the Gulf of California, a 150,000-square-kilometer slice of water wedged between Mexico's west coast and the Baja Peninsula. When the project began in 1999, Sala says, researchers had little information about the distribution and abundance of the gulf's biological wealth, which is under increasing threat. "We had to start from scratch," he says.

To fill the gap, he and two students from Mexico's Autonomous University of South Baja California in La Paz made hundreds of dives at 84 spots along the gulf's coast, surveying sea life and documenting habitat types. They also interviewed local fishers for information about the spawning sites of seven economically important species of fish and looked carefully for nursery areas. Back at Sala's lab, another trio of researchers fed the information into a computer model designed to achieve preset goals.

In this case, Sala's team proposed a network that would protect all coral reefs, sea-grass beds, and known spawning sites, at least 50% of coastal mangroves, and at least 20% of all other habitat types—in a minimal area. To allow sea life to flow from one site to another, they decreed that no reserve should be more than 100 kilometers from the next one in the chain.

With those rules in place, the software—based on code developed by Ian Ball of Australia's Antarctic Division in Kingston, Tasmania, Hugh Possingham of the University of East Queensland in Brisbane, and NCEAS scientists—then spent hours sorting through thousands of possible combinations. The winning map, Sala reported at the meeting, showed that 18 reserves covering just 12% of the marine habitat could do the trick. As a bonus, it protected even more mangroves and other habitats—from sandy bottoms to submerged cliffs—than Sala's rules called for.

Sala's team wasn't finished, however. Knowing that reserve plans can founder on opposition from commercial anglers and other interests, it incorporated data on fishing boat activity collected by the World Wildlife Fund (WWF), one of the project's partners. The software identified several potential conflict zones, then reconfigured the network to avoid heavily fished areas but still satisfy the conservation goals, Sala said.

"It's a really elegant project" that is sure to influence other reserve planning projects, says marine policy expert Liz Lauck of the Wildlife Conservation Society in New York City. Most impressive, says coral specialist Jeremy Jackson of Scripps, is that the job took less than 3 years and cost only $400,000, provided by funders including the Moore Family and Tinker foundations. "It shows how quickly you can gather useful information," he says.

How Sala's findings will play in Mexico, however, remains to be seen. WWF and other groups are working with government officials to develop a long-term conservation plan for the gulf, and Sala's work is just one piece of the puzzle. Still, says Juan Carlos Barrera of WWF-Mexico in Hermosillo, Sonora, "the ability to consider social and economic factors along with ecological concerns is very helpful."

Data dive. Researcher Gustavo Paredes takes a sea-life survey used to design a marine reserve network for the Gulf of California.

Other researchers are pursuing similar work. Leanne Fernandes of Australia's Great Barrier Reef Marine Park Authority reported that her agency has turned to related software to help identify a network that will protect 70 "representative" bioregions along the reef. "The idea isn't to come up with the [ecologically sound] solution and [send] it in to the minister but to have a plan that already takes into account the concerns of the many stakeholders," she says.

Meanwhile, in the Bahamas, a team led by Dan Brumbaugh of the American Museum of Natural History in New York City hopes to build a dynamic model to finger the shifting social and biological forces that determine a reserve network's fate. Backed by a 5-year, $2.5 million grant from the National Science Foundation's Biocomplexity in the Environment program, Brumbaugh has assembled social and biological scientists from nearly a dozen institutions. A key question they hope to answer is whether networks designed to win community support can work as well over time as those focused on ecological goals.

The project demonstrates how marine reserve advocates, traditionally biologists, have begun to incorporate economic and social concerns into their thinking, says Brumbaugh. Successful efforts to design and evaluate reserves depend on "finding people who are willing to play nice with each other and overcome disciplinary suspicions," he adds. And a little silicon-based helper doesn't hurt, either.

–DAVID MALAKOFF

TECHNOLOGY

Microchips That Never Forget

Magnetic memory promises computers that turn on instantly, smarter gadgets, and a revolution in chip design—if it can elbow its way into the market

It's 2 a.m. and you're at your computer typing up that 20-page report that's due in 6 hours. You're on your eighth cup of coffee and your fourth candy bar, and just maybe you'll finish in time to take a shower before dashing off to work. You haven't saved your document in hours. Then you accidentally kick the electrical cord and unplug your computer.

No problem. Plug the cord back into the wall socket, and the machine instantly blinks back to life. It also remembers every t you've crossed and i you've dotted, so you continue to type as if nothing has happened.

This fanciful scenario could become a reality in the not-too-distant future, thanks to magnetic memory that can store information even when it loses power. The first commercial prototype chips should hit the market within 2 years. But magnetoresistive random access memory (MRAM) may not only allow your computer to turn on and off instantly without forgetting what it was doing; it could also reshape all of electronics.

