Design of Analog CMOS IC,
by B. Razavi

Chapter 1

Introduction to Analog Design

1.1 Why Analog?

It was in the early 1980s that many experts predicted the demise of analog circuits. Digital signal processing algorithms were becoming increasingly more powerful while advances in integrated-circuit (IC) technology provided compact, efficient implementation of these algorithms in silicon. Many functions that had traditionally been realized in analog form were now easily performed in the digital domain, suggesting that, with enough capability in IC fabrication, all processing of signals would eventually occur digitally. The future looked quite bleak to analog designers and they were seeking other jobs.

But, why are analog designers in such great demand today? After all, digital signal processing and IC technologies have advanced tremendously since the early 1980s, making it possible to realize processors containing millions of transistors and performing billions of operations per second. Why did this progress not confirm the earlier predictions?

While many types of signal processing have indeed moved to the digital domain, analog circuits have proved *fundamentally* necessary in many of today's complex, high-performance systems. Let us consider a few applications where it is very difficult or even impossible to replace analog functions with their digital counterparts regardless of advances in technology.

Processing of Natural Signals Naturally occurring signals are analog—at least at a macroscopic level. A high-quality microphone picking up the sound of an orchestra generates a voltage whose amplitude may vary from a few microvolts to hundreds of millivolts. The photocells in a video camera produce a current that is as low as a few electrons per microsecond. A seismographic sensor has an output voltage ranging from a few microvolts for very small vibrations of the earth to hundreds of millivolts for heavy earthquakes. Since all of these signals must eventually undergo extensive processing in the digital domain, we observe that each of these systems consists of an analog-to-digital converter (ADC) and a digital signal processor (DSP) [Fig. 1.1(a)]. The design of ADCs for high speed, high precision, and low power dissipation is one of many difficult challenges in analog design.

In practice, the electrical version of natural signals may be prohibitively small for direct digitization by the ADC. The signals are also often accompanied by unwanted, out-of-band

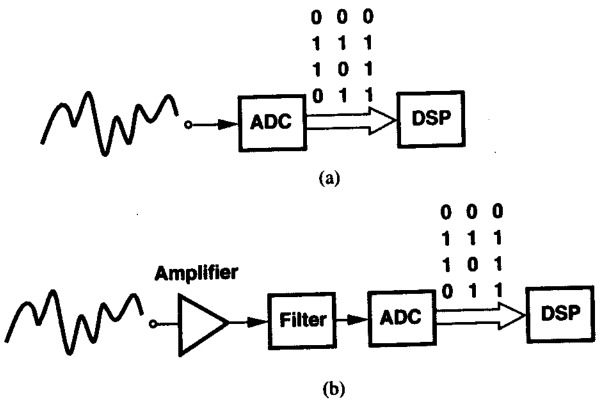

Figure 1.1 (a) Digitization of a natural signal, (b) addition of amplification and filtering for higher sensitivity.

interferers. The front end of Fig. 1.1(a) may therefore be modified as shown in Fig. 1.1(b), where an amplifier boosts the signal level and an analog filter suppresses the out-of-band components. The design of high-performance amplifiers and filters is also a topic of active research today.

Digital Communications Binary data generated by various systems must often be transmitted over long distances. For example, computer networks in large office buildings may transmit the data over cables that are hundreds of meters long.

What happens if a high-speed stream of binary data travels through a long cable? As illustrated in Fig. 1.2, the signal experiences both attenuation and "distortion," no longer resembling a digital waveform. Thus, a receiver similar to that of Fig. 1.1(b) may be necessary here.

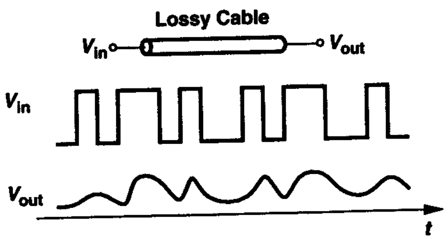

Figure 1.2 Attentuation and distortion of data through a lossy cable.

