


Chapter 1

Introduction



Layout design—or *physical design*, as it is also known in the industry and referred to throughout the book—is the final step in the design process for an electronic circuit. It aims to produce all information necessary for the fabrication process to follow. In order to achieve this, all components of the logical design, such as cells and their connections, must be generated in a geometric format (typically, as collections of rectangles), which is used to create the microscopic devices and connections during fabrication.

This chapter gives a sound introduction to the technologies, tasks and methodologies used to design the layout of an electronic circuit. With this basic design knowledge as a foundation, the subsequent chapters then delve deeper into specific constraints and aspects of physical design, such as semiconductor technologies (Chap. 2), interfaces, design rules and libraries (Chap. 3), design flows and models (Chap. 4), design steps (Chap. 5), analog design specifics (Chap. 6), and finally reliability measures (Chap. 7).

In Sect. 1.1, we introduce several of the most common fabrication technologies for electronic systems. The central topic of this book is the physical design of integrated circuits (aka *chips*, ICs) but hybrid technologies and printed circuit boards (PCBs) are also considered. In Sect. 1.2 of our introduction, we examine in more detail the significance and peculiarities of this related branch of modern electronics—also known as *microelectronics*. In Sect. 1.3, we then consider the physical design of both integrated circuits and printed circuit boards with a specific emphasis on their primary design steps. After these opening sections, we close the introductory chapter in Sect. 1.4 by presenting our motivation for this book and describing the organization of the chapters that follow.

1.1 Electronics Technologies

All electronic circuits comprise *electronic devices* (transistors, resistors, capacitors, etc.) and the metallic interconnects that electrically connect them. There is, however, a wide range of different fabrication technologies available to physically realize such electronic devices; these technologies can be classified into three main groups:

- *Printed circuit board technology*, which can be subdivided in
 - Through-hole technology (THT),
 - Surface-mount technology (SMT),
- *Hybrid technology*, often subdivided in
 - Thick-film technology,
 - Thin-film technology,
- *Semiconductor technology*, subdivided in
 - Discrete semiconductor components,
 - Integrated circuits.

To each of these technologies are added myriad extra features and custom designs for different use cases. Take, for example, automobile electronics, where a very high degree of robustness is required, or cellphones, where extreme compactness is a key requirement. We next examine the most important of these technologies in further detail.

1.1.1 Printed Circuit Board Technology

The *printed circuit board (PCB)* is the most widely used technology for electronic packaging. It mechanically supports and electrically connects electronic devices that are typically soldered onto the PCB.

Circuit Board

The basic element is an electrically isolated circuit board, known as the substrate core, and typically made of glass-fiber-reinforced epoxy resin. Everybody has seen this green board at some point. Papers stabilized with phenolic resins are also an option. (This approach was particularly widespread in the early years of electronics.) Paper-based circuit boards are only suitable for very low-spec applications, are rarely used anymore, and are not covered further in this book.

The circuit board has two primary functions: (i) to provide the foundation upon which the electronic devices are physically mounted, and (ii) to provide a surface upon which the interconnects for electrically connecting the devices can be constructed.

The interconnects are etched out of a metallic layer that has been applied to the substrate surface. The metal layer is made of copper and can be applied to one side

of the circuit board, or to both sides. Copper is the material of choice here, as it has several very beneficial properties: (i) it is an excellent electrical conductor; (ii) it lends itself well to etching; and (iii) it can be soldered easily. (Devices are soldered in place on the board, and simultaneously the chip pins are electrically connected to the interconnects on the board with solder.)

Fabricating Interconnects

How interconnects are fabricated is visualized in Fig. 1.1 and explained below with reference to steps (a) to (i) in the figure. The substrate core that is the foundation of the board is sometimes also referred to as the *carrier substrate*, as it “carries” (i.e., “holds” in place) the electronic devices and interconnects.

A substrate core is first coated with copper, and then a layer of photoresist is applied (a to c). The photoresist has a special property, namely that after exposure to light, it can be dissolved with a fluid called the *developer*. In the next step we use a *mask*, which is a (transparent) glass plate onto which the image of the desired interconnect has been applied as an opaque layer to the bottom surface of the glass plate (shown in black in Fig. 1.1d).

This mask is then positioned over the PCB (d) and exposed with light (illuminated sectors in yellow and shaded sectors in gray, Fig. 1.1e). The shaded area produces an image of the desired interconnect on the PCB. The resist at the exposed areas thus becomes dissolvable (the areas in pale blue in Fig. 1.1f); this exposed resist can then be dissolved and washed away with the developer. The remaining unexposed




Fig. 1.1 Sectional view of the creation of interconnects on a printed circuit board (PCB) by photolithography and subsequent etching

areas retain their photoresist, which protects the copper layer underneath against etching in the next step (g), such that the etching agent only removes the copper on the unprotected areas. We say the resist “masks” the etching. After etching, only the copper remains at the prior unexposed areas; the remaining photoresist is washed away with a suitable fluid (h). Through this process, we have created in the copper layer the interconnect structure that was patterned on the mask (i).

When only a few PCBs are required, such as for prototypes, the interconnects are sometimes not formed by etching, but by mechanically milling the metal layer.

Multi-Layer Printed Circuit Boards

Boards can be constructed of several stacked substrate cores, and in this case are called *multilayer PCBs*. Figure 1.2 shows an example of a multilayer board, consisting of three cores and six routing layers (the top and bottom of each of the three cores). The cores are glued together with a bonding agent, which also acts as an electrical isolator between the opposing copper layers of neighboring substrate cores, to prevent short circuits.

Plated-through contacts, known as *vias*, are then used to electrically connect different routing layers. Vias are created at the beginning of the manufacturing process by drilling holes through the core layers. The via walls are later coated with copper to make them electrically conductive. These vias are labeled as *buried*, *blind* and *through-hole vias* depending on where the holes are located (see Fig. 1.2).

Mounting Technologies

Two different technologies are primarily used for mounting components on PCBs:

- Through-hole technology (THT), and
- Surface-mount technology (SMT).

THT makes use of devices that have leads to make the electrical contacts. These leads are inserted in holes, which are through-hole vias, and soldered in place (see Fig. 1.2, left). With SMT, the devices instead have metal pads for connecting to the surface of the board (shown in black in Fig. 1.2). Associated with these mounting technologies—THT and SMT—are respectively *through-hole devices (THDs)* and *surface-mount(ed) devices (SMDs)*.




Fig. 1.2 Cross-section of a multilayer board with six routing layers

The two mounting technologies can also be mixed. Component placement systems can handle SMDs much more easily than THDs. In addition, much higher packing densities can be achieved with SMDs, as they are smaller and can be mounted on both sides of a PCB. These advantages make surface-mount technology (SMT) the more widely used approach today.

Besides discrete devices, integrated circuits (ICs) can also be mounted on PCBs. In general, however, they must be “packaged” in an enclosure. Unpackaged ICs (aka *bare dies*) are sometimes mounted directly on PCBs; in this case, however, the stability of the connection is critical due to the different thermal expansion rates of semiconductors and boards.

1.1.2 Hybrid Technology

Hybrid technology is an approach in which some of the electrical components are physically separate devices that are then mounted on the carrier substrate, while other components are created directly on the carrier substrate during fabrication. Thus, we get the name “hybrid technology”.

A variety of carrier substrates are used in hybrid technology to construct the substrate core; ceramic, glass and quartz are commonly used. SMDs can be mounted on these carrier substrates; however, THDs cannot be mounted, as through holes for mounting purposes are not drilled in these substrates.

Interconnect traces are produced differently on hybrid-technology substrates than on PCBs. The two technologies typically used are *thick-film technology* and *thin-film technology*. In thick-film technology, conductive pastes are applied in a silk-screen printing process and then burnt in. In thin-film technology, the conductive material is vaporized or sputtered¹ onto the substrate. Interconnects are subsequently formed using a photolithographic process much like the process described above for PCBs.

