Toward Giga-Scale Silicon Integrated Circuits

Colin A. Warwick
Ran-Hong Yan
Young Ock Kim
Abbas Ourmazd

Silicon integrated circuits now permeate our daily lives, in ways that seem surprising. Today, a "notepad" computer is small enough to be used on an airplane to perform tasks that, a few years ago, would have required a machine heavier than the combined weights of all the passengers on board. In reality, however, silicon technology has followed unremitting and predictable trends since its inception over 30 years ago. Here, we outline these trends, assess their implications, and examine their potential limits. Beyond the question of limits, however, lies the important role that silicon technology can play in aggressive companies like AT&T. In an essential sense, silicon is the embodiment of intellectual property. Appropriate, concurrent engineering-based use of this technology enables vertically integrated companies to erect barriers against rapid duplication, which will extend the initial, highly profitable stages of the product cycle. As such, silicon technology can be a formidable weapon in rapidly evolving competitive markets.

Introduction

ENIAC — formally dedicated in 1946, and generally regarded as the first electronic digital computer — was built to calculate firing tables for artillery shells. Its three-ton bulk included 18,800 vacuum tubes, and calculated a ballistic trajectory faster than a shell could reach its target. Today, calculation of trajectories is just *one* of the functions performed by the Hobbit™ microprocessor in AT&T's hand-held personal communicator. The trajectories calculated by the Hobbit microprocessor are the motions of a pen, as the user writes on a "notepad." These motions are "recognized" by the Hobbit microprocessor as characters, and stored, using about 30,000 times more computing power than the ENIAC could deliver. Handwritten character recognition lightens our burden, because most people find writing more convenient than typing.[1] And it literally lightens our burden, by packing enormous computing power into a 2.2-lb. "notepad." In many ways, silicon integrated circuit technology and software have combined to bring previously unimaginable computing power into our lives. This continuing revolution promises to mass-produce products and services now encountered only in science fiction movies and comic books. (See Panel 1 for definitions of abbreviations, acronyms, and terms.)

The dramatic improvements in silicon technology have been sustained for over 30 years. Within a decade, circuits with a billion elements or more (known as *giga-scale integration*) will be available in the marketplace. To the customer, these improvements in silicon technology offer affordable solutions, rather than solutions that require computing power available only at great cost, or not available at all. All this has been made possible by rapid increases in the number of circuit elements per chip; equally rapid reductions in the price per element; and less rapid, but important, improvements in speed and power. To the technologist, these advances have stemmed from head-long reductions in the defect density and *minimum feature size* (i.e., the smallest component on a chip, usually a transistor gate), and less rapid, but significant, increases in device complexity and die size.

To the manufacturer, these developments have been accompanied not only by continual growth in market size, but also by

Panel 1. Abbreviations, Acronyms, and Terms

chip — a silicon integrated circuit after packaging
CMOS — complementary metal-oxide-semiconductor
die — a silicon integrated circuit before packaging
DRAM — dynamic random access memory
ENIAC — electronic numerical integrator and computer
Gbit/s — gigabit(s) per second
GHz — gigahertz
gm — gram
MIPS — millions of instructions per second
MOS — metal-oxide semiconductor
MOSFET — metal-oxide semiconductor field-effect transistor
mph — mile(s) per hour
mV — millivolt
mW — milliwatt
NiCd — nickel cadmium
S/m — siemen(s) per meter
VSELP — verbal sum-excited linear prediction

breathtaking escalations in the cost of manufacturing. These factors have precipitated the formation of consortia to share costs, and have fueled the search for alternative manufacturing paradigms. At the same time, aggressive, vertically integrated companies have used access to technology to erect barriers against duplication in rapidly evolving, competitive markets, thereby extending the initial, highly profitable phase of the product cycle.

In this paper, we examine some trends in silicon technology, outline its potential limits, and examine emerging products and services that may affect these trends. We conclude with a brief discussion of the way silicon technology can bring significant competitive advantage to vertically integrated companies. We illustrate general concepts by giving specific examples, rather than providing a detailed survey of a complex field.