In order to improve the quality of communication, the above system may incorporate "multi-level"—rather than binary—signals. For example, if, as shown in Fig. 1.3, every two consecutive bits in the sequence are grouped and converted to one of four levels, then

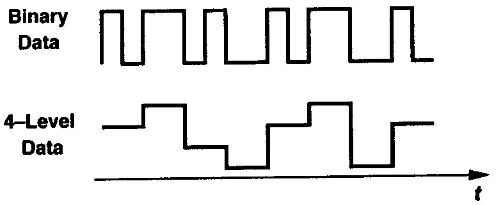

Figure 1.3 Use of multi-level signalling to reduce the required bandwidth.

each level is twice as long as a bit period, demanding only *half* the bandwidth required for transmission of the binary stream. Utilized extensively in today's communication systems, multi-level signals necessitate a digital-to-analog converter (DAC) in the transmitter to produce multiple levels from the grouped binary data and an ADC in the receiver to determine which level has been transmitted. The key point here is that increasing the number of levels relaxes the bandwidth requirements while demanding a higher precision in the DAC and the ADC.

Disk Drive Electronics The data stored magnetically on a computer hard disk is in binary form. However, when the data is read by a magnetic head and converted to an electrical signal, the result appears as shown in Fig. 1.4. The amplitude is only a few millivolts, the noise content is quite high, and the bits experience substantial distortion.

Digital
Storage

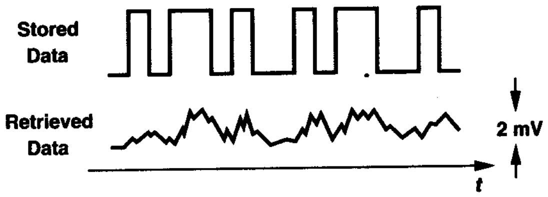

Figure 1.4 Data stored in and retrieved from a hard disk.

Thus, as illustrated in Fig. 1.1, the signal is amplified, filtered, and digitized for further processing. Depending on the overall system architecture, the analog filter in this case may in fact serve to remove a significant portion of the noise and the distortion of the signal. The design of each of these building blocks poses great challenges as the speed of computers and their storage media continues to increase every year. For example, today's disk drives require a speed of 500 Mb/s.

Wireless Receivers The signal picked up by the antenna of a radio-frequency (RF) receiver, e.g., a pager or a cellular telephone, exhibits an amplitude of only a few microvolts and a center frequency of 1 GHz or higher. Furthermore, the signal is accompanied by large

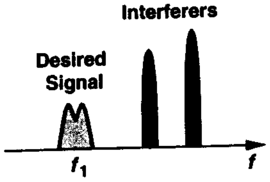

Figure 1.5 Signal and interferers received by the antenna of a wireless receiver.

interferers (Fig. 1.5). The receiver must therefore amplify the low-level signal with minimal noise, operate at a high frequency, and withstand large unwanted components. Note that these requirements are necessary even if the desired signal is not in "analog" form. The trade-offs between noise, frequency of operation, tolerance of interferers, power dissipation, and *cost* constitute the principal challenge in today's wireless industry.

Optical Receivers For transmission of high-speed data over very long distances, cables generally prove inadequate because of their limited bandwidth and considerable attenuation. Thus, as illustrated in Fig. 1.6, the data is converted to light by means of a laser diode and transmitted over an optical fiber, which exhibits an extremely wide band and a very low loss. At the receive end, the light is converted to a small electrical current by a photodiode. The receiver must then process a low-level signal at a very high speed, requiring low-noise, broadband circuit design. For example, state-of-the-art optical receivers operate in the range of 10 to 40 Gb/s.

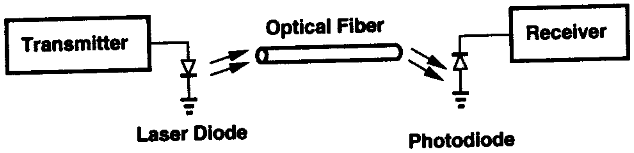

Figure 1.6 Optical fiber system.

Sensors Mechanical, electrical, and optical sensors play a critical role in our lives. For example, video cameras incorporate an array of photodiodes to convert an image to current and ultrasound systems use an acoustic sensor to generate a voltage proportional to the amplitude of the ultrasound waveform. Amplification, filtering, and A/D conversion are essential functions in these applications.

