Crisscrossing assemblies of defect-prone nanowires

CROSSBAR NANOCOMPUTERS

could succeed today's silicon-based circuits

By Philip J. Kuekes,
Gregory S. Snider and
R. Stanley Williams

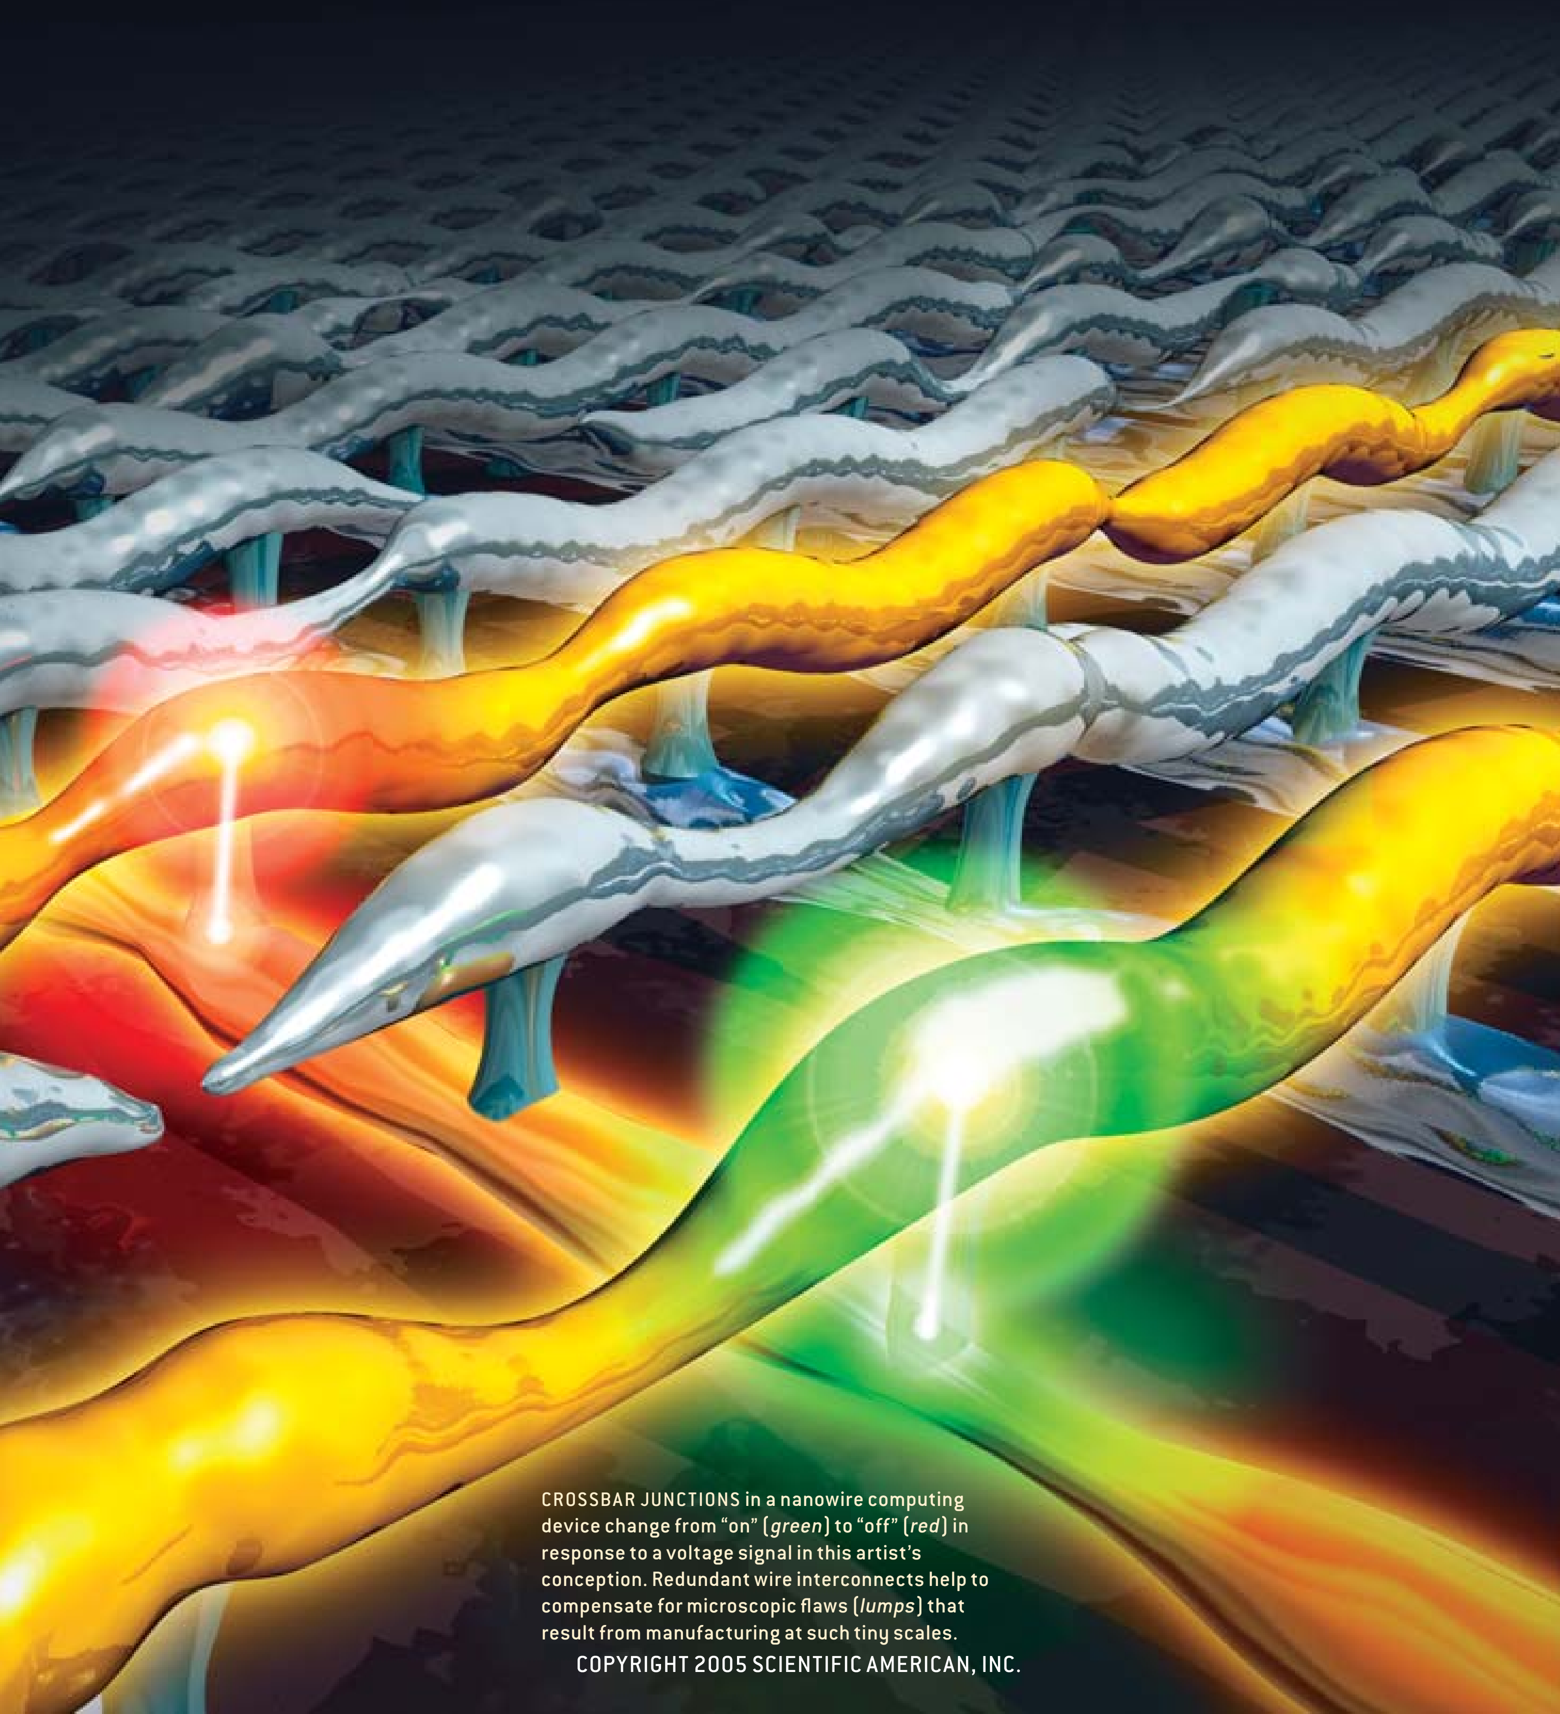

CROSSBAR JUNCTIONS in a nanowire computing device change from "on" (*green*) to "off" (*red*) in response to a voltage signal in this artist's conception. Redundant wire interconnects help to compensate for microscopic flaws (*lumps*) that result from manufacturing at such tiny scales.

IN A LITTLE OVER HALF A CENTURY, THE NUMBER OF TRANSISTORS ON A SILICON CHIP has grown from just one to nearly a billion—an accomplishment celebrated as Moore's Law. By greatly enhancing digital machines' ability to crunch numbers, execute logical operations and store data, this unprecedented manufacturing success has enabled revolutionary changes in our day-to-day lives while spawning one of the planet's largest and most influential industries.

