Are We Really Ready for VLSI^2?

Gordon E. Moore
Intel Corporation

A tremendous interest in VLSI is all around us. There is much talk of electron-beam and X-ray lithography tools to achieve VLSI's submicron structures. In all of the VLSI discussions, the implication is that it will allow us to keep on enjoying the same kind of fantastic low-cost advantages previous IC technologies have provided us in electronic products. Perhaps this may become true, but if the semiconductor industry had a million-transistor technology like VLSI, I'm not so sure it would know what to do with it. Besides products containing memory devices, it isn't clear what future electronic products that take advantage of VLSI will be.

Examples abound of products with decreases in cost from 10 to 100,000 fold, made possible by progress in semiconductor integration levels. Each increase in integration level has opened up new applications, and in several instances developed completely new industries. As semiconductor device technology evolved from discrete, to small-scale, to medium-scale, and through large-scale integration levels, product advantages have multiplied. Doesn't it seem a matter of straightforward calculation that an order-of-magnitude increase in IC device complexity should result in many of the same product advantages? Perhaps, if the products are memory related.

Memory is certainly one function that can be used in large chunks, assuming that the cost/bit will be low enough to make this possible. Single-chip microcomputers could be extended with more memory on the chip. But even here, memory modularity at some size becomes important, thus limiting the amount of memory usefully incorporated on chip.

Beyond memory, I haven't the slightest idea on how to take advantage of VLSI. In fact, the semiconductor industry is not now process-technology limited for non-memory products. How to best make use of the processing technology is really what the problem is.

Criteria for Success

Several things are required to produce a successful product, and processing technology is only one of them. (Successful product means a product that can be sold at an acceptable price to both maker and user.) Figure 1 illustrates the process of creating an LSI IC product. Each of the blocks in the figure is made up of a number of complex factors. For example, the "design" block includes the design of the process as well as that of the product. Process design requires a description of the processing sequence, the layout rules, and the electrical description of the elements of which it is composed. Product design of a complex structure requires logic and circuit designs, mask layout, and design verification. Any one of the aforementioned factors can be a formidable barrier.

At some point in the past, each of the blocks listed in Figure 1 had been a limiting factor in the success of semiconductor devices. For instance, during the first decade of the transistor, the main limitation in its successful implementation was no less than processing technology. The technology for diffusion and for making contacts had to be developed to make the transistor a reality, a device whose electrical requirements were fairly easy to define. Similarly, in the early days of the integrated circuit, processing technology was also the limiting factor to success. Features such as isolation structures had to be developed to make the IC a reality. Probably, the classic case is that of the insulated gate field effect transistor -- a device which a group at Bell Laboratories was trying to make when in the process they got hung up on something called surface states, thus leading to the invention of the point-contact and junction transistors. It wasn't until 15 years later that the semiconductor industry learned how to manufacture a stable MOS device, and even later before it understood why.

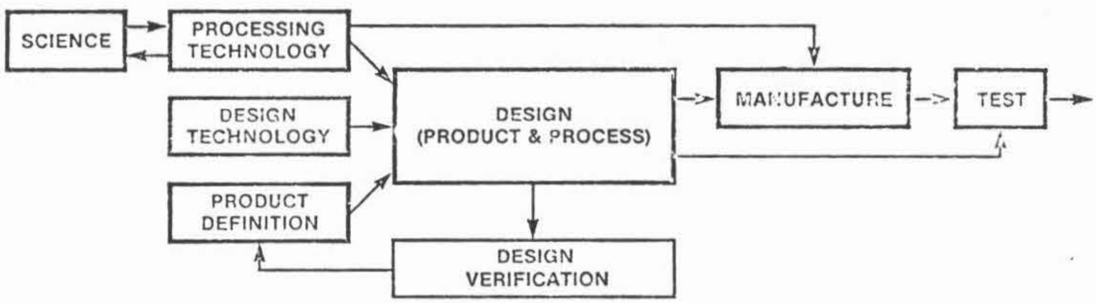

Figure 1.