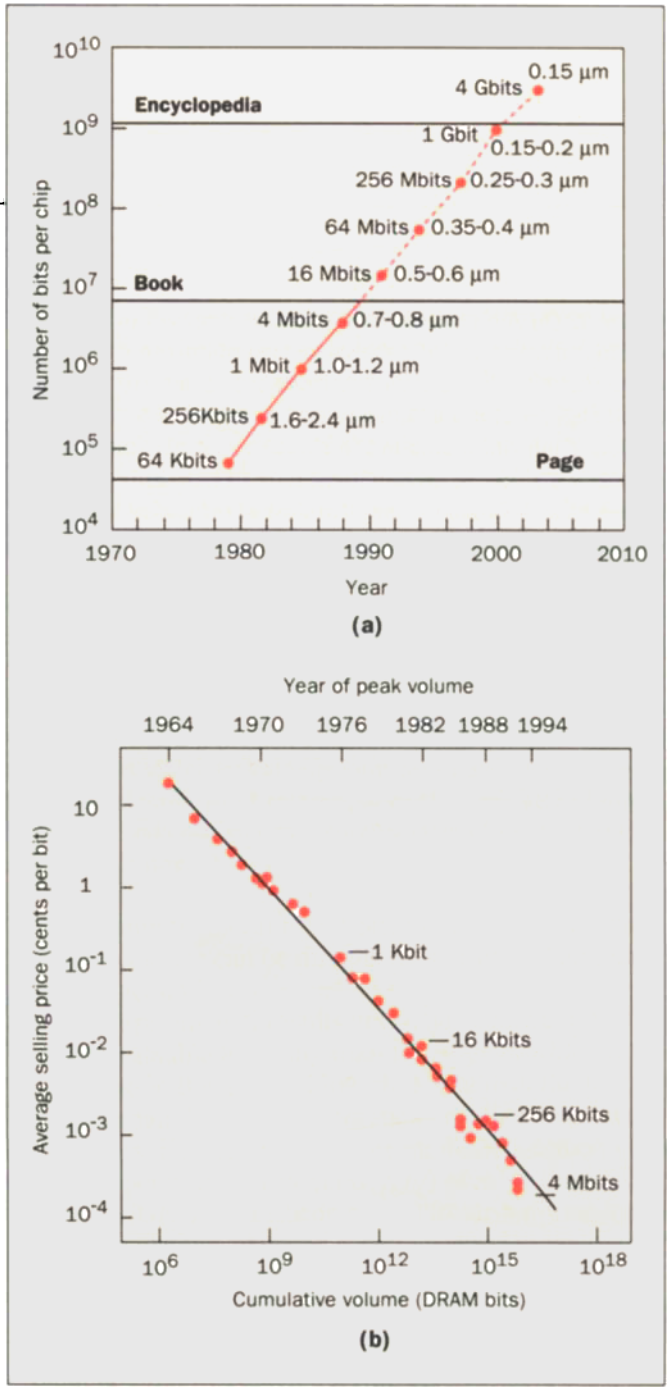

Figure 1. (a) Number of bits stored on a DRAM silicon chip per year. The dotted line shows extrapolation of present trends. (b) Price per bit, plotted versus cumulative experience, as measured by total number of bits manufactured.

Trends and Their Potential Limits

Since their introduction to the marketplace, silicon integrated circuits have brought about exponential increases in transistor density, equally rapid reductions in cost, and important, but less dramatic, improvements in performance and power consumption. In 30 years, as Figure 1 shows, the number of transistors per chip has increased one million times, and the price per transistor has dropped 100,000 times. Concurrently, power consumption has been reduced 100,000-fold, and transistor speed has been increased.

