CHAPTER 1

The Big Picture

1. WHAT IS A CHIP?

A *chip* (*integrated circuit* or *IC*) is a miniaturized electronic circuit that is manufactured on the surface of a thin substrate of semiconductor material.

Functionally, a chip is a hardware component that can perform certain desired functions. For example, a simple chip might be designed to perform a simple function of a logic NOR (such as the 4000-series CMOS, dual 3-input NOR gate and NOT gate shown in Figure 1.1), a simple operational amplifier, or an *analog-to-digital converter (ADC)*. However, a complex *system on chip (SoC)* performs much more complicated tasks (see Figure 1.2). Examples include those for video decoders, cellular phones, network routers, or general-purpose CPUs for personal computers.

Structurally, chips are manufactured on a semiconductor material called silicon. Basic components such as transistors, diodes, resistors, inductors, and capacitors are constructed on the silicon. Those basic components make up the chip, simple or complex. Simple chips may only contain hundreds of those basic components, whereas complex chips may contain hundreds of millions of those components. Since 1959 (the year that the first integrated circuit-related patent was filed by Jack Kilby), several terms has been created to reflect the status of integrated-circuit development: *small-scale integration (SSI)* for tens of transistors on a chip, *medium-scale integration (MSI)* for hundreds of transistors per chip, *large-scale integration (LSI)* with tens of thousands of transistors per chip, and *very large scale integration (VLSI)* with hundreds of thousands of transistors. *Ultra large scale integration (ULSI)* and system-on-chip (SoC) are the latest terms to cover the modern, ultracomplicated chips with billions of transistors on a single chip. All chips are roughly classified as one of three types: purely digital, analog, or mixed-signal.

Application-wise, chips can be designed to target various applications: video/graphic, audio, communications, networking, general-purpose personal computing, supercomputing, automotive, industry control, medical instrument, and military.

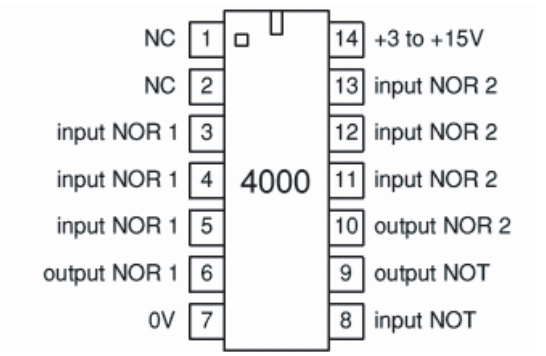

Figure 1.1. 4000-series CMOS, dual 3-input NOR gate and NOT gate.

The majority of the today's chips are designed for processing signals or manipulating information. Among the tasks performed are collecting, transporting, presenting, processing, or manipulating all kinds of information. And today, information plays a vital role in our daily lives. There are billions of billions of bits of information generated every day to support the normal operations of human society. Every single one of those bits must be processed by some kind of chip. Thus, it is not a surprise that the semicon-

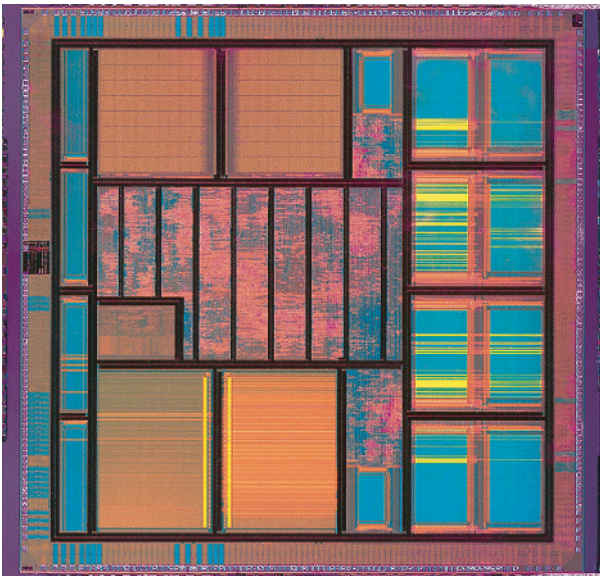

Figure 1.2. System-on-chip example.

ductor chip is built into our life. In addition to this information-processing chip, there are other types of chips that interface with our activities directly by driving electrical, mechanical, or optical components that result in something that we can see, hear, feel, or smell.

