
1 Introduction

The scaling of semiconductor process technologies has been continuing for more than four decades. Advancements in process technologies are the fuel that has been moving the semiconductor industry. In response to growing customer demand for enhanced performance and functionality at reduced cost, a new process technology generation has been introduced by the semiconductor industry every two to three years during the past four decades [1]. Both the performance and the complexity of integrated circuits have grown dramatically since the invention of the integrated circuit in 1959. Microphotographs of the first monolithic integrated circuit (Fairchild Semiconductor, 1959), the first microprocessor (Intel 4004, 1971), and a recent microprocessor (Intel Pentium 4, 2002) are shown in Figure 1.1.

Technology scaling reduces the delay of the circuit elements, enhancing the operating frequency of an integrated circuit (IC) [1]–[5]. The density and number of transistors on an IC are increased by scaling the feature size. By utilizing this growing number of available transistors in each new process technology, novel circuit techniques and microarchitectures can be employed, further enhancing the performance of the ICs beyond the levels made possible by simply scaling (or shrinking) a previous generation [1]–[7]. The price for these performance and functional enhancements has traditionally been increased design complexity and power consumption. The generation, distribution, and dissipation of power are now at the forefront of current problems faced by IC designers.

Historically, circuit techniques and architectures employed during the evolution of the IC have followed two different paths. For a group of technologies, enhancing speed has been at the core of the design process. This class of ICs represents the high end of the performance spectrum. In this high end arena, increasing clock frequency and die size and the widespread use of power-hungry circuit techniques and microarchitectures (with continuously increasing levels of speculative execution often translated into an inefficient use of energy) have increased power consumption many fold over the years [2], [3], [7]. Until recently, the removal of heat in high performance ICs was handled by inexpensive packaging solutions, passive heat sinks, and air fans. With the power dissipation of ICs rising well above 100 W, however, more expensive packaging and cooling solutions such as liquid cooling or refrigeration hardware will soon be required [2]–[10]. Issues related to power dissipation and heat removal are likely to be the primary cause of the end to the trend of continuously decreasing price to performance ratios of high performance ICs.

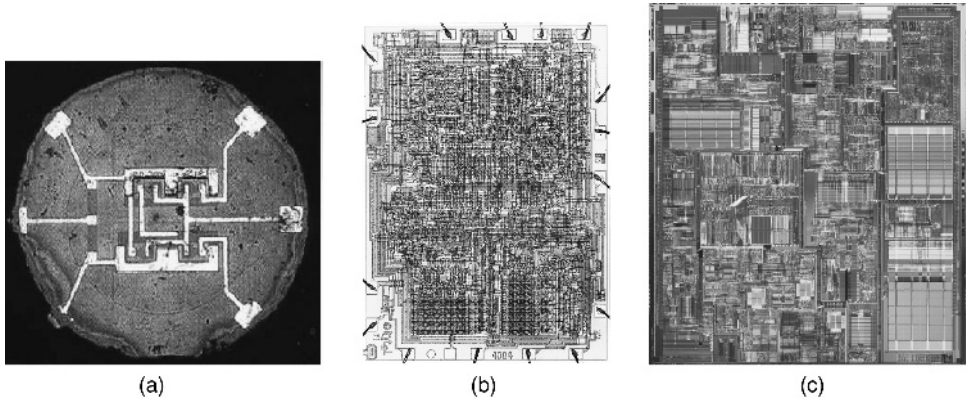


Figure 1.1 Microphotographs of three landmark ICs from the evolution of the IC technology (the sizes of the dies are not to scale). (a) The first monolithic integrated circuit, Fairchild Semiconductor (1959). (b) The first microprocessor, Intel 4004 (1971). (c) A recent Intel Pentium 4 microprocessor (2002)

Another important group of ICs has emerged as a result of customer demand for miniaturization and portability. Portable devices, until recently, represented the low end of the performance spectrum with power constraints always dominating over speed [4], [6], [9]. Extended battery life and reduced system cost constraints drove the portable equipment design process until the 1990s. However, since the 1990s, strong customer demand has been growing for higher performance (for high speed computing and data transfer) and a wider variety of applications in portable equipment. Today, people expect from their portable devices almost the same computing capability as a desktop system.

While the performance of mobile devices continues to advance at a fast pace in accordance with general semiconductor technology trends, the evolution of battery technologies has progressed at a much slower pace [4], [9], [11]. Before rechargeable battery technologies evolved to offer sufficient energy in a miniaturized volume, standard disposable alkaline battery technology was the popular power solution. Frequent battery purchases coupled with the inconvenience of carrying replacement batteries increased the market demand for a rechargeable battery solution. Nickel-cadmium (Ni-Cd) chemistry (invented in 1899 [11]) became the battery supply for portable devices toward the end of the 1980s. Ni-Cd was replaced by nickel-metal-hydride (Ni-M-H) chemistry during the mid-1990s. Ni-M-H batteries offer twice the energy density with faster charging times as compared to Ni-Cd batteries [11]. Lithium-ion (Li-ion) batteries (first introduced in the early 1990s) gradually replaced the Ni-M-H technology toward the end of the last decade. Li-ion, with enhanced energy density characteristics as compared to both Ni-Cd and Ni-M-H batteries, is the most widely used battery technology today [11].