An interesting example of sensors is the accelerometers employed in automobiles to activate air bags. When the vehicle hits an obstacle, the drop in the speed is measured as acceleration and, if exceeding a certain threshold, it triggers the air bag release mechanism. Modern accelerometers are based on a variable capacitor consisting of a fixed plate and a deflectable plate [Fig. 1.7(a)]. The deflection and hence the value of the capacitor are proportional to the acceleration, requiring a circuit that accurately measures the change in capacitance. The design of such interface circuits is quite difficult because for typical accelerations, the interplate capacitance may change by less than 1%, demanding a high precision in the measurement. In practice, the structure of Fig. 1.7(b) is used to provide two capacitors that change in opposite directions, reducing the task to the measurement of the *difference* between two capacitances rather than the absolute value of one.

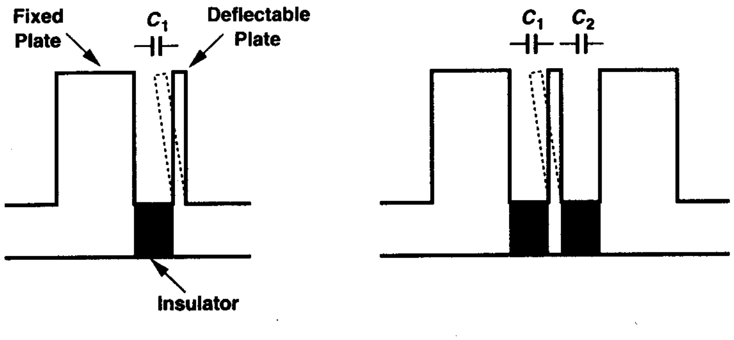

(a) (b)

Figure 1.7 (a) Simple accelerometer, (b) differential accelerometer.

Microprocessors and Memories Today's microprocessors and memories draw upon a great deal of analog design expertise. Many issues related to the distribution and timing of data and clocks across a large chip or among chips mandate that high-speed signals be viewed as analog waveforms. Furthermore, nonidealities in signal and power interconnects on the chip as well as package parasitics require a solid understanding of analog design. In addition, semiconductor memories employ high-speed "sense amplifiers" extensively, necessitating many analog techniques. For these reasons, it is often said "high-speed digital design is in fact analog design."

The foregoing applications demonstrate the wide and inevitable spread of analog circuits in modern industry. But, why is analog design difficult? We make the following observations. (1) Whereas digital circuits entail primarily one trade-off between speed and power dissipation, analog design must deal with a multi-dimensional trade-off consisting of speed, power dissipation, gain, precision, supply voltage, etc. (2) With the speed and precision required in processing analog signals, analog circuits are much more sensitive to noise, crosstalk, and other interferers than are digital circuits. (3) Second-order effects in devices influence the performance of analog circuits much more heavily than that of digital circuits. (4) The design of high-performance analog circuits can rarely be automated, usually requiring that every device be "hand-crafted." By contrast, many digital circuits are automatically synthesized and laid out. (5) Despite tremendous progress, modeling and simulation of many effects in analog circuits continue to pose difficulties, forcing the designers to draw upon experience and intuition when analyzing the results of a simulation. (6) An important thrust in today's semiconductor industry is to design analog circuits in mainstream IC technologies used to fabricate digital products. Developed and characterized for digital applications, such technologies do not easily lend themselves to analog design, requiring novel circuits and architectures to achieve a high performance.

ny Integrated?

The idea of placing multiple electronic devices on the same substrate was conceived in the late 1950s. In 40 years, the technology has evolved from producing simple chips containing a handful of components to fabricating memories accommodating more than one billion transistors as well as microprocessors comprising more than 10 million devices. As Gordon Moore (one of the founders of Intel) predicted in the early 1970s, the number of transistors per chip has continued to double approximately every one and a half years. At the same time, the minimum dimension of transistors has dropped from about 25 μm in 1960 to about 0.18 μm in the year 2000, resulting in a tremendous improvement in the speed of integrated circuits.

