



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	100
Number of Logic Elements/Cells	238
Total RAM Bits	3200
Number of I/O	80
Number of Gates	3000
Voltage - Supply	4.5V ~ 5.5V
Mounting Type	Through Hole
Operating Temperature	-40°C ~ 100°C (TJ)
Package / Case	120-BCPGA
Supplier Device Package	120-CPGA (34.54x34.54)
Purchase URL	https://www.e-xfl.com/product-detail/xilinx/xc4003e-4pg120i

XC4000E and XC4000X Series Compared to the XC4000

For readers already familiar with the XC4000 family of Xilinx Field Programmable Gate Arrays, the major new features in the XC4000 Series devices are listed in this section. The biggest advantages of XC4000E and XC4000X devices are significantly increased system speed, greater capacity, and new architectural features, particularly Select-RAM memory. The XC4000X devices also offer many new routing features, including special high-speed clock buffers that can be used to capture input data with minimal delay.

Any XC4000E device is pinout- and bitstream-compatible with the corresponding XC4000 device. An existing XC4000 bitstream can be used to program an XC4000E device. However, since the XC4000E includes many new features, an XC4000E bitstream cannot be loaded into an XC4000 device.

XC4000X Series devices are not bitstream-compatible with equivalent array size devices in the XC4000 or XC4000E families. However, equivalent array size devices, such as the XC4025, XC4025E, XC4028EX, and XC4028XL, are pinout-compatible.

Improvements in XC4000E and XC4000X

Increased System Speed

XC4000E and XC4000X devices can run at synchronous system clock rates of up to 80 MHz, and internal performance can exceed 150 MHz. This increase in performance over the previous families stems from improvements in both device processing and system architecture. XC4000 Series devices use a sub-micron multi-layer metal process. In addition, many architectural improvements have been made, as described below.

The XC4000XL family is a high performance 3.3V family based on 0.35 μ SRAM technology and supports system speeds to 80 MHz.

PCI Compliance

XC4000 Series -2 and faster speed grades are fully PCI compliant. XC4000E and XC4000X devices can be used to implement a one-chip PCI solution.

Carry Logic

The speed of the carry logic chain has increased dramatically. Some parameters, such as the delay on the carry chain through a single CLB (T_{BYP}), have improved by as

much as 50% from XC4000 values. See [“Fast Carry Logic” on page 18](#) for more information.

Select-RAM Memory: Edge-Triggered, Synchronous RAM Modes

The RAM in any CLB can be configured for synchronous, edge-triggered, write operation. The read operation is not affected by this change to an edge-triggered write.

Dual-Port RAM

A separate option converts the 16x2 RAM in any CLB into a 16x1 dual-port RAM with simultaneous Read/Write.

The function generators in each CLB can be configured as either level-sensitive (asynchronous) single-port RAM, edge-triggered (synchronous) single-port RAM, edge-triggered (synchronous) dual-port RAM, or as combinatorial logic.

Configurable RAM Content

The RAM content can now be loaded at configuration time, so that the RAM starts up with user-defined data.

H Function Generator

In current XC4000 Series devices, the H function generator is more versatile than in the original XC4000. Its inputs can come not only from the F and G function generators but also from up to three of the four control input lines. The H function generator can thus be totally or partially independent of the other two function generators, increasing the maximum capacity of the device.

IOB Clock Enable

The two flip-flops in each IOB have a common clock enable input, which through configuration can be activated individually for the input or output flip-flop or both. This clock enable operates exactly like the EC pin on the XC4000 CLB. This new feature makes the IOBs more versatile, and avoids the need for clock gating.

Output Drivers

The output pull-up structure defaults to a TTL-like totem-pole. This driver is an n-channel pull-up transistor, pulling to a voltage one transistor threshold below V_{cc} , just like the XC4000 family outputs. Alternatively, XC4000 Series devices can be globally configured with CMOS outputs, with p-channel pull-up transistors pulling to V_{cc} . Also, the configurable pull-up resistor in the XC4000 Series is a p-channel transistor that pulls to V_{cc} , whereas in the original XC4000 family it is an n-channel transistor that pulls to a voltage one transistor threshold below V_{cc} .

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.

Set/Reset

An asynchronous storage element input (SR) can be configured as either set or reset. This configuration option determines the state in which each flip-flop becomes operational after configuration. It also determines the effect of a Global Set/Reset pulse during normal operation, and the effect of a pulse on the SR pin of the CLB. All three set/reset functions for any single flip-flop are controlled by the same configuration data bit.

The set/reset state can be independently specified for each flip-flop. This input can also be independently disabled for either flip-flop.

The set/reset state is specified by using the INIT attribute, or by placing the appropriate set or reset flip-flop library symbol.

SR is active High. It is not invertible within the CLB.

Global Set/Reset

A separate Global Set/Reset line (not shown in Figure 1) sets or clears each storage element during power-up, re-configuration, or when a dedicated Reset net is driven active. This global net (GSR) does not compete with other routing resources; it uses a dedicated distribution network.

Each flip-flop is configured as either globally set or reset in the same way that the local set/reset (SR) is specified. Therefore, if a flip-flop is set by SR, it is also set by GSR. Similarly, a reset flip-flop is reset by both SR and GSR.




Figure 2: Schematic Symbols for Global Set/Reset

GSR can be driven from any user-programmable pin as a global reset input. To use this global net, place an input pad and input buffer in the schematic or HDL code, driving the GSR pin of the STARTUP symbol. (See Figure 2.) A specific pin location can be assigned to this input using a LOC attribute or property, just as with any other user-programmable pad. An inverter can optionally be inserted after the input buffer to invert the sense of the Global Set/Reset signal.

Alternatively, GSR can be driven from any internal node.

Data Inputs and Outputs

The source of a storage element data input is programmable. It is driven by any of the functions F', G', and H', or by the Direct In (DIN) block input. The flip-flops or latches drive the XQ and YQ CLB outputs.

