Welcome to E-XFL.COM

Understanding Embedded - FPGAs (Field Programmable Gate Array)

Embedded - FPGAs, or Field Programmable Gate Arrays, are advanced integrated circuits that offer unparalleled flexibility and performance for digital systems. Unlike traditional fixed-function logic devices, FPGAs can be programmed and reprogrammed to execute a wide array of logical operations, enabling customized functionality tailored to specific applications. This reprogrammability allows developers to iterate designs quickly and implement complex functions without the need for custom hardware.

Applications of Embedded - FPGAs

The versatility of Embedded - FPGAs makes them indispensable in numerous fields. In telecommunications,

Details

Product Status	Obsolete
Number of LABs/CLBs	1296
Number of Logic Elements/Cells	3078
Total RAM Bits	41472
Number of I/O	129
Number of Gates	36000
Voltage - Supply	3V ~ 3.6V
Mounting Type	Surface Mount
Operating Temperature	-40°C ~ 100°C (TJ)
Package / Case	160-BQFP Exposed Pad
Supplier Device Package	160-PQFP (28x28)
Purchase URL	https://www.e-xfl.com/product-detail/xilinx/xc4036xl-1hq160i

Table 1: XC4000E and XC4000X Series Field Programmable Gate Arrays

Device	Logic Cells	Max Logic Gates (No RAM)	Max. RAM Bits (No Logic)	Typical Gate Range (Logic and RAM)*	CLB Matrix	Total CLBs	Number of Flip-Flops	Max. User I/O
XC4002XL	152	1,600	2,048	1,000 - 3,000	8 x 8	64	256	64
XC4003E	238	3,000	3,200	2,000 - 5,000	10 x 10	100	360	80
XC4005E/XL	466	5,000	6,272	3,000 - 9,000	14 x 14	196	616	112
XC4006E	608	6,000	8,192	4,000 - 12,000	16 x 16	256	768	128
XC4008E	770	8,000	10,368	6,000 - 15,000	18 x 18	324	936	144
XC4010E/XL	950	10,000	12,800	7,000 - 20,000	20 x 20	400	1,120	160
XC4013E/XL	1368	13,000	18,432	10,000 - 30,000	24 x 24	576	1,536	192
XC4020E/XL	1862	20,000	25,088	13,000 - 40,000	28 x 28	784	2,016	224
XC4025E	2432	25,000	32,768	15,000 - 45,000	32 x 32	1,024	2,560	256
XC4028EX/XL	2432	28,000	32,768	18,000 - 50,000	32 x 32	1,024	2,560	256
XC4036EX/XL	3078	36,000	41,472	22,000 - 65,000	36 x 36	1,296	3,168	288
XC4044XL	3800	44,000	51,200	27,000 - 80,000	40 x 40	1,600	3,840	320
XC4052XL	4598	52,000	61,952	33,000 - 100,000	44 x 44	1,936	4,576	352
XC4062XL	5472	62,000	73,728	40,000 - 130,000	48 x 48	2,304	5,376	384
XC4085XL	7448	85,000	100,352	55,000 - 180,000	56 x 56	3,136	7,168	448

* Max values of Typical Gate Range include 20-30% of CLBs used as RAM.

Note: *All functionality in low-voltage families is the same as in the corresponding 5-Volt family, except where numerical references are made to timing or power.*

Description

XC4000 Series devices are implemented with a regular, flexible, programmable architecture of Configurable Logic Blocks (CLBs), interconnected by a powerful hierarchy of versatile routing resources, and surrounded by a perimeter of programmable Input/Output Blocks (IOBs). They have generous routing resources to accommodate the most complex interconnect patterns.

The devices are customized by loading configuration data into internal memory cells. The FPGA can either actively read its configuration data from an external serial or byte-parallel PROM (master modes), or the configuration data can be written into the FPGA from an external device (slave and peripheral modes).

XC4000 Series FPGAs are supported by powerful and sophisticated software, covering every aspect of design from schematic or behavioral entry, floor planning, simulation, automatic block placement and routing of interconnects, to the creation, downloading, and readback of the configuration bit stream.

Because Xilinx FPGAs can be reprogrammed an unlimited number of times, they can be used in innovative designs where hardware is changed dynamically, or where hardware must be adapted to different user applications. FPGAs are ideal for shortening design and development cycles, and also offer a cost-effective solution for production rates well beyond 5,000 systems per month.

Taking Advantage of Re-configuration

FPGA devices can be re-configured to change logic function while resident in the system. This capability gives the system designer a new degree of freedom not available with any other type of logic.

Hardware can be changed as easily as software. Design updates or modifications are easy, and can be made to products already in the field. An FPGA can even be re-configured dynamically to perform different functions at different times.

Re-configurable logic can be used to implement system self-diagnostics, create systems capable of being re-configured for different environments or operations, or implement multi-purpose hardware for a given application. As an added benefit, using re-configurable FPGA devices simplifies hardware design and debugging and shortens product time-to-market.

Detailed Functional Description

XC4000 Series devices achieve high speed through advanced semiconductor technology and improved architecture. The XC4000E and XC4000X support system clock rates of up to 80 MHz and internal performance in excess of 150 MHz. Compared to older Xilinx FPGA families, XC4000 Series devices are more powerful. They offer on-chip edge-triggered and dual-port RAM, clock enables on I/O flip-flops, and wide-input decoders. They are more versatile in many applications, especially those involving RAM. Design cycles are faster due to a combination of increased routing resources and more sophisticated software.

Basic Building Blocks

Xilinx user-programmable gate arrays include two major configurable elements: configurable logic blocks (CLBs) and input/output blocks (IOBs).

- CLBs provide the functional elements for constructing the user's logic.
- IOBs provide the interface between the package pins and internal signal lines.

Three other types of circuits are also available:

- 3-State buffers (TBUFs) driving horizontal longlines are associated with each CLB.
- Wide edge decoders are available around the periphery of each device.
- An on-chip oscillator is provided.

Programmable interconnect resources provide routing paths to connect the inputs and outputs of these configurable elements to the appropriate networks.

The functionality of each circuit block is customized during configuration by programming internal static memory cells. The values stored in these memory cells determine the logic functions and interconnections implemented in the FPGA. Each of these available circuits is described in this section.

Configurable Logic Blocks (CLBs)

Configurable Logic Blocks implement most of the logic in an FPGA. The principal CLB elements are shown in Figure 1. Two 4-input function generators (F and G) offer unrestricted versatility.   Most combinatorial logic functions need four or fewer inputs. However, a third function generator (H) is provided. The H function generator has three inputs. Either zero, one, or two of these inputs can be the outputs of F and G; the other input(s) are from outside the CLB. The CLB can, therefore, implement certain functions of up to nine variables, like parity check or expandable-identity comparison of two sets of four inputs.

Each CLB contains two storage elements that can be used to store the function generator outputs. However, the storage elements and function generators can also be used independently. These storage elements can be configured as flip-flops in both XC4000E and XC4000X devices; in the XC4000X they can optionally be configured as latches. DIN can be used as a direct input to either of the two storage elements. H1 can drive the other through the H function generator. Function generator outputs can also drive two outputs independent of the storage element outputs. This versatility increases logic capacity and simplifies routing.

Thirteen CLB inputs and four CLB outputs provide access to the function generators and storage elements. These inputs and outputs connect to the programmable interconnect resources outside the block.

Function Generators

Four independent inputs are provided to each of two function generators (F1 - F4 and G1 - G4). These function generators, with outputs labeled F' and G', are each capable of implementing any arbitrarily defined Boolean function of four inputs. The function generators are implemented as memory look-up tables. The propagation delay is therefore independent of the function implemented.

A third function generator, labeled H', can implement any Boolean function of its three inputs. Two of these inputs can optionally be the F' and G' functional generator outputs. Alternatively, one or both of these inputs can come from outside the CLB (H2, H0). The third input must come from outside the block (H1).

Signals from the function generators can exit the CLB on two outputs. F' or H' can be connected to the X output. G' or H' can be connected to the Y output.

A CLB can be used to implement any of the following functions:

- any function of up to four variables, plus any second function of up to four unrelated variables, plus any third function of up to three unrelated variables[1]
- any single function of five variables
- any function of four variables together with some functions of six variables
- some functions of up to nine variables.

Implementing wide functions in a single block reduces both the number of blocks required and the delay in the signal path, achieving both increased capacity and speed.

The versatility of the CLB function generators significantly improves system speed. In addition, the design-software tools can deal with each function generator independently. This flexibility improves cell usage.

---

1. When three separate functions are generated, one of the function outputs must be captured in a flip-flop internal to the CLB. Only two unregistered function generator outputs are available from the CLB.

Supported CLB memory configurations and timing modes for single- and dual-port modes are shown in Table 3.

XC4000 Series devices are the first programmable logic devices with edge-triggered (synchronous) and dual-port RAM accessible to the user. Edge-triggered RAM simplifies system timing. Dual-port RAM doubles the effective throughput of FIFO applications. These features can be individually programmed in any XC4000 Series CLB.

Advantages of On-Chip and Edge-Triggered RAM

The on-chip RAM is extremely fast. The read access time is the same as the logic delay. The write access time is slightly slower. Both access times are much faster than any off-chip solution, because they avoid I/O delays.

Edge-triggered RAM, also called synchronous RAM, is a feature never before available in a Field Programmable Gate Array. The simplicity of designing with edge-triggered RAM, and the markedly higher achievable performance, add up to a significant improvement over existing devices with on-chip RAM.

Three application notes are available from Xilinx that discuss edge-triggered RAM: "*XC4000E Edge-Triggered and Dual-Port RAM Capability*," "*Implementing FIFOs in XC4000E RAM*," and "*Synchronous and Asynchronous FIFO Designs*." All three application notes apply to both XC4000E and XC4000X RAM.

Table 3: Supported RAM Modes

	16 x 1	16 x 2	32 x 1	Edge-Triggered Timing	Level-Sensitive Timing
Single-Port	√	√	√	√	√
Dual-Port	√			√	

RAM Configuration Options

The function generators in any CLB can be configured as RAM arrays in the following sizes:

- Two 16x1 RAMs: two data inputs and two data outputs with identical or, if preferred, different addressing for each RAM
- One 32x1 RAM: one data input and one data output.

One F or G function generator can be configured as a 16x1 RAM while the other function generators are used to implement any function of up to 5 inputs.

Additionally, the XC4000 Series RAM may have either of two timing modes:

- Edge-Triggered (Synchronous): data written by the designated edge of the CLB clock. WE acts as a true clock enable.
- Level-Sensitive (Asynchronous): an external WE signal acts as the write strobe.

The selected timing mode applies to both function generators within a CLB when both are configured as RAM.

The number of read ports is also programmable:

- Single Port: each function generator has a common read and write port
- Dual Port: both function generators are configured together as a single 16x1 dual-port RAM with one write port and two read ports. Simultaneous read and write operations to the same or different addresses are supported.

RAM configuration options are selected by placing the appropriate library symbol.