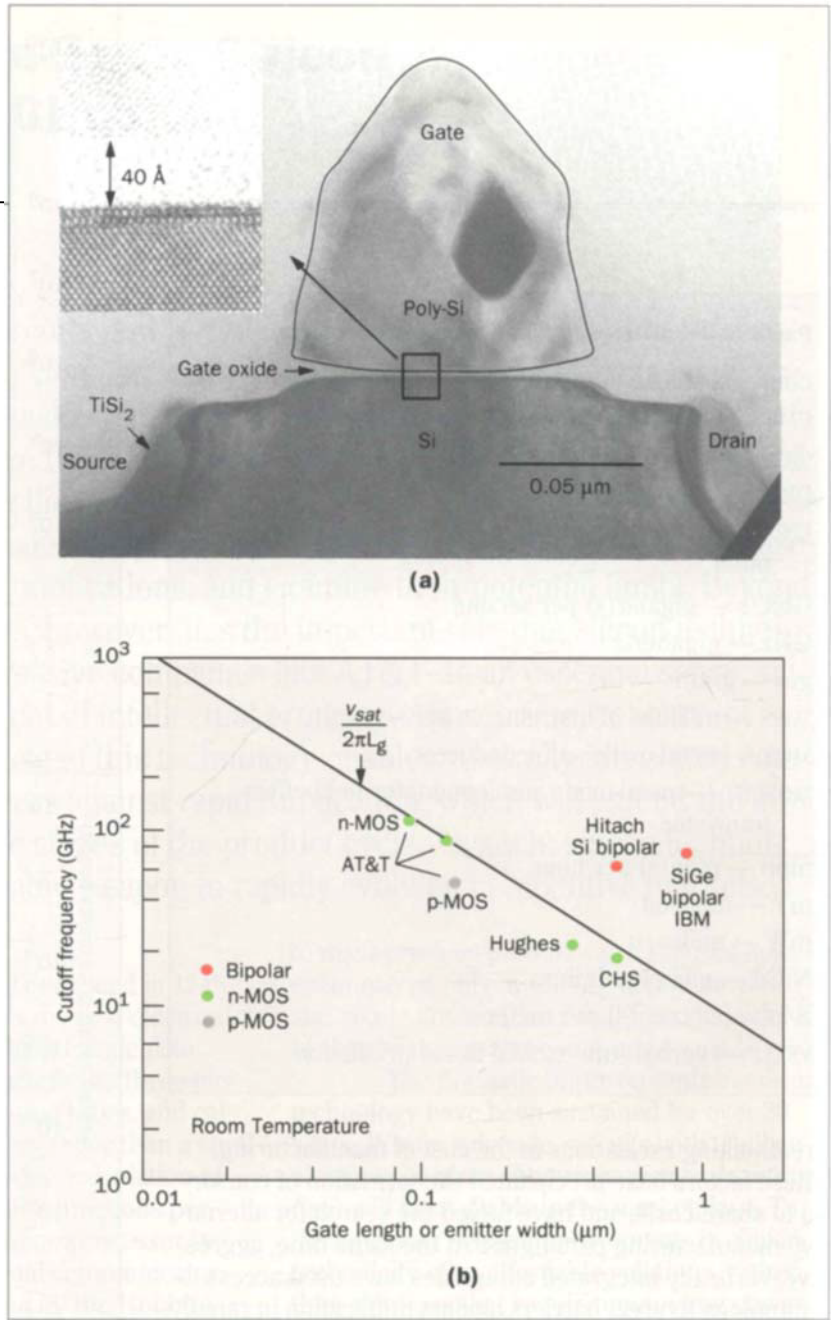

Figure 2. (a) A high-resolution transmission electron micrograph of a cross-section of an n-channel MOS transistor. The inset shows a lattice image of the channel region of this device. Each white blob in the silicon substrate represents a pair of atom columns. The channel is less than 400 atoms long. (b) Plot of intrinsic device speed (or cut-off frequency f_T) versus gate length or emitter width for silicon n-MOSFETs, p-MOSFETs, and silicon-based bipolar transistors. The line represents a simple estimate of the intrinsic n-MOSFET speed, based on fundamental considerations. The agreement between experimental points and this estimate provides confidence that the expected intrinsic speed can be approached in practice.[3,4,17–22]

To place these developments in a more familiar context, we can equate the physical characteristics and price of a car with those of a computer, we can liken the capacity of a car to the memory of a computer, and the speed of a car to the number of instructions a computer can carry out in a second. If we extend this analogy, improvements in the automotive industry equivalent to those in computers since 1946 would have resulted in a present-day car that weighs 2 ounces, costs $40, has 50,000 cubic feet of luggage space, and uses only one gallon of gasoline per 1,740,000 miles to travel at 1,350,000 miles per hour.

Because of these spectacular advances, it is not surprising that every generation of scientists and engineers is fascinated with identifying the factors that may slow down and, ultimately, stop the "silicon rush." However, such exercises define the boundary of contemporary knowledge as much as they predict the eventual limits of silicon technology. Here, we outline the factors that set fundamental limits on silicon technology, acknowledging that practical considerations may prevent us from reaching these fundamental limits.

Shrinking the Transistor

Because complementary metal-oxide-semiconductor (CMOS) devices and circuits consume little power, they are the most likely candidates for mainstream, ultra-dense circuits, in which power consumption

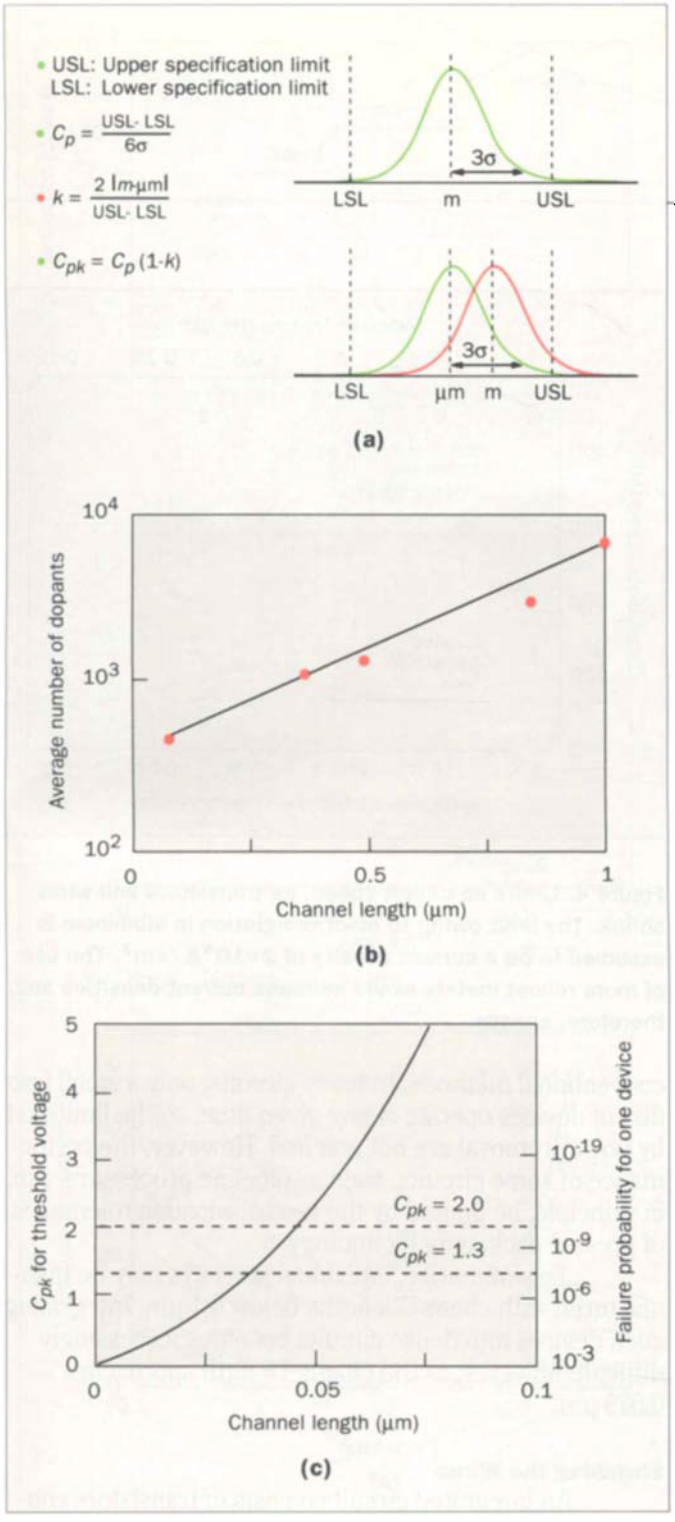

Figure 3. (a) Process capability indices (C_{pk}) are used to quantify the deviations in the outcome of a statistical process from the design value. A statistical distribution is characterized in terms of its width σ and its mean *m*. The C_{pk} parameter is a measure of the width of the distribution, and the deviation of its mean *m*, from the design value μ. C_{pk} values below 1.3 are considered unacceptable for most purposes. (b) Number of dopant atoms in the channel region of a silicon MOSFET versus channel length. (c) Effect of statistical variations in channel doping on the C_{pk} parameter for threshold voltage of deep submicron silicon MOSFETs. These statistical variations are a fundamental property of how dopant atoms are introduced into the device. At small channel lengths, they cause significant variations in the threshold at which a device switches. C_{pk} values of less than 1.3, reached at channel lengths of about 0.03 μm, are thought to indicate an unacceptable level of variation from device to device.

