Superconducting Electronics: Toward an Ultrafast Computer

The ultimate speed of any computer is set by the speed of light, the fastest any electronic signal can travel. Thus, even if a computer is equipped with logic and memory circuits that operated instantaneously, its speed would be reduced to the time it takes a signal to travel from one component to another. Therein lies the rub because today's computers are built around semiconductor electronics, and semiconductors generate heat that must be dissipated before the devices overheat. The problem of heat dissipation puts a limit on how closely semiconductor devices can be spaced and hence also on the maximum speed of operation. Although heat dissipation is not yet a serious problem, some say that the day it will become one is not far off.

For this reason, a few investigators, principally at IBM's Yorktown Heights Research Center, are looking into an altogether new technology, superconducting electronics, as a way to overcome the apparent limitations of semiconductors. Although no computer made of superconducting devices is expected for many years, recent progress is encouraging, and an outline of what such a computer might be like is emerging.

IBM has the biggest superconducting electronics program with more than 50 professional scientists devoting their time to superconductors. Another entrant in the field is Bell Laboratories, which is interested in ultrafast switching machinery for the telephone system. The Bell Labs effort is substantially smaller than that at IBM, however. Other groups in the United States and in Japan are also engaging in small-scale programs.

Superconductors, of course, are those marvelous metals that lose all resistance to the flow of electricity at cryogenic temperatures. (The highest temperature at which superconductivity has been observed in any material is just over 23°K.) Electronic switches, which are the basic ingredients of logic and memory circuits and are made from superconductors, are right now much faster than comparable semiconductor devices, although projected improvements in the latter may well narrow the difference. The real advantage of superconductors is the low rate at which they generate heat. It is this property that one day may allow computer components to be packed so tightly together that full advantage can be taken of their high switching speeds—that is, an ultrafast computer may be possible.

A common figure of merit for switching devices is the product of the time it takes a signal to run through the circuit and the power required to operate it. According to Wilhelm Anacker, manager of the IBM superconducting electronics effort, individual superconducting switches now made there have a speed-power product some 10,000 times more advantageous than that of high-speed silicon devices. (Silicon is used in all semiconductor microelectronic circuits.) The superconductors used at IBM, however, must be cooled in liquid helium to a temperature of 4.2°K. Helium is a poorer conductor of heat than is Freon, a liquid coolant used to keep at least one of the most advanced computers now available from overheating. Allowing for this difference in cooling capacity, says Anacker, still leaves about a 50-fold advantage to the superconductor when many switches are packed closely together.

The basis of all the hope for superconducting electronics is the Josephson junction, named after Brian Josephson of the University of Cambridge, who shared a Nobel Prize for working out the properties of junctions between superconductors. A Josephson junction consists of two superconducting metals (usually thin films) separated in the region where they overlap by a thin layer of a nonsuperconducting material. Often the material is an oxide of one of the superconductors and is a few nanometers thick. Josephson's theory predicted that, below a certain critical value, the electrical current flowing between the two superconductors by way of the oxide layer would itself be a superconducting current with no voltage drop across the junction, whereas above the critical value, it would be a normal current with a voltage appearing across the junction. On top of all this, the critical or transition current could be adjusted downward by the application of a magnetic field.

All of these predictions were confirmed experimentally by Phillip Anderson and John Rowell of Bell Laboratories and by others. Then in the mid-1960's Juri Matisoo of IBM recognized that the transition between superconducting and normal states in a Josephson junction could be the basis of an electrical switch. Matisoo's switching procedure was to cause a transition between the two states by adjusting the critical current with an applied magnetic field to values above and below that of a current flowing through the junction. Test devices worked, and thus the IBM superconducting electronics program was born. Although Rowell had come up with similar ideas even earlier, work on superconducting electronics did not get started there until the early 1970's, according to Theodore Fulton of Bell Labs. Researchers at Bell, however, have followed a slightly different path than that at IBM. Rather than using a magnetic field to bring about switching, they introduce a control current directly into the junction that is high enough to send it into the normal state.

The construction of superconducting logic and memory circuits bears a striking similarity to the way semiconductor microelectronic circuits are made. Several sequences consisting of vapor deposition, wet chemical, and thermal heat treatments are used to build up a multilayer structure. The pattern of superconducting switches and conductors that connect them into a logic or memory circuit is defined by the same photolithographic process as that perfected by semiconductor manufacturers.