Choosing a RAM Configuration Mode

The appropriate choice of RAM mode for a given design should be based on timing and resource requirements, desired functionality, and the simplicity of the design process. Recommended usage is shown in Table 4.

The difference between level-sensitive, edge-triggered, and dual-port RAM is only in the write operation. Read operation and timing is identical for all modes of operation.

Table 4: RAM Mode Selection

	Level-Sensitive	Edge-Triggered	Dual-Port Edge-Triggered
Use for New Designs?	No	Yes	Yes
Size (16x1, Registered)	1/2 CLB	1/2 CLB	1 CLB
Simultaneous Read/Write	No	No	Yes
Relative Performance	X	2X	2X (4X effective)

RAM Inputs and Outputs

The F1-F4 and G1-G4 inputs to the function generators act as address lines, selecting a particular memory cell in each look-up table.

The functionality of the CLB control signals changes when the function generators are configured as RAM. The DIN/H2, H1, and SR/H0 lines become the two data inputs (D0, D1) and the Write Enable (WE) input for the 16x2 memory. When the 32x1 configuration is selected, D1 acts as the fifth address bit and D0 is the data input.

The contents of the memory cell(s) being addressed are available at the F' and G' function-generator outputs. They can exit the CLB through its X and Y outputs, or can be captured in the CLB flip-flop(s).

Configuring the CLB function generators as Read/Write memory does not affect the functionality of the other por-

XILINX®

Dual-Port Edge-Triggered Mode

In dual-port mode, both the F and G function generators are used to create a single 16x1 RAM array with one write port and two read ports. The resulting RAM array can be read and written simultaneously at two independent addresses. Simultaneous read and write operations at the same address are also supported.

Dual-port mode always has edge-triggered write timing, as shown in Figure 3.

Figure 6 shows a simple model of an XC4000 Series CLB configured as dual-port RAM. One address port, labeled A[3:0], supplies both the read and write address for the F function generator. This function generator behaves the same as a 16x1 single-port edge-triggered RAM array. The RAM output, Single Port Out (SPO), appears at the F function generator output. SPO, therefore, reflects the data at address A[3:0].

The other address port, labeled DPRA[3:0] for Dual Port Read Address, supplies the read address for the G function generator. The write address for the G function generator, however, comes from the address A[3:0]. The output from this 16x1 RAM array, Dual Port Out (DPO), appears at the G function generator output. DPO, therefore, reflects the data at address DPRA[3:0].

Therefore, by using A[3:0] for the write address and DPRA[3:0] for the read address, and reading only the DPO output, a FIFO that can read and write simultaneously is easily generated. Simultaneous access doubles the effective throughput of the FIFO.

The relationships between CLB pins and RAM inputs and outputs for dual-port, edge-triggered mode are shown in Table 6. See Figure 7 on page 16 for a block diagram of a CLB configured in this mode.
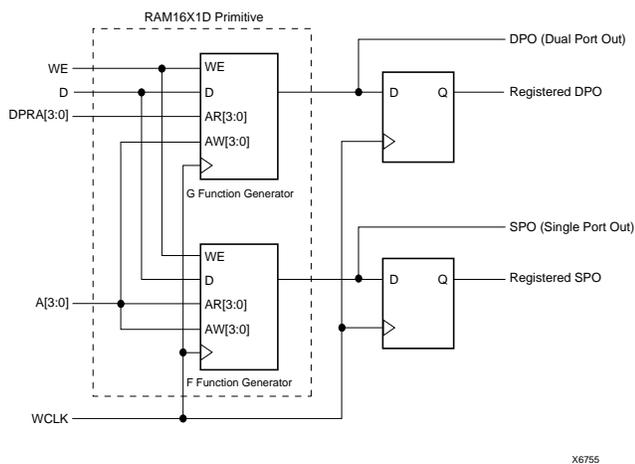

Figure 6: XC4000 Series Dual-Port RAM, Simple Model

Table 6: Dual-Port Edge-Triggered RAM Signals

RAM Signal	CLB Pin	Function
D	D0	Data In
A[3:0]	F1-F4	Read Address for F, Write Address for F and G
DPRA[3:0]	G1-G4	Read Address for G
WE	WE	Write Enable
WCLK	K	Clock
SPO	F'	Single Port Out (addressed by A[3:0])
DPO	G'	Dual Port Out (addressed by DPRA[3:0])

Note: The pulse following the active edge of WCLK (T_{WPS} in Figure 3) must be less than one millisecond wide. For most applications, this requirement is not overly restrictive; however, it must not be forgotten. Stopping WCLK at this point in the write cycle could result in excessive current and even damage to the larger devices if many CLBs are configured as edge-triggered RAM.

Single-Port Level-Sensitive Timing Mode

Note: Edge-triggered mode is recommended for all new designs. Level-sensitive mode, also called asynchronous mode, is still supported for XC4000 Series backward-compatibility with the XC4000 family.

Level-sensitive RAM timing is simple in concept but can be complicated in execution. Data and address signals are presented, then a positive pulse on the write enable pin (WE) performs a write into the RAM at the designated address. As indicated by the "level-sensitive" label, this RAM acts like a latch. During the WE High pulse, changing the data lines results in new data written to the old address. Changing the address lines while WE is High results in spurious data written to the new address—and possibly at other addresses as well, as the address lines inevitably do not all change simultaneously.

The user must generate a carefully timed WE signal. The delay on the WE signal and the address lines must be carefully verified to ensure that WE does not become active until after the address lines have settled, and that WE goes inactive before the address lines change again. The data must be stable before and after the falling edge of WE.

In practical terms, WE is usually generated by a 2X clock. If a 2X clock is not available, the falling edge of the system clock can be used. However, there are inherent risks in this approach, since the WE pulse must be guaranteed inactive before the next rising edge of the system clock. Several older application notes are available from Xilinx that discuss the design of level-sensitive RAMs.

However, the edge-triggered RAM available in the XC4000 Series is superior to level-sensitive RAM for almost every application.

6

XILINX®
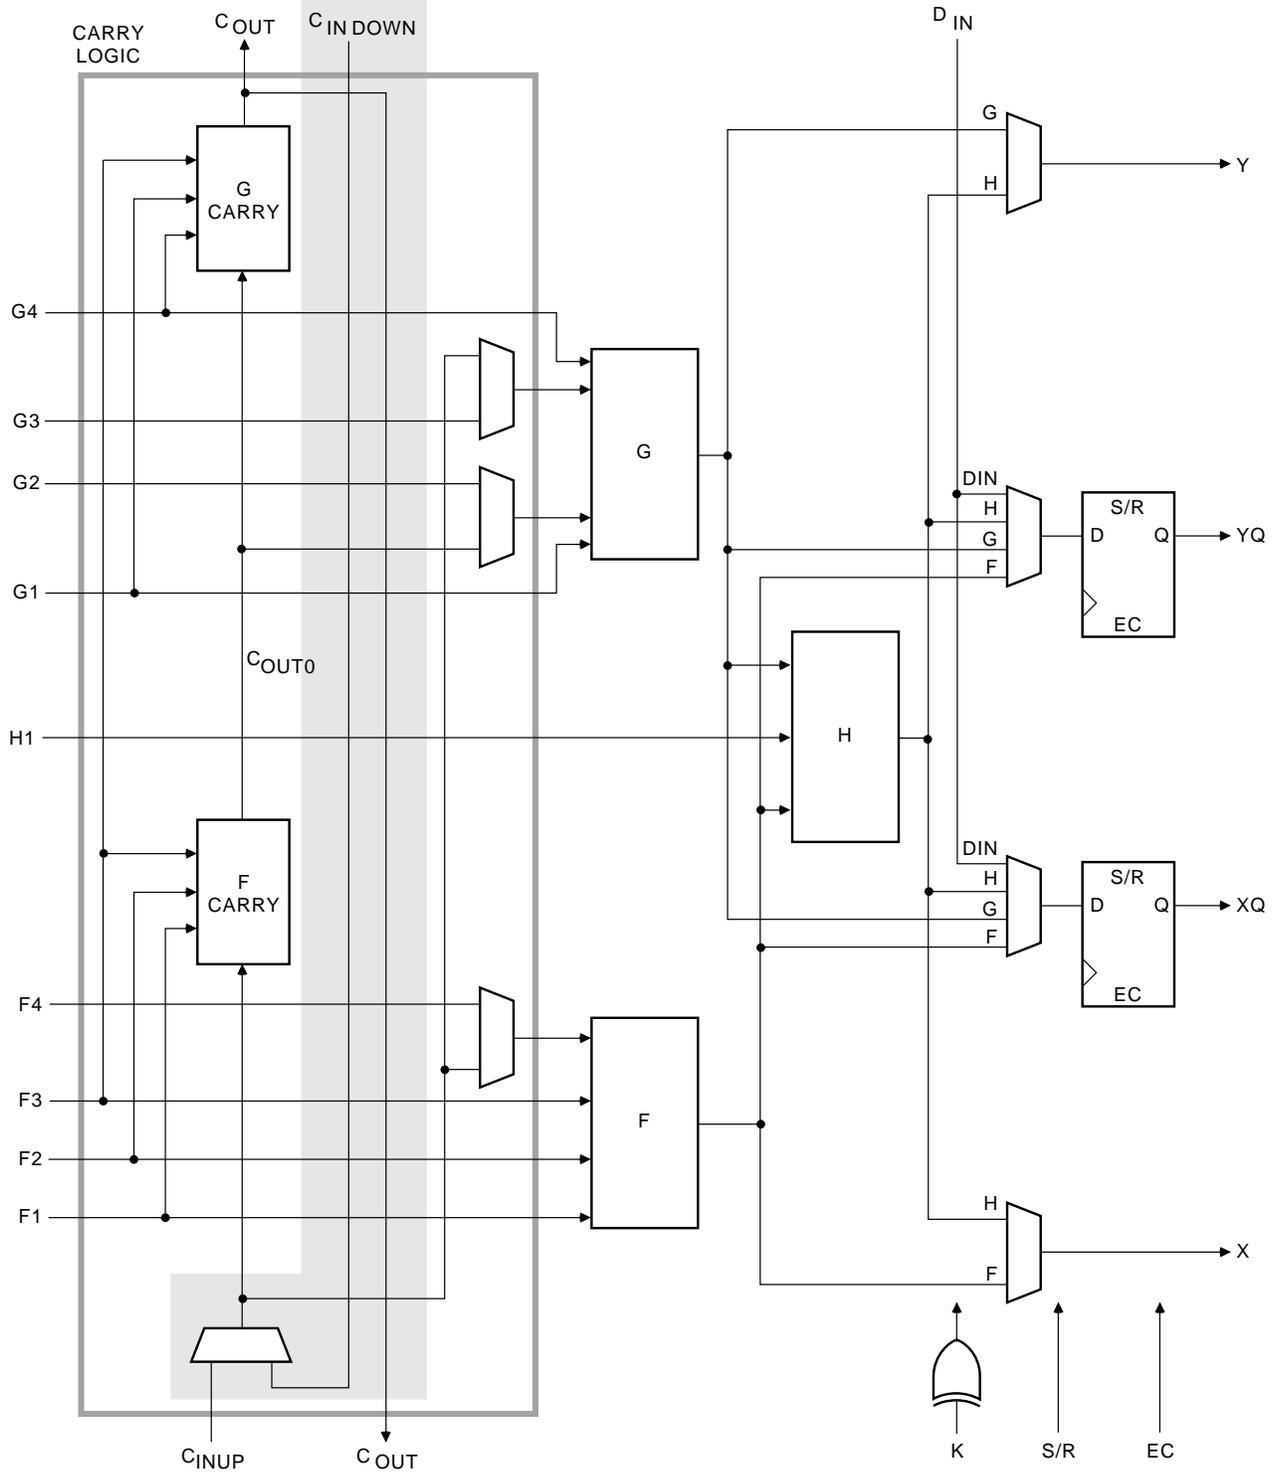
Figure 13: Fast Carry Logic in XC4000E CLB (shaded area not present in XC4000X)

Additional Input Latch for Fast Capture (XC4000X only)

The XC4000X IOB has an additional optional latch on the input. This latch, as shown in Figure 16, is clocked by the output clock — the clock used for the output flip-flop — rather than the input clock. Therefore, two different clocks can be used to clock the two input storage elements. This additional latch allows the very fast capture of input data, which is then synchronized to the internal clock by the IOB flip-flop or latch.