The electrical conductivity of the deposited layers can be adjusted over a large range, which allows electrical resistors to be produced along with interconnects using these technologies. Resistivity values can be finely tuned with a laser; resistance can be increased by trimming the outer regions of the deposited interconnect back perpendicular to the current flow until the desired resistance value is reached.

Interconnect crossings and capacitors can also be realized, such as by alternately stacking conductive and insulating layers. Capacitors can also be created by intertwining comb-like interconnects within a metal layer. An example of this printed capacitor is given in Fig. 1.3.

Example: LTCC Technology

A widely used variant of the thick-film hybrids is *LTCC* technology, which stands for *low temperature co-fired ceramics*. The LTCC fabrication process is representative of

¹ Sputtering is a physical process whereby microscopic particles of a solid material are ejected from its surface by bombarding the solid with high-energy ions.




Fig. 1.3 Fabrication of an LTCC hybrid circuit with printed capacitor, inductor and resistor (LTCC: low temperature co-fired ceramics)

many other technology variants and we will discuss it below. The LTCC fabrication process is depicted in Fig. 1.3a–h.

LTCC technology does not use a prepared ceramic substrate. Instead, fabrication starts with films, in which the ceramic mass in powder form is bonded with other substances. These films are called *green sheets* (a). As we will see below, these films will be made rigid during subsequent processing steps, to produce a board that will become a component in the larger system. First, holes for vias are punched in the sheets (b). These are then filled with conductive paste (c). The interconnect geometries are formed with conductive paste in a silk-screen printing process (d). The green sheets are then stacked on top of one another and laminated by heating them up slightly. They are thus bonded together (e). The stack is then cut to size, pressed together and fired and sintered in an oven (f). Some of the additives escape from the sheet stack during this process and the hybrid shrinks and is sintered to

a ceramic plate, which can contain several interconnect layers in its core. Pressure continues to be applied to the laminated stack during burnout and sintering, so that the shrinking effect is almost completely in the z-axis and the lateral dimensions are not affected. Resistors and conductive surfaces for contacting SMDs and ICs are then printed onto the surface and burnt in (g).

Finally, the SMDs and ICs are mounted (h). The SMDs are fixed in place with a conductive glue or by reflow soldering. The ICs can be mounted unpackaged as bare dies because of the similar thermal expansion rates of carrier substrate (ceramic) and the dies' silicon. They are electrically connected using bond wires running from contact surfaces on the IC, so-called *pads*, to contact surfaces on the carrier.

Among the benefits of hybrid technology over printed circuit technology are: (i) greater mechanical stability (when exposed to extreme shocks and vibrations in automobiles, for example); (ii) higher packing density (due to the use of bare dies); and (iii) better dissipation of thermal losses.

Thermal dissipation is mainly achieved with LTCCs by mounting the hybrid on a heat sink to maximize the thermal coupling over the entire hybrid surface. Good heat dissipation from the top of the hybrid to its bottom can be further improved by deploying so-called *thermal vias*—these are plated-through contacts designed specifically to transmit heat.

While hybrid technology has advantages as noted above (e.g., greater mechanical stability, packing density, and thermal dissipation), it has the disadvantage of higher manufacturing costs, as compared to PCBs.

1.1.3 Semiconductor Technology

With the technologies we have discussed thus far, the electronic devices must be wholly or partially added from an external source. With semiconductor technology, on the other hand, an entire electronic circuit can be built as a single unit. Here, all electronic devices and all electrical connections are created in the fabrication process itself. This type of circuit is fully integrated—this is where the name *integrated circuit* (*IC*) originates—in a monolithic semiconductor die. These single, small flat pieces, comprised mostly of silicon, are also called “chips”.

We can also use semiconductor technology to construct purely discrete (i.e., single) electronic devices. Typical examples are diodes or active devices, such as transistors and thyristors, for driving very large currents in power electronics. If we examine these devices more closely, we find that they are composed of many similar devices connected in parallel on the chip. Protective circuitry that remains hidden

from view, but which supports the device's performance characteristics, is also often integrated in the chip.

What are Semiconductors?—Physical Aspects of Semiconductor Materials

While semiconducting materials can conduct electrical current, their electrical resistivity at room temperature is quite high. However, their conductivity increases exponentially with rising temperature. This thermal characteristic is very different from standard interconnects (metals) and is a key property of semiconductors. We will now delve deeper into the underlying physics.

Free moving charge carriers are required for a current flow. In solids, these charge carriers are electrons. Hence, the question is: “How do we get enough ‘free’ electrons in semiconductors?” Electrons orbit the atomic nucleus, and their energy levels increase the further they are away from the nucleus. It is also a fact that they can only have certain energy states called *shells* (electron shells) that expand into so-called *bands* in constellations of many atoms. The most exterior band in a substance is also called the *valence band*. If electrons in the valence band (so-called *valence electrons*) can be sufficiently energized to enable them to jump to the next higher band, they can move freely there and thus increase the electrical conductivity of the material. Therefore, this band is also labeled as *conduction band*.

Valence and conduction bands are particularly close together in conductive materials, such as metals; they can even overlap (see the orange area in Fig. 1.4). In this case, many electrons can jump to the conduction band. Hence, metals are excellent conductors. In the case of insulators, on the other hand, the energy gap ΔE between the valence and the conduction band (the so-called *band gap* or *band distance*) is so large that it is an almost insurmountable threshold, and there are practically no electrons in the conduction band (blue area in Fig. 1.4).

The main characteristic of semiconductors is that their band gap lies between these two extremes (central green area in Fig. 1.4). On the one hand, this gap is so large that only very few electrons in the valence band have enough energy at




Fig. 1.4 Materials in the “conductor”, “semiconductor” and “insulator” categories as a function of the band gap between valence and conduction band. The values are given for typical semiconductor materials at 300 K. (SiC can take values between 2.4 eV and 3.3 eV, depending on the formed crystal lattice. The value for crystal lattice “6H” is shown.)

room temperature to reach the conduction band. On the other hand, the conduction band is close enough that a temperature rise on the order of a few hundred degrees kelvin above room temperature provides enough energy to increase the number of free electrons, and thus the conductivity, by many orders of magnitude.

The conductivity increases not only because of the free electrons in the conduction band, but also due to the creation of electron holes—known as *defect electrons* or *holes*—in the valence band. These holes can be easily filled with valence electrons from neighboring atoms; the holes then disappear, generating new ones in the “electron-delivering” atoms. A current flow—known as *hole conduction*—can then take place as a result of this chain movement of valence electrons. However, generating free-charge carriers (electrons and holes) by thermal means is not our main concern here, rather we want to highlight the underlying physics for a better understanding of our primary intention.

Let us now take silicon as a typical example (Fig. 1.5). For silicon, the band gap between the lower edge of the conduction band E_C and the top edge of the valence band E_V is given by $\Delta E = E_C - E_V = 1.1$ eV. Silicon is *4-valent*—in other words, the valence band of a silicon atom contains four electrons. If we replace a silicon atom with an atom of a 5-valent element, such as phosphor, arsenic or antimony, the extra electron will not “fit” into the valence band of the surrounding silicon crystal. It has an energy level E_D , which is just slightly below the silicon conduction band; it is so close to this band that at room temperature it has enough energy to allow it to enter this conduction band (Fig. 1.5, left). Hence, introducing 5-valence impurity atoms into the silicon allows us to increase its conductivity.

