Cambridge University Press
978-0-521-88267-5 - Digital Integrated Circuit Design: From VLSI Architectures to CMOS Fabrication
Hubert Kaeslin and ETH Zurich
Excerpt
More information

Chapter 1

Introduction to Microelectronics

1.1 Economic impact

Let us begin by relating the worldwide sales of semiconductor products to the world's gross domestic product (GDP).[1] In 2005, this proportion was 237 GUSD out of 44.4 TUSD (0.53%) and rising.

Assessing the significance of semiconductors on the basis of sales volume grossly underestimates their impact on the world economy, however. This is because microelectronics is acting as a technology driver that enables or expedites a range of other industrial, commercial, and service activities. Just consider

- The computer and software industry,
- The telecommunications and media industry,
- Commerce, logistics, and transportation,
- Natural science and medicine,
- Power generation and distribution, and — last but not least —
- Finance and administration.

Microelectronics thus has an enormous economic leverage as any progress there spurs many, if not most, innovations in "downstream" industries and services.

A popular example . . .

After a rapid growth during the last three decades, the electric and electronic content of passenger cars nowadays makes up more than 15% of the total value in simpler cars and close to 30% in well-equipped vehicles. What's more, microelectronics is responsible for the vast majority of improvements that we have witnessed. Just consider electronic ignition and injection that have subsequently been combined and extended to become electronic engine management. Add to that anti-lock brakes and anti-skid stability programs, trigger circuits for airbags, anti-theft equipment,

[1] The GDP indicates the value of all goods and services sold during some specified year.

Cambridge University Press
978-0-521-88267-5 - Digital Integrated Circuit Design: From VLSI Architectures to CMOS Fabrication
Hubert Kaeslin and ETH Zurich
Excerpt
More information

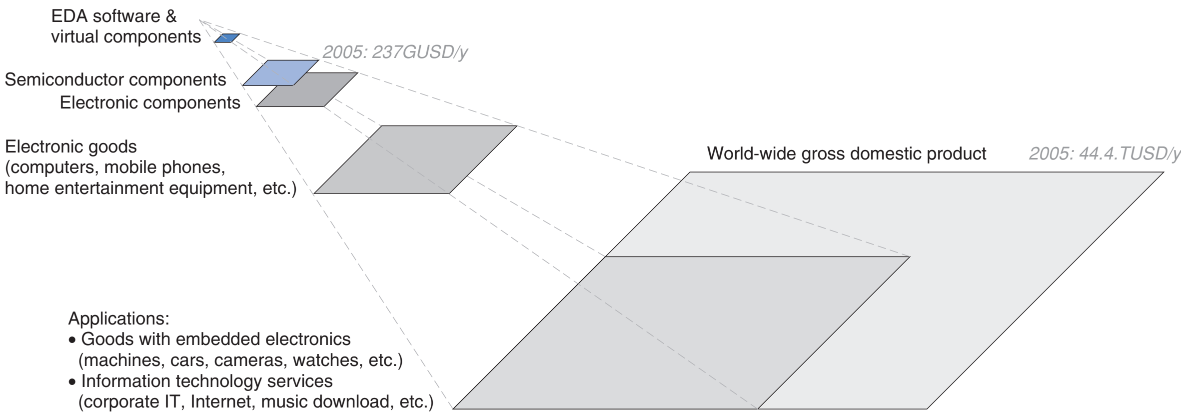

Fig. 1.1 Economic leverage of microelectronics on "downstream" industries and services.

automatic air conditioning, instrument panels that include a travel computer, remote control of locks, navigation aids, multiplexed busses, electronically controlled drive train and suspension, audio/video information and entertainment, and upcoming night vision and collision avoidance systems. And any future transition to propulsion by other forms of energy is bound to intensify the importance of semiconductors in the automotive industry even further.

Forthcoming innovations include LED illumination and headlights, active "flywheels", hybrid propulsion, electronically driven valve trains, brake by wire, drive by wire, and, possibly, 42 V power supply to support the extra electrical load.

. . . and its less evident face

Perhaps less obvious but as important are the many contributions of electronics to the processes of development, manufacturing, and servicing. Innovations behind the scenes of the automotive industry include computer-aided design (CAD) and finite element analysis, virtual crash tests, computational fluid dynamics, computer numeric-controlled (CNC) machine tools, welding and assembly robots, computer-integrated manufacturing (CIM), quality control and process monitoring, order processing, supply chain management, and diagnostic procedures.

