Are VLSI Microcircuits Too Hard to Design?

Engineers may find themselves unable to design the complicated microcircuits made possible by advancing miniaturization technology

Although there are substantial technological obstacles to be overcome, semiconductor experts confidently predict that microcircuits with many millions of transistors will be manufacturable in the next decade. Except for one thing. There are few systems in the world, electronic or otherwise, that consist of a million or more parts. Organizing a million transistors on a silicon chip into a useful circuit is a formidable undertaking, and the tools for designing such a complex device do not yet exist. The integrated circuit industry is on a collision course with complexity. Finding faster and less expensive ways to design microcircuits will be necessary if the industry is to match its projected growth from sales of about $11 billion last year to $50 billion or more annually by 1990.

This is the last of five Research News articles on microelectronics.

Microprocessors are the most intricate integrated circuits now being made (although they may have fewer transistors than the largest computer memory chips). It already takes a team of a half-dozen design engineers a year or two to develop a state-of-the-art microprocessor. The next generation of integrated circuits, which goes by the name of VLSI for very large scale integration, is where the design problem becomes acute. Says Dennis Buss of Texas Instruments, the Dallas-based company that sells more microcircuits than any other in the world, "the dominant VLSI factor is that it is too hard to design." At best with existing computer aids, one person can in a year's time design about a thousand gates in an integrated circuit at a cost of approximately $100 per gate. (A gate is the basic computer logic circuit and consists of just a few transistors. There are several types of gates, one for each type of logic operation.) By 1990, 1 million gates on a chip may be possible, but even if the cost per gate decreased to $10, the $10 million necessary to design the entire microcircuit would still be out of the question, says Buss.

A look at present integrated circuit design practice helps explain why the process is so lengthy. A microprocessor, which is the central processor unit of a computer shrunk to the size of one microcircuit, consists of several subunits, such as the circuitry that performs arithmetic operations, registers that temporarily hold data or instructions, and controllers that move instructions and data between the microprocessor and memory or various peripheral accessories (sensors, displays, and keyboards). A logic designer builds up these subunits according to the principles of computer design. Then a circuit designer converts the logic design into specific electronic circuits. And a layout designer works out the placement of logic gates on the silicon chip that corresponds to the desired circuits. Finally, a draftsman draws an accurate circuit layout for use in the microfabrication process. Moreover, the design procedure is often iterative: bugs are found and corrected; modifications are made to improve performance; and bugs introduced by the modifications are detected. Even after a chip is made and tested, engineers may find that it does not operate quite the way they would like, and another round of modifications progresses through each stage of the design loop. The process quickly becomes time consuming and expensive.

The expense of designing complex intergrated circuits has for several years contributed to a certain amount of friction between semiconductor makers and those customers who want special-purpose chips for use in the electronic equipment they manufacture. Because of the tremendous expense of designing complex integrated circuits, manufacturers prefer to develop only products that can be sold in large volumes, computer memories being the outstanding example. Chips that cannot be sold in large quantities will not repay their development costs, but microprocessors, memories, and other standard circuits are not always the most effective components for an electronic system. At an annual symposium that reviews the state of the industry (held in May in New Orleans by Rosen Research, Inc., a New York firm that analyzes the electronics and small computer business), the issue of custom chips was raised once again by several participants.

In response, Charles Sporck, president of National Semiconductor in Santa Clara, California, and Jerry Sanders, chairman of Advanced Micro Devices in neighboring Sunnyvale, each encouraged electronic equipment and computer manufacturers to set up their own design and fabrication facilities for integrated circuits crucial to their products but needed in quantities too small to be economical for semiconductor companies to produce. This is exactly what many organizations that are large enough to afford it have done. But smaller firms cannot afford this approach. There is a large need for custom integrated circuits that the semiconductor industry is not adequately meeting, admits Gordon Moore, chairman of the Intel Corporation.

What can be done to bring the problem of managing complexity in VLSI microcircuits down to a level at which it can be dealt with? Some of the possibilities are extensions of ideas already incorporated in some integrated circuits. They also involve sacrificing some performance (speed, number of transistors on the chip, and so forth) in order to make the design process simpler. A computer memory chip is largely a rectangular array of memory cells, each storing one bit of information. A microprocessor or a digital logic circuit, in contrast, looks more akin to a bowl of spaghetti. Although memory microcircuits are far from trivial to design, forecasters predict that the ability to store more and more bits on a chip will be limited by miniaturization technology. It is in microprocessors and digital logic circuits that the design problem lies. One solution is to make logic circuits more regular. Two ways of doing this are already in use.