Instead of implanting 5-valence impurity atoms, we can also implant 3-valence impurity atoms, such as boron, indium or aluminum, to increase the conductivity. In this case, the impurity atom is at energy level E_A , which is just slightly above the valence band edge E_V of silicon; hence, this impurity atom can very easily accept a fourth electron from a neighboring silicon atom (Fig. 1.5, right). The number of holes in the silicon valence band is thus increased. These holes can be viewed as positive charge carriers that are available for current flow, as they can move freely.




Fig. 1.5 Generating free-charge carriers and thus current flow in a semiconductor (silicon). The creation of free charge carriers by doping with donors is shown on the left and with acceptors on the right, resulting in conduction by electrons (left) and holes (right)

Implanting impurities into a substance is called *doping*. Five-valence impurity atoms are called *donors*, as they release an electron (into the conduction band). They are listed in the column to the right of silicon in the periodic table. Three-valence impurity atoms are known as *acceptors* because of their ability to accept valence electrons from neighboring atoms. They are found in the column to the left of silicon in the periodic table.

A semiconductor containing donors is called *n-doped* and one containing acceptors is called *p-doped*. In areas of n- and p-doping, the additional electrons are trapped by the additional holes. They are said to be *recombined*. Donors and acceptors effectively neutralize each other here.

Any residual donors or acceptors are key to conductivity. The semiconductor is said to be *n-conductive* when donors are in surplus, as the current flow is mostly due to electrons, i.e., negative charge carriers. The semiconductor is said to be *p-conductive* when there are surplus acceptors, as here it is mainly the holes—that can be viewed as positive charge carriers—that cause the current flow. The more abundant charge carriers are called *majority carriers*, while the correspondingly less abundant charge carriers are called *minority carriers*.

The Use of Semiconductors

Extremely pure semiconductor material in monocrystalline form is required for producing integrated circuits. All atoms must be physically arranged in a continuous regular structure. This type of structure must be manufactured as it does not occur naturally. It is produced in the form of crystal ingots that are then machined into very thin slices, or *wafers*, which are used for chip fabrication. One wafer can hold a huge number of chips—many hundreds to tens of thousands, depending on the chip and wafer sizes; the chips are all fabricated together on the wafer. At the end of the fabrication process, individual rectangular chips, or *dies*, are cut from the wafer.

Figure 1.6 shows a finished wafer under a microscope. The dies have been separated from one another and are held in place by an adhesive foil (so-called “blue tape”).

The most widely used material in the semiconductor industry is silicon. Other semiconducting materials are utilized for special purposes in the industry. Typical examples include: gallium arsenide (GaAs) and silicon germanium (SiGe) in RF circuits; and gallium nitride (GaN) and silicon carbide (SiC) in power electronics. Two 4-valent elements are combined both in SiC and SiGe; a 3-valent element is combined with a 5-valent element in GaAs and GaN. The resulting crystalline structure again behaves like a 4-valent element.

Integrated Devices and Interconnect Traces

Integrated devices are produced by successively doping a wafer in different ways. The doping operations can be altered in the following ways: (i) different impurities (there are generally several different types of donors and acceptors); (ii) different concentrations (the number of impurity atoms per unit volume); (iii) the penetration depth (up to some μm), and (iv) the doping location.




Fig. 1.6 Wafer under microscope; one chip is contacted by two probes for testing

Integrated devices can be produced using simple processes having less than 10 doping operations. In more complex processes there can be more than 20 doping procedures. For example, an additional layer of the base material is often deposited on the wafer surface between doping with impurity atoms. (This process is called *epitaxy*.)

The doping process is conceptually similar to the photolithographic process for PCBs presented in Sect. 1.1.1, in that masks are again used to selectively alter the composition of the surface of an area. For PCBs we masked areas to create interconnect traces (using subsequent etching, etc.). Here we also use masks, but ones that are far more detailed, with feature sizes that are on the order of nanometers, to selectively deposit atoms on the surface of the wafer, effectively implementing different doped areas.

Different types of electronic devices, such as transistors, diodes and resistors, can be created by implementing differently doped areas. However, care needs to be taken here. When we talk about “devices” on chips, we should be aware that we are only talking about different sections of a monolithic semiconductor die. These sections are designed in such a way that the electrical interaction of the doped areas contained in each section produces the behavior of the desired electronic device. In contrast to the devices on a printed circuit board (PCB), the devices on a chip are never isolated from one another. This can always cause the devices on a chip to interact. These interactions must be considered in the design flow as will be discussed in detail in Chap. 7.




Fig. 1.7 Sectional view of an IC chip of a typical NPN transistor **a** at the beginning of the process, **b** after “front-end-of-line” (FEOL) and **c** after “back-end-of-line” (BEOL) in the semiconductor process. n-doped regions are drawn in blue, p-doped regions in red. Metallic layers are brown and insulating layers ocher

Figure 1.7 shows a sectional view of a chip. A wafer is slightly less than 1 mm thick. The electrically active parts, however, are located in a very thin layer on one of the two chip surfaces. This area of approximately 1 to 2% of the wafer thickness is depicted in Fig. 1.7; the figure also shows the fabrication stages, which are explained below.

First, the electronic devices are constructed on the wafer surface using doping operations. Semiconductor manufacture begins with a *raw wafer* (see Fig. 1.7a). All doping is carried out in this so-called “front-end-of-line” (FEOL) part of IC fabrication² and an epitaxial layer is applied as well, if required. The result is exemplified with an NPN bipolar transistor³ in Fig. 1.7b. The regions doped with impurities are shown in color. We always show n-conductive regions in blue and p-conductive regions in red in this book. The colors indicate that the raw wafer (a) was originally p-doped and the epitaxial layer (b, on top of the raw wafer) n-doped.

Second, the electrical devices are interconnected by constructing metallic and insulating layers. This process is commonly labeled as “back-end-of-line” (BEOL). Here too a photolithographic process involving masks is used to create interconnects. This process is also conceptually similar to the way interconnects were created on PCBs (as presented in Sect. 1.1.1), but here the masks are again far more detailed with nanometer-sized features. Insulating layers (shown in ocher) and metallic layers (brown) are alternately stacked on top of one another and structured in the back-end-of-line (BEOL) part of fabrication. *Interconnects* and *through contacts* are formed in this step.

²While the term “front-end-of-line” (FEOL) refers to the first portion of any IC fabrication where the individual devices are patterned, “back-end-of-line” (BEOL) comprises the subsequent deposition of metal interconnect layers. Both are discussed in Chap. 2.

³Bipolar transistors are devices whose operation depends on the two types of charge carriers (electrons and holes). We will cover these devices and their operation more fully in Chap. 6.

The result of the BEOL is shown in Fig. 1.7c for two interconnect layers. The devices are electrically connected on the silicon surface through *contact holes* (*contacts*) in the bottom insulating layer. The through contacts between neighboring metal layers are known as *vias* in IC chips as well as in PCBs. (Please note that we use the term *through contact* to label any vertical connection between layers. Referring to ICs, we further differentiate between *contacts* that connect the devices to the (first) metal layer and *vias* that connect (two) metal layers.)

As illustrated in Fig. 1.7, photolithography is used to create all structures on the chip, regardless of whether it is delineating doping areas in the FEOL (front-end-of-line), or the formation of vias and interconnects in the BEOL (back-end-of-line). Photolithography plays a key role in fabrication, as we also saw in the production of PCBs in Sect. 1.1.1. We shall cover the process steps in semiconductor fabrication in detail in Chap. 2.

1.2 Integrated Circuits

1.2.1 Importance and Characteristics

Since the first integrated circuits (ICs) appeared in the 1960s, microelectronics has developed at a breathtaking pace. It has long become a key technology across all our technological advances. It has a huge effect on all our lives and will continue to do so. What are the drivers behind this massive relentless creative power? We will now dig deeper into the important properties of ICs that continue to propel these developments, and try to answer this question.

