appendix

E

COMMERCIAL DEVICES

In Chapter 3 we described the three main types of programmable logic devices (PLDs): simple PLDs, complex PLDs, and field-programmable gate arrays (FPGAs). This appendix describes some examples of commercial PLD products.

E.1 SIMPLE PLDS

Simple PLDs (SPLDs) include PLAs, PALs, and other similar types of devices. Some major manufacturers of SPLD products are listed in Table E.1. The first and second columns show the company name and some of the SPLD products it offers. Data sheets that describe each product can be obtained from the World Wide Web (WWW), using the locator given in the third column in the table.

E.1.1 THE 22V10 PAL DEVICE

PAL devices are among the most commonly used SPLDs. They are offered in a range of sizes and are identified by a part number of the form *NNXMM−S*. The digits *NN* specify the total number of input and output pins; the digits *MM* give the number of pins that can be used as outputs. The letter *X* gives additional information, such as whether the PAL contains flip-flops. The final digit, *S*, specifies the *speed grade*. This value represents the

Table E.1 Commercial SPLD Products.

Manufacturer	SPLD Products	WWW Locator
Altera	Classic	http//www.altera.com
Atmel	PAL	http//www.atmel.com
Lattice	ispGAL	http//www.latticesemi.com

propagation delay from an input pin on the PAL to an output pin, assuming that the flip-flop, if present, is bypassed.

An example of a commonly used PAL is the 22V10 [1], which is depicted in Figure E.1. There are 11 input pins that feed the AND plane, and an additional input that can also serve as a clock input. The OR gates are of variable size, ranging from 8 to 16 inputs. Each output pin has a tri-state buffer, which allows the pin to optionally be used as an input pin.

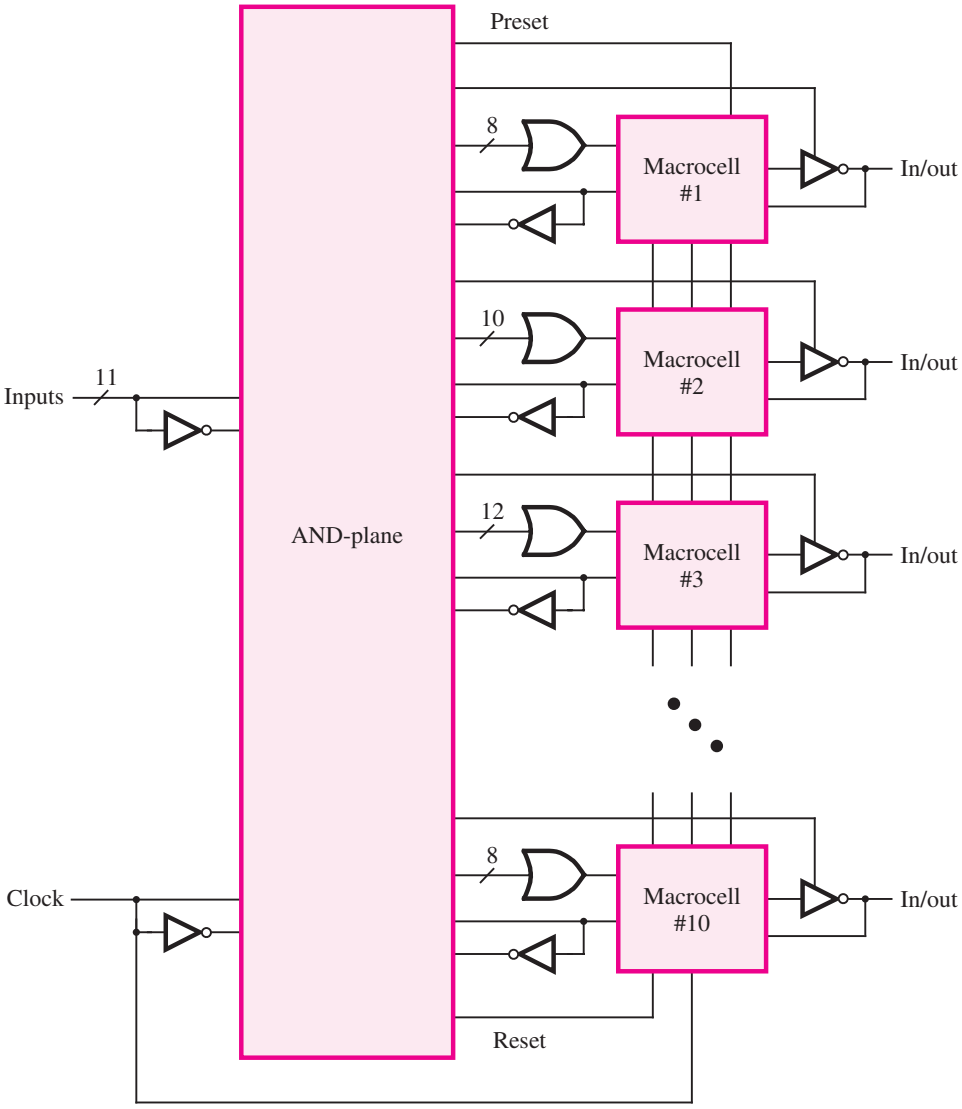

Figure E.1 The 22V10 PAL device.

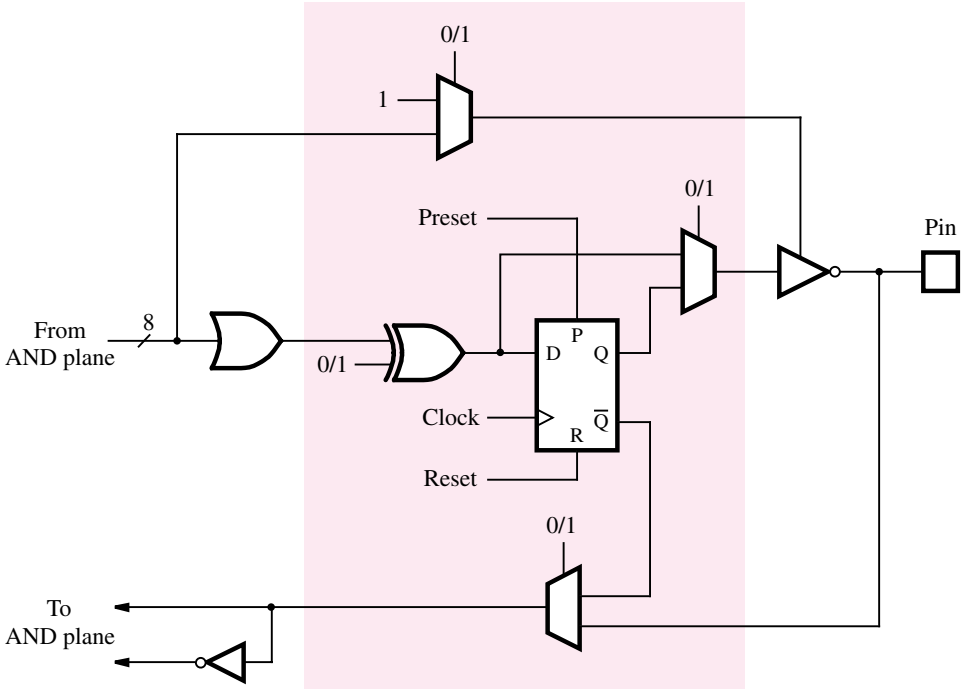

Figure E.2 The 22V10 macrocell.

