



Welcome to [E-XFL.COM](https://www.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	256
Number of Gates	36000
Voltage - Supply	4.5V ~ 5.5V
Mounting Type	Surface Mount
Operating Temperature	-40°C ~ 100°C (TJ)
Package / Case	304-BFQFP Exposed Pad
Supplier Device Package	304-PQFP (40x40)
Purchase URL	https://www.e-xfl.com/product-detail/xilinx/xc4036ex-4hq304i

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.

Input Thresholds

The input thresholds of 5V devices can be globally configured for either TTL (1.2 V threshold) or CMOS (2.5 V threshold), just like XC2000 and XC3000 inputs. The two global adjustments of input threshold and output level are independent of each other. The XC4000XL family has an input threshold of 1.6V, compatible with both 3.3V CMOS and TTL levels.

Global Signal Access to Logic

There is additional access from global clocks to the F and G function generator inputs.

Configuration Pin Pull-Up Resistors

During configuration, these pins have weak pull-up resistors. For the most popular configuration mode, Slave Serial, the mode pins can thus be left unconnected. The three mode inputs can be individually configured with or without weak pull-up or pull-down resistors. A pull-down resistor value of 4.7 k Ω is recommended.

The three mode inputs can be individually configured with or without weak pull-up or pull-down resistors after configuration.

The PROGRAM input pin has a permanent weak pull-up.

Soft Start-up

Like the XC3000A, XC4000 Series devices have "Soft Start-up." When the configuration process is finished and the device starts up, the first activation of the outputs is automatically slew-rate limited. This feature avoids potential ground bounce when all outputs are turned on simultaneously. Immediately after start-up, the slew rate of the individual outputs is, as in the XC4000 family, determined by the individual configuration option.

XC4000 and XC4000A Compatibility

Existing XC4000 bitstreams can be used to configure an XC4000E device. XC4000A bitstreams must be recompiled for use with the XC4000E due to improved routing resources, although the devices are pin-for-pin compatible.

Additional Improvements in XC4000X Only

Increased Routing

New interconnect in the XC4000X includes twenty-two additional vertical lines in each column of CLBs and twelve new horizontal lines in each row of CLBs. The twelve "Quad Lines" in each CLB row and column include optional repowering buffers for maximum speed. Additional high-performance routing near the IOBs enhances pin flexibility.

Faster Input and Output

A fast, dedicated early clock sourced by global clock buffers is available for the IOBs. To ensure synchronization with the regular global clocks, a Fast Capture latch driven by the early clock is available. The input data can be initially loaded into the Fast Capture latch with the early clock, then transferred to the input flip-flop or latch with the low-skew global clock. A programmable delay on the input can be used to avoid hold-time requirements. See "IOB Input Signals" on page 20 for more information.

Latch Capability in CLBs

Storage elements in the XC4000X CLB can be configured as either flip-flops or latches. This capability makes the FPGA highly synthesis-compatible.

IOB Output MUX From Output Clock

A multiplexer in the IOB allows the output clock to select either the output data or the IOB clock enable as the output to the pad. Thus, two different data signals can share a single output pad, effectively doubling the number of device outputs without requiring a larger, more expensive package. This multiplexer can also be configured as an AND-gate to implement a very fast pin-to-pin path. See "IOB Output Signals" on page 23 for more information.

Additional Address Bits

Larger devices require more bits of configuration data. A daisy chain of several large XC4000X devices may require a PROM that cannot be addressed by the eighteen address bits supported in the XC4000E. The XC4000X Series therefore extends the addressing in Master Parallel configuration mode to 22 bits.

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 variables¹
- 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.

tions of the CLB, with the exception of the redefinition of the control signals. In 16x2 and 16x1 modes, the H' function generator can be used to implement Boolean functions of F', G', and D1, and the D flip-flops can latch the F', G', H', or D0 signals.

Single-Port Edge-Triggered Mode

Edge-triggered (synchronous) RAM simplifies timing requirements. XC4000 Series edge-triggered RAM timing operates like writing to a data register. Data and address are presented. The register is enabled for writing by a logic High on the write enable input, WE. Then a rising or falling clock edge loads the data into the register, as shown in **Figure 3**.




