, 40% is semiconductors and 60% is passives (boards, etc.). Up to 40 microprocessors with 500 Kbytes of total program ROM are in a single car. By the year 2000, the value of electronics in the average North American car is expected to rise to about $2000. In the future, bus-based wiring is expected to replace the conventional point-to-point wiring harnesses to reduce costs.

Microprocessors

Digital Equipment Corp. and Sun Microsystems made the biggest splashes in the microprocessor session with two papers each. In addition to setting clock rate and transistor count records, these companies set records for the most authors. DEC's VAX implementation paper had 30 authors while SUN's SuperSPARC paper had 34. DEC had the lowest transistor per author ratio of 43K/author while SuperSPARC had a respectable 91K/author.

DEC's latest VAX implementation, called NVAX, comes in the form of a 100-MHz, 1.3-million-transistor single chip that integrates a 2-KB direct-mapped I-cache, a 2-way set-associative 8-KB D-cache, an integer unit, a floating-point unit, a 96-entry fully-associative TLB, read and write buffers, and a second-level cache controller with multiprocessor ownership cache coherency support. DEC claims a 2.4-fold improvement in clocks-per-instruction over previous VAX implementations through added pipeline parallelism, larger caches, and compiler improvements. After patching a few microcode bugs with the on-chip control-store patching CAM, the 100-MHz chips have been benchmarked at 50 SPECmarks—an impressive achievement for the super-CISC VAX architecture.

DEC's real bombshell, however, came in the form of the 200-MHz Alpha processor implementation (see µPR 3/4/92, p. 1). This 1.68-million-transistor chip integrates an 8-KB I-cache and an 8-KB D-cache, both direct-mapped, a 64-bit integer unit, a floating-point unit, two TLBs, a write-buffer, and bus-interface unit with some second-level cache control. The chip operates at 3.3 volts but still dissipates a coffee-warming 30 watts at 200 MHz. This implementation can issue up to two instructions in a cycle, yielding a peak rate of 400 MIPS. Alpha joins the R4000 as one of the first implementations of a 64-bit architecture.

At somewhat the opposite end of the implementation spectrum from Alpha is the SuperSPARC from Sun and TI. Not much new information was presented at ISSCC, but the clock rate—40 MHz for the current prototype systems—and power consumption—8 watts—were revealed. The presenter said that chips were running "some tests" at 50 MHz, and this is likely to be the target for initial production. SuperSPARC integrates an astounding 3.1 million transistors. Usually, a high transistor count at ISSCC indicates that a chip is experimental, but not in this case. As with other high-end single-chip microprocessors, it integrates the entire processor subsystem: 20-KB I-cache, 16-KB D-cache, integer unit, floating-point unit, write buffer, and TLB. While the Alpha implementation is superscalar, it cannot issue two integer instructions in the same clock cycle. SuperSPARC can issue up to three instructions at once, two of which can be integer ALU instructions, and

the two integer instructions can even be dependent (see µPR 12/4/91, p. 1). The combination of larger, set-associative caches and its more general superscalar capability may bring SuperSPARC close to Alpha in sustained performance, at least on integer programs.

Sun also disclosed details about SuperSPARC's companion external cache controller. At 2.2 million transistors, this is another monster chip. Interestingly, this chip had a claimed operating frequency of 50 MHz, so it is likely that Sun and TI are expecting to raise the SuperSPARC CPU to at least that frequency. Since both of these chips are being fabricated in a new, very aggressive technology, there may be some process or circuit design problems limiting the clock frequency of SuperSPARC.

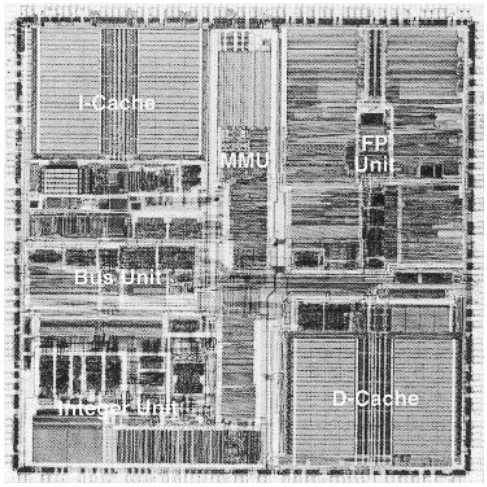

Die photo of the 3.1-million-transistor, 16×16 mm SuperSPARC.