We said in section 3.6.2 that the circuitry between an OR gate and an output in a PAL is usually called a macrocell. Figure E.2 shows one of the macrocells in the 22V10 PAL. It connects the OR gate shown to one input on an XOR gate, which feeds a D flip-flop. Since the other input to the XOR gate can be programmed to be 0 or 1, it can be used to complement the OR-gate output. A 2-to-1 multiplexer allows bypassing of the flip-flop, and the tri-state buffer can be either permanently enabled or connected to a product term from the AND plane. Either the $\overline{Q}$ output from the flip-flop or the output of the tri-state buffer can be connected to the AND plane. If the tri-state buffer is disabled, the corresponding pin can be used as an input.

E.2 COMPLEX PLDS

Names of Complex PLDs (CPLDs) manufacturers, some of the products they offer, and WWW locators are listed in Table E.2. An example of a widely used CPLD family, the Altera MAX 7000 [2], is described in the next section.

Table E.2 Commercial CPLD Products.

Manufacturer	CPLD Products	WWW Locator
Altera	MAX 3000, 7000, MAX II	http//www.altera.com
Atmel	ATF	http//www.atmel.com
Lattice	ispLSI, MachXO	http//www.latticesemi.com
Xilinx	XC9500, CoolRunner-II	http//www.xilinx.com

E.2.1 ALTERA MAX 7000

The MAX 7000 CPLD family includes chips that range in size from the 7032, which has 32 macrocells, to the 7512, which has 512 macrocells. There are two main variants of these chips, identified by the suffix *S*. If this letter is present in the chip name, as in 7128S, then the chip is in-system programmable. But if the suffix is absent, as in 7128, then the chip has to be programmed in a programming unit.

The overall structure of a MAX 7000 chip is illustrated in Figure E.3. There are four dedicated input pins; two of these can be used as global clock inputs, and one can be used

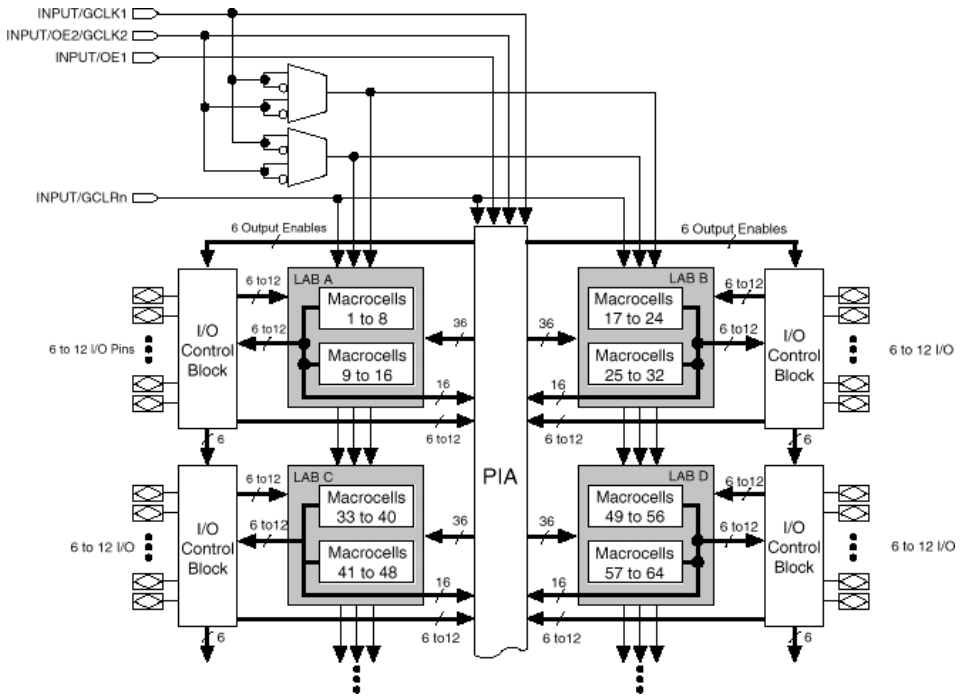

Figure E.3 MAX 7000 CPLD (courtesy of Altera).

as a global reset for all flip-flops. Each shaded box in the figure is called a *logic array block* (*LAB*), which contains 16 macrocells. Each LAB is connected to an *I/O control block*, which contains tri-state buffers that are connected to pins on the chip package; each of these pins can be used as an input or output pin. Each LAB is also connected to the *programmable interconnect array* (*PIA*). The PIA consists of a set of wires that span the entire device. All connections between macrocells are made using the PIA.

Figure E.4 shows the structure of a MAX 7000 macrocell. There are five product terms that can be connected through the *product term select matrix* to an OR gate. This OR gate can be configured to use only the product terms needed for the logic function being implemented in the macrocell. If more than five product terms are required, additional product terms can be "shared" from other macrocells, as described below. The OR gate is connected through an XOR gate to a flip-flop, which can be bypassed.

Figure E.5 shows how product terms can be shared between macrocells. The OR gate in a macrocell includes an extra input that can be connected to the output of the OR gate in the macrocell above it. This feature is called *parallel expanders* and is used for logic functions with up to 20 product terms. If even more product terms are needed, then a feature called *shared expanders* is used. As shown in the lower shaded box in Figure E.4, one of the product terms in a macrocell is inverted and fed back to the product term array. If the inputs to this product terms are used in their complemented form, then using DeMorgan's theorem, a sum term is produced. A shared expander can be used by any macrocell in the same LAB.

Each specific MAX 7000 device is available in a range of speed grades. These grades specify the propagation delay from an input pin through the PIA and a macrocell to an output pin. For example, the chip named 7128S-7 has a propagation delay of 7.5 ns. If the logic function implemented uses parallel or shared expanders, the propagation delay is increased.

