New Superconductors for a Supercomputer

A reorganization at IBM fuels rumors that the company is switching from lead alloys to niobium for its Josephson junction supercomputer project

Without IBM in this field, no one would be anywhere. We are riding on their backs.—DANIEL PROBER, Yale University.

The supercomputer of the 1990's could be built around Josephson junction switches made from superconductors rather than transistors made from semiconductors. In the United States, IBM is carrying so much of the load in developing the technology for such a machine that any sign of change in the company's long-standing and well-funded research in this area causes everyone else concern. One such change, occurring earlier this year, was a realignment of the management of IBM's Yorktown Heights, New York, laboratory, which among other things placed the Josephson junction electronics program organizationally closer to semiconductor research. A semiconductor expert has also been put in charge of the superconductors.

Although the new man, Joseph Logue, says the ultimate goals of IBM's Josephson junction program are unchanged, in the near term, he will be heading up an effort to determine once and for all whether the materials problem, one of several that are blocking a commercial product, can be overcome. The prevailing sentiment among researchers outside the company is that this effort probably will involve a switch from the soft, lead alloy superconductors that have been the mainstay of research to harder—but more difficult to work with—niobium superconductors. As it happens, physicists and materials scientists in several laboratories have recently reported encouraging results that make niobium (or other refractory superconductors) look more promising than it had before.

All the electrical signals flowing through the circuitry of a computer are periodically synchronized by an internal clock that has a characteristic period or cycle time. Supercomputers of the future may have cycle times of 1 nanosecond or less. This is more than 25 times as fast as today's large, general-purpose computers can run and is about 10 times as fast as the Cray-1, a special-purpose scientific computer.

The speed of light and thermodynamics conspire to prevent arbitrarily decreasing the cycle time. Electrical signals can travel no faster than light, so all the components of a computer (even one made of infinitely fast switches) must be squeezed into a volume about 10 centimeters on a side or smaller to achieve a 1-nanosecond cycle time. Moreover, the minute power needed to run a transistor is dissipated as heat. The 10 to 100 million transistors in the logic and memory circuits of a supercomputer made of semiconductors would emit so much heat that a machine this small would fry in its own fat. The lure of Josephson junctions, which are at present only twice as fast as the best semiconductor switches, is that they use only one-thousandth the energy and can be packed more closely together. Short of radically rethinking how computers are organized, IBM and others have argued, superconductors are the best bet for an ultrafast computer.

Nonetheless, until recently, only IBM has been willing to make the considerable investment of time and money needed to develop the all-new Josephson junction technology in the face of the mountainous body of knowledge already existing about semiconductors. Although the company prefers not to give out figures, outsiders estimate that the IBM program, which started with the 1965 experiments of Juri Matisoo, encompasses about 140 professionals. Other research efforts, at Bell Laboratories, the Sperry Research Center, the National Bureau of Standards in Boulder, and elsewhere, are each less than one-tenth this size. One of the causes of the recent rumors about changes at IBM is that the company, which had been building up its staff with new Ph.D.'s from university groups doing superconductivity research, has stopped hiring. Logue admits that a few scientists have even been taken off the Josephson junction program, but says that stories that the personnel had been reduced by 25 percent or more are "inaccurate."

IBM is no longer the only major player in the game. In Japan, the Ministry of International Trade and Industry (MITI) has launched another one of its infamous (to American electronics and computer companies) government-industry cooperative research ventures. This one, involving four firms, is an 8-year, $300-million (half from the government and half from the companies) program to produce a supercomputer by the end of the decade. Josephson junctions are one of three alternative technologies being considered. (Silicon, the presently used semiconductor, and gallium arsenide, a faster but more complicated semiconductor, are the others.) Martin Nisenoff of the Naval Research Laboratory estimates that the size of the Josephson junction research component of the MITI program will be comparable to that of IBM shortly. Commented one observer, not entirely humorously, in bringing Josephson junction technology up to its present high state, "IBM may have opened Pandora's box, and out jumped the Japanese."

The logic and memory circuits that make up a Josephson junction computer are built up from SQUID's (Superconducting Quantum Interference Devices) or other similar devices. These, in turn, consist of two or more Josephson junctions arranged in electrical parallel. The most common form of Josephson junction for computer applications is the tunneling junction, which consists of two strips of superconductor (a few micrometers wide) separated by a thin (20 to 40 angstroms) layer of insulator where they overlap. Circuits are built up on silicon substrates in much the same way (by photolithographic and other processes) as are silicon integrated circuits. But in Josephson junction circuits, the silicon substrate plays no electrical role.