Once the basic process steps were in place, progress in making ICs in ever more complex structures moved along rapidly (Figure 2), in an exponential fashion. The curve in Figure 2 is essentially the envelope of IC complexity growth. Points indicated in the figure are a sprinkling of the most complex circuit types available commercially at the time indicated. Most of the circuits introduced fall well below this curve. I expect a change in slope to occur at about the present time. From the doubling of the curve annually for the first 15 years or so, the slope drops to about one half its previous value, to a doubling once every two years. This is the rate of complexity growth than can be predicted for the future.

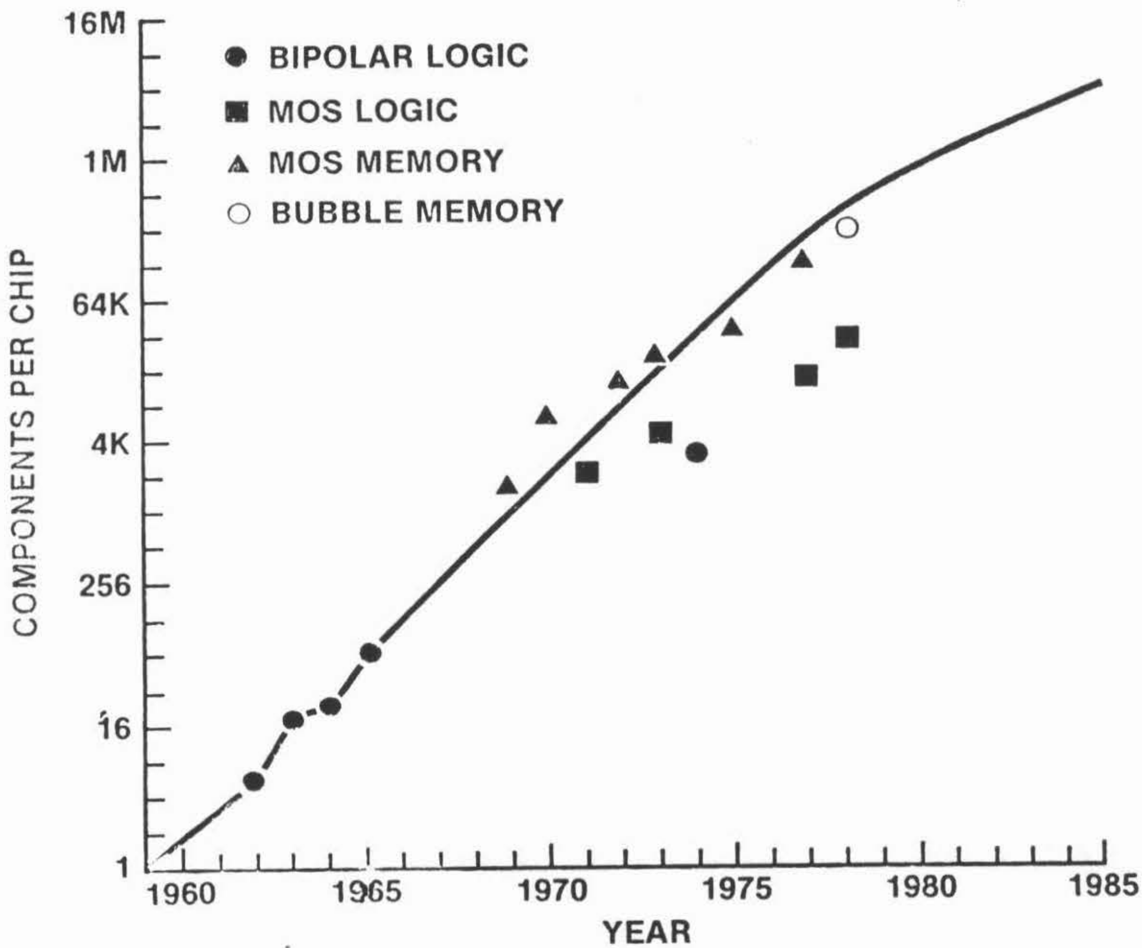

Figure 2.

