

# M I C R O P R O C E S S O R

www.MPRonline.com

THE INSIDER'S GUIDE TO MICROPROCESSOR HARDWARE

## Z-RAM SHRINKS EMBEDDED MEMORY

*Innovative Silicon's Tiny DRAM Cells Alter the Memory Equation*

By Tom R. Halfhill {10/25/05-03}

Many of today's microprocessors seem like memory chips with embedded logic—their on-chip caches and memories account for most of the transistors and sometimes most of the die. External memory hasn't kept pace with processor speeds, so it's easier to improve performance by

slapping down more cache arrays than by designing a more sophisticated processor core. In addition, complex processor cores use more power when they're active, and they leak more current when they're not.


Even with the advent of multicore processors, the long-term trend is toward more on-chip memory. This is certainly true of PC and server processors, because their general-purpose workloads benefit greatly from larger multilevel caches. But it's also true of embedded processors, ASICs, SoCs, DSPs, and microcontrollers. Figure 1 shows how industry forecasters expect embedded memory to dominate the die area of future processors.

L1 cache arrays will continue to use SRAM for the near future, because it's the fastest memory technology and it scales readily to smaller fabrication processes. However, each SRAM bit-cell requires four to six transistors, so it's too costly for larger amounts of embedded memory in microprocessors and microcontrollers. Flash memory requires fewer transistors than SRAM does and is non-volatile, but it's very slow—often too slow for direct execution. Embedded DRAM (eDRAM) is another option. Although it's not as fast as SRAM, it requires only one transistor per bit-cell, and it's much faster than flash memory.

The problem with eDRAM lies in scaling to smaller fabrication processes. As Figure 2 shows, the deep trench required for each bit-cell's capacitor distorts the aspect ratio of the cell's structure and doesn't scale well beyond the 90nm node. In addition, changing the state of an eDRAM

bit-cell requires more power to overcome the cell's relatively high capacitance, and the deep trench is incompatible with the extra substrate in a silicon-on-insulator (SOI) process.

Earlier this year, a Swiss startup, Innovative Silicon, announced a new embedded-memory technology to meet all these challenges. Innovative Silicon calls its technology Z-RAM, because each one-transistor bit-cell requires zero capacitors. Z-RAM isn't fast enough to replace SRAM, but it's



**Figure 1.** Embedded memory already accounts for more than half the die area of typical microprocessors and SoCs, and it will soon overwhelm the silicon devoted to logic. Sources: Innovative Silicon, the Semiconductor Industry Association (SIA), and the International Technology Roadmap for Semiconductors (ITRS 2000). \*Forecasts.



**Figure 2.** Conventional embedded DRAM (eDRAM) requires a deep-trench capacitor structure in addition to the transistor for each bit-cell. The capacitor trench makes the bit-cell much taller than its width and poses a problem for advanced fabrication processes.

faster than conventional eDRAM and doesn't need a separate capacitor to temporarily hold the cell's binary state. Instead, Z-RAM exploits an inherent electrical effect of SOI technology for that purpose. In a technical presentation at **Fall Processor Forum 2005**, Innovative Silicon's president and CEO, Mark-Eric Jones, explained how Z-RAM works and made a strong argument that Z-RAM is the logical alternative for embedding memory in future microprocessors.

### Exploiting the Cinderella Effect


Jones became CEO of Innovative Silicon after leaving MoSys, another company offering one-transistor embedded

memory for microprocessors. (See *MPR 9/13/99-05*, "MoSys Explains 1T-SRAM Technology.") However, Z-RAM works on an entirely different principle than MoSys 1T-SRAM, and the companies have slightly different business strategies. For now, at least, Innovative Silicon concentrates on licensing Z-RAM for embedded memory in microprocessors and SoCs. MoSys licenses 1T-SRAM to memory manufacturers as well as to processor vendors.

Z-RAM exploits an electrical characteristic of SOI transistors known as the floating-body effect or history effect. (Innovative Silicon also calls it the Cinderella effect, because Z-RAM transforms a disadvantage into an advantage.) SOI transistors still have some residual capacitance, even though the purpose of the extra insulating layer beneath the gate dielectric is to reduce capacitance so the transistor can switch states more rapidly or consume less power when changing states. The residual capacitance is generally considered parasitic. But in 1990, a researcher at Belgium's Interuniversity Microelectronics Center found a way of using the floating-body effect to temporarily store a bit in a one-transistor memory cell. Unfortunately, isolating each bit-cell within a memory array was a problem—writing one cell often corrupted others.

In 2001, Swiss engineer Dr. Pierre Fazan published a landmark paper showing how to apply the floating-body effect to a memory array. Fazan, formerly with Micron Technology, cofounded Innovative Silicon a year later and is the CTO. His coinventor, Dr. Serguei Okhonin, is the company's chief scientist. Their patented innovation is to use the residual capacitance of an SOI transistor to store the bit-cell's binary state and prevent read/write operations from altering other cells in a memory array. The natural gain effect of the transistor's gate amplifies the residual capacitance, so Z-RAM uses the transistor for gain as well as for a switch. Figure 3 illustrates the principles of the Cinderella effect.