Following IBM's lead, most researchers making experimental devices have used lead alloys (lead-indium-gold and lead-bismuth) as the superconducting elements in the circuits. Investigators have found that junctions made of these materials with a lead-indium (native) oxide as the insulator or tunneling barrier exhibit good switching characteristics that are reproducible enough to allow making superconducting integrated circuit chips with hundreds or thousands of Josephson junctions on them. Switching refers to a transition between nonresistive and resistive states that represent the 0's and

1's of binary digital circuits. In the nonresistive state, electrical current flowing through the tunneling barrier between the superconductors meets no resistance. In the resistive state, it does, and a voltage appears across the junction.

Logue emphasizes that there is a considerable difference between making experimental circuits in a research laboratory and producing large numbers of them in a factory. For this reason, IBM is developing an experimental production line for Josephson circuits at its East Fishkill, New York, facility. One example of a device characteristic that must be the same to within a few percent for every junction on a chip and the same for every chip is the critical current, the maximum current a Josephson junction will carry before it switches from a nonresistive to a resistive state. The critical current varies so strongly with the thickness of the oxide insulating layer that researchers must be able to regulate the thickness to a single atomic layer. Logue says that IBM's present goal is to learn whether the production process can be controlled well enough "to have a meaningful long-term program."

Observers speculate that this kind of thing is exactly what Logue was brought into the Josephson junction program to do. He has a reputation within the company for successfully converting research projects into commercial technologies. "He knows how these problems have been solved before," was the way one researcher put it. Logue has also closed down projects that were not working out, such as an early 1960's superconducting electronics program involving devices called cryotrons.

An additional problem that Logue will have to wrestle with is the choice of materials. Despite the success of lead alloy superconductors in Josephson circuits, their mechanical properties are far from ideal. A weak point of Josephson junction electronics is that superconductors must be cooled to a temperature in the vicinity of that of liquid helium. The refrigeration itself is not the problem; it is the cycling between cryogenic and room temperatures that is the villain. Because the coefficients of thermal expansion of silicon, the lead alloys, and the other materials in a Josephson circuit differ considerably, large stresses build up during heating and cooling. Because the substrate is the most massive element, the other materials must conform to the contraction and expansion of silicon. The lead alloys, being mechanically soft, are highly susceptible to damage.

Masanori Murakami and others at the Yorktown Heights laboratory have shown that deformation of the lead alloy material by the stresses due to thermal contraction and expansion causes a rupture of the oxide layer and failure of the junction. The IBM investigators have also found that they could greatly enhance the resistance of Josephson circuits to thermal cycling damage by depositing the lead alloy superconductor at the temperature of liquid nitrogen (77 K). The low-temperature deposition caused thin (0.2 micrometer) lead alloy films to have a grain size comparable to their thickness. The boundaries between the crystalline grains act as traps for certain imperfections (dislocations), which agglomerate under stress and cause the rupturing. With small grain sizes, there are fewer dislocations in each grain, and deformation is reduced. With small grain sized lead alloy superconductors, Josephson circuits can be cycled about 500 times between room temperature and liquid helium temperature before failures begin to appear, whereas 100 percent of larger grain sized superconductors have failed by this time.

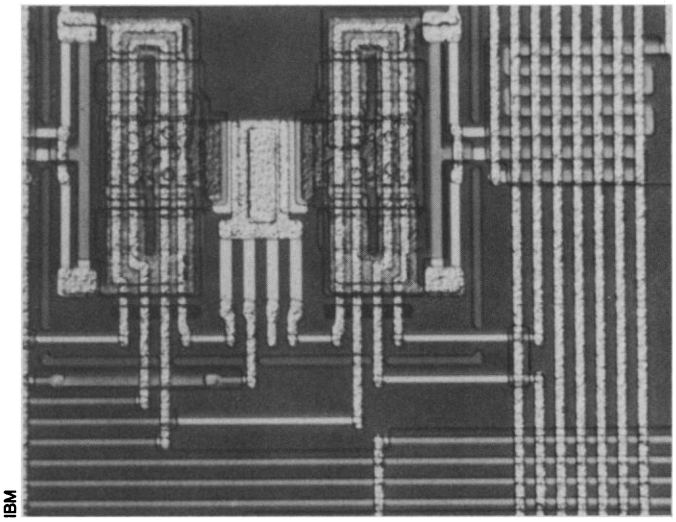

Josephson logic

The two devices at the top-left and top-center are SQUIDS which IBM calls three-junction interferometers. The lead-alloy Josephson junctions are under the four tiny circles near the center of each interferometer. The 2.5-micrometer-wide conducting lines connect the two interferometers and another device to make an AND logic circuit.