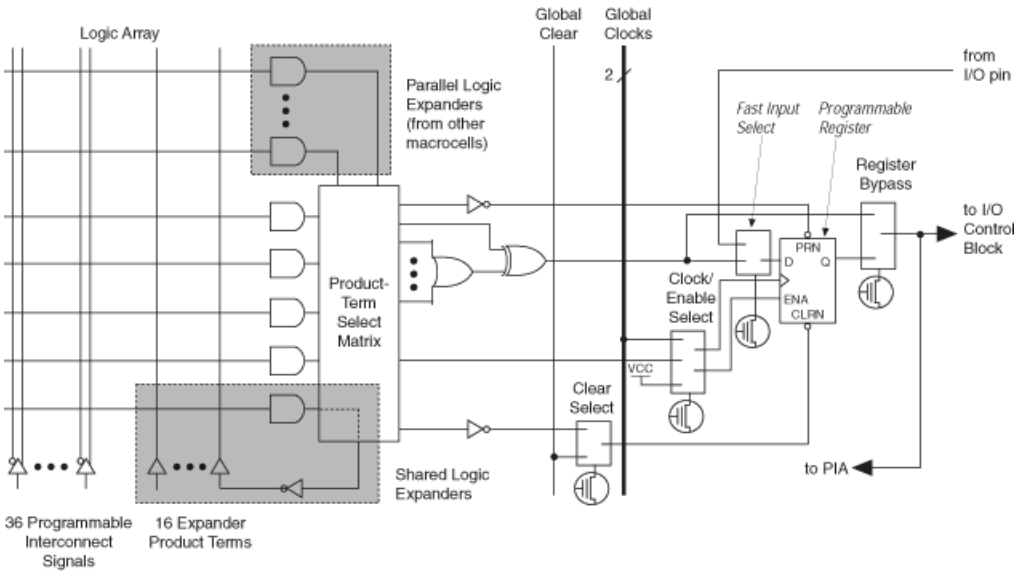

Figure E.4 MAX 7000 macrocell (courtesy of Altera).

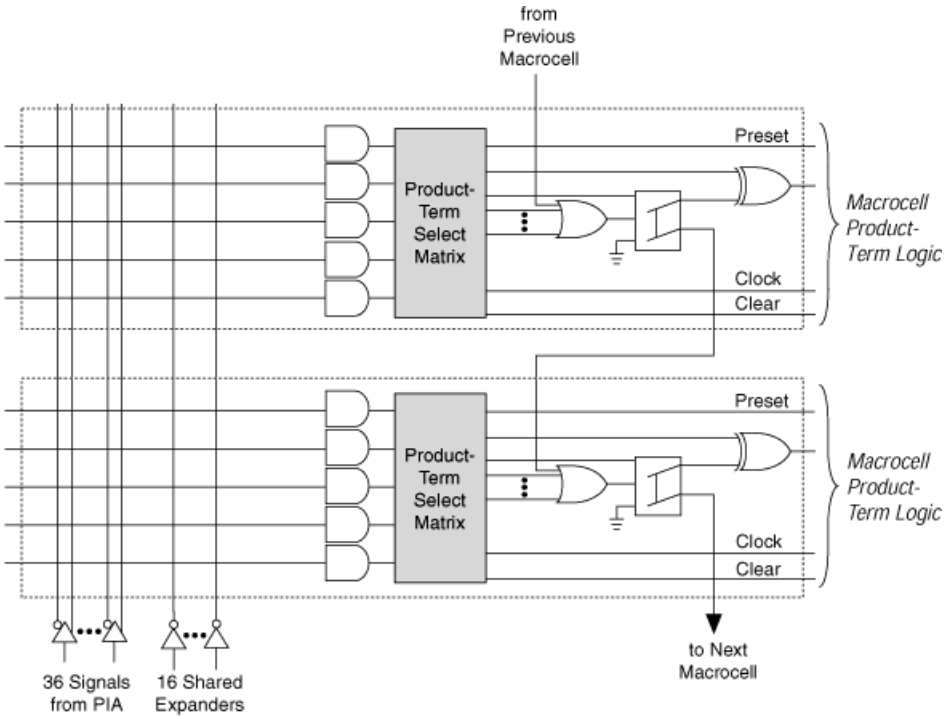

Figure E.5 Parallel Expanders (courtesy of Altera).

E.3 FIELD-PROGRAMMABLE GATE ARRAYS

Table E.3 lists the names of FPGA manufacturers, some of their products, and their WWW locators. This section describes examples of FPGAs produced by Altera and Xilinx.

E.3.1 ALTERA FLEX 10K

Figure E.6 shows the structure of the FLEX 10K chip [3]. It contains a collection of logic array blocks (LABs), where each LAB comprises eight *logic elements* based on lookup tables (LUTs). In addition to LABs, the chip also contains embedded array blocks (EABs), which are SRAM blocks that can be configured to provide memory blocks of various aspect ratios (see section 10.1.3). The LABs and EABs can be interconnected using the row and column *interconnect wires*. These wires also provide connections to the input and output pins on the chip package.

Figure E.7 shows the contents of a LAB. It has a number of inputs that are provided from the adjacent row interconnect wires to a set of local interconnect wires inside the

Table E.3 Commercial FPGA Products.

Manufacturer	FPGA Products	WWW Locator
Actel	MX, SX, eX	http//www.actel.com
Altera	Stratix (II/III), Cyclone (II/III)	http//www.altera.com
	FLEX 10K, APEX 20K	
Lattice	ECP2/M, SC	http//www.latticesemi.com
Xilinx	Virtex-(4/5), Virtex-II(Pro)	http//www.xilinx.com
	Spartan-3 (A/E), XC4000	

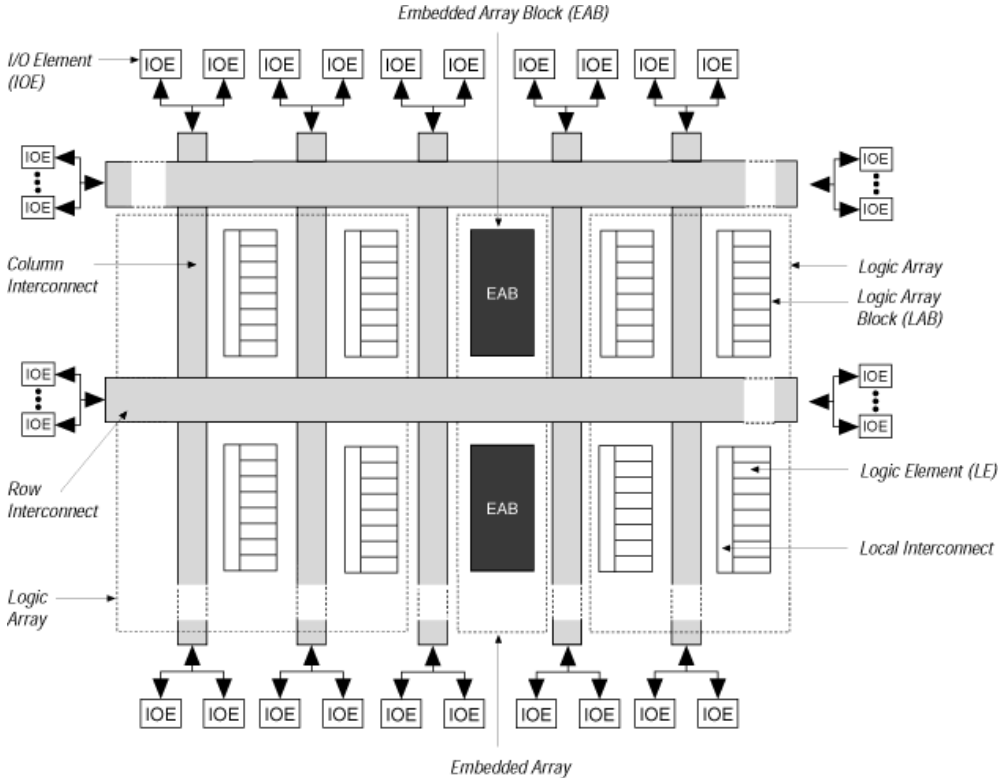

Figure E.6 FLEX 10K FPGA (courtesy of Altera).