The SuperSPARC cache controller supports up to 2 Mbytes of direct-mapped, second-level cache and interfaces to either the MBus or the XBus. MBus is the standard SPARC multiprocessor interconnection bus for small- to medium-scale multiprocessors and will be used in desktop workstations. XBus, which is derived from the DynaBus developed at Xerox PARC, is a more sophisticated multiprocessor bus and will be used in large-scale servers. XBus can operate at either TTL levels or the new, lower-voltage-swing Gunning-transceiver logic levels for faster operation. In addition, a multiprocessor system can incorporate multiple XBus channels; this cache controller can connect to up as many as four XBus interface chips (which means one SuperSPARC could reside on four shared-memory buses at the same time). Sun will be able to create sophisticated, high-performance—and expensive—multiprocessor servers.

Fujitsu talked about a vector coprocessor (it is not intended to stand alone) on a chip. The architecture of this chip builds on Fujitsu's experience in supercomputer architecture. The chip has four vector registers, each with room for 256 double-precision operands, a mask register to control which vector elements participate in operations, a cross-bar switch to route vectors from registers to execution units, and four pipelines: add, multiply, divide, and load/store. The chip operates at 70 MHz and has a peak performance of 149 MFLOPS (double precision) and 249 MFLOPS (single precision). Unfortunately, the external interface is only 64 bits wide, which is not wide enough to keep the execution units busy. In a realistic situation, computing a vector result requires reading each of two vectors separately, computing the result, and writing the result vector, leaving the execution units idle much of the time, waiting for the single bus. More or wider external buses would enable better utilization of the execution units. Fujitsu claims to have at least a couple of design wins and expects this chip to play a role in the first teraFLOP massively parallel machine. This will be their first production 0.5 micron chip but will only be available at 50 MHz in 1992.

Hitachi disclosed a small (8.1 × 8.0 mm) but exciting 1-million-transistor microprocessor with a peak performance level of 1000 MIPS. This chip contains two dual-issue superscalar integer processors with five-stage pipelines, four 1-KB, two-way set-associative primary caches, four 64-entry, two-way set-associative TLBs, and one 8-KB, two-way set-associative secondary cache. The chip is small and fast (250 MHz) because it is fabricated in an astounding 0.3-micron BiCMOS process: metal-3 has a very tight 0.55 micron width and 0.75 micron spacing. The ALU performs an addition in only 1.2 ns.

SRAMs

Until hundreds of kilobytes of cache can be integrated along with a processor, discrete SRAMs for external cache will be an important part of high-performance computer systems. As density continues to increase, SRAMs will be increasingly important in battery-powered computer applications because of their low standby current and simple interface. The SRAM papers at this ISSCC indicate that the technology is keeping pace with the demands of high speed and low power.

Hitachi described two SRAM technologies. The first SRAM was only a 4-Kb test chip, but its technology is impressive. Fabricated in a 0.4-micron, 4-poly, 2-metal process, this chip operates with a power supply of only 1 volt. The designers claim that a 4-Mb part using this technology would require only 0.73 µA standby current at 1 volt. Its 250-ns access is slow at 1 volt, but the access time drops quickly to around 20 ns at 3 volts.

The second SRAM from Hitachi sports a 7-ns access time in a 1-Mb density. Using a 0.3-micron process, the die size is relatively small at 3.96 × 7.4 mm. At 3 volts and 100 MHz, the part consumes only 47 mA.

Toshiba and NEC both described 4-Mb SRAMs fabricated in 0.5-micron BiCMOS. Toshiba's part operates at 3.3 volts with a TTL interface and delivers a 9-ns typical access time. NEC's chip offers the flexibility of either an ECL or a TTL interface. With ECL I/O, the

access time is 6 ns; with TTL I/O, access is slightly slower at 8 ns.

The first 16-Mb SRAM designs were revealed by both Fujitsu and NEC. Both chips use a 0.4-micron process with four polysilicon and two metal connection layers. Continuing the trend toward lower voltages, both operate at 3.3 volts. The Fujitsu chip claims a 15-ns access time while NEC boasts a faster 12-ns access. Power consumption is 55 mA for the Fujitsu part and 90 mA for the NEC chip, both at 30 MHz. Both chips are huge at 225 mm^2.