Vendors have responded to the continuous market demand for greater functionality and higher processing speed while continuing to decrease the physical size and weight of portable devices. Batteries are, therefore, required to offer increasing amounts of energy while occupying smaller volumes as semiconductor technologies progress [11]. Today, the lack of a low cost, small volume, and lightweight battery technology with a higher energy density as compared to the Li-ion technology is a primary limitation to further advancements in portable IC technologies.

Traditional circuits and architectures in high performance ICs, because of the power-hungry characteristics of these technologies, are not applicable to those ICs designed for portable systems. Alternatively, circuits and architectures that have been developed for portable devices, because of the typical low throughput characteristics of these technologies, are not effective in high performance ICs. Today, the IC industry is experiencing a shift in requirements at both the high performance and portability ends of the market. Power dissipation is no longer a secondary issue in high performance ICs. Similarly, enhancing throughput is as important as lowering the power, area, and weight in many portable devices. Energy-efficient semiconductor devices, circuit techniques, and microarchitectures are necessary to maintain the pace of expansion that the semiconductor industry has been enjoying for the past forty years [1]–[12].

In retrospect, the invention of the transistor in 1947 can be seen as the first step toward low power electronics. Operation of a vacuum tube requires hundreds of volts of anode voltage and a few watts of power. In comparison, a transistor operates at a higher speed and at a significantly lower supply voltage and consumes orders of magnitude smaller power. Similarly, the invention of the IC in the late 1950s can be seen as the first step toward low power microelectronics. ICs consume less power, are lower weight, and occupy smaller volume while offering the same functionality, with enhanced performance and reliability as compared to circuits composed of discrete devices [13], [14]. These trends that shaped the evolution of IC technology are reviewed in Section 1.1. An outline of this book is summarized in Section 1.2.

1.1 EVOLUTION OF INTEGRATED CIRCUITS

The monolithic IC was invented in 1959. The primary reasons for implementing certain functions as ICs were to lower the weight and size while enhancing the reliability and performance characteristics as compared to circuits composed of discrete devices [13]. ICs were an expensive technology during the 1960s, limiting the use of ICs to specific military applications with severe requirements of weight, size, and reliability. Gordon Moore noticed in 1965, only six years after the birth of the very first IC, that the unit costs of ICs were steadily decreasing as technology evolved and fabrication techniques matured [13], [14]. Moore saw that shrinking transistor sizes, increasing manufacturing yield, and larger wafer and die sizes would make ICs increasingly cheaper, more powerful, and more plentiful. As Moore declared in 1965, ‘the future of integrated electronics turned out to be the future of electronics itself’ [13]. Advances in IC technology enabled the so-called ‘information age’ that is experienced today. A timeline of some of the key events that led to the invention and advancements of IC technologies is provided in Figure 1.2.

The general form of Moore’s law is depicted in Figure 1.3 [13]. As more components are added to an IC at a particular process technology generation (or technology node), the relative manufacturing cost per component decreases (assuming that the same semiconductor material and the same package are used to incorporate additional components) [13]. However, as more components are integrated onto the same die, the complexity (at the circuits, physical design, and process levels) increases, degrading yield. There is, therefore, an optimum number of components per IC that minimizes the total manufacturing cost at any generation in the evolution of an IC technology [13]. The unit price of a transistor decreases as the device dimensions scale, defect densities are reduced, and wafer and die sizes grow [13], [14]. The

4 INTRODUCTION

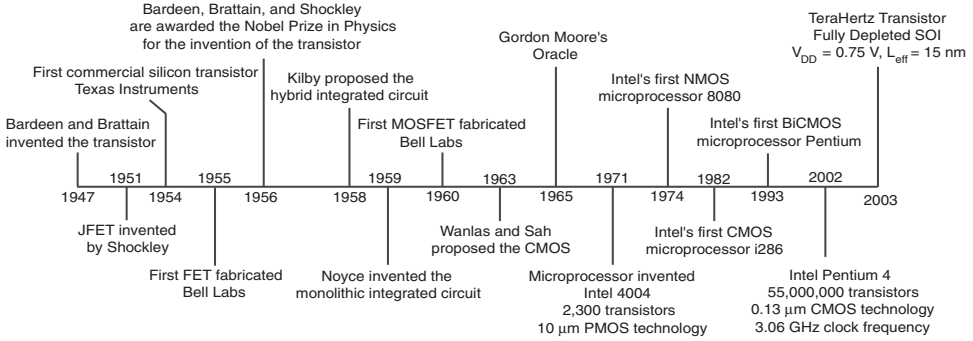


Figure 1.2 A timeline of some of the key events during the evolution of semiconductor technologies

optimum number of transistors that minimizes the total manufacturing cost, therefore, increases from one technology generation to the next as shown in Figure 1.3. The total number of transistors that can be integrated onto a piece of semiconductor material has increased by more than a million times since the mid-1960s, verifying the trends Moore observed in 1965. What began as an observation has become both the compass and engine, setting the bar for the semiconductor industry over the past four decades.