Finally, an indivisible part of the chip is the associated software. Software enables the chip to perform certain specific tasks. Software tells the chip *when* and how to do *what.* Without software, the chip is useless, just like a human without brain. Well-developed software can perfect the chip's feature sets, can prolong the chip's life, and can make the difference between success and failure.

When used, a chip is packed in a package, which is mounted on a *printed circuit board (PCB)* and installed in an end-equipment system.

In summary, a chip is an entity that has a large number of transistors integrated into it. Constructing circuits in this manner is an enormous improvement over the manual assembly of circuits using discrete electronic components. Two primary advantages are cost and performance. Cost is low because the components within a chip are built as one unit and not constructed one transistor at a time. Performance is high because the integrated transistors switch quicker and consume less power due to the fact that the components are small and close together.

2. WHAT ARE THE REQUIREMENTS OF A SUCCESSFUL CHIP DESIGN?

In the field of modern VLSI circuit design, constructing a chip from concept to silicon is an ultracomplicated task that involves many factors. For a successful project, the chip must be:

- Structurally correct to achieve its intended design functions
- Functionally correct at the designed clock speed in various working environments (voltage, temperature, and process corner)
- Reliable throughout its life (e.g., 100k hours or eleven years)
- Manufacturing-friendly

Further, it must be built such that:

- It can be handled safely in an assembly line and various other environments without being damaged (e.g., it is protected from *electrostatic discharge* or *ESD* and *latch-up*).
- It can be packaged economically.

- It stays within its power budget.
- Cost is minimized.
- It is manufactured within its time schedule.

And, then, finally, there must be an existing or potential market for this chip.

3. WHAT ARE THE CHALLENGES IN TODAY'S VERY DEEP SUBMICRON (VDSM), MULTIMILLION GATE DESIGNS?

Designing a system-on-chip (tens of millions of gates and larger) in a very deep submicron (90 nm and below) environment is a task of solving many complicated, interdependent problems at once. The design/implementation/verification methodology required is a dynamic development since the challenges involved are ever-changing as the process technology continuously advances. The most outstanding challenges today are listed below:

- *Timing closure.* Timing closure is often the most difficult task in designing a chip owing to the fact that a logic gate's timing behavior (or speed) varies greatly at different temperatures, supply voltages, and process conditions under which the device is built and operated. Moreover, a logic gate's speed is also affected by the drive and load environment surrounding the logic gate. Timing closure means that the chip must run at a designed speed (represented by clock frequency) reliably under all conditions. This is not an easy task to achieve, especially when the process shrinks to even finer geometries and wire delays become more dominating in the overall delay equation.
- *Design verification.* Modern SoC devices contain a large number of components on board, such as processors, memories, on-chip busses, special function macros, and so on. The task of design verification is to ensure that the components work together faultlessly as designed. The magnitude of difficulty involved in this task increases dramatically as integration levels continuously grow and design sizes correspondingly increase.
- *Design integrity.* Design integrity includes *cross talk, IR drop, electromigration (EM), gate oxide integrity (GOI), electrostatic discharge (ESD),* and *latch-up protection.* The chip must be free of these problems before delivery for field application. These issues will become increasingly difficult to resolve as process technology advances.

- *Design for Testability.* The design must be testable for production defects. This testability must be built into the chip. As process geometry continually shrinks, new defect mechanisms constantly surface. As a result, design for testability is a subject investigated unceasingly by process scientists, design engineers, and tool developers.
- *Power budgeting and management.* Modern SoC chips can support more functions and perform tasks at higher speeds. Consequently, they tend to use much more power. In consideration of chip packaging, heat dissipation, and battery life, the chip's power consumption must be reduced.
- *Packaging.* As chips bear more I/Os and consume more power and I/O signals travel at higher speeds, chip packaging becomes more challenging.
- *Design reuse.* Characteristic of the SoC approach is the integration of components, rather than the design of individual components. The more components that can be reused from previous projects, or from other sources, the lower the development costs and the faster the project execution pace.
- *Hardware/software codesign.* Traditionally, software development cannot start until the hardware (the chip) is available. A new methodology, or design environment, is needed to solve this problem.
- *Clock management and distribution.* As a chip's clock speed increases and its clock structure becomes more complex, clock-related design issues will become more challenging.
- *Leakage current management and control.* As process geometries shrink below 90 nm, device leakage current increases dramatically. This problem has moved from backstage to front stage.
- *Design for manufacturability.* As process geometries shrink, device manufacturing requires more rigorous control. This fact imposes additional constraints on the chip design process.