Figure 3: Edge-Triggered RAM Write Timing

Complex timing relationships between address, data, and write enable signals are not required, and the external write enable pulse becomes a simple clock enable. The active edge of WCLK latches the address, input data, and WE sig-

nals. An internal write pulse is generated that performs the write. See **Figure 4** and **Figure 5** for block diagrams of a CLB configured as 16x2 and 32x1 edge-triggered, single-port RAM.

The relationships between CLB pins and RAM inputs and outputs for single-port, edge-triggered mode are shown in **Table 5**.

The Write Clock input (WCLK) can be configured as active on either the rising edge (default) or the falling edge. It uses the same CLB pin (K) used to clock the CLB flip-flops, but it can be independently inverted. Consequently, the RAM output can optionally be registered within the same CLB either by the same clock edge as the RAM, or by the opposite edge of this clock. The sense of WCLK applies to both function generators in the CLB when both are configured as RAM.

The WE pin is active-High and is not invertible within the CLB.

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.

Table 5: Single-Port Edge-Triggered RAM Signals

RAM Signal	CLB Pin	Function
D	D0 or D1 (16x2, 16x1), D0 (32x1)	Data In
A[3:0]	F1-F4 or G1-G4	Address
A[4]	D1 (32x1)	Address
WE	WE	Write Enable
WCLK	K	Clock
SPO (Data Out)	F' or G'	Single Port Out (Data Out)

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.




Figure 7: 16x1 Edge-Triggered Dual-Port RAM

Figure 8 shows the write timing for level-sensitive, single-port RAM.

The relationships between CLB pins and RAM inputs and outputs for single-port level-sensitive mode are shown in Table 7.

Figure 9 and Figure 10 show block diagrams of a CLB configured as 16x2 and 32x1 level-sensitive, single-port RAM.

Initializing RAM at Configuration

Both RAM and ROM implementations of the XC4000 Series devices are initialized during configuration. The initial contents are defined via an INIT attribute or property

attached to the RAM or ROM symbol, as described in the schematic library guide. If not defined, all RAM contents are initialized to all zeros, by default.

RAM initialization occurs only during configuration. The RAM content is not affected by Global Set/Reset.

Table 7: Single-Port Level-Sensitive RAM Signals

RAM Signal	CLB Pin	Function
D	D0 or D1	Data In
A[3:0]	F1-F4 or G1-G4	Address
WE	WE	Write Enable
O	F' or G'	Data Out





Figure 8: Level-Sensitive RAM Write Timing



X6699

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

X9013

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



Figure 21: Open-Drain Buffers Implement a Wired-AND Function



x5994

Figure 25: High-Level Routing Diagram of XC4000 Series CLB (shaded arrows indicate XC4000X only)


Table 14: Routing per CLB in XC4000 Series Devices

	XC4000E		XC4000X	
	Vertical	Horizontal	Vertical	Horizontal
Singles	8	8	8	8
Doubles	4	4	4	4
Quads	0	0	12	12
Longlines	6	6	10	6
Direct Connects	0	0	2	2
Globals	4	0	8	0
Carry Logic	2	0	1	0
Total	24	18	45	32

Programmable Switch Matrices

The horizontal and vertical single- and double-length lines intersect at a box called a programmable switch matrix (PSM). Each switch matrix consists of programmable pass transistors used to establish connections between the lines (see Figure 26).

For example, a single-length signal entering on the right side of the switch matrix can be routed to a single-length line on the top, left, or bottom sides, or any combination thereof, if multiple branches are required. Similarly, a double-length signal can be routed to a double-length line on any or all of the other three edges of the programmable switch matrix.



X6600

Figure 26: Programmable Switch Matrix (PSM)

Single-Length Lines

Single-length lines provide the greatest interconnect flexibility and offer fast routing between adjacent blocks. There are eight vertical and eight horizontal single-length lines associated with each CLB. These lines connect the switching matrices that are located in every row and a column of CLBs.

Single-length lines are connected by way of the programmable switch matrices, as shown in Figure 28. Routing connectivity is shown in Figure 27.

Single-length lines incur a delay whenever they go through a switching matrix. Therefore, they are not suitable for routing signals for long distances. They are normally used to conduct signals within a localized area and to provide the branching for nets with fanout greater than one.