IBM

But how many cycles are enough? A computer, once running, would likely be kept cold indefinitely. Logue says that the desirable number of cycles is still a big unknown, and he is "nervous" about the performance to date. IBM is, therefore, looking at other materials, while continuing to work with lead alloys. Some researchers say that sooner or later a switch to niobium or some other refractory material will have to be made. Theodore Fulton of Bell Laboratories asserts that the feeling in the research community is that the only question is when to make the change—now, or after further progress toward a demonstration computer made with lead alloy Josephson junctions.

Niobium has a slightly higher transition temperature (maximum temperature at which it still exhibits superconductivity) of 9.2 K than the lead-indium-gold alloy (7 K). What is more important is that it is a refractory metal with a very high melting temperature and a stiffer resistance to deformation. A good deal of IBM's work with niobium Josephson junctions has been at the company's Zürich, Switzerland, laboratory. Last year, Ronald Broom of Zürich, Stanley Raider of Yorktown Heights, and their co-workers reported that Josephson junctions made from niobium superconductors could be cycled between room and liquid helium temperatures indefinitely. Two kinds of devices were tested. The first consisted of a niobium base (bottom) electrode, niobium oxide (Nb_2O_5) insulator, and lead alloy counter (top) electrode. Four chips with a total of more than 48,000 Josephson junctions exhibited no failures in 1400 cycles. The second device type had a niobium counter electrode. Fewer of these junctions have been tested, but again so far none have failed. The all-niobium devices have poorer electrical behavior than the niobium-lead hybrids, so researchers cannot yet say they fully understand niobium.

It is the niobium oxide insulator that seems to be at fault. To start with, niobium oxide has a dielectric constant that is about three times larger than that of lead-indium oxide. The high dielectric constant translates into a high capacitance junction and therefore a slower switching action. (The high dielectric constant also slows the propagation of electrical signals through transmission lines consisting of superconductors separated by niobium oxide insulators.) IBM's persistent pursuit of lead alloys is said to be due to the company's unwillingness to settle for the poorer electrical performance of the mechanically superior niobium superconductor.

A Computer Is More Than Chips

Making any computer, much less a supercomputer, from superconducting Josephson junctions, involves considerably more than perfecting these devices. It also takes more than the ability to make densely packed integrated circuits with from several hundred devices per chip (logic) to 16,000 or more devices per chip (memory). For a supercomputer with a 1-nanosecond cycle time (see story), engineers have to find a way to pack more than 10,000 chips in a volume 10 centimeters on a side in such a way that the heat generated by the incessant switching action of 100 million Josephson junctions can be dissipated and in a way that electrical signals can travel anywhere they need to in one cycle time or less.

A very significant step in this direction has been reported by Mark Ketchen, Bernard van der Hoeven, and 25 other authors at IBM's Yorktown Heights Laboratory.* The researchers have demonstrated a "cross-section model" with a cycle time of 3.7 nanoseconds. A cross-section model is not a complete computer or even a special-purpose signal processor. But it does have all the kinds of

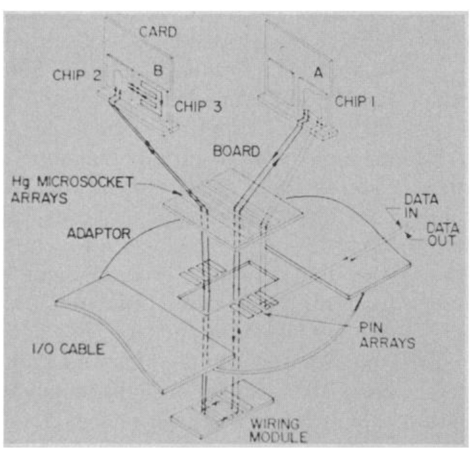

IBM

circuits that a signal processor would have. IBM's cross-section model simulates the longest path that is likely to occur in a signal processor. A second important feature of the cross-section model is that the manner of arranging and connecting (packaging) the superconducting integrated circuit chips is representative of what would be done in a full-scale computer.

As explained by Ketchen and van der Hoeven, the cross-section model consists of four identical integrated circuit chips. The chips are silicon substrates 6.35 millimeters on a side with Josephson junction circuits made with IBM's well-developed lead-alloy technology. The chips are organized into four identical quadrants that can execute different functions. The cross-section model is arranged so that in operation one quadrant from each of three chips is activated. It is possible to vary the connections between the chips and in this way carry out eight different "experiments" with the same four chips.