According to Charles Kircher of IBM, a finished superconducting switch might consist of some eight layers having thicknesses ranging from a few hundred nanometers to 1 micrometer. The basic elements of the switch, the superconductors, are a lead alloy with about 12 percent indium and 4 percent gold for the lower superconductor and a lead-bismuth alloy for the upper. The thin oxide separating them at the junction is primarily indium oxide. Silicon monoxide and niobium pentoxide are used as insulators to electrically isolate portions of the circuit from each other. A layer of superconducting niobium serves as a ground plane, and the whole assembly sits on a silicon substrate. Over the assembly is a network of "control lines" made of the lead-indium-gold alloy. The control lines generate the magnetic field used to initiate switching; current in the control lines gives rise to the field.

A major step along the way to making logic and memory circuits from superconductors was finding a way to reliably grow the thin oxide layer between the superconducting metals in the Josephson junction. In 1971, IBM's James Greiner found one way to solve this problem by

0036-8075/78/0818-0602$00.50/0

developing a glow discharge process. The thickness was controlled by the pressure of oxygen in the discharge; for each value of the pressure, there was a unique thickness at which the discharge removed oxide as fast as it formed it.

Another advance occurred 4 years later when Dennis Herrell and his colleagues at IBM fabricated a logic device called a 4-bit multiplier. The multiplier circuit contained 45 superconducting switches. The smallest dimension in the plane of any microelectronic circuit is usually the width of the conductors connecting the individual switches. In the multiplier, the conductor width was 25 micrometers. Recently, a similar logic device, a 1-bit adder, having 25 micrometer line widths and nine switches was made by Fulton and his co-workers at Bell Labs.

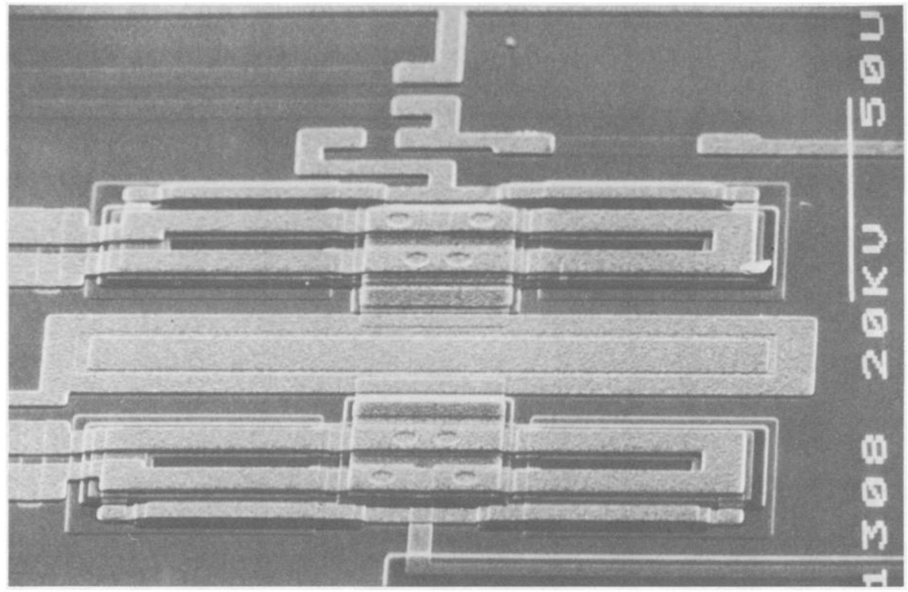

Fig. 1. Scanning electron micrograph of a Josephson junction logic circuit. In logic circuits, individual junctions, which act as electrical switches, are combined to perform logic functions. An OR circuit puts out a current pulse if there are current pulses in one or more of its input lines. An AND circuit, such as shown in the figure, puts out a current pulse only if all its input lines have pulses in them. These and similar logic circuits are then combined to carry out arithmetic and other operations. The Josephson junctions are the eight small (5-micrometer diameter) circular depressions near the center of the micrograph. [Source: R. T. Miller, IBM Corp.]

Earlier this year, Melvin Klein, Herrell, and Arthur Davidson of IBM reported on logic circuits with conductor widths of 5 micrometers, which is comparable to the minimum dimension in semiconductor microelectronic devices (Fig. 1). Although these circuits were less complex than the earlier multiplier circuit, they embodied a modified form of the Josephson junction that will, say the IBM researchers, enable further miniaturization and concomitant increases in speed. Depending on the type of logic circuit, the present devices passed signals ranging from 40 to 60 picoseconds. This is so fast that the switching times could be determined only by averaging over a string of 11 devices in series and is ten times faster than comparable silicon devices with 5 micrometer line widths.