To use this Fast Capture technique, drive the output clock pin (the Fast Capture latching signal) from the output of one of the Global Early buffers supplied in the XC4000X. The second storage element should be clocked by a Global Low-Skew buffer, to synchronize the incoming data to the internal logic. (See Figure 17.) These special buffers are described in "Global Nets and Buffers (XC4000X only)" on page 37.

The Fast Capture latch (FCL) is designed primarily for use with a Global Early buffer. For Fast Capture, a single clock signal is routed through both a Global Early buffer and a Global Low-Skew buffer. (The two buffers share an input pad.) The Fast Capture latch is clocked by the Global Early buffer, and the standard IOB flip-flop or latch is clocked by the Global Low-Skew buffer. This mode is the safest way to use the Fast Capture latch, because the clock buffers on both storage elements are driven by the same pad. There is no external skew between clock pads to create potential problems.

To place the Fast Capture latch in a design, use one of the special library symbols, ILFFX or ILFLX. ILFFX is a transparent-Low Fast Capture latch followed by an active-High input flip-flop. ILFLX is a transparent-Low Fast Capture latch followed by a transparent-High input latch. Any of the clock inputs can be inverted before driving the library element, and the inverter is absorbed into the IOB. If a single BUFG output is used to drive both clock inputs, the software automatically runs the clock through both a Global Low-Skew buffer and a Global Early buffer, and clocks the Fast Capture latch appropriately.

Figure 16 on page 21 also shows a two-tap delay on the input. By default, if the Fast Capture latch is used, the Xilinx software assumes a Global Early buffer is driving the clock, and selects MEDDELAY to ensure a zero hold time. Select
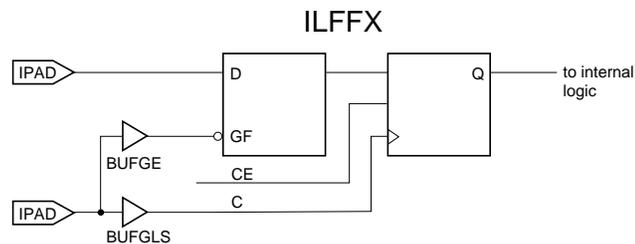
the desired delay based on the discussion in the previous subsection.

IOB Output Signals

Output signals can be optionally inverted within the IOB, and can pass directly to the pad or be stored in an edge-triggered flip-flop. The functionality of this flip-flop is shown in Table 11.

An active-High 3-state signal can be used to place the output buffer in a high-impedance state, implementing 3-state outputs or bidirectional I/O. Under configuration control, the output (OUT) and output 3-state (T) signals can be inverted. The polarity of these signals is independently configured for each IOB.

The 4-mA maximum output current specification of many FPGAs often forces the user to add external buffers, which are especially cumbersome on bidirectional I/O lines. The XC4000E and XC4000EX/XL devices solve many of these problems by providing a guaranteed output sink current of 12 mA. Two adjacent outputs can be interconnected externally to sink up to 24 mA. The XC4000E and XC4000EX/XL FPGAs can thus directly drive buses on a printed circuit board.

By default, the output pull-up structure is configured as a TTL-like totem-pole. The High driver is an n-channel pull-up transistor, pulling to a voltage one transistor threshold below Vcc. Alternatively, the outputs can be globally configured as CMOS drivers, with p-channel pull-up transistors pulling to Vcc. This option, applied using the bitstream generation software, applies to all outputs on the device. It is not individually programmable. In the XC4000XL, all outputs are pulled to the positive supply rail.

Table 11: Output Flip-Flop Functionality (active rising edge is shown)

Mode	Clock	Clock Enable	T	D	Q
Power-Up or GSR	X	X	0*	X	SR
Flip-Flop	X	0	0*	X	Q
	_/‾	1*	0*	D	D
	X	X	1	X	Z
	0	X	0*	X	Q

Legend:
X	Don't care
_/‾	Rising edge
SR	Set or Reset value. Reset is default.
0*	Input is Low or unconnected (default value)
1*	Input is High or unconnected (default value)
Z	3-state

Figure 17: Examples Using XC4000X FCL

or clear on reset and after configuration. Other than the global GSR net, no user-controlled set/reset signal is available to the I/O flip-flops. The choice of set or clear applies to both the initial state of the flip-flop and the response to the Global Set/Reset pulse. See "Global Set/Reset" on page 11 for a description of how to use GSR.

JTAG Support

Embedded logic attached to the IOBs contains test structures compatible with IEEE Standard 1149.1 for boundary scan testing, permitting easy chip and board-level testing. More information is provided in "Boundary Scan" on page 42.

Three-State Buffers

A pair of 3-state buffers is associated with each CLB in the array. (See Figure 27 on page 30.) These 3-state buffers can be used to drive signals onto the nearest horizontal longlines above and below the CLB. They can therefore be used to implement multiplexed or bidirectional buses on the horizontal longlines, saving logic resources. Programmable pull-up resistors attached to these longlines help to implement a wide wired-AND function.

The buffer enable is an active-High 3-state (i.e. an active-Low enable), as shown in Table 13.

Another 3-state buffer with similar access is located near each I/O block along the right and left edges of the array. (See Figure 33 on page 34.)

The horizontal longlines driven by the 3-state buffers have a weak keeper at each end. This circuit prevents undefined floating levels. However, it is overridden by any driver, even a pull-up resistor.

Special longlines running along the perimeter of the array can be used to wire-AND signals coming from nearby IOBs or from internal longlines. These longlines form the wide edge decoders discussed in "Wide Edge Decoders" on page 27.

Three-State Buffer Modes

The 3-state buffers can be configured in three modes:

- Standard 3-state buffer
- Wired-AND with input on the I pin
- Wired OR-AND

Standard 3-State Buffer

All three pins are used. Place the library element BUFT. Connect the input to the I pin and the output to the O pin. The T pin is an active-High 3-state (i.e. an active-Low enable). Tie the T pin to Ground to implement a standard buffer.

Wired-AND with Input on the I Pin

The buffer can be used as a Wired-AND. Use the WAND1 library symbol, which is essentially an open-drain buffer. WAND4, WAND8, and WAND16 are also available. See the *XACT Libraries Guide* for further information.

The T pin is internally tied to the I pin. Connect the input to the I pin and the output to the O pin. Connect the outputs of all the WAND1s together and attach a PULLUP symbol.

Wired OR-AND

The buffer can be configured as a Wired OR-AND. A High level on either input turns off the output. Use the WOR2AND library symbol, which is essentially an open-drain 2-input OR gate. The two input pins are functionally equivalent. Attach the two inputs to the I0 and I1 pins and tie the output to the O pin. Tie the outputs of all the WOR2ANDs together and attach a PULLUP symbol.

Three-State Buffer Examples

Figure 21 shows how to use the 3-state buffers to implement a wired-AND function. When all the buffer inputs are High, the pull-up resistor(s) provide the High output.

Figure 22 shows how to use the 3-state buffers to implement a multiplexer. The selection is accomplished by the buffer 3-state signal.

Pay particular attention to the polarity of the T pin when using these buffers in a design. Active-High 3-state (T) is identical to an active-Low output enable, as shown in Table 13.

Table 13: Three-State Buffer Functionality

IN	T	OUT
X	1	Z
IN	0	IN

$$Z = D_A \bullet D_B \bullet (D_C + D_D) \bullet (D_E + D_F)$$
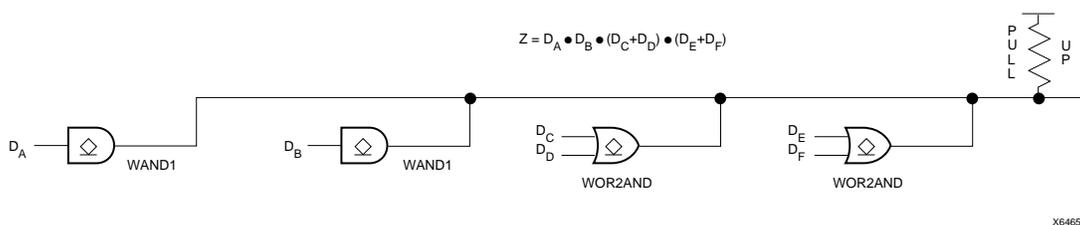
Figure 21:   Open-Drain Buffers Implement a Wired-AND Function

$$Z = D_A \cdot \overline{A} + D_B \cdot \overline{B} + D_C \cdot \overline{C} + D_N \cdot \overline{N}$$
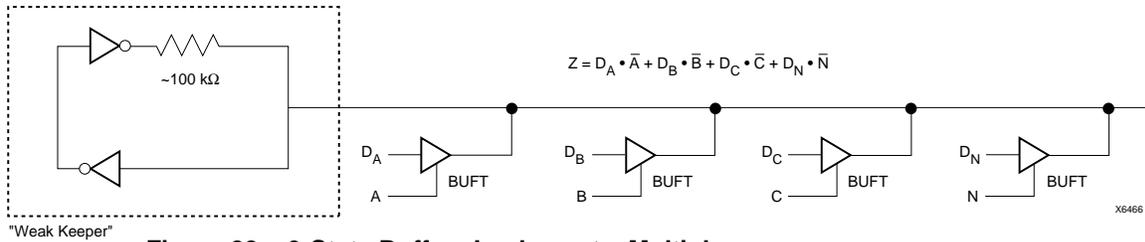
Figure 22:   3-State Buffers Implement a Multiplexer

Wide Edge Decoders

Dedicated decoder circuitry boosts the performance of wide decoding functions. When the address or data field is wider than the function generator inputs, FPGAs need multi-level decoding and are thus slower than PALs. XC4000 Series CLBs have nine inputs. Any decoder of up to nine inputs is, therefore, compact and fast. However, there is also a need for much wider decoders, especially for address decoding in large microprocessor systems.

