Very Large Scale Integration

Designing "Street Maps" of North America

Theoretically, a million transistors (the equivalent of the entire works of a good-sized computer) can be put on a silicon chip one-tenth the size of a postage stamp. In fact, the technology for manufacturing transistors small enough for such Very Large Scale Integration (VLSI) actually already exists. One of its pioneers and most creative developers is Caltech's Carver Mead, the Gordon and Betty Moore Professor of Computer Science.

VLSI is the result of the marriage of the computer and semiconductor technologies, both of which have experienced phenomenally rapid evolution in the last three decades. Computer technology developed out of the pioneering work of John Von Neumann and others in the late 1940s. The semiconductor industry was born at about the same time with the invention, or discovery, of the transistor — a small low-power amplifier — by Walter Brattain, John Bardeen, and William Shockley. It soon became clear that transistors could replace the comparatively cumbersome tubes, resistors, and wires, and do everything a vacuum circuit could do in a computer — store one bit of information or combine two bits to make a logic function. And they could do it using a fraction of the space and energy, which also translated into a fraction of the cost.

In 1960 the microelectronic revolution got going in earnest with the birth of the integrated circuit, so called because silicon, a conductor of electricity that was one of the components of the first transistors, integrates the circuit itself with the technology that makes the transistor. Since conducting layers on the surface of the silicon can be used to interconnect transistors, the silicon can act as its own circuit board.

The first integrated circuit had 12 transistors and did just one of the elementary computing functions. Ten years later circuits had been so scaled down that silicon chips with a thousand transistors were being manufactured. Today several hundred thousand transistors can be put on a chip one-quarter inch on a side — about the size of a thumbtack.

The process of photolithography makes it possible to manufacture integrated circuits. The integrated circuit is built up layer by layer — transistors, wiring, and contacts — into what is almost a three-dimensional architecture. Circuit designers (who work on a larger scale) create masks for each layer. An oxidized silicon wafer, which can contain many chips at one time in a four-inch-diameter space, is coated with a thin film of photoresist (a light-sensitive material). With the mask laid over it, the photoresist is exposed to ultraviolet radiation, which makes it resistant to a solvent, thus leaving the pattern of the mask etched into that layer. Another thin film is then laid on that layer and another pattern photoengraved on it, and so on for several layers.

About ten years ago Mead decided to try to calculate the physical limits of this technology. How small could you make a transistor and still expect it to function in computer machinery? His astonishing prediction was that transistors could be made a thousand times smaller than those being manufactured at the time. Transistors this small — with the potential of a million or more on a chip — have been built since that prediction and proved to function, but there's a problem. They are incredibly complex. Even if the physical limits do allow a million transistors and the accompanying wiring on a tiny chip, how can human beings cope with something that intricate? The complexity of VLSI is so enormous that it completely overwhelms any other difficulties. Silicon chips no longer present problems in physics but rather in computer science — how to make sure such a system works.

Mead often uses an analogy originated by his Caltech colleague, Charles Seitz, to explain just how complex these integrated circuits become as you scale down the size of the wires and transistors and scale up the size of the chip. If you blow up the scale of one of these integrated circuits so that the distance between the wires is equal to a city block, then you can imagine the whole chip as a street map. In the early 1960s, when the chips were about a millimeter across, or forty-thousandths of an inch, with wires about two-thousandths of an inch apart, the chips would have had to be expanded by a factor of four million to attain city-block size between wires. Then you would end up with something like a street map of Pasadena — a small city where it's not too difficult to remember where everything is and how to get around.

The next round of technology (where things were about two years ago with chips five millimeters across) brought something like a map of Los Angeles, where it's a bit harder to remember how to find everything, and even the map itself is rather unwieldy. And yet to come are chips a centimeter across with wires the size of two wavelengths of visible light. This would translate into an urban street

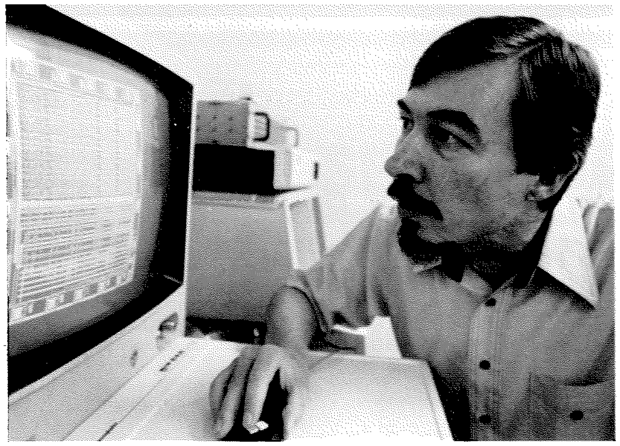

Carver Mead

map the size of California and Nevada. The analogy goes still further. When we reach the physical limits of transistor size that Mead predicted ten years ago, the silicon chip will resemble a city street map the size of North America. Of course, no one has done this yet, but Mead believes there is every reason to expect that it can and will be done.