The projected slowdown in IC complexity growth is caused by the semiconductor industry's loss of one of the principal factors that has allowed it in the past to increase complexities: the ability to pack more and more elements on a chip's surface by eliminating non-functional chip areas. The latest IC devices indicated on the graph of Figure 2 represent the densist ICs with the smallest amount of non-functional areas on their chips.

A Repetition of Earlier Problems

Note the gap between 1965 and 1968 in Figure 2. This gap existed because it was difficult at the time to identify any semiconductor products whose complexity came close to the limit of the time. This condition did not arise out of a lack of effort (in fact, this was a period of intense activity), but out of a problem of product definition, the very same problem the semiconductor industry is now facing as VLSI technology comes into existence. It was difficult at the time to define semiconductor products that fit the criteria for success and were near the limits of device complexity.

Two major problems faced the semiconductor industry then, as it tried to partition digital systems into complex blocks: interconnections and product uniqueness. The former problem arose from the fact that the number of leads for circuit increased so rapidly with the increase in circuit components that it went well beyond the packaging capability of that era. The latter problem resulted because the blocks tended to become unique with a resulting explosion of different part types, each required in small quantities. This condition was not conducive to successful semiconductor products.

Thus, a crisis of product definition existed. The semiconductor industry was unable to define products of high complexity that were useful in sufficiently large numbers of applications to justify their designs, and that were packagable with the available technology.

A variety of attempts to solve the problems were explored. Computer designers were asked to partition their systems into functional elements to minimize the interconnection problem. Efforts were made to confront the parts-number explosion directly. I remember at that time having discussions on how to design, manufacture and test several hundred new part types every week, in volumes of perhaps only 10 to 100 of each type. Several techniques evolved with approaches that today might be called gate arrays, wherein customized layers of metal interconnections were used on standardized diffused wafers.

The powerful computer design aids required to handle the large number of part numbers were slow in coming. Only recently have successful results been obtained. For example, IBM recently described a fantastic system utilizing direct-electron-beam writing on the silicon wafer, and a highly automatic line to handle the problem of making small quantities of a very large number of different IC designs.

In general, such efforts to solve the semiconductor industry's problems of the 1965-1968 era were not successful. The product definition crisis persisted and limited IC complexity through the mid sixties. Two things broke the crisis for the semiconductor component manufacturer, though not necessarily for the mainframe computer manufacturer; the development of the calculator and the advent of semiconductor memory devices.

The calculator was a simple system that could be partitioned into about four 40-pin IC packages, making the interconnection problem tractable. Since it was made in large quantities, sufficiently large quantities identical of components used within the calculator were manufactured to justify design costs.

As for memory, it is a universal function that can be used at the highest level of integration available. With the use of on-chip decoding, the number of leads was reduced to match available packages. What remained was for semiconductor memory to be cost competitive with established technologies for it to blossom.

Thus, the interconnection and product definition problems of the past were not necessarily solved. They were simply circumvented. The semiconductor industry developed a different set of markets in which it could keep itself busy, postponing the solution of its previous problems.

The MicroProcessor Smooths the Way

Just as the calculator and memory enabled the semiconductor industry to continue making more complex devices for certain applications, the microprocessor extended the range of use. With its general purpose architecture one could program the microprocessor to perform in a wide variety of applications providing a solution for the product definition problem.

Thus, during the 1970s, the semiconductor industry kept developing more complex memory chips to track the complexity curve in Figure 2, with microprocessor products following closely behind. Large-computer manufacturers were left to solve their own problems of part number proliferation and low-volume uses, often through the use of components with lower levels of integration. Thus modern LSI technology has not eliminated predecessor technologies of small-scale and medium scale integration. For example, the number of bipolar semiconductor devices produced continues to grow rapidly, from about 850 million circuits in 1972, to about 1.5 billion in 1974, down to a little over 1 billion during 1975-1976, and up again to about 2.5 billion last year, worldwide. The availability of high levels of device complexity has not resulted in the complete replacement of less-complex devices. Co-existence is more often the case. Even a company devoted to making LSI IC products finds that it cannot use the capability for complexity in all its products. The complexity of products introduced by the Intel Corporation, for example, over the last two years is shown in Figure 3, and can be compared to the limits of Figure 2.