4. WHAT MAJOR PROCESS TECHNOLOGIES ARE USED IN TODAY'S DESIGN ENVIRONMENT?

The mainstream process technology used in today's chip design/manufacturing environment is complementary-symmetry metal oxide semiconductor (CMOS) technology. Other technologies include bipolar, biCMOS, silicon on insulator (SOI), and gallium arsenide (GaAs).

In CMOS technology, *complementary symmetry* refers to the fact that a

CMOS circuit uses symmetrical pairs of p-type and n-type MOSFET transistors for logic functions. Originally, the phrase *metal oxide semiconductor* was a reference to the metal gate electrode placed on top of an oxide insulator. However, in today's CMOS processes, instead of metal, the gate electrode is comprised of a different material, polysilicon. Nevertheless, the name CMOS remains in use for the modern descendants of the original process. Today, in terms of dollar amount, the majority of integrated circuits manufactured are CMOS circuits. This is due to three characteristics of CMOS devices: high noise immunity, low static power, and high density.

The CMOS process has consistently advanced to smaller feature sizes over the years, allowing more circuitry to be packed in one chip, as described by Moore's law. This is the empirical observation made in 1965 by Gordon E. Moore (cofounder of Intel Corporation) that the number of transistors on an integrated circuit for minimum component cost doubles approximately every 24 months. Although Moore's law was initially made in the form of an observation and forecast, it has gradually served as a goal for the entire semiconductor industry. During the past several decades, it has driven semiconductor manufacturers to invest enormous resources for specified increases in processing power that were presumed to be soon attained by one or more of their competitors. In this regard, it can be viewed as a self-fulfilling prophecy.

As feature sizes shrink, costs per unit decrease, circuit speeds increase, and power consumption drops. Therefore, there is fierce competition among the manufacturers to use finer geometries. The status of current processes and the anticipated progress over the next few years is described and documented by the *International Technology Roadmap for Semiconductors (ITRS)*. Currently (in 2006), process geometries have dropped well below one micron. Today, 90 nm technology has been widely used for commercial products. In the near future, it is believed that the 65 nm technology will move front stage. And, tomorrow, 45 nm and 32 nm technologies will take the lead.

Copper has replaced aluminum for wire interconnect signal propagation material due to its improved electric conductivity. The interconnecting metal level has also been increased from two to six or even to seven or more. The power supply voltage for semiconductor chips has continually dropped, due to the shrinking of transistor size, to the current level of 1.1 V.

Table 1.1 presents typical data for each CMOS technology node, in which L_{drawn} represents the minimum transistor channel length and V_{DD} is the supply voltage for transistors. Lower V_{DD} can reduce the transistors' power usage. Density measures the number of logic gates that can be packed into one square millimeter of silicon. Unit gate capacitive indicates

Table 1.1. Typical metrics for CMOS technologies

	180 nm	130 nm	90 nm	65 nm	45 nm
L_{drawn} (nm)	180	95	60	50	40
Metal Level	4–5	5–6	5–6	6–7	7–8
Density (kgates/mm^2)	~ 70	~ 140	~ 250	~ 650	~ 1200
V_{DD} (V)	1.8	1.5	1.2	1.2	1.1
V_T (V)	~ 0.5	~ 0.45	~ 0.45	~ 0.4	~ 0.35
Unit gate capacitive (fp/μm^2)	~ 8	~ 10	~ 10	~ 10	~ 10
Metal resistance (ohms/square)	~ 0.1	~ 0.1	~ 0.12	~ 0.2	~ 0.3
Minimum metal width (μm)	0.25	0.175	0.13	0.1	0.07
Metal pitch (μm)	0.5	0.35	0.27	0.21	0.14

the capacitive loading of transistors, which has a great impact on the logic gate's speed. Metal level is the number of metal layers used for interconnecting. When more metal layers are used, higher densities can be achieved (but the cost is greater). Metal resistance measures the quality of the metal as interconnect material. Minimum metal width and metal pitch indicate the minimum width of metal that is allowed in a chip layout and how close metal can be placed to metal. These two parameters, together with L_{drawn}, primarily determine the gate density of the technology. V_T is the threshold voltage that controls when the NMOS or PMOS transistor switches. The level of V_T has great impact on noise margin and leakage current.