As more and more transistors are packed onto silicon integrated circuits (ICs) during the next decade and a half, the lengths of the smallest chip features will shrink to nearly the molecular scale. Even the most optimistic proponents of ICs believe that major innovations will be required to reach the ultimate operating limit of the silicon transistor: a length for functional features around 10 nanometers (nm), or about 30 atoms long. Finding alternative technologies that can further shrink computing devices is crucial to maintaining technological progress. But because of the silicon IC's amazingly successful track record, the performance bar for any successor is so high it will take at least a decade to develop candidates that will be available when they are needed.

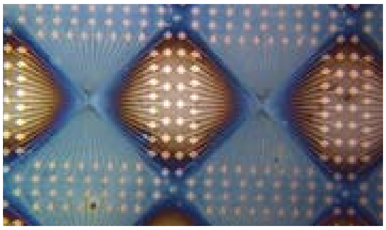

CROSSBAR memories with their test pads.

> Random fluctuations make it impossible to build a perfect machine from nanocomponents.

Researchers worldwide are exploring several exciting alternatives. Quantum computing, for instance, is a novel technique that takes advantage of "spooky" quantum-mechanical properties to process information. Quantum computing is, however, decades away from realization, and even then it remains unclear how useful it would be for most applications. Many research groups are therefore searching for a midterm alternative that could be ready for commercialization in about 10 years. To be economically viable, such a technology must share a great deal with the existing IC processor infrastructure, including critical items like fabrication foundries and software platforms.

Our research team at Hewlett-Packard (HP) Laboratories views the crossbar architecture as the most likely path forward. A crossbar consists of one set of parallel nanowires (less than 100 atoms wide) that cross over a second set. A material that can be stimulated electrically to conduct either more electricity or less is sandwiched between the two sets of wires. The resulting interwire junctions form a switch at each intersection between crossing wires that can hold its "on" or "off" status over time.

Crossbars offer several benefits: The regular pattern of crisscrossing nanowires makes manufacturing relatively straightforward, especially compared with the complex structures of microprocessors. Its arraylike composition provides clear ways to instill defect tolerance in circuits. The structure can be built using a wide range of substances and processes, which provides tremendous flexibility in adapting existing designs to new materials. Finally, this single geometry can provide memory, logic and interconnection, making it very adaptable.

Overview/*Nanoelectronics*

- Moving beyond today's silicon integrated chip technology will require shrinking logic and memory circuits to the scale of a few nanometers. Large arrays of intersecting nanowires called crossbars provide the basis for one of the best candidate technologies for nanocomputing success.
- The nanowires that comprise crossbars are so small that atomic defects and flaws in their manufacture are unavoidable and serious. Building redundancy into the circuitry and using coding theory techniques compensate for the many imperfections.

(Criss)crossing Over

OUR TEAM'S JOURNEY toward this avenue of research began in 1995, when one of us (Williams) moved to HP from the chemistry department at the University of California, Los Angeles. Though not a computer expert, he did know a few things about electronics: one, that a computer's circuits had to be perfect to operate correctly and, two, that random atomic fluctuations at room temperature and above (caused by entropy) would make it impossible to build a perfect machine from billions of components, each composed of only a few atoms. Even

SLIM FILMS (*preceding pages*); G.-Y. JUNG AND B. BAUSK *HP Labs* (*this page*)

ON AND OFF AT THE CROSSROADS

The key component of the crossbar architecture is a nanoscale switch that can be turned "on" or "off" by applying an appropriate voltage to the wires it connects.

In Hewlett-Packard (HP) Laboratories's version, the switch is formed at the junction between two crossing nanowires that are separated by a single monolayer of molecules. The switch starts out in a high-resistance state, blocking the flow of electrons between its two nanowires ("off," *red highlights, below*). But when a large enough voltage of the appropriate polarity is placed across it (*indicated by the yellow and orange nanowires*), the switch changes abruptly to a much lower resistance state, allowing electrons to flow more easily ("on," *green highlights*). The switch stays in this low-resistance state until a large enough negative voltage makes it revert to its original state. As long as the voltage is maintained between these positive and negative thresholds, the switch remains in the state in which it was last set. Some switches the authors have examined have retained their set states for more than three years so far. If the switches can be toggled back and forth many times, they are reconfigurable and can be used in a random-access memory or a reprogrammable logic circuit.