Each chip is bonded to a square silicon "card" 15 millimeters on a side that is analogous to the familiar printed circuit boards in electronic equipment of all types. Bonding is by way of miniature solder blobs. On the card are superconducting transmission lines that function like miniature cables to carry signals at nearly the speed of light and without smearing them out. The transmission lines connect the two chips on a card or lead to a second card that carries the remaining two chips. The two cards, which lie 3.35 millimeters apart and parallel to one another, are attached to a silicon "board" by an array of micropins that sit in feet on one edge of the cards. The micropins in the card feet are electrically attached by way of mercury pools embedded in the board to a matching set of pins in a wiring module on the opposite side. The wiring module is also silicon with lead alloy superconducting transmission lines. It is by changing wiring modules with different wiring patterns and by varying external electrical input signals that the researchers could select which quadrant of which chip was used in a particular experiment.

This "card on board" package sits in a bath of liquid helium. Connection to the rest of the world at room temperature is by way of a silicon adapter that the package is attached to and an 80-conductor flexible cable that runs from the adapter out of the refrigerator to power supplies, input-output devices, and test equipment. A full-scale computer might consist of 20 or so pairs of boards arranged back to back with wiring modules in between, and nearly 1400 cards with eight chips each.

In a typical experiment, the three parts of the data path are a master chip, a memory support chip, and a simulated memory chip. There are three input signals that control the operation of the master chip. If the inputs are correct, a signal travels from the master chip to the memory support chip to the memory chip and returns in one cycle time. The output of the master chip depends on the combination and the timing of the input signals. Since the inputs are spread over two cycles, it actually takes three cycle times to obtain an output. The IBM researchers measured a minimum cycle time of 3.7 nanoseconds.

Also incorporated into the four chips are various other circuits. One type of circuit is a power supply to convert incoming alternating current into a regulated alternating voltage with a certain waveform. Other circuits include test devices to measure the waveforms of the signal as it passes through various parts of the cross-section model.

Ketchen and van der Hoeven say that the main value of the cross-section model experiment was to generate a list of problems that need to be addressed before a more complete system can be attacked. The researchers encountered many difficulties in getting the cross-section model to work; it was by no means an exercise of wiring a few chips together. Heading the list of problems was the inability to achieve a uniformity in the electrical characteristics of the Josephson junctions. Apparently, IBM is not yet at the point where it can be sure that any given four chips, two cards, and board, when assembled into a cross-section model, will work together as a system. What the cross-section model does demonstrate is that, once the proper chips, cards, and board are found, the system works: a signal can be gotten through it without distortions or extra time delays. This, observers say, is a big step.—A.L.R.

*M. B. Ketchen *et al.*, *IEEE Electron Device Letters*, *EDL-2* (1981), p. 262.

Two results in the last year, one from IBM's Zürich laboratory and one from Cornell University, suggest one way to overcome the capacitance problem. Since the capacitance of a Josephson junction is proportional to the junction area, the investigators have made very small junctions. The limitation in making junctions small is the resolution of the photolithography process by which patterns are imprinted onto integrated circuits. The thickness of films can, however, be controlled to much smaller dimensions because the thickness is determined by the time a deposition process is allowed to proceed. An old idea is to use the edge of a film rather than its top as the surface to be oxidized in forming the junction. This approach was used by Lawrence Jackel, Richard Howard, Evelyn Hu, and their colleagues at Bell Laboratories, Holmdel, New Jersey, with lead-based Josephson junctions. Independently, Broom, Albert Oosenbrug, and Wilhelm Walter of IBM made niobium–niobium oxide–lead Josephson junctions on the edges of niobium films that were 0.3 micrometer thick. The widths of the superconducting strips are usually about 2.5 micrometers, so that the area of the junction was reduced by a factor of about 10.

At Cornell, Robert Buhrman, Alan Kleinsasser (now at IBM), and their co-workers had access to the National Research and Resource Facility for Submicron Structures and could therefore concentrate on making even smaller junctions. With widths of the superconducting strips of about 1.5 micrometers and film thickness of 0.15 micrometer, they made junctions with areas of a few times 10^{-9} square centimeter. With the use of electron beam lithography, instead of optical photolithography, they reduced junction areas to a few times 10^{-10} square centimeter. The Cornell researchers used a lead-bismuth alloy as the counter electrode. So far there are no reports of all niobium "edge" junctions.

