

Summary

This Application Note contains additional information that may be of use when designing with the XC3000 class of LCA devices. This information supplements the data sheets, and is provided for guidance only.

Xilinx Family

XC3000/XC3000A/XC3000L/XC3100/XC3100A

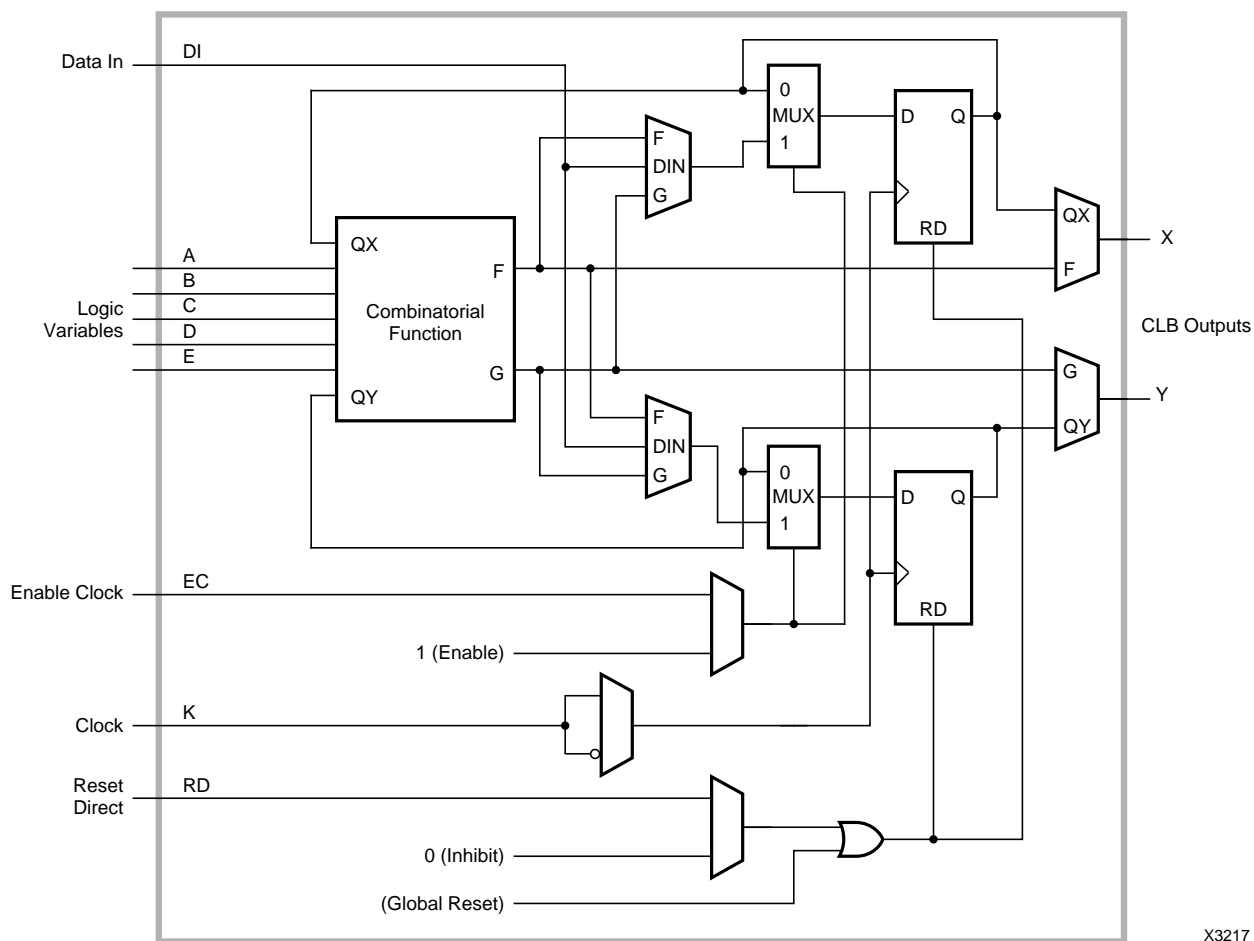
Introduction

The background information provided in this Application Note supplements the XC3000, XC3000A, XC3000L and XC3100 data sheets. It covers a wide range of topics, including a number of electrical parameters not specified in the data sheets, and unless otherwise noted, applies to all four families. These additional parameters are sufficiently accurate for most design purposes;

unlike the parameters specified in the data sheets, however, they are not worst-case values over temperature and voltage, and are not 100% production tested. They can, therefore, not be guaranteed.

Configurable Logic Blocks

The XC3000/XC3100 CLB, shown in Figure 1, comprises a combinatorial function generator and two D-type flip-flops. Two output pins may be driven by either the



X3217

Figure 1. Configurable Logic Block (CLB)

function generators or the flip-flops. The flip-flop outputs may be routed directly back to the function generator inputs without going outside of the CLB.

The function generator consists of two 4-input look-up tables that may be used separately or combined into a single function. Figure 2 shows the three available options. Since the CLB only has five inputs to the function generator, inputs must be shared between the two look-up tables.