Bipolar refers to an electric circuit made of bipolar junction transistors, which were the devices of choice in the design of discrete and integrated circuits before the 1980s. It offers high speed, high gain, and low output impedance. However, its use has declined in favor of CMOS technology due to its high power consumption and large size.

BiCMOS is a technology that integrates bipolar and CMOS together to take advantage of the high input impedance of CMOS and the low output impedance and high gain of bipolar. A typical example of a BiCMOS circuit is a two-stage amplifier, which uses MOS transistors in the first stage and bipolar transistors in the second. However, BiCMOS as a fabrication process is not nearly as mature as either Bipolar or CMOS. It is very difficult to fine-tune both the bipolar and MOS components without adding extra fabrication steps and, consequently, increasing the cost.

Silicon on insulator (SOI) is a layered structure consisting of a thin layer of silicon fabricated on an insulating substrate. This process reduces the amount of electrical charge that a transistor must move during a switching operation and thus increases circuit speed and reduces switching energy (an improvement over CMOS). Moreover, SOI devices are inherently latch-up

resistant, and there is a significant reduction in transistor leakage current, which makes this technology an attractive choice for low-power circuit design. However, the production of SOI chips requires restructured CMOS fabrication methods and facilities. Thus, it costs more to produce SOI chips, so they are generally used for high-end applications.

As we move toward 45 nm and 32 nm nodes, *multigate FETS (MuGFETs)* are increasingly being considered as a necessary alternative to keep pace with Moore's law. MuGFET is the general term for the class of devices that gain extra component width by allowing vertical active gates. FinFETs and trigates are examples of these devices. This new technology relies heavily on using high-quality, very thin SOI wafers as a starting material. Another popular SOI technology is *silicon on sapphire (SOS)*, which is used for special radiation-hardening applications in the military and aerospace industries.

Gallium arsenide (GaAs) is a semiconductor that has some electrical properties that are superior to silicon's. It has higher saturated electron velocity and higher electron mobility, allowing it to function at much higher frequencies. GaAs devices generate less noise than silicon devices. Also, they can be operated at higher power levels than the equivalent silicon device because they have higher breakdown voltages. These properties make GaAs circuitry ideal for mobile phones, satellite communications, microwave point-to-point links, and radar systems. However, high fabrication costs and high power consumption have made GaAs circuits unable to compete with silicon CMOS circuits in most applications.

5. WHAT ARE THE GOALS OF NEW CHIP DESIGN?

When a company makes a decision to invest in a project to create a product (designing a chip), the ultimate goal is to generate maximum profit from this investment. The approach to pursuing this goal is by conducting business "faster, better, and cheaper."

Faster means that the new chip must operate faster than its predecessors or faster than similar chips produced by competitors, which requires it to perform specific tasks in less time.

Better refers to the fact that the chip must support more functions (do more) than its predecessors.

Cheaper means that the cost of developing and manufacturing the new chip must be kept to a minimum.

This desire to develop something "faster, better, and cheaper" has motivated scientists and engineers working in this field to make enormous tech-

nical strides and will continue to drive them as they work to create superior products. This will, in turn, make our lives more enjoyable.

Perhaps the only exceptions to this ruthless pursuit are projects that are research oriented or not for profit or those produced for the government. In these cases, cheaper is not a concern. Therefore, faster and better can be pursued at whatever cost is required and on whatever schedule is demanded.

6. WHAT ARE THE MAJOR APPROACHES OF TODAY'S VERY LARGE SCALE INTEGRATION (VLSI) CIRCUIT DESIGN PRACTICES?

The major approaches for modern chip design practice follow:

- Custom design
- Field programmable gate array (FPGA)
- Standard cell-based design (ASIC)
- Platform/structured ASIC

In the *custom design* approach, each individual transistor is designed and laid out manually. The main advantage of this method is that the circuit is highly optimized for speed, area, or power. This design style is only suitable for very high performance circuitries, however, due to amount of manual work involved.