Two fast feed-through paths are available, as shown in Figure 1. A two-to-one multiplexer on each of the XQ and YQ outputs selects between a storage element output and any of the control inputs. This bypass is sometimes used by the automated router to repower internal signals.

Control Signals

Multiplexers in the CLB map the four control inputs (C1 - C4 in Figure 1) into the four internal control signals (H1, DIN/H2, SR/H0, and EC). Any of these inputs can drive any of the four internal control signals.

When the logic function is enabled, the four inputs are:

- EC — Enable Clock
- SR/H0 — Asynchronous Set/Reset or H function generator Input 0
- DIN/H2 — Direct In or H function generator Input 2
- H1 — H function generator Input 1.

When the memory function is enabled, the four inputs are:

- EC — Enable Clock
- WE — Write Enable
- D0 — Data Input to F and/or G function generator
- D1 — Data input to G function generator (16x1 and 16x2 modes) or 5th Address bit (32x1 mode).

Using FPGA Flip-Flops and Latches

The abundance of flip-flops in the XC4000 Series invites pipelined designs. This is a powerful way of increasing performance by breaking the function into smaller subfunctions and executing them in parallel, passing on the results through pipeline flip-flops. This method should be seriously considered wherever throughput is more important than latency.

To include a CLB flip-flop, place the appropriate library symbol. For example, FDCE is a D-type flip-flop with clock enable and asynchronous clear. The corresponding latch symbol (for the XC4000X only) is called LDCE.

In XC4000 Series devices, the flip flops can be used as registers or shift registers without blocking the function generators from performing a different, perhaps unrelated task. This ability increases the functional capacity of the devices.

The CLB setup time is specified between the function generator inputs and the clock input K. Therefore, the specified CLB flip-flop setup time includes the delay through the function generator.

Using Function Generators as RAM

Optional modes for each CLB make the memory look-up tables in the F' and G' function generators usable as an array of Read/Write memory cells. Available modes are level-sensitive (similar to the XC4000/A/H families), edge-triggered, and dual-port edge-triggered. Depending on the selected mode, a single CLB can be configured as either a 16x2, 32x1, or 16x1 bit array.

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-Sens itive	Edge-Trigg ered	Dual-Port Edge-Trigg ered
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)

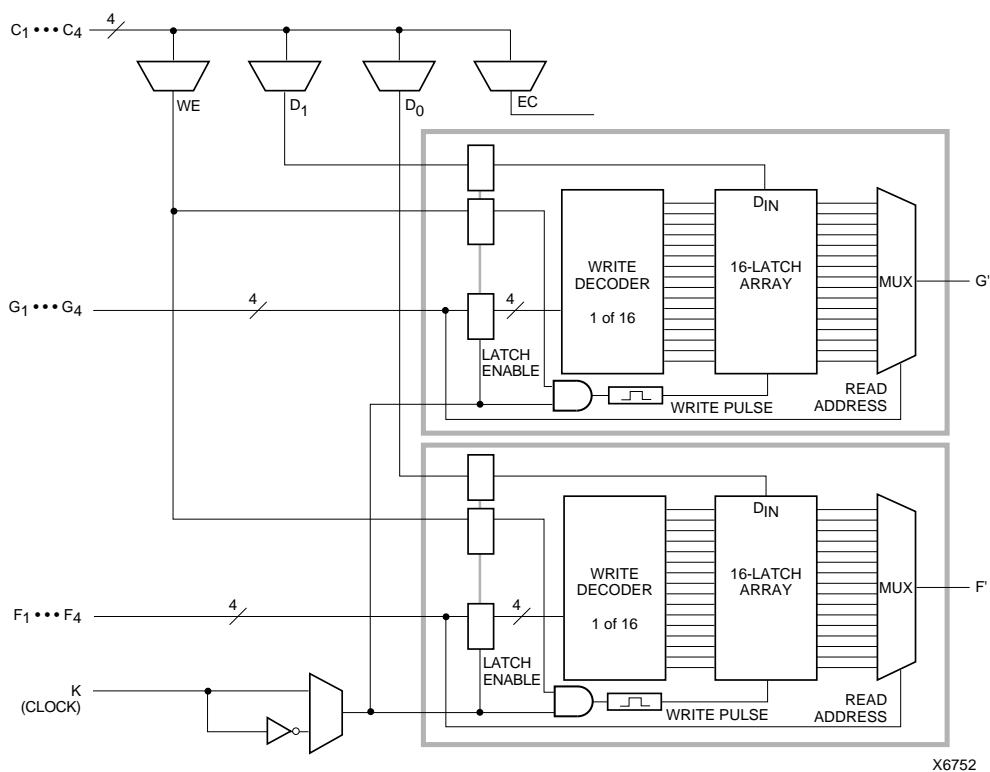
RAM Inputs and Outputs

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-



X6752

Figure 4: 16x2 (or 16x1) Edge-Triggered Single-Port RAM



X6754

Figure 5: 32x1 Edge-Triggered Single-Port RAM (F and G addresses are identical)

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.

Fast Carry Logic

Each CLB F and G function generator contains dedicated arithmetic logic for the fast generation of carry and borrow signals. This extra output is passed on to the function generator in the adjacent CLB. The carry chain is independent of normal routing resources.

Dedicated fast carry logic greatly increases the efficiency and performance of adders, subtractors, accumulators, comparators and counters. It also opens the door to many new applications involving arithmetic operation, where the previous generations of FPGAs were not fast enough or too inefficient. High-speed address offset calculations in micro-processor or graphics systems, and high-speed addition in digital signal processing are two typical applications.

The two 4-input function generators can be configured as a 2-bit adder with built-in hidden carry that can be expanded to any length. This dedicated carry circuitry is so fast and efficient that conventional speed-up methods like carry generate/propagate are meaningless even at the 16-bit level, and of marginal benefit at the 32-bit level.