An XC4000 Series FPGA has four programmable decoders located on each edge of the device. The inputs to each decoder are any of the IOB I1 signals on that edge plus one local interconnect per CLB row or column. Each row or column of CLBs provides up to three variables or their compliments., as shown in Figure 23. Each decoder generates a High output (resistor pull-up) when the AND condition of the selected inputs, or their complements, is true. This is analogous to a product term in typical PAL devices.

Each of these wired-AND gates is capable of accepting up to 42 inputs on the XC4005E and 72 on the XC4013E. There are up to 96 inputs for each decoder on the XC4028X and 132 on the XC4052X.   The decoders may also be split in two when a larger number of narrower decoders are required, for a maximum of 32 decoders per device.

The decoder outputs can drive CLB inputs, so they can be combined with other logic to form a PAL-like AND/OR structure. The decoder outputs can also be routed directly to the chip outputs. For fastest speed, the output should be on the same chip edge as the decoder. Very large PALs can be emulated by ORing the decoder outputs in a CLB. This decoding feature covers what has long been considered a weakness of older FPGAs. Users often resorted to external PALs for simple but fast decoding functions. Now, the dedicated decoders in the XC4000 Series device can implement these functions fast and efficiently.

To use the wide edge decoders, place one or more of the WAND library symbols (WAND1, WAND4, WAND8, WAND16). Attach a DECODE attribute or property to each WAND symbol. Tie the outputs together and attach a PUL-

LUP symbol. Location attributes or properties such as L (left edge) or TR (right half of top edge) should also be used to ensure the correct placement of the decoder inputs.
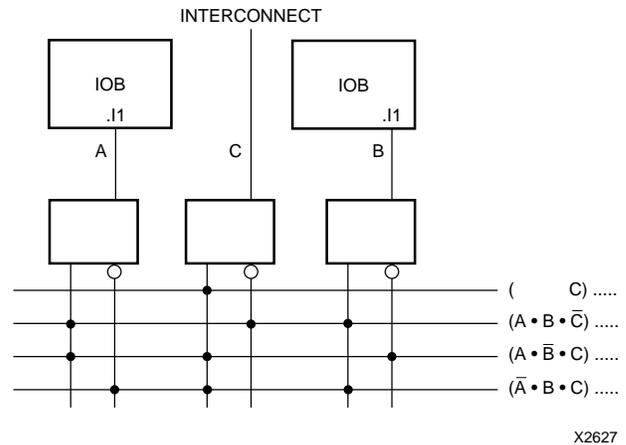
Figure 23:   XC4000 Series Edge Decoding Example
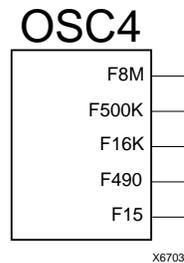
Figure 24:   XC4000 Series Oscillator Symbol

On-Chip Oscillator

XC4000 Series devices include an internal oscillator. This oscillator is used to clock the power-on time-out, for configuration memory clearing, and as the source of CCLK in Master configuration modes. The oscillator runs at a nominal 8 MHz frequency that varies with process, Vcc, and temperature. The output frequency falls between 4 and 10 MHz.

The oscillator output is optionally available after configuration. Any two of four resynchronized taps of a built-in divider are also available. These taps are at the fourth, ninth, fourteenth and nineteenth bits of the divider. Therefore, if the primary oscillator output is running at the nominal 8 MHz, the user has access to an 8 MHz clock, plus any two of 500 kHz, 16kHz, 490Hz and 15Hz (up to 10% lower for low-voltage devices). These frequencies can vary by as much as -50% or +25%.

These signals can be accessed by placing the OSC4 library element in a schematic or in HDL code (see Figure 24).

The oscillator is automatically disabled after configuration if the OSC4 symbol is not used in the design.

Programmable Interconnect

All internal connections are composed of metal segments with programmable switching points and switching matrices to implement the desired routing. A structured, hierarchical matrix of routing resources is provided to achieve efficient automated routing.

The XC4000E and XC4000X share a basic interconnect structure. XC4000X devices, however, have additional routing not available in the XC4000E. The extra routing resources allow high utilization in high-capacity devices. All XC4000X-specific routing resources are clearly identified throughout this section. Any resources not identified as XC4000X-specific are present in all XC4000 Series devices.

This section describes the varied routing resources available in XC4000 Series devices. The implementation software automatically assigns the appropriate resources based on the density and timing requirements of the design.

Interconnect Overview

There are several types of interconnect.

- CLB routing is associated with each row and column of the CLB array.
- IOB routing forms a ring (called a VersaRing) around the outside of the CLB array. It connects the I/O with the internal logic blocks.

- Global routing consists of dedicated networks primarily designed to distribute clocks throughout the device with minimum delay and skew. Global routing can also be used for other high-fanout signals.

Five interconnect types are distinguished by the relative length of their segments: single-length lines, double-length lines, quad and octal lines (XC4000X only), and longlines. In the XC4000X, direct connects allow fast data flow between adjacent CLBs, and between IOBs and CLBs.

Extra routing is included in the IOB pad ring. The XC4000X also includes a ring of octal interconnect lines near the IOBs to improve pin-swapping and routing to locked pins.

XC4000E/X devices include two types of global buffers. These global buffers have different properties, and are intended for different purposes. They are discussed in detail later in this section.

CLB Routing Connections

A high-level diagram of the routing resources associated with one CLB is shown in Figure 25. The shaded arrows represent routing present only in XC4000X devices.

Table 14 shows how much routing of each type is available in XC4000E and XC4000X CLB arrays. Clearly, very large designs, or designs with a great deal of interconnect, will route more easily in the XC4000X. Smaller XC4000E designs, typically requiring significantly less interconnect, do not require the additional routing.

Figure 27 on page 30 is a detailed diagram of both the XC4000E and the XC4000X CLB, with associated routing. The shaded square is the programmable switch matrix, present in both the XC4000E and the XC4000X. The L-shaped shaded area is present only in XC4000X devices. As shown in the figure, the XC4000X block is essentially an XC4000E block with additional routing.

CLB inputs and outputs are distributed on all four sides, providing maximum routing flexibility. In general, the entire architecture is symmetrical and regular. It is well suited to established placement and routing algorithms. Inputs, outputs, and function generators can freely swap positions within a CLB to avoid routing congestion during the placement and routing operation.

QUAD

DOUBLE

SINGLE

DOUBLE

LONG

DIRECT

FEEDBACK

LONG

F4 C4 G4

YQ
Y

G1
C1
F1

CLB

G3
C3
F3

K
X
XQ

F2 C2 G2

QUAD LONG GLOBAL LONG DOUBLE SINGLE DOUBLE LONG GLOBAL DIRECT FEEDBACK
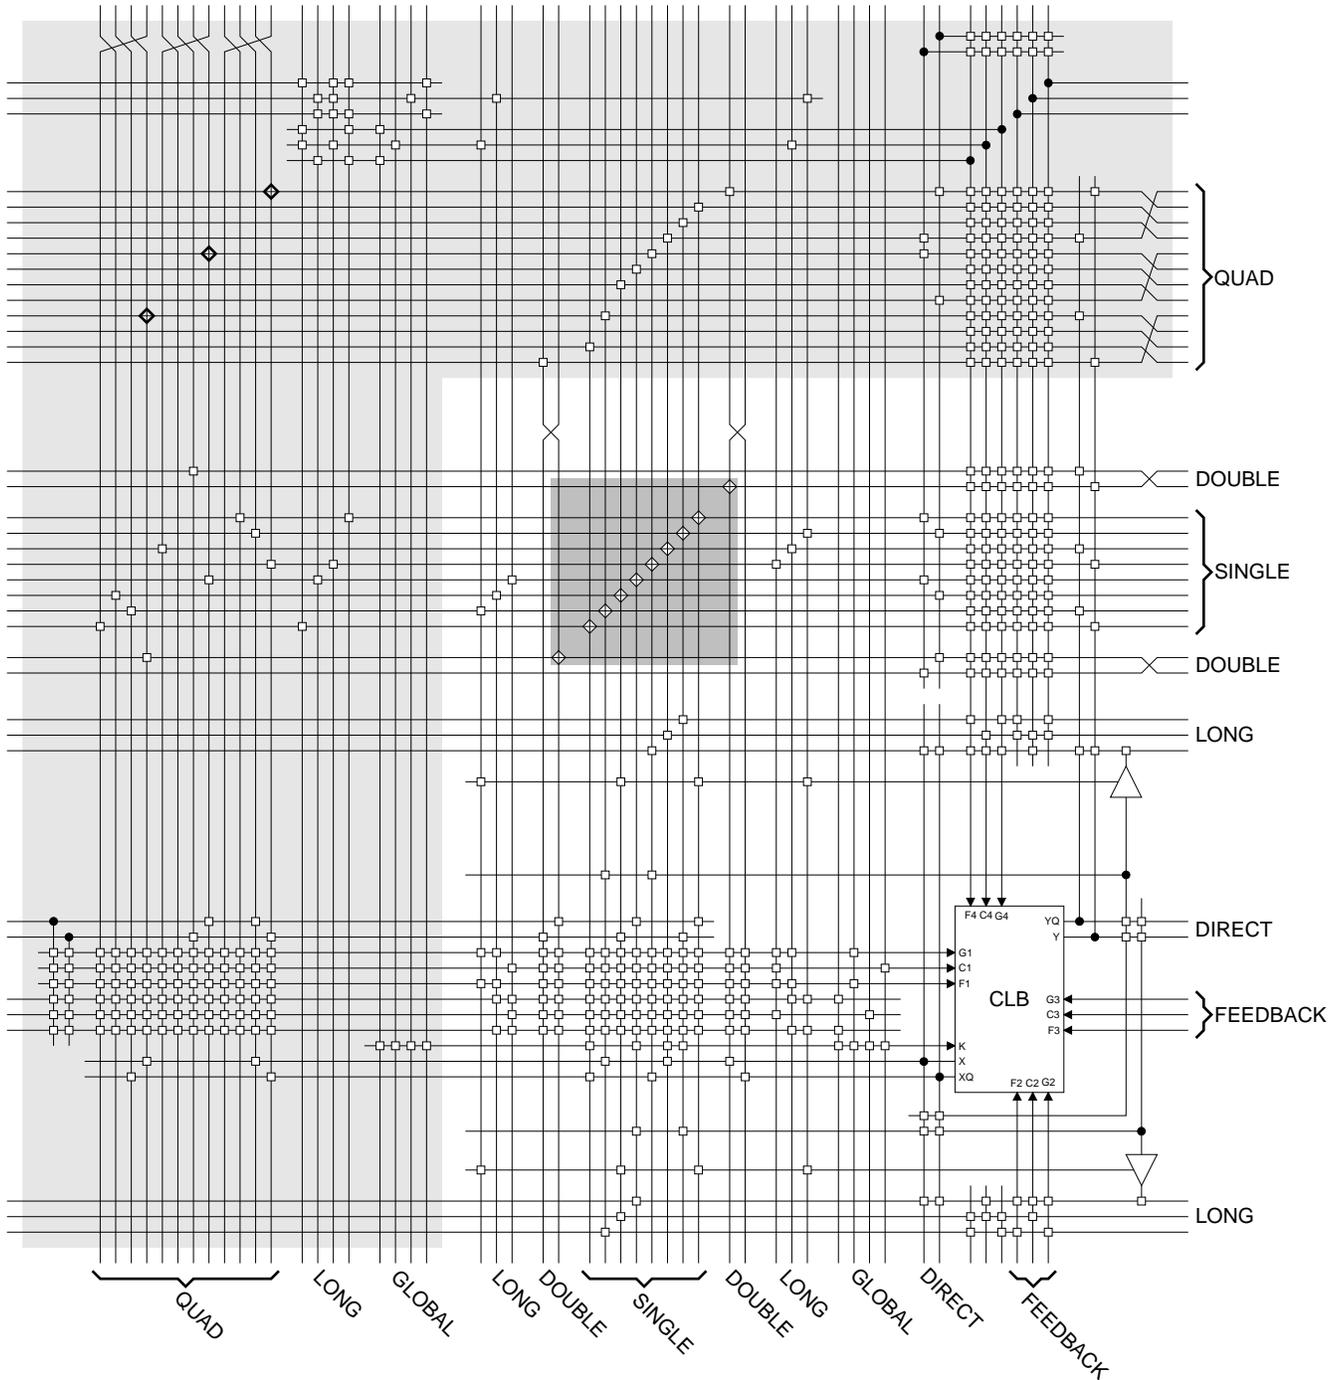
☐ Common to XC4000E and XC4000X