In the FG mode, the function generator provides any two 4-input functions of A, B and C plus D or E; the choice between D and E is made separately for each function. In the F mode, all five inputs are combined into a single 5-input function of A, B, C, D and E. Any 5-input function may be emulated. The FGM mode is a superset of the F mode, where two 4-input functions of A, B, C and D are multiplexed together according to the fifth variable, E.

In all modes, either of the B and C inputs may be selectively replaced by QX and QY, the flip-flop outputs. In the FG mode, this selection is made separately for the two look-up tables, extending the functionality to any two functions of four variable chosen from seven, provided two of the variables are stored in the flip-flops. This is particularly useful in state-machine-like applications.

In the F mode, the function generators implement a single function of five variables that may be chosen from seven, as described above. The selection of QX and QY is constrained to be the same for both look-up tables. The FGM mode differs from the F mode in that QX and QY may be selected separately for the two look-up tables, as in the FG mode. This added flexibility permits the emulation of selected functions that can include all seven possible inputs.

The automatic logic-partitioning software in the XACT development system only uses the FG and F modes. However, all three modes are available with manual partitioning, which may be performed in the schematic. If FG or F modes are required, it is simply a matter of including in the schematic CLBMAPs that define the inputs and outputs of the CLB.

The FGM mode is only slightly more complicated. Again, a CLBMAP must be used, with the signal that multiplexes between the two 4-input functions locked onto the E pin. The CLB will be configured in the FGM mode if the logic is drawn such that the gates forming the multiplexer are shown explicitly with no additional logic merged into them.

The two D-type flip-flops share a common clock, a common clock enable, and a common asynchronous reset signal. An asynchronous preset can be achieved using the asynchronous reset if data is stored in active-low form; the Low created by reset corresponds to the bit being asserted. The flip-flops cannot be used as latches.

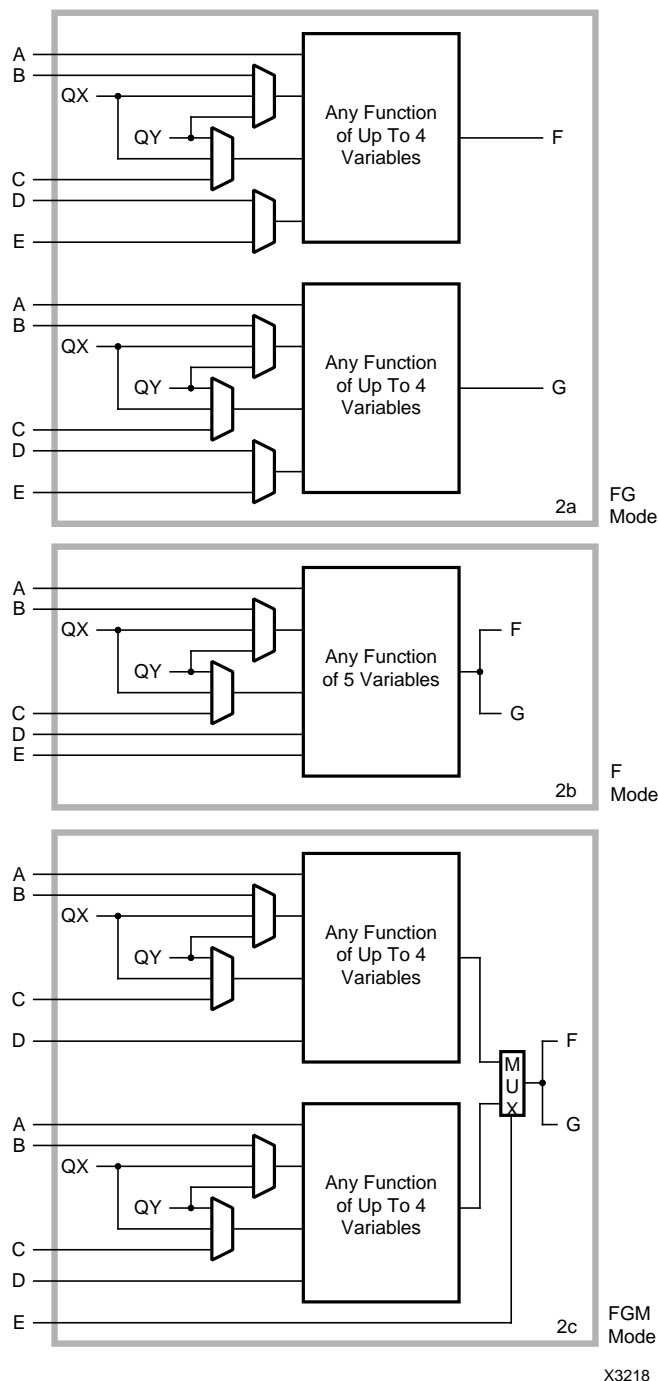


Figure 2. CLB Logic Options

If input data to a CLB flip-flop is derived directly from an input pad, without an intervening flip-flop, the data-pad-to-clock-pad hold time will typically be non-zero. This hold time is equal the delay from the clock pad to the CLB, but may be reduced according to the 70% rule, described later in the IOB Input section of this Application Note. Under this rule, the hold time is reduced by 70% of the delay from the data pad to the CLB, excluding the CLB

Table 1. Longline to CLB Direct Access

Longline	CLB								TBUF
	A	B	C	D	E	K	EC	RD	T
Left Most Vertical (GCLK)						X			
Left Middle Vertical		X					X	X	X
Right Middle Vertical			X		X				
Right Most Vertical (ACLK)						X			
Upper Horizontal				X					
Lower Horizontal	X							X	

set-up time. The minimum hold time is zero, even when applying the 70% rule results in a negative number.