Figure 28: Single- and Double-Length Lines, with Programmable Switch Matrices (PSMs)

Double-Length Lines

The double-length lines consist of a grid of metal segments, each twice as long as the single-length lines: they run past two CLBs before entering a switch matrix. Double-length lines are grouped in pairs with the switch matrices staggered, so that each line goes through a switch matrix at every other row or column of CLBs (see [Figure 28](#)).

There are four vertical and four horizontal double-length lines associated with each CLB. These lines provide faster signal routing over intermediate distances, while retaining routing flexibility. Double-length lines are connected by way of the programmable switch matrices. Routing connectivity is shown in [Figure 27](#).

Quad Lines (XC4000X only)

XC4000X devices also include twelve vertical and twelve horizontal quad lines per CLB row and column. Quad lines are four times as long as the single-length lines. They are interconnected via buffered switch matrices (shown as diamonds in [Figure 27 on page 30](#)). Quad lines run past four CLBs before entering a buffered switch matrix. They are grouped in fours, with the buffered switch matrices staggered, so that each line goes through a buffered switch matrix at every fourth CLB location in that row or column. (See [Figure 29](#).)

The buffered switch matrixes have four pins, one on each edge. All of the pins are bidirectional. Any pin can drive any or all of the other pins.

Each buffered switch matrix contains one buffer and six pass transistors. It resembles the programmable switch matrix shown in [Figure 26](#), with the addition of a programmable buffer. There can be up to two independent inputs




Figure 29: Quad Lines (XC4000X only)

and up to two independent outputs. Only one of the independent inputs can be buffered.

The place and route software automatically uses the timing requirements of the design to determine whether or not a quad line signal should be buffered. A heavily loaded signal is typically buffered, while a lightly loaded one is not. One scenario is to alternate buffers and pass transistors. This allows both vertical and horizontal quad lines to be buffered at alternating buffered switch matrices.

Due to the buffered switch matrices, quad lines are very fast. They provide the fastest available method of routing heavily loaded signals for long distances across the device.

Longlines

Longlines form a grid of metal interconnect segments that run the entire length or width of the array. Longlines are intended for high fan-out, time-critical signal nets, or nets that are distributed over long distances. In XC4000X devices, quad lines are preferred for critical nets, because the buffered switch matrices make them faster for high fan-out nets.

Two horizontal longlines per CLB can be driven by 3-state or open-drain drivers (TBUFs). They can therefore implement unidirectional or bidirectional buses, wide multiplexers, or wired-AND functions. (See [“Three-State Buffers” on page 26](#) for more details.)

Each horizontal longline driven by TBUFs has either two (XC4000E) or eight (XC4000X) pull-up resistors. To activate these resistors, attach a PULLUP symbol to the long-line net. The software automatically activates the appropriate number of pull-ups. There is also a weak keeper at each end of these two horizontal longlines. This

circuit prevents undefined floating levels. However, it is overridden by any driver, even a pull-up resistor.

Each XC4000E longline has a programmable splitter switch at its center, as does each XC4000X longline driven by TBUFs. This switch can separate the line into two independent routing channels, each running half the width or height of the array.

Each XC4000X longline not driven by TBUFs has a buffered programmable splitter switch at the 1/4, 1/2, and 3/4 points of the array. Due to the buffering, XC4000X longline performance does not deteriorate with the larger array sizes. If the longline is split, the resulting partial longlines are independent.

Routing connectivity of the longlines is shown in [Figure 27 on page 30](#).

Direct Interconnect (XC4000X only)

The XC4000X offers two direct, efficient and fast connections between adjacent CLBs. These nets facilitate a data flow from the left to the right side of the device, or from the top to the bottom, as shown in [Figure 30](#). Signals routed on the direct interconnect exhibit minimum interconnect propagation delay and use no general routing resources.

The direct interconnect is also present between CLBs and adjacent IOBs. Each IOB on the left and top device edges has a direct path to the nearest CLB. Each CLB on the right and bottom edges of the array has a direct path to the nearest two IOBs, since there are two IOBs for each row or column of CLBs.

The place and route software uses direct interconnect whenever possible, to maximize routing resources and minimize interconnect delays.