The idea of integrating electronic circuits on a single piece of semiconductor material was first expressed towards the end of the 1950s by Jack Kilby [3] and Robert Noyce [5] independently of one another. The first commercial semiconductor chip was produced in 1961: it was a logical memory element (called a *flipflop*) with four transistors and five resistors [1].

This was the birth of microelectronics and the beginning of the modern computing era. From that time forth, semiconductor technology went from strength to strength, accompanied by unabating IC miniaturization. This miniaturization is the driver for a series of effects, that are mutually supportive, and whose cumulate effect, upon closer scrutiny, continues to amaze.

The on-going reduction in the footprint of individual devices means the chips use less power, operate faster, and can accommodate an increasing number of functions as the individual devices can be packed more densely. These are all logical developments that are easy to understand. Why the extra functions should be cheaper is not so easy to fathom though. We can explain it as follows. It is a fact that semiconductor technologies and chip space become more expensive with increasing miniaturization. Nonetheless, the extra costs are more than recouped as individual functions require

less chip space due to the miniaturization. You get more bang for your buck—effectively “more functionality for the same price”—with each new chip generation.

There is another effect that is less obvious but is nonetheless a very important aspect of the chip success story. Ever increasing integration density in chips has a very positive effect on the reliability of electronic systems: every incremental reduction in the number of solder points and every discrete device you can do without reduces the probability of a system failure. An entire chip is only a single device as far as the downtime risk goes. (Recall that integrated devices are essentially only sections of a semiconductor chip.) The semiconductor chip thus represents a single device, and a single point of failure; as a result, systems built using such semiconductor chips have fewer possible points of failure, which leads to higher reliability.

Let us try to imagine the implications of these effects: If you split the electronics from a modern mobile phone into discrete (i.e., individual) electronic components, and place them on PCBs, for example, you would need a huge industrial building to hold them. This monstrous “device” would not only be unwieldy and therefore useless, it would also be prohibitively expensive and highly unreliable. If a single device or connection in this hypothetical system failed, the system would fail as well. It would likely be permanently down, which would mean that (on the bright side) at least you wouldn’t need the entire power plant required to run it!

Let us list these six effects again: continued improvements in microelectronics make electronic systems smaller, faster, more economical, more intelligent, cheaper and more reliable.⁴ Normally, these performance characteristics cannot all be improved together—as evidenced in other technology sectors, such as the automotive industry, for example. They typically hamper each other, and engineers must find the best compromise for each use case. In contrast, all these performance characteristics have improved in unison in microelectronics, which explains its persistent and sustainable success.

1.2.2 Analog, Digital and Mixed-Signal Circuits

Today’s chips are highly complex. The first thing to be aware of is that most of them contain both digital and analog circuits. Not only do these two types of circuits operate fundamentally differently from one another, they differ greatly in their suitability for existing design flows and semiconductor technologies.

Let us look at digital circuits first. They are much more technically amenable than their analog counterparts because they handle exclusively discrete signal values. These typically are binary signals, with only two differentiable values, which are

⁴We should mention here that downscaling to lower technology nodes in semiconductor fabrication has reached a point where aging effects are becoming increasingly critical. One of the more acute concerns is interconnect degradation by migration effects, in which the electrical current flowing through the IC can slowly erode the miniature physical structures. Preventive measures for these effects are needed especially in the physical design flow. We shall deal with this topic fully in Chap. 7.

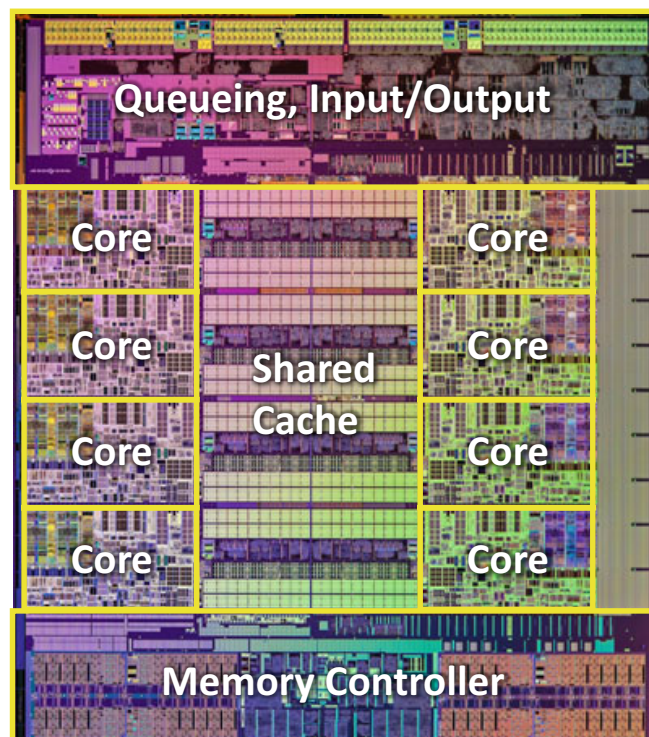
normally interpreted as the dual binary digits “1” and “0”, or the logical values “true” and “false”.

Digital logic can be implemented electrically with two (arbitrary) voltage levels. These given voltage levels must only be reached within a given tolerance. A range is defined between the logic states so that they can be clearly demarcated. Flawless operation can be achieved by waiting for a sufficiently long period to ensure that all logic states have settled at a defined value. (This is accomplished by setting the clock rate accordingly.) Hence, the requirements for these devices (typically unipolar CMOS⁵ transistors) to switch logic states can be low.

Modern ICs in digital electronics integrate numerous cores and the necessary peripherals directly on the chip itself. These chips can contain more than ten billion transistors today (2020). Figure 1.8 shows the Intel® “i7 Haswell-E™” from 2014 as an (historic) example. The chip, which was fabricated in a 22 nm node, has a surface area of 355 mm². It contains eight cores and consists of a total of 2.6 billion transistors [6]. These chips often look like images of cities taken from satellites orbiting the earth; it is amazing to realize that the level of complexity similar to that of a huge city (which would stretch over an entire continent!) has been precisely manufactured on a piece of silicon the size of your fingernail.

Aside from digital circuitry, electronic systems require analog circuit devices, too. They mediate between the abstract digital data processing and our real world—with its myriad physical parameters whose values and periodicities change *continuously*

Fig. 1.8 Intel microprocessor in 22 nm technology node with eight core processors



⁵CMOS is an acronym for “complementary metal oxide silicon”. CMOS technology incorporates two complementary n-type and p-type unipolar transistors, which are manufactured as metal-oxide-silicon layers. We will cover them in detail in Chap. 2.

in contrast to *discrete* digital states. This “task sharing” is analogous to biological organisms. Here, in addition to a brain for information processing, every organism needs (i) sensory organs to scan the environment, (ii) internal energy supplies, and (iii) appendages to physically act on the environment. Similarly, in any mechatronic or electronic system there is a need for analog circuitry, to support its inherent digital information processing. These systems (i) scan analog sensor inputs and convert them to digital data, (ii) supply the system with current and voltage, and (iii) act upon the computed data to control external displays, speakers, valves, electric motors, and the like. All these tasks are performed by circuits with one thing in common: they handle and produce analog signals.

It should be noted that the CMOS technology used in digital circuits can be used as well for many analog tasks. In addition, there are many applications where devices with special performance characteristics are required. Among these custom-designed devices are bipolar transistors—characterized by high cutoff voltages, robustness, usage of temperature dependence—, and special power transistors with very low resistance in the “on” state and the ability to conduct very large currents. These custom circuits were implemented using different, separate chips in the early semiconductor years. They were manufactured using semiconductor processes tailored for the respective devices.