The CLB pins to which Longlines have direct access are shown in Table 1. Note that the clock enable pin (EC) and the TBUF control pin are both driven from the same vertical Long Line. Consequently, EC cannot easily be used to enable a register that must be 3-stated onto a bus. Similarly, EC cannot easily be used in a register that uses the Reset Direct pin (RD).

Input/Output Blocks

The XC3000/XC3100 IOB, shown in Figure 3, includes a 3-state output driver that may be driven directly or registered. The polarities of both the output data and the 3-state control are determined by configuration bits. Each output buffer may be configured to have either a fast or a slow slew rate.

The IOB input may also be direct or registered. Additionally, the input flip-flop may be configured as a latch. When an IOB is used exclusively as an input, an optional pull-up resistor is available, the value of which is 40-150 k Ω . This resistor cannot be used when the IOB is configured as an output or as a bidirectional pin.

Unused IOBs should be left unconfigured. They default to inputs pulled High with the internal resistor.

Inputs

All inputs have limited hysteresis, typically in excess of 200 mV for TTL input thresholds and in excess of 100 mV for CMOS thresholds. Exceptions to this are the **PWRDWN** pin, and the **XTL2** pin when it is configured as the crystal oscillator input.

Experiments show that the input rise and fall times should not exceed 250 ns. This value was established through a worst-case test using internal ring oscillators to drive all I/O pins except two, thus generating a maximum of on-

chip noise. One of the remaining I/O pins was configured as an input, and tested for single-edge response; the other I/O was used as an output to monitor the response.

These test conditions are, perhaps, overly demanding, although it was assumed that the PC board had negligible ground noise and good power-supply decoupling. While conservative, the resulting specification is, in most instances, easily satisfied.

IOB input flip-flops are guaranteed to operate correctly without data hold times (with respect to the device clock-input pad) provided that the dedicated CMOS clock input pad and the GCLK buffer are used. The use of a TTL clock or a different clock pad will result in a data-hold-time requirement. The length of this hold time is equal to the delay from the actual clock pad to the GCLK buffer minus the delay from the dedicated CMOS clock pad to the GCLK buffer.

To ensure that the input flip-flop has a zero hold time, delay is incorporated in the D input of the flip-flop, causing it to have a relatively long set-up time. However, the set-up time specified in the data sheet is with respect to the clock reaching the IOB. Since there is an unavoidable delay between the clock pad and the IOB, the input-pad-to-clock-pad set-up time is actually less than the data sheet number.

Part of the clock delay can be subtracted from the internal set-up time. Ideally, all of the clock delay could be subtracted, but it is possible for the clock delay to be less than its maximum while the internal set-up time is at its maximum value. Consequently, it is recommended that, in a worst-case design, only 70% of the clock delay is subtracted.

The clock delay can only be less than 70% of its maximum if the internal set-up time requirement is also less than its maximum. In this case, the pad-to-pad set-up time actually required will be less than that calculated.

For example, in the XC3000-125, the input set-up time with respect to the clock reaching the IOB is 16 ns. If the delay from the clock pad to the IOB is 6 ns, then 70% of this delay, 4.2 ns, can be subtracted to arrive at a maximum pad-to-pad set-up time of ~12 ns.

The 70% rule must be applied whenever one delay is subtracted from another. However, it is recommended that delay compensation only be used routinely in connection with input hold times. *Delay compensation in asynchronous circuits is specifically not recommended.* In any case, the compensated delay must not become negative. If 70% of the compensating delay is greater than the delay from which it is deducted, the resulting delay is zero.

