Chips Seek High Bandwidth at ISSCC 99

IBM Mainframe CPU Shows PC Processors a Trick or Two

by Brian Case

The theme of the forty-sixth International Solid-State Circuits Conference, held February 15–17 in San Francisco, was high-bandwidth systems. With ever-shrinking process technologies come faster processors, denser memories, and complete systems on a single chip.

A much-discussed use of this emerging technology is to lower the cost of broadband communication by using highly integrated DSP-based chips. One of the plenary sessions discussed bringing high-bandwidth connectivity to the home and office, and a session on xDSL signal processors reported on cheap chips for multimegabit-per-second modems. But most users have yet to benefit from these technologies: deployment is spotty and expensive. We hope these chips influence products as quickly as the processors presented at ISSCC 99.

The conference set a record, with 173 technical papers (24 more than last year) in 25 sessions. Packing all this information into three days of presentations, however, required five parallel tracks, forcing most attendees to miss many talks of interest. The good news is that there was almost always an interesting presentation somewhere. The conference also offered eight evening panels, six tutorial lectures, and a short course titled "Fast Local Area Networks."

One-Gbit DRAM Waxes, SLDRAM Wanes

The conference lacked the appealing firsts of the past few years, such as the first 1G DRAM, the first 4G DRAM (see MPR 3/10/97, p. 9), and the first gigahertz microprocessor (see MPR 3/9/98, p. 9). But in reality, these firsts were only experiments and had little to do with real products.

For gigascale DRAMs, however, several companies reported progress toward manufacturability. Samsung revealed its progress in the development of a 1G DDR (double-data-rate) SDRAM. The chip is built in a 0.14-micron triple-well three-layer-metal CMOS process and operates from a 2.5-V supply. The die measures 349 mm^2 (29 × 12 mm); while still large, this is a significant improvement over die sizes of 1G DRAMs presented in past years, as Table 1 shows. The Samsung chip is packaged in a 60-pin flip-BGA and offers a ×16 data interface. Internally, the DRAM has eight banks and achieves a transfer rate of 333 Mbits/s per pin, or 666 MBytes/s for the chip.

ISSCC 1G DRAMs		CMOS Process	Die Size
1995	Hitachi	0.16µ	715 mm^2
	NEC	0.25µ	936 mm^2
1996	Mitsubishi	0.15µ	582 mm^2
	Samsung	0.16µ	652 mm^2
1997	OKI	0.16µ	543 mm^2
1998	Fujitsu	0.18µ	505 mm^2
1999	IBM	0.175µ	390 mm^2
	NEC	0.18µ	547 mm^2
	Samsung	0.14µ	349 mm^2

Table 1. One-gigabit DRAM chips show a steady, if not monotonic, decrease in die size.

At 547 mm^2 (17.8 × 30.8 mm), NEC's latest 1G chip represents a huge improvement over the company's first 1G effort, but it is still significantly larger than either Samsung's or IBM's device. NEC's process appears less aggressive than Samsung's but comparable to IBM's. On the other hand, using a single metric to compare processes is often misleading.

NEC's chip uses a 60-pin CSP (chip-scale package) with 4 × 15 bumps at a pitch of 1.27 mm. The die has bond-out options for ×4, ×8, and ×16 data-bus organizations. At 250 Mbits/s per pin, this chip offers a bandwidth of only 500 MBytes/s, but in a comparable process, NEC's part would probably at least equal the performance of Samsung's chip. Supply voltage for NEC's device is also 2.5 V.

IBM's 1G DRAM is slightly larger than Samsung's but offers better performance. In its ×16 organization, the part yields 800 MBytes/s, or 1.6 GBytes/s with its ×32 version. The number of internal banks depends on the data-bus organization: there are 16 banks in ×16 mode but only half as many in ×32 mode. Like the other 1G DRAMs, this 390-mm^2 chip (14.3 × 27.3 mm) uses a 2.5-V supply. IBM chose an 88-pin TSOP-II package measuring 16.0 × 29.5 mm.