This fast carry logic is one of the more significant features of the XC4000 Series, speeding up arithmetic and counting into the 70 MHz range.

The carry chain in XC4000E devices can run either up or down. At the top and bottom of the columns where there are no CLBs above or below, the carry is propagated to the right. (See Figure 11.) In order to improve speed in the high-capacity XC4000X devices, which can potentially have very long carry chains, the carry chain travels upward only, as shown in Figure 12. Additionally, standard interconnect can be used to route a carry signal in the downward direction.


Figure 13 on page 19 shows an XC4000E CLB with dedicated fast carry logic. The carry logic in the XC4000X is similar, except that COUT exits at the top only, and the signal CINDOWN does not exist. As shown in Figure 13, the carry logic shares operand and control inputs with the function generators. The carry outputs connect to the function generators, where they are combined with the operands to form the sums.

Figure 14 on page 20 shows the details of the carry logic for the XC4000E. This diagram shows the contents of the box labeled "CARRY LOGIC" in Figure 13. The XC4000X carry logic is very similar, but a multiplexer on the pass-through carry chain has been eliminated to reduce delay. Additionally, in the XC4000X the multiplexer on the G4 path has a memory-programmable 0 input, which permits G4 to directly connect to COUT. G4 thus becomes an additional high-speed initialization path for carry-in.

The dedicated carry logic is discussed in detail in Xilinx document XAPP 013: "Using the Dedicated Carry Logic in


XC4000." This discussion also applies to XC4000E devices, and to XC4000X devices when the minor logic changes are taken into account.

The fast carry logic can be accessed by placing special library symbols, or by using Xilinx Relationally Placed Macros (RPMs) that already include these symbols.



X6687

Figure 11: Available XC4000E Carry Propagation Paths



X6610

Figure 12: Available XC4000X Carry Propagation Paths (dotted lines use general interconnect)

Table 8: Supported Sources for XC4000 Series Device Inputs

Source	XC4000E/EX Series Inputs		XC4000XL Series Inputs
	5 V, TTL	5 V, CMOS	3.3 V CMOS
Any device, V _{CC} = 3.3 V, CMOS outputs	✓	Unreliable Data	✓
XC4000 Series, V _{CC} = 5 V, TTL outputs	✓		✓
Any device, V _{CC} = 5 V, TTL outputs (V _{oh} ≤ 3.7 V)	✓		✓
Any device, V _{CC} = 5 V, CMOS outputs	✓	✓	✓

XC4000XL 5-Volt Tolerant I/Os

The I/Os on the XC4000XL are fully 5-volt tolerant even though the V_{CC} is 3.3 volts. This allows 5 V signals to directly connect to the XC4000XL inputs without damage, as shown in Table 8. In addition, the 3.3 volt V_{CC} can be applied before or after 5 volt signals are applied to the I/Os. This makes the XC4000XL immune to power supply sequencing problems.

Registered Inputs

The I1 and I2 signals that exit the block can each carry either the direct or registered input signal.

The input and output storage elements in each IOB have a common clock enable input, which, through configuration, can be activated individually for the input or output flip-flop, or both. This clock enable operates exactly like the EC pin on the XC4000 Series CLB. It cannot be inverted within the IOB.

The storage element behavior is shown in Table 9.

Table 9: Input Register Functionality (active rising edge is shown)

Mode	Clock	Clock Enable	D	Q
Power-Up or GSR	X	X	X	SR
Flip-Flop		1*	D	D
	0	X	X	Q
Latch	1	1*	X	Q
	0	1*	D	D
Both	X	0	X	Q

Legend:

X 

SR

0*

1*

Don't care
Rising edge

Set or Reset value. Reset is default.

Input is Low or unconnected (default value)

Input is High or unconnected (default value)

Optional Delay Guarantees Zero Hold Time

The data input to the register can optionally be delayed by several nanoseconds. With the delay enabled, the setup time of the input flip-flop is increased so that normal clock routing does not result in a positive hold-time requirement. A positive hold time requirement can lead to unreliable, temperature- or processing-dependent operation.

The input flip-flop setup time is defined between the data measured at the device I/O pin and the clock input at the IOB (not at the clock pin). Any routing delay from the device clock pin to the clock input of the IOB must, therefore, be subtracted from this setup time to arrive at the real setup time requirement relative to the device pins. A short specified setup time might, therefore, result in a negative setup time at the device pins, i.e., a positive hold-time requirement.

When a delay is inserted on the data line, more clock delay can be tolerated without causing a positive hold-time requirement. Sufficient delay eliminates the possibility of a data hold-time requirement at the external pin. The maximum delay is therefore inserted as the default.

The XC4000E IOB has a one-tap delay element: either the delay is inserted (default), or it is not. The delay guarantees a zero hold time with respect to clocks routed through any of the XC4000E global clock buffers. (See "Global Nets and Buffers (XC4000E only)" on page 35 for a description of the global clock buffers in the XC4000E.) For a shorter input register setup time, with non-zero hold, attach a NODELAY attribute or property to the flip-flop.

The XC4000X IOB has a two-tap delay element, with choices of a full delay, a partial delay, or no delay. The attributes or properties used to select the desired delay are shown in Table 10. The choices are no added attribute, MEDDELAY, and NODELAY. The default setting, with no added attribute, ensures no hold time with respect to any of the XC4000X clock buffers, including the Global Low-Skew buffers. MEDDELAY ensures no hold time with respect to the Global Early buffers. Inputs with NODELAY may have a positive hold time with respect to all clock buffers. For a description of each of these buffers, see "Global Nets and Buffers (XC4000X only)" on page 37.