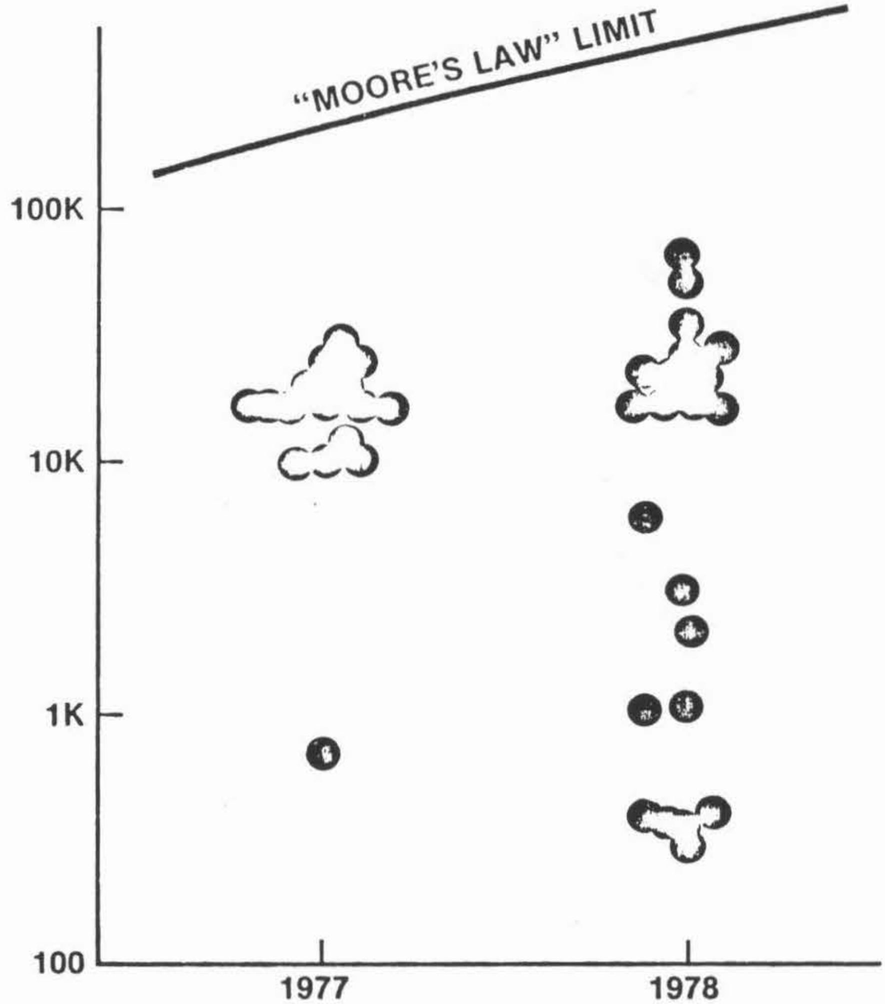

Figure 3: Complexity of Intel's Semiconductor Product Introductions for 1977 and 1978

Note that few of the products depicted in Figure 3 are close to the "Moore's Law" limit of the same figure, many of which miss it by large factors. The most complex circuits tend to be memories, with simpler ones being microcomputer peripherals.

In Figure 3, microprocessor and complex peripheral devices tend to group around the same level of complexity. This is the level that the semiconductor industry can presently define for useful products. Although similar devices two to three times more complex can be made, a definition of the products they would constitute is needed first. Thus we come full circle to our dilemma: how to best make use of our capability for ever more complex devices such as VLSI ICs, by properly defining such products.