This almost total penetration has been made possible by a long-running drop of **cost per function**. Historically, costs have been dropping at a rate of 25% to 29% per year according to [1]. While computing, telecommunication, and entertainment products existed before the advent of microelectronics, today's anywhere, anytime information and telecommunication society would not have been possible without it; just compare the electronic devices in fig.1.2.

Observation 1.1. *Microelectronics is the enabler of information technology.*

On the positive side, microelectronics and information technology improve speed, efficiency, safety, comfort, and pollution control of industrial products and commercial processes, thereby bringing competitive advantages to those companies that take advantage of them.

CAMBRIDGE

Cambridge University Press
978-0-521-88267-5 - Digital Integrated Circuit Design: From VLSI Architectures to CMOS Fabrication
Hubert Kaeslin and ETH Zurich
Excerpt
More information

Fig. 1.2 Four electronic products that take advantage of microelectronics opposed to analogous products that do not. The antiquated devices operate with vacuum tubes, discrete solid-state devices, and other electronic components but include no large-scale integrated circuits. Also observe that, were it not for display size and audio volume, one might replace all four devices with Apple's iPhone that has brought seamless system integration to even higher levels (photos courtesy of Alain Kaeslin).

On the negative side, the rapid progress, most of which is ultimately fueled by advances in semiconductor manufacturing technology, also implies a rapid obsoletion of hardware and software products, services, know-how, and organizations. A highly cyclic economy is another unfortunate trait of the semiconductor industry [2].

1.2 Concepts and terminology

An **integrated circuit** (IC) is an electronic component that incorporates and interconnects a multitude of miniature electronic devices, mostly transistors, on a single piece of semiconductor material, typically **silicon**.[2] Many such circuits are jointly manufactured on a thin semiconductor wafer with a diameter of 200 or 300 mm before they get cut apart to become (naked) **dies**. The sizes of typical dies range between a pinhead and a large postage stamp. The vast majority of ICs, or **(micro)chips** as they are colloquially referred to, gets individually encapsulated in a hermetic package before being soldered onto **printed circuit boards** (PCB).

The rapid progress of semiconductor technology in conjunction with marketing activities of many competing companies — notably trademark registration and eye catching — has led to a plethora of terms and acronyms, the meaning of which is not consistently understood by all members of the microelectronics community. This section introduces the most important terms, clarifies what they mean, and so prepares the ground for more in-depth discussions.

Depending on perspective, microchips are classified according to different criteria.

1.2.1 The Guinness book of records point of view

In a world obsessed with records, a prominent question is “How large is that circuit?”

Die size is a poor metric for design complexity because the geometric dimensions of a circuit greatly vary as a function of technology generation, fabrication depth, and design style.

Transistor count is a much better indication. Still, comparing across logic families is problematic as the number of devices necessary to implement some given function varies.[3]

Gate equivalents attempt to capture a design’s hardware complexity independently from its actual circuit style and fabrication technology. One gate equivalent (GE) stands for a two-input NAND gate and corresponds to four MOSFETs in static CMOS; a flip-flop takes roughly 7 GEs. Memory circuits are rated according to storage capacity in bits. Gate equivalents and memory capacities are at the basis of the naming convention below.

[2] This is a note to non-Angloamerican readers made necessary by a tricky translation of the term silicon.

English	German	French	Italian	meaning
silicon	Silizium	silicium	silicio	Si, the chemical element with atomic number 14
silicone	Silikon	silicone	silicone	a broad family of polymers of Si with hydrocarbon groups that comprises viscous liquids, greases, and rubber-like solids

[3] Consistent with our top-down approach, there is no need to know the technicalities of CMOS, TTL, and other logic families at this point. Interested readers will find a minimum of information in appendix 1.6.

circuit complexity	GEs of logic + bits of memory
small-scale integration (SSI)	1–10
medium-scale integration (MSI)	10–100
large-scale integration (LSI)	100–10 000
very-large-scale integration (VLSI)	10 000–1 000 000
ultra-large-scale integration (ULSI)	1 000 000 . . .

Clearly, this type of classification is a very arbitrary one in that it attempts to impose boundaries where there are none. Also, it equates one storage bit to one gate equivalent. While this is approximately correct when talking of static RAM (SRAM) with its four- or six-transistor cells, the single-transistor cells found in dynamic RAMs (DRAMs) and in ROMs cannot be likened to a two-input NAND gate. A better idea is to state storage capacities separately from logic complexity and along with the memory type concerned, e.g. 75 000 GE of logic + 32 kibit SRAM + 512 bit flash ≈ 108 000 GE overall complexity.[4]

One should not forget that circuit complexity per se is of no merit. Rather than coming up with inflated designs, engineers are challenged to find the most simple and elegant solutions that satisfy the specifications given in an efficient and dependable way.