Table 10: XC4000X IOB Input Delay Element

Value	When to Use
full delay (default, no attribute added)	Zero Hold with respect to Global Low-Skew Buffer, Global Early Buffer
MEDDELAY	Zero Hold with respect to Global Early Buffer
NODELAY	Short Setup, positive Hold time

Output Multiplexer/2-Input Function Generator (XC4000X only)

As shown in [Figure 16 on page 21](#), the output path in the XC4000X IOB contains an additional multiplexer not available in the XC4000E IOB. The multiplexer can also be configured as a 2-input function generator, implementing a pass-gate, AND-gate, OR-gate, or XOR-gate, with 0, 1, or 2 inverted inputs. The logic used to implement these functions is shown in the upper gray area of [Figure 16](#).

When configured as a multiplexer, this feature allows two output signals to time-share the same output pad; effectively doubling the number of device outputs without requiring a larger, more expensive package.

When the MUX is configured as a 2-input function generator, logic can be implemented within the IOB itself. Combined with a Global Early buffer, this arrangement allows very high-speed gating of a single signal. For example, a wide decoder can be implemented in CLBs, and its output gated with a Read or Write Strobe Driven by a BUFGE buffer, as shown in [Figure 19](#). The critical-path pin-to-pin delay of this circuit is less than 6 nanoseconds.

As shown in [Figure 16](#), the IOB input pins Out, Output Clock, and Clock Enable have different delays and different flexibilities regarding polarity. Additionally, Output Clock sources are more limited than the other inputs. Therefore, the Xilinx software does not move logic into the IOB function generators unless explicitly directed to do so.

The user can specify that the IOB function generator be used, by placing special library symbols beginning with the letter "O." For example, a 2-input AND-gate in the IOB function generator is called OAND2. Use the symbol input pin labelled "F" for the signal on the critical path. This signal is placed on the OK pin — the IOB input with the shortest delay to the function generator. Two examples are shown in [Figure 20](#).



Figure 19: Fast Pin-to-Pin Path in XC4000X



Figure 20: AND & MUX Symbols in XC4000X IOB

Other IOB Options

There are a number of other programmable options in the XC4000 Series IOB.

Pull-up and Pull-down Resistors

Programmable pull-up and pull-down resistors are useful for tying unused pins to Vcc or Ground to minimize power consumption and reduce noise sensitivity. The configurable pull-up resistor is a p-channel transistor that pulls to Vcc. The configurable pull-down resistor is an n-channel transistor that pulls to Ground.

The value of these resistors is 50 kΩ – 100 kΩ. This high value makes them unsuitable as wired-AND pull-up resistors.

The pull-up resistors for most user-programmable IOBs are active during the configuration process. See [Table 22 on page 58](#) for a list of pins with pull-ups active before and during configuration.

After configuration, voltage levels of unused pads, bonded or un-bonded, must be valid logic levels, to reduce noise sensitivity and avoid excess current. Therefore, by default, unused pads are configured with the internal pull-up resistor active. Alternatively, they can be individually configured with the pull-down resistor, or as a driven output, or to be driven by an external source. To activate the internal pull-up, attach the PULLUP library component to the net attached to the pad. To activate the internal pull-down, attach the PULLDOWN library component to the net attached to the pad.

Independent Clocks

Separate clock signals are provided for the input and output flip-flops. The clock can be independently inverted for each flip-flop within the IOB, generating either falling-edge or rising-edge triggered flip-flops. The clock inputs for each IOB are independent, except that in the XC4000X, the Fast Capture latch shares an IOB input with the output clock pin.

Early Clock for IOBs (XC4000X only)

Special early clocks are available for IOBs. These clocks are sourced by the same sources as the Global Low-Skew buffers, but are separately buffered. They have fewer loads and therefore less delay. The early clock can drive either the IOB output clock or the IOB input clock, or both. The early clock allows fast capture of input data, and fast clock-to-output on output data. The Global Early buffers that drive these clocks are described in ["Global Nets and Buffers \(XC4000X only\)" on page 37](#).

Global Set/Reset

As with the CLB registers, the Global Set/Reset signal (GSR) can be used to set or clear the input and output registers, depending on the value of the INIT attribute or property. The two flip-flops can be individually configured to set

IOB inputs and outputs interface with the octal lines via the single-length interconnect lines. Single-length lines are also used for communication between the octals and double-length lines, quads, and longlines within the CLB array.

Segmentation into buffered octals was found to be optimal for distributing signals over long distances around the device.

Global Nets and Buffers

Both the XC4000E and the XC4000X have dedicated global networks. These networks are designed to distribute clocks and other high fanout control signals throughout the devices with minimal skew. The global buffers are described in detail in the following sections. The text descriptions and diagrams are summarized in [Table 15](#). The table shows which CLB and IOB clock pins can be sourced by which global buffers.

In both XC4000E and XC4000X devices, placement of a library symbol called BUFG results in the software choosing the appropriate clock buffer, based on the timing requirements of the design. The detailed information in these sections is included only for reference.

Global Nets and Buffers (XC4000E only)

Four vertical longlines in each CLB column are driven exclusively by special global buffers. These longlines are in addition to the vertical longlines used for standard interconnect. The four global lines can be driven by either of two types of global buffers. The clock pins of every CLB and IOB can also be sourced from local interconnect.

Two different types of clock buffers are available in the XC4000E:

- Primary Global Buffers (BUFGP)
- Secondary Global Buffers (BUFGS)

Four Primary Global buffers offer the shortest delay and negligible skew. Four Secondary Global buffers have slightly longer delay and slightly more skew due to potentially heavier loading, but offer greater flexibility when used to drive non-clock CLB inputs.

The Primary Global buffers must be driven by the semi-dedicated pads. The Secondary Global buffers can be sourced by either semi-dedicated pads or internal nets.

Each CLB column has four dedicated vertical Global lines. Each of these lines can be accessed by one particular Primary Global buffer, or by any of the Secondary Global buffers, as shown in [Figure 34](#). Each corner of the device has one Primary buffer and one Secondary buffer.

IOBs along the left and right edges have four vertical global longlines. Top and bottom IOBs can be clocked from the global lines in the adjacent CLB column.