☐ XC4000X only

☐ Programmable Switch Matrix

Figure 27:   Detail of Programmable Interconnect Associated with XC4000 Series CLB

XILINX®

X5995
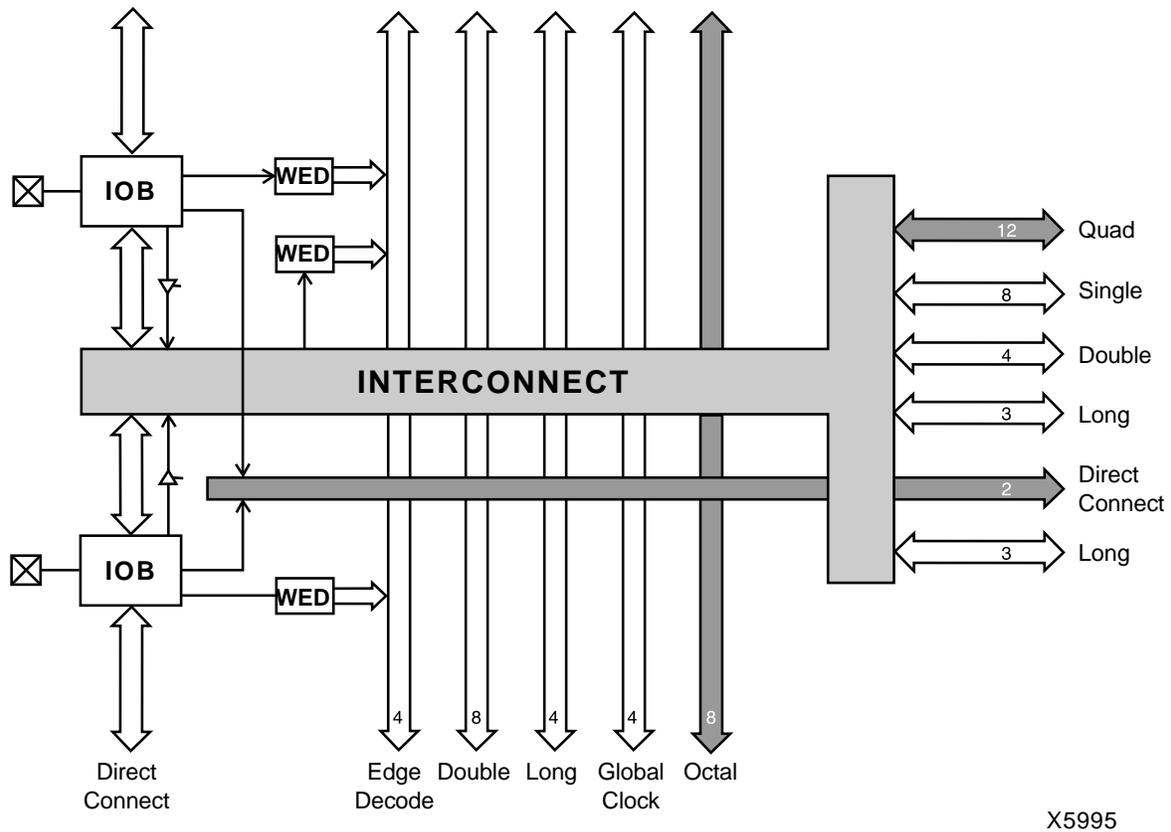
Figure 31: High-Level Routing Diagram of XC4000 Series VersaRing (Left Edge)
WED = Wide Edge Decoder, IOB = I/O Block (shaded arrows indicate XC4000X only)

6

Segment with nearest buffer
connects to segment with furthest buffer

X9015
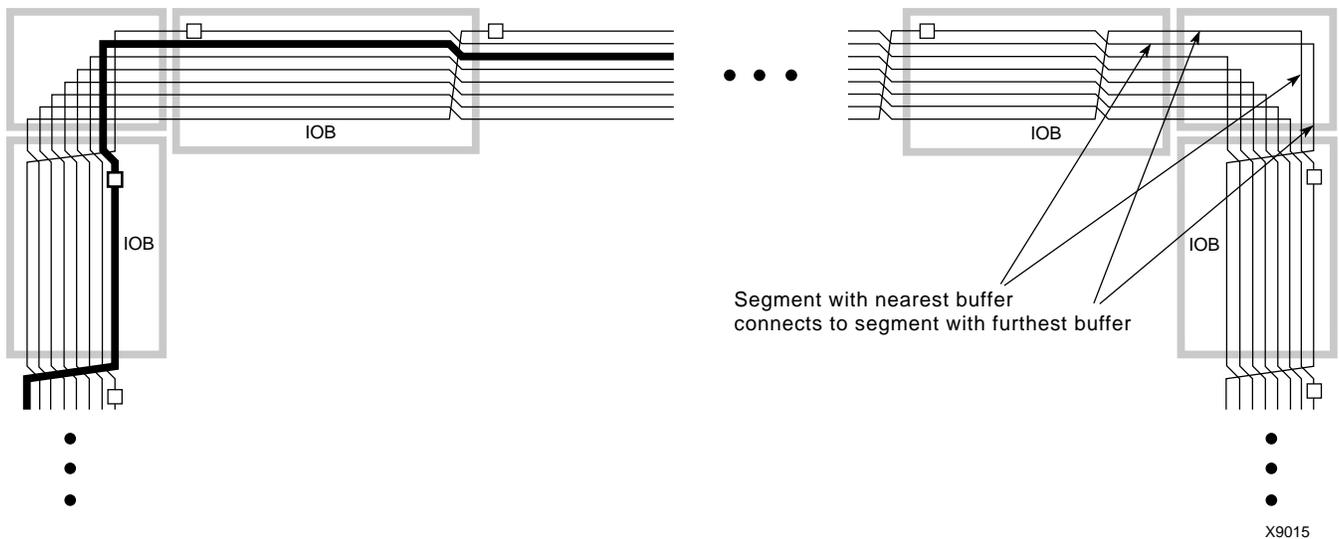
Figure 32: XC4000X Octal I/O Routing

Global Nets and Buffers (XC4000X only)

Eight vertical longlines in each CLB column are driven by special global buffers. These longlines are in addition to the vertical longlines used for standard interconnect. The global lines are broken in the center of the array, to allow faster distribution and to minimize skew across the whole array. Each half-column global line has its own buffered multiplexer, as shown in Figure 35. The top and bottom global lines cannot be connected across the center of the device, as this connection might introduce unacceptable skew. The top and bottom halves of the global lines must be separately driven — although they can be driven by the same global buffer.

The eight global lines in each CLB column can be driven by either of two types of global buffers. They can also be driven by internal logic, because they can be accessed by single, double, and quad lines at the top, bottom, half, and quarter points. Consequently, the number of different clocks that can be used simultaneously in an XC4000X device is very large.

There are four global lines feeding the IOBs at the left edge of the device. IOBs along the right edge have eight global lines. There is a single global line along the top and bottom edges with access to the IOBs. All IOB global lines are broken at the center. They cannot be connected across the center of the device, as this connection might introduce unacceptable skew.

IOB global lines can be driven from two types of global buffers, or from local interconnect. Alternatively, top and bottom IOBs can be clocked from the global lines in the adjacent CLB column.

Two different types of clock buffers are available in the XC4000X:

- Global Low-Skew Buffers (BUFGLS)
- Global Early Buffers (BUFGE)

Global Low-Skew Buffers are the standard clock buffers. They should be used for most internal clocking, whenever a large portion of the device must be driven.

Global Early Buffers are designed to provide a faster clock access, but CLB access is limited to one-fourth of the device. They also facilitate a faster I/O interface.

Figure 35 is a conceptual diagram of the global net structure in the XC4000X.

Global Early buffers and Global Low-Skew buffers share a single pad. Therefore, the same IPAD symbol can drive one buffer of each type, in parallel. This configuration is particularly useful when using the Fast Capture latches, as described in "IOB Input Signals" on page 20. Paired Global

Early and Global Low-Skew buffers share a common input; they cannot be driven by two different signals.

Choosing an XC4000X Clock Buffer

The clocking structure of the XC4000X provides a large variety of features. However, it can be simple to use, without understanding all the details. The software automatically handles clocks, along with all other routing, when the appropriate clock buffer is placed in the design. In fact, if a buffer symbol called BUFG is placed, rather than a specific type of buffer, the software even chooses the buffer most appropriate for the design. The detailed information in this section is provided for those users who want a finer level of control over their designs.

If fine control is desired, use the following summary and Table 15 on page 35 to choose an appropriate clock buffer.

- The simplest thing to do is to use a Global Low-Skew buffer.
- If a faster clock path is needed, try a BUFG. The software will first try to use a Global Low-Skew Buffer. If timing requirements are not met, a faster buffer will automatically be used.
- If a single quadrant of the chip is sufficient for the clocked logic, and the timing requires a faster clock than the Global Low-Skew buffer, use a Global Early buffer.

Global Low-Skew Buffers

Each corner of the XC4000X device has two Global Low-Skew buffers. Any of the eight Global Low-Skew buffers can drive any of the eight vertical Global lines in a column of CLBs. In addition, any of the buffers can drive any of the four vertical lines accessing the IOBs on the left edge of the device, and any of the eight vertical lines accessing the IOBs on the right edge of the device. (See Figure 36 on page 38.)

IOBs at the top and bottom edges of the device are accessed through the vertical Global lines in the CLB array, as in the XC4000E. Any Global Low-Skew buffer can, therefore, access every IOB and CLB in the device.