Driven by primarily the memory and microprocessor market, integrated-circuit technologies have also embraced analog design extensively, affording a complexity, speed, and precision that would be impossible to achieve using discrete implementations. Analog and mixed analog/digital integrated circuits containing tens of thousands of devices now routinely appear in consumer products. We can no longer build a discrete prototype to predict the behavior and performance of modern analog circuits.

hy CMOS?

The idea of metal-oxide-silicon field-effect transistors (MOSFETs) was patented by J. E. Lilienfeld in the early 1930s—well before the invention of the bipolar transistor. Owing to fabrication limitations, however, MOS technologies became practical much later, in the early 1960s, with the first several generations producing only n-type transistors. It was in the mid-1960s that complementary MOS (CMOS) devices (i.e., both n-type and p-type transistors) were introduced, initiating a revolution in the semiconductor industry.

CMOS technologies rapidly captured the digital market: CMOS gates dissipated power only during switching and required very few devices, two attributes in sharp contrast to their bipolar or GaAs counterparts. It was also soon discovered that the dimensions of MOS devices could be scaled down more easily than those of other types of transistors. Furthermore, CMOS circuits proved to have a lower fabrication cost.

The next obvious step was to apply CMOS technology to analog design. The low cost of fabrication and the possibility of placing both analog and digital circuits on the same chip so as to improve the overall performance and/or reduce the cost of packaging made CMOS technology attractive. However, MOSFETs were quite slower and noisier than bipolar transistors, finding limited application.

How did CMOS technology come to dominate the analog market as well? The principal force was device scaling because it continued to improve the speed of MOSFETs. The intrinsic speed of MOS transistors has increased by more than three orders of magnitude in the past 30 years, becoming comparable with that of bipolar devices even though the latter

have also been scaled (but not as fast). Multi-gigahertz analog CMOS circuits are now in production.

1.4 Why This Book?

The design of analog circuits itself has evolved together with the technology and the performance requirements. As the device dimensions shrink, the supply voltage of integrated circuits drops, and analog and digital circuits are fabricated on one chip, many design issues arise that were unimportant only a decade ago. Such trends demand that the analysis and design of circuits be accompanied by an in-depth understanding of their advantages and disadvantages with respect to new technology-imposed limitations.

Good analog design requires intuition, rigor, and creativity. As analog designers, we must wear our engineer's hat for a quick and intuitive understanding of a large circuit, our mathematician's hat for quantifying subtle, yet important effects in a circuit, and our artist's hat for inventing new circuit topologies.

This book describes modern analog design from both intuitive and rigorous angles. It also fosters the reader's creativity by carefully guiding him/her through the evolution of each circuit and presenting the thought process that occurs during the development of new circuit techniques.

1.5 General Concepts

1.5.1 Levels of Abstraction

Analysis and design of integrated circuits often require thinking at various levels of abstraction. Depending on the effect or quantity of interest, we may study a complex circuit at device physics level, transistor level, architecture level, or system level. In other words, we may consider the behavior of individual devices in terms of their internal electric fields and charge transport [Fig. 1.8(a)], the interaction of a group of devices according to their electrical characteristics [Fig. 1.8(b)], the function of several building blocks operating as a unit [Fig. 1.8(c)], or the performance of the system in terms of that of its constituent subsystems [Fig. 1.8(d)]. Switching between levels of abstraction becomes necessary in both understanding the details of the operation and optimizing the overall performance. In fact, in today's IC industry, the interaction between all groups, from device physicists to system designers, is essential to achieving a high performance and a low cost. In this book, we begin with device physics and develop increasingly more complex circuit topologies.

1.5.2 Robust Analog Design

Many device and circuit parameters vary with the fabrication process, supply voltage, and ambient temperature. We denote these effects by PVT and design circuits such that their performance remains in an acceptable range for a specified range of PVT variations. For example, the supply voltage may vary from 2.7 V to 3.3 V and the temperature from 0° to 70°. Robust analog design in CMOS technology is a challenging task because device parameters vary significantly from wafer to wafer.

