VLSI AND TECHNOLOGICAL INNOVATIONS

Carver A. Mead

California Institute of Technology
Pasadena, CA 91125, USA

Rather than innovation in general, I will discuss what I believe to be the most important opportunity since the industrial revolution, rivalling it in significance. This unique circumstance is created by the emerging Very Large Scale Integrated (VLSI) technology, with which enormously complex digital electronic systems can be fabricated on a single chip of Silicon one-tenth the size of a postage stamp. Out of it systems will be created which radically change our modes of communication, commerce, education, entertainment, science and the underlying rate of cultural evolution. The quality of human life can be improved in remarkable ways by these changes. Electronics creates no noxious by-products and uses only miniscule amounts of energy. It can accomplish tasks which were previously energy intensive, and dangerous or degrading to human workers. There is no doubt that this electronic revolution will take place.

For 20 years the US semiconductor industry has been the world's prime example of innovative excellence. Each new round of technology was carried by a new wave of start-up companies financed with venture capital. These in turn became the mainstays of their market only to spawn another brood as a new opportunity unfolded. Each turn of this wheel of fortune added new markets, formed new capital, and created additional jobs. The spirit of those heady times is by no means dead but the basic nature of the game has changed. System design, not technology, is now the area in which small firms outshine their giant mentors. Implementation of such designs, however, requires a silicon wafer-fabrication capability unaffordable by any single small enterprise. A fabrication facility serving many such firms can set the industry on a course even more exciting than that of earlier times. In what follows I have tried to describe the status of the field and how a simple, decisive act can assure a dynamic new round of innovation within it.

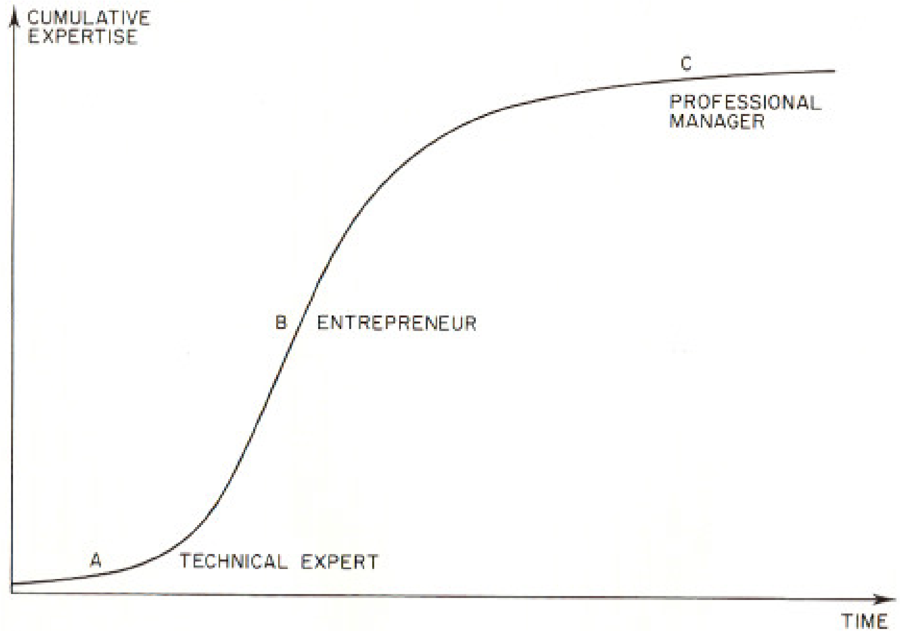
CUMULATIVE
EXPERTISE
A
TECHNICAL EXPERT
B
ENTREPRENEUR
C
PROFESSIONAL
MANAGER
TIME

INNOVATION

I have found the following simple view helpful in understanding how innovation has worked in the past, and how we might encourage it in the future. 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 Fig. 1. 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 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) in Fig. 1, 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 applications, 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 Fig. 1, progress becomes logarithmic in time. Few changes in the market share "pecking order" occur. Manufacturing methods are refined even 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 required credentials. Professional managers and large firms dominate the business.