The Global Low-Skew buffers can be driven by either semi-dedicated pads or internal logic.

To use a Global Low-Skew buffer, instantiate a BUFGLS element in a schematic or in HDL code. If desired, attach a LOC attribute or property to direct placement to the designated location. For example, attach a LOC=T attribute or property to direct that a BUFGLS be placed in one of the two Global Low-Skew buffers on the top edge of the device, or a LOC=TR to indicate the Global Low-Skew buffer on the top edge of the device, on the right.

6

Table 16: Pin Descriptions (Continued)

Pin Name	I/O During Config.	I/O After Config.	Pin Description
$\overline{CS0}$, CS1, $\overline{WS}$, $\overline{RS}$	I	I/O	These four inputs are used in Asynchronous Peripheral mode. The chip is selected when $\overline{CS0}$ is Low and CS1 is High. While the chip is selected, a Low on Write Strobe ($\overline{WS}$) loads the data present on the D0 - D7 inputs into the internal data buffer. A Low on Read Strobe ($\overline{RS}$) changes D7 into a status output — High if Ready, Low if Busy — and drives D0 - D6 High. In Express mode, CS1 is used as a serial-enable signal for daisy-chaining. $\overline{WS}$ and $\overline{RS}$ should be mutually exclusive, but if both are Low simultaneously, the Write Strobe overrides. After configuration, these are user-programmable I/O pins.
A0 - A17	O	I/O	During Master Parallel configuration, these 18 output pins address the configuration EPROM. After configuration, they are user-programmable I/O pins.
A18 - A21 (XC4003XL to XC4085XL)	O	I/O	During Master Parallel configuration with an XC4000X master, these 4 output pins add 4 more bits to address the configuration EPROM. After configuration, they are user-pro-grammable I/O pins. (See Master Parallel Configuration section for additional details.)
D0 - D7	I	I/O	During Master Parallel and Peripheral configuration, these eight input pins receive con-figuration data. After configuration, they are user-programmable I/O pins.
DIN	I	I/O	During Slave Serial or Master Serial configuration, DIN is the serial configuration data input receiving data on the rising edge of CCLK. During Parallel configuration, DIN is the D0 input. After configuration, DIN is a user-programmable I/O pin.
DOUT	O	I/O	During configuration in any mode but Express mode, DOUT is the serial configuration data output that can drive the DIN of daisy-chained slave FPGAs. DOUT data changes on the falling edge of CCLK, one-and-a-half CCLK periods after it was received at the DIN input. In Express modefor XC4000E and XC4000X only, DOUT is the status output that can drive the CS1 of daisy-chained FPGAs, to enable and disable downstream devices. After configuration, DOUT is a user-programmable I/O pin.
Unrestricted User-Programmable I/O Pins			
I/O	Weak Pull-up	I/O	These pins can be configured to be input and/or output after configuration is completed. Before configuration is completed, these pins have an internal high-value pull-up resis-tor (25 kΩ - 100 kΩ) that defines the logic level as High.

Boundary Scan

The 'bed of nails' has been the traditional method of testing electronic assemblies. This approach has become less appropriate, due to closer pin spacing and more sophisti-cated assembly methods like surface-mount technology and multi-layer boards. The IEEE Boundary Scan Standard 1149.1 was developed to facilitate board-level testing of electronic assemblies. Design and test engineers can imbed a standard test logic structure in their device to achieve high fault coverage for I/O and internal logic. This structure is easily implemented with a four-pin interface on any boundary scan-compatible IC. IEEE 1149.1-compati-ble devices may be serial daisy-chained together, con-nected in parallel, or a combination of the two.

The XC4000 Series implements IEEE 1149.1-compatible BYPASS, PRELOAD/SAMPLE and EXTEST boundary scan instructions. When the boundary scan configuration option is selected, three normal user I/O pins become ded-icated inputs for these functions. Another user output pin becomes the dedicated boundary scan output. The details

of how to enable this circuitry are covered later in this sec-tion.

By exercising these input signals, the user can serially load commands and data into these devices to control the driv-ing of their outputs and to examine their inputs. This method is an improvement over bed-of-nails testing. It avoids the need to over-drive device outputs, and it reduces the user interface to four pins. An optional fifth pin, a reset for the control logic, is described in the standard but is not implemented in Xilinx devices.

The dedicated on-chip logic implementing the IEEE 1149.1 functions includes a 16-state machine, an instruction regis-ter and a number of data registers. The functional details can be found in the IEEE 1149.1 specification and are also discussed in the Xilinx application note XAPP 017: "*Bound-ary Scan in XC4000 Devices.*"

Figure shows a simplified block diagram of the XC4000E Input/Output Block with boundary scan implemented. XC4000X boundary scan logic is identical.

Figure 41 on page 44 is a diagram of the XC4000 Series boundary scan logic. It includes three bits of Data Register per IOB, the IEEE 1149.1 Test Access Port controller, and the Instruction Register with decodes.

XC4000 Series devices can also be configured through the boundary scan logic. See "Readback" on page 55.

Data Registers

The primary data register is the boundary scan register. For each IOB pin in the FPGA, bonded or not, it includes three bits for In, Out and 3-State Control. Non-IOB pins have appropriate partial bit population for In or Out only. $\overline{PRO-GRAM}$, CCLK and DONE are not included in the boundary scan register. Each EXTEST CAPTURE-DR state captures all In, Out, and 3-state pins.

The data register also includes the following non-pin bits: TDO.T, and TDO.O, which are always bits 0 and 1 of the

data register, respectively, and BSCANT.UPD, which is always the last bit of the data register. These three boundary scan bits are special-purpose Xilinx test signals.

The other standard data register is the single flip-flop BYPASS register. It synchronizes data being passed through the FPGA to the next downstream boundary scan device.

The FPGA provides two additional data registers that can be specified using the BSCAN macro. The FPGA provides two user pins (BSCAN.SEL1 and BSCAN.SEL2) which are the decodes of two user instructions. For these instructions, two corresponding pins (BSCAN.TDO1 and BSCAN.TDO2) allow user scan data to be shifted out on TDO. The data register clock (BSCAN.DRCK) is available for control of test logic which the user may wish to implement with CLBs. The NAND of TCK and RUN-TEST-IDLE is also provided (BSCAN.IDLE).
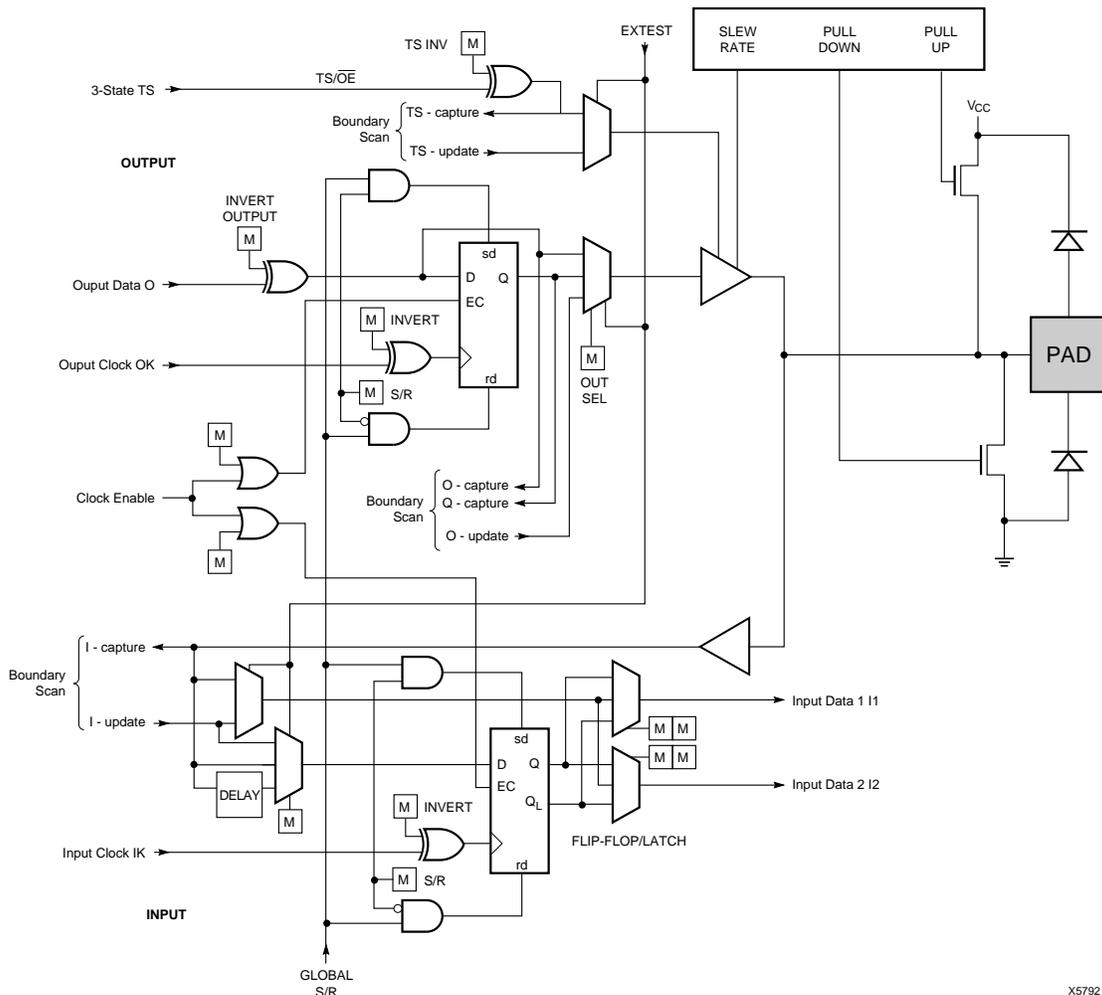

Figure 40: Block Diagram of XC4000E IOB with Boundary Scan (some details not shown). XC4000X Boundary Scan Logic is Identical.

XILINX®

is passed through and is captured by each FPGA when it recognizes the 0010 preamble. Following the length-count data, each FPGA outputs a High on DOUT until it has received its required number of data frames.

After an FPGA has received its configuration data, it passes on any additional frame start bits and configuration data on DOUT. When the total number of configuration clocks applied after memory initialization equals the value of the 24-bit length count, the FPGAs begin the start-up sequence and become operational together. FPGA I/O are normally released two CCLK cycles after the last configuration bit is received. Figure 47 on page 53 shows the start-up timing for an XC4000 Series device.

The daisy-chained bitstream is not simply a concatenation of the individual bitstreams. The PROM file formatter must be used to combine the bitstreams for a daisy-chained configuration.

Multi-Family Daisy Chain

All Xilinx FPGAs of the XC2000, XC3000, and XC4000 Series use a compatible bitstream format and can, therefore, be connected in a daisy chain in an arbitrary sequence. There is, however, one limitation. The lead device must belong to the highest family in the chain. If the chain contains XC4000 Series devices, the master normally cannot be an XC2000 or XC3000 device.