The first is the programmable logic array (PLA). A PLA consists of two kinds of gates arrayed in a very specific manner. There are two arrays connected together; one is made entirely of gates of the first type, while the other is of the

0036-8075/80/0711-0258$01.25/0

second type. The PLA is called programmable not because a computer program directs its operation, but because the logic function of the array is determined by which devices are interconnected within the two arrays of identical gates. Because the gates are laid out in a regular pattern, the main design problem is to decide which gates to connect. However, the penalty paid for a simpler design process is that a PLA takes up more space than a dedicated logic circuit. Since the cost of fabricating an integrated circuit is determined almost entirely by how much area it occupies, there is a considerable premium on keeping areas down.

Nonetheless, PLA's are now quite common in microprocessors for certain functions such as circuitry to control the execution of instructions in the program that the microprocessor is running. In addition to being easier to design, PLA's offer a certain amount of flexibility. Simply by changing the interconnections between the gates, the function of the PLA can be altered. In this way, an engineer could make major changes in a microprocessor without going to an all-new design, says Robert Sproull of Carnegie-Mellon University.

The second approach to regular logic structures sounds superficially quite similar to PLA's but in fact is quite different. Master slices or gate arrays provide a way to manufacture custom integrated circuits at a considerable saving in design time. Companies that make large, high-speed computers, for example, may need a thousand or more special-purpose microcircuits in order to squeeze the maximum performance out of their machines, and the cost of designing each from scratch would be prohibitive. Richard Petritz, president of the Inmos Corporation, a British company located in Colorado Springs, has estimated that the time from initial design to production can be reduced by as much as a factor of 6 when recourse is made to master slices.

The master slice idea is actually several years old. After an initial wave of interest, the concept found limited application, as only IBM used it successfully in designing digital logic circuits for its newest computers. Now several semiconductor companies are offering master slice chips. As manufactured at IBM, master slices consist of columns of identical "cells" comprising several interconnected transistors. The columns are separated by channels filled with the strips of evaporated metal that serve as wires in an integrated circuit. All master slices are identical until the last stages in the manufacturing process, when the cells are connected to the wires in the wiring channels. The particular pattern of interconnections determines what kind of logic circuit the master slice becomes. According to James McGroddy of IBM's Yorktown Heights laboratory, the process there is largely automated. The designer indicates to the computer what function the circuit is to perform, and the computer determines the pattern of interconnections. The company can design hundreds of unique circuits per year. A variation on the master slice theme has been developed by Motorola, which also leaves the interconnections between transistors in a cell unspecified until the last stage of fabrication. A computer library stores the interconnection patterns for some 85 different cell configurations.

Because of the tremendous expense of designing complex integrated circuits, manufacturers prefer to develop only products that can be sold in large volumes.

Master slices and PLA's are similar in that the function of the completed array is determined by the pattern of interconnections. But the rather rigid structure of the PLA limits what functions it can perform. "You wouldn't want to use a PLA for just anything," says McGroddy. Master slices are more versatile but are regarded primarily as a partial solution to the custom part problem by those in the semiconductor industry because they use space on the silicon chip inefficiently.

Increasing the geometric regularity is not the only approach to simplifying integrated circuit design. Dividing a chip into functional subunits which become the building blocks with which the microcircuit is constructed is another concept, also being used in a limited way. In the earlier days, the building block was the transistor itself. In master slices, the building block is the multitransistor cell. More complex building blocks are the logical next step. Breaking up a microcircuit into functional subunits is an idea borrowed from the discipline of computer programming. Paul Penfield of MIT points out that large programs may contain a million or more instructions and thus are comparable in complexity to the projected VLSI microcircuits. The ideas developed by programmers over the years for dealing with complexity, if not directly applicable to integrated circuit design, provide a starting point for attacking the problem. Structured programming, for example, is a technique that divides a complex program into modules, writable by separate groups of programmers, that are connected in highly specific ways. The strictures on intermodule communication ensure that the overall program will work when the modules are connected. A similar method could be applied to microcircuit design, says Sproull, provided that account is taken of the need not only to logically organize the circuit but also to spatially organize it so that it fits on a chip.

