Making Single Electrons Compute

As chipmakers shrink features on chips, electrons begin spilling over the barriers designed to confine them. So now researchers are devising a new generation of transistors that manipulate electrons one at a time

Big may be beautiful. But in the realm of microelectronics, small is where the cash is. Chipmakers have followed the money for decades, shrinking features on chips to make computers ever more powerful. And today, researchers are hard at work on new chip-patterning schemes that will be able to etch features far smaller than those carved by today's lithographic tools (*Science*, 13 December 1996, p. 1834). But if chipmakers succeed, they will be faced with an entirely different challenge: how to redesign microelectronic devices so they will function at ultrasmall dimensions. When key features in conventional transistors are shrunk too much, they can no longer control the stop-start flow of electrons that represents digital 1s and 0s. The electrons spill over barriers that are supposed to confine them, disrupting the chip's functioning.

In the brave new ultrafine world, researchers will need to devise a different way of corralling electrons, says Kosta Likharev, a physicist at the State University of New York, Stony Brook. And for many, the solution will lie in devices that manipulate electrons one at a time. Although still in an embryonic stage, research on these so-called single-electron devices "is picking up a lot of steam," says Haroon Ahmed, a professor of microelectronics at Cambridge University in the United Kingdom. Their tiny size means chipmakers could pack orders of magnitude more processing and memory circuits on computer chips than is currently possible. The devices also would require a fraction of the power consumed by today's transistors, which could extend the lifetime of laptop computer batteries, among myriad other benefits. Indeed, single-electron devices may be "the ultimate point humans can reach" in the drive to make power-saving devices, says JawShen Tsai, a physicist at NEC Fundamental Research Labs in Tsukuba, Japan.

Likharev and other researchers around the globe have been experimenting with single-electron devices since the late 1980s. In principle, these devices aren't that different from a group of conventional devices known as metal-on-oxide field effect transistors, or MOSFETs. In a MOSFET, the center of the action is a semiconducting region called a channel, which either permits or blocks the flow of electrons between two electrically conducting metal pads. These pads, called the source and drain electrodes, sit on either side of the channel and are accompanied by a third electrode, called a gate, located above the channel. The channel is normally in a nonconducting state, so that electrons cannot travel from the source to the drain electrodes. But when a voltage is applied to the gate electrode, the electrical conductivity of the channel rises, allowing electrons to flow across it: The transistor flips from "off" to "on."

As the components shrink, however, problems arise. When the distance between the source and drain electrodes is less than about 10 nanometers—a size well below the 250-nanometer gap in today's standard devices—electrical charges spontaneously leap across the channel—a catastrophic failure that prevents the device from being switched to its "off" state.

Island-hopping. A vanishingly small island placed in the middle of a device's semiconducting channel compels electrons to hopscotch single file between the source and drain electrodes.

Researchers are working to redesign conventional transistor architecture to keep lawless electrons from making the leap by forcing them to move in an orderly, single-file line. To make these so-called single-electron transistors (SETs), the main class of single-electron devices, researchers replace the channel between the source and drain with one or more metal or semiconductor islands—sometimes called "quantum dots"—each surrounded by a moat of insulating, or nonconducting, material. The key is that the islands are vanishingly small: Because electrons naturally repel one another, they resist crowding on the island and are compelled to hopscotch one at a time across the one or more islands to get from the source to the drain.

The moat isolates the island so that electrons can't jump onto it without a small push. They don't get that push directly. Rather, when a voltage is applied to the gate, the conductivity of the island, in effect, is raised, thereby lowering the amount of energy an electron needs to hop on.