In other DRAM developments, Hyundai, Mosaid, Siemens, and Vanguard described a 72-Mbit SLDRAM (formerly known as SyncLink DRAM) as a proof of concept. The chip achieves 800 MBytes/s with an 18-bit (16 plus parity) data interface. Unfortunately, this is only half the bandwidth offered by Rambus's Direct RDRAM, and rather than fight a three-way battle, the SLDRAM consortium announced that it has thrown in the towel and will put its weight behind the DDR II (Double-Data-Rate II) specification. This leaves DDR II as the only serious alternative to Direct RDRAM; these two technologies will compete for space on tomorrow's motherboard, and though it has not emerged as fast as promised, DRDRAM still looks as if it will dominate.

Graphics Chips Clip Desktop CPUs

Although desktop and server microprocessors continue to offer satisfying performance increases and some new features, their fundamental instruction sets and microarchitectures are changing more slowly than in the past (Intel's SSE and IA-64/Merced notwithstanding). Accordingly, ISSCC offered some interesting papers on x86 chips, Hewlett-Packard's PA-8500, and mainframe CPUs from IBM.

Perhaps stealing the show was the PlayStation 2000 chip from Sony and Toshiba (see MPR 4/19/99, p. 1), which was described in the Multimedia Processors session. This single device integrates 10 floating-point multiply-accumulate units, four floating-point dividers, an MPEG-2 decoder, a 10-channel DMA unit, and a two-issue MIPS CPU. The chip was touted to achieve up to 5.2 GFLOPS at 250 MHz (the product is expected to ship at 6.2 GFLOPS at 300 MHz).

Not to be outdone, Fujitsu described a VLIW geometry processor with an on-chip PCI/AGP bus bridge; the bus bridge allows the chip to connect directly to the host PC's AGP bus and offer a standard PCI connection to a rendering chip. The Fujitsu chip computes 6.5 million polygons per second. The key to this performance is a VLIW architecture that packs four operations per instruction, and, like all modern graphics processors, includes many SIMD operations in the instruction set. Using SIMD, the chip can execute up to eight simultaneous FP calculations; at the preliminary clock rate of 312 MHz, the chip will achieve 2.5 GFLOPS.

Although this device offers fewer features and lower peak performance than the PlayStation chip, it is also much more reasonable in size. At 84 mm^2 in a 0.21-micron process, it is just one-third the size of the Sony/Toshiba chip.

The industry's first single-chip MPEG-2 audio/video/system encoder (see also MPR 10/26/98, p. 5) was described by iCompression (*www.icompression.com*). As reported at the Microprocessor Forum last fall, the device accomplishes this with two stack-based DSPs, which were designed in house (they are not a Sun-derived design) that execute Java bytecodes. The 172-mm^2 chip is implemented in a 0.35-micron process; in a more modern process, this chip could be small and inexpensive enough to make MPEG-2 encoding a standard PC feature.

General-Purpose Processors Fail To Excite

In the general-purpose microprocessors session, AMD described the K7 and the K7 floating-point unit, Intel revealed some circuit-design aspects of Pentium III, and Motorola presented its 450-MHz PowerPC with copper interconnect. Stepping away from clock-speed one-upsmanship, IBM gave two papers on the technology used in its Series 390 mainframes.

HP revealed few new details about the PA-8500, and though it is old news, the 1.5M of on-chip cache on this huge (469 mm^2) die is still impressive at a time when highly integrated PC processors have less than one-fifth as much. The HP paper presented one implementation detail in an unusual way: the designers say that RC models show the on-chip clock signals travel at 90% of the speed of light; since other designers don't quote the speed of signals this way, we don't know if this speed is exceptional or not. The chip is packaged in a 544-pin ceramic LGA (land-grid array) and achieves a 500-MHz operating frequency at 85° C with a 2.0-V supply (but so far only 440 MHz in real systems).

AMD presented the forthcoming K7, but only a few new details were revealed. A couple of tidbits left out of the Microprocessor Forum talk last October (see MPR 10/26/98, p. 1) were the size of the integer file and the FP-register file. At ISSCC, AMD disclosed that the K7's FP-register file has 88 entries, each 90 bits with five read and five write ports. The 32-bit-wide integer file has 24 entries, each with nine read and eight write ports.