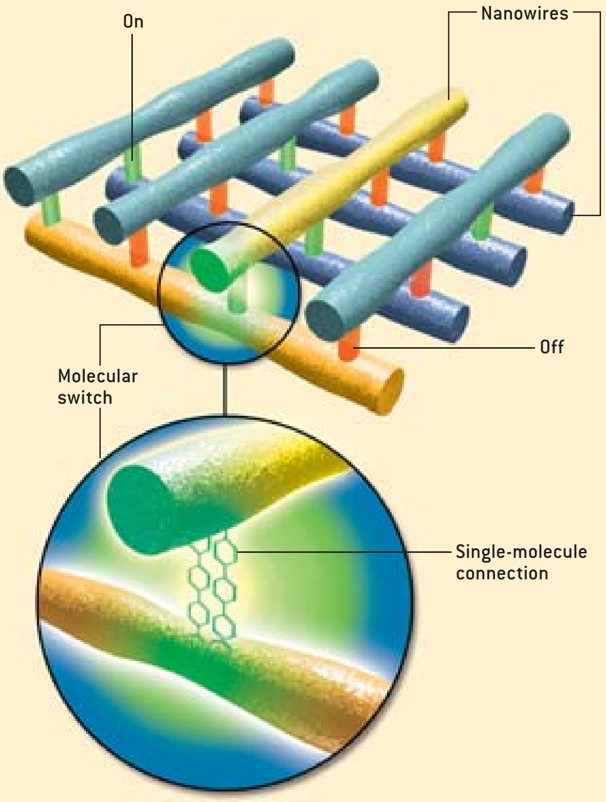

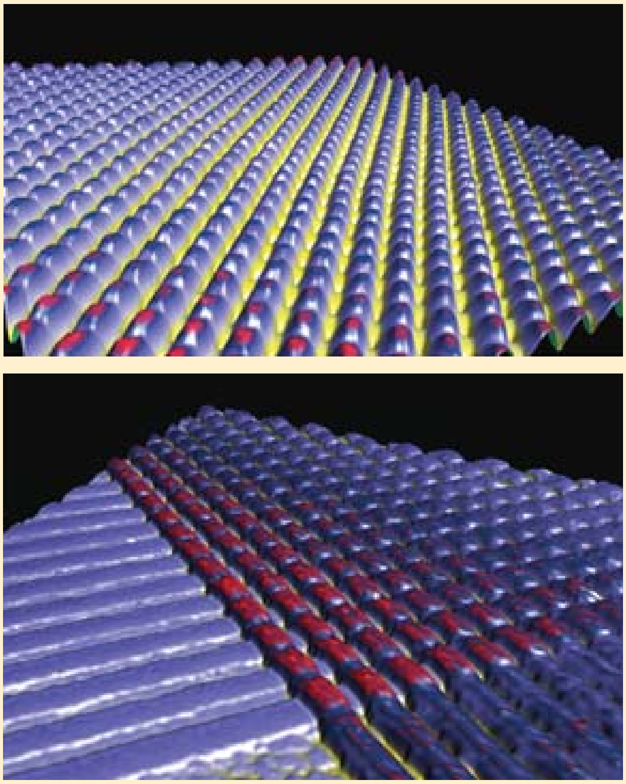

PROTOTYPE ARRAY for a crossbar computing device, depicted in an atomic-force micrograph (*top*), has 34 nanowires (each 30 nanometers wide) intersecting with 34 others. The detail (*bottom*) shows how one set of nanowires crosses over the other set. A junction of two nanowires is smaller than a typical virus.

JANA BRENNING (*illustrations*); COURTESY OF W. WU AND G. MEDEIROS-RIBEIRO *HP Labs* (*micrographs*)

atomic-scale irregularities impose significant variations on the size of nanodevices, which can destroy their electrical properties. Consequently, some sizable fraction of the tiny devices will not work. It was natural for Williams to conclude that nanoelectronics were therefore impossible and that his research at HP should focus on other technologies.

A chance meeting the following year with an HP computer architect (Kuekes) changed that perception dramatically and set the pair on an unexpected path. Kuekes told Williams about a supercomputer called Teramac that he and others (including Snider) had built. Teramac operated perfectly, even though about 220,000 of its components (approximately 3 percent of the total) were defective. The trick, Kuekes said, was that the supercomputer's design had significant redundancy in its interconnect circuitry. After all the flaws were located and catalogued, programs run on the computer were compiled to avoid the broken parts, essentially by routing around the defects via the extra connections.

Williams saw at once that Teramac's tolerance of defects provided a way to construct computers that operate perfectly despite huge numbers of "broken" nanoscale parts. That summer Williams and visiting U.C.L.A. chemist James R. Heath worked on applying the concepts of nanoparticle assembly (assembling complex structures out of tiny building blocks) to computers. After much discussion with Kuekes and Snider about the defect tolerance of chemically assembled computing

BUILD TOP-DOWN OR BOTTOM-UP?

The field of nanoscale fabrication is extremely active today, with many competing techniques under study. These approaches can be classified into two categories: top-down and bottom-up (*below*). The former examples resemble conventional IC manufacturing methods that use photolithography followed by chemical etching or deposition of materials to create the desired features. The latter approaches are based on extensions of chemical or biochemical processes by which atoms or molecules self-assemble into a desired configuration because of their planned, inherent properties. Most investigators in this field agree that some combination of the two approaches will be required to build future nanoscale circuits.