The situation has changed in recent years. Semiconductor technologies, in which *all types* of devices needed for an integrated system can be manufactured on a single chip, have been available since the 1990s. Examples of these “mixed technologies” are *BICMOS* (bipolar transistors and CMOS) and *BCD* (bipolar transistors, CMOS and DMOS⁶).


Due to today’s high degree of integration, the combination of digital and analog parts of a circuit chip is standard practice. Now most chips are *mixed-signal* chips; they are fabricated in CMOS or BICMOS depending on the specification. If they include power transistors as well, then they are called *smart power* ICs which are fabricated in the above-mentioned BCD technology.

Figure 1.9 shows a smart power chip from 2018 for an automotive control unit from Robert Bosch GmbH®. All system functions are integrated in this single chip: analog circuits for sensor scanning (“Sense”); internal power supply (“Supply”); power stages for actuator control (“Act”); digital information processing (“Think”) implemented with *standard cells*.⁷ A chip containing all these different types of electronic modules is called an *SOC* (*system on chip*) [2]. The chip illustrated in Fig. 1.9 is manufactured in a BCD technology in the 130 nm node, with a surface area of 34 mm² and containing 164,000 devices in the analog circuit parts and some

⁶DMOS stands for “double diffused metal oxide silicon”. This is a fabrication technology for unipolar transistors that switch large currents in power electronics. We can implement exceptionally low on-state resistances on the order of mΩ with DMOS transistors.

⁷Designing with standard cells is a very efficient and hence popular design flow for ICs. We will introduce and discuss standard cells and design flows in Chap. 4.

Fig. 1.9 Bosch smart power chip in 130 nm BCD technology (BCD: bipolar transistors, CMOS and DMOS) for automotive electronics



3 million transistors in the digital part (yellow box). The external power supply is 14 V, and the chip breakdown voltage⁸ is 60 V.

1.2.3 Moore's Law and Design Gaps

We have seen how semiconductor downscaling, the relentless push to smaller chip structures, is central to the microelectronics evolution. So-called *technology nodes* (also referred to as *process nodes*) are categorized according to the smallest structure size that can be implemented reliably and reproducibly on a wafer. There is however no universally accepted definition of the scale in question here. Determining and specifying this feature size differs from manufacturer to manufacturer. For example, via holes, or the smallest permissible wire width, or the smallest permissible active length of a unipolar transistor (defined as the distance between the transistor's source and drain), are some of the smallest structures found on chips. Nevertheless, we can at least estimate the approximate size of these dimensions from the smallest feature size specification.

Fabrication processes for semiconductors are very complex, and every effort is made to avoid tinkering with established production processes. Reducing the feature size is a huge undertaking. Hence, miniaturization is not a continuous process; rather, it takes place in clearly defined steps. Experience has shown that transitioning to a smaller structure size is economically viable if the number of manufacturable devices

⁸As we will further explain in Chaps. 6 and 7, the breakdown voltage of an insulating material defines the maximum electric field that the material can withstand without breaking down, i.e., without conducting some amount of electricity.

per unit surface area can approximately be doubled. This means that the devices' surface areas must be halved without loss of functionality.

As described above, CMOS transistors used in digital circuits have the lowest functional requirements—which is also the reason why this type of transistor lends itself so well to miniaturization. Since the late 1970s (when CMOS technology matured), the surface areas of CMOS transistors have been successfully halved every two to three years, while their internal structures have remained unchanged.⁹

Hence, economical miniaturization steps have regularly been achieved with structural downscaling by a factor of $(1/\sqrt{2})$. This process is also called *shrinking*. Technology nodes have scaled approximately in this way since “1 μm processes” (that is, processes that allow a 1 μm structure size) first appeared. As mentioned earlier, these technological milestones are referred to as “process nodes” or “technology nodes”.

Figure 1.10 shows the technology node evolution for different semiconductor processes on a logarithmical scale since 1970. CMOS technology (the brown/red plot) has clearly been the main driver in these advances. The plot is based on when the first microcontroller chips became available in the respective feature sizes. The semiconductor processes for other applications (shown in blue) follow this “leading edge technology” at different time intervals. All curves plot the averaged long-term trend based on real data. These plots are not meant to give precise information on dates and feature sizes, but serve as a guide to the trends.

As we have seen, the miniaturization depicted in Fig. 1.10 has paved the way for integrating ever increasing numbers of devices on a single chip and, hence, more and more functions as well. This progression is also shown in Fig. 1.11, also beginning




Fig. 1.10 Plots of the smallest manufacturable feature sizes over time for different technology nodes. Marked points indicate typical “process nodes”

⁹This statement applies to structure sizes greater than approximately 20 nm. For smaller sizes, a different internal structure is required for unipolar transistors. FinFETs—which are beyond the scope of this book—are used in this case.




Fig. 1.11 Approximate growth rates for devices per IC (black, left scale) and design productivity in devices per person-year (red, right scale) for digital chips (above) and smart power ICs which contain both analog and digital parts (bottom). The digital design gap, the analog design gap and the gap between analog- and digital-design productivity (red vertical arrow) are also shown

in the 1970s. The chart shows the exponential increase in the number of devices integrated on a chip (black curves, left scale).¹⁰

The director of research and development at Fairchild Semiconductor Inc., Gordon Moore stated right at the beginning of the chip-miniaturization narrative that the number of devices on a chip doubled every year and he plotted his observation in a similar chart back in 1965 [4]. He predicted in his publication as well that this trend would continue into the foreseeable future. This exponential trend was then given the name “*Moore’s Law*” in the early 1970s—the first microcontrollers emerged at that time—when it became clear that the trend had by then become established. (In 1975, Moore revised the forecast to doubling every two years, a compound annual growth rate of 41.4%.)

In Sect. 1.2.1, we discussed the amazing effects of miniaturization that have caused Moore’s Law to still apply today. So far, we have discussed all this in the context of the end user and the chip fabricators. But there is one important aspect we haven’t touched on yet, and this brings us to the topic of this book: there is much more to be done before these amazing little chips can be manufactured and ultimately used—they need to be *designed* first!

Designing ICs, i.e., laying out their specific physical dimensions, is a huge challenge. The first digital chips used for switching logic gates were drawn in circuit diagrams and the photomasks for them were designed manually or with simple

¹⁰We use “devices/IC” and not, like most other authors, “transistors/IC” as a unit of scale, as many other types of devices besides transistors are used in mixed processes. The data for “transistors/IC” and “devices/IC” is almost the same for logic chips, though.

drafting software. This type of manual design was quickly superseded in the 1980s because it was too inefficient. Accompanying the exponential growth in complexity in microelectronics, significant work was done in academia and industry to provide IC designers with powerful software tools and innovative design techniques. This field is known as *electronic design automation*, or *EDA* for short.

The effectiveness of digital IC designers could be increased considerably with EDA. While the design process for integrated logic circuits is currently highly automated (i.e., many steps are routinely performed by software programs), the work involved in designing logic chips is nevertheless increasing all the time. This conundrum can be quantified and visualized by considering the number of devices on a chip and the total effort required for its development measured in person-years, and then calculating the quotients. This metric is called *design productivity*. It is plotted in red in Fig. 1.11 and refers to the scale on the right. Although the increase in design productivity is exponential, as well, the rate of increase falls far short of Moore's law. In other words, the mean IC complexity and the design productivity continuously drift apart. This phenomenon is known as *design gap*.

The design gap in digital logic design has been widely written about. It is visualized in Fig. 1.11 by the upper shaded area and the brown vertical arrow, and is one of the toughest and most urgent problems in microelectronics. A major effect of the design gap, along with the rise in IC development costs, is that the number of designers working in chip design must be constantly increased, as the design lead time cannot be extended due to market pressures. Upwards of 1,000 engineers—and they are often spread around the globe—are required today in a typical project team to launch a new computer chip on the market.