The reason for this rule is shown in Figure 47 on page 53. Since all devices in the chain store the same length count value and generate or receive one common sequence of CCLK pulses, they all recognize length-count match on the same CCLK edge, as indicated on the left edge of Figure 47. The master device then generates additional CCLK pulses until it reaches its finish point F. The different families generate or require different numbers of additional CCLK pulses until they reach F. Not reaching F means that the device does not really finish its configuration, although DONE may have gone High, the outputs became active, and the internal reset was released. For the XC4000 Series device, not reaching F means that readback cannot be ini-

tiated and most boundary scan instructions cannot be used.

The user has some control over the relative timing of these events and can, therefore, make sure that they occur at the proper time and the finish point F is reached. Timing is controlled using options in the bitstream generation software.

XC3000 Master with an XC4000 Series Slave

Some designers want to use an inexpensive lead device in peripheral mode and have the more precious I/O pins of the XC4000 Series devices all available for user I/O. Figure 44 provides a solution for that case.

This solution requires one CLB, one IOB and pin, and an internal oscillator with a frequency of up to 5 MHz as a clock source. The XC3000 master device must be configured with late Internal Reset, which is the default option.

One CLB and one IOB in the lead XC3000-family device are used to generate the additional CCLK pulse required by the XC4000 Series devices. When the lead device removes the internal RESET signal, the 2-bit shift register responds to its clock input and generates an active Low output signal for the duration of the subsequent clock period. An external connection between this output and CCLK thus creates the extra CCLK pulse.
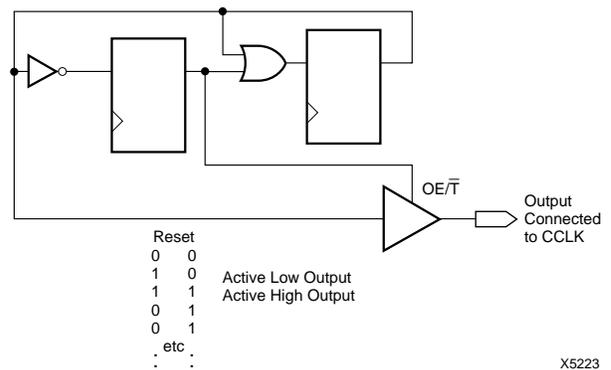

Figure 44: CCLK Generation for XC3000 Master Driving an XC4000 Series Slave

used), and if RAM is present, the RAM content must be unchanged.

Statistically, one error out of 2048 might go undetected.

Configuration Sequence

There are four major steps in the XC4000 Series power-up configuration sequence.

- Configuration Memory Clear
- Initialization
- Configuration
- Start-Up

The full process is illustrated in Figure 46.

Configuration Memory Clear

When power is first applied or is reapplied to an FPGA, an internal circuit forces initialization of the configuration logic. When Vcc reaches an operational level, and the circuit passes the write and read test of a sample pair of configuration bits, a time delay is started. This time delay is nominally 16 ms, and up to 10% longer in the low-voltage devices. The delay is four times as long when in Master Modes (M0 Low), to allow ample time for all slaves to reach a stable Vcc. When all INIT pins are tied together, as recommended, the longest delay takes precedence. Therefore, devices with different time delays can easily be mixed and matched in a daisy chain.

This delay is applied only on power-up. It is not applied when re-configuring an FPGA by pulsing the PROGRAM pin


Polynomial: X16 + X15 + X2 + 1
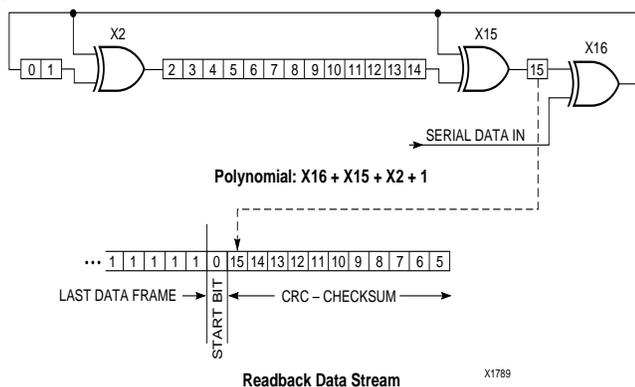
**Figure 45: Circuit for Generating CRC-16**
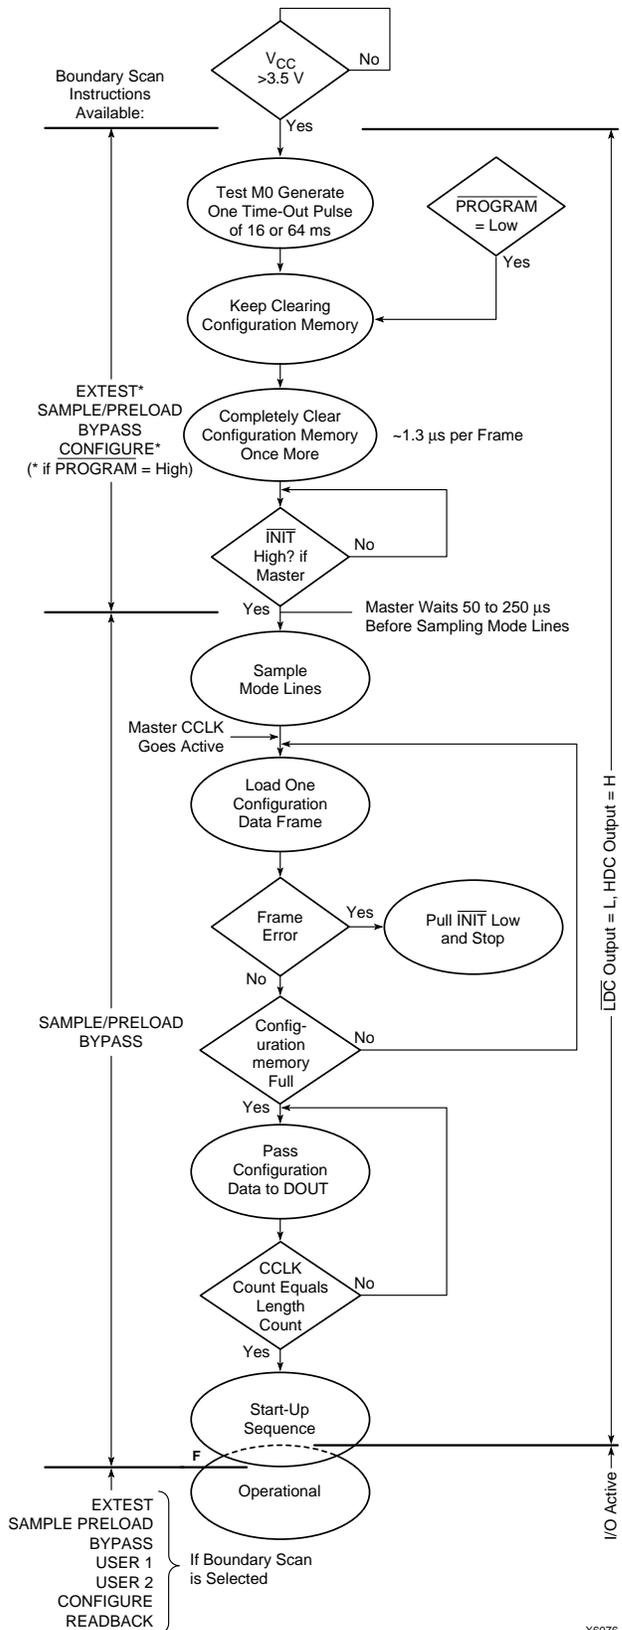
Figure 46: Power-up Configuration Sequence

Start-up from a User Clock (STARTUP.CLK)

When, instead of CCLK, a user-supplied start-up clock is selected, Q1 is used to bridge the unknown phase relationship between CCLK and the user clock. This arbitration causes an unavoidable one-cycle uncertainty in the timing of the rest of the start-up sequence.

DONE Goes High to Signal End of Configuration

XC4000 Series devices read the expected length count from the bitstream and store it in an internal register. The length count varies according to the number of devices and the composition of the daisy chain. Each device also counts the number of CCLKs during configuration.

Two conditions have to be met in order for the DONE pin to go high:

- the chip's internal memory must be full, and
- the configuration length count must be met, *exactly*.

This is important because the counter that determines when the length count is met begins with the very first CCLK, not the first one after the preamble.

Therefore, if a stray bit is inserted before the preamble, or the data source is not ready at the time of the first CCLK, the internal counter that holds the number of CCLKs will be one ahead of the actual number of data bits read. At the end of configuration, the configuration memory will be full, but the number of bits in the internal counter will not match the expected length count.

As a consequence, a Master mode device will continue to send out CCLKs until the internal counter turns over to zero, and then reaches the correct length count a second time. This will take several seconds [$2^{24} * $ CCLK period] — which is sometimes interpreted as the device not configuring at all.

If it is not possible to have the data ready at the time of the first CCLK, the problem can be avoided by increasing the number in the length count by the appropriate value. The *XACT User Guide* includes detailed information about manually altering the length count.

Note that DONE is an open-drain output and does not go High unless an internal pull-up is activated or an external pull-up is attached. The internal pull-up is activated as the default by the bitstream generation software.

Release of User I/O After DONE Goes High

By default, the user I/O are released one CCLK cycle after the DONE pin goes High. If CCLK is not clocked after DONE goes High, the outputs remain in their initial state — 3-stated, with a 50 kΩ - 100 kΩ pull-up. The delay from DONE High to active user I/O is controlled by an option to the bitstream generation software.

Release of Global Set/Reset After DONE Goes High

By default, Global Set/Reset (GSR) is released two CCLK cycles after the DONE pin goes High. If CCLK is not clocked twice after DONE goes High, all flip-flops are held in their initial set or reset state. The delay from DONE High to GSR inactive is controlled by an option to the bitstream generation software.

Configuration Complete After DONE Goes High

Three full CCLK cycles are required after the DONE pin goes High, as shown in Figure 47 on page 53. If CCLK is not clocked three times after DONE goes High, readback cannot be initiated and most boundary scan instructions cannot be used.

Configuration Through the Boundary Scan Pins

XC4000 Series devices can be configured through the boundary scan pins. The basic procedure is as follows:

- Power up the FPGA with $\overline{INIT}$ held Low (or drive the $\overline{PROGRAM}$ pin Low for more than 300 ns followed by a High while holding $\overline{INIT}$ Low). Holding $\overline{INIT}$ Low allows enough time to issue the CONFIG command to the FPGA. The pin can be used as I/O after configuration if a resistor is used to hold $\overline{INIT}$ Low.
- Issue the CONFIG command to the TMS input
- Wait for $\overline{INIT}$ to go High
- Sequence the boundary scan Test Access Port to the SHIFT-DR state
- Toggle TCK to clock data into TDI pin.

The user must account for all TCK clock cycles after INIT goes High, as all of these cycles affect the Length Count compare.

For more detailed information, refer to the Xilinx application note XAPP017, "*Boundary Scan in XC4000 Devices.*" This application note also applies to XC4000E and XC4000X devices.