Another Perspective

The product definition problem can be shown from a different perspective, by looking at the amount of effort required for product definition, design, and layout (in person-months), starting with the first planar transistor of 1959 and projecting into the future (Figure 4). This design effort is plotted on a logarithmic graph

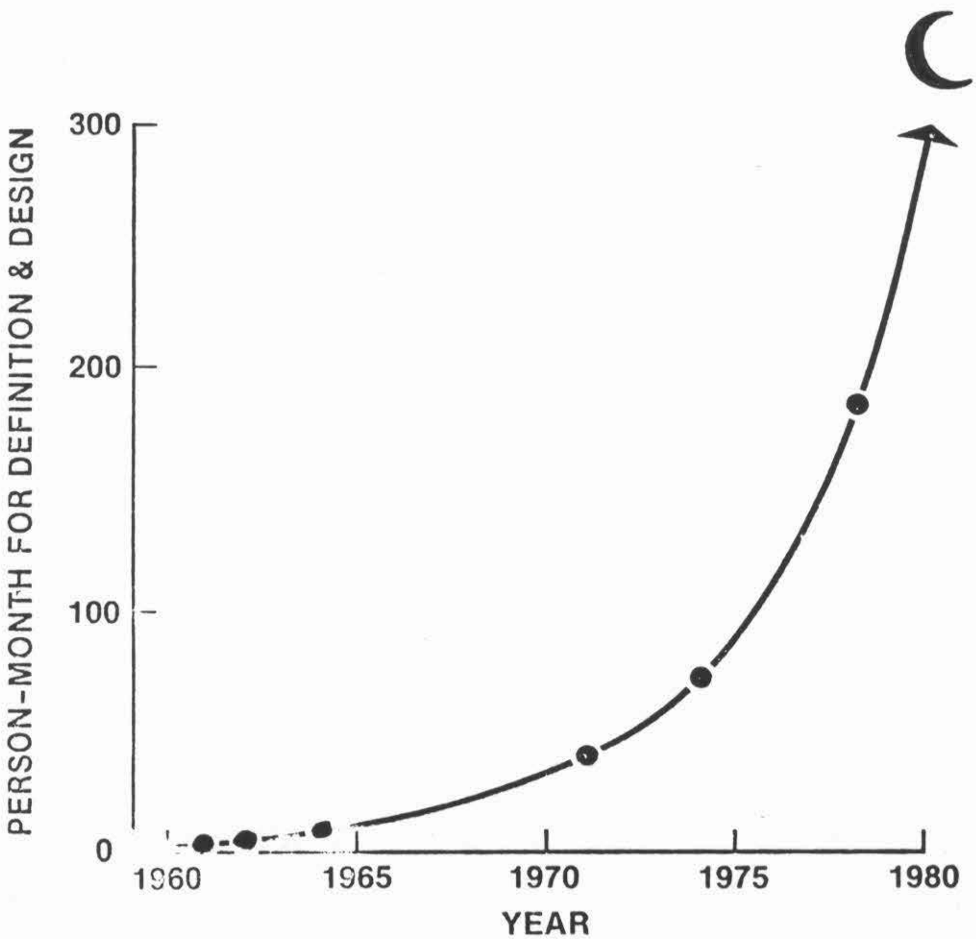

Figure 4.