THE ARENA

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 (Setz, 1979) presented in Fig. 2. At several points in the evolution of the technology,

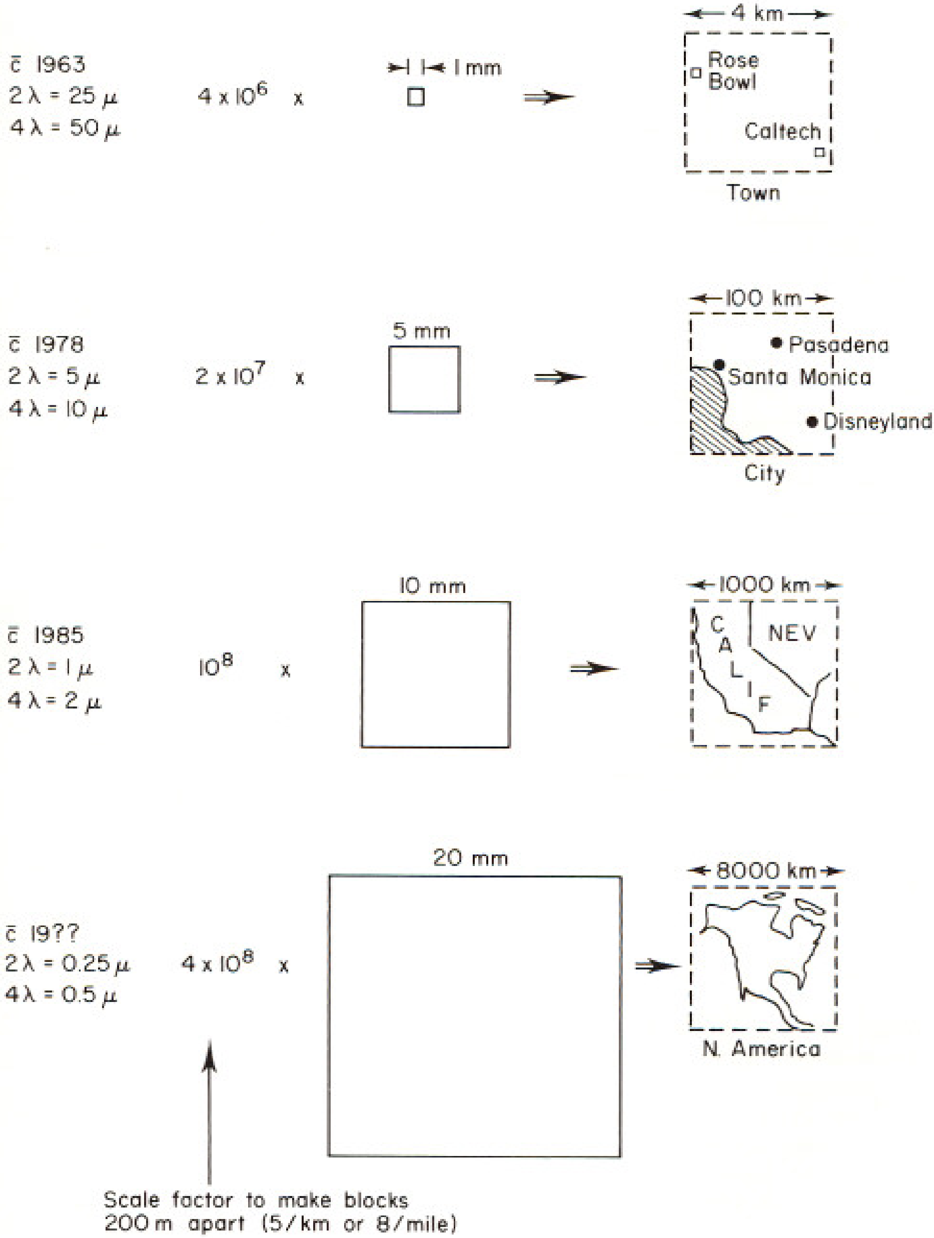

c̄ 1963
2 λ = 25 μ
4 λ = 50 μ
4 x 10^6 x
1 mm
4 km
Rose Bowl
Caltech
Town
c̄ 1978
2 λ = 5 μ
4 λ = 10 μ
2 x 10^7 x
5 mm
100 km
Pasadena
Santa Monica
Disneyland
City
c̄ 1985
2 λ = 1 μ
4 λ = 2 μ
10^8 x
10 mm
1000 km
C A L I F
NEV
c̄ 19??
2 λ = 0.25 μ
4 λ = 0.5 μ
4 x 10^8 x
20 mm
8000 km
N. America
Scale factor to make blocks
200 m apart (5/km or 8/mile)

Fig. 2.

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 carrying electrical signals instead of cars. In the mid 1960s, 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 using a 5μ tech-technology is comparable to the entire Los Angeles basin. By the time a 1μ technology is solidly in place, designing a chip will be comparable to planning a street network covering all of California and Nevada at urban densities. The ultimate $\frac{1}{4}$μ technology will likely be capable of producing chips with the complexity of 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. In order to realize the full potential