High performance microprocessors currently represent the front end of the market demand for enhanced performance and functionality. No IC technology has witnessed the employment of more aggressive semiconductor process technologies, circuits, and architectures as compared to high performance microprocessors [1], [2]. The high performance microprocessor and high density random access memory (RAM) industries have, historically, led the advances in semiconductor technology and hence encountered the side effects of the technology evolution before any other portion of the semiconductor industry.

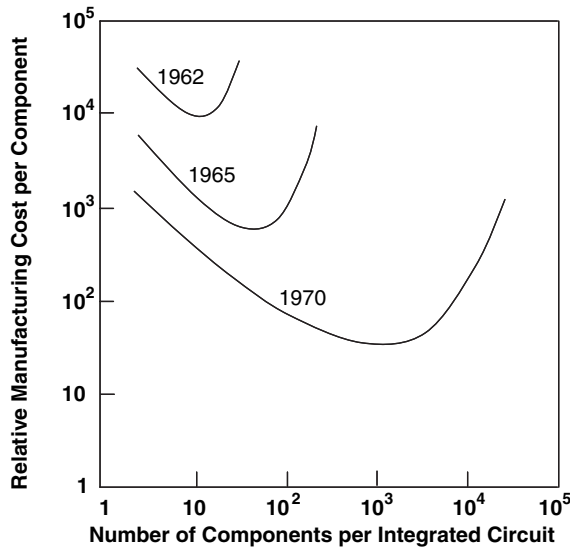


Figure 1.3 The general form of Moore' law [13]

Table 1.1 Technological Trends of High Performance Microprocessors

Vendor	DEC Alpha			AMD			IBM/MOTOROLA PowerPC		
Microprocessor	21064	21164	21264	K5	K6	K7	750	7400	7450
Technology (μm)	0.75	0.5	0.35	0.5	0.3	0.25	0.29	0.2	0.18
Frequency (MHz)	200	300	600	75	233	500	266	400	667
Die area (mm^2)	234	299	314	251	162	184	67	83	106
Transistor count (millions)	1.68	9.3	15.2	4.3	8.8	22	6.35	10.5	33
Supply voltage (V)	3.3	3.3	2.2	3.5	3.3	1.6	2.6	1.8	1.6
Supply current (A)	9.1	15.2	32.7	3.3	9.2	26.3	3	6.3	11.9
Power (W)	30	50	72	11.6	30.2	42	7.9	11.3	19
Power Density (W/cm^2)	12.8	16.7	22.9	4.6	18.6	22.8	11.8	13.6	17.9

The focus of this section is on the advancements of high performance microprocessor technologies. The technological trends in the evolution of the lead Intel microprocessors will be examined. The choice of the lead microprocessor product line of Intel Corporation is due to the significant role that the company has played in the semiconductor industry during the past 35 years. Similar technological trends can also be observed in other leading vendor product lines. Common trends in the characteristics of some of the technological parameters among different microprocessor generations for three vendors are listed in Table 1.1.

The primary force shaping the IC evolution is the advancing fabrication technology that permits technology scaling [1]–[9]. The feature size of the transistors and interconnect have continually been scaled, increasing the integration density in each new process technology generation. The minimum feature size of the transistors in the lead Intel microprocessors has decreased from $10\ \mu\text{m}$ in 1971 to $0.13\ \mu\text{m}$ in 2002 as shown in Figure 1.4. The second primary development behind the IC evolution is the reduction of defect densities due to the maturing fabrication technology, thereby making larger dies (individual ICs or chips) economical. Die areas have grown steadily by about 14% per year from 1971 to 1995 as shown in Figure 1.5. Starting in the mid-1990s, however, limits to further increases in die size became necessary due to concerns about increasing power consumption and soaring fabrication and packaging costs [3], [5]. As a result of the reduced physical dimensions of the transistors and the increased die area, the total number of transistors in the lead Intel microprocessors has increased by twenty four thousand times over the past three decades, as shown in Figure 1.4.

The increasing number of transistors per IC in each new process technology generation offers more tools for enhancing circuit performance and functionality. The propagation delays are reduced as the physical dimensions of the transistors are scaled. Enhancements related to technology scaling coupled with advances in circuits and microarchitectures (such as deeper pipelining, superscalar, and out-of-order execution) have significantly increased the performance of ICs [1]–[8]. As shown in Figure 1.6, the operating frequency of the lead Intel microprocessors has increased by more than twenty eight thousand times since the introduction of the first microprocessor (Intel 4004) in 1971.

Ideal scaling theory suggests shrinking all of the voltages, currents, and physical dimensions and increasing all of the doping concentrations by the same scaling factor (λ) to maintain constant electric fields within a device [40], [56]. Historically, however, the voltages and currents have been scaled at a lower rate as compared to the physical dimensions. The electric




Figure 1.4 Scaling of the minimum feature size and the increasing total number of transistors within each lead Intel microprocessor




Figure 1.5 Die area of lead Intel microprocessors




Figure 1.6 Operating frequency and supply voltage of lead Intel microprocessors

fields within the devices have, therefore, significantly increased. An important reason for the reluctance to scale the voltages and currents as rapidly as the physical dimensions has been the beneficial effect of increasing electric fields on device performance [40], [56]. An equally important reason for the slow pace of supply voltage scaling has been the need to maintain high noise margins for maintaining reliability in a difficult-to-control noisy on-chip environment.