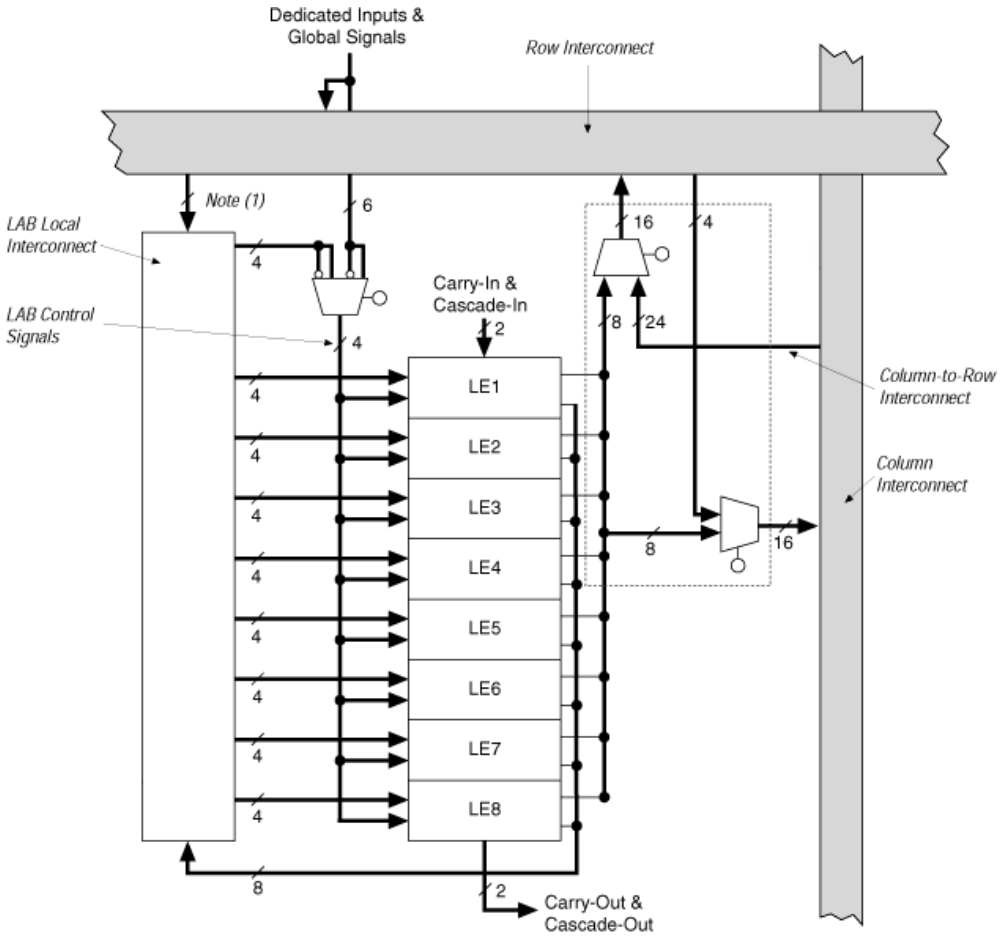

Figure E.7 FLEX 10K logic array block (courtesy of Altera).

LAB. These local wires are used to make connections to the inputs of the logic elements, and the logic element outputs also feed back to the local wires. Logic element outputs also connect to the adjacent row and column wires. The structure of a logic element is depicted in Figure E.8. The element has a four-input LUT and a flip-flop that can be bypassed. For implementation of arithmetic adders, the four-input LUT can be used to implement 2 three-input functions, namely, the sum and carry functions in a full-adder.

The structure of an EAB is depicted in Figure E.9. It contains 2048 SRAM cells, which can be used to provide memory blocks that have a range of aspect ratios: 256 $\times$ 8, 512 $\times$ 4, 1024 $\times$ 2, and 2048 $\times$1 bits. The address and data inputs to the memory block are provided from a set of local interconnect wires. These inputs, as well as a write enable for the memory block, can optionally be stored in flip-flops. Figure E.9 shows that the number of address and data inputs connected to the memory block varies depending on the aspect

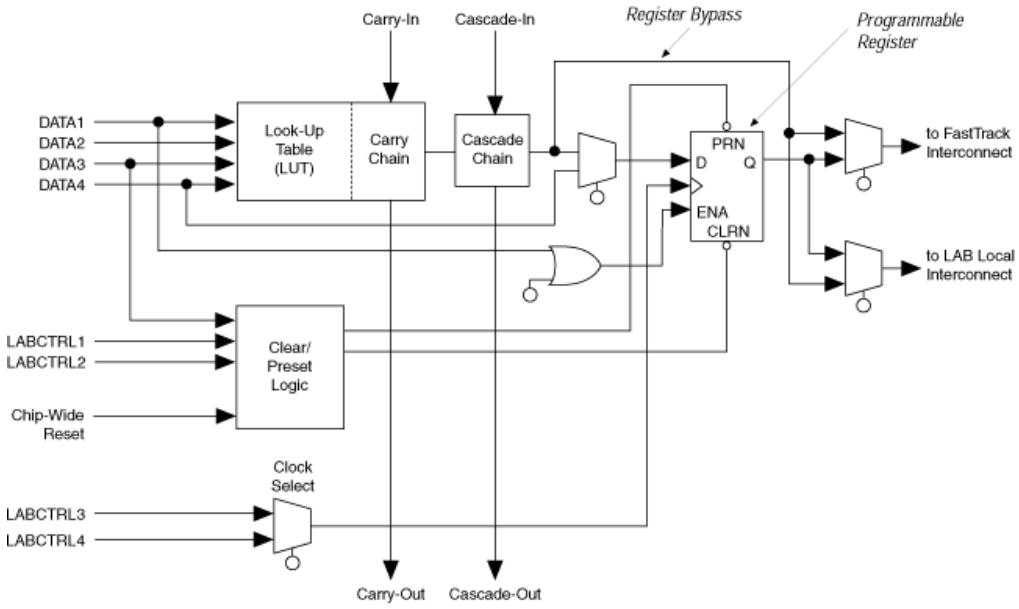

Figure E.8 FLEX 10K logic element (courtesy of Altera).

ratio being used. The data outputs can also optionally be stored in flip-flops. For large memory blocks it is possible to combine multiple EABs.

Configuration of EABs is done using predesigned modules, such as those in the LPM library. For example, the module named *lpm_ram_dq* can be used to specify an SRAM block, and *lpm_rom* can be used for a ROM block. These modules can be imported into a schematic or instantiated in code using a language such as VHDL. It is possible to specify initial data to be loaded into the memory block when the FPGA chip is programmed. This is done by creating a special type of file, called a *memory initialization file*, that is associated with the *lpm_ram_dq* or *lpm_rom* module. Complete details on using these modules can be found in the Quartus II documentation.