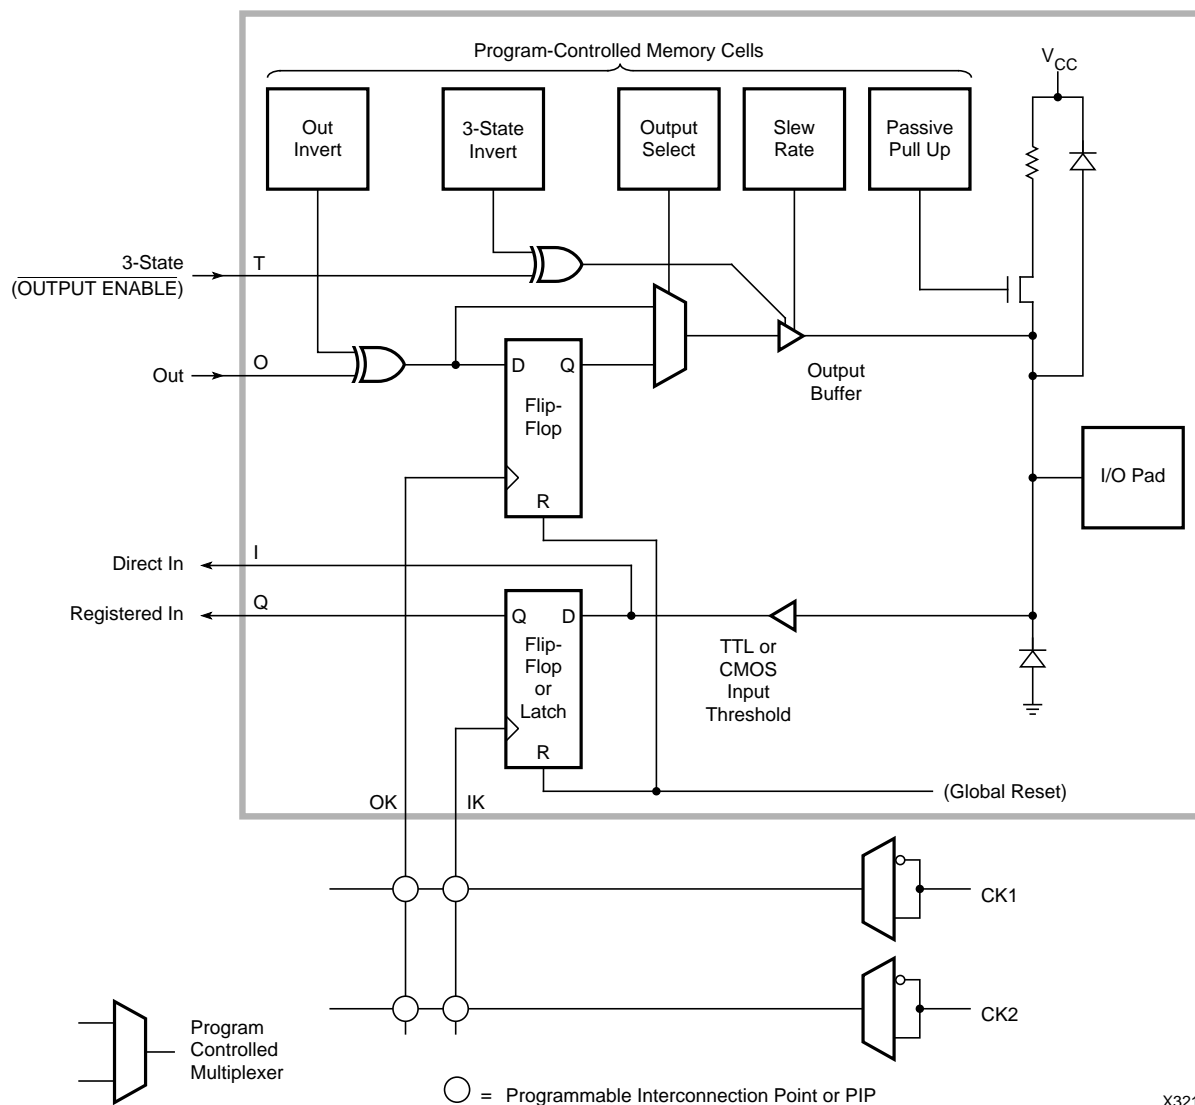


Figure 3. Input/Output Block (IOB)

The 70% rule in no way defines the absolute minimum values delays that might be encountered from chip to chip, and with temperature and power-supply variations. It simply indicates the relative variations that might be found within a specific chip over the range of operating conditions.

Typically, all delays will be less than their maximum, with some delays being disproportionately faster than others. The 70% rule describes the spread in the scaling factors; the delay that decreases the most will be no less than 70% of what it would have been if it had scaled in proportion to the delay that decreased the least. In particular, in a worst-case design where it is assumed that any delay might not have scaled at all, and remains at its maximum value, other delays will be no less than 70% of their maximum.

Outputs

All XC3000/XC3100 LCA outputs are true CMOS with n-channel transistors pulling down and p-channel transistors pulling up. Unloaded, these outputs pull rail-to-rail. Some additional ac characteristics of the output are listed in Table 2. Figures 4 and 5 show output current/voltage curves for typical XC3000 and XC3100 devices. See other Xilinx product family output characteristic on pages 8-6, 8-7 and 9-23.

Output-short-circuit-current values are given only to indicate the capability to charge and discharge capacitive loads. In accordance with common industry practice for other logic devices, only one output at a time may be short circuited, and the duration of this short circuit to V_{CC} or ground may not exceed one second. Xilinx does not recommend a continuous output or clamp current in




Figure 4. Output Current/Voltage Characteristics for XC3020

excess of 20 mA on any one output pin. The data sheet guarantees the outputs for no more than 4 mA at 320 mV to avoid problems when many outputs are sinking current simultaneously.

The active-High 3-state control (T) is the same as an active-Low output enable (\overline{OE}). In other words, a High on the T-pin of an OBUFZ places the output in a high impedance state, and a Low enables the output. The same naming convention is used for TBUFs within the LCA device.

I/O Clocks

Internally, up to eight distinct I/O clocks can be used, two on each of the four edges of the die. While the IOB does not provide programmable clock polarity, the two clock lines serving an IOB can be used for true and inverted clock, and the appropriate polarity connected to the IOB. This does, however, limit all IOBs on that edge of the die to using only the two edges of the one clock.

IOB latches have active-Low Latch Enables; they are transparent when the clock input is Low and are closed when it is High. The latch captures data on what would otherwise be the active clock edge, and is transparent in the half clock period before the active clock edge.