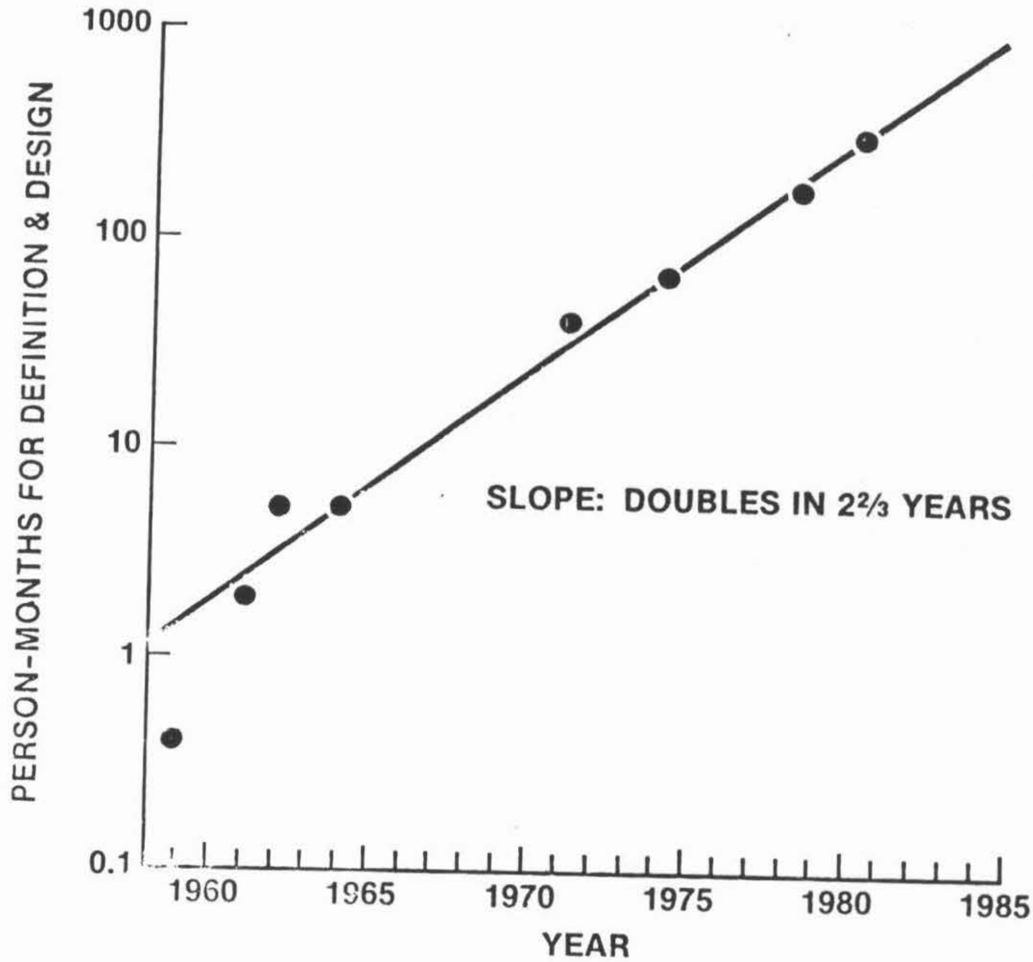

Figure 5.

in Figure 5. As can be seen from Figures 4 and 5, its growth is exponential, doubling every 2 and 2/3 years.

If it is assumed that the cost per person-month is inflating at 10 percent per year (a conservative figure considering the need for increased computer support, etc.), then the costs double every two years. We should keep in mind that device complexity is also doubling every two years, resulting in a constant cost per element to define, design, and layout complex ICs.

This cost can be contrasted with the manufacturing costs that are approximately independent of device complexity. Whereas once manufacturing costs were dominant and exceeded those of design, the situation is now reversing, with design costs becoming dominant. The implication is clear: product definition and design technology are where work is really needed. And the kinds of answer the semiconductor companies will come up with in response to these challenges will depend on the nature of their businesses.

The component supplier must have large markets across which he can amortize his high design costs. This requires high-level standardization, either at the processor level or at the very large system level. This will limit the breadth of VLSI's impact as shown in Figure 6. Only memory devices may utilize maximum complexity. Discrete devices, MSI and LSI logic functions, and LSI will remain important in future systems.

The principal capability for defining and designing LSI and VLSI products is in the hands of the systems suppliers. If product definition and design will become the important factor of the future and I believe that it will, then the systems companies may have the advantage in VLSI's success. They also have the desire to preserve existing structures such as large cumulative software investments.

The result is that a structural change is occurring in the semiconductor industry. On the one hand, component suppliers, as always, are pushing for standard products that are useful in large numbers across a broad spectrum of applications. On the other hand, an increasing number of systems companies or captive suppliers are becoming more skilled in the technology of making complex ICs. Such companies are expanding their in-house processing capabilities and are using them successfully. A few years ago, I maintained that there were only two successful captive suppliers in the world. Today, there are clearly many more.

According to the most recent compilation from Dataquest Corporation, the number of worldwide component suppliers in the semiconductor industry between 1975 and 1979 dropped by about

10 percent. On the other hand, system companies with in-house captive suppliers -- not simply R & D laboratories, but companies making products for use in their own equipment -- grew from 19 to 43 during the same period of time. Clearly, the industry is changing.