Field programmable gate arrays (FPGAs) are semiconductor devices that are comprised of programmable logic components and programmable interconnects. The programmable logic components are programmed to duplicate the functionality of basic logic gates, such as AND, OR, XOR, or NOT gates, or more complex combinational functions, such as decoders, or certain simple math functions. Structurally, the FPGA approach is a chip implementation methodology in which the base layers are premanufactured. When implemented in FPGA, only metal layers need be programmed.

In the past, this design approach was reserved primarily for emulations and prototypes. However, there are increasingly more FPGA-based products surfacing as this method gradually becomes mature and efficient. FPGA could offer an attractive alternative for low-volume commercial products because it has a lower *nonrecurring engineering (NRE)* cost. It also has a shorter time to market and can be reprogrammed in the field to fix bugs and so on. However, compared to ASIC, its unit cost can be much higher. Thus, for large-volume products, the ASIC approach is the better

choice. Furthermore, due to its structure, FPGA performance is often inferior to that of ASIC.

Standard cell methodology is a method of designing *application-specific integrated circuits (ASICs)* with mostly digital content. A standard cell is group of transistor and interconnect structures that provides a Boolean logic function (e.g., AND, OR, XOR, XNOR, or inverter) or a storage function (flip-flop or latch). The cell's Boolean logic function is called its *logical view.* Its functional behavior is captured in the form of a truth table or Boolean algebraic equation (for combinational logic) or a state transition table (for sequential logic). From a manufacturing perspective, the layout of the standard cell (an abstract drawing of polygons) is the critical view. Layout is organized into base layers, which correspond to the structures of the transistor devices, and interconnect layers (metal layers), which join the terminals of the transistor formations. In design practice, the layout view is the lowest level of design abstraction.

The invention of logic synthesis and place and route tools has enabled the standard cell design style or ASIC approach. In this approach, the standard cells and other preassembled macro cells are grouped together to form an ASIC library. The chip functions are achieved by the cells in the library and through logic synthesis and physical place and route. In this approach, as contrasted to FPGA and platform ASIC, a mask is required for every layer, including the base and metal layers. The NRE cost associated with ASIC is often high due to the design, verification, implementation, and mask cost. However, for very large volume, high NRE costs could be offset by relatively low manufacturing costs.

Standard-cell ASIC methodology together with semiconductor process advances are the two major factors that have enabled ASIC chips to scale from simple, single-function ICs of several thousand gates to complex SoC devices of many million gates.

A platform-structured ASIC approach falls between an ASIC and a FPGA. It is an ASIC approach based on a preassembled *platform.* Inside the various platforms that a vender offers, certain special functions are already predesigned and verified. Random logic functions can be achieved by programming the metal layers in certain areas reserved for that purpose. Users can select a desired platform based on their needs. The main advantage of this platform-based ASIC is that the platform is already preassembled, which saves the mask cost of base layers. The only expense is the metal-programmable layers. Another benefit is that the verification costs can be significantly lower than those of an ASIC because the major functions on the platform might be preverified. As for performance, a platform ASIC often cannot match an ASIC; the design is not fully optimized as in the case of

an ASIC. However, performance should be significantly better than that of an FPGA. In summary, the platform ASIC trades the high performance of an ASIC with shorter time to market and lower development cost.

The platform ASIC approach is gaining momentum due to its relatively lower NRE cost as compared to an ASIC. But for very large volume products, its unit cost could be higher than that of ASIC.

7. WHAT IS STANDARD CELL-BASED, APPLICATION-SPECIFIC INTEGRATED CIRCUIT (ASIC) DESIGN METHODOLOGY?

Standard cell methodology is a chip design approach that is based on preassembled library cells. The standard cells and macros, such as memories, I/Os, special-function cells, phase lock loops (PLLs), and so on, associated with this library are already designed, laid out, and verified in a predetermined process node. These cells are completely characterized and logical, timing, physical, and electrical models are already created and properly packed in the library. After a design is created in register transfer level (RTL) format, it can be mapped into those preassembled cells through a logic synthesis process by sophisticated synthesis algorithms. The resultant netlist from this logic synthesis step is then fed into a physical implementation process, which includes the place and route steps.