The emerging MRAM combines the best features of the currently existing electronic memory technologies, says Saied Tehrani, an electrical engineer at Motorola in Tempe, Arizona. It therefore could potentially replace all of them. "MRAM really has the potential to be a universal memory," Tehrani says. MRAM bits can also mix with the transistors in standard silicon chips, so the technology could allow chip designers to put an entire computer on a single chip, making portable devices such as cell phones and personal data assistants far more powerful.

But it isn't certain that MRAM can topple the reigning champion of computer memory, an electronic technology called dynamic random access memory (DRAM), says Bob Buhrman, a physicist at Cornell

CREDIT: E. SALA

University in Ithaca, New York. "I think that if DRAM and MRAM were starting at the same point, MRAM would win," Buhrman says. But DRAM has a huge head start, he says, and it isn't clear whether MRAM can catch up and compete economically. DRAM may have drawbacks, but chip manufacturers can pump it out for a few tenths of a cent per megabit, and production is a multibillion-dollar industry.

Researchers, investors, and consumers may soon find out if MRAM can live up to its promise, as major electronics companies are nudging the fledgling technology out of the lab and into production. Motorola plans to introduce a 4-megabit chip by the end of next year, and IBM will introduce a 256-megabit chip in 2004. Meanwhile Honeywell, Hewlett-Packard, and several other companies have their own MRAM programs. "I wouldn't be too surprised if someone gets something out a little quicker" than IBM and Motorola, says Jim Daughton, an electrical engineer at NVE Corp. in Eden Prairie, Minnesota.

A Bit of This and That

Today's computers rely mainly on **dynamic random access memory** (DRAM), which stores bits of information by charging or discharging myriad tiny capacitors. DRAM sets the standard for packing the most information into the least amount of space, with the smallest bits currently measuring just over a tenth of a square micrometer in area. But the capacitors leak charge, so DRAM must be refreshed thousands of times a second.

For some tasks, computers use faster and more expensive **static random access memory** (SRAM). Each bit of SRAM consists of a small network of transistors that flips between two stable conditions. SRAM doesn't require constant refreshing, but it takes up a lot more space. And like DRAM, SRAM forgets everything as soon as it's turned off.

On the other hand, **Flash memory** can hold information without power. A Flash bit consists of a pair of transistors separated by a tunneling barrier that can be made more or less permeable by briefly applying an electric field. Flash wears out after it's been rewritten roughly a million times, so it is used mainly in cell phones and other portable devices, although computers use it to store some internal settings.

An emerging electronic technology called **ferroelectric random access memory** (FeRAM) works much like DRAM but uses a capacitor that can be set to a high or low value. FeRAM retains information without power and could compete directly with magnetoresistive RAM. However, it may be relatively expensive, difficult to manufacture, and hard to integrate with standard chip technology.

–A.C.

High-tech layer cake

RAM serves as a scratch pad on which a number-crunching chip, or processor, keeps data and instructions, including its own operating system. Such memory is called "random access" because the processor can reach into any part of it at any time. Most current RAM can't hold information without power, which is why your computer writes everything to the hard drive before it shuts down. The machine must retrieve that information when it turns back on, which is why it comes on slowly.

Researchers have been trying for decades to produce magnetic random access memory devices that avoid those hassles. In the 1980s, Honeywell developed a technology that exploited the fact that electricity flows most easily through a magnet if the current flows in the direction of the magnetism. But that effect, known as anisotropic magnetoresistance, is small—a few percent—so the memory proved slow, says Daughton, who headed the research. Then in 1988, physicists discovered a bigger effect called gigantic magnetoresistance (GMR) in films in which a layer of nonmagnetic metal lies between two magnetic layers. The resistance is low when the outer layers are magnetized in the same direction but increases dramatically if the layers are magnetized in opposite directions. However, GMR memory devices tend to be bulky or hard to read because the total resistance of a bit of film is low.

The new MRAM devices largely grew out of a 5-year effort funded by the Defense Advanced Research Projects Agency (DARPA). In 1996, DARPA began supporting IBM, Motorola, Honeywell, and academic researchers to develop novel magnetic materials and devices. IBM researchers realized that a tiny high-tech layer cake known as a "magnetic tunnel junction" would provide resistance changes up to 30%, even beefier than the GMR effect. A memory device in which every bit consisted of a junction would be faster and more compact than a GMR memory, they reasoned.

