VLSI AND TECHNOLOGICAL INNOVATION

Carver A. Mead

Professor of Computer Science,
Electrical Engineering and
Applied Physics

California Institute of Technology

VLSI relies on a range of disciplines for its successful implementation. Two of the most important of these are still in their infant stages.

A. Design methodologies to manage complexity.

B. Architecture of ultra concurrent machines.

Innovation in infant disciplines occurs most rapidly and successfully when a large number of small groups proceed independently under the motivation of market opportunity. In a few years, a substantial fraction of the engineering work force will have a working knowledge of LSI design. At the same time, fabrication areas are becoming more and more capital intensive. What is needed is a clean, standard interface between a multitude of small diverse VLSI design groups and a few state-of-the-art fabrication suppliers. A proposal for such an interface is presented in this article.

[Note: This article elaborates on only one of the several topics in Prof. Mead's talk at the conference. --ed]

The electronics and computer industry of the future will look radically different than it does today. Using the past as a guide, we can guess with reasonable certainty the course of future evolution. Figure 1 shows the evolution of the various components of today's silicon manufacturing business. At the bottom are the discrete transistors, diodes, rectifiers, etc. They still form a substantial fraction of the entire semiconductor business. Above them lies the small scale, medium scale and large scale standard parts integrated circuit business. This business, dealing in large volumes of standard catalogue items, will always exist and in fact will grow in the future. Riding above it, however, is a rapidly increasing segment dominated by VLSI designed by those who will take it to the end user market. This is the true world of VLSI. It will not compete directly with the other branches of the semiconductor industry just as memory manufacturers do not compete with rectifier manufacturers.

VLSI is a statement about system complexity, not about transistor size or circuit performance. VLSI defines a technology capable of creating systems so complicated that coping with the raw complexity overwhelms all other difficulties. From this definition, we can see that the way in which the industry responds to VLSI must, in fact, be different from the way it has historically evolved through its other phases.

The complexity scale implied by the new technology can be appreciated from the analogy presented in Figure 2 (1). At several points in the evolution of the technology, a typical chip has been scaled up to make the spacing between conductors equal to one city block. The circuit can then be thought of as a multi-level road network. In the mid 1960's, the complexity of a chip was comparable to that of the street network of a small town. Most people can navigate such a network by memory without difficulty. Today's microprocessor is comparable to the entire Los Angeles basin. By the time a 1μ technology is solidly in place, a chip design will be comparable to planning a street network covering all of California and Nevada at urban densities. The ultimate $\frac{1}{4}\mu$ technology will be capable of producing chips whose complexity rivals an urban network covering the entire North American Continent.

Designers are just now beginning to face complexity as a central and dominant issue of the next stage of evolution. They have not yet begun to face the capabilities that such a technological revolution brings to us. The evolution of the component fields which make up the present VLSI discipline are shown schematically in Figure 3. What is plotted here is the number of new ideas, weighted by their importance, as a function of time. Each component discipline undergoes a period of exponential growth when each new idea spawns several others. Later, a period of linear growth ensues while the interstices between the fundamental ideas are being filled in. Later, a logrithmic law ensues in which ideas are being ground finer and finer but very little conceptually new content is added. By now, the number of dramatically new ideas being added to the device physics area is small. Fabrication technology has essentially all the fundamental knowledge that will be required. Circuit and logic design have some cleverness left but that too will soon saturate. The large system design methodology is still