This structure combines the standard integer registers into an integrated future file/register file (IFFRF). A future file is used in an out-of-order processor to avoid the associative lookup in a reorder buffer when accessing register values; in AMD's implementation, the IFFRF stores both architectural state (normal register file for retired state) and speculative state (future file) and is thus roughly double the size of a simple register file. Because a future file doesn't eliminate the need for a reorder buffer and the K7 allows up to 72 active instructions, there is a large reorder buffer in addition to the IFFRF. AMD chose this organization because speed is improved compared with associative lookup in a reorder buffer.

AMD separately described the K7's floating-point unit. One interesting detail is that 80% of the FPU is implemented in 150,000 standard cells.

Your Data Is Safe With IBM

IBM presented some real meat in two papers on a CMOS implementation of the Series 390 architecture. While everyone else is focusing on clock rate, issue rate, and a slightly faster external interface, IBM is hard at work improving reliability and robustness. The CPUs are not the fastest uniprocessors, but the resulting system justifies the adage "no one ever got fired for buying IBM."

The G5 microprocessor is the heart of an S/390 mainframe (or multiprocessor server, as IBM now styles them); up to 12 of the chips are used in a single system. Implemented in 0.25-micron CMOS, the device runs at a respectable 600 MHz but is used at 500 MHz in the 10+2 SMP configuration (10 processors for the workload, two for I/O and system-level maintenance). IBM claims measured performance of this configuration is more than 1,000 S/390 MIPS. The G5 chip weighs in at 215 mm^2 and 25 million transistors, with the 256K unified on-chip L1 cache and other memory arrays accounting for 18 million.

As figure 1 shows, the chip contains two complete processors (the L1 is shared), but not for superscalar or multithreaded performance enhancement. Instead, the processors run in lockstep, and the results are compared on a per-cycle basis. If a mismatch is detected, the processors are cleared, the microarchitectural state of both is reloaded from a checkpoint copy kept in the comparison unit (the R-Unit in Figure 1), and the processors are restarted. If the recovery fails, the state of this bad chip is unloaded, and eventually another processor will pick up where the failed one left off. Now that's fault tolerance!

One indication that this is, after all, IBM, is the dual-format floating-point unit. The FPU supports both IEEE-standard and legacy IBM hexadecimal-format floating-point

operands. Although IBM's hex format has been shown to be numerically inferior to binary formats, the chip actually directly implements only the hex format; IEEE operands are converted to hexadecimal, which has extra internal precision, and the results are converted back to IEEE. This conversion increases the latency of IEEE operations, halving the issue rate to every other cycle, indicating that the chip is intended mainly for legacy code.

A complete G5 system uses eight custom level-two (L2) cache chips and two system-controller chips, which also hold L2 cache tags. Each 1M cache chip is a 273-mm^2 die with 59 million transistors. Even the MCM used to connect the chips is impressive: it has 75 layers with about 600 meters of wire, about 10,500 nets, and 2,664 signal I/O pins (4,224 total pins). When it comes to mainframes, IBM doesn't scrimp. Just for good measure, the MCM holds two cryptographic coprocessors in addition to the 12 microprocessors, eight L2 cache chips, and two system-controller chips. Anticipating the question, IBM claims the crypto chips facilitate e-commerce.

The system bandwidth is well balanced. Each microprocessor starts with 8 GBytes/s of data bandwidth. The system provides 16 GBytes/s of I/O bandwidth: 4 GBytes/s for each of four I/O adapters. The memory bandwidth is also 16 GBytes/s. The 12 processors are grouped into two 6-processor nodes with a node-to-node bandwidth of 16 GBytes/s. IBM claims its tests confirm that the bandwidth goals have all been met. Realistic S/390 workloads use only about 25% of the bandwidth, so there is still room for this system design to grow.