At HP, our team uses imprint lithography to create the crossbars. We and our collaborators employ electron beam lithography to construct molds for the circuits. Although this process is slow and costly, we can make duplicates of the final product, which then are used to stamp out large quantities of circuits, much as vinyl LP records were made. A thin layer of a polymer or polymer precursor coats a substrate, the mold is pressed into this soft layer, and the impressed pattern hardens under exposure to heat or ultraviolet light. The advantage of this approach is that electron-beam lithography can fabricate arbitrary wire geometries on the mold. The drawback is that the present resolution of the features in a set of parallel wires is limited to roughly 30-nanometer half-pitch (half the distance between the centers of two wires, a standard industry measure), although we are working on a number of techniques to improve on this performance. *—P.J.K., G.S.S. and R.S.W.*

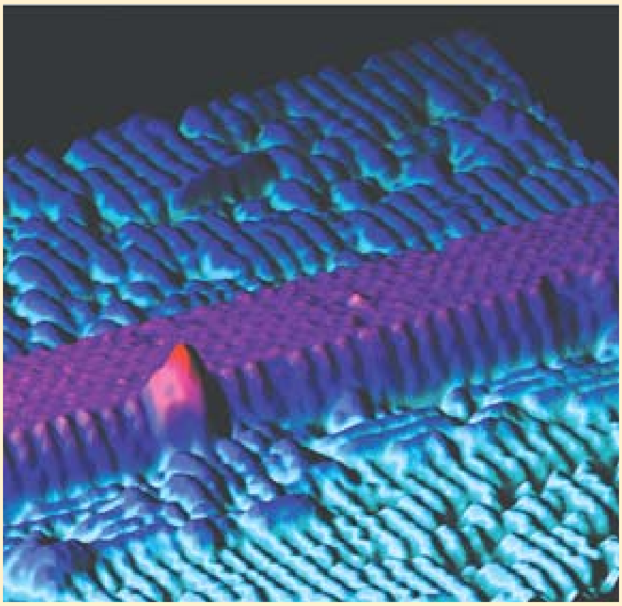

ATOMIC FLAWS AND DEFECTS appear in this scanning tunneling microscope image of a nanowire of erbium silicide that was grown on a silicon surface using a chemical (bottom-up) method. Bumps on the surface of the wire, which is approximately three nanometers (or 10 atoms) wide, are individual atoms. The bulge on the side of the nanowire is a defect where the width changes from 10 atoms wide to only nine.

D.A.A. OHLBERG AND G. MEDEIROS-RIBEIRO *HP Labs*

systems, Williams and Heath wrote a paper about the topic as an educational exercise. To the surprise of all involved, it was taken seriously and eventually published in *Science* in 1998.

Rapid Results Required

THAT SAME YEAR Bruce E. Gnade and William L. Warren, then program directors at the Defense Advanced Research Projects Agency (DARPA), recognized that effective architecture was critical for developing the new nanoscale device technologies the agency was supporting. At the time, interest in molecular electronics was enjoying a resurgence, years after it had first been proposed in 1974 by Avi Aviram of IBM and Mark A. Ratner of Northwestern University. It was not, however, until the early 1990s that Mark A. Reed of Yale University and James M. Tour of Rice University actually started measuring the electrical properties and synthesizing new molecules for electronics. Gnade and Warren understood that electronic devices without an architecture to link them into a useful circuit were mere intellectual curiosities. Their challenge to the research community to define a workable architecture for molecular devices kick-started the research efforts of many groups and encouraged the formation of several significant collaborations.

Our HP/U.C.L.A. team immediately embraced that challenge, but we faced a dilemma. The Teramac-inspired architecture that we had proposed would have required five years to develop, but DARPA wanted tangible results (as a 16-bit memory device) in only two. Heath, Kuekes and Williams brainstormed during the next several weeks to come up with a concept that could meet the deadline. Kuekes and Williams were aware of HP's magnetic random-access memory project and understood that the simple crossbar structure on which it was based was the ultimate abstraction of the Teramac configuration.

Heath pointed out that a crossbar "looked like a crystal" and that it should therefore be possible to build such a system chemically. What was needed was some way to connect each pair of intersecting wires in the crossbar with a switch that could be turned on and off at will. Williams suggested that an electrochemically active material sandwiched between the wires should make it possible to change the electrical resistance of the contacts substantially and reversibly, by applying the appropriate voltages across the two nanowires. That is, the switch would be closed by electrochemically shrinking the quantum-mechanical "tunneling" gap that the electrons have to jump across to get from one electrode to the other. Apply-

