

ISSCC '95 Foreshadows Gigachip Era

DRAMs Reach One-Gigabit—Processors Exceed 1 GIPS

by Brian Case

As if mimicking the rate of advancement in the IC technology upon which most of the conference is based, the proceedings of the 42nd International Solid State Circuits Conference again showed substantial growth. Held in San Francisco on February 15–17, ISSCC expanded its program to a full three days with four parallel tracks (up from two-and-a-half days with four tracks last year). ISSCC presentations typically represent research projects that are often years away from resulting in products. The microprocessors session, however, has been the major exception since the “first public disclosure rule” has been relaxed in recent years.

The program covered the usual breadth of topics—including memories, DSPs, microprocessors, analog circuits, and gate arrays—but the most significant development was the arrival of the “gigachips.” Of interest to MPR readers are NEC’s and Hitachi’s first 1-Gbit (i.e., 128-Mbyte) DRAM chips, DEC’s 300-MHz/1.2-GIPS 21164 Alpha microprocessor, and a 2.4-GOPS multi-processor DSP chip from UC Berkeley. Also appearing in “giga” scale were serial-communications circuits, optical preamplifiers, and sample-and-hold circuits.

The microprocessor session contained seven papers describing chips relevant to the merchant market, with all but the P6 previously disclosed (ISSCC was the first disclosure of the P6 design). Still, some interesting details were presented, and some new insight about Intel’s P6 was gained from a demonstration of a working system.

1-Gbit DRAMs Require Huge Die Size

The big news, of course, was the 1-Gbit chips, but since 16-Mbit chips are still on the leading edge of the mass market, DRAMs with lower density and/or other useful features were also described at the conference.

NEC described its 1-Gbit DRAM both in a regular paper and at a press dinner. Some NEC representatives were wearing tie tacks featuring the 1-Gbit die, which were almost wider than the tie being tacked. With chip dimensions of 25.8×36.3 mm (936 mm^2 !), the tie tacks were simply gaudy.

NEC uses a clever approach to increase yield of the 1-Gbit part. As shown in Figure 1, a wafer of 1-Gbit die using a conventional approach has few die sites and few—if any—good die. NEC’s approach fabricates individual 256-Mbit macros. Each of these macros can be used either as a standalone 256-Mbit chip or as any of the four corners of a 1-Gbit chip. The two wafers in Figure 1 have the same pattern of defects, but NEC’s approach allows the wafer to yield three 1-Gbit and many 256-Mbit chips instead of only one 1-Gbit chip.

The access-time characteristics of NEC’s 1-Gbit DRAM make it appropriate for file applications instead of main memory. The first access takes a leisurely 150 ns, but thereafter serial accesses can occur at a 10-ns rate. The chip uses a $\times 32$ organization to deliver bursts of data at 400 Mbytes/s. Fabrication is in a 0.25-micron process, the power supply is from 2.0 to 2.5 V, and operating current is 68 mA at 2 V and 100 MHz.

Hitachi used a more aggressive process to fabricate its 1-Gbit part, resulting in both higher frequency operation and a smaller die size. Using a 0.16-micron process (three-metal), the Hitachi chip measures 19.3×37.1 mm for a die area of “only” 715 mm^2 . RAS access time is 33 ns for the first access, but subsequent burst accesses take only 4.5 ns, a clock rate of 220 MHz. The chip interface uses a $\times 16$ organization, so the burst data-transfer rate is 440 Mbytes/s, 10% higher than that of the NEC part.