Researchers interested in understanding the physics behind single-electron hopping built SETs as early as 1987. These early devices had relatively large islands about 100 nanometers across, because islands of this size were relatively simple to construct. The devices had to be cooled to near absolute zero to force electrons to travel in single file, however. That's because there's plenty of space on such large islands for several electrons to coexist, so it doesn't take much extra energy for source-side electrons to overcome the small repulsion of their brethren on the island and hop on, too. Indeed, at everyday temperatures, additional electrons can pick up this boost just by absorbing a little heat from their surroundings. But by cooling the devices and robbing the electrons of the extra heat, researchers were able to coax electrons to flow on and off the islands one at a time.

Ultracold temperatures, however, aren't practical for making real-world computer chips. So researchers have recently focused their efforts on making smaller islands. This approach began to pay off last year. Over the course of several months, two research teams—one led by physicist Katsumi Murase at NTT's Basic Research Laboratories in Atsugi, Japan, and the other by Steven Chou at the University of Minnesota, Minneapolis—reported creating room-temperature devices with crystalline silicon islands just 5 to 10 nanometers across. Both groups cut islands in the silicon substrate in part by using a beam of electrons. Other teams at Hitachi Central Research Laboratory in Tokyo, as well as researchers at Japan's Electrotechnical Laboratory (ETL) in Tsukuba and at Stanford University in Palo Alto, California, have constructed room-temperature SETs with materials such as polycrystalline silicon and titanium.

While such demonstrations are impressive, says Likharev, researchers must shrink the islands still further to make the devices function better. The NTT and Minnesota teams, for example, still see only a very small repulsive effect between successive electrons, so, at room temperature, several often crowd onto the island at one time. Currently, these and other teams are trying to shrink the size of their islands further, to between 1 to 4 nanometers. To accomplish this, the researchers are trying everything from improving conventional electron-beam lithography, to etching the features with atomic-imaging microscopes and even growing tiny semiconductor or metal nanocrystals and then wiring them up with electrodes.

Shrunken memories

Researchers around the world are pursuing similar strategies in an effort to build the next generation of devices used in a type of computer memory known as random access memory, or RAM. Today, devices that make up RAM come in a variety of forms, but most face problems related to those of conventional MOSFETS as their dimensions continue to shrink. At the heart of these futuristic memory chips are single-electron devices that store a bit of information as the presence or absence of a single electron, or a small group of electrons, on a microscopic island.

Although similar to SETs, single-electron memory devices have one crucial difference: The semiconductor or metal islands do not replace the channel between the source and drain electrodes, as in SETs, but sit between the channel and the gate electrode that lies above it. In this case, the source and drain electrodes are kept far enough apart to prevent electrons from making spontaneous jumps. And electrons forced onto the island then influence the ability of the gate electrode to open and close the channel between the source and drain.

When there are no electrons on the island, only a half a volt or so of current applied to the gate electrode triggers electrons to flow from the source to the drain; that's the uncharged, or 0, state. To create a charged, or 1, state, a larger voltage pulse is applied between the gate and the source and drain, which pushes an electron from the channel through a thin insulating layer and onto the island; a thick insulating layer between the island and the gate electrode ensures that electrons do not hop onto the gate. The island-bound electron's own electric field then counteracts part of the field coming from the gate electrode, thereby reducing or completely shutting off the flow of electrons from the source to the drain. By tracking how much current is moving from the source to the drain for a given voltage applied to the gate, researchers can detect the presence of the lone electron on an island.

In 1993, Kazuo Yano and his colleagues at Hitachi reported the first single-electron memory devices that worked at room temperature. And last February, the Hitachi team reported at the IEEE International Solid State Circuits Conference in San Francisco that they had wired a series of such devices into a simple circuit. Much of the reason for the Hitachi group's success is that it chose to construct the channel in its devices out of a thin film of polycrystalline silicon. As electrons travel through such films, they typically snake along a single path of least resistance along the borders between crystalline grains. This effectively narrows the channel for the electrons, making it easier for electrons on the storage island to halt the flow, much as it's easier to dam a narrow creek than a wide river.