XC4000E/EX/XL Program Readback Switching Characteristic Guidelines

Testing of the switching parameters is modeled after testing methods specified by MIL-M-38510/605. All devices are 100% functionally tested. Internal timing parameters are not measured directly. They are derived from benchmark timing patterns that are taken at device introduction, prior to any process improvements.

The following guidelines reflect worst-case values over the recommended operating conditions.
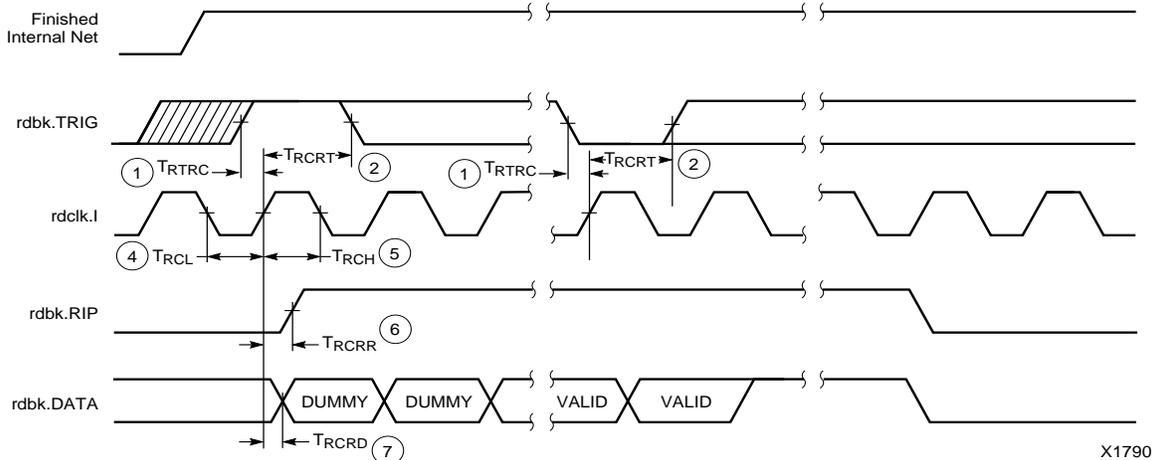
X1790

E/EX

	Description		Symbol	Min	Max	Units
rdbk.TRIG	rdbk.TRIG setup to initiate and abort Readback	1	T_{RTRC}	200	-	ns
	rdbk.TRIG hold to initiate and abort Readback	2	T_{RCRT}	50	-	ns
rdclk.1	rdbk.DATA delay	7	T_{RCRD}	-	250	ns
	rdbk.RIP delay	6	T_{RCRR}	-	250	ns
	High time	5	T_{RCH}	250	500	ns
	Low time	4	T_{RCL}	250	500	ns

Note 1: Timing parameters apply to all speed grades.
Note 2: If rdbk.TRIG is High prior to Finished, Finished will trigger the first Readback.

XL

	Description		Symbol	Min	Max	Units
rdbk.TRIG	rdbk.TRIG setup to initiate and abort Readback	1	T_{RTRC}	200	-	ns
	rdbk.TRIG hold to initiate and abort Readback	2	T_{RCRT}	50	-	ns
rdclk.1	rdbk.DATA delay	7	T_{RCRD}	-	250	ns
	rdbk.RIP delay	6	T_{RCRR}	-	250	ns
	High time	5	T_{RCH}	250	500	ns
	Low time	4	T_{RCL}	250	500	ns

Note 1: Timing parameters apply to all speed grades.
Note 2: If rdbk.TRIG is High prior to Finished, Finished will trigger the first Readback.

Configuration Timing

The seven configuration modes are discussed in detail in this section. Timing specifications are included.

Slave Serial Mode

In Slave Serial mode, an external signal drives the CCLK input of the FPGA. The serial configuration bitstream must be available at the DIN input of the lead FPGA a short setup time before each rising CCLK edge.

The lead FPGA then presents the preamble data—and all data that overflows the lead device—on its DOUT pin.

There is an internal delay of 0.5 CCLK periods, which means that DOUT changes on the falling CCLK edge, and the next FPGA in the daisy chain accepts data on the subsequent rising CCLK edge.

Figure 51 shows a full master/slave system. An XC4000 Series device in Slave Serial mode should be connected as shown in the third device from the left.

Slave Serial mode is selected by a <111> on the mode pins (M2, M1, M0). Slave Serial is the default mode if the mode pins are left unconnected, as they have weak pull-up resistors during configuration.
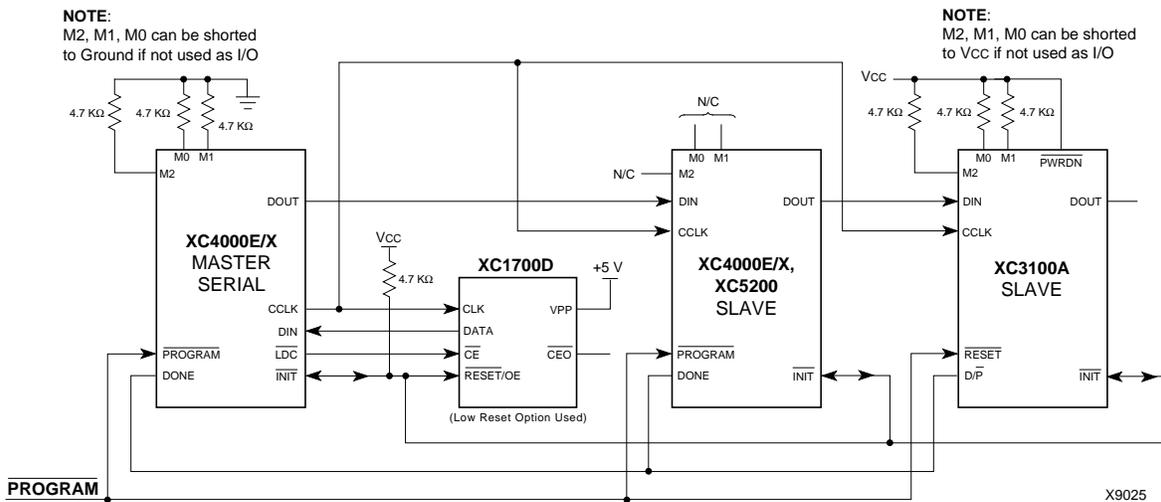

Figure 51: Master/Slave Serial Mode Circuit Diagram
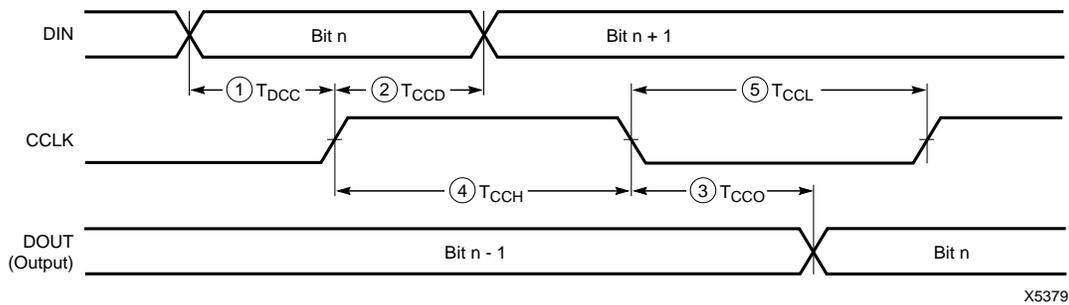
	Description		Symbol	Min	Max	Units
CCLK	DIN setup	1	T_{DCC}	20		ns
	DIN hold	2	T_{CCD}	0		ns
	DIN to DOUT	3	T_{CCO}		30	ns
	High time	4	T_{CCH}	45		ns
	Low time	5	T_{CCL}	45		ns
	Frequency		F_{CC}		10	MHz

Note: Configuration must be delayed until the INIT pins of all daisy-chained FPGAs are High.

Figure 52: Slave Serial Mode Programming Switching Characteristics

XILINX®
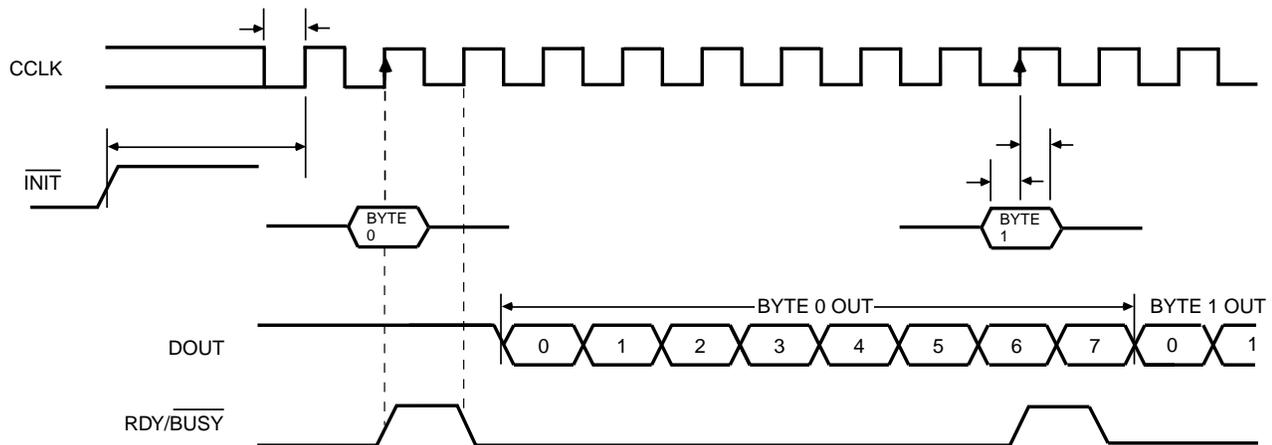
X6096

	Description	Symbol	Min	Max	Units
CCLK	$\overline{\text{INIT}}$ (High) setup time	T_{IC}	5		μs
	D0 - D7 setup time	T_{DC}	60		ns
	D0 - D7 hold time	T_{CD}	0		ns
	CCLK High time	T_{CCH}	50		ns
	CCLK Low time	T_{CCL}	60		ns
	CCLK Frequency	F_{CC}		8	MHz

Notes: 1. Peripheral Synchronous mode can be considered Slave Parallel mode. An external CCLK provides timing, clocking in the **first** data byte on the **second** rising edge of CCLK after $\overline{\text{INIT}}$ goes High. Subsequent data bytes are clocked in on every eighth consecutive rising edge of CCLK.

2. The RDY/$\overline{\text{BUSY}}$ line goes High for one CCLK period after data has been clocked in, although synchronous operation does not require such a response.

3. The pin name RDY/$\overline{\text{BUSY}}$ is a misnomer. In Synchronous Peripheral mode this is really an ACKNOWLEDGE signal.

4. Note that data starts to shift out serially on the DOUT pin 0.5 CCLK periods after it was loaded in parallel. Therefore, additional CCLK pulses are clearly required after the last byte has been loaded.

Figure 57: Synchronous Peripheral Mode Programming Switching Characteristics

6