Matsushita, which detailed a 256-Mbit chip last year, focused instead on low-power and long-retention DRAMs this year. Its paper described a 16-Mbit chip that requires less than $0.5 \mu\text{A}$ of retention current per megabyte. According to the company, a 20-Mbyte RAM disk using 10 of these chips

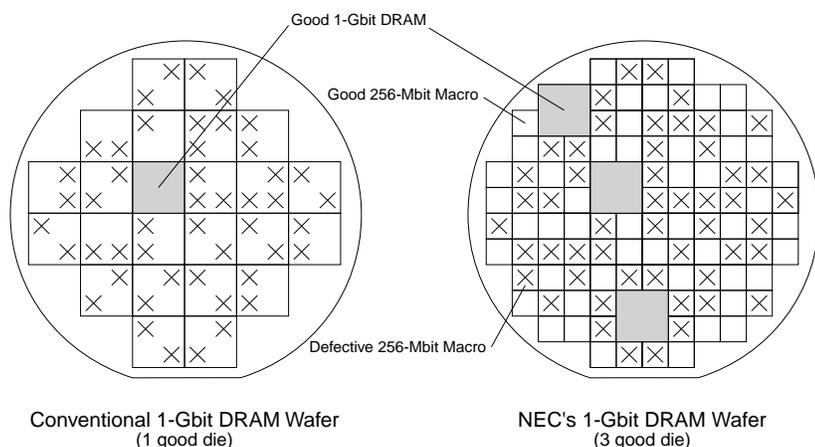


Figure 1. NEC’s clever approach to building 1-Gbit DRAMs uses 256-Mbit macros that can be configured either as standalone 256-Mbit chips or as any of the four corners of a 1-Gbit chip. Yield of 1-Gbit chips is increased from one to three.

would retain its data for 2.5 years with a single button-sized 190-mAh lithium battery. The chip can operate from a supply voltage of 1.8 V to 3.6 V, and die size in the 0.5-micron (three-poly, one-polycide, two-metal) process is a manufacturable 104 mm² (6.5 × 15.9 mm). RAS access time at 1.8 V is 41 ns. This is a DRAM chip that seems ready for real applications, and it will be even faster and smaller in a 0.35- or 0.25-micron process.

In addition to its 1-Gbit chip, Hitachi talked about a 64-Mbit DRAM that it jointly developed with Texas Instruments. This chip tests a hierarchical word-line scheme to deal with the fact that word-line pitch becomes a dominating factor in die size as memory cell size shrinks. This test chip is built in a 0.25-micron process (three-metal) and has a die size of 144 mm² (15.9 × 9.1 mm). The chip accepts a 3.3-V supply but generates 2.2-V power for internal operation, consumes 130 mA active power, and has a fast access time of 29 ns.

Mitsubishi, like Matsushita, is developing low-power chips for emerging high-volume, battery-operated devices. Mitsubishi described a 1.2-V, 16-Mbit DRAM. At this voltage, the chip consumes only 13 mA and delivers a RAS access time of 49 ns. Standby current is 0.5 μA. Using a 0.4-micron process (four-poly, two-metal), it has a die area of 132 mm² (16.7 × 7.9 mm).

Hyundai described its 256-Mbit chip that uses wave-pipelining methods to achieve high-speed operation. Wave pipelining uses the known delay characteristics of the circuit to achieve pipelined efficiency without pipeline latches. One set of signals is sent through the logic path (the column-access path in this case), and, before it propagates to the end of the circuit and is latched, another set of signals is started. This self-timed operation results in a 150-MHz burst data access time. With a ×16 organization, the chip achieves a 300-Mbyte/s data transfer rate. Hyundai's chip is fabricated in a 0.3-micron process, operates from 2.5 V or 3.3 V, and has a RAS access time of 33 ns.

Flash Memories Show Progress

Flash-memory development continues to show rapid progress, indicating that chip companies still expect flash to be important. Intel, Toshiba, Samsung, and SunDisk/Matsushita all described 32-Mbit flash chips. Intel's novel chip stores two data bits per flash-memory cell (a technique used in the microcode ROM of the 8086). This Intel device is not simply a flash chip but a flash system on a chip. As with some other Intel flash chips, this device contains a small RISC processor to handle on-chip functions.

Storing two bits per cell requires four voltage levels. The processor, using special sensing circuitry, applies programming pulses to memory cells until they are initialized to the desired voltage. Thus, programming time can vary, but write throughput is typically 10 ms/byte.