How would you plan, lay out, and manage a city the size of North America? The key is design, says Mead. And he's been saying it for ten years, having foreseen that problems of managing the complexity were sure to arise. With a chip the size of Los Angeles, industry realized it also.

Until then almost all the rapid advances in integrated circuit technology had taken place in industry. Where were the universities, the "cutting edge," all this time? Most of academia was left far behind, but at Caltech there were some people in the background doing what universities are good at (and industry often is not) — looking very far ahead, doing research whose outcome is uncertain, taking risks. Industry was thinking about immediate costs, not long-range problems.

But while industry wasn't looking, those costs changed. Along with a profound reduction in the overall cost of computation, which is radically affecting society in many ways, technology has also changed the relative costs of the parts of the integrated circuit. In the earlier days of integrated circuits the logic elements were the expensive part; as more and more computing functions were put on a single chip, whose cost didn't change essentially (about $10), the cost of the individual logic functions decreased dramatically. For the same $10 you now have many orders of magnitude greater computing power than a decade ago. Now the costly element has become the wiring — the time and energy it takes to communicate among the increasingly more numerous units of an integrated circuit. Scaling down the size of the wires has also resulted in increased resistance and therefore increased delay.

Another time factor — the time in person-months that it takes to design a very complicated chip — has increased exponentially. And as complexity continues to increase, this design cost will grow out of reach of even the largest companies, Mead maintains.

Industry hasn't changed its method of design much since the time when there were 12 transistors on a chip. Most semiconductor companies design each of the transistors and the interconnections individually, by hand, a process that at the present level of complexity takes many tens of person-years for complicated microprocessors. The end product of what Mead calls the "spaghetti school" of design is an almost impenetrable maze of random wiring. Other firms, particularly the computer companies, have tried to use computers to simplify the process, arranging the transistors in regular rows with wiring laid down on top. Although chips designed in this way look more rational, they do not efficiently solve the problem of keeping communication distance (and time) at a minimum. As chips get larger, the area used for wiring and the time and energy spent sending signals around the chip increases tremendously.

At Caltech Mead and his colleagues discarded the traditional design methods and started from scratch to restructure the way integrated circuits are designed — to devise a new approach that would exploit the potential of VLSI and would cope with the complexity by using orderly, simplified floorplans, keeping interconnection paths as short as possible to save time and energy consumption. Essential to their approach is the concept of locality — placing elements that communicate with each other (both logic and memory elements) close together so that their messages don't have to travel back and forth over "long" distances across the chip. To accomplish this, Mead's group has developed hierarchical design — progressively splitting the whole system into smaller, simpler parts, or modules, that are independent of each other and of the whole system and that communicate with each other only at well-defined points. As the name implies, hierarchical design approaches the problem from the top down, like a reverse tree (indeed trees and leaves are designations in many of the design ideas devised at Caltech).

Mead's structured approach applies hierarchical design to the particular constraints of VLSI systems — implementing the design in the many-layered construction of the actual chip, placing modules with similar functions next to each other in regular patterns like a tiled floor. Algorithms also simplify the design task, and computers help to determine the optimum arrangement of modules, making it several orders of magnitude simpler than the traditional design methods.

One particularly successful system, devised by graduate student Dave Johannsen, makes a chip design possible in a few minutes rather than in person-years. It's a silicon compiler, a computer program that performs most of the implementation computation to turn out a mask set for all the various layers of the integrated circuit. The designer can specify what functions the modules, or blocks, are to perform, and the computer does the rest — figuring out the circuitry within each block and between the blocks. (The program is called "Bristle Blocks" from the appearance of a rectangular module with interconnections sticking out all over it.) With the Bristle Block program a designer can, in effect, design a chip by moving around the building blocks, leaving the complicated interconnections to the computer.