and removal of dissipated heat are major concerns. At present, the scientific community thinks that it will be possible to design and manufacture very small metal-oxide semiconductor (MOS) transistors, with channel lengths approaching 0.075 μm.[2] More recent experimental results, illustrated in Figure 2, have demonstrated deep submicron n- and p-MOS field-effect transistors (MOSFETs), which display excellent dc and ac characteristics at room temperature.[3,4] Such devices are so small that their lengths can be directly measured by counting atomic spacings.

Experimental results of the type shown in Figure 2 are valuable for many reasons, three of which we will discuss in this paper. First, they establish the performance domain for future generations of silicon devices and circuits. For example, the recent demonstration of an intrinsic speed (cut-off frequency, f_T) of more than 100 gigahertz (GHz) for a silicon n-MOSFET strongly suggests that extrapolations of speed based on simple arguments will remain valid for several future generations of silicon devices (see Figure 2b). Second, the experience gained in manufacturing these devices allows us to assess the extent to which processes based on current technology can be used to manufacture devices with minimum feature sizes as small as 0.1 μm. This significant finding identifies areas in which radically new technologies may, or may not, be needed. Finally, such results allow us to conclude that deep submicron silicon devices will operate well at room temperature, obviating the need (suggested in earlier work[2]) to design them for a temperature-controlled enclosure. This is particularly important, since the design of many future systems will be shaped by weight and portability requirements, and because battery-powered, temperature-controlled enclosures would be very heavy.

Various fundamental considerations, however, set upper bounds on the density (per unit area) of

devices that may be integrated into circuits. A computer model may tell us the limit for idealized, that is, *perfect*, transistors, but any real process used to manufacture devices produces results that vary statistically from device to device. A small probability exists that a few devices will differ significantly from the intended design. In very dense circuits, with transistor counts approaching one billion, this can cause enough device failures to limit chip yield. Expressing the statistical variations in terms of dimensionless C_{pk} parameters, for a particular distribution of devices, gives the width of the distribution and the deviation of its mean from the design value (see Figure 3a). Values of C_{pk} parameters that fall below 1.3 should reduce yield significantly. For giga-scale circuits, whose manufacture is likely to include as many as 850 steps, a C_{pk} of 1.3 for each step implies a yield loss of about 8 percent from process variations alone. Defects in the wafer may further reduce the yield.

Many processes used in silicon technology are fundamentally random in nature, subject to statistical fluctuations that place fundamental upper bounds on circuit yield. For example, careful control of the number of dopants in each silicon device ensures satisfactory circuit performance. Although channel doping *concentrations* are high in ultra-small MOSFETs, the number of dopant atoms in the channel region of a device can be very small (see Figures 3b and 3c). When these dopants are introduced by processes such as implantation or diffusion, the number of atoms incorporated into each device is subject to statistical fluctuations, which become significant for small numbers of dopants. In Figure 3c, the fluctuations became large enough to reduce the C_{pk} value for the MOSFET threshold voltage to less than 1.3, at channel lengths approaching 0.03 μm. Such effects conspire to limit the circuit yield — and, therefore, the maximum device density — to perhaps a few billion transistors per centimeter squared.

Even in circuits less dense than these upper bounds, more practical considerations, such as the need to dispose of heat generated in these circuits, limit the device density per unit area. By reducing the device size, we encounter excessive unwanted currents emanating from partly open switches (subthreshold device leakage). Power dissipates substantially below channel lengths of about 0.1 μm, even with devices in a quiescent state. Also, very dense circuits generate amounts of heat that become increasingly difficult to remove using

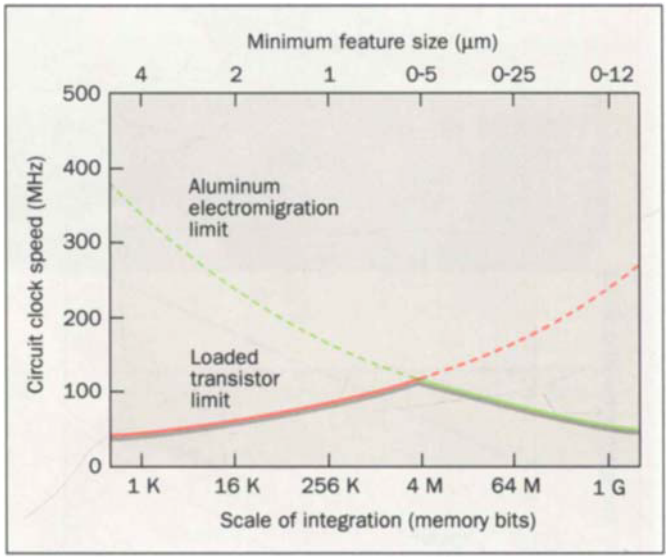

Figure 4. Limits on circuit speed, as transistors and wires shrink. The limit owing to electromigration in aluminum is assumed to be a current density of 2×10^5 A/cm^2. The use of more robust metals would increase current densities and, therefore, speeds.

conventional methods. In many circuits, only a small fraction of devices operate at any given time, so the limits set by power removal are not reached. However, the performance of some circuits, such as pipeline processors, can, in principle, be limited by the heat dissipation tolerances of present packaging technology.