Figure 30: XC4000X Direct Interconnect

I/O Routing

XC4000 Series devices have additional routing around the IOB ring. This routing is called a VersaRing. The VersaRing facilitates pin-swapping and redesign without affecting board layout. Included are eight double-length lines spanning two CLBs (four IOBs), and four longlines. Global lines and Wide Edge Decoder lines are provided. XC4000X devices also include eight octal lines.

A high-level diagram of the VersaRing is shown in [Figure 31](#). The shaded arrows represent routing present only in XC4000X devices.

[Figure 33 on page 34](#) is a detailed diagram of the XC4000E and XC4000X VersaRing. The area shown includes two IOBs. There are two IOBs per CLB row or column, therefore this diagram corresponds to the CLB routing diagram shown in [Figure 27 on page 30](#). The shaded areas represent routing and routing connections present only in XC4000X devices.

Octal I/O Routing (XC4000X only)

Between the XC4000X CLB array and the pad ring, eight interconnect tracks provide for versatility in pin assignment and fixed pinout flexibility. (See [Figure 32 on page 33](#).)

These routing tracks are called octals, because they can be broken every eight CLBs (sixteen IOBs) by a programmable buffer that also functions as a splitter switch. The buffers are staggered, so each line goes through a buffer at every eighth CLB location around the device edge.

The octal lines bend around the corners of the device. The lines cross at the corners in such a way that the segment most recently buffered before the turn has the farthest distance to travel before the next buffer, as shown in [Figure 32](#).

The top and bottom Global Early buffers are about 1 ns slower clock to out than the left and right Global Early buffers.

The Global Early buffers can be driven by either semi-dedicated pads or internal logic. They share pads with the Global Low-Skew buffers, so a single net can drive both global buffers, as described above.

To use a Global Early buffer, place a BUFGE 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 BUFGE be placed in one of the two Global Early buffers on the top edge of the device, or a LOC=TR to indicate the Global Early buffer on the top edge of the device, on the right.

Power Distribution

Power for the FPGA is distributed through a grid to achieve high noise immunity and isolation between logic and I/O. Inside the FPGA, a dedicated Vcc and Ground ring surrounding the logic array provides power to the I/O drivers, as shown in [Figure 39](#). An independent matrix of Vcc and Ground lines supplies the interior logic of the device.

This power distribution grid provides a stable supply and ground for all internal logic, providing the external package power pins are all connected and appropriately de-coupled. Typically, a 0.1 μ F capacitor connected between each Vcc pin and the board's Ground plane will provide adequate de-coupling.

Output buffers capable of driving/sinking the specified 12 mA loads under specified worst-case conditions may be capable of driving/sinking up to 10 times as much current under best case conditions.

Noise can be reduced by minimizing external load capacitance and reducing simultaneous output transitions in the same direction. It may also be beneficial to locate heavily loaded output buffers near the Ground pads. The I/O Block output buffers have a slew-rate limited mode (default) which should be used where output rise and fall times are not speed-critical.



Figure 39: XC4000 Series Power Distribution

Pin Descriptions

There are three types of pins in the XC4000 Series devices:

- Permanently dedicated pins
- User I/O pins that can have special functions
- Unrestricted user-programmable I/O pins.

Before and during configuration, all outputs not used for the configuration process are 3-stated with a 50 k Ω - 100 k Ω pull-up resistor.

After configuration, if an IOB is unused it is configured as an input with a 50 k Ω - 100 k Ω pull-up resistor.

XC4000 Series devices have no dedicated Reset input. Any user I/O can be configured to drive the Global Set/Reset net, GSR. See ["Global Set/Reset" on page 11](#) for more information on GSR.

XC4000 Series devices have no Powerdown control input, as the XC3000 and XC2000 families do. The XC3000/XC2000 Powerdown control also 3-stated all of the device

I/O pins. For XC4000 Series devices, use the global 3-state net, GTS, instead. This net 3-states all outputs, but does not place the device in low-power mode. See ["IOB Output Signals" on page 23](#) for more information on GTS.

Device pins for XC4000 Series devices are described in [Table 16](#). Pin functions during configuration for each of the seven configuration modes are summarized in [Table 22 on page 58](#), in the "Configuration Timing" section.