Several logic circuits are needed to carry out a computer operation, such as adding numbers. Thus a few nanoseconds would be taken up for each operation in a computer with such fast switches. One of the things needed to operate at such speeds is a power supply that can switch current on and off at this rate. Recently, Frank Fang of IBM and Herrell developed such a power supply. It makes use of superconducting switches located on the same chip as the memory circuits and can hold the voltage constant to within 1 millivolt.

Logic circuits make up the central processor unit of a computer, where computation is actually carried out. Two other primary components are a fast and a slow memory. The fast memory, also known as a cache memory, contains the information most often needed by the central processor, including the basic instructions by which the computer operates. In a properly designed system, it will have the asked-for data about 90 percent of the time. The slow memory is much larger than the fast, is the "main memory" of the computer, and holds most of the data used in computation. According to Anacker, IBM plans to make superconducting versions of all three components.

The company has not yet announced any results on the fast memory. Researchers at the IBM Zurich Research Laboratory have the responsibility for the slow memory. At the same time that the logic circuit results were reported, Peter Wolf, Theordor Mohr, and their co-workers from Zurich announced the construction of a portion of a memory that eventually will store about 16,000 binary bits of information. An important parameter of computer memories is the access time, the time between the receipt of a request for information and its release. The access time measured for the completed portion of the memory, which contains about 4500 switches, was 7 nanoseconds. The minimum linewidth of the memory circuit, which also uses the new Josephson junction configuration, named a Josephson interferometer by IBM, was 2.5 micrometers.

Unsolved is the major engineering task of deciding how to arrange the thousands of switches on a chip and the conductors connecting them that make up the superconducting circuit and how to arrange the chips themselves in such a way that the high speed can be fully utilized. IBM's Alan Brown and Matisoo described to *Science* one approach being tried for superconducting circuits. The key is silicon, oddly enough; also important is the fact that superconducting wires can carry large currents even when the wires have small cross sections. One reason for the importance of a large current-carrying capability is that the magnetic fields that cause the switching are generated by electrical currents. As the cross section of a nonsuperconducting wire decreases, its resistance increases and a lower current and field result. Since superconductors have no resistance to the current pulses in the circuits, this limitation associated with miniaturization is irrelevant.

Just as each logic and memory circuit containing a thousand or more switches is made on a silicon substrate, each substrate chip in turn is to be mounted on a silicon wafer. Wafers used by the semiconductor industry today are typically 7 to 10 centimeters in diameter and can therefore hold many chips, which are usually rectangular and about 6 millimeters on a side. For the purposes of a superconducting computer, the wafers would be trimmed to squares 5 centimeters on a side. The wafers would then be inserted in a silicon box, called by electronics engineers an egg crate. The squared wafers would be arrayed in parallel planes. Spacing between wafers might be a millimeter or less before heat removal becomes a problem. A high-performance superconducting computer, exclusive of input-output and mass storage devices, might consist of eight egg crates, each 5 centimeters on a side.

All the interconnecting wires between chips on a wafer and between wafers would be made of the same superconductors used for conductors within the circuits on one chip. The choice of silicon as the primary structural material is based on the desire to make the computer from the same material as much as possible in order to minimize stresses and strains induced by the thermal cycling between room temperature and that of liquid helium. Since silicon's properties make it a good substrate material for each circuit, its use for the rest of the box seemed mandated as well.

Packing so many switches (about 10^8) so tightly in a three-dimensional arrangement is not possible with silicon semiconductor electronics, yet such close packing is needed for proposed super-

Rickettsiae: A New Vaccine for Rocky Mountain Spotted Fever

Rocky Mountain spotted fever (RMSF) is not one of the most widespread of diseases. Last year, the United States had 1115 cases, a record number; that is a mere pittance compared to more prevalent diseases, such as hepatitis and influenza, but nonetheless a sharp increase from the 774 cases reported in 1974. The majority of the victims are children, most of them in the so-called "tick-belt" states of Maryland, Virginia, North Carolina, South Carolina, and Georgia. As many as 10 percent of the victims of RMSF die, and in another 20 percent the kidneys, liver, and nervous system may be permanently damaged. The number of cases of RMSF may be markedly reduced, however, by use of a new vaccine that promises much better protection than has previously been available.