The building block idea is being used in embryonic form by Texas Instruments in its TMS-1000 microcomputer. (A microcomputer is a microprocessor together with memory and circuitry to facilitate communication with input/output devices, all on one chip.) A microcomputer has two types of memory: a random access memory in which information can be stored and from which it can be retrieved and a read-only memory whose content is fixed. Since small microcomputers most often are used for only one purpose (run only one program), the program is conveniently stored in a read-only memory. To change the function of the microcomputer, all that is necessary is to change the read-only memory. However, some applications may involve long programs and others short. Some applications may require storage of more data than others. The company therefore does not sell just a single version of its microcomputer. It offers 14 versions, which are differentiated by the size of the random access and read-only memories and by the input/output circuitry, James Fischer of Texas Instruments reported at the previous New Orleans semiconductor forum last year.

But, asks David Hodges of the University of California at Berkeley, why stop there? An important characteristic of a computer is the width (number of bits) of a computer word. Just as memories are treated as a functional block whose size can be varied by specifying certain parameters in the computer program that holds the information about the memory design, so it might be possible to have a design program for an arithmetic logic unit whose word width is variable. In fact, a whole set of generic units within the computer could be designed so that the characteristics of the machine could be fixed by specifying a set of parameters in the design program.

Intel's Moore has gathered data on the increasing time it takes integrated circuit companies to design complex chips. The curve is an exponential one, paralleling the exponential increase in the number of transistors per chip. Hodges says that his form of the functional block approach will not reduce the time to design *the first chip* because the generic subunits still have to be generated in the conventional way. It is in designing related chips in the same family that great savings of time accrue because the hard work has to be done only once.

It is possible to combine the functional building block and master slice approaches. A chip would then consist of an array of building blocks (which would not necessarily all be alike) that were connected toward the end of the manufacturing process. The manner in which the blocks were interconnected would determine the function of the microcircuit. To ensure that the completed circuit functions as designed, however, it is preferable to make the interconnections only after the chip is made, says Allan Anderson of MIT's Lincoln Laboratory.

As circuits become more miniaturized, they are more susceptible to defects that prevent them from functioning properly—one defect, and the whole chip must be discarded. "Yield" is the fraction of chips produced that are defect-free. With the prospect of 1 million transistors on a chip, it becomes feasible to put on more transistors than will actually be used. If, during testing, certain portions of the chip are found to be defective, the unused transistors can be activated to replace those that do not work properly. Both IBM and Bell Laboratories use this principle in their largest random access memory chips. The semiconductor industry as a whole has not yet adopted this philosophy because redundancy takes up space on a chip and because yields are still high enough to make it worthwhile to gamble that enough defect-free chips will be produced to be economical.

The Lincoln Laboratory approach is to subdivide a complex VLSI circuit into smaller subunits of perhaps 2000 gates each. One advantage of doing this is that subunits of this size are still relatively easy to design. Another advantage is that yields for chips of this size can be reasonably large, perhaps 50 percent. A large VLSI microcircuit with 1 million gates would require 500 working subunits. Thus, the Lincoln Laboratory procedure would be to make a chip with 1000 subunits. Testing of the completed microcircuit would determine which of the subunits functioned properly and these would be interconnected. Anderson and his co-workers call this approach restructurable logic. In this way, very complex systems would also be customizable by means of adjusting the interconnection patterns. Anderson says that eventually it may be possible for certain applications, such as those in the military where low weight and power consumption are important, to vary the interconnection pattern during a computation. In this way, the same subunits could be used more than once and for different purposes.

Finding ways to make interconnec-

A Way to End the IC Designer Shortage

Universities have largely been left outside the microelectronics revolution. The number of laboratories in American universities that can design and fabricate a complex integrated circuit (IC) can be counted on the fingers of one hand. Moreover, the highly competitive companies in the semiconductor industry have guarded their hard-won secrets of microcircuit design so closely that the outside world, including the universities, does not know all the principles by which chips are designed.

At the same time, there is a severe shortage of engineers capable of designing advanced integrated circuits. One set of statistics making the rounds of electronics meetings has it that there are about 1500 device designers in the U.S. industry, whereas the number of major league professional athletes (football, basketball, baseball, and hockey) is half again as large (2243). Not yet organized as well as professional athletes, design engineers may not be doing as well financially (observers say that creative circuit designers command salaries of $60,000 a year), but competition among the companies for their services is nonetheless intense. Gordon Moore of the Intel Corporation recently estimated that between now and 1990 the semiconductor industry worldwide would require 20 percent of all U.S. electrical engineers who graduated during that period. For Ph.D.'s, the demand would be even higher.