With an earlier integrated circuit technology, there were microprocessors (the arithmetic and logic units) and there were memories — separate functions on separate circuits, communicating by a "bus" (usually a cable of wires providing common transportation for data). VLSI makes it possible and efficient to have the two functions together on one chip — many processors and many memories or many processors with a common memory. This leads to the possibility of concurrent processing — lots of calculations going on at once instead of in sequence, the way current computers work. Mead and his colleagues have developed a number of design patterns, including various arrays and trees, to facilitate the fewest and shortest possible interconnections among the units.

The Caltech group is not trying to hide its ideas. In fact, to emphasize the importance of partnership between industry and academia and avoid repetition of the early years of integrated circuits when industry was preoccupied with the immediate future, Caltech's Silicon Structures Project involves a number of industrial sponsors in a working relationship. These include IBM, Xerox, Burroughs, Hewlett-Packard, Digital Equipment Corporation, Intel, and Honeywell, with more on a waiting list for the informal "think tank." Each of the companies sends a scientist to Caltech for a year to work on design ideas and methodologies with Mead's group; thus the participating firms have contact with really innovative research, and the Institute, in return, gets a better understanding of industry's problems.

Another factor that is returning universities to leadership in integrated circuit technology is Mead's original VLSI course, developed at Caltech. Out of that course came the only textbook in the field, *Introduction to VLSI Systems*, written by Mead and Lynn Conway of the Xerox Palo Alto Research Center. This extraordinary course has already been adopted at MIT, Stanford, Berkeley, Carnegie-Mellon, Washington University-St. Louis, USC, UCLA, and the universities of Florida, Washington, Illinois, Rochester, Utah, and Colorado. Reflecting Mead's simplifying approach to design, the course also applies this simplification to instruction, providing the minimum of basic information about fabrication technology, logic design techniques, and system architecture. By limiting instruction to the key concepts, from the underlying physics to the complete VLSI systems, and eliminating all the rest of the "mental baggage," the course is turning out designers at a surprising rate — and they can walk right out of the classroom and start to work.

One reason for this quickly acquired skill is the "learning by doing" feature of the course. Students work on projects involving architecture, design, layout, and testing of real integrated circuit systems that are then actually manufactured — the chips of a whole class on a single silicon wafer. Originally the Caltech class had to beg for space on commercial lines to get student designs produced. Now, however, at Caltech's instigation the Advanced Research Projects Agency (ARPA) has funded a fabrication plant expressly for the innovative designs coming out of universities.

The ARPA-funded silicon "foundry" is a prototype of what Mead considers essential to the future of the industry — the division of labor between the designers and the fabricators, between product creation and product replication. While there are still hard problems to be solved in the technology of processing, the outcome is predictable; industry knows what has to be done and knows that it can be done. Where the startling advances will come now is in the area of design. And the sophisticated developments in design are coming out of the universities and small new firms that have no access to manufacturing. They have insufficient capital to begin making their own chips because fabrication has reached such a capital-intensive stage.

Mead envisions the semiconductor industry in the future with an analogy to writers and printing companies. Designers should create the circuits, and other firms (the silicon foundries) would "print" them, since, as in printing, an unlimited number of system designs can be reproduced by a single process. Only if access to these foundries is provided by well-capitalized firms can the high level of innovation in computer electronics continue.

Here again Mead's streamlining and standardizing approach to design proves necessary. If the rules of the game are simplified and well defined, designers and manufacturers will have a "clean interface" (with requirements of geometric design rules, standard data format, and standard test chip) and will be able to communicate even though their functions are separated. If the designers can generate the complete layout for the chips (and by the new method they can), then the only information that must be transferred to the "printer" is the patterns representing the various layers. Despite the close cooperation of Institute and industry in the Silicon Structures Project, industry has not wholeheartedly embraced all of Mead's revolutionary ideas. It has been to some extent unwilling to trade off some of the things that Mead's approach demands (for instance, maximum number of transistors on a chip) for a simpler design method, and the large semiconductor companies are not enthusiastic about processing competitors' designs.

Although both the computer and semiconductor industries have been heading toward the same goal, neither has adjusted to the innovations of the other. But there is no question that eventually both are going to have a totally different structure that will require working together. Mead sees the universities, which stress the underlying unity of what sometimes seem to be disparate disciplines, as marriage brokers. Arranging the marriage has been left to the universities, and — although industry may not yet recognize it — if Mead is right, the ceremony has already taken place. □