A global buffer should be specified for all timing-sensitive global signal distribution. To use a global buffer, place a BUFGP (primary buffer), BUFGS (secondary buffer), or BUFG (either primary or secondary buffer) 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=L attribute or property to a BUFGS symbol to direct that a buffer be placed in one of the two Secondary Global buffers on the left edge of the device, or a LOC=BL to indicate the Secondary Global buffer on the bottom edge of the device, on the left.

Table 15: Clock Pin Access

	XC4000E		XC4000X			Local Inter-connect
	BUFGP	BUFGS	BUFGLS	L & R BUFGE	T & B BUFGE	
All CLBs in Quadrant	√	√	√	√	√	√
All CLBs in Device	√	√	√			√
IOBs on Adjacent Vertical Half Edge	√	√	√	√	√	√
IOBs on Adjacent Vertical Full Edge	√	√	√	√		√
IOBs on Adjacent Horizontal Half Edge (Direct)				√		√
IOBs on Adjacent Horizontal Half Edge (through CLB globals)	√	√	√	√	√	√
IOBs on Adjacent Horizontal Full Edge (through CLB globals)	√	√	√			√

L = Left, R = Right, T = Top, B = Bottom




Figure 36: Any BUFGLS (GCK1 - GCK8) Can Drive Any or All Clock Inputs on the Device

Global Early Buffers

Each corner of the XC4000X device has two Global Early buffers. The primary purpose of the Global Early buffers is to provide an earlier clock access than the potentially heavily-loaded Global Low-Skew buffers. A clock source applied to both buffers will result in the Global Early clock edge occurring several nanoseconds earlier than the Global Low-Skew buffer clock edge, due to the lighter loading.

Global Early buffers also facilitate the fast capture of device inputs, using the Fast Capture latches described in “**IOB Input Signals**” on page 20. For Fast Capture, take a single clock signal, and route it through both a Global Early buffer and a Global Low-Skew buffer. (The two buffers share an input pad.) Use the Global Early buffer to clock the Fast Capture latch, and the Global Low-Skew buffer to clock the normal input flip-flop or latch, as shown in **Figure 17** on page 23.

The Global Early buffers can also be used to provide a fast Clock-to-Out on device output pins. However, an early clock in the output flip-flop IOB must be taken into consideration when calculating the internal clock speed for the design.

The Global Early buffers at the left and right edges of the chip have slightly different capabilities than the ones at the top and bottom. Refer to **Figure 37**, **Figure 38**, and **Figure 35** on page 36 while reading the following explanation.

Each Global Early buffer can access the eight vertical Global lines for all CLBs in the quadrant. Therefore, only one-fourth of the CLB clock pins can be accessed. This restriction is in large part responsible for the faster speed of the buffers, relative to the Global Low-Skew buffers.




Figure 37: Left and Right BUFGEs Can Drive Any or All Clock Inputs in Same Quadrant or Edge (GCK1 is shown. GCK2, GCK5 and GCK6 are similar.)

The left-side Global Early buffers can each drive two of the four vertical lines accessing the IOBs on the entire left edge of the device. The right-side Global Early buffers can each drive two of the eight vertical lines accessing the IOBs on the entire right edge of the device. (See **Figure 37**.)

Each left and right Global Early buffer can also drive half of the IOBs along either the top or bottom edge of the device, using a dedicated line that can only be accessed through the Global Early buffers.

The top and bottom Global Early buffers can drive half of the IOBs along either the left or right edge of the device, as shown in **Figure 38**. They can only access the top and bottom IOBs via the CLB global lines.




Figure 38: Top and Bottom BUFGEs Can Drive Any or All Clock Inputs in Same Quadrant (GCK8 is shown. GCK3, GCK4 and GCK7 are similar.)

Table 20: XC4000E Program Data

Device	XC4003E	XC4005E	XC4006E	XC4008E	XC4010E	XC4013E	XC4020E	XC4025E
Max Logic Gates	3,000	5,000	6,000	8,000	10,000	13,000	20,000	25,000
CLBs (Row x Col.)	100 (10 x 10)	196 (14 x 14)	256 (16 x 16)	324 (18 x 18)	400 (20 x 20)	576 (24 x 24)	784 (28 x 28)	1,024 (32 x 32)
IOBs	80	112	128	144	160	192	224	256
Flip-Flops	360	616	768	936	1,120	1,536	2,016	2,560
Bits per Frame	126	166	186	206	226	266	306	346
Frames	428	572	644	716	788	932	1,076	1,220
Program Data	53,936	94,960	119,792	147,504	178,096	247,920	329,264	422,128
PROM Size (bits)	53,984	95,008	119,840	147,552	178,144	247,968	329,312	422,176

- Notes:
- Bits per Frame = (10 x number of rows) + 7 for the top + 13 for the bottom + 1 + 1 start bit + 4 error check bits
 Number of Frames = (36 x number of columns) + 26 for the left edge + 41 for the right edge + 1
 Program Data = (Bits per Frame x Number of Frames) + 8 postamble bits
 PROM Size = Program Data + 40 (header) + 8
 - The user can add more "one" bits as leading dummy bits in the header, or, if CRC = off, as trailing dummy bits at the end of any frame, following the four error check bits. However, the Length Count value **must** be adjusted for all such extra "one" bits, even for extra leading ones at the beginning of the header.