Logic synthesis is the process of transforming the chip's RTL description into a technology-dependent gate netlist by using the library's logical view. In contrast to RTL description, which only contains functional information, the gate netlist is the standard cell representation of the design at the component level. It is comprised of gate instances and the port connectivity among these instances. The primary requirement for the task of logic synthesis is ensuring the mathematical equivalency between the synthesized gate netlist and the original RTL description.

The process of *place* is the first step in creating the chip in a physical domain. It determinates the physical locations of each individual cell in the netlist based on design constraints. Placement is a complicated process that is very algorithm intensive and time-consuming. The quality of the placement work has a preeminent impact on the chip's performance. The following *route* process is also critical. It creates the physical wire connections for the signal and power nets that are defined in the logic connectivity of the netlist. It is a very complicated process whose goals include meeting the design speed target, minimizing the total wire length, and avoiding the design rule violations.

After the place and route steps, the resultant physical entity is checked against various rules, such as the process manufacturing rules (foundry design rules) and design integrity and reliability criteria. This physical entity is also checked logically to ensure that it matches the design intention defined in the original RTL code. After these rigorous checks, the final layout is sent to the mask shop for the creation of photomasks. This is called *tapeout,* the final step in this standard cell-based ASIC (application-specific integrated circuit) design approach.

Currently, the standard cell-based ASIC approach is the main design methodology for commercial products, especially for large digitally dominated designs. A majority of SoC projects are carried out with this implementation approach.

8. WHAT IS THE SYSTEM-ON-CHIP (SOC) APPROACH?

SoC or *system on chip* is the design approach of integrating the components of an electronic system into a single chip. In the past, chips could only perform dedicated simple functions, such as simple logic operations, decoding/encoding operations, analog-to-digital conversion, digital-to-analog conversion, and so on. As time went by, more and more functions were integrated into a single chip. This integration trend is so significant that it has reached the point where a single chip can perform the functions of an entire electronic system, such as an MPEG decoder, a network router, or a cellular phone. As a result, a colorful name was created for such chips: system on chip (SoC). SoC designs often consume less power, cost less, and are more reliable than the multichip systems that they are designed to replace. Furthermore, assembly cost is reduced due to the fact that there are fewer packages in the system.

The key to the SoC approach is integration. By integrating increasingly more preassembled and verified blocks, which have dedicated functions, into one chip, a sophisticated system is created in a timely and economical fashion. Figure 1.3 is a block diagram of a SoC that shows the various blocks on a chip. As seen in the figure, integrating predesigned and verified blocks into a large chip is the essence of SoC approach.

A typical SoC chip has one or more microprocessors or microcontrollers on board, the brain of the SoC chip. The on-chip processor (e.g., an RISC controller) coordinates the activities inside the chip. In some cases, a dedicated DSP engine, which targets algorithm-intensive signal processing tasks, may also be found on a SoC chip. Having a large number of memory blocks is another characteristic of a SoC chip. These mem-