FLEX 10K chips are available in sizes ranging from the 10K10 to 10K250, which offer about 10,000 and 250,000 equivalent logic gates, respectively. Specific chips are available in various speeds, indicated using a suffix letter, such as *A*, as in 10K10A, and a speed grade, as in 10K10A-1. Unlike PALs and CPLDs, the speed grade for an FPGA does not specify an actual propagation delay in nanoseconds. Instead, it represents a relative speed within the device family. For instance, the 10K10-1 is a faster chip than the 10K10-2. The actual propagation delays in implemented circuits can be examined using a timing simulator CAD tool.

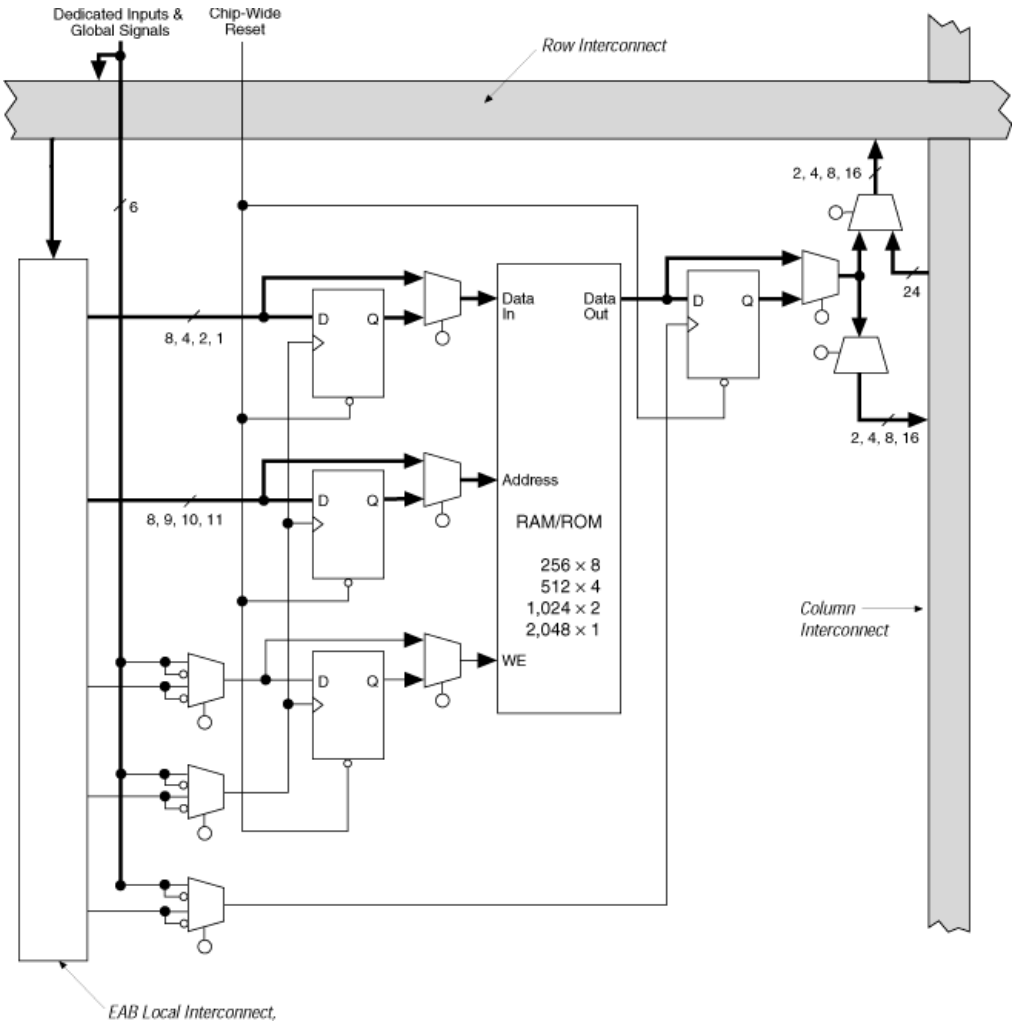

Figure E.9 Embedded array block (courtesy of Altera).

E.3.2 XILINX XC4000

The structure of a Xilinx XC4000 chip [4] is similar to the FPGA structure shown in Figure 3.35. It has a two-dimensional array of *configurable logic blocks* (*CLBs*) that can be interconnected using the vertical and horizontal routing channels. Chips range in size from the XC4002 to XC40250, which have about 2000 and 250,000 equivalent logic gates, respectively. As shown in Figure E.10, a CLB contains 2 four-input LUTs; hence it can implement any two logic functions of up to four variables. The output of each of these

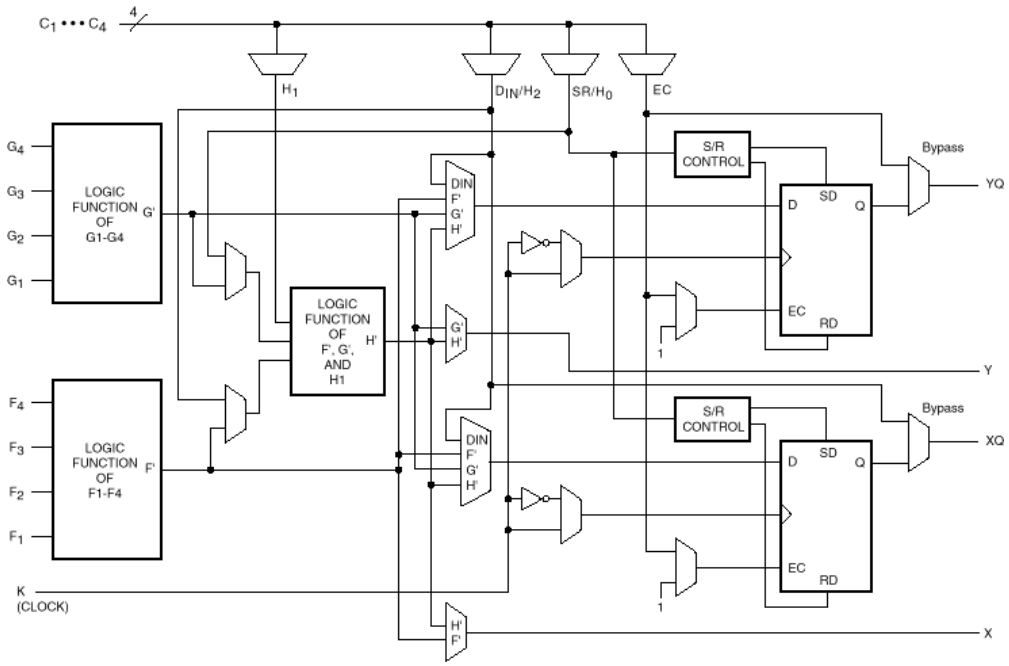

Figure E.10 XC4000 configurable logic block (courtesy of Xilinx).

LUTs can optionally be stored in a flip-flop. The CLB also contains a three-input LUT connected to the 2 four-input LUTs, which allows implementation of functions with five or more variables.