1.2.2 The marketing point of view

In this section, let us adopt a market-oriented perspective and ask
"How do functionality and target markets relate to each other?"

General-purpose ICs

The function of a general-purpose IC is either so simple or so generic that the component is being used in a multitude of applications and typically sold in huge quantities. Examples include gates, flip-flops, counters, and other components of the various 7400 families but also RAMs, ROMs, microcomputers, and most digital signal processors (DSPs).

Application-specific integrated circuits

Application-specific integrated circuits (ASICs) are being specified and designed with a particular purpose, equipment, or processing algorithm in mind. Initially, the term had been closely associated with **glue logic**, that is with all those bus drivers, decoders, multiplexers, registers, interfaces, etc. that exist in almost any system assembled from highly integrated parts. ASICs have evolved from substituting a single package for many such ancillary functions that originally had to be dispersed over several SSI/MSI circuits.

Today's highly-integrated ASICs are much more complex and include powerful systems or subsystems that implement highly specialized tasks in data and/or signal processing. The term

[4] Kibi- (ki), mebi- (Mi), gibi- (Gi), and tebi- (Ti) are binary prefixes recommended by various standard bodies for 2^{10}, 2^{20}, 2^{30}, and 2^{40} respectively because the more common decimal SI prefixes kilo- (k), mega- (M), giga- (G) and tera- (T) give rise to ambiguity as $2^{10} \neq 10^3$. As an example, 1 MiByte = 8 Mibit = $8 \cdot 2^{20}$ bit.

system-on-a-chip (SoC) has been coined to reflect this development. Overall manufacturing costs, performance, miniaturization, and energy efficiency are key reasons for opting for ASICs.

Still from a marketing point of view, ASICs are subdivided further into application-specific standard products and user-specific ICs.

Application-specific standard product (ASSP). While designed and optimized for a highly specific task, an application-specific standard product circuit is being sold to various customers for incorporation into their own products. Examples include graphics accelerators, multimedia chips, data compression circuits, forward error correction devices, ciphering/deciphering circuits, smart card chips, chip sets for cellular radio, serial-ATA and Ethernet interfaces, wireless LAN chips, and driver circuits for power semiconductor devices, to name just a few.[5]

User-specific integrated circuit (USIC). As opposed to ASSPs, user-specific ICs are being designed and produced for a single company that seeks a competitive advantage for their products; they are not intended to be marketed as such. Control of innovation and protection of proprietary know-how are high-ranking motivations for designing circuits of this category. Parts are often fabricated in relatively modest quantities.

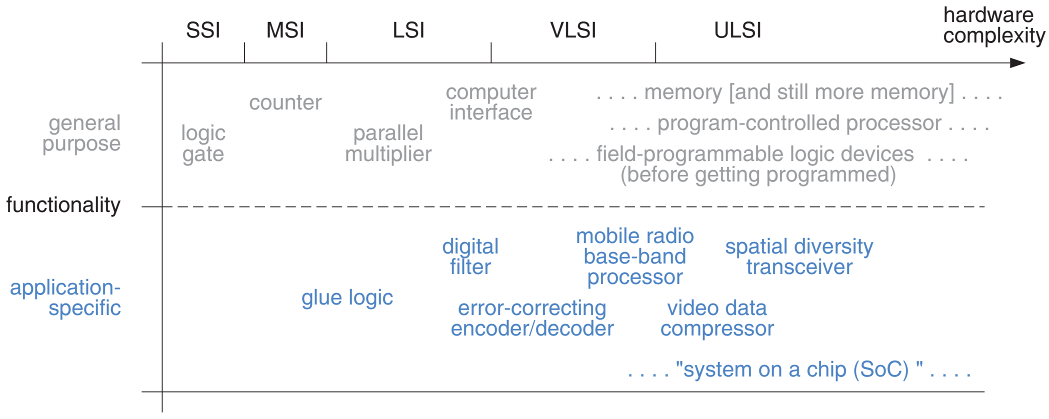

Fig. 1.3 ICs classified as a function of functionality and hardware complexity.