As for my original question, whether or not the semiconductor industry is ready for VLSI, the conclusion is that for maximum advantage, both suppliers of components and systems must address the problems of product definition and design. In fact, unless we address and solve these problems, as we look back on the VLSI era, we may only be able to say, "Thanks for the memories."

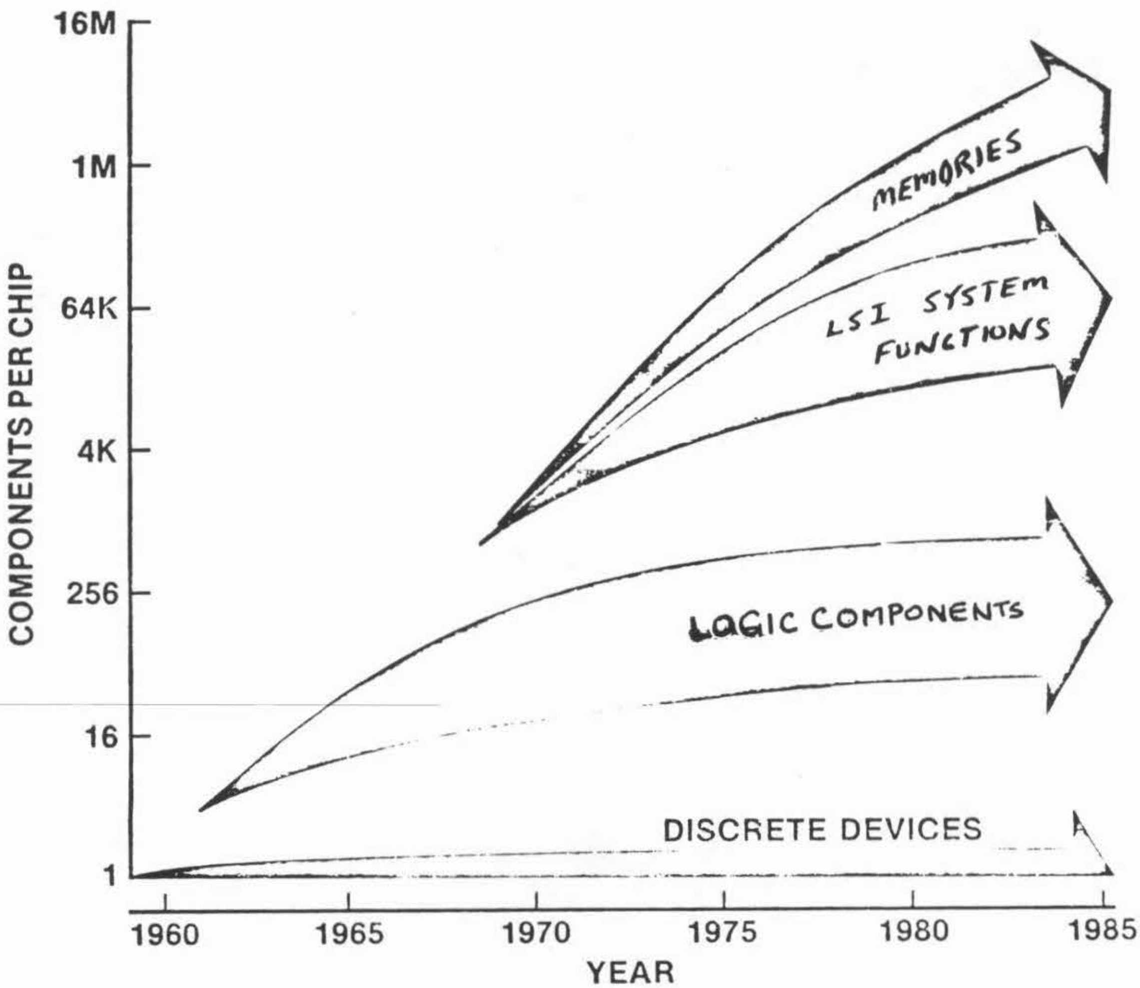

Figure 6.