Electronic Memory Goes High Rise

THE processing speed of the average commercial desktop computer is increasing at a good clip. At the same time, memory latency—or the time it takes for a computer's central processing unit to grab a piece of information stored in random access memory (RAM)—is increasing much less quickly. One problem is that in conventional dynamic random access memory (DRAM) or static random access memory (SRAM), each line in a two-dimensional memory array is managed by one switch. The farther a piece of memory is from the switch or central processing unit, the longer it takes to retrieve it. This problem is compounded by the complicated code and data structures in modern software applications that require access to a vast amount of random memory.

Reducing memory latency is one of several major challenges to developing the massively parallel computers of the Accelerated Strategic Computing Initiative (ASCI) beyond the next generation. ASCI computers are a key component of the Department of Energy's program for stewardship of our nation's stockpile of nuclear weapons. Combined with nonnuclear experiments, simulations of nuclear implosions and other phenomena in three dimensions are needed so that scientists can assure that the stockpile remains safe and secure without underground nuclear testing. The current ASCI White computer, the most powerful in the world, operates at 12 teraops (trillion operations per second). The next-generation ASCI computers will operate in the 30 to 70 teraops range. However, the imbalance between microprocessor speed and the delivery time of information to the processor will hamper further ASCI development and severely limit performance.

Scientists at Livermore think that magnetic random access memory (MRAM) might be one solution to the memory latency problem. They are working on integrating a diode switch on top of the magnetic tunneling junctions that make MRAM work, creating more efficient vertical RAM. Mathematical physicist Charles Cerjan, who is leading work on MRAM at Livermore, says, "To decrease memory latency, we want to get away from the flat suburban sprawl of today's memory systems and create high-density 'skyscrapers' of RAM."

In addition to speeding up information access, MRAM offers several other advantages. It is immune to radiation damage, consumes little power, and continues to function over wide temperature ranges. Unlike most other forms of RAM, magnetic RAM is nonvolatile, which is to say that it retains its memory even after power is removed. For example, after the explosion of the space shuttle Challenger in 1986, NASA was able to retrieve the shuttle's magnetic memory—still readable—from the bottom of the Atlantic Ocean.

Magnetic memory has been around for a long time as cassette tapes and disk drives. But until recently, fast, high-density MRAM was not possible because the access times to read and write data were inferior to those in semiconductor-based memory. Great strides have been made in manufacturing thin-film multilayers, which are key to MRAM's operation. Challenges remain, however. Primary among them are finding the right combination of multilayers

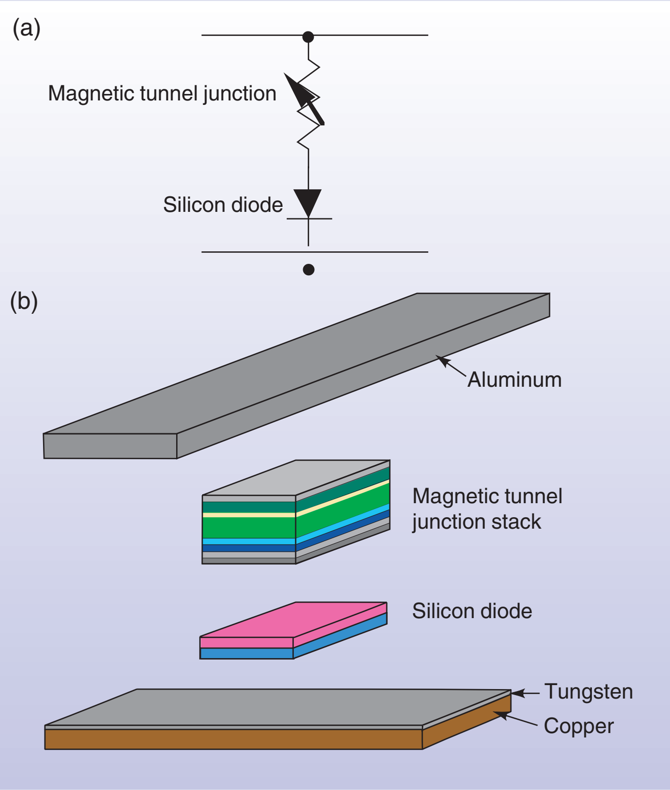

(a) A simplified circuit diagram showing the configuration of a magnetic memory cell containing a magnetic tunnel junction. (b) An expanded view of the multilayer stack and diode configuration.

to maximize their performance and attaching a microprocessor to the magnetic material to bring memory and processing as close together as possible.