The disease, which is characterized by chills, headaches, high fever, and, at the outset, a rash on wrists and ankles, is caused by the tick-borne parasite *Rickettsia rickettsii*. Rickettsiae are small microorganisms that resemble viruses in that they can reproduce only in living cells of other organisms. The first vaccine against RMSF was prepared in 1924 by Roscoe R. Parker and R. R. Spencer of the United States Public Health Service, who crushed ticks in phenol and injected a suspension of the product. In 1938, Herald R. Cox, who was then also with the Public Health Service, developed a technique for growing rickettsiae in the yolk sacs of chicken eggs. The parasite could be harvested, killed with formaldehyde, and extracted with ether to yield a vaccine. A commercial product based on this technique was first marketed by Lederle Laboratories in 1948. A new study sponsored by the Food and Drug Administration and conducted in 1973 by Herbert L. DuPont, now at the University of Texas Medical School at Houston, Richard B. Hornick, and colleagues at the University of Maryland showed that neither the tick nor the egg vaccine conferred immunity in man, and the Lederle product was subsequently withdrawn from the market.

The new vaccine has been developed over the course of 8 years by Richard H. Kenyon and his associates at the U.S. Army Research Institute of Infectious Diseases at Ft. Detrick, Maryland. Their chief goals were to increase the number of rickettsiae obtained from cultures and to eliminate the egg yolk lipids and proteins that were a contaminant in the commercial vaccine. In most individuals given the vaccine, these contaminants often produce swelling and tenderness at the injection site and in people who are allergic to eggs can produce a much more severe reaction.

After screening several tissue culture systems to see which yielded the highest concentration of rickettsiae, they began growing the parasites in monolayer cultures of cells from duck embryos. The principal problem with this and all other culture systems is that rickettsiae are very sensitive to the penicillin and other antibiotics commonly used to prevent contamination of the cultures (the parasite is less sensitive to the antibiotics in man); contamination can thus be prevented only by scrupulous care in handling the cultures. After incubation, the parasite-infected cells are broken up by alternate freezing and thawing followed by sonication; the cellular debris is removed by centrifugation at low speed, and the rickettsiae, which remain in the supernatant are then killed with formaldehyde.

Kenyon found that substantial quantities of rickettsiae are present in the suspension, whereas a microscopic examination of the commercial vaccine revealed no intact parasites. He and his colleagues showed that injection of the suspension into guinea pigs induces formation of a high concentration of serum antibody to the parasite. They also found, however, that only limited quantities of pathogen-free duck embryo cells are available.

They thus replicated the work with chicken embryo cells and found that they obtained only slightly fewer rickettsiae. They then tested vaccines produced from both duck and chicken embryo cultures in rhesus monkeys and found that two doses of either vaccine, administered at 15- or 30-day intervals (the common procedure for killed-bacteria or killed-virus vaccines) provide complete protection against a subsequent challenge with the parasite.

They have since tested a dilute version of the chicken embryo vaccine in 16 human volunteers. They observed that a small quantity of antibody was induced, and that there were no apparent side effects. They are now planning to try the undiluted vaccine in ten more humans to see whether it confers protection. If that attempt is successful, Kenyon says, some further testing will be required, but the vaccine could be available for widespread use within 2 to 3 years if there is sufficient demand.

How large that demand might be is not clear. Before the Lederle vaccine was withdrawn from the market, the company sold about 25,000 doses per year. Use of that vaccine was somewhat limited, however, because of the local reactions to immunization and because of the relatively poor immunity it conferred. A better vaccine with fewer or no side effects might find a much larger market.

A major target might be children who live in the tick-belt states and who play in wooded areas. Another might be backpackers and hikers, who represent a growing segment of the population. The Army may also find it useful to immunize recruits because many of its training camps are located in tick-belt states. The population at risk thus seems large enough to support commercial production, despite the relatively small number of cases reported each year. No manufacturer has expressed interest in such production yet, but that would seem to be only a matter of time.

—THOMAS H. MAUGH II

conducting computers that operate with characteristic times of nanoseconds or less. In 1 nanosecond, light in vacuum travels 30 centimeters; an electrical signal moves only a few times more slowly in these circuits.

Mention of superconductors invariably calls attention to the need for cryogenic refrigeration systems to cool them to their operating temperatures and to the expense and inconvenience of such systems. Brown says that he has worked out the characteristics of a refrigeration system of the type needed to cool an object 10 by 10 by 30 centimeters. A closed cycle refrigerator—that is, one which collects helium gas formed when the liquid evaporates and reliquefies it—could be about the size of an office desk. The power required to operate the refrigerator would be about 2 kilowatts, which is much less than the power requirement of 500 kilowatts of one existing supercomputer, the Cray-1 made by Cray Research, Inc. The Cray-1 occupies a cylindrical volume about 1.8 meters in diameter and about the same height (*Science*, 27 January, p. 304).