Reading a Z-RAM bit-cell is similar to reading any DRAM cell. The chip applies a small pulse to the cell's transistor and uses a sense amplifier to compare the resulting



**Figure 3.** How Z-RAM works. Passing a current through the channel builds up excess positive charges in the transistor, which decreases the threshold voltage and the current flowing through the channel. This action writes binary 1 to the bit-cell. Applying a bias to the transistor channel removes the holes through the junction, writing binary 0. Innovative Silicon refers to positive charging as "impact ionization" and to negative charging as "hole removal." In either case, the transistor's capacitance retains the bit value between periodic refresh cycles.

current flow with the current in a reference cell. Figure 4 shows a schematic of a Z-RAM cell and a graph illustrating the way the cell responds to the sense pulse, depending on whether the stored bit is 0 or 1.

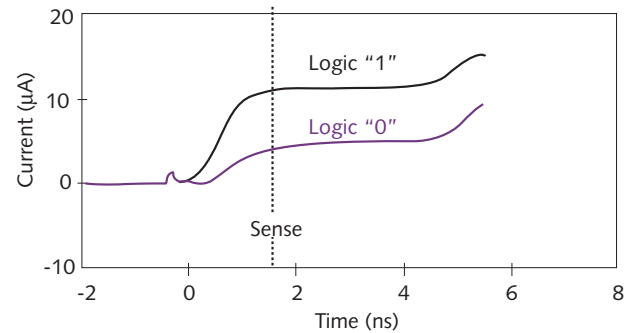

Although Z-RAM isn't as fast as SRAM, it's faster than conventional DRAM, and each cell requires only one transistor instead of the four to six transistors in an SRAM cell. Z-RAM test chips fabricated in a 90nm SOI process have demonstrated read/write speeds faster than 3ns, and performance will scale with future process shrinks. Innovative Silicon estimates that Z-RAM arrays fabricated in a 90nm process can achieve densities of 0.30mm<sup>2</sup> to 0.40mm<sup>2</sup> per megabit—twice the density of conventional eDRAM and five times denser than six-transistor SRAM. By enabling designs that are more tightly packed, Z-RAM can reduce wire delays, thus improving the microprocessor's overall performance. Figure 5 illustrates the physical structure of a one-transistor, zero-capacitor Z-RAM bit-cell.

### Numerous Advantages of Z-RAM

Discarding the trench capacitor of a conventional DRAM cell gives Z-RAM several advantages. First, there's no trench to interfere with the SOI substrate, so it's compatible with all SOI fabrication processes, including partially depleted SOI, fully depleted SOI, and future FinFET technology (double-gate MOSFET transistors on an SOI substrate). Second, Z-RAM scales more readily to smaller fabrication processes than DRAM does, because the Z-RAM cells can shrink in all dimensions, whereas the deep trench of a DRAM cell severely limits vertical scaling. In fact, DRAM scaling becomes progressively worse at smaller geometries, because the capacitor structure must grow proportionally larger to maintain adequate charge storage.

Third, the absence of a separate structure for the capacitor makes Z-RAM twice as dense as conventional DRAM and about five times denser than SRAM. Fourth, Z-RAM read/write performance is faster than conventional DRAM, because SOI transistors have less capacitance to overcome when changing states. Fifth, the simpler structure of a Z-RAM cell and its exploitation of a former disadvantage (parasitic capacitance) could make an SOI chip as inexpensive to manufacture as a bulk CMOS chip—or even less expensive than bulk CMOS.

That last advantage—lower cost—depends on how much Z-RAM is embedded in a processor. Normally, SOI is 10–15% more expensive than bulk CMOS at the same scale. However, Z-RAM can significantly reduce the total die area of a chip in comparison with the same processor using conventional eDRAM and bulk CMOS. If an SOI chip embeds enough Z-RAM, it starts getting cheaper to manufacture than a bulk CMOS chip with eDRAM. As the amount of memory embedded in microprocessors continues to climb, Z-RAM could make SOI more affordable than bulk CMOS, especially if the forecasts in Figure 1 are anywhere in the ballpark. The breakpoint occurs when memory occupies about 18% of the die.



**Figure 4.** Detecting the difference between a stored 0 or 1 in a Z-RAM bit-cell is similar to sensing the value of a conventional DRAM bit-cell. The transistor's current flow diverges sharply at the threshold of the sense pulse, as this graph illustrates. At the top is a schematic of the cell.