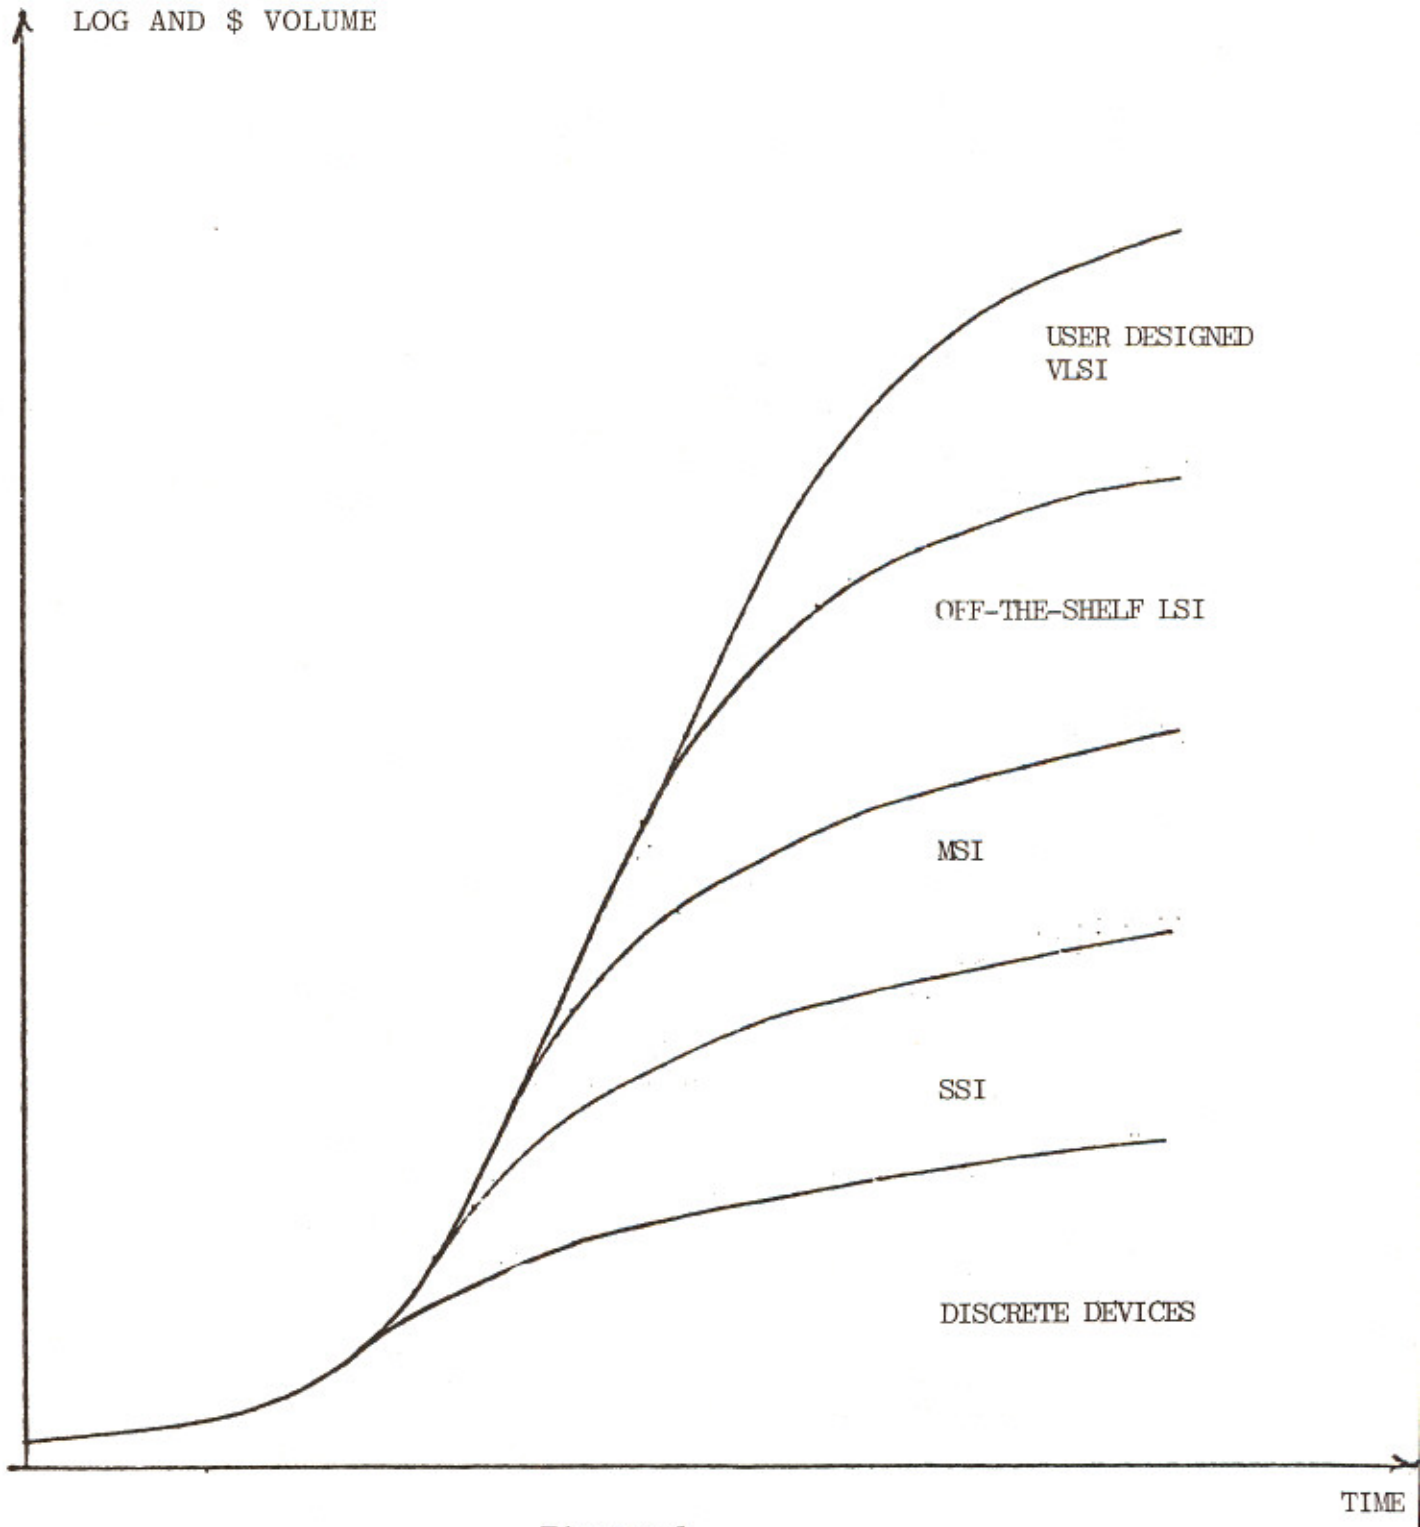
LOG AND $ VOLUME
USER DESIGNED VLSI
OFF-THE-SHELF LSI
MSI
SSI
DISCRETE DEVICES
TIME