Table 2. Additional AC Output Characteristics

AC Parameters	Fast*	Slow*
Unloaded Output Slew Rate	2.8 V/ns	0.5 V/ns
Unloaded Transition Time	1.45 ns	7.9 ns
Additional rise time for 812 pF	100 ns	100 ns
normalized	0.12 ns/pF	0.12 ns/pF
Additional fall time for 812 pF	50 ns	64 ns
normalized	0.06 ns/pF	0.08 ns/pF

* Fast and Slow refer to the output programming option.




Figure 5. Output Current/Voltage Characteristics for XC3142

Routing

Horizontal Longlines

As shown in Table 3, there are two horizontal Longlines (HLLs) per row of CLBs. Each HLL is driven by one TBUF for each column of CLBs, plus an additional TBUF at the left end of the Longline. This additional TBUF is convenient for driving IOB data onto the Longline. In general, the routing resources to the T and I pins of TBUFs are somewhat limited.

Optionally, HLLs can be pulled up at either end, or at both ends. The value of each pull-up resistor is 3-10 k Ω .

In addition, HLLs are permanently driven by low-powered latches that are easily overridden by active outputs or pull-up resistors. These latches maintain the logic levels on HLLs that are not pulled up and temporarily are not driven. The logic level maintained is the last level actively driven onto the line.

When using 3-state HLLs for multiplexing, the use of fewer than four TBUFs can waste resources. Multiplexers with four or fewer inputs can be implemented more efficiently using CLBs.

Table 3. Number of Horizontal Longlines

Part Name	Rows x Columns	CLBs	Horizontal Longlines	TBUFs per HLL
XC3020	8 x 8	64	16	9
XC3030	10 x 10	100	20	11
XC3042	12 x 12	144	24	13
XC3064	16 x 14	224	32	15
XC3090	20 x 16	320	40	17
XC3195	22 x 22	484	44	23

Vertical Longlines

There are four vertical Longlines per routing channel: two general purpose, one for the global clock net and one for the alternate clock net.

Clock Buffers

XC3000/XC3100 devices each contain two high-fan-out, low-skew clock-distribution networks. The global-clock net originates from the GCLK buffer in the upper left corner of the die, while the alternate clock net originates from the ACLK buffer in the lower right corner of the die.

The global and alternate clock networks each have optional fast CMOS inputs, called TCLKIN and BCLKIN, respectively. Using these inputs provides the fastest path from the PC board to the internal flip-flops and latches. Since the signal bypasses the input buffer, well-defined CMOS levels must be guaranteed on these clock pins.

To specify the use of TCLKIN or BCLKIN in a schematic, connect an IPAD symbol directly to the GCLK or ACLK symbol. Placing an IBUF between the IPAD and the clock buffer will prevent TCLKIN or BCLKIN from being used.

The clock buffer output nets only drive CLB and IOB clock pins. *They do not drive any other CLB inputs.* In rare cases where a clock needs to be connected to a logic input or a device output, a signal should be tapped off the clock buffer input, and routed to the logic input. This is not possible with clocks using TCLKIN or BCLKIN.

The clock skew created by routing clocks through local interconnect makes safe designs very difficult to achieve, and this practice is not recommended. In general, the fewer clocks that are used, the safer the design. High fan-out clocks should always use GCLK or ACLK. If more than two clocks are required, the ACLK net can be segmented into individual vertical lines that can be driven by PIPs at the top and bottom of each column. Clock signals routed through local interconnect should only be considered for individual flip-flops.

General Information

Recovery from Reset

Recovery from Reset is not specified in Xilinx data sheets because it is very difficult to measure in a production environment. The following values may be assumed for all XC3000/XC3100 devices and speed grades.

- The CLB can be clocked immediately (<0.2 ns) after the end of the internal Reset Direct signal (RD).
- The CLB can be clocked no earlier than 25 ns (worst case) after the release of an externally applied Global Reset signal, i.e., after the rising edge of the active-Low signal.

Configuration and Start-up

Until the chip goes active after configuration, all I/O pins not involved in the configuration process remain in a high-impedance state with weak pull-up resistors; all internal flip-flops and latches are held reset. Multiple LCA devices hooked up in a daisy chain will all go active simultaneously on the same CCLK edge. This is well documented in the data sheets.

Not documented, however, is how the internal combinatorial logic comes alive during configuration: As configuration data is shifted in and reaches its destination, it activates the logic and also “looks at” the IOB inputs. Even the crystal oscillator starts operating as soon as it receives its configuration data. Since all flip-flops and latches are being held reset, and all outputs are being held in their high-impedance state, there is no danger in this “staggered awakening” of the internal logic. The operation of the logic prior to the end of configuration is even useful; it ensures that clock enables and output enables are correctly defined before the elements they control become active.

Once configuration is complete, the LCA device is activated. This occurs on a rising edge of CCLK, when all outputs and clocks that are enabled become active simultaneously. Since the activation is triggered by CCLK, it is an asynchronous event with respect to the system clock. To avoid start-up problems caused by this asynchronism, some designs might require a reset pulse that is synchronized to the system clock.

The circuit shown in Figure 5 generates a short Global Reset pulse in response to the first system clock after the end of configuration. It uses one CLB and one IOB, and also precludes the use of the LDC pin as I/O.