Similar to the logic elements in the FLEX 10K FPGAs described in section E.3.1, the CLB can be configured for efficient implementation of adder modules. In this mode each four-input LUT in the CLB implements both the sum and carry functions of a full-adder. Also, instead of implementing logic functions, the CLB can be used as a memory module. Each four-input LUT can serve as a 16 $\times$ 1 memory block, or both four-LUTs can be combined into a 32 $\times$ 1 memory block. Multiple CLBs can be combined to form larger memory blocks.

The CLBs are interconnected using the wires in the routing channels. Wires of various lengths are provided, from wires that span a single CLB to wires that span the entire device. The number of wires in a routing channel varies for each specific chip.

E.3.3 ALTERA APEX 20K

The Altera APEX 20K [5] family is the next generation product following the FLEX 10K. The logic element (LE), which is an optimized version of the one depicted in Figure E.8, contains a four-input LUT and a flip-flop. Chips range in sizes from 1200 to 51,840 LEs.

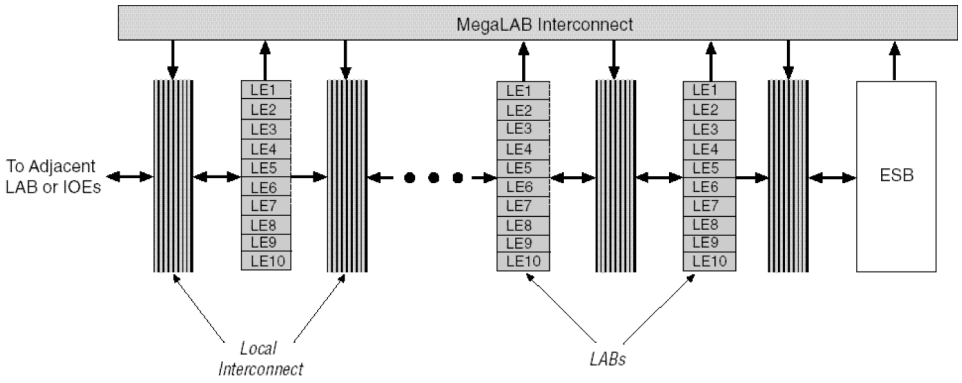

Figure E.11 APEX 20K MegaLAB (courtesy of Altera).

Each APEX device contains logic elements (LUTs), memory blocks, and IO cells. The LEs are arranged into LABs similar, to the structure depicted in Figure E.7, with ten LEs per LAB. The LABs are further grouped into MegaLABs, with up to 24 LABs in a MegaLAB. As shown in Figure E.11, the MegaLAB contains wires to interconnect the LABs, and it also contains a memory block, called the embedded system block (ESB). Similar to the EAB shown in Figure E.9, the ESB supports memory blocks with various aspect ratios. An APEX device comprises either two or four columns of MegaLABs; the number of MegaLABs per column varies for each device.

E.3.4 ALTERA STRATIX

Stratix [6] is Altera's FPGA product that supersedes the APEX family. Figure E.12 shows the architecture of a Stratix device. Each chip comprises columns of resources of various types. The LAB columns house logic elements arranged into LABs that have ten LEs per LAB. Each LE contains a four-input LUT and a register, and can be configured in a variety of modes, including a fast arithmetic mode. There are a number of types of wiring resources in a Stratix chip. Connections within a LAB are made using fast local resources, such as a carry chain that runs downward in each column. For connections from one LAB to other resources there exist short nearest-neighbor connections, wires that span four columns or rows, and longer wires.

In addition to LAB columns, Stratix devices contain three other types of columns. The M512 columns consist of memory blocks with 512 bits each, and the M4K columns contain larger memory blocks with 4K bits per block. Each of the M512 and M4K blocks support implementations of memories with various aspect ratios. Stratix devices also include very large memory blocks called MegaRAMs, each of which contains 512K bits of memory.

Finally, there are columns that comprise Digital Signal Processing (DSP) blocks. Each of these blocks includes hardware multiplier and adder circuits that allow fast multiplication

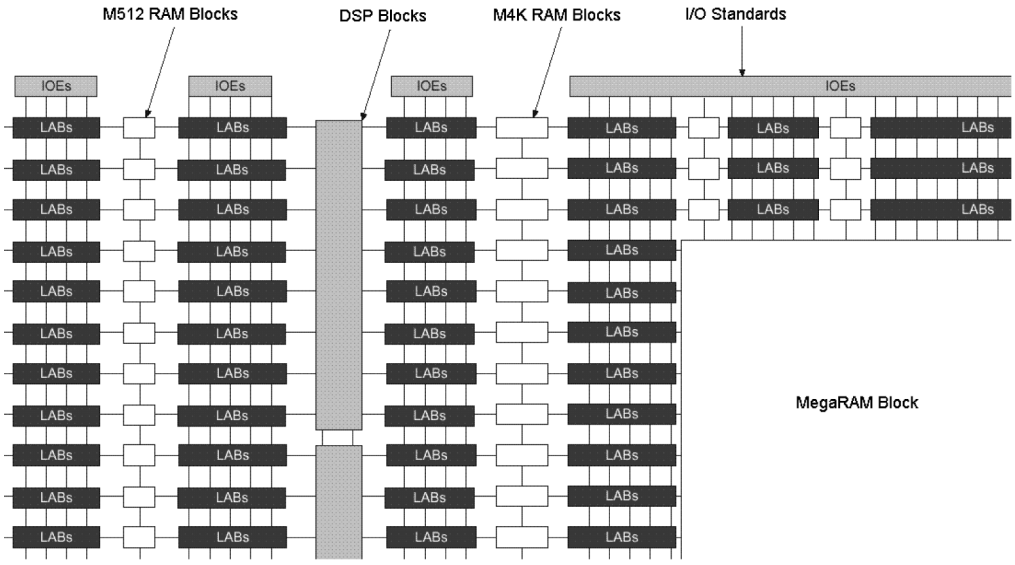

Figure E.12 Stratix LAB, DSP, and memory blocks (courtesy of Altera).

and accumulation (summing) of data. These blocks provide efficient implementation of the types of circuits used in digital signal processing applications.

Stratix chips are available in sizes from 10,570 to 79,040 logic elements and over seven Mbits of memory.

E.3.5 ALTERA CYCLONE, CYCLONE II, AND CYCLONE III

Cyclone [7] FPGAs are based on the Stratix architecture, but are intended for low-cost applications. There are three generations of these devices, called Cyclone, Cyclone II, and Cyclone III. A Cyclone chip has the same basic structure as that shown in Figure E.12, with a four-input LUT logic element that has dedicated arithmetic circuitry and a programmable flip-flop. The types of memory blocks provided in these devices are M4K in Cyclone and Cyclone II, and M9K in Cyclone III. Cyclone II and III devices also include DSP blocks. Cyclone devices range in size from 2910 to 119,088 logic elements and 4 Mbits of memory.

An example of a commercial product that includes a Cyclone II device is the DE2 Development and Education board from Altera, which is described in Appendix D.