Figure 1

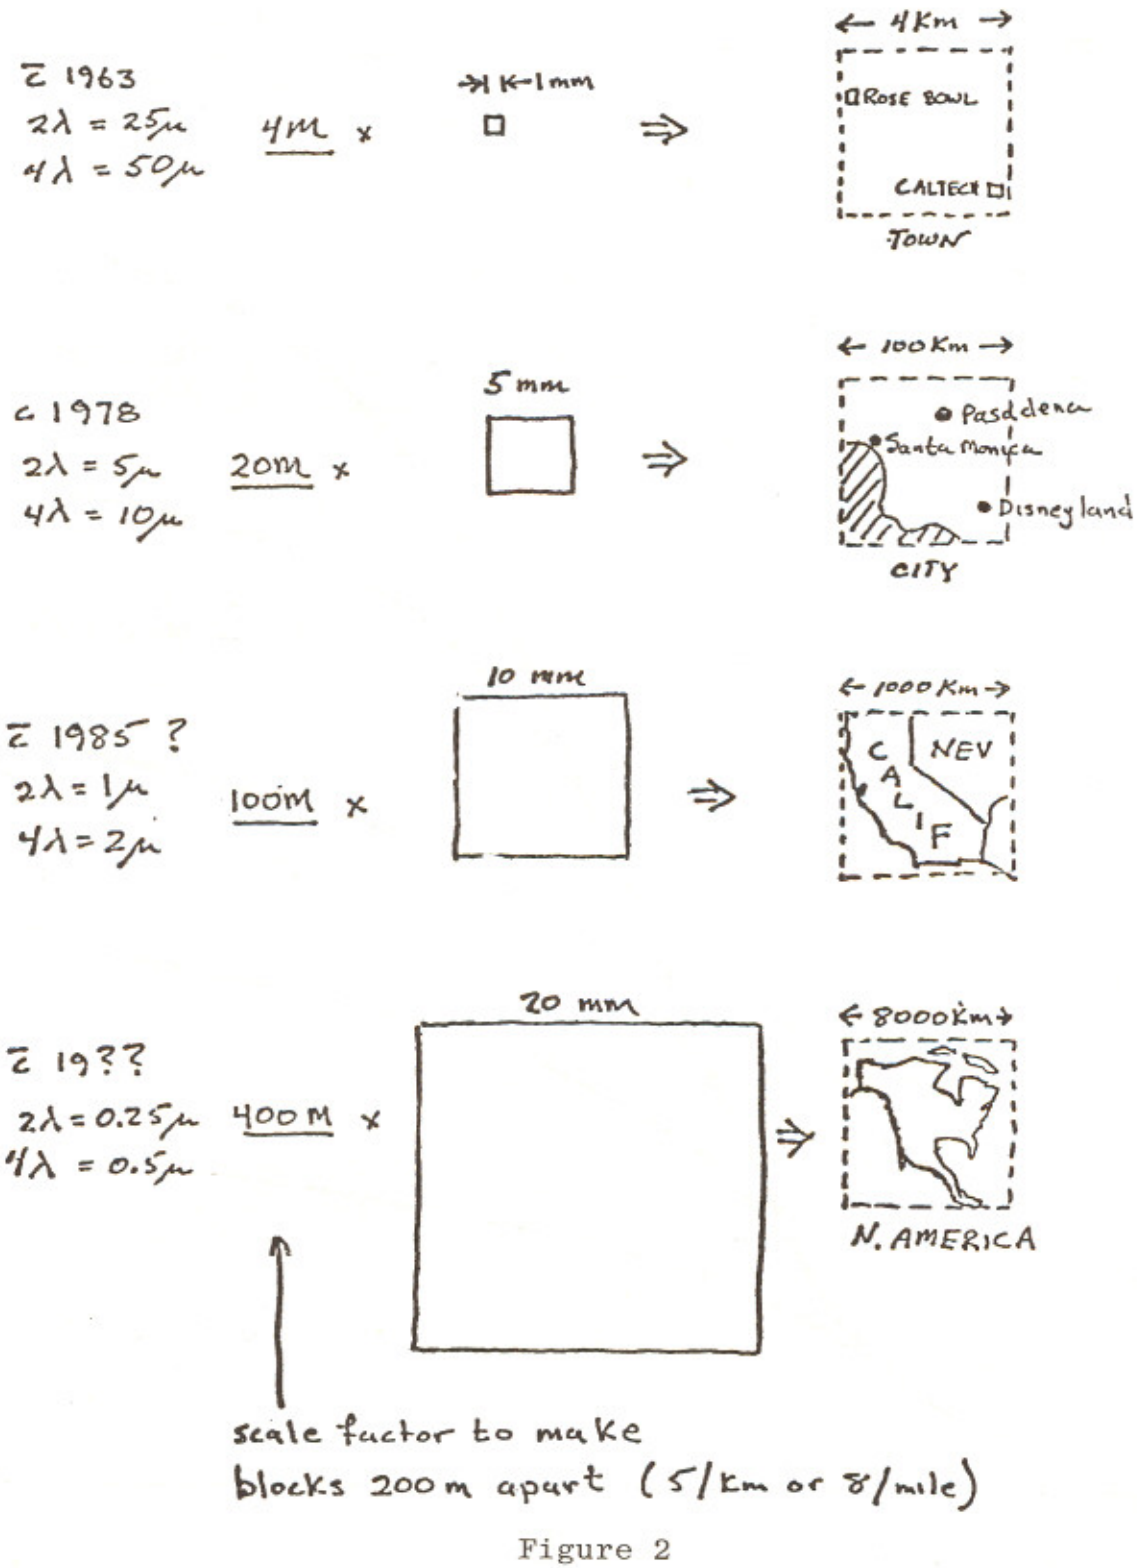
~ 1963
2λ = 25μ
4λ = 50μ
4M x
1K-1mm
4Km
ROSE BOWL
CALTECH
TOWN
c 1978
2λ = 5μ
4λ = 10μ
20M x
5mm
100Km
Pasadena
Santa Monica
Disneyland
CITY
~ 1985 ?
2λ = 1μ
4λ = 2μ
100M x
10 mm
1000Km
NEV
CALIF
~ 19??
2λ = 0.25μ
4λ = 0.5μ
400M x
20 mm
8000Km
N. AMERICA
scale factor to make
blocks 200m apart (5/km or 8/mile)