We want to address another similar problem now—the *analog design gap*. This gap started emerging around the turn of the century. It is now a critical issue that affects all chips with both digital and analog circuitry. These are mixed-signal and smart power chips that—as already mentioned—make up most of all current chip designs. These ICs also comply with Moore's Law, and although the growth in absolute device numbers lags that of digital (logic) chips, the growth rates are similar.

The increase in device numbers is primarily due to the growing digital circuitry in these mixed-signal designs. The situation for smart power ICs is shown in the lower solid black line in Fig. 1.11. More than 90% of the devices in a modern smart power IC are digital. The highly automated EDA processes tailored for digital design (indicated by the upper red line) are ideally suited for designing these digital sections.

In contrast, the situation with the design of analog circuit devices is very different: their numbers are growing too, but at a slower pace (dotted line). As described in Sect. 1.2.2, analog signals are signals of continuous amplitude and time that must be undistorted as much as possible. IC designers must consider the effects of multiple noise sources on their design: noise can disturb analog signals and cause malfunctions. Designers use analog design methods and take steps in the physical layout design to suppress noise as much as possible. As there are a multitude of different sources of interference, they must consider the physical interactions in their entirety and apply all available design options. This design challenge is so complex that it is very difficult to model mathematically. Hence, the search for an automated

solution has proved unsuccessful thus far. To this day, designing analog ICs is mainly based on designer experience and has remained a primarily manual task.

About 90% of the entire development costs of mixed signal and smart power ICs are due to the analog parts, although, based on the number of devices they contain, these parts are only a very small percentage (typically <10%, indicated by the gray vertical arrow in Fig. 1.11) of the chip. This means that analog design productivity is two to three orders of magnitude lower than digital design productivity. This is clearly shown in Fig. 1.11 by the distance between the red lines (red vertical arrow).

Analog circuit design for mixed signal and smart power chips has become a bottleneck, and urgent action is needed to improve the analog design flow in order to prevent further widening of the analog design gap (shown by the lower shaded area and the brown vertical arrow). We shall propose some measures to tackle this issue in physical design at the end of Chap. 4.

1.3 Physical Design

1.3.1 Main Design Steps

A greatly simplified schematic of the design flow for electronic systems is given in Fig. 1.12. It starts with a *specification*, where the desired system functions and performance characteristics for the intended task are set out. Besides the standard signal representations in time and frequency domains, other characterizations, provided as text, figures, tables, and the like, are included in order to specify the design goal as accurately and comprehensively as possible. The specification describes *what* the system is intended to do, with a focus on inputs and outputs; for example, part of a

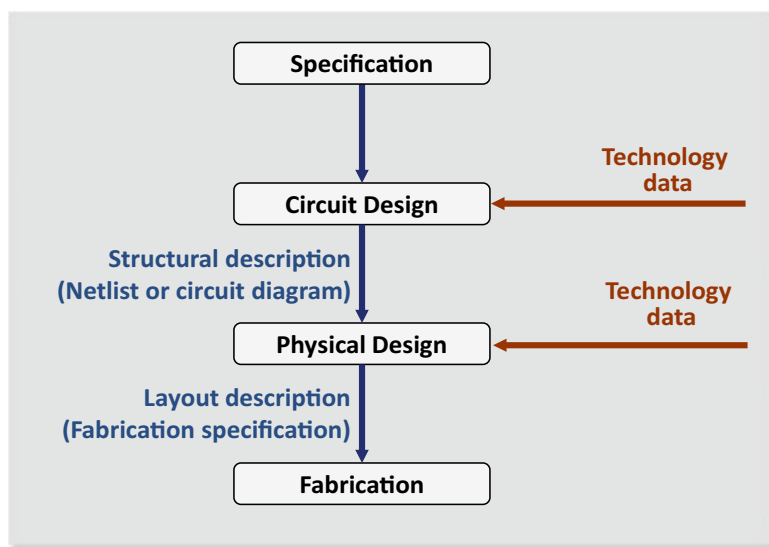


Fig. 1.12 A simplified schematic of the main design steps for an electronic circuit

specification might be: “The system shall receive two digital inputs, using pins 0–7 and 8–15, at a frequency of 1.2 GHz, and produce their multiplicative product as a single 16-bit digital output using pins 16–31, in no more than 5 clock cycles.” The specification describes the task to be performed, but does not define *how* it is to be done. (How the system provides or performs the required functionality is determined as part of the design process in subsequent steps.)

In general, the electronic system is designed to the specification in two major steps, which are outlined next.

Circuit Design

During the circuit design, the functional specification is used to create an electrical circuit that correctly implements the required functionality. For digital circuits this design process will likely include multiple levels of top-down decomposition, where higher-level functionality is broken down iteratively into simpler and simpler lower-level functions, each of which can finally be implemented by a functional unit (for example, an AND adder, a comparator, or a register). The result of this design process is captured in a structural description of the electronic system.

Thus, a *structural description* of the electronic system, such as an IC, is produced as an output of *circuit design*. All the required electrical functional units of the circuit and the electrical connections between them (*nets*) are listed in the structural description. Each net represents an electrical connection—in essence, a wire—that connects an output of a first functional unit to an input of one or more secondary functional units. Due to the complexity of modern electronic systems, this so-called “netlist” is typically organized in a hierarchical tree structure. As such, it contains functional units in the form of standard electronic devices as well as so-called *circuit blocks* or *function blocks* specifying sections of the circuit as subsets of the integrated system.

This structural description, generated during circuit design, serves as the input data to the next major step, the physical design of the circuit. The structural description can be in the form of text as a *netlist*, or in graphical form as a *circuit diagram* (also known as a *schematic diagram* or a *schematic*).

A circuit diagram is a pictorial representation of a netlist, in which the functional units are depicted as symbols and the nets as lines. A simple circuit diagram is drawn as an example on the right in Fig. 1.13. The circuit comprises four basic electronic devices (two resistors R1, R2, two capacitors C1, C2) and a function block (an operational amplifier, shown as a triangle symbol). A circuit diagram and a netlist are effectively equivalent representations of a structural description. Which of the two formats is used depends on the application at hand.

The circuit diagram is produced by graphics-based data entry in a schematic editor in predominantly manual design styles—often used for analog ICs or for PCBs. Symbols for resistors, transistors, etc. stored in *symbol libraries* are loaded and placed on the circuit diagram; the devices are then connected (Fig. 1.13, right). The circuit structure is stored as a netlist in the design tool.




Fig. 1.13 Circuit design: from the specification (left) to the circuit diagram (right)

In highly automated design styles, such as digital IC design, the circuit design data is generated by synthesis procedures and is output as a netlist. A circuit diagram is often not required in these cases.

Physical Design

The purpose of the physical design (aka layout design) is to produce a *fabrication specification* from the structural description of a circuit. The circuit is then manufactured (that is, it is physically realized) based on this specification. *Optimization goals* are pursued and given boundary conditions must be met by this transformation from a structural description to a fabrication specification (“layout”). We discuss these next.

These boundary conditions can be classified into process- and project-specific ones. *Process*-specific boundary conditions describe the options and boundaries of the deployed fabrication technology. They are included in the *technology data* in Fig. 1.12 and must be considered in all designs realized in the given fabrication technology. *Project*-specific boundary conditions, on the other hand, only apply to the product being developed. They are produced during circuit design as another result supplementing the structural description and are generally known as *constraints*. They are instructions for the IC designer concerning the proper functioning of a circuit or the required reliability. In Chap. 4, we shall examine optimization goals and constraints in the context of different design models and design styles.

Along with the fabrication specifications, the result of physical design also defines the exterior appearance—in particular, the dimensions—of the designed electronic system and its component parts. Therefore, the design result is also known as a *layout*, as it is a (physical) interpretation or layout of the (abstract) structural description. The physical design is therefore also called *layout design*.