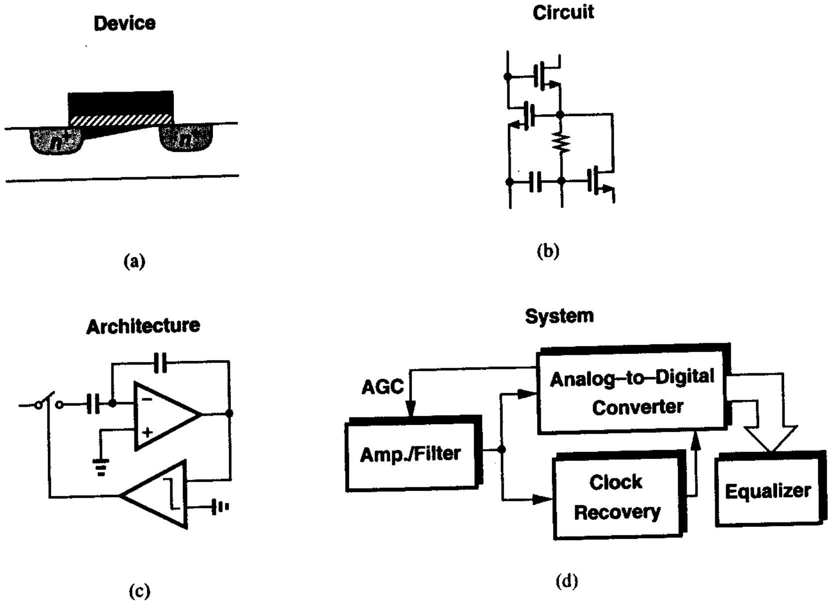

Figure 1.8 Abstraction levels in circuit design: (a) device level, (b) circuit level, (c) architecture level, (d) system level.

1.5.3 Notations

The voltages and currents in integrated circuits typically contain a bias component and a signal component. While it is desirable to employ a notation that distinguishes between these quantities, in practice other difficulties arise. For example, if the drain bias current of a transistor is denoted by I_D and the drain signal current by i_D, then the Laplace transform of i_D, $I_D(s)$, may be confused with I_D unless it is always accompanied by s. Furthermore, it is confusing to write the low-frequency gain of a circuit as $v_{out}/v_{in} = -g_m R_D$ and the high-frequency gain as $V_{out}/V_{in} = -g_m R_D/(1 + R_D C_L s)$.

In this book, we denote most voltages and currents by uppercase letters, making it clear from the context which component they represent. For example, I_D, V_{GS}, and V_X denote bias, signal, or bias+signal quantities. For input and output voltages, we use V_{in} and V_{out}, respectively.

Chapter 2

Basic MOS Device Physics

In studying the design of integrated circuits, one of two extreme approaches can be taken: (1) begin with quantum mechanics and understand solid-state physics, semiconductor device physics, device modeling, and finally the design of circuits; (2) treat each semiconductor device as a black box whose behavior is described in terms of its terminal voltages and currents and design circuits with little attention to the internal operation of the device. Experience shows that neither approach is optimum. In the first case, the reader cannot see the relevance of all of the physics to designing circuits, and in the second, he/she is constantly mystified by the contents of the black box.

In today's IC industry, a solid understanding of semiconductor devices is essential, more so in analog design than in digital design because in the former, transistors are not considered as simple switches and many of their second-order effects directly impact the performance. Furthermore, as each new generation of IC technologies scales the devices, these effects become more significant. Since the designer must often decide which effects can be neglected in a given circuit, insight into device operation proves invaluable.

In this chapter, we study the physics of MOSFETs at an elementary level, covering the bare minimum that is necessary for basic analog design. The ultimate goal is still to develop a circuit model for each device by formulating its operation, but this is accomplished with a good understanding of the underlying principles. After studying many analog circuits in Chapters 3 through 13 and gaining motivation for a deeper understanding of devices, we return to the subject in Chapter 16 and deal with other aspects of MOS operation.

We begin our study with the structure of MOS transistors and derive their I/V characteristics. Next, we describe second-order effects such as body effect, channel-length modulation, and subthreshold conduction. We then identify the parasitic capacitances of MOSFETs, derive a small-signal model, and present a simple SPICE model. We assume that the reader is familiar with such basic concepts as doping, mobility, and *pn* junctions.