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



X9024

Figure 47: Start-up Timing

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 [$2^{24} * \text{CCLK period}$] — 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 k Ω - 100 k Ω 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 $\overline{\text{INIT}}$ held Low (or drive the $\overline{\text{PROGRAM}}$ pin Low for more than 300 ns followed by a High while holding $\overline{\text{INIT}}$ Low). Holding $\overline{\text{INIT}}$ 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 $\overline{\text{INIT}}$ Low.
- Issue the CONFIG command to the TMS input
- Wait for $\overline{\text{INIT}}$ 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.

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.


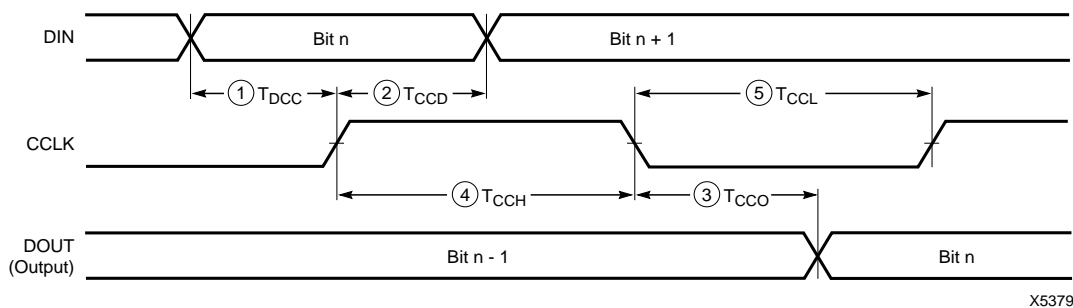


Figure 51: Master/Slave Serial Mode Circuit Diagram



	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 $\overline{\text{INIT}}$ pins of all daisy-chained FPGAs are High.

Figure 52: Slave Serial Mode Programming Switching Characteristics

Master Parallel Modes

In the two Master Parallel modes, the lead FPGA directly addresses an industry-standard byte-wide EPROM, and accepts eight data bits just before incrementing or decrementing the address outputs.

The eight data bits are serialized in the lead FPGA, which then presents the preamble data—and all data that overflows the lead device—on its DOUT pin. There is an internal delay of 1.5 CCLK periods, after the rising CCLK edge that accepts a byte of data (and also changes the EPROM address) until the falling CCLK edge that makes the LSB (D0) of this byte appear at DOUT. This 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.

The PROM address pins can be incremented or decremented, depending on the mode pin settings. This option allows the FPGA to share the PROM with a wide variety of microprocessors and micro controllers. Some processors must boot from the bottom of memory (all zeros) while others must boot from the top. The FPGA is flexible and can load its configuration bitstream from either end of the memory.

Master Parallel Up mode is selected by a <100> on the mode pins (M2, M1, M0). The EPROM addresses start at 00000 and increment.

Master Parallel Down mode is selected by a <110> on the mode pins. The EPROM addresses start at 3FFFF and decrement.

Additional Address lines in XC4000 devices

The XC4000X devices have additional address lines (A18-A21) allowing the additional address space required to daisy-chain several large devices.

The extra address lines are programmable in XC4000EX devices. By default these address lines are not activated. In the default mode, the devices are compatible with existing XC4000 and XC4000E products. If desired, the extra address lines can be used by specifying the address lines option in bitgen as 22 (bitgen -g AddressLines:22). The lines (A18-A21) are driven when a master device detects, via the bitstream, that it should be using all 22 address lines. Because these pins will initially be pulled high by internal pull-ups, designers using Master Parallel Up mode should use external pull down resistors on pins A18-A21. If Master Parallel Down mode is used external resistors are not necessary.

All 22 address lines are always active in Master Parallel modes with XC4000XL devices. The additional address lines behave identically to the lower order address lines. If the Address Lines option in bitgen is set to 18, it will be ignored by the XC4000XL device.

The additional address lines (A18-A21) are not available in the PC84 package.