Both the Zürich and Cornell groups had to take extreme care in constructing their Josephson devices. It appears that Nb_2O_5 is not the only niobium oxide that can form. The other possibilities are NbO and NbO_2, which are not insulators. In general, all three species form. Raider and Robert Drake of IBM have described an elaborate technique for minimizing the formation of the undersirable oxide species. A critical step involves surface cleaning of the niobium by radio-frequency plasma etching before forming the oxide layer. Plasma etching is a technique whereby ions in the plasma created by a radio-frequency electric field strike the sample and remove surface atoms in the process. The niobium-based Josephson junctions that survived thermal cycling so well were made by this process. Apparently, however, when a niobium counter electrode is deposited, the highly reactive niobium attacks the oxygen in the barrier layer and thereby degrades its properties, and this accounts for the poorer electrical performance of all-niobium junctions with niobium oxide insulator as compared to niobium-lead hybrids.

One reason that the cleaning procedure is so critical is that the lower electrode of a junction is subjected to a variety of processing steps before the actual tunneling barrier is formed and the counter electrode deposited. For example, in conventional processing, photoresist, a light- or electron-sensitive polymer, is applied for the photolithographic patterning. At the Sperry Research Center, Sudbury, Massachusetts, Harry Kroger, Lawrence Smith, and Don Jillie have developed a process which avoids any such "intentional" contamination. They form all Josephson junctions before any patterning steps. Individual junctions are subsequently isolated from each other by anodization (oxidation). Their selective niobium anodization process or SNAP is combined with the use of a deposited silicon tunneling barrier layer. In this way, they simultaneously overcome the deleterious effects of the undesired niobium oxide species and they get a Josephson junction with a capacitance that is even lower than that of lead alloy junctions.

Kroger argues that the SNAP process actually brings Josephson junction fabrication more in line with the well-developed silicon integrated circuit production procedures. In silicon integrated circuits, the ease of formation of silicon dioxide and its use as a mask for pattern making and as an insulator for electrical isolation of one part of a device from another is probably the single biggest reason why microcircuits are made of silicon and not some other semiconductor. In the SNAP process, Nb_2O_5 plays a similar role. The other oxide species cause no problems here because the oxide layer is thick and robust, not thin and delicate as in the tunnel junction barrier.

With the use of a silicon tunnel barrier layer in place of niobium oxide, the Sperry researchers have made all-niobium Josephson junctions with good electrical properties. There are not a lot of data yet on thermal cycling behavior, but presumably the devices will exhibit the resistance to failure demonstrated at IBM. The group is also making Josephson junction integrated circuits. Rather than the SQUID's that IBM uses, Kroger, Smith, and Jillie are working with a somewhat different kind of switch with the colorful name JAWS (Josephson-Atto-Weber Switch) that was introduced by Fulton, S. S. Pei, and L. N. Dunkelberger at Bell Laboratories.

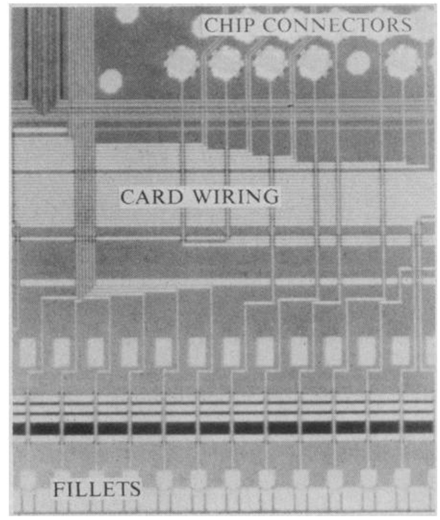

Chip carrier

In a computer, integrated circuit chips would be attached to "cards" by way of the solder blobs at the top of the photograph. Wiring connects different chips on the same card together or to chips on other cards by way of the fillets at the bottom of the card.

Also at Bell Laboratories, Michael Gurvitch and John Rowell have been experimenting with an approach similar to that at Sperry. Rowell says that he, Gurvitch, and others were making some basic studies of tunneling behavior in niobium-aluminum multilayers, and their results led to an investigation of thin aluminum layers (10 angstroms) on niobium. Gurvitch and Rowell then found that the aluminum layer, when oxidized, formed an excellent tunneling barrier layer in niobium–aluminum oxide–lead junctions and found encouraging results in niobium–aluminum oxide–niobium junctions. They also found that a second layer of aluminum just under the niobium counter electrode (symmetrical structure) improved the properties of the all-niobium junctions. The researchers use a SNAP procedure similar to Sperry's.

When asked if the Japanese goal of a supercomputer by the end of the decade was feasible if the machine were to be made of Josephson junctions, IBM's Logue smiled and said "no comment." When and if a superconducting computer does make an appearance, researchers seem confident that it will use Josephson junctions based on niobium (or some other refractory material), not lead alloys.—ARTHUR L. ROBINSON