During Configuration, LDC is asserted Low and holds the D-input of the flip-flop High, while Q is held Low by the internal reset, and RESET is kept High by internal and external pull-up resistors. At the end of configuration, the LDC pin is unasserted, but D remains High since the function generator acts as an R-S latch; Q stays Low, and RESET is still pulled High by the external resistor. On the first system clock after configuration ends, Q is clocked High, resetting the latch and enabling the output driver, which forces RESET Low. This resets the whole chip until the Low on Q permits RESET to be pulled High again.

The whole chip has thus been reset by a short pulse instigated by the system clock. No further pulses are generated, since the High on LDC prevents the R-S latch from becoming set.

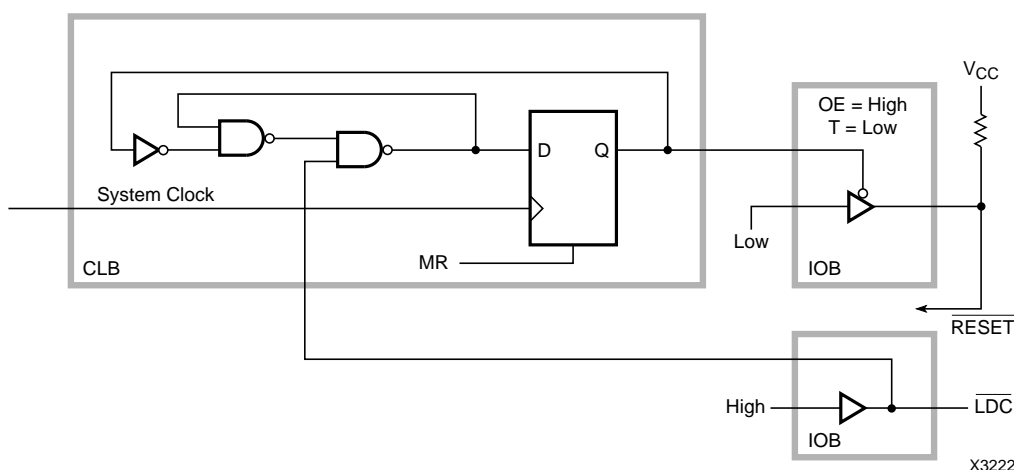


Figure 5. Synchronous Reset

Power Dissipation

As in most CMOS ICs, almost all LCA power dissipation is dynamic, and is caused by the charging and discharging of internal capacitances. Each node in the device dissipates power according to the capacitance in the node, which is fixed for each type of node, and the frequency at which the particular node is switching, which can be different from the clock frequency. The total dynamic power is the sum of the power dissipated in the individual nodes.

While the clock line frequency is easy to specify, it is usually more difficult to estimate the average frequency of other nodes. Two extreme cases are binary counters, where half the total power is dissipated in the first flip-flop, and shift registers with alternating zeros and ones, where the whole circuit is exercised at the clocking speed.

A popular assumption is that, on average, each node is exercised at 20% of the clock rate; a major EPLD vendor uses a 16-bit counter as a model, where the effective percentage is only 12%. Undoubtable, there are extreme cases, where the ratio is much lower or much higher, but 15 to 20% may be a valid approximation for most normal designs. Note that global clock lines must always be entered with their real, and obviously well-known, frequency.

Consequently, most power consumption estimates only serve as guidelines based on gross approximations. Table 4 shows the dynamic power dissipation, in mW per MHz, for different types of XC3000 nodes. While not precise, these numbers are sufficiently accurate for the calculations in which they are used, and may be used for any XC3000/XC3100 device. Table 5 shows a sample power calculation.

Table 4. Dynamic Power Dissipation

	XC3020	XC3090	
One CLB driving three local interconnects	0.25	0.25	mW/MHz
One device output with a 50 pF load	1.25	1.25	mW/MHz
One Global Clock Buffer and line	2.00	3.50	mW/MHz
One Longline without driver	0.10	0.15	mW/MHz

Table 5. Sample Power Calculation for XC3020

Quantity	Node	MHz	mW/MHz	mW
1	Clock Buffer	40	2.00	80
5	CLBs	40	0.25	50
10	CLBs	20	0.25	50
40	CLBs	10	0.25	100
8	Longlines	20	0.10	16
20	Outputs	20	1.25	500
				Total Power ~800

Crystal Oscillator

XC3000 and XC3100 devices contain an on-chip crystal oscillator circuit that connects to the ACLK buffer. This circuit, Figure 6, comprises a high-speed, high-gain inverting amplifier with its input connected to the dedicated XTL2 pin, and its output connected to the XTL1 pin. An external biasing resistor, R1, with a value of 0.5 to 1 M Ω is required.

A crystal, Y1, and additional phase-shifting components, R2, C1 and C2, complete the circuit. The capacitors, C1 and C2, in parallel form the load on the crystal. This load is specified by the crystal manufacturer, and is typically 40 pF. The capacitors should be approximately equal: 20 pF each for a 40 pF crystal.

Either series- or parallel-resonant crystals may be used, since they differ only in their specification. Crystals constrain oscillation to a narrow band of frequencies, the width of which is $\ll 1\%$ of the oscillating frequency; the exact frequency of oscillation within this band depends on the components surrounding the crystal. Series-resonant crystals are specified by their manufacturers according to the lower edge of the frequency band, parallel-resonant crystals according to the upper edge.