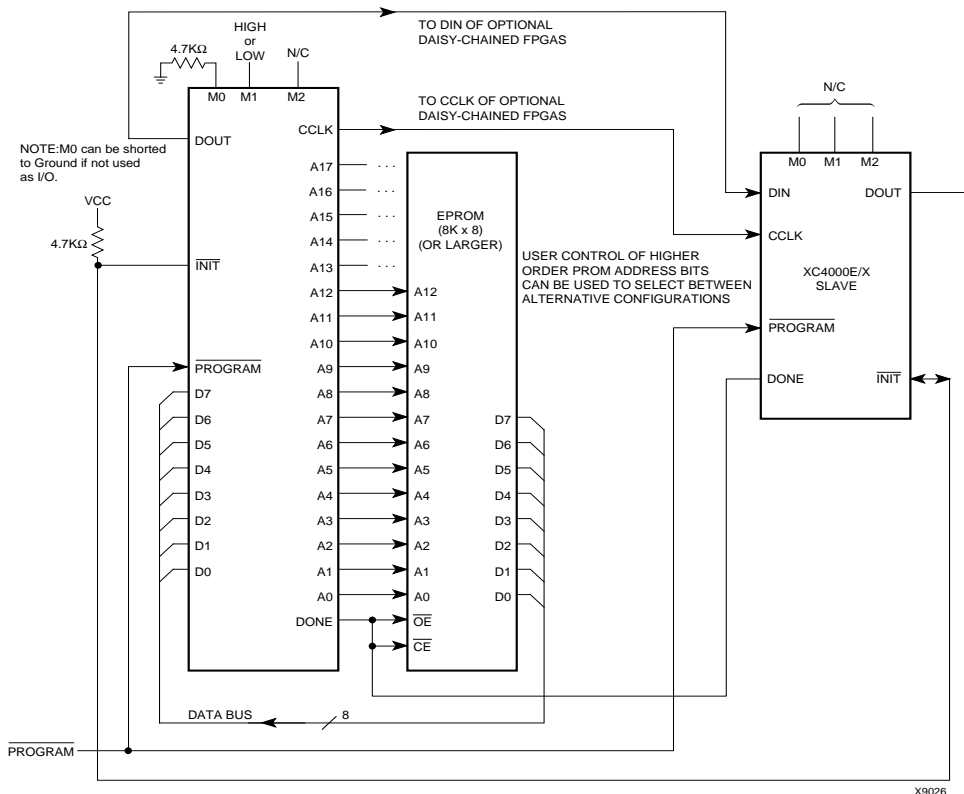



Figure 54: Master Parallel Mode Circuit Diagram




	Description	Symbol	Min	Max	Units
RCLK	Delay to Address valid	1 T_{RAC}	0	200	ns
	Data setup time	2 T_{DRC}	60		ns
	Data hold time	3 T_{RCD}	0		ns

Notes: 1. At power-up, V_{cc} must rise from 2.0 V to V_{cc} min in less than 25 ms, otherwise delay configuration by pulling PROGRAM Low until V_{cc} is valid.
2. The first Data byte is loaded and CCLK starts at the end of the first \overline{RCLK} active cycle (rising edge).
This timing diagram shows that the EPROM requirements are extremely relaxed. EPROM access time can be longer than 500 ns. EPROM data output has no hold-time requirements.

Figure 55: Master Parallel Mode Programming Switching Characteristics

Configuration Switching Characteristics



X1532

Master Modes (XC4000E/EX)

Description		Symbol	Min	Max	Units
Power-On Reset	M0 = High	T_{POR}	10	40	ms
	M0 = Low	T_{POR}	40	130	ms
Program Latency		T_{PI}	30	200	μ s per CLB column
CCLK (output) Delay		T_{ICCK}	40	250	μ s
CCLK (output) Period, slow		T_{CCLK}	640	2000	ns
CCLK (output) Period, fast		T_{CCLK}	80	250	ns

Master Modes (XC4000XL)

Description		Symbol	Min	Max	Units
Power-On Reset	M0 = High	T_{POR}	10	40	ms
	M0 = Low	T_{POR}	40	130	ms
Program Latency		T_{PI}	30	200	μ s per CLB column
CCLK (output) Delay		T_{ICCK}	40	250	μ s
CCLK (output) Period, slow		T_{CCLK}	540	1600	ns
CCLK (output) Period, fast		T_{CCLK}	67	200	ns

Slave and Peripheral Modes (All)

Description		Symbol	Min	Max	Units
Power-On Reset		T_{POR}	10	33	ms
Program Latency		T_{PI}	30	200	μ s per CLB column
CCLK (input) Delay (required)		T_{ICCK}	4		μ s
CCLK (input) Period (required)		T_{CCLK}	100		ns