The period of technology scaling, since the invention of the first IC, is divided into two primary eras depending upon the characteristics of the supply voltage in a scaled technology as compared to a preceding technology generation. The supply voltage in the first three Intel microprocessor generations was 12 V as shown in Figure 1.6. Starting with the 3 μm technology node, the supply voltage was reduced to 5 V. IC supply voltages were maintained at 5 V until the 0.8 μm technology node was commercialized during the early 1990s (see Figure 1.6). At the 0.8 μm technology node, supply voltage scaling became an essential part of the technology scaling process due to transistor reliability and power consumption concerns [3]–[5], [15]. The era (until 1993 in the case of Intel) during which supply voltage scaling was not necessarily a part of technology scaling is called the constant voltage scaling era. The technology scaling era (after 1993 in the case of Intel), during which supply voltage scaling occurs with scaling of the other device parameters, is called the constant field scaling era [3], [4], [15]. Constant field scaling arises from the concept that the supply voltage for a new technology is ideally chosen to maintain constant electric fields between the terminals of a transistor [15]. The need to slow the rate of increase in power consumption became an increasingly important factor in supply voltage scaling toward the end of the 1990s. Today, the requirements for lowering the power consumption and improving device reliability




Figure 1.7 Maximum power consumption of lead Intel microprocessors

together with circuit speed determine the rate of supply voltage scaling in each new technology generation [1]–[5], [7], [8], [11], [15], [16].

An increase in the operating frequency and die size (due to the greater number of transistors for the additional circuitry and novel microarchitectures) not only enhances the speed, but also increases the power consumption [1]–[6], [8]–[10], [16], [17]. As shown in Figure 1.7, the power consumption of the lead Intel microprocessors has been increasing over the past 30 years. The technology in the first two Intel microprocessor generations was p-channel metal oxide semiconductor (PMOS). Starting with the Intel 8080, n-channel metal oxide semiconductor (NMOS) became the preferred technology due to the speed and area advantages of NMOS transistors as compared to PMOS transistors. NMOS circuits, however, suffered from higher static DC power consumption and lower noise margins [18], [20]. By the end of the 1970s, scaling of NMOS technology became increasingly difficult as the low noise margins of the NMOS circuits did not permit supply voltage scaling to accompany scaling of the feature size [18]. The increasing number of transistors operating at higher clock frequencies at a high supply voltage coupled with the intrinsic static DC power consumption of the NMOS circuits set the stage for the end of a decade-long dominance of NMOS as the technology of choice. As shown in Figure 1.8, the power density of the last NMOS Intel microprocessor (the i8086 that was commercialized in 1978) is similar to the power density of a kitchen hot plate. The packaging and cooling technologies available at the beginning of the 1980s were quite limited, permitting no further technological advances that would lead to an increase in power dissipation.

The complementary metal oxide semiconductor (CMOS) circuit topology (first proposed in 1963 by Wanlass and Sah [165]) was adapted by the IC industry in the early 1980s due to the




Figure 1.8 Power density trends of lead Intel microprocessors

intrinsically lower power consumption and enhanced scaling characteristics of CMOS as compared to NMOS [18]–[20]. The higher noise margins in CMOS circuits made possible supply voltage scaling that accelerated in the 1990s, enhancing both transistor reliability and energy efficiency. CMOS became the preferred circuit topology in the lead Intel microprocessors starting with the i286 (introduced in 1982). The transition from NMOS to CMOS reduced both the power consumption and the power density of the Intel microprocessors as shown in Figures 1.7 and 1.8, respectively [5].

The reduction in the power dissipation of high performance microprocessors due to the transition to CMOS, however, provided only temporary relief. Maintaining the approach of employing higher clock frequencies coupled with power-hungry circuits and highly speculative architectures in order to achieve enhanced performance, the power consumption and power density of the post-NMOS era (i.e., CMOS and BiCMOS) ICs were, once again, pushed to higher levels. As illustrated in Figures 1.7 and 1.8, respectively, both the power consumption and power density of the lead Intel microprocessors (with the exception of the first generation Pentium 3) have been increasing since the introduction of the second generation CMOS microprocessor (i386) in 1985. As depicted in Figure 1.8, the power density of current high performance microprocessors has greatly exceeded the power density of the heating coil of a kitchen hot plate [2], [3], [5], [21].