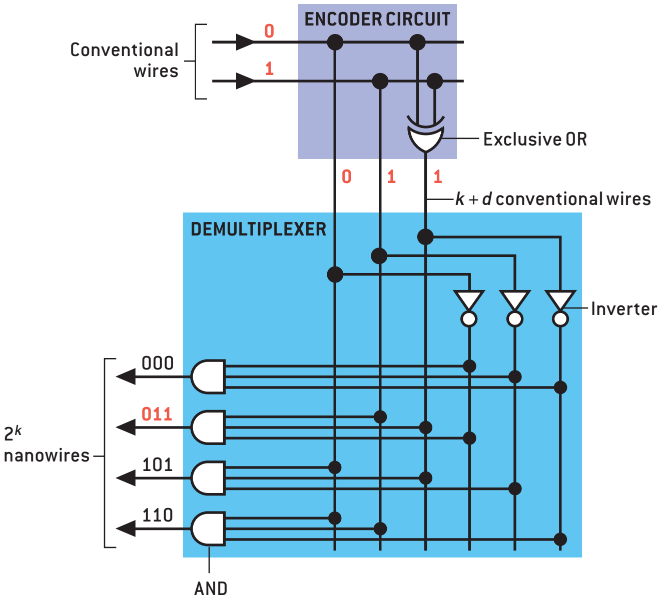

DEMULTIPLEXER enables conventional wires on silicon chips to control a far greater number of nanowires. If k is the number of conventional wires, the multiplexer can control 2^k nanowires. An additional d conventional wires provide sufficient redundancy for the control to work despite broken connections between nanowires and conventional wires. In this simplified diagram, $k = 2$ and $d = 1$; two micron-scale wires control four nanowires with one bit of redundancy . In this example, the conventional wires input the switch address 01 (*red*), to which the encoder circuit adds a redundant bit, yielding the coded address 011. The coded address then activates the nanowire in the demultiplexer designated 011 [*see next page for further explanation*].

ing the opposite voltage bias to widen the tunneling gap and raise the electrical resistance would reopen the switch.

Heath provided the material we needed. He introduced our collaboration to molecules that had been designed by J. Fraser Stoddart, then a new U.C.L.A. faculty member, to operate as electrochemically actuated mechanical switches. The concept was that anything that will change shape between two wires should also affect the ability of electrons to tunnel from one wire to another. A key step was persuading a very busy Stoddart to chemically modify his molecules, which he had christened rotaxanes, to make them oily. This alteration enabled Heath to place a small drop of rotaxanes on a water surface so they would spread out to form a film one molecule thick, which was transferred onto a substrate (a process called the Langmuir-Blodgett technique) on which the bottom set of wires had been formed. After that, we deposited the top set of wires by evaporating metal through a mask, which completed the circuit. These early experiments led to several U.S. patent applications, a proposal to DARPA and another paper in *Science*.

Making the Cut

DESPITE CONSIDERABLE SKEPTICISM on the part of the research community, our crossbar and electrochemical switch concept was accepted by DARPA for its two-year trial, along with several others. Early in this effort the Heath and Stoddart research groups demonstrated that rotaxane molecules sandwiched between electrodes could indeed toggle between high- and low-resistance states. We and others, including Charles Lieber's group at Harvard University as well as the Reed and Tour groups, have since seen a range of different nanoscale switching mechanisms. The diverse observations and approaches generated some confusion within the broader research community, and the various switching phenomena have yet to be sorted out, but the existence of electrical switching is today widely recognized. Dozens of research teams across the globe are now working to develop robust nanoscale electrical switches based on atoms or molecules [*see box on page 80*].

Using the crossbar structure, our U.C.L.A. partners became the first group to demonstrate a working 16-bit memory for the DARPA program in 2000. Their success encouraged the agency to fund a successor program with a far more ambitious goal: the fabrication of a 16-kilobit memory with a density of 100 billion bits per square centimeter. This objective sets the bar extremely high, because it requires fabrication capabilities that are not expected to become available to the semiconductor industry until around 2018.

Our group at HP continues to invent new types of circuits based on the crossbar—notably, defect-tolerant memories and different families of logic circuits. Interesting modifications of the original architectural concept also have been developed by André DeHon of the California Institute of Technology, who collaborates with the Lieber group, and Konstantin K. Likharev of Stony Brook University. Although the crossbar and switch architecture started off as the dark-horse candidate in the DARPA challenge, it has now been adopted and adapted by many research groups worldwide, including those of Masakazu Aono of the National Institute for Materials Science in Japan and Rainer Waser of the Research Center Jülich in Germany.

To understand the crossbar approach, we must discuss the nature of the switch and crossbar structure, the fabrication of crossbars from nanoscale wires [*see box on opposite page*] and the possibility of building reliable circuits from unreliable components.

THE AUTHORS

PHILIP J. KUEKES, GREGORY S. SNIDER and *R. STANLEY WILLIAMS* develop next-generation computing technologies at the Quantum Science Research (QSR) program at Hewlett-Packard Laboratories in Palo Alto, Calif. Kuekes devises novel ideas in the areas of computation, electronic circuits and devices, and quantum information research. The chief architect of the QSR program has designed and built leading-edge computers for more than 30 years. Snider, currently a consultant with HP, is exploring ways to improve the architectural design of nanoelectronics. He has worked previously on logic circuit design, compilers, operating systems, logic synthesis, digital signal processing, computer security and networking systems. Williams, as HP Senior Fellow and director of HP's QSR program, guides the multidisciplinary team that designs, builds and tests new nanocircuits. In the past Williams focused on solid-state chemistry and physics, but his primary interest now is the study of the intersection of nanoscience and information technology.