1.2.3 The fabrication point of view

Another natural question is
"To what extent is a circuit manufactured according to user specifications?"

[5] Microprocessors that have their instruction sets, input/output capabilities, memory configurations, timers, and other auxiliary features tailored to meet specific needs also belong to the ASSP category.

Full-custom ICs

Integrated circuits are manufactured by patterning multiple layers of semiconductor materials, metals, and dielectrics. In a full-custom IC, all such layers are patterned according to user specifications. Fabricating a particular design requires wafers to go through all processing steps under control of a full set of lithographic **photomasks** all of which are made to order for this very design, see fig.1.4. This is relevant from an economic point of view because mask manufacturing is a dominant contribution to non-recurring VLSI fabrication costs. A very basic CMOS process featuring two layers of metal requires some 10 to 12 fabrication masks, any additional metal layer requires two more masks. At the time of writing (late 2007), one of the most advanced CMOS processes comprises 12 layers of metal and involves some 45 lithography cycles.

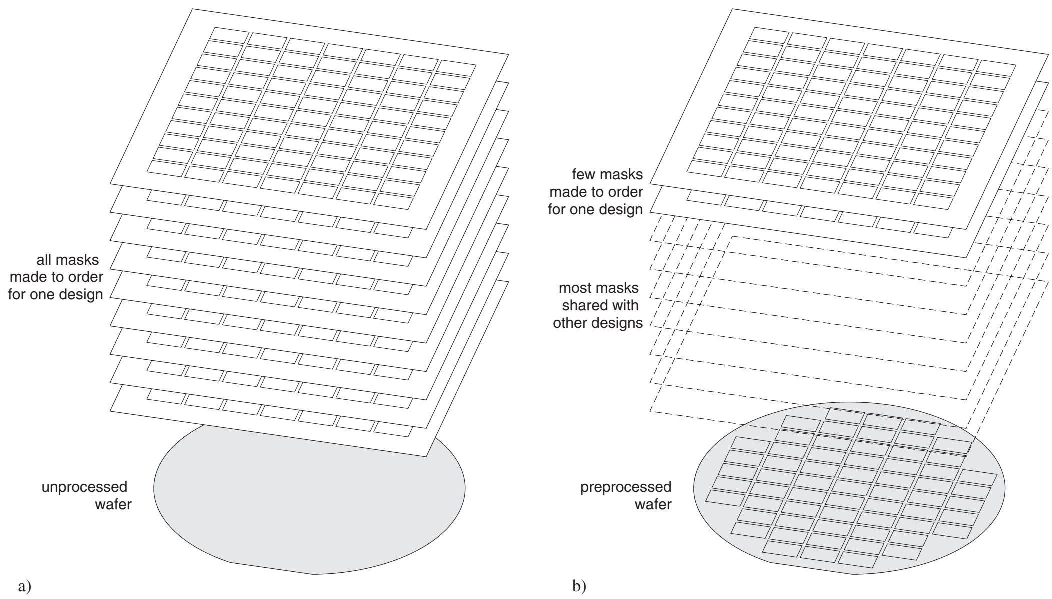

Fig. 1.4 Full-custom (a) and semi-custom (b) mask sets compared.

Semi-custom ICs

Only a small subset of fabrication layers is unique to each design. Customization starts from preprocessed wafers that include large quantities of prefabricated but largely uncommitted primitive items such as transistors or logic gates. These so-called **master wafers** then undergo a few more processing steps during which those primitives get interconnected in such a way as to complete the electrical and logic circuitry required for a particular design. As an example, fig.1.5 shows how a logic gate is manufactured from a few pre-existing MOSFETs by etching open contact holes followed by deposition and patterning of one metal layer.

In order to accommodate designs of different complexities, vendors make masters available in various sizes ranging from a couple of thousands to millions of usable gate equivalents. Organization and customization of semi-custom ICs have evolved over the years.