The temperature of a die is controlled to maintain proper operation of the circuitry compliant with technical specifications [5], [10], [22]. Thermal management of high performance ICs has become increasingly difficult due to the continuously increasing power dissipation and power density in each new process technology generation [2]–[5], [10], [12], [16], [17], [21], [23]. Within a few technology generations, traditional cooling solutions such as low cost heat sinks and air flow fans will become ineffective for thermal management [2]–[5], [10], [22]. If the current trend in the rate of increase in the power levels continues, ICs




Figure 1.9 Increasing contribution of leakage currents to the total power consumption of the lead Intel microprocessors [161]

will consume thousands of watts of power in the near future [2], [5], [21]. The power density of a high performance microprocessor will, within the next decade, exceed the power density levels encountered in typical rocket nozzles [2], [5]. Low cost cooling solutions that can handle power densities in excess of nuclear reactors or rocket nozzles do not presently exist for ICs. As acknowledged by many designers and researchers, excessive power dissipation has emerged as the single greatest jeopardy to further advances in IC technologies [1]–[14].

Dynamic switching power consumption has typically been the dominant source of power consumption in CMOS ICs. Recently, however, leakage power has become a significant portion of the total power consumption in high complexity CMOS ICs, as illustrated in Figure 1.9. Ideally, an MOS switch has infinite input impedance. Similarly, an ideal cut-off transistor has infinite drain-to-source resistance. However, in reality, an active transistor has a finite input impedance and a cut-off transistor has a finite channel resistance, producing gate oxide and subthreshold leakage current, respectively. Due to the aggressive scaling of the threshold voltages and the thickness of the gate dielectric layer in order to enhance device speed, modern MOSFETs no longer resemble, even remotely, an ideal switch. As illustrated in Figure 1.9, subthreshold and gate oxide leakage currents will become the dominant source of power consumption in the near future.

Another important challenge directly linked to the advances of semiconductor technologies is maintaining the reliability of scaled CMOS circuits. The reliability of CMOS ICs has degraded due to scaling the device and interconnect dimensions and the on-chip voltage levels. Error-free operation of CMOS circuits has become increasingly challenging as IC technologies evolve. CMOS ICs have become more sensitive to noise while on-chip noise levels continue to rise with each new technology generation. Various sources of noise in a microprocessor are schematically illustrated in Figure 1.10. The clock distribution network acts as a source of noise to the surrounding circuitry and interconnect lines. On-chip clock generators inject considerable amounts of noise into the substrate. Similarly, a monolithic

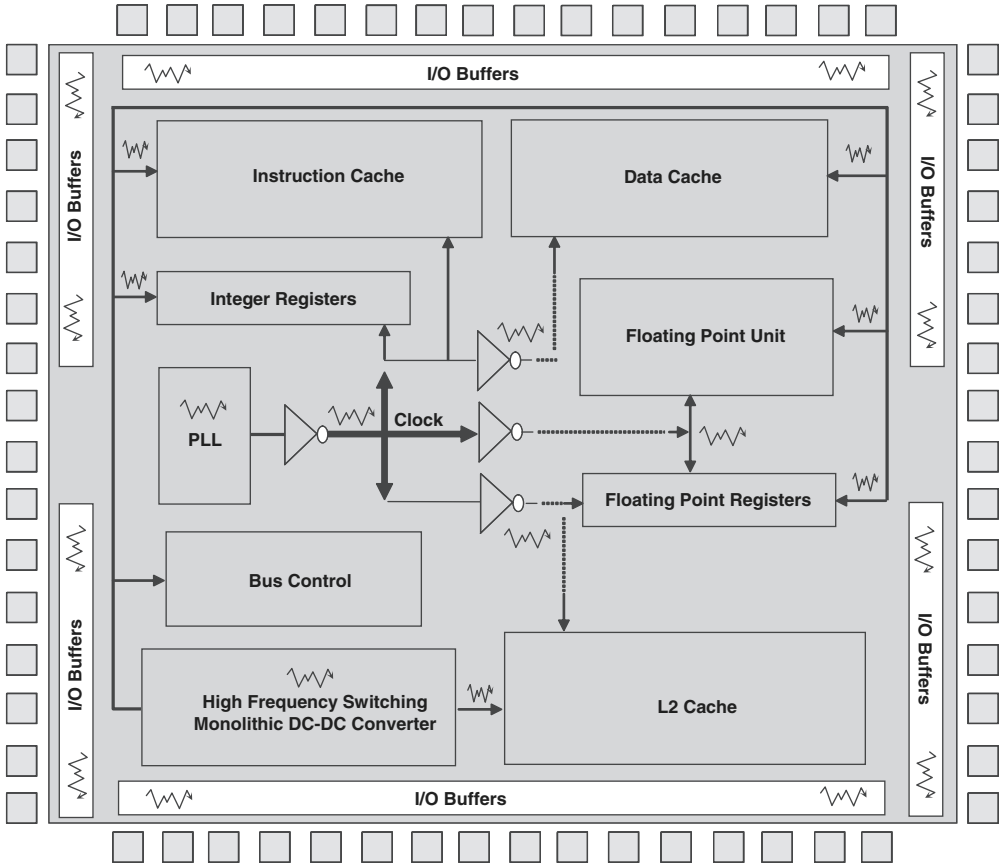


Figure 1.10 Various sources of noise in a microprocessor

switching DC–DC converter can produce significant noise on a microprocessor die, as illustrated in Figure 1.10.