Table 21: XC4000EX/XL Program Data

Device	XC4002XL	XC4005	XC4010	XC4013	XC4020	XC4028	XC4036	XC4044	XC4052	XC4062	XC4085
Max Logic Gates	2,000	5,000	10,000	13,000	20,000	28,000	36,000	44,000	52,000	62,000	85,000
CLBs (Row x Column)	64 (8 x 8)	196 (14 x 14)	400 (20 x 20)	576 (24 x 24)	784 (28 x 28)	1,024 (32 x 32)	1,296 (36 x 36)	1,600 (40 x 40)	1,936 (44 x 44)	2,304 (48 x 48)	3,136 (56 x 56)
IOBs	64	112	160	192	224	256	288	320	352	384	448
Flip-Flops	256	616	1,120	1,536	2,016	2,560	3,168	3,840	4,576	5,376	7,168
Bits per Frame	133	205	277	325	373	421	469	517	565	613	709
Frames	459	741	1,023	1,211	1,399	1,587	1,775	1,963	2,151	2,339	2,715
Program Data	61,052	151,910	283,376	393,580	521,832	668,124	832,480	1,014,876	1,215,320	1,433,804	1,924,940
PROM Size (bits)	61,104	151,960	283,424	393,632	521,880	668,172	832,528	1,014,924	1,215,368	1,433,852	1,924,992

- Notes:
- Bits per frame = (13 x number of rows) + 9 for the top + 17 for the bottom + 8 + 1 start bit + 4 error check bits.
 Frames = (47 x number of columns) + 27 for the left edge + 52 for the right edge + 4.
 Program data = (bits per frame x number of frames) + 5 postamble bits.
 PROM size = (program data + 40 header bits + 8 start bits) rounded up to the nearest byte.
 - The user can add more "one" bits as leading dummy bits in the header, or, if CRC = off, as trailing dummy bits at the end of any frame, following the four error check bits. However, the Length Count value must be adjusted for all such extra "one" bits, even for extra leading "ones" at the beginning of the header.

Cyclic Redundancy Check (CRC) for Configuration and Readback

The Cyclic Redundancy Check is a method of error detection in data transmission applications. Generally, the transmitting system performs a calculation on the serial bitstream. The result of this calculation is tagged onto the data stream as additional check bits. The receiving system performs an identical calculation on the bitstream and compares the result with the received checksum.

Each data frame of the configuration bitstream has four error bits at the end, as shown in [Table 19](#). If a frame data error is detected during the loading of the FPGA, the con-

figuration process with a potentially corrupted bitstream is terminated. The FPGA pulls the $\overline{\text{INIT}}$ pin Low and goes into a Wait state.

During Readback, 11 bits of the 16-bit checksum are added to the end of the Readback data stream. The checksum is computed using the CRC-16 CCITT polynomial, as shown in [Figure 45](#). The checksum consists of the 11 most significant bits of the 16-bit code. A change in the checksum indicates a change in the Readback bitstream. A comparison to a previous checksum is meaningful only if the readback data is independent of the current device state. CLB outputs should not be included (Read Capture option not

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 V_{CC} 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 V_{CC} . When all \overline{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 $\overline{PROGRAM}$ pin



Figure 45: Circuit for Generating CRC-16



Figure 46: Power-up Configuration Sequence

Low. During this time delay, or as long as the $\overline{\text{PROGRAM}}$ input is asserted, the configuration logic is held in a Configuration Memory Clear state. The configuration-memory frames are consecutively initialized, using the internal oscillator.

At the end of each complete pass through the frame addressing, the power-on time-out delay circuitry and the level of the $\overline{\text{PROGRAM}}$ pin are tested. If neither is asserted, the logic initiates one additional clearing of the configuration frames and then tests the $\overline{\text{INIT}}$ input.

Initialization

During initialization and configuration, user pins HDC , $\overline{\text{LDC}}$, $\overline{\text{INIT}}$ and DONE provide status outputs for the system interface. The outputs $\overline{\text{LDC}}$, $\overline{\text{INIT}}$ and DONE are held Low and HDC is held High starting at the initial application of power.

The open drain $\overline{\text{INIT}}$ pin is released after the final initialization pass through the frame addresses. There is a deliberate delay of 50 to 250 μs (up to 10% longer for low-voltage devices) before a Master-mode device recognizes an inactive $\overline{\text{INIT}}$. Two internal clocks after the $\overline{\text{INIT}}$ pin is recognized as High, the FPGA samples the three mode lines to determine the configuration mode. The appropriate interface lines become active and the configuration preamble and data can be loaded. Configuration

The 0010 preamble code indicates that the following 24 bits represent the length count. The length count is the total number of configuration clocks needed to load the complete configuration data. (Four additional configuration clocks are required to complete the configuration process, as discussed below.) After the preamble and the length count have been passed through to all devices in the daisy chain, DOUT is held High to prevent frame start bits from reaching any daisy-chained devices.

A specific configuration bit, early in the first frame of a master device, controls the configuration-clock rate and can increase it by a factor of eight. Therefore, if a fast configuration clock is selected by the bitstream, the slower clock rate is used until this configuration bit is detected.

Each frame has a start field followed by the frame-configuration data bits and a frame error field. If a frame data error is detected, the FPGA halts loading, and signals the error by pulling the open-drain $\overline{\text{INIT}}$ pin Low. After all configuration frames have been loaded into an FPGA, DOUT again follows the input data so that the remaining data is passed on to the next device.

Delaying Configuration After Power-Up

There are two methods of delaying configuration after power-up: put a logic Low on the $\overline{\text{PROGRAM}}$ input, or pull the bidirectional $\overline{\text{INIT}}$ pin Low, using an open-collector (open-drain) driver. (See [Figure 46 on page 50](#).)

A Low on the $\overline{\text{PROGRAM}}$ input is the more radical approach, and is recommended when the power-supply

rise time is excessive or poorly defined. As long as $\overline{\text{PROGRAM}}$ is Low, the FPGA keeps clearing its configuration memory. When $\overline{\text{PROGRAM}}$ goes High, the configuration memory is cleared one more time, followed by the beginning of configuration, provided the $\overline{\text{INIT}}$ input is not externally held Low. Note that a Low on the $\overline{\text{PROGRAM}}$ input automatically forces a Low on the $\overline{\text{INIT}}$ output. The XC4000 Series $\overline{\text{PROGRAM}}$ pin has a permanent weak pull-up.