Building on Success

Livermore's current work on MRAM is the successor to a magnetic ultrahigh-density read head for computer disk drives, for which Cerjan and other Livermore researchers won an R&D 100 Award in 1996 (see *S&TR*, October 1996, pp. 24–25). Using Livermore's expertise in thin, multilayer films and microfabrication technologies, they developed a layered sensor that is smaller and offers better performance than conventional magnetic sensors. It has alternating magnetic and nonmagnetic layers, each less than 5 nanometers thick. The total thickness of the sensor is only about 100 nanometers, or about one-thousandth the width of a human hair.

Working with two industrial partners to commercialize the device, the Livermore researchers modified their original design into a magnetic tunnel junction, in which an insulating barrier is inserted between two different types of magnetic material. When a current is applied perpendicularly to the layers, it "tunnels" through the insulator. The relative magnetic field alignment of the two separated magnetic materials induces either a low- or high-resistance current path. This resistance difference can be identified as a stored bit of information, either a 0 or a 1. Already, these individual memory elements have comparable or superior performance to any reported in the literature.

Overcoming Incompatibilities

Experiments are under way to attach a diode switch to this magnetic sandwich, the first step before attempting to attach a microprocessor. Standard mechanisms for processing semiconductor and magnetic materials are normally incompatible. Semiconductor materials require high temperatures to fabricate parts, but magnetic materials lose their magnetic properties if heated above approximately 300°C. The team is experimenting with low-temperature laser annealing, a technique developed at Livermore. Because the thermal penetration depth of the laser annealing process is relatively shallow, a diode switch can apparently be fabricated on a multilayer stack of magnetic materials without adversely affecting the storage characteristics of the adjacent memory device. Says Cerjan, "We have been able to heat just one layer of the amorphous silicon so that the metal beneath it is not damaged."

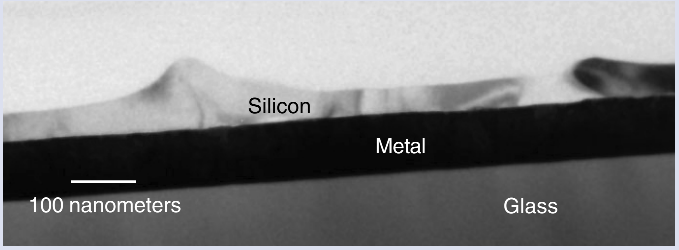

A cross section of polycrystalline silicon grown on a metal substrate by laser annealing, as seen through a transmission electron microscope.

Within the next year, the team plans to integrate the diode switches and relatively large MRAM cells on 4 by 4 arrays of 10-micrometer cells. If they are successful in installing a diode switch in the magnetic tunnel junction cells, they will have produced individually addressable random access memory elements. The next step will be to make these devices even smaller to achieve higher density and hence improved performance. Further development will require the participation of a commercial partner to ensure that the design is practical and that the devices can be readily manufactured.

Spin-Offs from Spin Electronics

Aligning the magnetization of the two layers in magnetic tunnel junction alters the spin polarization of the conducting electrons. The alterations in spin polarization, in turn, affect the overall tunneling probability and hence the magnetoresistive ratio (that is, the ratio of change in electrical resistance when a magnetic field is applied). To date, the team has measured resistance ratios as high as 25 percent, which is large enough to make MRAM elements competitive with semiconductor-based memory.

These developments have prompted Livermore researchers to investigate new classes of so-called spintronic materials, which put both the charge and spin of the conductive electrons to work. A potential application of these new materials would be in secure quantum communication and perhaps, in the future, in quantum computers. Quantum communication is inherently "unbreakable"—any attempt to intercept it destroys the signal. Quantum computers are considerably farther down the road. Replacing the linear computers of today, quantum computers would entangle all functions and solve them at once. Such computers sound like science fiction today but may be reality some day.

—Katie Walter

Key Words: magnetic random access memory (MRAM), magnetic tunnel junctions, Accelerated Strategic Computing Initiative (ASCI).

For further information contact Charles Cerjan (925) 423-8032 (cerjan1@llnl.gov).