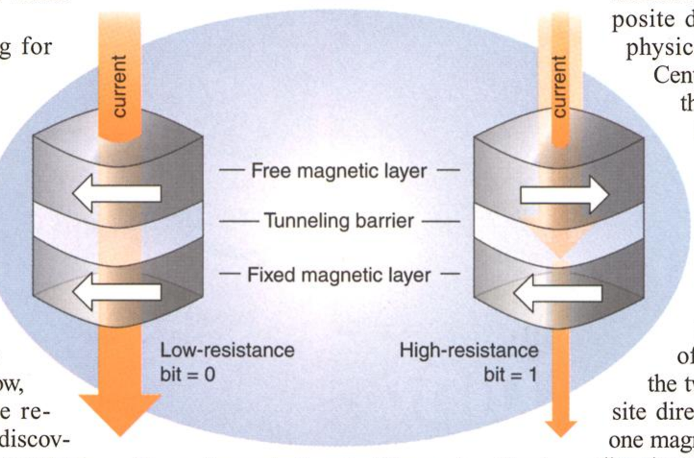

The going gets tough. When a junction's magnetizations cross, less current can tunnel through the barrier.

Motorola also switched to tunnel junctions. (Honeywell, the military, and others continue to develop GMR-based memory, in part because it may better withstand radiation and other hazards.)

The heart of a magnetic tunnel junction consists of two layers of magnetic material, such as nickel iron and cobalt iron, that sandwich a very thin insulator, typically a layer of aluminum oxide only a few atoms thick. Current flows down through the layers, and thanks to quantum mechanics, it meets less resistance when the two magnetic layers are magnetized in the same direction, and more resistance when they're magnetized in opposite directions. The two states serve to encode a 0 or a 1.

The difference in resistance arises because a magnetized material essentially carries two unequal but independent currents of electrons, with their spins polarized in opposite directions, says Stuart Parkin, a physicist at IBM's Almaden Research Center in San Jose, California. When the magnetic layers are magnetized in the same direction, the larger current on one side of the insulating layer is polarized in the same direction as the larger current on the other side. That alignment allows electrons in the larger current to quantum-mechanically tunnel into and out of the barrier with relative ease. If the two layers are magnetized in opposite directions, then the larger current in one magnetic layer is polarized in the same direction as the smaller current in the other, and that smaller current simply cannot accommodate all the electrons burrowing through the insulating barrier (see figure). "The majority electrons can tunnel into the barrier from one side very easily," Parkin says, "but they can't get out the other side."

Crucially, the magnetization of a junction's top layer does not change direction too easily. When it points one way, it resists magnetic fields trying to flip it the other way, much as a stiff light switch resists the push of a finger. Only when the fields exceed a certain threshold does the magnetization realign. That's why the memory device retains information when it's turned off.

To fashion a memory device, millions of junctions are arranged in a square grid with the lower magnetic layers of all the junctions magnetized in the same fixed direction. Parallel conducting stripes called "bit lines" run on top of the junctions, and below the junctions run stripes called "word lines." Each junction sits at the intersection of a bit line and word line. When currents run through the two lines, they create magnetic fields that add together to flip the magnetism of the junction's upper layer in the desired direction (see figure). To measure the junction's resistance, a smaller current passes from the bit line, through the junction, and out through a transistor.

Tunnel junction MRAM boasts a combination of properties that should enable it to take on all of the leading electronic technologies (see sidebar, p. 247). For example, MRAM's ability to retain information without power gives it an edge on the two types of memory used most in computers—DRAM and static random access memory (SRAM)—both of which forget everything as soon as the lights go out. MRAM should also be able to keep pace with the faster SRAM and pack information nearly as densely as the more compact DRAM. On the other hand, MRAM should be much faster and more durable than Flash, a type of electronic memory that can hold its state without power and is used primarily in cell phones and portable devices as well as computers.

MRAM also possesses several novel properties that give it even greater potential, says Jimmy Zhu, an electrical engineer at Carnegie Mellon University in Pittsburgh, Pennsylvania. For example, MRAM bits can be etched into standard chips. So not only should MRAM allow designers to do away with hard drives, it also should enable them to put a processor and its memory on the same slab of silicon. "You can put a whole computer on a single chip," Zhu says, and that could mean putting a fully functioning computer inside your cell phone or personal data assistant.

MRAM bits might even mix with the transistors of traditional electronics to produce chips that reconfigure themselves as they run, says Bill Black, an electrical engineer at Xilinx Inc. in Ames, Iowa. For example, a tiny circuit that multiplies two numbers at one moment might change on the fly into a circuit that divides two numbers. Such morphing chips could radically alter chip design and even the relation between hardware and software.

Hurdles ahead

Before tunnel junction MRAM can live up to that promise, manufacturers must show that they can meet several technical challenges while cranking out loads of chips with few failures and at low cost, as both Motorola and IBM are now trying to do. For example, the high and low resistance values must be nearly the same from junction to junction. But those resistances vary exponentially with the thickness of the tunneling barrier, so chipmakers have to ensure that this exquisitely thin layer is smooth and uniform across the entire chip. Techniques developed in the last decade make this possible, but researchers must show that those techniques work reliably at high volumes.