An important source of noise in CMOS ICs is interconnect coupling noise [158], [159]. Due to increasing device densities, interconnect lines are physically closer with each new technology generation, as illustrated in Figure 1.11. The resistance of the interconnect lines increases as the width of the interconnect lines is reduced with technology scaling. Due to the increasing resistance of the interconnect lines, the delay due to the interconnect rather than the gate delays dominates the propagation delay characteristics in current CMOS circuits, as illustrated in Figure 1.12. The higher interconnect resistance also increases the parasitic power dissipation. In order to limit the increased resistance of the interconnect lines, the height of the interconnect lines is scaled at a much smaller rate as compared to the width with each new technology generation. The aspect ratio, therefore, increases significantly, thereby increasing the coupling capacitance between adjacent interconnect lines on the same metal layer. Similarly, due to the vertical scaling trend of adding more metal layers to CMOS fabrication processes, the intermetal layer coupling capacitances also tend to increase. Noise generated on a quiescent line (victim line) due to the coupling capacitance with a nearby line (aggressor




Figure 1.11 Effect of technology scaling on the physical geometries of the interconnect lines

line) can cause erroneous transitions, degrade speed, produce excessive power consumption, and cause a circuit to malfunction.

The power distribution network is another important source of noise in deeply scaled nanometer CMOS ICs. An important factor exacerbating the on-chip noise issue is the increasing current demand of modern ICs. While the power consumption of the ICs continues to increase, the supply voltages have been reduced, as shown in Figure 1.6. The supply current, therefore, increases, as shown in Figure 1.13. Increased current demand of ICs

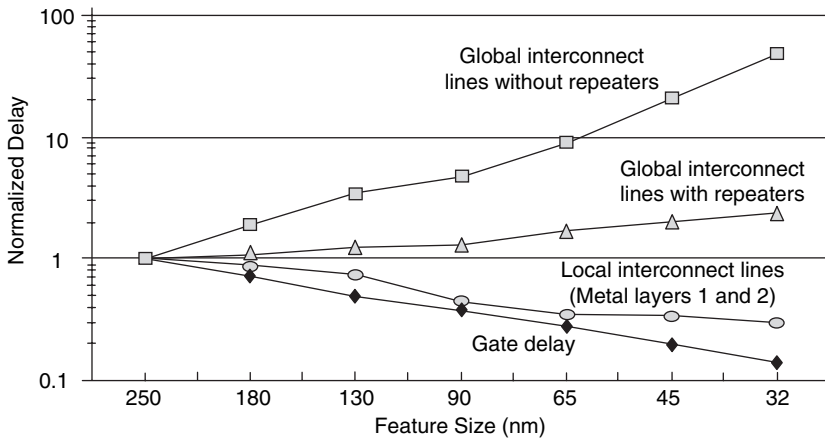


Figure 1.12 Effect of technology scaling on interconnect and gate delays [162]




Figure 1.13 Increasing current demand of lead Intel microprocessors

coupled with scaled wire dimensions create metal migration and voltage drop problems within the power distribution network [21], [24], [25].

Power supply noise has both low frequency and high frequency components [158]. The low frequency component of the power supply noise is due to the resistive IR drops on the printed circuit board, within the package, and along the on-chip power grid. The tolerance of ICs to voltage fluctuations in the power supply grid is typically reduced while the resistance of the interconnect is increased with technology scaling [18], [19], [24]. The resistive voltage drop of the power distribution grid, therefore, has become an increasingly important concern for maintaining performance and reliability. Alternatively, the high frequency components of the power supply noise are due to the inductance of the printed circuit board planes, package, and the on-chip power grid. Current slew rates typically increase due to the higher operating frequencies as well as the growing current demand in each new technology generation. Simultaneous switching noise ($L di/dt$) due to the inductance of a power distribution grid affects the supply voltage, thereby degrading the performance and possibly causing circuit malfunctions [25], [164].

Power generation, delivery, and dissipation are primary limitations to further advancements of IC technologies [1]–[9], [16]–[18], [21], [23]. In order to continue to reduce the unit cost of an IC while simultaneously enhancing the performance and functionality, radical changes are required in the way ICs have been designed during the past three decades. Higher performance at all costs is no longer an option. Novel energy-efficient devices, circuits, microarchitectures, and macroarchitectures must be developed to lower the rate of increase in the power consumed by next generation ICs.

1.2 OUTLINE OF THE BOOK

Several new techniques for the design of low power and high performance ICs are described in this book. Particular emphasis is placed on issues related to the scaling of the supply and threshold voltages in high performance ICs.

An analysis of power dissipation-related problems faced by the semiconductor industry starts with identifying the sources of power consumption. The primary sources of power consumption in CMOS ICs are described in Chapter 2. Specifically, dynamic, short-circuit, leakage, and static DC power components are individually described.