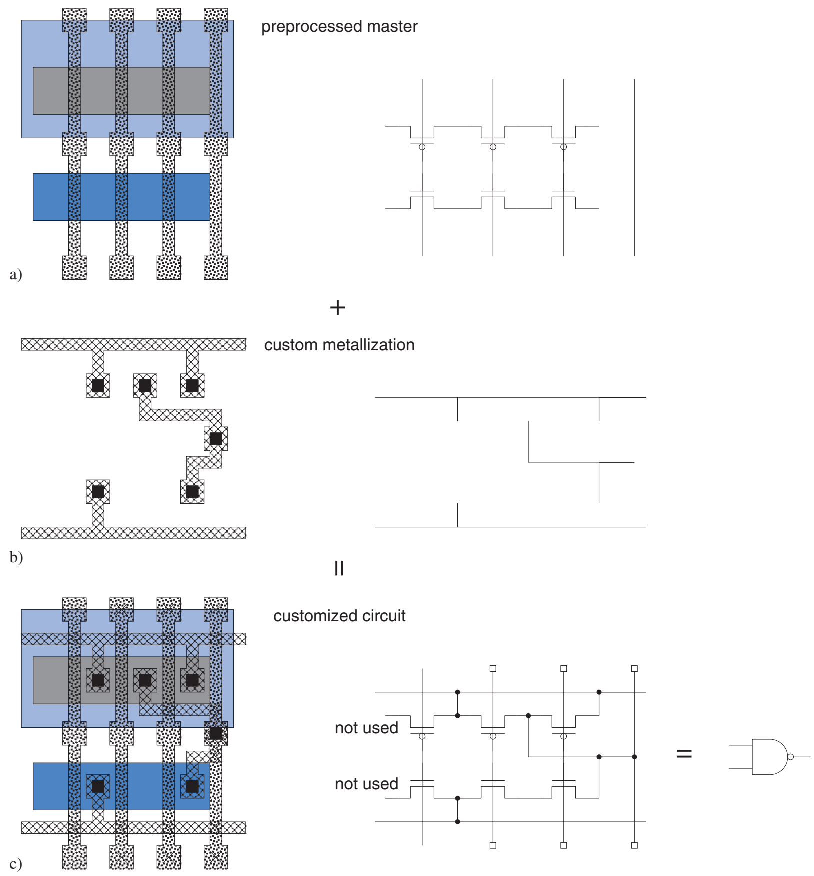

Fig. 1.5 Customization of a gate array site (simplified). A six-pack of prefabricated MOS transistors (a), metal pattern with contact openings (b), and finished 2-input NAND gate (c).

Gate array, aka channeled gate array. Originally, sites of a few uncommitted transistors each were arranged in long rows that extended across most of the die's width. Metal lines were then used to connect the prefabricated transistors into gates and the gates into circuits. The number of custom photomasks was twice that of metal layers made to order. As long as no more than two layers of metal were available, special routing channels had to be set aside in between to accommodate the necessary intercell wiring, see fig.1.6a.

Sea-of-gates. When more metals became available in the early 1990s, those early components got displaced by channelless sea-of-gate circuits because of their superior layout density. The availability of higher-level metals allowed for routing over gates and bistables customized on the layers underneath, so dispensing with the waste of routing channels, see fig.1.6b. More metals further made it possible to insulate adjacent transistors electrically where needed, doing away with periodic gaps in the layout. Sea-of-gates also afforded more flexibility for accommodating highly repetitive structures such as RAMs and ROMs.

Structured ASIC. A decade later, the number of metal layers had grown to a point where it became uneconomical to customize them all. Instead, transistors are prefabricated and preconnected into small generic subcircuits such as NANDs, MUXes, full-adders, and bistables with the aid of the lower layers of metal. Customization is confined to interconnecting those subcircuits on the top two to four metal layers. What's more, the design process is accelerated as supply and clock distribution networks are largely prefabricated.

Fabric. Exploding mask costs and the limitations of sub-wavelength lithography currently work against many custom-made photomasks. The idea behind fabrics is to standardize the metal layers as much as possible. A subset of them is patterned into fixed segments of predetermined lengths which get pieced together by short metal straps, aka jumpers, on the next metal layer below or above to obtain the desired wiring. Customization is via the vertical contact plugs, called vias, that connect between two adjacent layers.