Figure 2

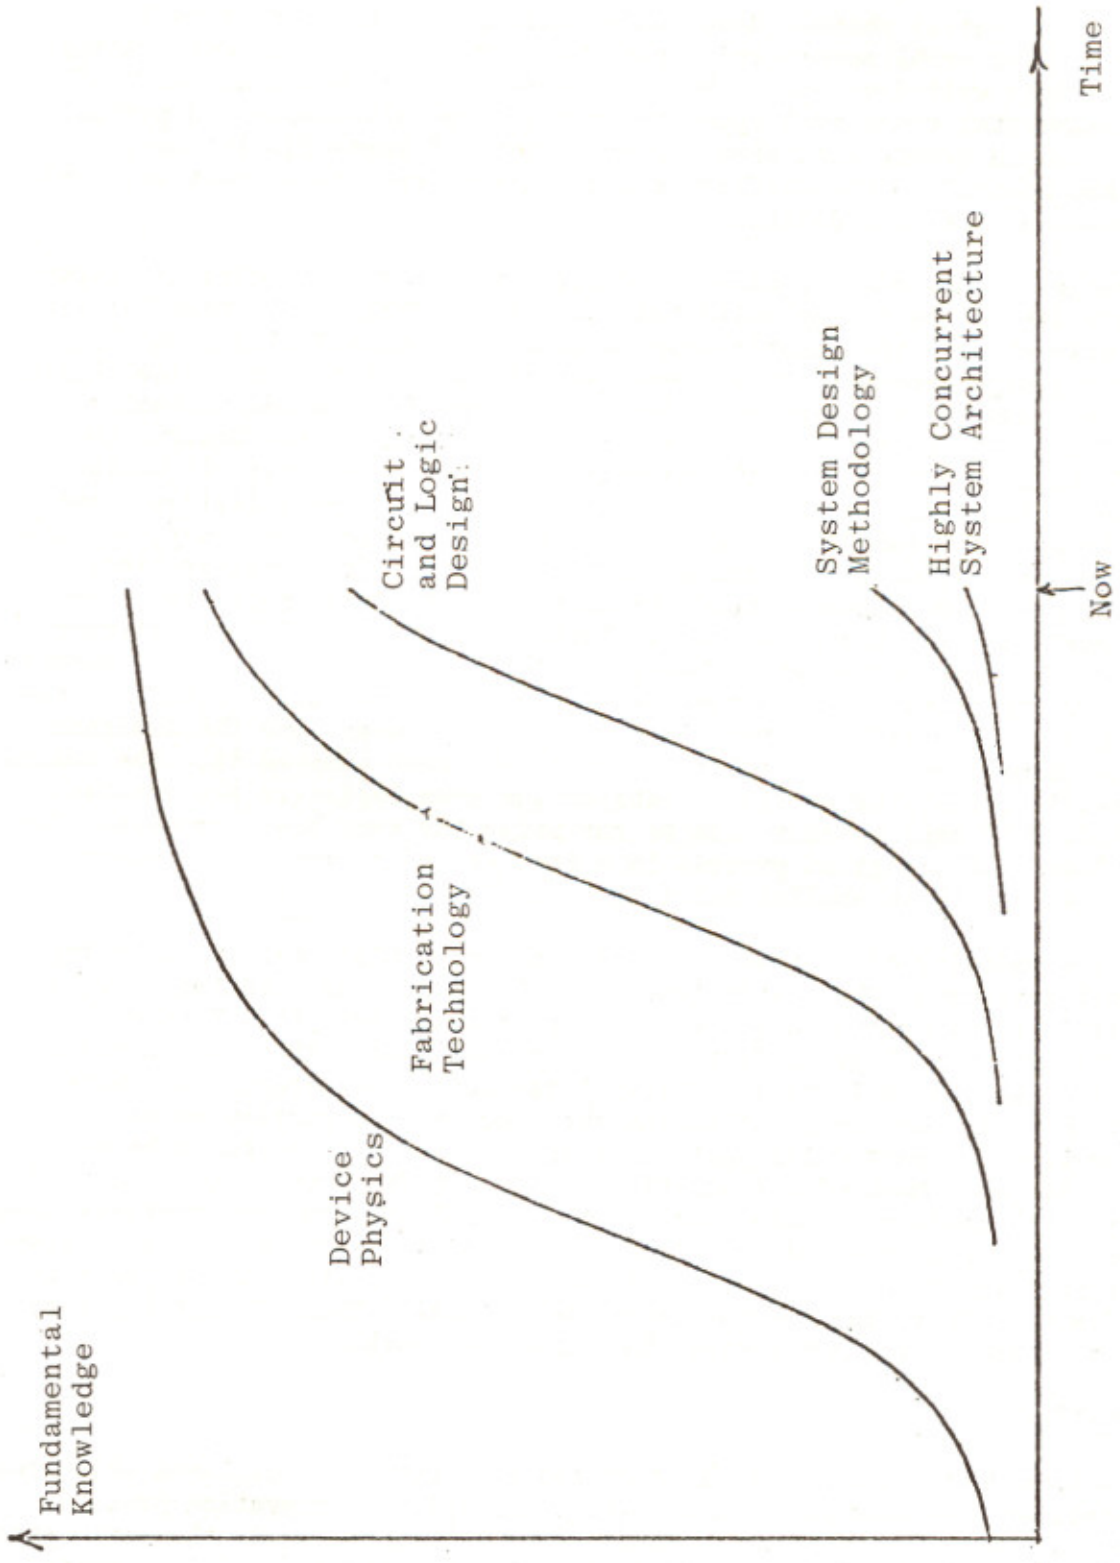
Fundamental Knowledge
Device Physics
Fabrication Technology
Circuit and Logic Design
System Design Methodology
Highly Concurrent System Architecture
Now
Time

Figure 3

in its exponential phase. Many fundamental ideas have yet to be discovered. The architecture and algorithms for highly concurrent systems is even less well developed. Only a few results are known and much of the fundamental conceptual apparatus needs to be discovered. A period of very rapid growth lies ahead of us in both of these disciplines. They are central to the difference between VLSI and the current way semiconductor devices are designed.

The range of knowledge required to design integrated circuits has expanded greatly as their complexity has increased. This requirement has dramatically changed the relationship between the manufacturing technology and the design process. In the earliest times, getting the device physics right was most of the problem. The physics and fabrication technology were intimately entertwined. One person could oversee the design, the manufacturing process, and the testing as well. Later, circuit design became as important as the device design, but still one individual could work between the two disciplines. In many linear circuits today, a different process is used for each product, a heritage from earlier times. However, there has been a steady trend toward standard processes. The main early driving force was the evolution of families of logic elements such as TTL, ECL, CMOS, etc. Here, the designer could implement a number of logic functions with the same process. The trend toward standard processes and a simplified interface between that process and the designer has had many beneficial results for those who have adopted it. The design process is greatly simplified. Fabrication area logistics are greatly simplified if many products can be run using the same basic process. The maintenance of any given process is a complex and tedious job. Fewer processes result in smaller maintenance problems.