SLIM FILMS; SOURCE: WARREN ROBINETT

From Micro to Nano and Back

THE PHILOSOPHY BEHIND the nanoscale crossbar is that we must learn to live with its unavoidable imperfections and work around them. The "find and avoid" strategy of Teramac will work as long as it is possible to communicate with the nanowires. This, however, poses another question: How does one bridge the gaps in size and number of wires between nanoelectronics and the conventional-scale silicon ICs that will be required to control the crossbars? If only one-to-one connections can be made, nanoscale crossbars would afford no net advantage. We can solve that problem by making the electrical connections through a demultiplexer, a subcircuit that takes as an input a binary number (such as 1010) and selects a single nanowire that has that binary string as a unique identifier [*see illustration on preceding page*]. In our case, the demultiplexer is a special type of crossbar in which many nanowires connect to a small number of conventional wires.

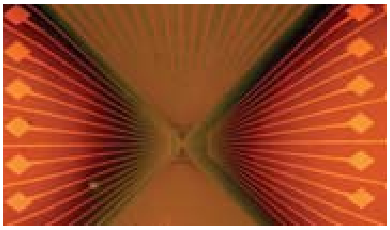

FAN-OUT WIRES from crossbar to test pads.

We managed to create logic functions without the use of transistors in a crossbar circuit.

The number of wires required to input a binary address is the same as the length of the digital names, but the quantity of nanowires that can be addressed equals the number of unique addresses. For example, a string four bits long (0000, 0001, 0010 and so on) can specify 16 unique addresses. Therefore, four micron-scale wires can control 16 nanowires. This fact is important because, to make building the nanoscale circuits worthwhile, one needs to be able to control a lot of nanocircuitry with little conventional electronics. In general, if *k* is the number of conventional wires feeding into the demultiplexer, it can individually control 2^k nanowires, which is very favorable exponential scaling.

A big problem occurs, however, if one of the connections between a nanowire in the demultiplexer and a conventional wire is broken. Then it is no longer possible to distinguish among the *k* different nanowires that share that defective bit in its address. (For instance, if the last bit in the string is broken, then 0000 and 0001 seem identical, as do 1110 and 1111 and other pairs.) Hence, one bad connection in the demultiplexer leads to the loss of all the nanocircuitry that is downstream from *k* nanowires—a catastrophic failure. This result appears to require that demultiplexers, which are half nanocircuitry, be perfect—a violation of our guiding principle that nanoelectronics will be defective.

We find the solution to this quandary in the field of coding theory, which engineers apply when transmitting digital information through noisy environments (as in orbital satellite communications). The general idea is: first, break up a message into blocks of binary data—strings of 0s and 1s. Each block is then extended by adding some extra bits to make a larger block, the code. The extra bits are calculated with an algebraic expression utilizing the bits in the original message block as inputs. When this larger message is sent through the air or some other noisy environment, a few of the bits in the coded message may be garbled (some 1s are turned into 0s and vice versa). By running the code backward on the receiving end, however, the original message can be recovered exactly (as long as the number of flipped bits does not overwhelm the code).

With the guidance of Gadiel Seroussi, Ronnie Roth and Warren Robinett of HP, our team has applied this concept to protect our nanowires from broken connections in the demultiplexer. Rather than numbering the nanowires consecutively, we use an extended address space in which the number of wires entering the demultiplexer is larger than the minimum number needed to specify each nanowire uniquely (by an additional *d* wires). In this case, it turns out that each nanowire can have several broken connections to the conventional wires, yet the demultiplexer can still successfully address all the nanowires. The degree of redundancy required depends on the probability of connection defects; a relatively small amount of redundancy (about 40 percent) can improve the effective manufacturing yield of a demultiplexer from 0.0001 to 0.9999 if the defect rate of the connections in the demultiplexer is 0.01.

Making Memories

SINCE THAT FIRST 16-bit memory device, both the Heath group and our team at HP have demonstrated 64-bit memories at 62-nm half-pitch (which is half the distance between the centers of two adjacent wires, a standard semiconductor industry measure) in 2002 and last year a one-kilobit crossbar at 30-nm half-pitch, using different approaches for the wires and switches. (In comparison, the half-pitch of the most advanced semiconductor IC in 2005 is 90 nm.) Each nanowire in these demonstration memories was connected to an individual contact. We wrote a bit as a 1 (low-resistance) or a 0 (high-resistance) simply by applying a bias voltage that exceeded the threshold for directly toggling the desired switch across its two wires. As long as the voltage threshold for recording a 1 or a 0 is relatively sharp and the variation in "write" voltages among the junctions in the array is less than one half the switching voltage, this procedure ensures that only the desired bit in the array is written (and no others are accidentally written or erased). We read the bit stored in the switch by applying a much lower volt-