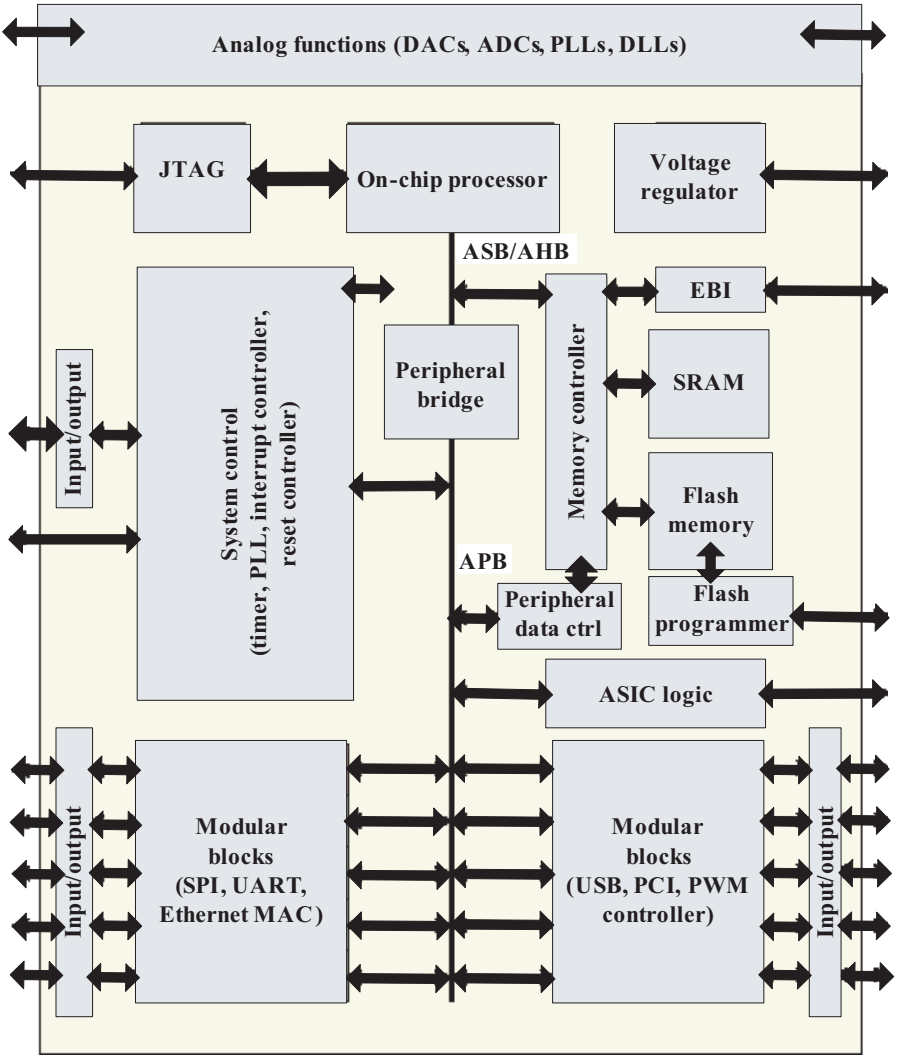

Figure 1.3. A SoC block diagram.

ories (ROM, RAM, EEPROM, and Flash) support the SoC's software functions. Another indispensable component of a SoC chip is the timing source, which includes an oscillator and phase lock loop (PLL). It is almost always true that one or more PLLs are found on any SoC chip since most SoC designs are based on synchronous design principle, and clocks are the key design feature.

A SoC needs external interfaces, such as industry standard USB, Firewire, Ethernet, and UART, to communicate with the outside world. A direct memory access (DMA) controller can be used to route data directly between the external interfaces and memories, bypassing the on-chip processor and thereby increasing the data throughput.

If a SoC is designed to interface with devices that have direct contact with human activities, some analog components, such as ADC or DAC, are essential. In some cases, on-chip voltage regulators and power management circuits can be found in a SoC as well.

To tie the components of a SoC together, an on-chip bus architecture is required for internal data transferring. This is either a proprietary bus or an industry-standard bus such as the AMBA bus from ARM. *Network on a chip (NoC)* is a new approach to SoC design. In an NoC system, modules such as processor cores, memories, and specialized IP blocks exchange data using a network as a public-transportation subsystem. The network is constructed from multiple point-to-point data links interconnected by switches such that messages are relayed from any source module to any destination module over several links by making routing decisions at the switches.

The NoC approach brings a networking solution to on-chip communication and provides notable improvements over conventional bus systems. From the viewpoint of physical design, the on-chip interconnect dominates both the dynamic power dissipation and performance of deep submicron CMOS technologies. If a signal is required across the chip, it may require multiple clock cycles when propagated in wires. A NoC link, on the other hand, can reduce the complexity of designing long interconnecting wires to achieve predictable speed, low noise, and high reliability due to its regular, well-controlled structure. From the viewpoint of system design and with the advent of multicore processor systems, a network is a natural architectural choice. A NoC can provide separation between the tasks of computation and communication, support modularity, and IP reuse via standard interfaces, efficiently handle synchronization issues, and serve as a platform for system test.