The resistor R2 controls the loop gain and its value must be established by experimentation. If it is too small, the oscillation will be distorted; if it is too large, the oscillation will fail to start, or only start slowly. In most cases, the value of R2 is non-critical, and typically is 0 to 1 k Ω .

Once the component values have been chosen, it is good practice to test the oscillator with a resistor (~ 1 k Ω) in series with the crystal. If the oscillator still starts reliably, independent of whether the power supply turns on quickly or slowly, it will always work without the resistor.



Figure 6. Crystal Oscillator

Table 6. Third Harmonic Crystal Oscillator Tank-Circuit

Frequency (MHz)	LC Tank				
	L (μ H)	C2 (pF)	Freq (MHz)	R2 (Ω)	C1 (pF)
32	1	60	20.6	430	23
35	1	44	24.0	310	23
49	1	31	28.6	190	23
72	1	18	37.5	150	12

For operation above 20 to 25 MHz, the crystal must be operated at its third harmonic. The capacitor C2 is replaced by a parallel-resonant LC tank circuit tuned to $\sim 2/3$ of the desired frequency, i.e., twice the fundamental frequency of the crystal. Table 6 shows typical component values for the tank circuit.

See pages 9-30 and 9-31 for a more detailed description of crystal oscillators.

CCLK Frequency Variation

The on-chip R-C oscillator that is brought out as CCLK also performs several other internal function. It generates the power-on delay, $2^{16} = 65,536$ periods for a master, $2^{14} = 16,384$ periods for a slave or peripheral device. It generates the shift pulses for clearing the configuration array, using one clock period per frame, and it is the clock source for several small shift registers acting as low-pass filters for a variety of input signals.

The nominal frequency of this oscillator is 1 MHz with a max deviation of +25% to -10%. The clock frequency, therefore, is between 1.25 MHz and 0.5 MHz. In the XC4000 family, the 1-MHz clock is derived from an internal 8-MHz clock that also can be used as CCLK source.

Xilinx circuit designers make sure that the internal clock frequency does not get faster as devices are migrated to smaller geometries and faster processes. Even the newest and fastest Xilinx FPGA is compatible with the oldest and slowest device ever manufactured. The CCLK frequency is fairly insensitive to changes in V_{CC} , varying only 0.6% for a 10% change in V_{CC} . It is, however, very temperature dependent, increasing 40% as the temperature drops from 25°C to -30°C, Table 7.

Table 7. CCLK Frequency Variation on a Sample Device

V_{CC}	Temp	Frequency
4.5 V	25°C	687 kHz
5.0 V	25°C	691 kHz
5.5 V	25°C	695 kHz
4.5 V	-30°C	966 kHz
4.5 V	+130°C	457 kHz

Metastable Recovery

Whenever a clocked flip-flop synchronizes an asynchronous input, there is a small probability that the flip-flop output will exhibit an unpredictable delay. This happens when the input transition not only violates the setup and hold-time specifications, but actually occurs within the tiny timing window where the flip-flop accepts the new input. Under these circumstances, the flip-flop can enter a symmetrically balanced transitory state, called metastable (meta = between).

While the slightest deviation from perfect balance will cause the output to revert to one of its two stable states, the delay in doing so depends not only on the gain bandwidth product of the circuit, but also on how perfect the balance is, and on the noise level within the circuit; the delay can, therefore, only be described in statistical terms.

The problem for the system designer is not the illegal logic level in the balanced state (it's easy enough to translate that to either a 0 or a 1), but the unpredictable timing of the final change to a valid logic state. If the metastable flip-flop drives two destinations with differing path delays, one destination might reflect the final data state while the other does not.

With the help of a mostly self-contained circuit on the demonstration board that is available to all Xilinx customers, Xilinx evaluated the XC3020-70 CLB flip-flop. The result of this evaluation shows the Xilinx CLB flip-flop to be superior in metastable performance to many popular MSI and PLD devices.

Statistically, when an asynchronous event with a frequency of approximately 1 MHz is being synchronized by a 10-MHz clock, the CLB flip-flop suffers an additional delay, of 4.2 ns once per hour, and 8.4 ns once per 1,000 years.

The frequency of occurrence of these metastable delays is proportional to the product of the asynchronous event frequency and the clock frequency. If, as an example, a

100-kHz event is synchronized by a 2-MHz clock, the above delays (besides being far more tolerable) will occur 50 times less often.

Since metastability can only be measured statistically, this data was obtained by configuring an XC3020 with eight concurrent detectors. Eight D-type flip-flops were clocked from a common high-speed source, and their D inputs driven from a common, lower frequency asynchronous signal, Figure 7. The output of each flip-flop fed the D inputs of two more flip-flops, one clocked half a clock period later and the second a full clock period later.

If a metastable event in the first flip-flop increased the output settling time to more than one-half clock period, the second two flip-flops would capture differing data. Thus, the occurrence of a long metastable delay could be detected using a simple comparator. Deliberate skew in the input data to the eight metastable circuits ensured that at most one metastable event could occur each clock. This permitted the eight detectors to be ORed into a single metastable event counter.