implied by such complexity, entirely new design methods and system organizations must be invented. A high rate of innovation is required to achieve leadership in this remarkable arena.

The evolution of the component fields which make up the present VLSI disciplines are shown schematically in Fig. 3. Progress in each depends upon those before it being well in place. 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 in its exponential phase. Many fundamental ideas have yet to be discovered. The organization and programming of highly concurrent systems are 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. It is here that the major innovative possibilities lie.

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

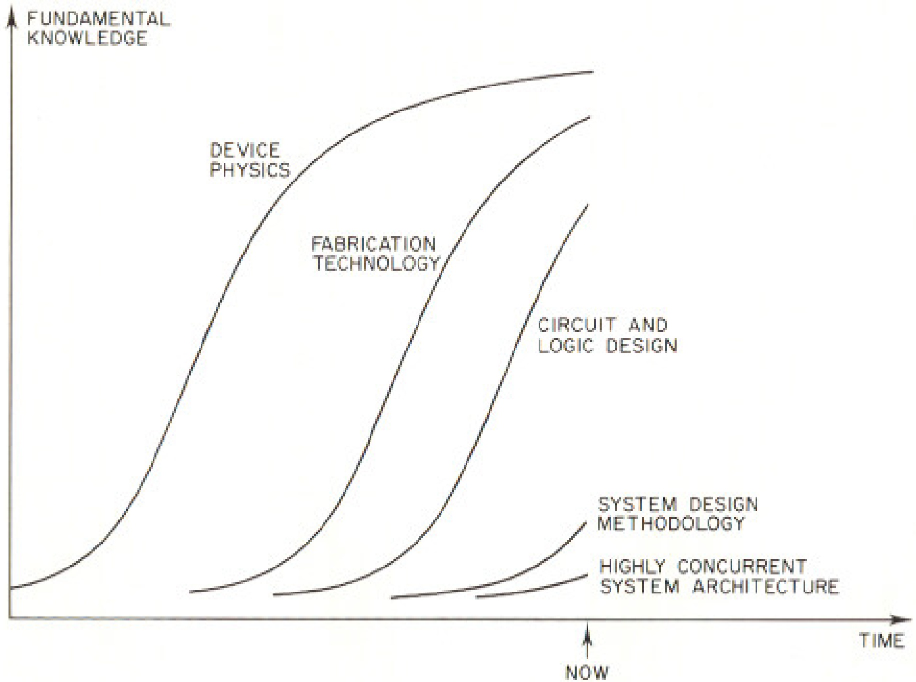
FUNDAMENTAL
KNOWLEDGE
DEVICE
PHYSICS
FABRICATION
TECHNOLOGY
CIRCUIT AND
LOGIC DESIGN
SYSTEM DESIGN
METHODOLOGY
HIGHLY CONCURRENT
SYSTEM ARCHITECTURE
TIME
NOW

share by virtue of its innovation. Existing dominant firms were then forced to retrofit these ideas these ideas into their own product lines. For the reasons mentioned above, I expect this trend to be even more apparent as VLSI evolves.