Using an open-collector or open-drain driver to hold $\overline{\text{INIT}}$ Low before the beginning of configuration causes the FPGA to wait after completing the configuration memory clear operation. When $\overline{\text{INIT}}$ is no longer held Low externally, the device determines its configuration mode by capturing its mode pins, and is ready to start the configuration process. A master device waits up to an additional 250 μs to make sure that any slaves in the optional daisy chain have seen that $\overline{\text{INIT}}$ is High.

Start-Up

Start-up is the transition from the configuration process to the intended user operation. This transition involves a change from one clock source to another, and a change from interfacing parallel or serial configuration data where most outputs are 3-stated, to normal operation with I/O pins active in the user-system. Start-up must make sure that the user-logic 'wakes up' gracefully, that the outputs become active without causing contention with the configuration signals, and that the internal flip-flops are released from the global Reset or Set at the right time.

[Figure 47](#) describes start-up timing for the three Xilinx families in detail. The configuration modes can use any of the four timing sequences.

To access the internal start-up signals, place the STARTUP library symbol.

Start-up Timing

Different FPGA families have different start-up sequences.

The XC2000 family goes through a fixed sequence. DONE goes High and the internal global Reset is de-activated one CCLK period after the I/O become active.

The XC3000A family offers some flexibility. DONE can be programmed to go High one CCLK period before or after the I/O become active. Independent of DONE , the internal global Reset is de-activated one CCLK period before or after the I/O become active.

The XC4000 Series offers additional flexibility. The three events — DONE going High, the internal Set/Reset being de-activated, and the user I/O going active — can all occur in any arbitrary sequence. Each of them can occur one CCLK period before or after, or simultaneous with, any of the others. This relative timing is selected by means of software options in the bitstream generation software.

The default option, and the most practical one, is for DONE to go High first, disconnecting the configuration data source and avoiding any contention when the I/Os become active one clock later. Reset/Set is then released another clock period later to make sure that user-operation starts from stable internal conditions. This is the most common sequence, shown with heavy lines in [Figure 47](#), but the designer can modify it to meet particular requirements.

Normally, the start-up sequence is controlled by the internal device oscillator output (CCLK), which is asynchronous to the system clock.

XC4000 Series offers another start-up clocking option, UCLK_NOSYNC. The three events described above need not be triggered by CCLK. They can, as a configuration option, be triggered by a user clock. This means that the device can wake up in synchronism with the user system.

When the UCLK_SYNC option is enabled, the user can externally hold the open-drain DONE output Low, and thus stall all further progress in the start-up sequence until DONE is released and has gone High. This option can be used to force synchronization of several FPGAs to a common user clock, or to guarantee that all devices are successfully configured before any I/Os go active.

If either of these two options is selected, and no user clock is specified in the design or attached to the device, the chip could reach a point where the configuration of the device is complete and the Done pin is asserted, but the outputs do not become active. The solution is either to recreate the bit-stream specifying the start-up clock as CCLK, or to supply the appropriate user clock.

Start-up Sequence

The Start-up sequence begins when the configuration memory is full, and the total number of configuration clocks

received since $\overline{\text{INIT}}$ went High equals the loaded value of the length count.

The next rising clock edge sets a flip-flop Q0, shown in [Figure 48](#). Q0 is the leading bit of a 5-bit shift register. The outputs of this register can be programmed to control three events.

- The release of the open-drain DONE output
- The change of configuration-related pins to the user function, activating all IOBs.
- The termination of the global Set/Reset initialization of all CLB and IOB storage elements.


The DONE pin can also be wire-ANDed with DONE pins of other FPGAs or with other external signals, and can then be used as input to bit Q3 of the start-up register. This is called “Start-up Timing Synchronous to Done In” and is selected by either CCLK_SYNC or UCLK_SYNC.

When DONE is not used as an input, the operation is called “Start-up Timing Not Synchronous to DONE In,” and is selected by either CCLK_NOSYNC or UCLK_NOSYNC.

As a configuration option, the start-up control register beyond Q0 can be clocked either by subsequent CCLK pulses or from an on-chip user net called STARTUP.CLK. These signals can be accessed by placing the STARTUP library symbol.

Start-up from CCLK

If CCLK is used to drive the start-up, Q0 through Q3 provide the timing. Heavy lines in [Figure 47](#) show the default timing, which is compatible with XC2000 and XC3000 devices using early DONE and late Reset. The thin lines indicate all other possible timing options.



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.




Figure 49: Readback Schematic Example

Readback Options

Readback options are: Read Capture, Read Abort, and Clock Select. They are set with the bitstream generation software.

Read Capture

When the Read Capture option is selected, the readback data stream includes sampled values of CLB and IOB signals. The rising edge of RDBK.TRIG latches the inverted values of the four CLB outputs, the IOB output flip-flops and the input signals I1 and I2. Note that while the bits describing configuration (interconnect, function generators, and RAM content) are *not* inverted, the CLB and IOB output signals *are* inverted.

When the Read Capture option is not selected, the values of the capture bits reflect the configuration data originally written to those memory locations.

If the RAM capability of the CLBs is used, RAM data are available in readback, since they directly overwrite the F and G function-table configuration of the CLB.

RDBK.TRIG is located in the lower-left corner of the device, as shown in [Figure 50](#).

Read Abort

When the Read Abort option is selected, a High-to-Low transition on RDBK.TRIG terminates the readback operation and prepares the logic to accept another trigger.

After an aborted readback, additional clocks (up to one readback clock per configuration frame) may be required to re-initialize the control logic. The status of readback is indicated by the output control net RDBK.RIP. RDBK.RIP is High whenever a readback is in progress.

Clock Select

CCLK is the default clock. However, the user can insert another clock on RDBK.CLK. Readback control and data are clocked on rising edges of RDBK.CLK. If readback must be inhibited for security reasons, the readback control nets are simply not connected.