Supply and threshold voltage scaling techniques, aimed at lowering power consumption and enhancing device reliability without degrading performance, are discussed in Chapter 3. The importance of supply voltage scaling is discussed from an energy efficiency point of view. As the supply voltage is reduced, the performance of an IC degrades due to reduced transistor currents [27]. Systems with multiple supply voltages can minimize the degradation in speed while reducing power by selectively lowering the supply voltages along non-critical delay paths [28]. Dynamic and static versions of multiple supply voltage IC design techniques are reviewed. Another alternative technique for reducing the impact of supply voltage scaling on circuit performance is threshold voltage scaling. During the past decade, threshold voltage scaling has accelerated together with scaling of the supply voltages. At reduced threshold voltages, however, subthreshold leakage currents increase. Supply voltage scaling when coupled with threshold voltage reduction, therefore, increases the leakage power while reducing the dynamic switching power. Multiple threshold voltage circuits reduce leakage currents while enhancing performance by selectively lowering the threshold voltages only on speed-critical paths [29]. Dynamic threshold voltage scaling (V_T -hopping) and multiple threshold voltage CMOS circuit techniques are reviewed in Chapter 3. Dynamic and static versions of multiple threshold voltage circuit techniques are also discussed in this chapter.

A significant issue with threshold voltage and device scaling is the increasing effect of die-to-die and within-die parameter variations on the speed and power dissipation characteristics of CMOS ICs. Die-to-die and within-die fluctuations of the critical dimensions (gate length, gate oxide thickness, and junction depletion width) effectively increase with technology scaling. Moreover, the sensitivity of the threshold voltage to variations in the critical dimensions is greater due to increasing short-channel effects as the gate length and threshold voltage are both reduced with technology scaling. Process variations cause ICs to exhibit different speed and power characteristics. The electrical characteristics of a CMOS circuit fabricated in a deep submicrometer process technology have become increasingly probabilistic (less deterministic). The number of individual dies that satisfy a target clock frequency and maximum power dissipation constraint is lower, degrading the yield. The increasing cost of fabricating deep submicrometer ICs is, therefore, further aggravated by lower yields caused by greater process variations. Challenges imposed by these parameter variations are also addressed in Chapter 3.

The generation and distribution of the energy required for the proper functioning of an IC are important challenges due to system-level power budget limitations and circuit reliability issues. Increasing supply currents together with reduced supply voltages degrade the energy efficiency and voltage quality of power generation and distribution networks in high performance ICs. Energy-efficient low voltage monolithic DC–DC conversion and voltage regulation techniques are developed in this book. Before presenting the details of these

monolithic DC–DC conversion techniques in the following chapters, a basic background to DC–DC conversion and a review of several widely employed types of low voltage DC–DC converters are presented in Chapter 4.

In single power supply microprocessors, the primary power supply is typically an external (non-integrated) buck converter. In a typical non-integrated switching DC–DC converter, significant energy is dissipated by the parasitic impedances of the interconnect among the non-integrated devices (the filter inductor, filter capacitor, power transistors, and pulse width modulation circuitry) [9], [26], [30], [31]. Moreover, the devices of a discrete DC–DC converter are typically fabricated in older technologies with poor parasitic impedance characteristics. Integrating a DC–DC converter onto a microprocessor can potentially lower the parasitic losses as the interconnect between (and within) the DC–DC converter and the microprocessor is reduced. Additional energy savings can be realized by utilizing advanced deep submicrometer fabrication technologies with lower parasitic impedances. The efficiency attainable with a monolithic DC–DC converter can therefore be higher than a non-integrated DC–DC converter [30]. An analysis of on-chip buck converters is presented in Chapter 5. A model of the parasitic impedances of a buck converter is described. With this model, a design space is determined that supports the integration of active and passive devices on the same die for a target technology. A monolithic, high efficiency, and high frequency switching DC–DC converter with an integrated inductor on the same die as a dual supply voltage microprocessor is demonstrated to be feasible.

The model presented in Chapter 5 provides an accurate representation of the parasitic losses of a full voltage swing DC–DC converter (with an error of less than 2.4% as compared to simulation). A high switching frequency is the key design parameter that enables the full integration of a high efficiency DC–DC converter. At these high switching frequencies, the energy dissipated in the power MOSFETs and gate drivers dominates the total losses of a DC–DC converter. The efficiency can, therefore, be enhanced by applying a variety of MOSFET power reduction techniques [31]. A low swing MOSFET gate drive technique is described in Chapter 6 that enhances the efficiency of a DC–DC converter. An advanced circuit model for low swing circuit optimization is also presented. The gate voltages and transistor sizes are included as independent parameters in this model. The optimum gate voltage swing of a power MOSFET that maximizes efficiency is shown to be lower than a standard full voltage swing. Lowering the input and output voltage swing of a power MOSFET gate driver is shown to be effective for enhancing the efficiency characteristics of a DC–DC converter.