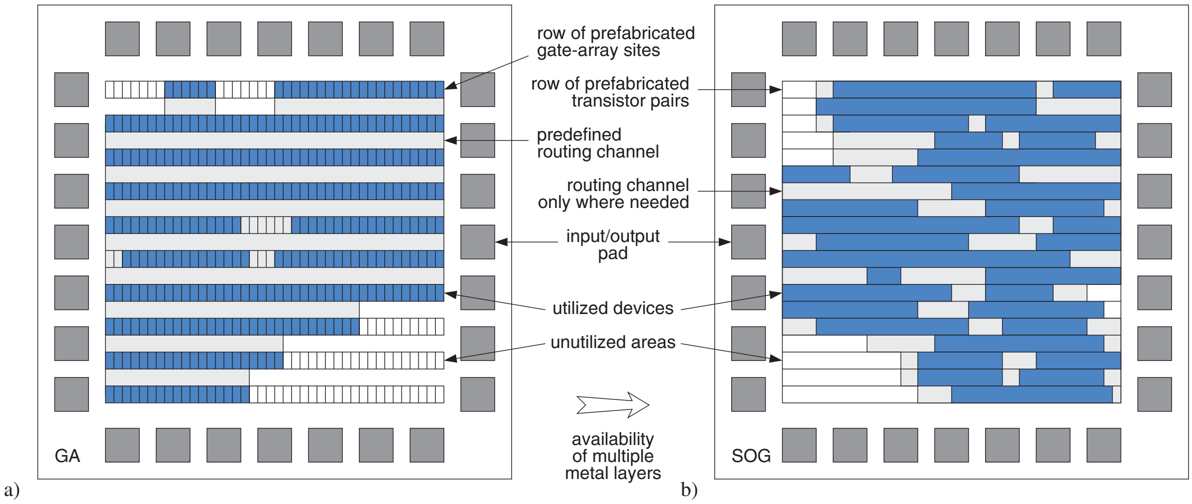

Fig. 1.6 Floorplan of channeled gate-array (a) versus channelless semi-custom circuits (b).

Due to the small number of design-specific photomasks and processing steps, semi-custom manufacturing significantly reduces the non-recurring costs as well as the turnaround time.[6] Conversely, prefabrication necessarily results in non-optimal layouts. Note the unused transistor pair in fig.1.5, for instance, or think of the extra parasitic capacitances and resistances caused by standardized wiring. Prefabrication also implies a self-restraint to fixed transistor geometries, thereby further limiting circuit density, speed, and energy efficiency. Lastly, not all semi-custom masters accommodate on-chip memories equally well.

[6] **Turnaround time** denotes the time elapsed from coming up with a finalized set of design data until physical samples become available for testing.

Incidentally, be informed that the concept of metal customization is also applied to analog and mixed-signal circuits. Prefabricated masters then essentially consist of uncommitted transistors (MOSFETs and/or BJTs) and of passive devices.[7]

Field-programmable logic

Rather than manufacturing dedicated layout structures, a generic part is made to assume a user-defined circuit configuration by purely electrical means. Field-programmable logic (FPL) devices are best viewed as "soft hardware". Unlike semi- or full-custom ASICs, FPL devices offer turnaround times that range from a few seconds to a couple of minutes; many product families even allow for in-system configuration (ISC).

The key to obtaining various gate-level networks from the same hardware resources is the inclusion of electrical links that can be done — and in many cases also undone — long after a device has left the factory. Four configuration technologies coexist today; they all have their roots in memory technology (SRAM, PROM, flash/EEPROM, and EPROM). For the moment, you can think of a programmable link as some kind of fuse.

A second dimension in which commercially available parts differ is the organization of on-chip hardware resources. **Field-programmable gate arrays** (FPGAs), for instance, resemble mask-programmed gate arrays (MPGAs) in that they are organized into a multitude of logic sites and interconnect channels. In this text, we will be using the term **field-programmable logic** (FPL) as a collective term for any kind of electrically configurable IC regardless of its capabilities, organization, and configuration technology.[8]

FPL was initially confined to glue logic applications, but has become an extremely attractive proposition for smaller volumes, for prototyping, when a short time to market is paramount, or when frequent modifications ask for agility. Its growing market share affords FPL a more detailed discussion in section 1.4. What also contributed to the success of FPL is the fact that many issues that must be addressed in great detail when designing a custom circuit are implicitly solved when opting for configuring an FPL device instead, just consider testability, I/O subcircuits, clock and power distribution, embedded memories, and the like.

Standard parts

By standard part, aka commercial off-the-shelf (COTS) component, we mean a catalog part with no customization of the circuit hardware whatsoever.

1.2.4 The design engineer's point of view

Hardware designers will want to know the answer to
"Which levels of detail are being addressed during a part's design process?"

[7] Microdul MD300 and Zetex 700 are just two examples.

[8] Referring to all such parts as "field-configurable" would be preferable as this better reflects what actually happens. This would also avoid confusion with program-controlled processors. Yet, the term "programmable" has gained so much acceptance in acronyms such as PLA, PAL, CPLD, FPGA, etc. that we will stay with it.