Moore suggested that the integrated circuit companies begin massive in-house training programs. However, a new strategy for teaching students how to design microcircuits developed by Carver Mead of the California Institute of Technology and Lynn Conway of the Xerox Palo Alto Research Center could simultaneously end the designer deficit and put universities back into the forefront of the microelectronics game. An extra feature of the Mead-Conway strategy, made possible by the Defense Advanced Research Projects Agency through its computer network ARPANET, Xerox, Hewlett-Packard, and Micro Mask, Inc. (a California company that makes the masks used in fabricating integrated circuits), has allowed students to get their designs implemented in chip form. Would-be designers can thereby test their ideas in the flesh, as it were, rather than contenting themselves with computer simulations.

According to Mead, the new design approach was brewed gradually over a 10-year period first at Caltech and then in a joint university-industry project with Xerox. (Cooperation was undoubtedly helped by the fact that W. R. Sutherland, manager of Xerox's systems science laboratory, and Ivan Sutherland, then chairman of Caltech's computer science department, are brothers). The first big test came in the fall of 1978, when Conway journeyed to MIT to teach a course in integrated circuit design. Conway says that the first 6 weeks of a single semester course were devoted to teaching the principles of the new design approach. Following the instruction phase, students began working on design projects. Nineteen projects were completed by the end of the course.

Circuit designs had to be put in a form that could be sent over the ARPANET. Meanwhile, at Xerox researchers had devised a computer program that "collated" the incoming circuit designs and organized them so that they all fit on one large chip. The output of the Xerox program was then sent to Micro Mask, which has a computer-controlled electron beam mask-making machine. The set of masks (it takes several to make one integrated circuit) from Micro Mask then went to Hewlett-Packard's Integrated Circuit Process-

tions after a chip has been fabricated is a major technological obstacle. Today there are chips called electrically erasable, programmable read-only memories. In one version of such a device, pulses of electrical current much larger than used in normal operation change the contents of a memory cell. In this way, the operating program of a microcomputer can be changed in the field. Anderson speculates that a similar technique in which either current pulses or bursts of laser light are used to make and break interconnections may be adopted.

Integrated circuit designers have had access to computer-aided design tools for several years, but it is generally agreed that for computers to make significant contributions to the design of complex VLSI microcircuits, some quite radical new thinking will be required. One would like, says Robert Kahn of the Defense Department's Advanced Research Projects Agency, the computer to take a more central role in the details of design and testing of integrated circuits and allow humans to concentrate on the more creative aspects of design. Another analogy with computer programming shows what the computer design aid could do.

Compilers are computer programs that translate programs written by users in so-called high-level languages, such as Fortran, Cobol, or Basic, into a form that the computer can understand, namely, a sequence of binary numbers that the computer interprets as instructions. An analogous sort of program could be envisioned which transforms a "high-level description" of an integrated circuit written by the designer into a detailed layout of every transistor and wire on a chip. Utopian as such an idea sounds, quite complicated integrated circuits have already been designed in this way.

At the California Institute of Technology, David Johanssen, a graduate student in Carver Mead's group, has devised a computer program that he calls "Bristle Blocks." At present the program is restricted to aiding in the design of a particular type of microprocessor. The microprocessor is partitioned into cells that are arranged as parallel columns on the silicon chip. The designer specifies in a high-level description that might take up only two or three typewritten pages what functions the cells are to perform. Using certain design principles developed by Mead of Caltech and Lynn Conway of Xerox's Palo Alto Research Center (see box), the computer program then generates the circuitry within each cell and locates the connection points where the cell must communicate with other cells. The name "Bristle Blocks" comes from the appearance of the rectangular cell covered with interconnections. The program then arranges the cells so that the proper connection points of adjacent cells meet. The program also designs a PLA that controls the execution of instructions in the microprocessor.

So far, says Johanssen, the Caltech group has not designed an entire microprocessor using "Bristle Blocks," but it is getting close. Once the functional descriptions of the cells are in hand, the program takes only 5 to 10 minutes to come up with a chip layout. But arriving

---

ing Laboratory, where the chips were fabricated and some preliminary tests were carried out. Conway says that by using the new design methods and by merging several projects into one mask set the average "effective" cost to design and get each project into chip form was a few thousand dollars. A typical figure in the semiconductor industry for designing and building a new complex integrated circuit is closer to $150,000, although part of this cost is for the salaries of the many specialists involved in optimizing the design. And the total time to design and fabricate the chips was a little over 2 months, whereas the industry timetable for developing complex circuits is 1.5 to 2 years. However, in industry, there may be more than one design cycle, again to optimize chip performance.