A Supercomputer Not the Only Option

What kind of computer would IBM, or anyone else, build around superconducting electronics? All the signs suggest a supercomputer; new technology appears in the Cadillac before the Chevrolet. IBM's current top-of-the-line computer, the 3033, has a characteristic time, called the cycle time, of 58 nanoseconds. The 3033 could be duplicated with superconducting electronics and have a cycle time of from 2 to 5 nanoseconds. The Cray-1 has a cycle time of 12.5 nanoseconds and is the leading example of what semiconductor technology can do. But the Cray-1 is a special-purpose computer for scientific computation that is not designed for the sort of general work that IBM is now oriented toward.

Making a prediction, however, is complicated by Anacker's estimate that at least five more years of research and an unknown period of product development lie ahead of a superconducting computer, if all the problems are solved. The technology underlying computers is a fast-moving one, and one well-known computer designer judged it a foolhardy act to try to foresee the shape of computers a decade or more into the future.

The semiconductor manufacturers are, for example, famous for their ability to cram every year more and more devices on a single small chip of silicon, thus offering more and more computing power at little or no increase in cost—the fuel of the electronics revolution. Further widely forecast advances should reduce the difference in the performance of super- and semiconductors.

An optimistic Anacker responds that advances in semiconductor technology should be paralleled by improvements in that of the superconductor at a rate sufficient to maintain their relative positions, and that eventually the heat dissipation limit will halt the advance of semiconductor technology. One skeptic takes a different point of view. Noting the gross disparity between the amount of effort going toward improvements in semiconductor technology and that directed at developing superconductors, he suggested it was hardly likely that the rate of progress in the two fields would be parallel.

Another problem in forecasting future computers is that the key to faster computation need not necessarily be faster switching devices. Rex Rice of the Fairchild Test Systems Group, San Jose, California, points out that computer designers have not yet begun to explore the limits of software (programs that tell the computer what to do) and architecture (organization of the component devices that determines how data are routed in order to carry out a particular operation). One example of a way to enhance the speed of a computer without going to faster switches is called parallel processing. In this mode of operation, several parts of a computer program are carried out simultaneously rather than sequentially, which requires alterations in both software and machine design.

A further peril of trying to guess how superconductors will impact computers lies in the observation that the advantages of high-speed switching do not necessarily lead only to supercomputers. It might be economically feasible to build lower performance computers using superconducting electronics. Although making a minicomputer seems out of the question because there is a certain minimum capital cost involved in using superconducting technology, Anacker says that anything from a medium-sized computer on up seems fair game for superconductors. One way to use superconducting electronics in lower performance machines is to make them from fewer circuits, thus vastly reducing the complexity of the computer. The expense of designing a computer is a major item in its final cost. The extreme speed of the few circuits used would permit serial processing of each binary bit of numbers being, for example, added together to be faster than parallel processing of the bits by semiconductor circuits that are more numerous but slower.

When, or if, a superconducting computer will be available is open to question. Some basic technology issues remain to be answered. Existing devices have a substantial failure rate when the number of thermal cycles between room and liquid helium temperature reaches several hundred. Although no one yet knows how often such cycling would occur, better performance in this regard is still needed. The ability to make reproducibly large arrays of devices on a single chip is yet to be shown, although the Zurich researchers say that the characteristics of their memory chips exhibit a range of values within about 10 percent of each other. This degree of reproducibility is said to be close to what is needed. But large-scale logic and fast memory arrays have yet to be made. One observer judged that no serious comparison between superconducting electronics and semiconductors was possible until the ability to make large batches of arrays was demonstrated. Another untested area is the miniature electrical connectors between wafers. These units are planned to be plug-ins similar to the circuit cards in today's electronics. But the electrical characteristics of such mechanical connectors at picosecond switching speeds is unknown. The inductance associated with the connectors could distort signals passing through.

Josephson junctions are not the first superconducting devices to be considered for computer applications. In the late 1950's, a class of switches called cryotrons was intensively studied at several laboratories. In part because the ability to make arrays of thousands of devices on a single chip did not then exist and in part because cryotrons turned out to be slow switches, the effort was ended on a bad note; superconductors were left with a poor reputation, at least among some electronics people.

IBM's success with the Josephson junction technology seems to be changing the situation. Stimulated by the progress to date, numerous academic and industrial groups are getting involved, if only to keep abreast of the technology. Especially encouraging to superconductivity researchers is that investigators are broadening their perspective to include kinds of instruments other than computers, such as analog-to-digital converters and microwave signal processors, as potential ways to use superconducting electronics. Thus, even if a superconducting computer fails to be developed in the foreseeable future, the effort at developing Josephson junction technology need not have been in vain.—Arthur L. Robinson