Due to the advantages of high voltage power delivery on a circuit board and monolithic DC–DC conversion, next generation low voltage and high power microprocessors are likely to require high input voltage, large step-down DC–DC converters monolithically integrated onto the same die. The voltage conversion ratios attainable with standard non-isolated switching DC–DC converter circuits are limited, however, due to MOSFET reliability issues. Provided that a DC–DC converter is integrated onto the same die as a microprocessor (fabricated in a low voltage nanometer CMOS technology), the range of input voltages that can be applied to a standard DC–DC converter circuit is further reduced. A standard non-isolated switching DC–DC converter topology such as the buck converter circuit discussed in Chapters 5 and 6 is, therefore, not suitable for providing a high voltage conversion ratio in future high performance ICs. Three cascode bridge circuits that can be used in monolithic DC–DC converters that provide high voltage conversion ratios are presented in Chapter 7. The circuits ensure that the voltages across the terminals of all of the MOSFETs in a DC–DC

converter are maintained within the limits imposed by available low voltage CMOS technologies.

In ICs with multiple supply voltages, signal transfer among the regions operating at different voltage levels requires specialized voltage interface circuits [32]. Another low power circuit technique that requires voltage-level conversion is low swing interconnect signaling. At each new IC generation, the relative amount of interconnect increases due to the greater number of transistors and the larger die size. In many recent systems, charging and discharging these interconnect lines can require more than 50% of the total power consumed on-chip. In certain programmable logic devices, more than 90% of the total power consumption is due to the interconnect wires [32]. Decreasing the signal voltage swing on the interconnect can significantly decrease the power consumption. In a low swing interconnect architecture, voltage-level converters are placed at the driver and receiver ends of the low swing interconnect to change the voltage levels. A bidirectional CMOS voltage interface circuit that drives high capacitive loads to full swing at high speed while consuming no static DC power is presented in Chapter 8. The propagation delay, power consumption, and power efficiency characteristics of this voltage interface circuit are compared to other interface circuits described in the literature. The voltage interface circuit offers significant power savings and lower propagation delay.

Domino logic circuit techniques are extensively applied in high performance microprocessors due to the superior speed and area characteristics of dynamic CMOS circuits as compared to static CMOS circuits. High speed operation of domino logic circuits is primarily due to the lower switching threshold voltage of dynamic circuits as compared to static gates. This property of a lower switching threshold voltage, however, makes domino logic circuits highly sensitive to noise as compared to static gates. On-chip noise becomes more severe with technology scaling and higher operating frequencies. Furthermore, the noise sensitivity of domino logic circuits increases with technology scaling. Error-free operation of domino logic circuits has, therefore, become a major challenge [33]. A variable threshold voltage keeper circuit technique is presented in Chapter 9 for simultaneous power reduction and speed enhancement of domino logic circuits. The threshold voltage of the keeper transistor is dynamically modified during circuit operation to reduce the contention current without sacrificing noise immunity. The variable threshold voltage keeper circuit technique is shown to enhance circuit evaluation speed by up to 60% while reducing power consumption by 35% as compared to a standard domino logic circuit. The keeper size can be increased while preserving the same delay or power characteristics as compared to a standard domino circuit since the contention current is reduced with this technique. The domino logic circuit technique offers 14.1%, 8.9%, or 11.9% higher noise immunity under the same delay, power, or power–delay product conditions, respectively, as compared to the standard domino logic circuit technique. Forward body biasing the keeper transistor is also described for improved noise immunity as compared to a standard domino circuit with the same keeper size. It is shown that by applying forward and reverse body bias circuit techniques, the noise immunity and evaluation speed of domino logic circuits are both enhanced.

The subthreshold leakage current of a domino logic circuit can vary dramatically with the voltage state of the dynamic and output nodes. A quantitative review of the subthreshold leakage current characteristics of standard low threshold voltage and dual threshold voltage domino logic circuits is presented in Chapter 10. Different subthreshold leakage current conduction paths which exist depending upon whether the dynamic node is charged or discharged are identified. A discharged dynamic node is preferable for reducing leakage

current in a dual threshold voltage circuit. Alternatively, a charged dynamic node is preferred for lower subthreshold leakage current in a standard low threshold voltage domino logic circuit with stacked pull-down devices, such as an AND gate. The effect of dual threshold voltage CMOS technologies on the noise immunity characteristics of domino logic circuits is also evaluated in Chapter 10.

A circuit technique is presented in Chapter 11 for exploiting the dynamic node voltage-dependent asymmetry of the subthreshold leakage current characteristics of domino logic circuits. Sleep switch transistors are used to place an idle dual threshold voltage domino logic circuit into a low subthreshold leakage state. The circuit technique enhances the effectiveness of a dual threshold voltage CMOS technology to reduce subthreshold leakage current by strongly turning off all of the high threshold voltage transistors. The sleep switch circuit technique significantly reduces the subthreshold leakage energy as compared to both standard low threshold voltage and dual threshold voltage domino logic circuits. A domino adder enters and leaves the low leakage sleep mode within a single clock cycle. The energy overhead of the circuit technique is low, justifying the activation of the sleep scheme by providing a net saving in total power consumption during short idle periods.

A summary of the themes and ideas presented in this book is provided in Chapter 12. It is emphasized that low power and reliability concerns will dominate at all levels of the design hierarchy as the end of the traditional speed-centric methodology approaches. Some of the opportunities that exist for low power and reliable IC and system design are revisited.