1.3.2 Physical Design of Integrated Circuits

As we have seen, in semiconductor technology all structures are constructed on a wafer using photolithography. Central to the photolithographic process are masks, which are used to iteratively create microscopic structures on the wafer—the individual electronic devices are created using FEOL processing, and the multi-layered

interconnects are created using BEOL processing, as described previously in reference to Fig. 1.7. The masks contain the production templates for the structures that are to be constructed on the wafer. The purpose of IC physical design is to generate these geometrical structures, which are subsequently used to create masks.

The result of the physical design process is an *IC layout*—a complete image of the chip to be fabricated, which is composed of a stacking or layering of these mask images. This IC layout defines the physical implementation of all circuit elements on an IC chip. Among these elements are (i) the internal structure of the devices, (ii) the devices' placement or arrangement on the chip, (iii) the interconnects, (iv) contacts and vias, and (v) the bond pads (located at the chip margin) for the electrical connection with the environment.

Figure 1.14 on the right side shows a small section of a chip layout, as it would appear in a graphics editor used in layout design. Each graphics element is assigned to a so-called *layer*, which generally corresponds to a mask. The designer uses layer-specific colors, line thicknesses and fill patterns to represent the graphics elements in the graphics editor to help distinguish which elements belong to which layer. For example, the yellow “Contact” layer in Fig. 1.14 shows the mask features used to make holes in the bottom insulating layer that lies directly above the silicon surface. These holes are filled with metal to generate an electrical connection between the silicon surface and the bottom metal layer (see also Fig. 1.7c). These connections are positioned where the electronic devices in silicon must be electrically contacted; this is also the reason why these holes are called *contacts* or *contact holes*.

The physical design of digital and analog IC circuits differs greatly, as we describe below.

Considering the physical design of digital circuits, standard elements with an immutable internal structure are stored as *cells* in a library. These cells, such as logic gates and memories, provide standard functions and their layouts have already been produced. More complex logic blocks are often stored in the library as *macro cells*, essentially larger cells that provide higher-level functions, such as an adder or multiplier, or even higher-level application-specific components, such as ones that implement e.g. communication protocols. Instances of these cells are placed on

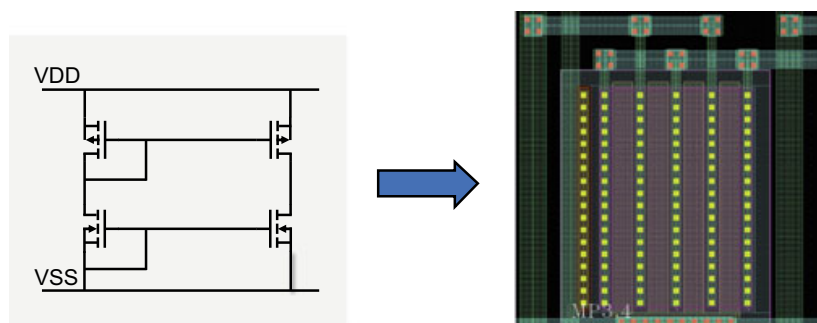


Fig. 1.14 Visualization of the physical design (layout design) of an integrated circuit: converting the structural description of a circuit (a circuit diagram in this case, left) to geometric data, which forms the basis for producing the masks (layout, right)

unoccupied surface areas in the first step in physical design—the *placement*. The interconnects, that make the electrical connections between the cells, are designed in the next *routing* step. Both steps are almost fully automated, which is why they are referred to as *layout synthesis*. A netlist that is suitable as input data for the software programs is used here. We shall examine the processes in (digital) layout synthesis and their application more closely in Chap. 4.

The situation for the physical design of analog circuits is totally different: here, as has been explained in Sect. 1.2.2, most of the design work is still performed manually to this day. Specifically, a circuit’s functionality depends very much on the individual device designs and how they are physically placed relative to one another.¹¹ The layout engineer must understand the circuit and how it works in order to make the right decisions regarding the physical design and placement of the devices. A circuit diagram of the structural description should therefore be drafted when designing analog ICs. This circuit structure allows designers to more easily understand the electrical relationships (the *circuit topology*) and thus the functionality implemented by the circuit.

The first step in designing physical analog circuits is to implement the devices. In this step, each device is sized according to the electrical parameters in the circuit diagram and adapted to other structural criteria for the specific use case. Generators—these are scripts that can automatically produce different layout variants using various parameters—are available for this task. Some of these parameters are defined by the circuit diagram. Other parameters allow the layout engineer to invoke further automated adaptations.

In the placement step, these (analog) devices are arranged on the chip to facilitate subsequent routing. Placement is performed almost entirely manually with the graphics editor due to the complexity of the task and the requirements to be considered. Automated routing software is not widely used in practice either for the subsequent routing step, as the expertise of the layout designer is often needed here for some critical parts of the layout. We shall go into the physical design of analog circuits in more depth in Chap. 6.

For both digital and analog physical design, the finished layout is stored as a graphics file and used to generate the masks, typically a separate mask for each layer. More information will be given on this topic in Chap. 3.

Once the physical design has been completed, it is essential to verify its correctness before committing the design to silicon in the manufacturing process that follows. Many automated verification algorithms are available to check the physical design result, i.e., the layout. Two of these algorithms are obligatory as they verify key quality requirements in chip development. Indeed, without them it would be impossible to flawlessly fabricate modern highly complex ICs. First, the layout is checked for compliance with technological constraints—specified as design rules—by the *design rule check (DRC)*. The manufacturability of a chip layout is confirmed with this verification technique. Second, the *electrical verification* (also called *layout versus*

¹¹ *Matching*—a technique used in analog chip layout—plays a key role here. We shall discuss this very important topic fully in Chap. 6.

schematic check, *LVS*) verifies that the specifications contained in the structural description have been correctly realized in the layout. Specifically, the LVS checks (i) that the devices are properly connected electrically; (ii) that the correct types of devices are in the layout; and (iii) that the devices are correctly parameterized. We shall examine the operation and use of these and other verification tools in Chaps. 3 and 5.

1.3.3 Physical Design of Printed Circuit Boards

In Sect. 1.3.2 we focused on the physical design of chips; these chips are then mounted on printed circuit boards (PCBs) and interconnected to realize the final system. Thus, the physical design of PCBs focuses on the design of a wiring substrate for mounting and electrically connecting electrical and electronic devices. In contrast to IC chip design, the devices are not an integral part of the technological fabrication process for the circuit, rather they are independently provided from an external source.

As explained in Sect. 1.1.1 and illustrated in Fig. 1.2, PCBs are made up of many conductive and non-conductive layers connected by vias. These layers are mapped into corresponding data structures in the design tool—in a similar manner to IC physical design. In addition to these layers used for connecting the devices, further layers are required for fabrication, such as solder resist and solder paste masks or placement imprint. The design of the graphics structures in all these layers is known as “PCB physical design” or “PCB layout design”.

The input data for the layout design of a PCB are specified by a circuit diagram. The following steps are required to produce a PCB layout:

- (1) Specifying the placement of the devices,
- (2) Defining the size of the board and number of the routing layers,
- (3) Specifying the location of the interconnects (wires) on the wiring substrate (the routing step) and arranging the via positions.

All the necessary structures for physically attaching and electrically connecting a device on the PCB are known collectively as a *footprint*, sometimes also called a *land pattern*. The footprint describes *pads* (contact surfaces) and the placement imprint as polygons; it also contains the necessary holes (for THDs) and through contacts (vias). Footprints are stored in *footprint libraries*.

Creating the *layout file* for the PCB is the first step in the layout process. Here, the devices are extracted from the circuit diagram and the matching footprints are loaded from the footprint library. The footprints are then placed—in other words, their locations on the PCB are defined (Fig. 1.15).