As the complexity of systems increases, the potential gain in achieving optimal designs at the system level greatly outweighs advantages to be achieved by customizing a process to a particular product. Even with today's LSI technology, a factor of a thousand to ten thousand is available if ways can be found to achieve large scale concurrency for system functions. By contrast, optimizing the process to a particular part or the design of a particular part to a specialized process may achieve a factor of two. Much of our experience with the development of software is directly applicable here, since both disciplines are fundamentally concerned with management of very large, very complex systems. A hard lesson has been learned in that arena; get the design correct at the highest level and don't yield to the temptation to suboptimize. There is nothing more useless than a very fast system which does not work.

Innovation

The semiconductor technology is composed of a set of disciplines which must be considered separately. In any given discipline, innovation proceeds along an S shaped curve such as that shown in Figure 4. In the early phases, marked (A) in the figure, progress is limited by the lack of fundamental ideas. A single good idea can make possible several other good ideas and hence the innovation rate is exponential. During this period, a single

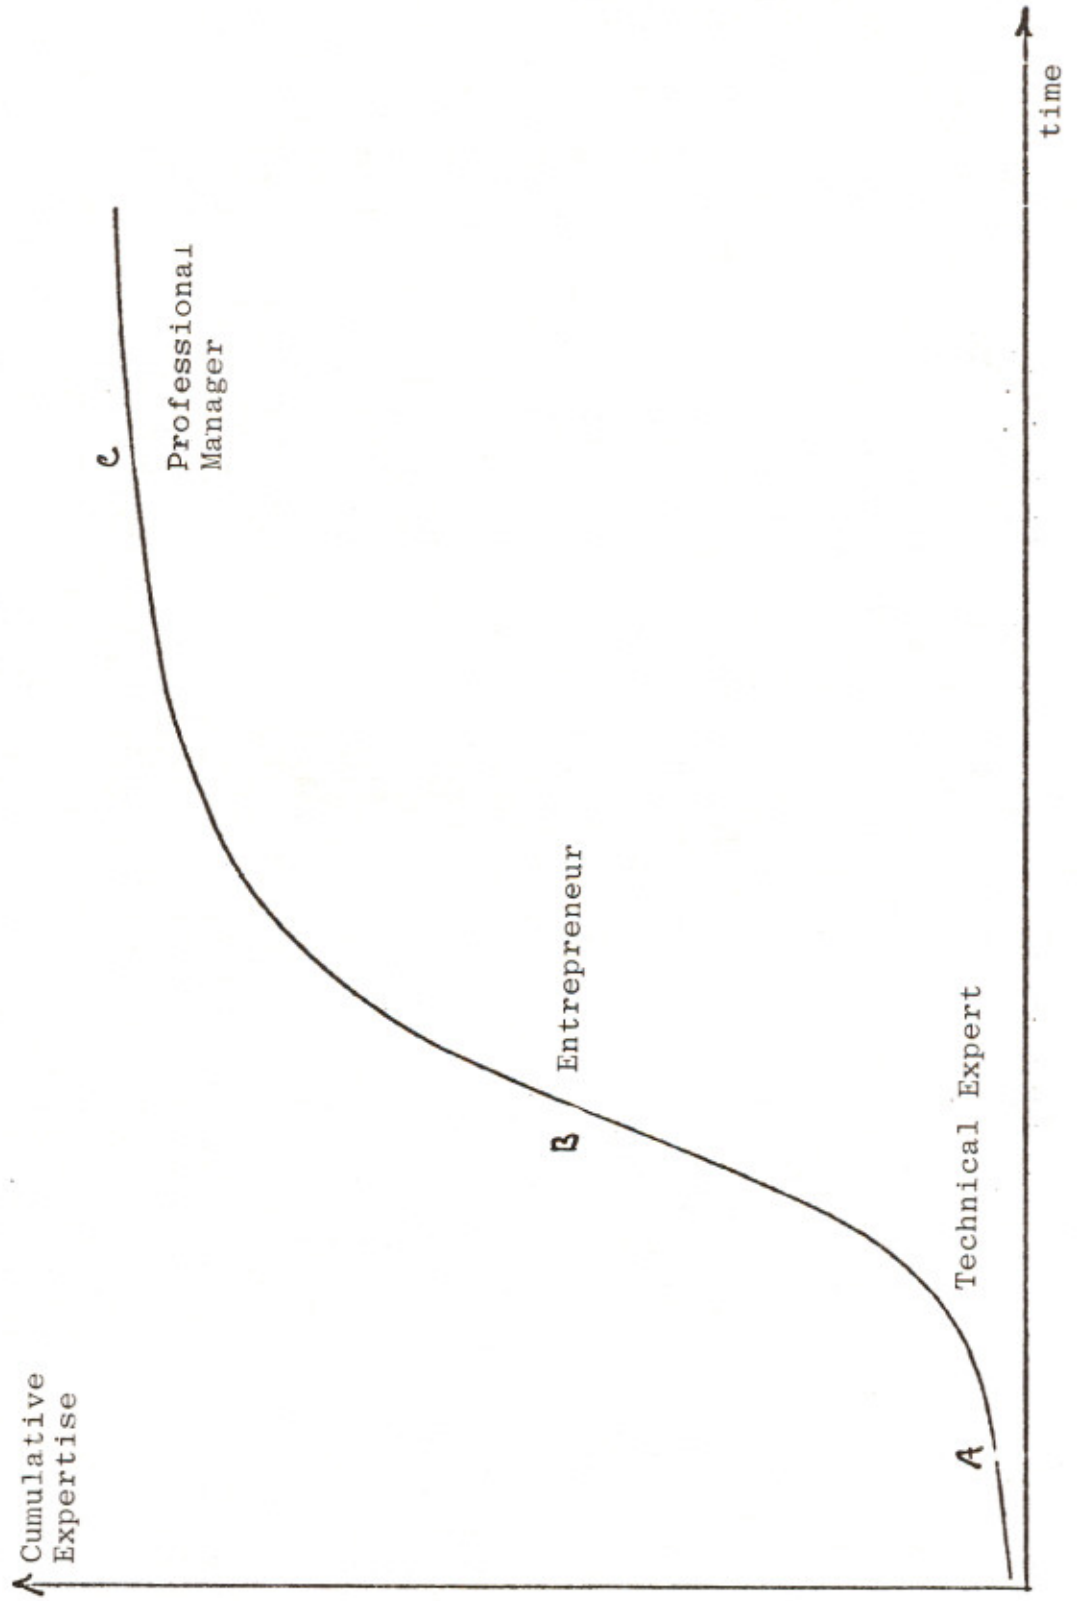
Cumulative Expertise
time
A
Technical Expert
B
Entrepreneur
C
Professional Manager