E.3.6 ALTERA STRATIX II AND STRATIX III

Stratix II [8] and Stratix III [9] FPGAs are the successor to the Stratix family. They offer device sizes from 15,600 to 338,000 logic elements and up to 16.7 Mbits of memory. Stratix II and Stratix III contain a more complex logic element than other FPGAs, called the Adaptive Logic Module (ALM). As shown in Figure E.13, the ALM comprises a combina-

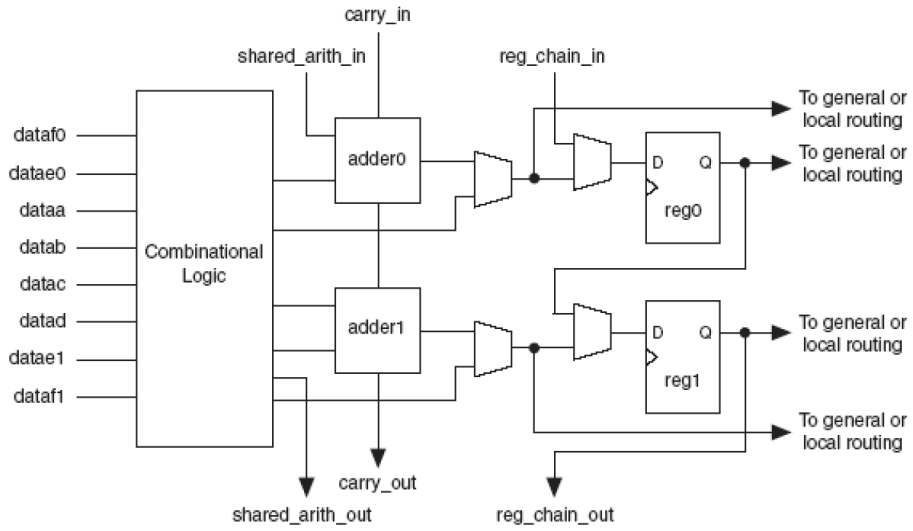

Figure E.13 The Stratix II Adaptive Logic Module.

tional logic circuit and two programmable flip-flops. The combinational logic circuit can be programmed as either one or two LUTs; it can implement a single logic function of up to seven inputs, or two functions of various sizes. Figure E.14 shows a few of the possible configurations of the ALM, such as realizing two four-input LUTs, a four-input LUT plus a five-input LUT, and so on. The Stratix III ALM has the option of being configured as a small memory block in addition to its use as a logic element [9].

E.3.7 XILINX VIRTEX

The Xilinx Virtex [10] FPGAs are the next generation family following the XC4000. As indicated in Figure E.15, each Virtex chip comprises logic resources called CLBs, and memory resources called Block RAMs (BRAMs). The CLB is an enhanced version of the XC4000 CLB shown in Figure E.10. As indicated in Figure E.16, the Virtex CLB is divided into two halves; each half is called a *slice*. Each slice contains two four-input LUTs, two registers, and dedicated arithmetic (carry chain) logic.

The BRAM blocks contain 4K bits of memory, and can be configured to support aspect ratios from 4096×1 to 256×16. The CLB and BRAM blocks can be interconnected by wires that span a single CLB, or longer distances. Virtex devices are available in sizes from 256 to 46,592 CLB slices.

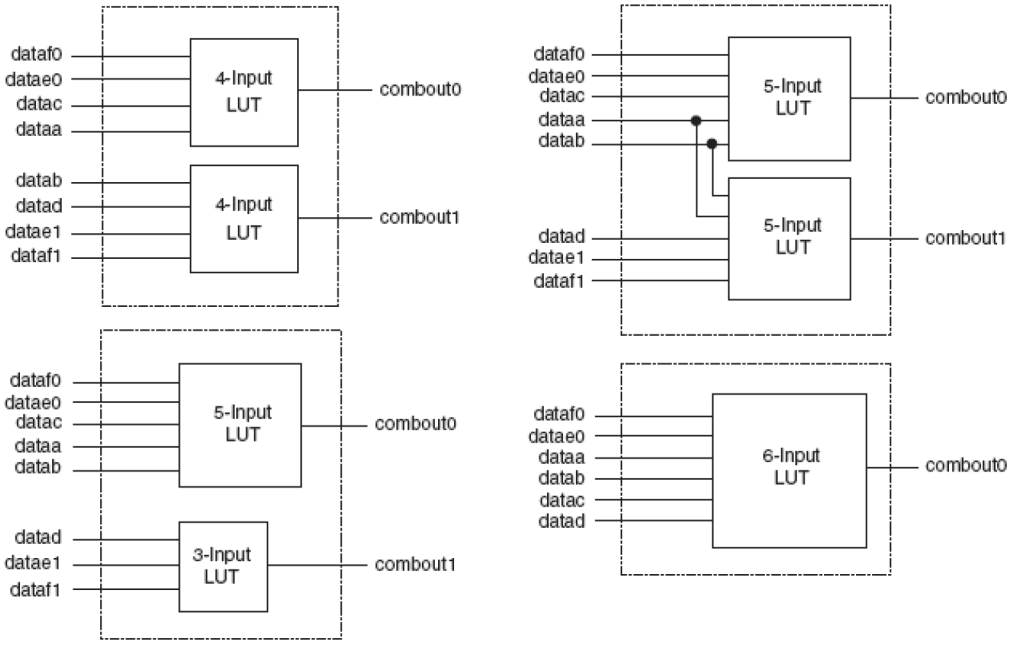

Figure E.14 Some of the modes of the Stratix II ALM.

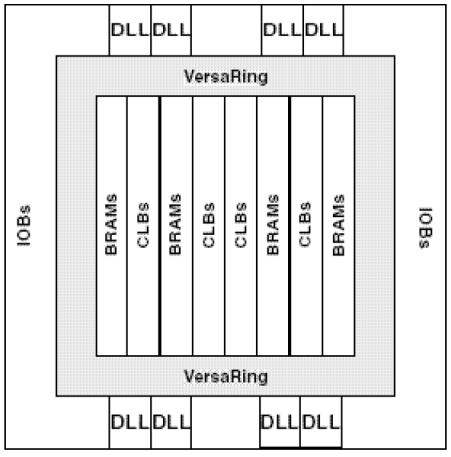

Figure E.15 Virtex FPGA (courtesy of Xilinx).

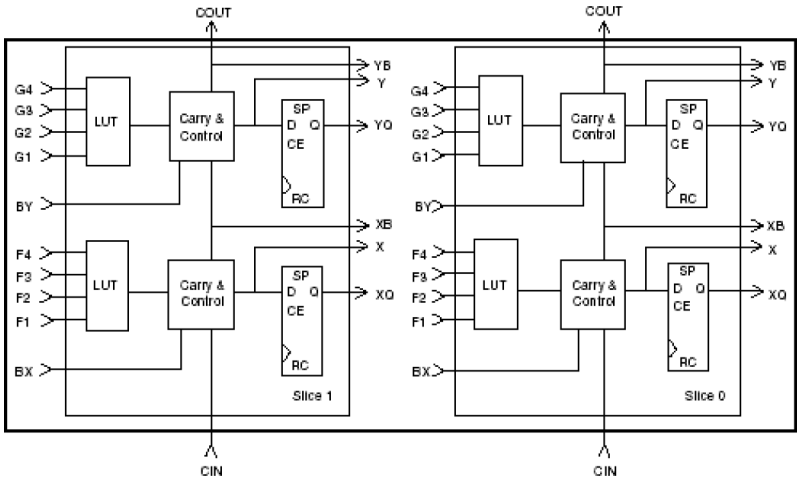

Figure E.16 Virtex logic block (courtesy of Xilinx).