The physical routes for the electrical connections (interconnects) between the devices have not yet been determined. They are shown symbolically as “rubber bands” (Fig. 1.15, right). Next, the shape and location of the interconnect tracks, and the layer to which they belong, are defined in the routing step (Fig. 1.16).

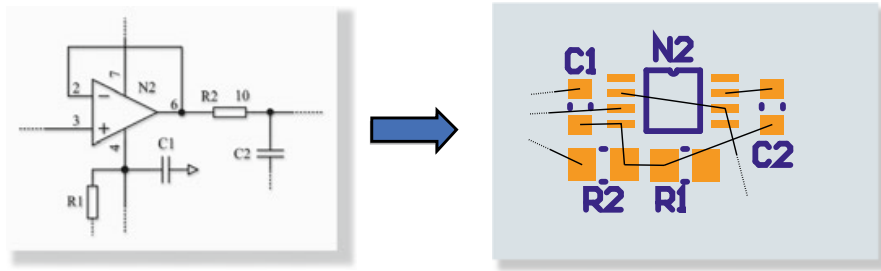


Fig. 1.15 First step in the PCB layout process: Specifying the placement of the devices on a PCB, originating from a circuit diagram (left)




Fig. 1.16 Second step in the PCB layout process: Specifying the routing of the interconnects which includes arranging the via positions and the layer allocation

After placement of the devices and the routing step, the layout design is verified. Specifically, the layout is checked for short circuits (electrical rule check, ERC) and for compliance with constraints necessary for proper functioning and manufacture (e.g., design rule check, DRC).

The required fabrication data is then exported. In contrast to an IC layout, where the entire finished layout is stored in one graphics file, different files and formats are needed for PCB manufacture. The fabrication data consists of a set of *Gerber files* that describe the conductor tracks for the individual layers, the solder resist and solder paste mask, and the placement imprint consisting of polygons. A *drill file* containing the diameters and coordinates of all drill holes in the PCB is required, as well. Finally, a *pick-and-place file* containing the location and alignment of the devices is generated for automated placement of the devices in the assembly process.

1.4 Motivation and Structure of This Book

As described throughout this chapter, design components are instantiated as geometric representations during physical (layout) design. In other words, all electronic devices, cells, gates, transistors, etc., are realized with fixed shapes and sizes, assigned locations (placement), and have appropriate wiring connections (routing) completed in metal layers. The result of layout design is a set of manufacturing specifications that must subsequently be verified prior to the actual manufacturing process.

The physical layout directly affects circuit performance, area, reliability, power consumption, and manufacturing yield. As such, the quality of layout design significantly influences the quality of the resulting electronic circuit, irrespective of whether it is an IC or a PCB.

The continuing miniaturization leads to increasing design problems that must be overcome by the layout designer, and include worsening parasitic disturbances, and more and more technological constraints. As a result, the demand for experienced layout designers continues to grow. At the same time, the need for new methodologies and tools in layout design is also increasing.

This book addresses all these challenges. It presents the fundamental knowledge of layout design from the ground up, from technological constraints to reliability requirements. Chapter by chapter, the book provides the awareness a layout designer must possess to convert a structural description produced during circuit design into the physical layout of transistors, cells, devices and wires on the surface of the chip or the board.

While all relevant aspects of layout design are covered (digital and analog, IC and PCB layout), the reader will notice in some parts a focus on analog layout design. This is due to the stronger need for manual work in analog design: here, the expert knowledge of the actual layout designer is more crucial than in the often fully-automated digital design flows. Nevertheless, this book is intended to provide the fundamentals of physical design irrespective of its specific application, as the basic knowledge is the same for all abstraction layers.

This chapter provided a basic and introductory grounding in the technologies, the tasks and the methodologies needed for designing the layout of an electronic circuit.

The following Chap. 2 introduces in detail the engineering know-how to transform silicon into devices and thus, integrated circuits. This knowledge is crucial for any IC layout designer, as the boundary conditions that must be considered during layout design result directly from the specific semiconductor technology subsequently applied for manufacturing the electronic circuit. This chapter should give the reader the requisite understanding of the technology for which the layout is targeted.

Chapter 3 describes the data interfaces of layout design. A layout designer must be aware of the implications of these “bridges to technology”. In this chapter, we first introduce circuit and layout data structures, such as netlists, layers and polygons. We also investigate the specific links that exist between layout design and targeted technology, such as mask data, design rules and libraries. A special emphasis is paid to the layout post-processing flow. Here we explain all the steps needed to transpose the layout data of an IC into the mask data (i.e., the specification for IC fabrication), including chip finishing, reticle layout, and graphic manipulation processes to achieve compliance of the graphic data with manufacturing requirements.

After covering technologies and the way the design process interfaces to them, physical design is dealt with in Chap. 4. Here, the flow, constraints and strategies of today’s state-of-the-art physical design are introduced. We investigate the various types of constraints, the design models and styles, and discuss the analog-digital

design gap. A detailed look into the specifics of analog design, including its outlook, are also presented. In summary, this chapter provides the basic knowledge any engineer must possess about physical design methodologies.

Due to its high complexity, physical design is split into several primary steps. These steps, which transform a netlist into optimized layout data, are dealt with one by one in Chap. 5. We first provide an overview on how to generate a netlist, either by using hardware description languages in digital design, or by deriving it from a schematic as is common in analog design. Then the physical design steps, such as partitioning, floorplanning, placement and routing, are presented in detail.

After physical design is completed, the layout must be fully verified to ensure correct electrical and logical functionality. Some problems found during physical verification can be tolerated if their impact on chip yield is negligible. In other cases, the layout must be modified, but these changes must be minimal and should not introduce new problems. These options for layout verification are discussed in Chap. 5 as well. We also touch on layout post-processing methodologies, such as resolution enhancement techniques (RET), that might impact physical design.

While the physical-design steps presented thus far are universal, analog circuits require additional layout techniques. Any analog layout designer must be fully aware of these analog design techniques, so we introduce the most common analog devices, cell generators, symmetry and matching principles in Chap. 6.

As reliability of our circuits is becoming a growing concern, the final Chap. 7 summarizes reliability aspects that are of relevance in layout design. We start by presenting reliability issues that can lead to *temporary* circuit malfunctions. We discuss in this context parasitic effects in the bulk of silicon, at its surface, and in the interconnect layers. Afterwards, we deal with the growing challenges of preventing ICs from *irreversible* damage. This requires the investigation of overvoltage events and migration processes, such as electromigration, thermal and stress migration. The goal of this chapter is to summarize the state of the art in reliability-driven design and related mitigating measures. This knowledge can be applied by a circuit designer to increase the reliability of the generated layout.

References

1. L. Berlin, *The Man Behind the Microchip: Robert Noyce and the Invention of Silicon Valley* (Oxford University Press, 2005), ISBN 978-019516343-8. <https://doi.org/10.1093/acprof:oso/9780195163438.001.0001>
2. R. Fischbach, J. Lienig, T. Meister, From 3D circuit technologies and data structures to interconnect prediction, in *Proceedings of 2009 International Workshop on System Level Interconnect Prediction (SLIP)* (2009), pp. 77–84. <https://doi.org/10.1145/1572471.1572485>
3. J. Kilby, Patent No. US3138743: Miniaturized electronic circuits. Patent filed Feb. 6, 1959, published June 23, 1964
4. G.E. Moore, Cramming more components onto integrated circuits. *Electronics* **38**(8), 114–117 (1965). <https://doi.org/10.1109/N-SSC.2006.4785860>
5. R.N. Noyce, Patent No. US2981877: Semiconductor device and lead structure. Patent filed June 30, 1959, published April 25, 1961
6. https://en.wikipedia.org/wiki/Transistor_count. Accessed 1 Jan 2020