Miscellany

As proof that the Japanese companies are gearing up for production of HDTV consumer equipment, three 2-million pixel (1920 × 1036 pixels) CCD image sensors were described. Toshiba's sensor features a 110 dB dynamic range and circuitry to completely eliminate image lag due to pixel capacitance. Chip size is 16.2 × 10.5 mm with a 1" diagonal image area of 14.0 × 7.8 mm. Sony's sensor is smaller with a 2/3" diagonal and has a narrower dynamic range at 70 dB. It may be appropriate for consumer video cameras. Like Toshiba's, NEC's sensor has a 1" diagonal, but its dynamic range is closer to the Sony part at 75 dB.

One of the most interesting chips described at this ISSCC was Hitachi's "System-Integrated ULSI Chip." Eleven 4 Mb DRAMs, six 64 kb SRAMs, and an 18 k gate array are integrated together on a single 50.4 × 38.2 mm "die." Only four of these chips fit on what appears to be a 5-inch wafer. The chip incorporates an interconnection scheme that allows bad DRAM segments to be disconnected. On this single chip, a system with 4 MB of RAM, some cache, and a small processor could be implemented. The DRAMs use 0.8-micron design rules, while 1.2-micron rules are used for the SRAM and gate array.

A group from Stanford talked about their work in wave pipelining. This technique calls for equalizing the delays of the signal paths between pipeline registers in a pipelined circuit. Because signals take approximately the same amount of time to propagate from one register to the next, more than one "wave" of data can be propagating between registers simultaneously. This allows the system to be clocked at a higher frequency than is allowed by the basic propagation delay of the logic between registers.

The Stanford circuit implements a 63-bit population counter function in BiCMOS current-mode logic that would operate at about 100 MHz if clocked conventionally. To test wave pipelining, the circuits were modified to equalize their lengths. The result is a chip that can operate correctly at about 250 MHz. The advantages of wave pipelining are shorter latency (less pipeline-register overhead), smaller area, and relaxed clock distribution. The disadvantages are that the pipeline stages cannot be "frozen" in a simple way since data will simply be lost if clocks are stopped and dynamic, precharged logic cannot be used as easily.

A group from UC Berkeley described an innovative multiprocessor chip designed to ease the task of real-time data path prototyping. The chip integrates eight simple 16-bit integer processors (EXUs). Two EXUs can be coupled to form a 32-bit data path. Each EXU has a six-entry register file, an eight-instruction "nanostore," a shifter, and an ALU. Connecting the EXUs is a crossbar switch. Real-time applications, such as DSPs, tend to have specialized data paths connected in specialized ways to meet the throughput requirements of the problem. This chip allows groups of processors to be connected with high bandwidth in arbitrary ways, allowing the simulation of a variety of real-time architectures. The chip would be much more useful if the processors had larger control stores and at least some of the processors had multipliers, but this is an innovative use of multiple processors on a single chip and points the way for further development. Together with the Hitachi chip, it appears that single-chip multiprocessors are beginning to get some attention.

Toshiba and NEC each described Flash EEPROMS using 0.6-micron process technology. Toshiba's 5-V-only technology is described as experimental but it could be used for both 4-Mb and 16-Mb densities. Access time is expected to be 58 ns typical.

The NEC part seems closer to commercial viability, although chip size is large at 6.3 × 18.5 mm. This chip can be organized as either 1 M × 16-bits or 2 M × 8-bits. The erase granularity is 512 words and takes 10 ms; programming time is 10 µs per word. This results in a sector erase/program time of about 15 ms. Access time is 58 ns random or 34 ns in page-mode. Such chips could be used to build fast solid-state disks and RAM cards.

Conclusions

This year's microprocessor session raised the level of achievement so high that future developments will have to be Herculean efforts to seem worthy. The 200+ MHz operating frequencies were previously only pursued in ECL or GaAs. SuperSPARC and its cache controller make any chip with less than 2 million transistors seem unaggressive. With such high frequencies and densities, the future holds excitement for microprocessor users.

The move to 0.5-micron technology seems to be well underway in Japan although it will be a while before commercial products are commonly fabricated in these processes. Along with the move to denser processes is the migration to lower-voltage operation: 3.3 volts now seems to be a necessity instead of a novelty.♦

Copies of the ISSCC Digest can be ordered for $125 from John Wuorinen; call him at 207/326-8811.