Y. CHEN AND B. BAUSK *HP Labs*

Groups Researching Crossbar Architectures

GROUP(S)	INSTITUTION(S)	SWITCH
J. R. Heath/ J. F. Stoddart	Caltech/U.C.L.A.	Rotaxane monolayer between silicon and titanium nanowires
C. Lieber/ A. DeHon	Harvard University/Caltech	Silicon nanowire field-effect transistors
M. Aono	National Institute for Materials Science, Japan	Silver sulfide ionic conductor (silver-based atomic switch)
R. Waser	Research Center Jülich, Germany	Defect motion in ferroelectric thin films
K. K. Likharev	Stony Brook University	Molecular single-electron transistor
Quantum Science Research	Hewlett-Packard Laboratories	Metal nanowire oxidation/reduction

age across the selected crossing wires and measuring the resistance at their junction. These initial results proved promising—in HP's 64-bit memory, the resistance ratio between 1 ("on") and 0 ("off") exceeded 100, making the bits very easy to read.

With the goal of nanoscale memory within reach (the DARPA challenge requires a half-pitch of 16 nm), our next big hurdle is to perform universal computation with nanoscale logic circuits. With Duncan R. Stewart at HP, we have configured crossbars to perform simple logic (Boolean AND and OR operations) by setting the resistance values of switches in a crossbar. The range of logic that can be performed, however, is limited without the NOT operation, or signal inversion, which changes a 1 to a 0 and a 0 to a 1. The wired logic functions also necessarily cause the voltage levels to trail off; if one tries to use too many in a series circuit, 1s and 0s would no longer be distinguishable and computation would not be possible.

In silicon ICs, transistors perform both signal restoration and inversion. This fact has motivated the Heath and Lieber groups to fabricate transistors made from silicon nanowires. We and DeHon have described logic circuits with a "tile and mosaic" topology that can be built using transistors and other elements that are fabricated into a crossbar. Because this approach employs current IC technology, however, eventually it suffers from the aforementioned limitations, so it does not offer an extension beyond Moore's Law. As an alternative, we are investigating signal inversion and restoration without transistors.

Our team is building an unusual form of crossbar logic circuit with arrays of switches and wired ANDs and ORs. In this case, the switches perform a latching operation, which we recently demonstrated with Stewart. We define the voltage level needed to turn a switch on as a 1 and that to turn it off as a 0. Any wire connected to the input of a switch will perforce set that switch to the wire's present logical state, thus transferring one bit of information from "logic" to "memory."

Once stored as a memory state, that bit can be employed in further logic operations by connecting the output wire from the switch to a voltage supply (in our case, a wire from the clock that controls the timing of the operations). This new connection can then be used to restore the voltage of the logic state to its desired value when it has degraded. Another trick is to switch the voltages representing a 1 and a 0 on the output wires, which inverts the logic signal. This change supplies the logical NOT operation and, combined with either ANDs or ORs, is sufficient to perform any computation. Hence, we managed to create the signal restoration and inversion functions without the use of transistors or their semiconductor properties in a crossbar logic circuit.

Beyond Silicon ICs

THE PATH TO universal computing beyond transistor integrated circuits is still highly uncertain, but the crossbar architecture has emerged during the past several years as a principal contender for a new computing paradigm. Much remains to be done. Three different areas of research must advance rapidly and together: architecture, device physics and nanomanufacturing. Ensuring good communications across disciplinary boundaries will be as challenging as solving the technical issues. Success will require multiple teams of researchers who are simultaneously competing against and cooperating with one another, such as the participants in the DARPA challenge.

Defect tolerance will be a necessary element of any future strategy for nanoelectronics. The crossbar architecture is ideal for implementing strategies based on finding and avoiding bad components and on coding theory to compensate for mistakes. Future circuits may actually be more robust than current electronics, even though they will start out with a high fraction of defective components. The built-in redundancy will make them resistant to forces (such as radiation exposure) that cause catastrophic failures in conventional circuits and instead enable their performance to degrade gracefully.

The quantum-mechanical nature of tunneling switches is suited for nanoscale circuits. As the feature sizes of devices shrink, the electrons in them behave more like quantum-mechanical bodies. Such switches should be able to scale down to nearly single-atom dimensions—which suggests just how far the future miniaturization of electronic circuitry might someday go.

MORE TO EXPLORE

Configurable Computing. John Villasenor and William H. Mangione-Smith in *Scientific American,* Vol. 276, No. 6, pages 66–71; June 1997.

A Defect-Tolerant Computer Architecture: Opportunities for Nanotechnology. J. R. Heath, P. J. Kuekes, G. S. Snider and R. S. Williams in *Science,* Vol. 280, pages 1716–1721; June 1998.

Computing with Molecules. Mark A. Reed and James M. Tour in *Scientific American,* Vol. 282, No. 6, pages 86–93; June 2000.

Feynman Lectures in Computation. Paperbound edition. Richard P. Feynman. Edited by Tony Hey and Robin W. Allen. Perseus Books Group, 2000.

The International Technology Roadmap for Semiconductors (ITRS) Web site is at **http://public.itrs.net/**