Figure 4

individual or small group of individuals can develop a viewpoint and contribute several crucial insights that set a field in an entirely new direction. It is the time during which progress is dependent upon a few visionaries within the field. During the central and most visible portion of the evolution, marked (B) on Figure 4, a linear region ensues. Here, the fundamental ideas are in place and innovation concerns itself with filling in the interstices between these ideas. Commercial exploitation abounds during this period. Specific designs, market application, manufacturing methods grow rapidly.

The field has not yet settled down at this point. Entrepreneurs backed by venture capital firms can have a large impact and achieve a dominant market share during this period. During the later stages of the evolution curve, marked (C) in Figure 4, progress becomes logrithmic in time. Manufacturing methods are refined ever further. More and more capital is expended to reduce the price of manufacturing. Here the business becomes capital intensive. Production know-how and financial expertise are the required credentials. Professional managers and large firms dominate the business.

Innovation proceeds most effectively in a large number of small groups. The problem faced by the semiconductor industry is apparent. Fabrication technology has reached its capital intensive phase. Design is still very early in its exponential phase. Historically, innovation in the industry has been spearheaded by small start-up firms and later taken up by large existing organizations. It is significant that the major suppliers of vacuum tubes did not become the major suppliers of transistors. The major suppliers of discrete transistors did not give us semiconductor memories. More recently, companies dominant in the semiconductor memory business did not bring us the multiplexed address random access memory. The microprocessor did not come from mainframe or minicomputer firms. Each of these innovations was brought to market fruition by a small start-up firm which rapidly gained market share by virtue of its innovation. Existent dominant firms were then forced to retrofit these ideas into their own product lines.

Each small group can no longer afford its own fabrication area. A start-up firm with a capital budget of one or two million dollars for a fabrication area was within the means of traditional venture capital sources. However, the same is not true for capital budgets of several tens of millions of dollars required for state-of-the-art fabrication lines in the near future. If innovation by a miriad of small groups and individuals is to carry us into the VLSI revolution, we must not expect these groups and individuals to provide their own fabrication facilities. The level of innovation required can be achieved only if fabrication is provided as a service by a few well capitalized firms.

Every time a qualitatively new element has been introduced into the industry, new business opportunities have been created. Small firms have obtained significant market shares in businesses previously dominated by large firms. The VLSI revolution we are facing is no exception. I fully expect a very large number of small firms, or small groups within larger firms, to create entirely new machine organizations and entirely new design methodologies. These will allow small, able groups to succeed in the varied market place

for systems in spite of historic dominance by capital intensive computer and semiconductor houses.

Evolution of a Product

A product which is implemented as one or more VLSI chips passes through five major evolutionary phases. These are conceptual, design, layout, pattern generation and wafer fabrication.

At some stages of the evolution, it is much more difficult to transfer the knowledge required for further progress than at others. Figure 5 shows a qualitative measure of the information required to transfer the product design from one group of people to another at any given point in the design or production process. At the left hand edge of the curve, the product conception could only be transferred to someone along with essentially complete knowledge of the customer base, the ecomonics of the business, the skills of the personnel in a particular company, cost constraints, timing constraints, etc. About half way through the design process, a block diagram could be transferred which might seem to contain only a small amount of information. However, along with the block diagram would be an enormous amount of context having to do with the miriad of special requirements not completely specified in the written specification. These performance and partitioning constraints are normally solved in an interactive manner during implementation. Examples are testing protocols used in systems developed by the same company, the way in which manufacturing constraints affect overall system design, etc. Once again, a very large amount of information accompanies a design transfer at this stage. Most of today's "custom" LSI designs are transferred to a semiconductor house at this awkward point.

There is an optimal point for the transfer of product. By the time a complete layout has been generated for the chip or chip set implementing a system function, the only information which needs to be transferred from one group to another is the patterns which represent the various layers. This point represents a true minimum in the total information transferred. For example, if one were to transfer a mask instead of data representing the pattern layers, not only would the patterns themselves need to be transferred, but also information which depends upon the details of wafer fabrication process; whether the process uses positive or negative photo resist, how much the lines or spaces of the various layers should be shrunk or expanded to compensate for abberitions introduced in a particular process, etc. None of these details need to be transferred if a data file representing the basic patterns is transferred from the designer to the factory. As the wafers proceed through the fabrication process, each layer is lithographed into the silicon and the amount of information needed for the next processing step decreases. Finally, when the wafers are finished, they may be returned to the designer as fully instantiated artifacts without any additional information.