RDBK.CLK is located in the lower right chip corner, as shown in [Figure 50](#).




Figure 50: READBACK Symbol in Graphical Editor

Violating the Maximum High and Low Time Specification for the Readback Clock

The readback clock has a maximum High and Low time specification. In some cases, this specification cannot be met. For example, if a processor is controlling readback, an interrupt may force it to stop in the middle of a readback. This necessitates stopping the clock, and thus violating the specification.

The specification is mandatory only on clocking data at the end of a frame prior to the next start bit. The transfer mechanism will load the data to a shift register during the last six clock cycles of the frame, prior to the start bit of the following frame. This loading process is dynamic, and is the source of the maximum High and Low time requirements.

Therefore, the specification only applies to the six clock cycles prior to and including any start bit, including the clocks before the first start bit in the readback data stream. At other times, the frame data is already in the register and the register is not dynamic. Thus, it can be shifted out just like a regular shift register.

The user must precisely calculate the location of the readback data relative to the frame. The system must keep track of the position within a data frame, and disable interrupts before frame boundaries. Frame lengths and data formats are listed in [Table 19](#), [Table 20](#) and [Table 21](#).

Readback with the XChecker Cable

The XChecker Universal Download/Readback Cable and Logic Probe uses the readback feature for bitstream verification. It can also display selected internal signals on the PC or workstation screen, functioning as a low-cost in-circuit emulator.

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

Synchronous Peripheral Mode

Synchronous Peripheral mode can also be considered Slave Parallel mode. An external signal drives the CCLK input(s) of the FPGA(s). The first byte of parallel configuration data must be available at the Data inputs of the lead FPGA a short setup time before the rising CCLK edge. Subsequent data bytes are clocked in on every eighth consecutive rising CCLK edge.

The same CCLK edge that accepts data, also causes the RDY/ $\overline{\text{BUSY}}$ output to go High for one CCLK period. The pin name is a misnomer. In Synchronous Peripheral mode it is really an ACKNOWLEDGE signal. Synchronous operation does not require this response, but it is a meaningful signal for test purposes. Note that RDY/ $\overline{\text{BUSY}}$ is pulled High with a high-impedance pullup prior to $\overline{\text{INIT}}$ going High.

The lead FPGA serializes the data and 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, 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.

In order to complete the serial shift operation, 10 additional CCLK rising edges are required after the last data byte has been loaded, plus one more CCLK cycle for each daisy-chained device.

Synchronous Peripheral mode is selected by a <011> on the mode pins (M2, M1, M0).

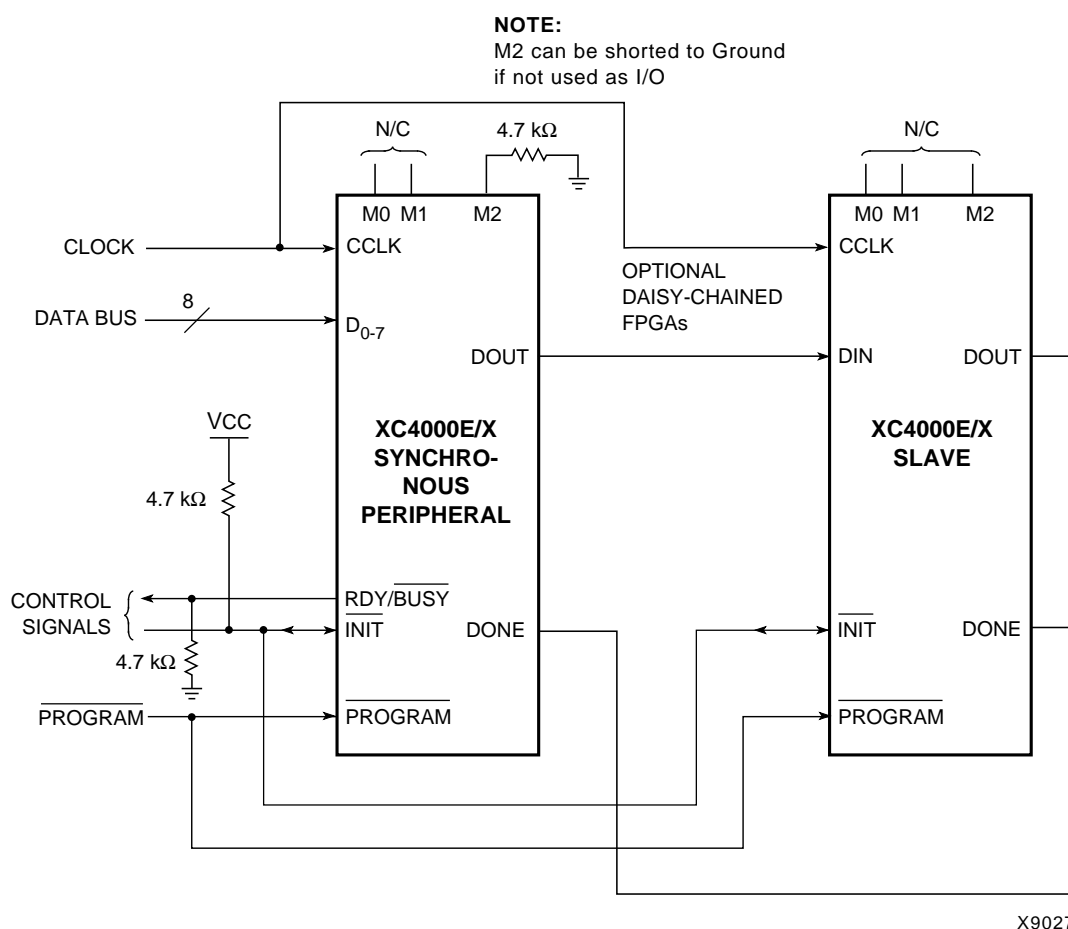


Figure 56: Synchronous Peripheral Mode Circuit Diagram