Sincere SOI Session

Silicon-on-insulator (SOI) process technology (see MPR 8/24/98, p. 8) got its own session this year, and some real developments were reported. IBM described two SOI PowerPC microprocessors, both using SOI and copper interconnect. In the first, a 32-bit processor, the use of SOI was estimated to bring a 20% performance improvement. A graph of the effects showed that the benefit was realized in custom circuits, not standard cells or interconnect. The second chip, a 64-bit PowerPC server chip from the AS400 division, realized 15–40% gains at the circuit level and 24–28% improvement in critical paths. A graph showed the higher levels of SOI performance gain usually obtained in complex circuits, such as an XOR gate, which was 40% faster in SOI.

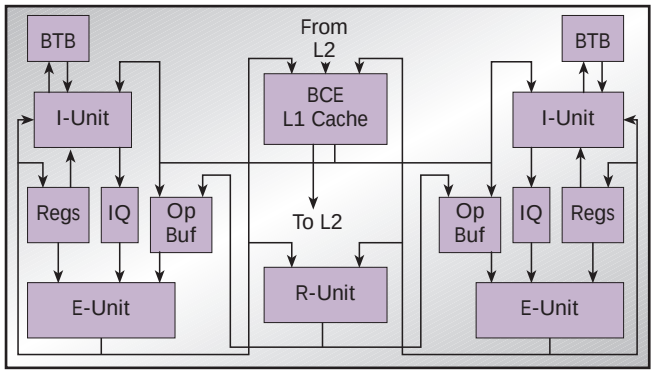

Figure 1. The register-checkpoint unit (R-Unit) of the IBM G5 S/390 chip holds the architectural state of the microprocessor, compares the results of the redundant execution cores, and stores a copy of the entire microarchitectural state that is reloaded into the mirrored cores after a mismatch.

Samsung presented the results of its SOI Alpha 21164 chip. The chip achieves the same 600-MHz clock speed as the standard CMOS 21164 but dissipates just 14 W using a 1.5-V supply, representing a dramatic improvement over the 25 W of the CMOS chip. The SOI chip is built in a 0.25-micron four-layer-metal SOI CMOS process and measures 209 mm^2.

Circuits Galore, But Diversions Available

Most of the conference presentations actually focused on circuits, but, as usual, a dose of the unusual was present. In a possible foreshadowing of post-silicon technology, Alan Johnson of the University of Pennsylvania revealed the "Electronics of Single-Wall Carbon Nanotubes." These SWNTs are formed by taking a narrow sheet of carbon and joining the opposite edges. Depending on the alignment of the atoms in the lattice, the resulting structure behaves like a metal or like a semiconductor. Some circuit elements have been demonstrated, and others are possible in theory. According to the paper, terahertz switching speeds are possible, making these devices hundreds of times faster than today's fastest CMOS gates.

Two presentations, one from the University of Michigan and one from North Carolina State, detailed neural stimulation ideas. Michigan's FINESS (fully integrated neuromuscular electrical stimulation system) receives power and data from an external 4-MHz RF signal. The single chip integrates a coil, a 4-V regulated power supply, and all the logic needed to decode the data from the 4-MHz carrier and stimulate a nerve. The chip derives 14.8 mW of power from the RF signal when performing stimulation.

The North Carolina research proposes "An Implantable Neuro-Stimulator Device for a Retinal Prosthesis." The impetus for this research is the fact that certain diseases destroy the retina's rods and cones, but the nerves can still respond to direct electrical stimulus. The team's chip uses an external coil but, like the Michigan chip, recovers power and data on-chip. The prototype can drive a 10 × 10 array of retinal electrodes at video data rates.

ISSCC (*www.isscc.org*) is mostly about circuits and the effective use of fabrication technology, but much of the material is still accessible to non-circuit designers. For example, you don't have to understand the operation of 450-MHz disk-drive read channels to understand that disk transfer rates should continue to increase in the near future. The bottom line is that the pace of innovation in IC fabrication and circuit technology shows no signs of slowing, and we can expect meaningful and exciting improvements in anything that uses integrated circuits. M