The minimum in information required to transfer between layout and pattern generation is no accident. This is a very special point in the evolution of a product. It is the end of the design process and the beginning of a

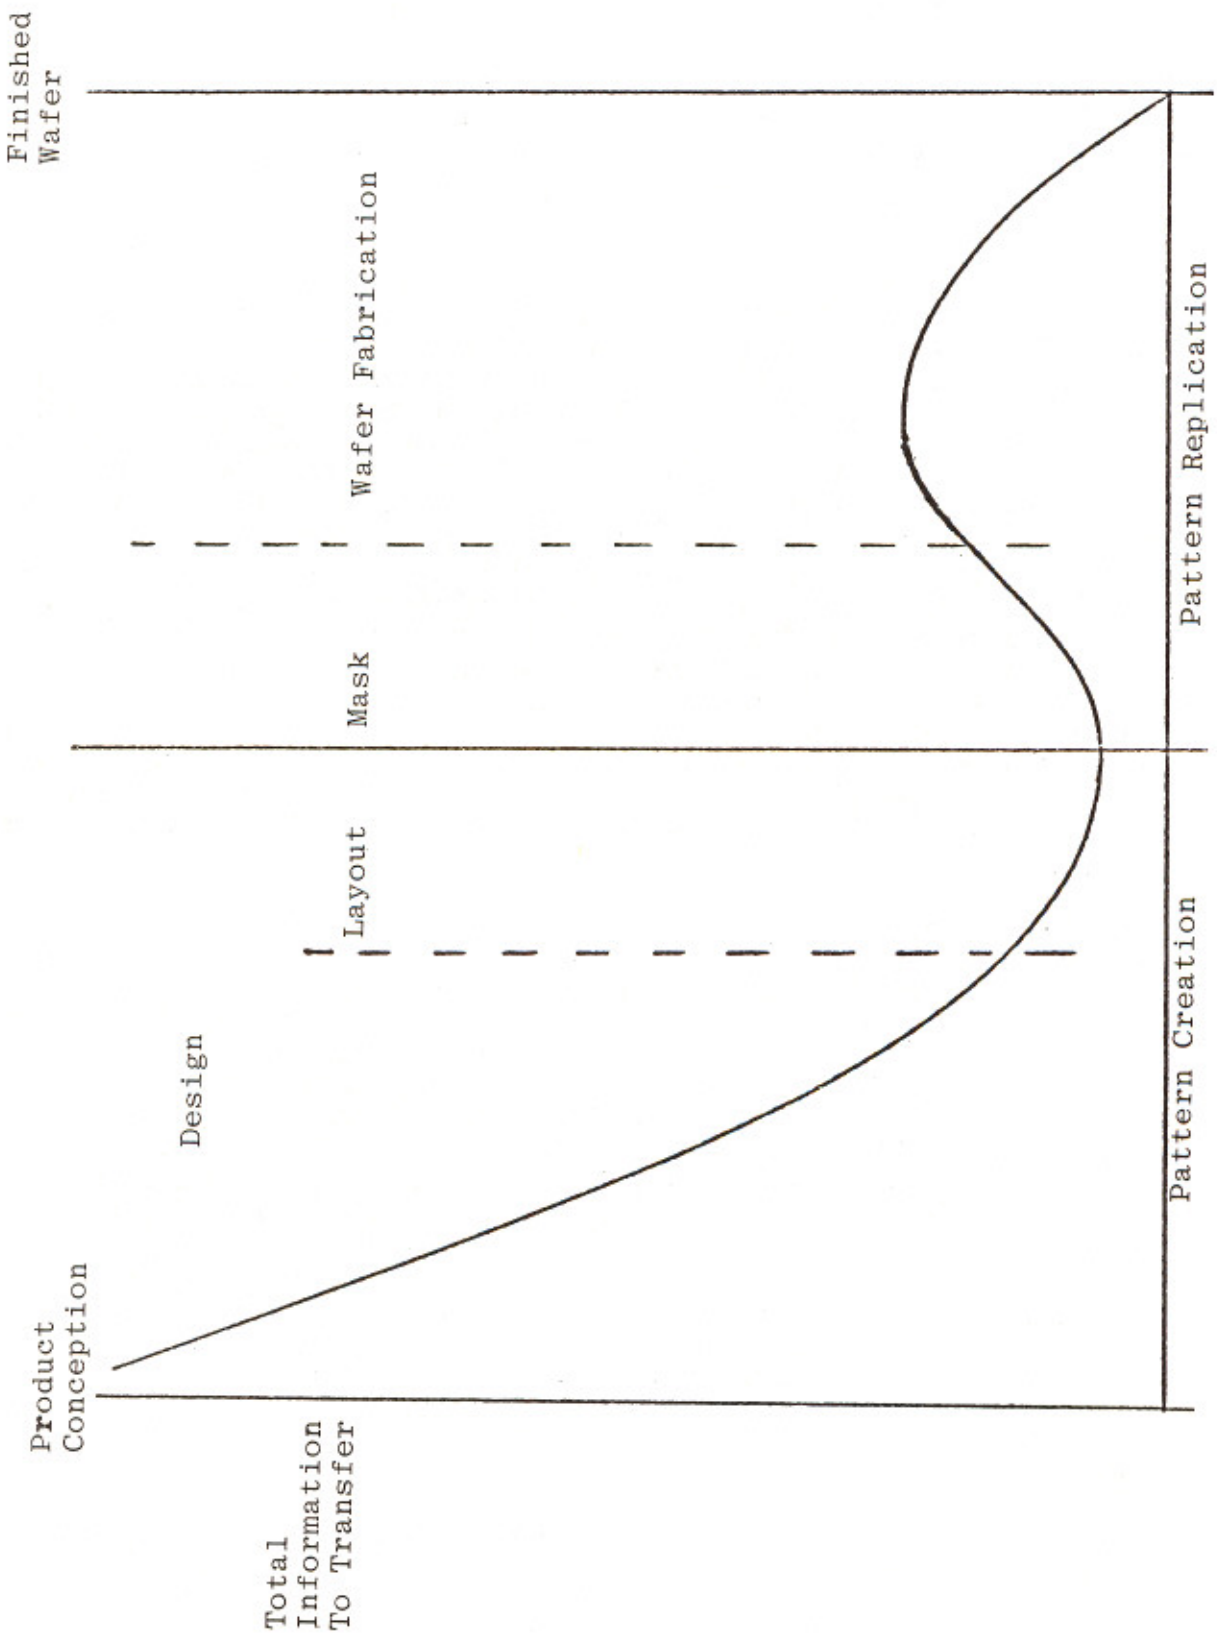
Finished Wafer
Wafer Fabrication
Mask
Layout
Design
Product Conception
Total Information To Transfer
Pattern Replication
Pattern Creation

Figure 5

pattern replication process. Everything to the left of that point has been involved with the specifics of a given product. Every action to the right of that line does not depend upon the specific product, but only upon the process by which the product will be replicated. It is thus the seam between product creation and product replication. To use a familiar analogy in the motion picture industry, those to the left hand side of the line are the producers, the stars, script writers and photographers. Those on the right hand side of the line are the film manufacturers and film processing laboratories. By analogy, we are led to ask what corresponds to the ASA number and color temperature specifications which are used to interface the two worlds of photography. Life would indeed be simple if such a clean interface could be formed between those creating designs and those printing them on wafers of silicon.