In the fall of 1979, the Mead-Conway design approach was being taught at 12 universities, which came up with 82 projects for implementation in chip form. One of the most complex projects was a complete microprocessor designed by Gerald Sussman and his co-workers at MIT (see story). And this spring, a third round of projects was completed at the same 12 universities. Among the 171 projects was one by Ronald Rivest and his colleagues at MIT, who designed a chip to implement a new encryption algorithm.

The interesting development is that neither students nor faculty members who have participated in the course had prior integrated circuit design experience. In fact, researchers like Sussman and Rivest were primarily acquainted with computers through having to write programs to run on the machines. What the Mead-Conway design approach seems to be able to do is tap a whole new community of scientists and engineers and quickly make microcircuit designers of them. Mead estimates that this year the universities will graduate 1000 engineers who will be qualified to design complex integrated circuits. If these people can move into the semiconductor industry and make contributions immediately, the much bemoaned shortage of circuit designers could be a thing of the past. How soon this will happen is not yet known. At their present stage of development, the Mead-Conway design principles require giving up a certain amount of circuit performance for ease of design. The most important characteristic surrendered is that of making maximum use of every square micrometer on the silicon chip surface (maximum possible number of transistors, for example). The semiconductor companies are not yet willing to make this sacrifice.

A large influx of new integrated circuit designers could have another effect on the semiconductor industry, says Douglas Fairbairn of VLSI Technology, Inc., in Los Gatos, California, who makes an analogy with the publishing business. In publishing, authors do the creative work (writing), printers make the books, while publishers coordinate activities at both ends and handle the marketing. In the semiconductor industry, one organization designs chips (the creative part in the analogy), fabricates them, and sells them. A large number of new designers might alter this arrangement, and a "dis-integrated" industry develop. Designers would become independent "authors" rather than employees of the semiconductor companies. A very few companies called "silicon foundries" would develop whose only role is to "print" chips. (Because chip processing is so capital-intensive, only a few firms would be feasible.) Already a few independent chip design groups have sprung up, and the so-called silicon foundry is a much discussed concept. Most observers think, however, that the role of such a dis-integrated semiconductor industry would be limited to making so-called custom chips, special designs sold in low quantities. The main products (computer memories, microprocessors, and other standard, high-volume parts) would continue to be made by integrated companies of the sort that dominate the industry now.—A.L.R.

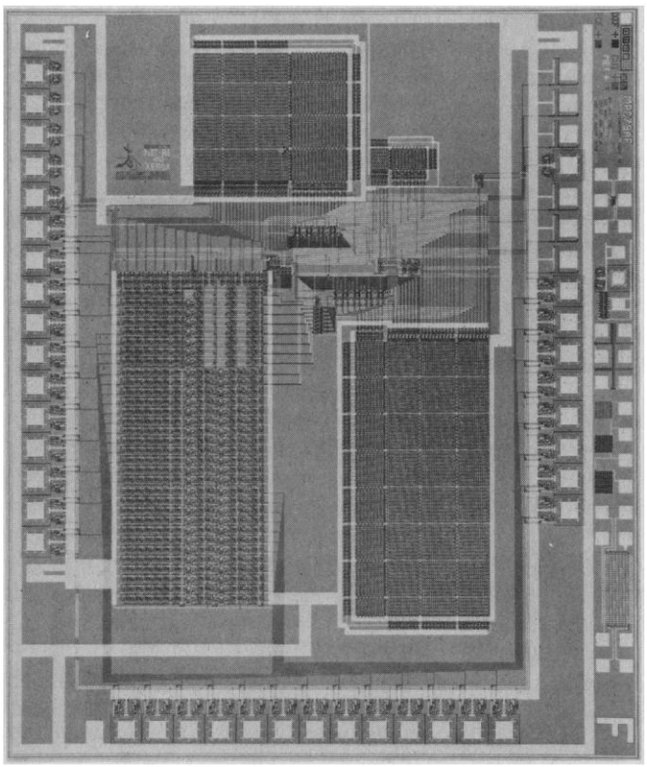

Microprocessor designed by Gerald Sussman, Jack Holloway, Guy Lewis Steele, Jr., and Alan Bell of MIT using artificial intelligence techniques. The microprocessor is geometrically more regular in appearance than commercial state-of-the-art chips, a trait that may be characteristic of many future microcircuits. [Source: Lynn Conway, Xerox Palo Alto Research Center]

at the functional descriptions is only half the job of designing a microprocessor, so "Bristle Blocks" by no means solves the entire problem.