E.3.8 XILINX VIRTEX-II, VIRTEX-II PRO, VIRTEX-4, AND VIRTEX-5

The Xilinx Virtex-II [11] Virtex-II Pro, Virtex-4, and Virtex-5 FPGAs [12] are the successors to the Virtex family. They are offered in sizes from 3168 to 331,776 logic elements and with more than 10.4 Mbits of memory. The logic elements are arranged into slices similar to the Virtex FPGAs (see Figure E.16), with four slices in a CLB. Some members of these families include one or more microprocessor cores within the chip, and have additional advanced features that are not present in Virtex-II.

E.3.9 XILINX SPARTAN-3

The Xilinx Spartan-3 [13] FPGAs are a low-cost version of the Virtex-II architecture. Similar to Virtex-II, the logic elements are arranged into CLBs that each have four slices, but not all slices have the same feature-set as in Virtex-II. Spartan-3 chips are available in sizes from 1728 to 74,880 logic elements and more than 1.8 Mbits of memory.

E.4 TRANSISTOR-TRANSISTOR LOGIC

Before the emergence of CMOS, the dominant technology was *transistor-transistor logic*, commonly referred to as *TTL*. Most digital systems built in the 1970s and 1980s were based on this technology. TTL circuits are available in relatively small sizes, known as small-scale integration (SSI) and medium-scale integration (MSI), as explained in section 3.5. A typical SSI chip contains just a few logic gates, with their inputs and outputs available on

the pins of the package. An MSI chip may comprise a somewhat larger circuit, such as a four-bit arithmetic and logic unit (ALU).

TTL technology is not as suitable for large-scale integration as CMOS technology, which has led to TTL's demise. However, its impact was so large that some aspects are still important today. In this section we consider these aspects.

Voltage Levels

TTL circuits use a 5-volt power supply. Any voltage in the range 0 to 0.8 V is interpreted as a logic 0 when applied to an input pin. A voltage in the range 2 to 5 volts is interpreted as a logic 1. Using the terminology from section 3.8, $V_{IL} = 0.8$ V and $V_{IH} = 2$ V. The maximum output voltage produced for logic 0 is $V_{OL} = 0.4$ V, and the minimum voltage produced for logic 1 is $V_{OH} = 2.4$ V. These parameters lead to the noise margins $NM_L = NM_H = 0.4$ V. Typical output voltages generated by a TTL circuit are 0.2 V for logic 0 and 3.6 V for logic 1.

When a new digital circuit is designed, it is often intended for use in an existing digital system. If different technologies are used to implement different parts of a system, it is essential to ensure that compatible voltage levels are used for signals in the interfaces between the different parts. While CMOS voltage levels are normally different from TTL levels, some CMOS chips, such as PLDs, can be configured to use TTL-compatible voltage levels on their input and output pins.

Input Connections

In CMOS circuits all inputs to a gate must always be driven to either logic value 0 or 1. Otherwise, the gate's output will have an unknown (usually tri-state) value. In the case of TTL circuits, an unconnected input behaves as if it were connected to a constant 1.

E.4.1 TTL CIRCUIT FAMILIES

TTL circuits are available in several designs that have different propagation speeds and power consumption. They have the same functional characteristics, defined by the specifications for the type of circuits known as the 7400 series, which is introduced in section 3.5. Actually, the 7400 label denotes a chip that comprises 4 two-input NAND gates. Other chips that contain different logic elements have the same prefix 74, but are identified by additional digits. For example, 7421 denotes a chip that comprises 2 four-input AND gates. Table E.4 presents the propagation delay and power dissipation characteristics of the various TTL families.

Standard TTL is based on the original specifications, and it was the first type of such circuits introduced in the 1960s. Subsequent versions provided various improvements. Faster circuits were developed, trading off increased power consumption for shorter propagation delays. Conversely, low-power circuits were developed, at the cost of longer propagation delays. Table E.4 gives the typical values that can be expected under normal operating conditions.

The maximum fan-out in TTL circuits is 10 in most cases, but it can be as high as 20 for the low-power types. The fan-in is determined by the number of inputs provided on a given chip.

Table E.4 TTL logic families.

Name	Designation	Propagation Delay (ns)	Power Dissipation (mW)
Standard	7400	9	10
Low power	74L00	33	1
High speed	74H00	6	22
Schottky	74S00	3	20
Low-power Schottky	74LS00	9	2
Advanced Schottky	74AS00	1.5	20
Advanced low-power Schottky	74ALS00	4	1
Fast	74F00	3	4

TTL gates can have different output configurations. In addition to the normal output configuration, there exist gates that have tri-state outputs or open-collector outputs. The purpose of a tri-state output is discussed in section 3.8.8. Gates with open-collector outputs are used when it is desirable to connect the outputs of two or more gates together directly. These gates are not damaged by such a connection, because each gate either drives the output to 0 or does not affect it at all. Connecting the outputs of several open-collector gates through a pull-up resistor to +5 V results in a circuit where the voltage at the output point is equal to +5 V if none of the gates produces an output of 0 and is equal to 0 if one or more gates produce the output of 0. A similar approach can be used with CMOS technology, resulting in open-drain gates.

We have not pursued TTL technology in any detail because of its diminished importance in today's design environment. An interested reader may consult numerous books that provide a detailed explanation. A particularly thorough reference is [14].

REFERENCES

1. Lattice Semiconductor, Simple PLDs Data Sheets, http://www.latticesemi.com
2. Altera Corporation, MAX 7000 CPLD Data Sheets, http://www.altera.com
3. Altera Corporation, FLEX 10K Data Sheets, http://www.altera.com
4. Xilinx Corporation, XC4000 FPGA Data Sheets, http://www.xilinx.com
5. Altera Corporation, APEX 20K Data Sheets, http://www.altera.com
6. Altera Corporation, Stratix FPGA Data Sheets, http://www.altera.com
7. Altera Corporation, Cyclone, Cyclone II, and Cyclone III FPGA Data Sheets, http://www.altera.com
8. Altera Corporation, Stratix II FPGA Data Sheets, http://www.altera.com

9. Altera Corporation, Stratix III FPGA Data Sheets, http://www.altera.com
10. Xilinx Corporation, Virtex FPGA Data Sheets, http://www.xilinx.com
11. Xilinx Corporation, Virtex-II FPGA Data Sheets, http://www.xilinx.com
12. Xilinx Corporation, Virtex-II Pro, Virtex-4/5 FPGA Data Sheets, http://www.xilinx.com
13. Xilinx Corporation, Spartan-3 FPGA Data Sheets, http://www.xilinx.com
14. A. S. Sedra and K. C. Smith, *Microelectronic Circuits*, 5th ed. (Oxford University Press: New York, 2003).