An Interface Proposal

As one might expect, the world of silicon is indeed more complicated than the world of film. However, not by as large a degree as the popular image would cause one to imagine. It is possible, with well developed standard processes, to establish a standard interface to almost all fabrication areas running that process. Such an interface requires a remarkably small amount of information to be passed across the boundary.

At Caltech, we have, over the past ten years, been working in colloboration with industry, and more recently with other universities, to develop such a clean interface to wafer fabrication. In the process of implementing 30 or so chip designs, we have interacted with ten different fabrication areas and six mask shops. Although the early interactions were very ad hoc in nature, there has recently emerged a clear vision of how such an interface can be made to work. We are convinced that a modicum of effort expended by those operating fabrication areas can drastically reduce the amount of effort required for user groups to transform designs into silicon. What is required for such an ideal interface to a standard wafer fabrication process? Such an interface consists of three specific, well defined objects.

1. Geometric Design Rules.

2. A Standard Data Format.

3. A Standard Test Chip.

A set of geometric design rules for nMOS silicon gate technology which allows designs to be run on any one of a large number of commercial fabrication areas is given in Figure 6 from reference 2. These rules have been defined in terms of a minimum length unit lambda, which can be selected to conform to the smallest dimensional tolerance in the process. It is thus scalable in such a way that it can follow the dimensional evolution of the process with time. In this way, changes in the geometric resolution of underlying fabrication steps can be taken advantage of without changing the chip design. The design rules used at Caltech and other universities have not changed in form in the last ten years. Over that period of time, the length unit lambda has changed more than a factor or two.

A standard data format is needed to transfer design files to a given fabrication area which uses a given pattern generator. Most output data formats are biased toward a particular output device. A university and industry group has recently developed an intermediate output language which is not biased to any particular output device or design system. It is known as the Caltech Intermediate Form (CIF) and is now used by a number of participating universities and industrial research organizations. The detailed description by Sproull and Lyon is given in reference 2.

The third component of a standard process interface is a standard test chip. This chip can be automatically inserted by the mask/fabrication supplier into the array of product chips at a number of places on every wafer. It contains patterns for process control and characterization of yield, reliability, circuit performance, and system performance. It must be the sole subject of the contract between a fabrication line and its users. In this way, the fabrication people are not blamed for design failures and vice versa. Recent excellent work at the National Bureau of Standards (3) and the Department of Defense (4) has brought this goal within reach. For each standard process, a standard test chip can be made available to all participating fabrication areas.

All semiconductor manufacturing organizations internally operate with the three pieces of interface lore discussed above. However, these objects are not common across corporate boundaries. Although those in many corporations are very similar, they are viewed as highly proprietary and are closely guarded secrets. Having such lore in the public domain is a key factor in assuming rapid innovation in the design of VLSI systems.

The use of a standard interface to a standard process has both costs and benefits. There is no doubt that standard processes do exist which are widely accepted within the industry. There is also no question that a standard interface to such processes could be achieved that would greatly simplify the interaction of designers, producers of design equipment, producers of pattern generation equipment, and managers of fabrication areas. Initially, the use of a standard process and common set of design rules results in lower density and speed than that possible when the product design and process are mutually optimized. Three factors minimize the penalty from this source. 1. Since the rules can be scaled, the product can be debugged with a process available immediately. By the time a product reaches the market, it can take advantage of a high density and higher performance. 2. Fabrication engineers are not distracted by a number of slightly different processes, and can concentrate on the evolution of a single, most advanced, standard process. 3. Design time can be dramatically decreased, getting product to market earlier. Product market life is larger, thus amortizing design cost over a larger number of units.

The central advantage of a clean interface is to allow designers to optimize design methodology and algorithms and architecture while fabrication engineers independently optimize fabrication processors. Such an interface has not been seen as important in the industry until recently because design

was not a major stumbling block in the road to implementing integrated circuits. The rapidly increasing cost of the design of complex systems, together with the enormous potential payoff to be achieved by the use of large scale concurrency in achieving system functions, means a major revolution in the semiconductor industry.

Opportunities exist for both operators and users of silicon fabrication facilities. Service facilities can be very profitable, since their costs are highly predictable and the market base is very broad. These facilities will become very much like the raw silicon wafer suppliers of today. High volume, high profit and low risk. Those firms engaged in the system design business will be completely different. They will be small, and must live by their wits in a constantly changing, enormously competitive industry. Many will be called, but few will be chosen. Those that in fact succeed will form the new cutting edge of an entirely reborn industry. It is an exciting future, but our ability as a nation to undertake such an adventure is dependent upon our willingness to create an available, state-of-the-art fabrication service, available to all in the field who need it.

REFERENCES

1. Seitz, Charles; "Self-Timed VLSI Systems"; from the Proceedings of Caltech Conference on VLSI, January, 1979.

2. Mead and Conway; Introduction to VLSI; to be printed by Addison-Wesley, limited printing by authors, 1978.

3. Buehler, M.G., Grant, S.D. and Thurber, W.R.; Journal of Electrochemical Society, vol. 125; p. 650; 1978.

4. Jurdenoek, R.T., Baire, Jr., H.F., Fromen, G.J.; IEEE Transactions on Elec. Dev. ED 25; p. 873, 1978.