To summarize, individual MOSFETs may be manufactured with channel lengths below 0.1 μm. Integrating such devices into dense circuits becomes increasingly difficult, however, as the channel length approaches 0.075 μm.

Shrinking the Wires

An integrated circuit consists of transistors connected by wires on the same chip. In addition to miniaturizing transistors, the manufacture of dense circuits requires shrinking the cross-section of the interconnecting wires. This trend places stringent demands on chip wiring. A typical wire in an integrated circuit must withstand a current density of up to 200,000 amperes per centimeter squared (A/cm^2), nearly 200 times greater than the maximum current density of household wiring. The passage of large currents through metal can cause

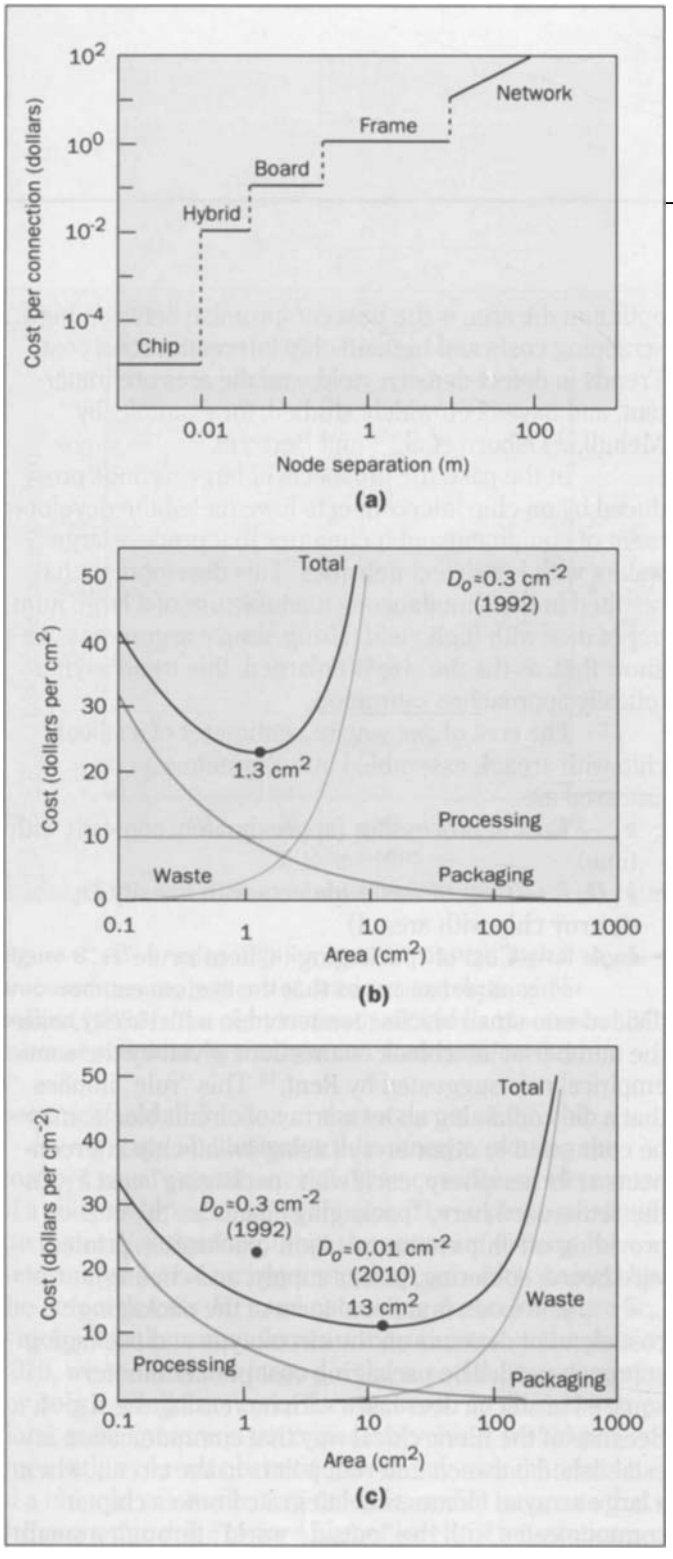

Figure 5. (a) The cost of interconnection versus distance between connected nodes.[8] Cost is dramatically reduced when the connection is between nodes on the same chip. (b) Plot of the cost per centimeter squared of a chip (a finished, packaged silicon die) and its constituent elements versus die area, at defect densities typical for 1992. The optimum die size is 1.3 cm^2. The cost increases rapidly as the die size departs from the optimum value. Reducing packaging costs by bringing more interconnects on chip would lower the cost, until a severe yield penalty is exacted by defects. (c) Same plot for the year 2010, using a projected defect density. The optimum die size is now 13 cm^2. The wide plateau in cost indicates insensitivity to die area for a broad range of sizes. Under these conditions, the economic advantage of bringing more interconnects on chip is essentially exhausted, and the optimum die size is more a matter of convenience.

atomic motion, known as *electromigration*. If this effect is not carefully controlled, a wire can develop physical voids and break. Because of their low resistivities and the ease with which they can be patterned into wires, aluminum and its alloys have been the materials of choice in silicon technology. However, aluminum can degrade rapidly in the presence of electromigration, particularly because it is placed under tensile stress as it is formed into on-chip wiring. Electromigration, which depends strongly on the current density that travels through the wire, places stringent limits on the current that can charge up stray capacitances and switch other transistors. Circuit speed is limited in proportion to miniaturization (see Figure 4), because the cross-section of the interconnecting wire decreases more rapidly than the current a miniaturized transistor can deliver.