Nonetheless, "Bristle Blocks" does accomplish about 3 man-years of work in a few minutes, and it would be hard to see how the semiconductor industry could be unreceptive to such a tool. Johanssen says he has talked to many industry people. The response has been of two types. Management is highly impressed and looks forward to reducing the cost of designing integrated circuits. However, layout designers are less enthusiastic. As with all the new design approaches, the Caltech program involves some compromises between performance and ease of design. Layout designers apparently are still primarily interested in "saving the last square micrometer of silicon chip area and the last nanosecond in speed."

Up to now, the people who design integrated circuits and those who design computers have belonged to quite separate communities (electrical engineers and computer scientists, for example). However, as the essentials of a computer get shrunk to chip size, the distinction becomes rather blurred. In fact, says Hodges at Berkeley, "there is something of a subconscious fight going on to establish which group will be doing the creative work in designing future computers." From one point of view, the integrated circuit people have a head start because they know how to design chips. From another viewpoint, the computer designers have an advantage because they have a backlog of accumulated expertise in computer science. The emergence of design programs like "Bristle Blocks" may upset this balance because it takes away the need of the integrated circuit designer to look at anything but the logic design. The program transforms that into a chip layout. A project under way at MIT under the direction of Gerald Sussman makes the point nicely.

Sussman does artificial intelligence research and he wants to understand how humans go about designing things from bridges to integrated circuits. One form of evidence that the human design process is understood is the ability to write a computer program that designs systems. Sussman and his co-workers have chosen VLSI microcircuits as the type of system they would like to write a program to design. Although their work is far from complete, the four group members have already constructed a program with which they designed in only 1 month a microprocessor with some 10,000 transistors. With projected improvements in their design program, they expect to reduce that time to a few days. The improvements should also permit designing a microprocessor with 50,000 transistors as easily as the existing program designed a 10,000-transistor chip. Sussman predicts that design tools based on artificial intelligence concepts should eventually permit integrated circuits of every kind except state-of-the-art chips, such as random access memories, to be designed from a high-level description of the desired functions written by software engineers (programmers), not circuit designers.

Most of the activity in figuring out faster ways to design integrated circuits has been directed toward making it easier to build computer logic circuits of the conventional type. But logic circuit concepts were formed when electrical switches (vacuum tube circuits) were expensive and the wires to connect them were negligible in cost. The philosophy was to design the logic so that there were as few switches as necessary. In digital integrated circuits, transistors play the role of switches. But the important point is that the cost of a transistor or a wire on an integrated circuit is determined only by how much area on the chip it occupies. On today's microcircuits, interconnections already account for as much as half of the area of a chip, with the remainder shared by transistors and the inactive areas that separate one device from another. Thus, most of the cost of today's integrated circuits is in the wires; that is, in communication.

In VLSI microcircuits the imbalance will get worse; as the number of transistors increases on a chip, the fraction of the chip area devoted to interconnections also increases, as Mead and Ivan Sutherland of Caltech pointed out several years ago. Mead concludes that a major focus of VLSI design should therefore be on minimizing interconnections not transistor switches, which are becoming negligible in cost as compared to the wires. Mead gives the example of master slices, which he characterizes as "images" of a logic diagram on a silicon chip. What is demanded to minimize interconnections is to alter the design process well before the stage of drawing a logic diagram is reached. Whereas engineers should focus on the creative part of VLSI design (the communication strategy), people work on logic diagrams and the computer does the wiring in master slice design. A complete reversal of priorities is needed.

All in all, the reality appears to be that satisfactory tools for designing the wondrous chips of the next decade do not yet exist. Until they do, VLSI will be unable to live up to its promise of continuing the tradition of integrated circuits that lower the cost of digital computation and thereby bring the benefits of automation and intelligence to myriad applications for which the expense is now too great. The semiconductor industry regards VLSI design as its number one scientific challenge.—ARTHUR L. ROBINSON