Of course, SOI has its own advantages, apart from enabling Z-RAM. Depending on a project's design goals, SOI can boost the processor's clock frequency by about 35% (because the lower-capacitance transistors switch states more quickly) or reduce power consumption by a like amount (because the lower-capacitance transistors require less energy to change states). If the primary reason for using SOI is to embed Z-RAM, the processor may get the extra performance of SOI essentially for free. If the primary reason



**Figure 5.** Z-RAM requires only one transistor, like conventional DRAM, but it doesn't need the deep-trench capacitor shown in Figure 2. As a result, it's compatible with SOI—indeed, it relies on the SOI floating-body effect to work—and it scales more easily to smaller fabrication processes.

### Price & Availability

Innovative Silicon is licensing Z-RAM technology now. An instance license allows customers to use unmodified Z-RAM macrocells in chip designs; a technology license allows customers to modify or develop their own Z-RAM cells; a compiler license allows customers to generate Z-RAM cells with a memory compiler. Innovative Silicon doesn't publicly disclose upfront licensing fees or chip royalties. For more information, visit [www.innovativesilicon.com](http://www.innovativesilicon.com).


for using SOI is to improve clock speed and/or save power, the processor can embed Z-RAM essentially for free.

And it gets better: Z-RAM adds no extra masks or manufacturing steps to a standard SOI fabrication process. In contrast, conventional eDRAM requires both additional masks and additional manufacturing steps. This difference could be crucial to the industry's adoption of Z-RAM, because the technology is compatible with existing design flows and with SOI processes at fabs and foundries worldwide.

Innovative Silicon says it has produced test silicon with Z-RAM at nine different fabs, using transistors with gate lengths as short as 40nm. Even the initial unoptimized Z-RAM cells—which measure 0.18 square microns in a 90nm process—are smaller than production-grade eDRAM cells fabricated in 90nm. The company estimates that optimized Z-RAM cells in 90nm will measure only 0.10 square microns. Figure 6 is a microphotograph of an actual Z-RAM FinFET bit-cell implemented in silicon.

### Practical Applications for Z-RAM

In theory, memory manufacturers could license Z-RAM to make DRAM chips twice as dense as conventional DRAM. However, conventional DRAM processes don't use SOI, so adopting Z-RAM would require memory manufacturers to expensively retool their fabs. For that reason, Innovative



**Figure 6.** This microphotograph shows a one-transistor Z-RAM FinFET bit-cell fabricated for test purposes.

Silicon has decided to concentrate on licensing Z-RAM for SoCs, embedded memory, general-purpose processors, and graphics processors at first—SOI is a more common option in logic processes. If the company can successfully enter those markets, memory manufacturers will provide the high-hanging, but perhaps juiciest, fruit.

Z-RAM isn't fast enough to replace SRAM in the L1 caches of microprocessors, but L2 and L3 caches could use it. Thanks to the transistor's gain effect, Z-RAM retains its state for about the same amount of time as eDRAM, even though a Z-RAM cell has much less capacitance. As a result, the refresh rate for Z-RAM is about the same as for eDRAM, but dynamic power is about 30% lower. It's possible to save even more power if the greater density of Z-RAM enables designers to embed more on-chip memory, thus reducing the power consumed by accessing off-chip memory.

Soft-error rates are the biggest remaining question. Z-RAM relies on the natural gain effect of its transistors, and the similar gain of multitransistor SRAM cells tends to amplify soft errors, especially as smaller fabrication processes pack the transistors more closely together. Although Innovative Silicon hasn't publicly disclosed soft-error data for Z-RAM, the company says initial tests look good and that error rates should be lower than for conventional DRAM or SRAM.

One reason that Z-RAM might have fewer soft errors than other memory technologies is that Z-RAM uses SOI, which typically reduces soft errors by about 5×. Another reason is that a Z-RAM bit-cell transistor also functions as the first step in the sense-amplifier chain, which may make it more resistant to interference. Innovative Silicon compares this effect to using a masthead amplifier on a radio or TV antenna instead of boosting gain in the receiver.

After three years of development, building on a dozen years of theoretical work, Z-RAM is finally moving out of the lab. Innovative Silicon has produced multimegabit test chips using 90nm SOI processes at Freescale Semiconductor and TSMC. Additional test chips are now being fabricated at 65nm. Innovative Silicon has begun sharing some early test results with potential customers under nondisclosure agreements and hopes to publish detailed reports in 2Q06. Interested parties should soon have enough information about actual memory densities, read/write speeds, power consumption, refresh rates, soft errors, and other performance characteristics to decide whether Z-RAM is ready for integration in production devices.

If test reports verify that the Cinderella effect and Z-RAM aren't fairy tales, *Microprocessor Report* expects this new technology to succeed, because it meshes smoothly with existing design and manufacturing practices. Z-RAM might even be attractive enough to lure SOI holdouts like Intel. ♦

To subscribe to *Microprocessor Report*, phone 480.483.4441 or visit [www.MPRonline.com](http://www.MPRonline.com)