Although several metals are more immune than aluminum to the effects of electromigration, and to a related effect known as stress voiding, their resistivities exceed those of aluminum and its alloys. Can present electromigration limits be circumvented by replacing aluminum with more robust, but also more resistive, metals? It is generally thought that increased resistance degrades circuit speed. Determining the effect of interconnect resistance on circuit speed is difficult, because circuits are highly complex and varied in character. However, the essence of many future circuits can be distilled into a simple model. This partitions the problem into two components — one generic to many circuits, and one circuit-specific. The generic component contains only a handful of physical constants, and the specific component varies weakly with circuit-specific parameters. In these conditions, it becomes possible to study aspects of circuit performance without allowing circuit-specific features to strongly affect the conclusions.

Such models suggest that more robust metals

can tolerate higher speeds in miniaturized circuits, despite their somewhat higher resistance. The speed of a giga-scale circuit is limited by the time needed to charge up the parasitic capacitances with the current available in a clock cycle. This limits the maximum capacitance — and, thereby, the wire length between nodes — to such small values that the wire resistance does not significantly affect the circuit speed.[5-7] Under these conditions, it might be possible to replace the low-resistance, though delicate, aluminum alloys with robust materials that are more immune to electromigration and stress voiding.

These qualitative arguments ignore important effects, such as crosstalk. They are insufficient, alone, to justify the significant changes in technology needed to replace aluminum in giga-scale circuits. A simple, generic model could be used, however, as a guide to make the complex choices we face as we develop and manufacture ultra-dense integrated circuits in silicon. Miniaturizing on-chip wiring presents challenges at least as formidable as those involved in making transistors smaller.

Growing the Die

As pointed out by John Mayo,[8] a major driving force toward higher levels of integration stems from the much lower cost of interconnecting two nodes on the same chip (see Figure 5a). We might, therefore, imagine continuing the push to levels of integration higher than device miniaturization allows, using larger die areas for each silicon chip. On the other hand, there is an economic limit to enlarging the die size, likely to be reached about the same time as the expected limits on shrinking integrated circuits.

When a die contains a defect, it must be scrapped. (Sometimes, it is feasible to locate the defects and repair the circuit, but the cost of this exercise often exceeds the value of the die. In some memory dice, defective sections can be replaced automatically with spare circuitry. In the following, the same arguments apply, provided the phrase "defect density" is used instead of "residual defect density after repair.") For each particular defect density, the probability that a die contains defects grows with the die area. The number of usable dice decreases with die area, resulting in a substantial yield penalty. At the same time, a larger die has a greater number of elements on chip, thus reducing the packaging cost. As noted by Murphy[9] and Noyce,[10] "The optimum die area is the best compromise between high scrapping costs and high off-chip interconnection cost." Trends in defect density, yield, and die area are important, and have been widely studied, for example, by Meindl,[11] Osburn et al.,[12] and Bertram.[13]

In the past, the prospects of large savings produced by on-chip interconnects have fueled the development of equipment and techniques that process large wafers with low-defect densities. This development has resulted in the simultaneous manufacture of a large number of dice with high yield. Using simple arguments, we show that, as the die size is enlarged, this trend asymptotically approaches saturation.

The cost of *one* square centimeter of a silicon chip with area A, assembled into a system, can be assessed as:

- k_1 — Cost of processing (approximately constant with time)
- $k_1 D_0 A$ — Cost of waste (defects with density D_0 destroy chip with area A)
- $k_2 A^{-1/2}$ — Cost of "packaging" ("Rent's rule").

This model assumes that the system can be divided into small blocks, connected in a hierarchy, with the number of interblock connections given by the semi-empirical rule suggested by Rent.[14] This "rule" implies that a die containing an $n \times n$ array of circuit blocks may be connected to other arrays using $4n$ off-chip interconnects at its periphery, each with "packaging" cost k_2. In the sense used here, "packaging" includes the cost of providing off-chip communication (packaging, printed wire board, soldering, power supply, and cooling).

The exact functional form of the packaging cost element depends on the circuit type and packaging approach used. The packaging cost per centimeter squared of silicon decreases with increasing die size A. Because of the hierarchical way that communication is established between different points in the circuit, when a large array of elements is integrated onto a chip, it communicates with the "outside world" through a small number of pins. We are familiar with this concept in the world of communication: The number of telephones in New Jersey and Florida is much larger than the number of lines connecting them. The cost of connecting two telephones is small, because they share the same line with many other nodes. In the same way, when many elements are integrated onto a chip, the cost of interconnecting them with the outside world (i.e., packaging)