The problem faced by the semiconductor industry is thus apparent. Fabrication technology has reached its capital intensive stage. Design is still very early in its exponential phase. Each small system group can no longer afford its own fabrication area. A start-up firm of ten years ago 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 myriad 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 by extending markets previously dominated by larger firms. The VLSI revolution we are facing is no exception. I fully expect a very large number of small 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.

The seeds of the new wave of innovation have already begun to take root. Much of the thrust is coming, not from the integrated circuit industry itself, but from the collaborative effort of faculty members in many university computer science departments and scattered individuals throughout the industry.

A new type of course is being offered at these universities that provides students of computer science and electrical engineering with a thorough introduction to integrated system architecture and design. The courses provide the minimum of basic information about devices, circuits, fabrication technology, logic design techniques, and system architecture, which is sufficient to enable the student to span fully the entire range of abstractions, from the underlying physics to complete VLSI digital computer systems. Only a surprisingly small set of carefully selected key concepts is necessary for this purpose. By minimizing the mental baggage carried along each step, the student emerges with a good overall understanding of the subject. The courses are based on a course I originated at Caltech in 1971, and on a new and highly integrated system design methodology that Lynn Conway (of the Xerox Palo Alto Research Center) and I present in a major new textbook on the subject (Mead and Conway, 1978/79).

A major portion of these courses is devoted to "learning by doing", with the student undertaking the architecture design, layout and testing of substantial LSI system project A typical student project might contain on the order of several hundred to several thousand transistors. The students typically describe their designs using a symbolic layout language. Only modest facilities are needed to support such courses, in addition to that commonly available at these universities. The universities do not usually have their own LSI fabrication facilities. The "multiproject chips" containing all the LSI projects for such a course are instead fabricated by commercial maskmaking and wafer fabrication firms, directly from the design produced by the students.

Throughout the evolution of courses and research projects, we have constantly been hampered by the difficulty of gaining access to wafer fabrication service. Although many industrial firms have been very helpful in this regard, their facilities are dedicated to their own proprietary products, and were made available to us on a special case, special favor basis. While these favors were crucial to getting the university programs going, they obviously cannot be the basis for supplying a multitude of emerging small system firms, nor can they effectively support the expanding group of participating universities. What is needed is an openly available service facility which will, for a fee, accept pattern files in a standard format (which we have already worked out) and "print" them on silicon wafers using a standard fabrication process.

The analogy with printing is a particularly apt one (Conway While wafer fabrication is the most sophisticated manufacturing technology ever undertaken by the human race, it is pattern independent. A limitless number of system designs can be replicated by a single process, much as books on any number of subjects can be replicated by a single printing press. Individual designers require access to fabrication in exactly the same way that individual authors require access to printing.

An open fabrication service should be established as the collaboration of an industrial partner with the university community. Funding arrangements should be left flexible, and a rigid format should not be prescribed. Once one or two of such facilities have been operating for a year or two, demand will increase and they should become profitable business enterprises in their own right. It is important that the only restriction placed on these facilities is that they make fabrication service openly and uniformly available. It is perhaps the only way to assure Freedom of the "Silicon Press" (Conway).

REFERENCES

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

Mead, Carver and Conway, Lynn (1978/79). "Introduction to VLSI Systems". Limited printing by the authors, 1978; Addison-Wesley, 1979.

Conway, Lynn. Personal communication.