Researchers must also ensure that they can flip precisely the bits they intend. When current runs through a bit line and a word line, only the bit at the intersection of the two—where their magnetic fields combine—must flip. But all the bits along the bit line feel a single, weaker magnetic field, as do all those along the word line. To prevent these bits from inadvertently flipping, researchers must exercise fine control over the strength of the magnetic fields and the size and shape of the bits. Moreover, as the junctions are made smaller, researchers must make sure that heat does not make them flip spontaneously.

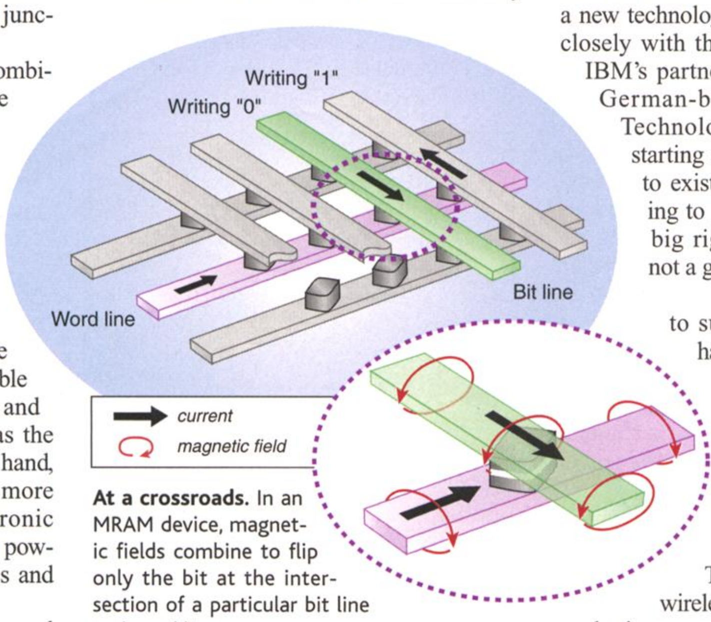

At a crossroads. In an MRAM device, magnetic fields combine to flip only the bit at the intersection of a particular bit line and word line.

CREDIT: ADAPTED FROM IBM/ALMADEN RESEARCH CENTER

Finally, if MRAM is going to succeed economically, production methods must allow it to follow the decades-long trend in which the size of transistors and other features on microchips shrinks by half every 18 months. Chip manufacturers need to see that MRAM has the potential to "scale" through several size reductions before they will be willing to invest in it, says Stuart Wolf, a physicist and manager of DARPA's magnetic materials and devices program in Arlington, Virginia. Wolf is confident that the technological problems can be solved, however. "The advances will come and it will scale," he says. "We just don't have all the answers right now."

More daunting may be the economic and even cultural barriers to working MRAM into commercially viable devices, says Xilinx's Black. "There's a little bit of a disconnect between the ultimate users of the products and the guys doing the research," Black says, "and I think that has served to slow progress."

However, Bill Gallagher, a physicist at IBM's Thomas J. Watson Research Center in Yorktown Heights, New York, and manager of the company's MRAM project, says his team is aware of the challenges of bringing a new technology to market and is working closely with the manufacturing experts at IBM's partner in MRAM development, German-based chipmaker Infineon Technologies. Gallagher envisions starting modestly by adding MRAM to existing devices. "We're not going to go in and replace something big right away," he says. "That's not a good strategy."

Motorola's Tehrani says that to succeed, manufacturers will have to find markets in which MRAM offers a performance advantage that justifies the extra expense of the chips, which at first are likely to cost more than the several cents per megabit that SRAM and Flash fetch. Those markets might include wireless communications, portable devices, and automotive applications, Tehrani says. "As we go through the learning curve in these markets," he says, "we'll see MRAM moving into other markets."

All agree that MRAM can take on DRAM only after it has established itself in other applications. For the moment, DRAM sits secure in its perch as the king of computer memory.

If MRAM is going to make it in the market, it must soon begin to pay its own way. DARPA support for the IBM and Motorola efforts will end this year, and chip manufacturers are likely to pay for more development only if the technology promises to turn a profit in the foreseeable future. "It's like your child," Gallagher says. "If it's going to make it, it's got to learn to walk on its own two feet." The bottom line is developers must soon turn magnetic bits into megabucks, and that may be as much an economic challenge as it is a scientific one.

–ADRIAN CHO

Adrian Cho is a freelance writer in Boone, North Carolina.