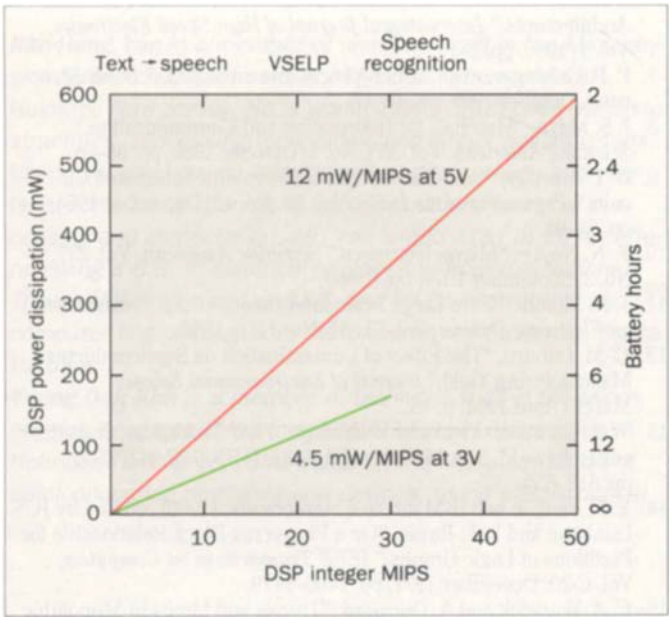

Figure 6. Power consumption of AT&T DSP-16 digital signal processor versus millions of (integer) instructions per second (MIPS). The right vertical axis shows the number of hours of operation on a 20-gm (AA size) NiCd battery.

becomes a small fraction of the total cost of the chip.

Figure 5b illustrates this concept, plotting the cost per square centimeter of three constituent elements of a silicon chip against its die size. These plots substantiate the earlier statement that the optimum die size is determined by balancing scrapping costs (determined by the defect density D_o) and packaging costs. In Figure 5c, which plots the same elements extrapolated to the year 2010, we see that, for large die sizes, packaging accounts for only a small fraction of the total cost. In these conditions, putting more interconnects on-chip[15] would produce little economic benefit. The flat plateau in the cost of a chip versus its die size (see Figure 5c) suggests that the economic drive toward higher levels of integration has been exhausted. Under these conditions, a range of die sizes will have similar costs, making the die size a matter of convenience rather than economics.

Such considerations lead us to expect various technological and economic limits, to be approached simultaneously for chips approximately 3 cm per side, containing over a billion elements, with minimum feature sizes as small as 0.1 μm. Dynamic random access memories of 4 gigabits per second (Gbits/s), static random access memories of 1 Gbit/s, and correspondingly complex logic circuits will be manufactured before these limits are reached.

New Market-Driven Trends

The evolutionary trends outlined above promise cheap, powerful processing, combined with enormous information storage capacity. While high-end applications will continue to harness this raw power for purely computational purposes, the mainstream is likely to use computational power to provide user-friendly features, such as graphical user interfaces.

At present, we access computers primarily through the keyboard. Most of us, who operate in the twilight zone between the "hunt-and-peck" typist and the "skilled" professional, find this a painful mode of interaction.[16] We type at an average rate of 30 to 50 words per minute, but speak about 150 words per minute. And while our reading speed averages about 360 words per minute, it requires our full attention. In contrast, we can listen to (and, perhaps, even comprehend) as many as 250 words per minute as we drive a car. It is likely that future computers will use part of their power to interact with us on our terms, rather than theirs, as we can already glimpse in the first computers with limited speech recognition and hand-written input capabilities.

We have so far outlined the drive toward integration and higher speed. However, recent developments suggest that the market is looking past density and speed toward portable, wireless products. Portability dictates low power consumption, even at the expense of some performance degradation. For example, a day's use (equivalent to three hours of continuous operation) on a fully charged, 20-gram (gm), AA-size nickel cadmium (NiCd) battery requires a system whose power consumption is less than 400 milliwatts (mW), which is about 20 times lower than the power consumption of a typical flashlight.

The power consumption of CMOS circuits scales approximately as fCV^2, where f is the clock frequency, C is the capacitance (primarily determined by circuit layout), and V is the supply voltage. Because the power consumption is most affected by voltage, market demands for low power place a strong downward pressure on the power supply voltage. These demands have already resulted in significant departures from the status quo in

power supply reduction, which were to stick with 12V, or later, with 5V, until device and material limitations made operation impossible without breakdown. In contrast, to reduce power consumption, power supply voltage has been reduced more quickly than warranted by such considerations alone.

The beneficial effects of this trend are illustrated in Figure 6, which shows the reduction in power consumption for the AT&T DSP-16 family of products. Migration from a 5V to a 3V power supply results in substantial power savings, and, more significantly, makes it possible to use standard algorithms on portable products.

Conclusion

Reductions in power consumption require measures that encompass system and chip architecture, circuit design, and technology. The rapid design and implementation of such measures by concurrent engineering can provide a significant competitive advantage to companies with applications knowledge. Access to and control of technology can also be a powerful barrier against competition. In 1992, Japan exported only 3 percent of the low-power components it manufactured, with the latest products not slated for export at all. In an essential sense, silicon is the embodiment of intellectual property, whose full potential can be achieved only in the products it makes possible. In this form, it is a weapon to be reckoned with in today's competitive market.

Acknowledgments

We acknowledge valuable discussions with W. F. Brinkman, J. C. Clemens, R. Cox, R. E. Howard, L. D. Jackel, H. D. Lin, L. Manchanda, L. Rabiner, B. Razavi, W. W. Troutman, and D. C. H. Yu.