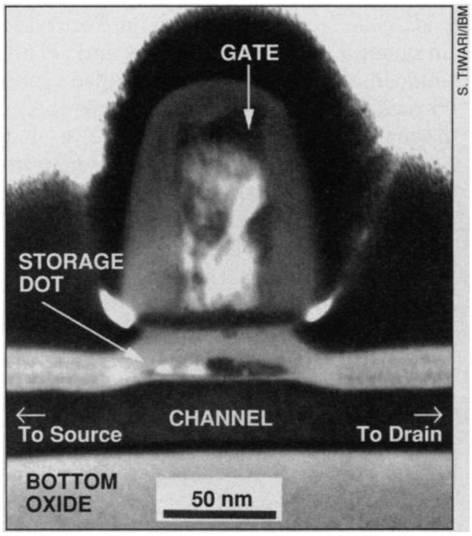

Home alone. New memory devices store a bit of information as the presence or absence of a single electron, or small group of electrons, on a microscopic storage dot.

Although Cambridge's Ahmed calls these demonstrations "very good work," he and others point out that polycrystalline devices are not likely to make it onto real-world, mass-produced chips. The main issue is reliability, says Minnesota's Chou. For chips to function, he explains, each of the millions of devices on them must behave in precisely the same way. Different polycrystalline devices behave differently, however, because the pattern of grains in the thin films is not uniform. That's likely to affect how current moves through the channel, and thus how much gate voltage is needed to open the bridge between the source and drain.

To get current to flow in a uniform manner, other teams have turned to a more orderly substance for their channels and islands—silicon in a continuous crystal lattice. But use of this material poses its own problems. Because in crystalline silicon electrons flow through the whole film rather than along a snaking path, to make working devices the teams have needed to narrow the channels and storage islands to ultrasmall dimensions. This fall, two teams reported making single-electron memory devices using electron-beam lithography to carve channels and islands measuring about 10 to 30 nanometers across. The first of these teams, led by Jeff Wesler and Sandip Tiwari at IBM's T. J. Watson Research Center in Yorktown Heights, New York, reported its results at the October meeting of the American Vacuum Society in Philadelphia, while the other—led by Chou at Minnesota—reported its results last month at the International Electron Devices Meeting in San Francisco.

These devices have functioned quite well in the lab, but, as Tiwari acknowledges, they still face a number of problems in becoming a real-world technology. For one, electron-beam lithography is far too slow to be useful on the factory floor, where lithographic systems must be able to etch millions of devices on chips an hour. Also, the devices are prone to giving false signals because electrons can get trapped in defects in the insulating material near the storage island, making it appear that an electron is sitting on an island when, in fact, none is.

Tiwari and his colleagues are attempting to overcome these and other obstacles by making memory devices with slightly larger features and a patch of several storage islands instead of just one. The multiple islands, notes Tiwari, can block conduction through a wider channel, which allows the researchers to design circuits with coarser features, which can be etched with higher speed techniques such as conventional optical lithography. Multiple islands also reduce false signals, because a lone, stray electron near one island wouldn't be able to halt conduction through the whole channel.

In making their multidot devices, Tiwari and his colleagues pattern the dots with a specialized vapor-deposition technique that allows them to lay down a layer of 3-nanometer crystalline silicon dots—separated by insulating silicon oxide—above a silicon channel. They etch other, larger features with conventional chip-manufacturing techniques, including high-speed optical lithography. The resulting devices work in essentially the same way as memory devices with single islands. Electrons forced onto dots counter the field from the gate electrode and prevent current from flowing from the source to drain. The IBMers published their results on these devices in the 4 March 1996 *Applied Physics Letters*.

These so-called multiple nanocrystal memory devices aren't technically single-electron devices because they store bits of data with a few dozen electrons. Nevertheless, they may prove to be the "middle step," says IBM chemist Christopher Murray, easing the difficult journey into a realm of vanishingly small electronic devices that manipulate charges one at a time.

–Robert F. Service