The chip is fabricated in a 0.6-micron process (one-poly, one-polycide, two-metal), has a die size of 152 mm² (13.7 × 11.1 mm), and is packaged in a 56-lead TSOP package.

Note, however, that Samsung's 32-Mbit chip is 37% smaller at 95 mm² in a 0.5-micron, two-poly/one-metal process. Thus, Intel will have to shrink its die considerably to compete on manufacturing cost.

P6 Voltage May Change

The microprocessors session contained seven papers describing the most interesting recent developments in RISC and x86 chips; all, however, have been described in detail in recent articles. The papers described:

- Hal/Fujitsu Sparc64 chip (*see 090301.PDF*)
- NexGen Nx586 chip (*see 080403.PDF*)
- IBM/Motorola PowerPC 620 (*see 081402.PDF*)
- Intel P6 (*see 090201.PDF*)
- Sun UltraSparc (*see 081301.PDF*)
- IBM/Motorola PowerPC 602 (*see 090203.PDF*)
- Digital 21164 (*see 081201.PDF*)

We gleaned some interesting information from an evening P6 demonstration given by Intel in a hotel room. Intel showed three different motherboards/daughter cards with P6 sockets. On the boards, an IDC-header-like connector was next to each P6 socket (or sockets for the multiprocessor boards). The connector holds a hefty DC-to-DC converter to provide the 2.9-V power for the P6. Since standard PC power supplies do not, of course, generate a 2.9-V supply, some sort of accommodation is to be expected in prototype motherboards. If the converter is required on production boards, however, the cost of using the P6 includes the chip *and* converter.

The power converter might be removable for a good reason: when Intel shifts production of the P6 to its 0.35-micron process, power supply voltage might need to be reduced even more, say to 2.5 volts. Dropping the voltage again would reduce power consumption significantly, allowing Intel to promote P6 for a broader range of applications, including laptop computers. If Intel is unsure about exactly when production will shift and so is telling its customers they will be getting a mixture of 2.9-V and 2.5-V parts, the removable converter makes sense.

Whatever the case, system vendors eyeing the market for P6 machines might want to start immediately locating a source of special voltage converters or P6-specific power supplies.

Other interesting microprocessor developments—not in the microprocessor session—included a couple of low-power chips. Toshiba described a 150-MIPS/W MIPS-architecture processor (the R3900, see MPR 2/16/95, p. 20) for PDA applications. Aiming at battery-operated consumer devices, NEC detailed a 0.9-V, 100-MHz 16-bit DSP core. Integrating a 16 × 16 hard-ware multiplier, a 32-bit adder, 1K of SRAM, and a PLL for clock generation, the core dissipates a miserly 4 mW of power. Built in an

aggressive 0.25-micron two-metal process, the core fits in only 2 mm² of chip area (1.4 × 1.4 mm).

3D Multichip Module

A paper from the University of Sheffield in the U.K. described a high-density multichip module (MCM) packaging technique. Referred to as MCM-V (for vertical), this three-dimensional package was used to build an extremely small integrated camera-and-image-processing system. Figure 2 shows a schematic cross-section of the experimental module.

The package is formed by first wire-bonding die to flexible layers of laminated film. Discrete components are surface mounted. The populated substrates can then be tested conventionally by using test pads on the edges of the film. Functional substrates are stacked and encapsulated in epoxy. The cured epoxy is then sawed to size. The sawing exposes connections on the faces of the three-dimensional package, which are plated with metal. Finally, the metal is patterned with a laser and etched to leave only the desired traces. The package itself can be attached to a PCB using a lead frame or wire bonding.

The system described in the paper incorporates a grayscale camera chip, four Transputer microprocessors, 64K of SRAM, a FIFO chip, an analog-to-digital converter, and 40 discrete components. The camera chip has a resolution of 312 × 287 pixels. Pixel values are quantized to 8 bits by the ADC, and the FIFO buffers data between the ADC and the Transputers. The first Transputer supervises the compression process and transfers data to the other processors via the built-in serial-communication links. Image size is reduced by decimation to QCIF resolution of 176 × 144 pixels.