References

1. *BIS Strategic Decisions*, Mobile Professional Market Segmentation Study, March 15, 1993.
2. G. A. Sai-Halasz, et al., "Design and Experimental Technology for 0.1-μm Gate-Length Low-Temperature Operation FETs," *IEEE Electronic Device Letters*, Vol. 8, No. 10, October 1987, pp. 463–466.
3. R. -H. Yan, et al., "High Performance 0.1-μm Room Temperature Silicon MOSFETs," *Technical Digest of the 1992 Symposium on VLSI Technology*, Seattle, Washington, June 1992, pp. 86–87.
4. K. F. Lee, et al., "0.1 μm P-Channel MOSFETs with 51 GHz f_T," *Technical Digest of the International Electron Devices Meeting, 1992*, San Francisco, California, December 1992, pp. 1012–1014.
5. *MICRO TECH 2000 Workshop Report*; National Advisory Committee on Semiconductors, 1555 Wilson Blvd, Suite 702, Arlington VA 22209.
6. R. H. Katz and J. L. Hennessy, "High Performances Microprocessor Architectures," *International Journal of High Speed Electronics*, Vol. 1, 1990, pp. 1–17.
7. P. P. Gelsinger, et al., "Microprocessors circa 2000," *IEEE Spectrum*, October 1989, pp. 43–47.
8. J. S. Mayo, "Materials for Information and Communication," *Scientific American*, Vol. 255, No. 4, October 1986, pp. 58–65.
9. B. T. Murphy, "Cost-Size Optima of Monolithic Integrated Circuits," *Proceeding of the IEEE*, Vol. 52, No. 12, December 1964, pp. 58–65.
10. R. N. Noyce, "Microelectronics," *Scientific American*, Vol. 237, No. 3, September 1977, pp. 63–69.
11. J. D. Meindl, "Ultra Large Scale Integration," *IEEE Transactions on Electronic Devices*, Vol. ED-31, No. 11, p. 1555.
12. C. M. Osburn, "The Effect of Contamination on Semiconductor Manufacturing Yield," *Journal of Environmental Science*, March/April 1964, p. 45.
13. W. J. Bertram, "Yield and Reliability," *VLSI Technology*, 2nd ed., edited by S. M. Sze, McGraw-Hill, 1988, ISBN 0-07-062735-5, pp. 612–655.
14. E. F. Rent, in two IBM internal memoranda, (1960), quoted by B. S. Landman and R. L. Russo, "On a Pin versus Block Relationship for Partitions of Logic Graphs," *IEEE Transactions on Computers*, Vol. C-20, December 1971, pp. 1469–1479.
15. C. A. Warwick and A. Ourmazd, "Trends and Limits in Monolithic Integration by Increasing the Die Area," *IEEE Transactions on Semiconductor Manufacturing*, Vol. 6, No. 3, August 1993, pp. 284–289.
16. R. W. Lucky, *Silicon Dreams*, St. Martin's Press, p. 151.
17. R. -H. Yan, et al., "89-GHz f_T Room-Temperature Silicon MOSFETs," *IEEE Electronic Device Letters*, Vol. 13, May 1992, pp. 256–258.
18. R. -H. Yan, et al., "Ultra-Deep Submicron Si MOSFETs with f_T Exceeding 100 GHz," *Proceedings on Ultrafast Electronics and Optoelectronics, 1993*, San Francisco, California, February 1993.
19. C. Raynaud, et al., "High-Frequency Performance of Submicrometer Channel-Length Silicon MOSFETs," *IEEE Electronic Device Letters*, Vol. 12, December 1991, pp. 667–669.
20. A. E. Schmitz, et al., "A Deep Submicron Microwave/Digital CMOS/SOS Technology," *IEEE Electronic Device Letters*, Vol. 12, January 1991, pp. 16–17.
21. M. Nanba, et al., "A 64GHz Si Bipolar Transistor Using In-Situ Phosphorus Doped Polysilicon Emitter Technology," *Technical Digest of the International Electron Devices Meeting, 1991*, pp. 443–446.
22. G. L. Patton, et al., "75-GHz f_T SiGe-Base Heterojunction Bipolar Transistors," *IEEE Electronic Device Letters*, Vol. 11, April 1990, pp. 171–173.

(Manuscript approved September 1993)

Colin A. Warwick *is a member of technical staff in the Microphysics Research Department at AT&T Bell Laboratories in Holmdel, New Jersey. He works on battery-powered digital circuits for personal communicators and studies trends in the semiconductor industry. Mr. Warwick joined AT&T in 1988, after receiving a B.S. in physics and a Ph.D. in physical sciences, both from the University of Oxford, the United Kingdom.*

Ran-Hong Yan *is a member of technical staff in the Microphysics Research Department at AT&T Bell Laboratories in Holmdel, New Jersey. He is working on 0.1-μm CMOS transistor structures, processing technology, and circuit design issues; low-power electronics, from processing to system; and very large-scale integration (VLSI) implementation for signal processing and computation. Mr. Yan joined AT&T in 1990, after receiving a B.S. in electrical engineering from the National Taiwan University, and an M.S. and Ph.D. in electrical and computer engineering from the University of California, Santa Barbara.*

Young Ock Kim *is a member of technical staff in the Microphysics Research Department at AT&T Bell Laboratories in Holmdel, New Jersey. She works on the characterization of semiconducting materials and devices, using electro-optical equipment, including high-resolution transmission and scanning electron microscopes. Ms. Kim received a Ph.D. in solid state physics from Arizona State University, Tempe, Arizona. She joined AT&T in 1988.*

Abbas Ourmazd *is head of the Microphysics Research Department at AT&T Bell Laboratories in Holmdel, New Jersey. He is responsible for deep submicron VLSI and low-power electronics. Mr. Ourmazd joined AT&T in 1984, after receiving a B.A. from Corpus Christi College, Oxford, United Kingdom, and a Ph.D. from Wolfson College, Oxford, United Kingdom, both in physics. Mr. Ourmazd was an Alexander von Humboldt Fellow at Gottingen University, Germany.*