As expected, no metastable events were observed at clock rates below 25 MHz, since a half clock period of 20 ns is adequate for almost any metastability-resolution delay plus the flip-flop set-up time. Increasing the clock rate to around 27 MHz brought a sudden burst of metastable events. Careful adjustment of the clock frequency gave repeatable, reliable measurements showing that a 500 ps decrease in the half clock period increased the frequency of metastable occurrences by a factor of 41.

To be conservative, to compensate for favorable conditions at room temperature and to avoid any possibility of overstating a good case, the measurements were interpreted as follows:

When capturing asynchronous data, the error rate decreases by a factor of 40 for every additional nanosecond of metastability-resolution delay that the system can tolerate.

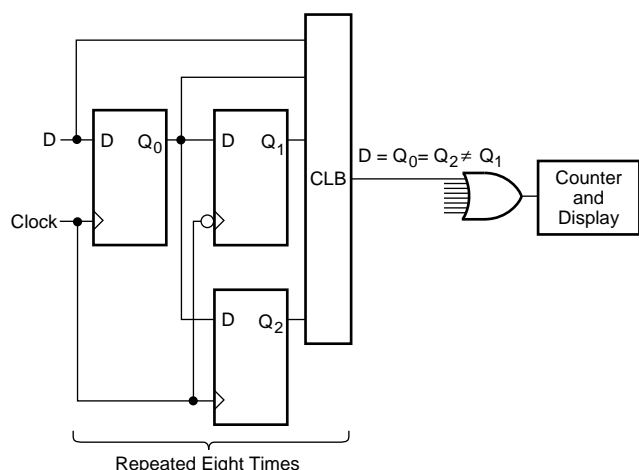



Figure 7. Metastable Measuring Circuit



X3226

Metastability Calculations

The Mean Time Between Failure (MTBF) can only be defined statistically. It is inversely proportional to the product of the two frequencies involved, the clock frequency and the average frequency of data changes, provided that these two frequencies are independent and have no correlation.

K1 is a factor that has the dimension of time, and describes the likelihood of going metastable. K2 is an exponent that describes the speed with which the metastable condition is being resolved.

$$1/\text{MTBF} = f_1 \cdot f_2 \cdot K_1 \cdot e^{-K_2 \cdot t}$$


MTBF in seconds

f1 and f2 in Hz

K1 = $1.5 \cdot 10^{-10}$ seconds (measured for XC3020-70)

K2 = $\ln 40 / \text{ns} = 3.69 \cdot 10^9$ per second (XC3020-70)

For a 10 MHz clock and approximately 1 MHz data rate, the table below gives the expected MTBF as a function of the acceptable extra delay at the output of the metastable-going flip-flop.



Extra Delay	MTBF
1.0 ns	27 milliseconds
4.2 ns	1 hour
6.7 ns	423 days
8.5 ns	890 years
10.0 ns	225,000 years
11.0 ns	9 million years
12.0 ns	360 million years

Figure 8. Metastable MTBF as a Function of Additional Acceptable Delay

Battery Back-up

Since Logic Cell Arrays are manufactured using a high-performance low-power CMOS process, they can preserve the configuration data stored in the internal static memory cells even during a loss of primary power. This is accomplished by forcing the device into a low-power non-operational state, while supplying the minimal current requirement of V_{CC} from a battery.

Circuit techniques used in XC3100 and XC4000 devices prevent I_{CC} from being reduced to the level need for battery back-up. Consequently, battery back-up should only be used for XC2000, XC3000 and XC3000A devices.

There are two primary considerations for battery backup which must be accomplished by external circuits.

- Control of the Power-Down (PWRDWN) pin
- Switching between the primary V_{CC} supply and the battery.

Important considerations include the following.

- Insure that PWRDWN is asserted logic Low prior to V_{CC} falling, is held Low while the primary V_{CC} is absent, and returned High after V_{CC} has returned to a normal level. PWRDWN edges must not rise or fall slowly.
- Insure "glitch-free" switching of the power connections to the LCA device from the primary V_{CC} to the battery and back.
- Insure that, during normal operation, the LCA V_{CC} is maintained at an acceptable level, $5.0 \text{ V} \pm 5\%$ ($\pm 10\%$ for Industrial and Military).

Figure 9 shows a power-down circuit developed by Shel Epstein of Epstein Associates, Wilmette, IL. Two Schottky diodes power the LCA from either the 5.2 V primary supply or a 3 V Lithium battery. A Seiko S8054 3-terminal power monitor circuit monitors V_{CC} and pulls PWRDWN Low whenever V_{CC} falls below 4 V.

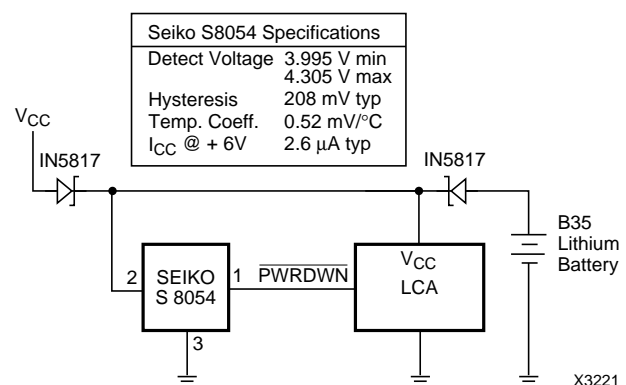


Figure 9. Battery Back-up Circuit