Just as the major hardware blocks are critical, so is the software of a SoC. The software controls the microcontroller, microprocessor, and DSP cores; the peripherals, and the interfaces to achieve various system functions.

One indispensable step in SoC development is *emulation.* Emulation is the process of using one system to perform the tasks in exactly the same way as another system, perhaps at a slower speed. Before a SoC device is sent out to fabrication, it must be verified by emulation for behavior analysis and making predications. During emulation, the SoC hardware is mapped onto an emulation platform based on a FPGA (or the likes) that

mimics the behavior of the SoC. The software modules are loaded into the memory of the emulation platform. Once programmed, the emulation platform enables both the testing and the debugging of the SoC hardware and the software.

In summary, the SoC approach is primarily focused on the integration of predesigned, preverified blocks, not on the design of individual components. In other words, the keyword is *integration,* not design.

9. WHAT ARE THE DRIVING FORCES BEHIND THE SOC TREND?

One of the major driving forces behind the SoC trend is cost. Integrating more functions into a single chip can reduce the chip count of a system and thus shrink the package and board cost. It could potentially lower the overall system cost and make the product more competitive. In today's consumer electronic market and in others, better price always provides advantage of gaining market share. During the past decade (from the late 1990s) or so, the SoC approach has been proven to be one of the most effective ways of reducing the cost of electronic devices.

The other forces behind this trend include pursuing higher chip performance or higher operating frequency. This is owing to the fact that SoC can eliminate interchip communication and shorten the distances among the on-chip components, which positively enhances the chip speed. In some cases, the demand for overall lower system power usage is also a factor for choosing the SoC approach. And, portability is another advantage of the SoC method. When a system is migrated from an old process to a new one, SoC can greatly reduce the workload compared to the transfer of several chips.

Overall, SoC chip implementation has enabled many technology innovations to reach the consumer in shorter and shorter time frames.

10. WHAT ARE THE MAJOR TASKS IN DEVELOPING A SOC CHIP FROM CONCEPT TO SILICON?

The process of developing a SoC chip from concept to silicon is divided into the following four tasks: design, verification, implementation, and software development.

Design often starts with marketing research and product definition and is followed by system design. It ends with RTL coding.

Verification is a means of ensuring that the chip can perform faithfully in functionality, according to its design specifications. It includes verification

at the system, RTL, and gate levels, and sometimes even at the transistor level. This bug-finding struggle continues until the chip is ready to ramp into production.

Implementation is the process of actually creating the hardware, which results in an entity that one can see and feel. It includes both the logical and physical implementations.

Software development is the process of programming the brain of the SoC (the on-chip processors), or arming the chip with intelligence.

11. WHAT ARE THE MAJOR COSTS OF DEVELOPING A CHIP?

There are two types of costs associated with the task of developing a VLSI chip: fixed costs and variable costs.

Fixed costs are also called *nonrecurring engineering (NRE)* costs. These refer to the one-time costs of researching, designing, and testing a new product. When developing a budget for a project and analyzing if a new product will be profitable, NRE must be considered. In the chip designing business, these costs include the engineering design cost (salaries, EDA tools licenses, CPU time, disk space, etc.) and mask (reticle) cost. Currently (2006), for 90 nm technology, the mask cost alone is in the neighborhood of one million dollars. This fixed cost is nonrefundable and is unrelated to product volume. To achieve profitability, then, the product must sell well enough to produce a return that covers at least the initial NRE and the materials and processing costs to make the initial material.

Variable cost is the cost of manufacturing, testing, and packaging the production chip. This cost is proportional to volume as every chip needs raw material for manufacture and tester's time for testing. This expense must be paid continually in order to maintain the product's manufacture.

The NRE or fixed cost represents a significant percentage of the overall cost of small-volume products. As volume grows, however, the fixed cost is gradually buried below the surface and the variable cost becomes dominant.

During the early phase of a project, the cost study is a very important, complex, and sensitive subject. Predicting or estimating the downstream expense and potential revenue-generating capability of a product with reasonable accuracy is not a trivial task. Sometimes, it is very hard to make a decision whether to kill or support a technically promising but market-unfriendly project as there might always be unknowns. The market is dynamic in nature; something unattractive today could be very popular tomorrow and vice versa.