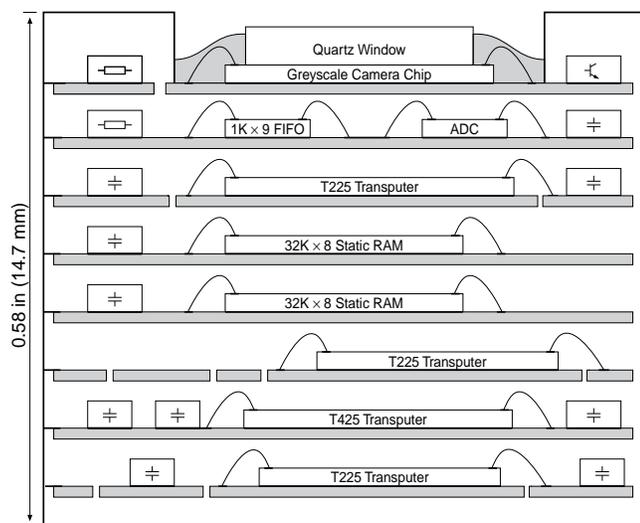


Figure 2. This MCM-V (vertical) package from the University of Sheffield consists of layers of flexible laminated substrate to which die are wire bonded. The system shown above crams four processors, 64K of SRAM, and 40 discrete components forming a three-frames-per-second JPEG compression system into 0.3 in³ of space.

In 0.3 cubic inches of space, the system implements enough processing power to perform JPEG image compression at a rate of three frames per second. The authors claim this packaging technology is low cost and achieves six times the density of advanced PCBs. One problem could be heat dissipation, but even if this packaging technology is restricted to low-power circuits, it could have wide applicability.

Other Papers Described Micromachining

As usual, ISSCC '95 included a sprinkling of unusual papers. Analog Devices and a team from the Center for Integrated Systems at Stanford described surface micromachined chips. The Analog Devices chips are accelerometers for automobile airbag systems. Surface micromachining—instead of bulk micromachining, which etches from both sides of a wafer—reduces the cost of the accelerometer.

The Stanford team used micromachining to suspend a temperature-regulated voltage reference in mid-air. Arms of oxide suspend a small amount of circuitry—including a heater—over a pit etched in the silicon substrate. The result is a voltage reference requiring 200 times less power than previous heated-substrate circuits and warms up 150 times faster.

A team from UC Berkeley talked about an experimental multiprocessor DSP chip. The chip contains 48 16-bit processing elements (PEs), each with a small local memory and connected by a two-level high-bandwidth network. The chip is fabricated in what now seems like ancient technology: 1.0-micron, two-metal CMOS. Even so, the die size is not too bad for 600,000 transistors, at 12 × 12 mm (144 mm²). Performance on some common DSP algorithms is compelling. For an eight-state Viterbi detector, 34 PEs achieve a 12.5-Mbyte/s channel data rate, which, according to the authors, is more than ten times faster than an advanced DSP.

Future Looks Bright for Chip Industry

With the announcement of the first experimental 1-Gbit DRAMs and the increasing use of 0.25-micron and even sub-0.2-micron processes, IC fabrication technology shows no signs of slowing its frantic pace of advancement. For example, the Hitachi 1-Gbit DRAM is fabricated using a direct EB (electron-beam) hybrid lithography, which eliminates the traditional optical reticle limit by allowing a die to be “stitched” together from several separate exposures. Even if widespread use of EB lithography is years away—the 1-Gbit DRAM generation is still about a decade in the future—larger die sizes plus smaller geometries guarantee a bright future for the chip—and microprocessor—industry. ♦

Copies of the ISSCC Digest of Papers are available for \$125; contact John Wuorinen at 207.326.8811.