



XC3000 Series Technical Information

XAPP 024 November 24, 1997 (Version 1.0)

Application Note By Peter Alfke and Bernie New

Summary

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

Xilinx Family

XC3000/XC3000A/XC3000L/XC3100/XC3100A/XC3100L

Contents

- CLBs
 - Function Generators
 - Flip-flops
 - Longline Access
- IOBs
 - Inputs
 - Outputs
- Routing
 - Horizontal Longlines
 - Bus contention
 - Vertical Longlines
 - Vertical Longlines
 - Clock Buffers
 - Vertical Longlines
 - Clock Buffers
- Power Dissipation
- Crystal Oscillator
- CCLK Frequency Stability and Low-time restriction
- Powerdown and Battery-Backup
- Configuration and Start-Up
- Reset
 - Beware of slow rise-time

Introduction

The background information provided in this Application Note supplements the XC3000, XC3000A, XC3000L, XC3100A and XC3100L 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 six 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](#), contains a combinatorial function generator and two D-type flip-flops. Two output pins may be driven by either the 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 variables 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.

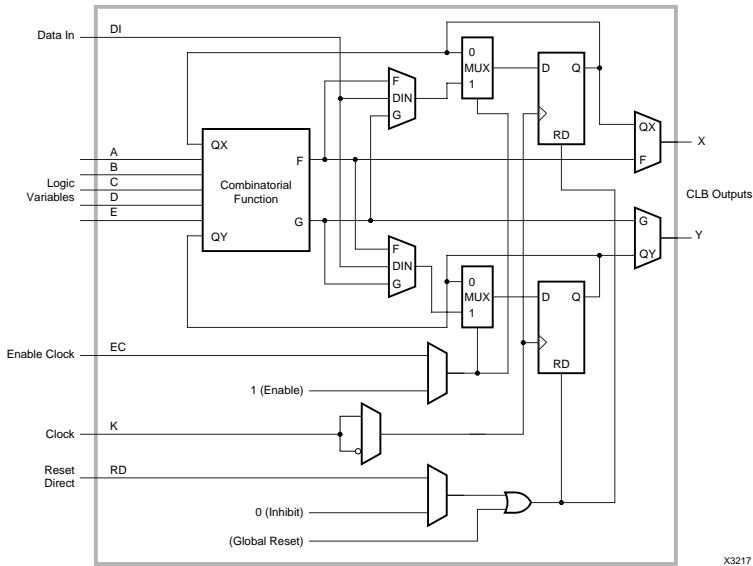


Figure 1: Configurable Logic Block (CLB)

Function Generator Avoids Glitches

The combinatorial logic in all CLBs is implemented as a function generator in the form of a multiplexer, built out of transfer gates. The logic inputs form the select inputs to this multiplexer, while the configuration bits drive the data inputs to the multiplexer.

The Xilinx circuit designers were very careful to achieve a balanced design with similar (almost equal) propagation delays from the various select inputs to the data output.

The delay from the data inputs to the output is, of course, immaterial, since the data inputs do not change dynamically. They are only affected by configuration.

This balanced design minimizes the duration of possible decoding glitches when more than one select input changes. Note that there can never be a decoding glitch when only one select input changes. Even a non-overlapping decoder cannot generate a glitch problem, since the node capacitance will retain the previous logic level until the new transfer gate is activated about a nanosecond later.

When more than one input changes “simultaneously,” the user should analyze the logic output for any possible intermediate code. If any such code permutation produces a different result, the user must assume that such a glitch might occur and must make the system design immune to it. The glitch might be only a few nanoseconds long, but that is long enough to upset an asynchronous design.

If none of the possible address sequences produces a different result, the user can be sure that there will be no glitch.

The designer of synchronous systems generally doesn't worry about such glitches, since synchronous designs are fundamentally immune to glitches on all signals except clocks or direct SET/RESET inputs.




Figure 2: CLB Logic Options

The automatic logic-partitioning software in the XACT^{step} 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 CLBMAP symbols 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.

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 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 to 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).

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	

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 kΩ. 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.

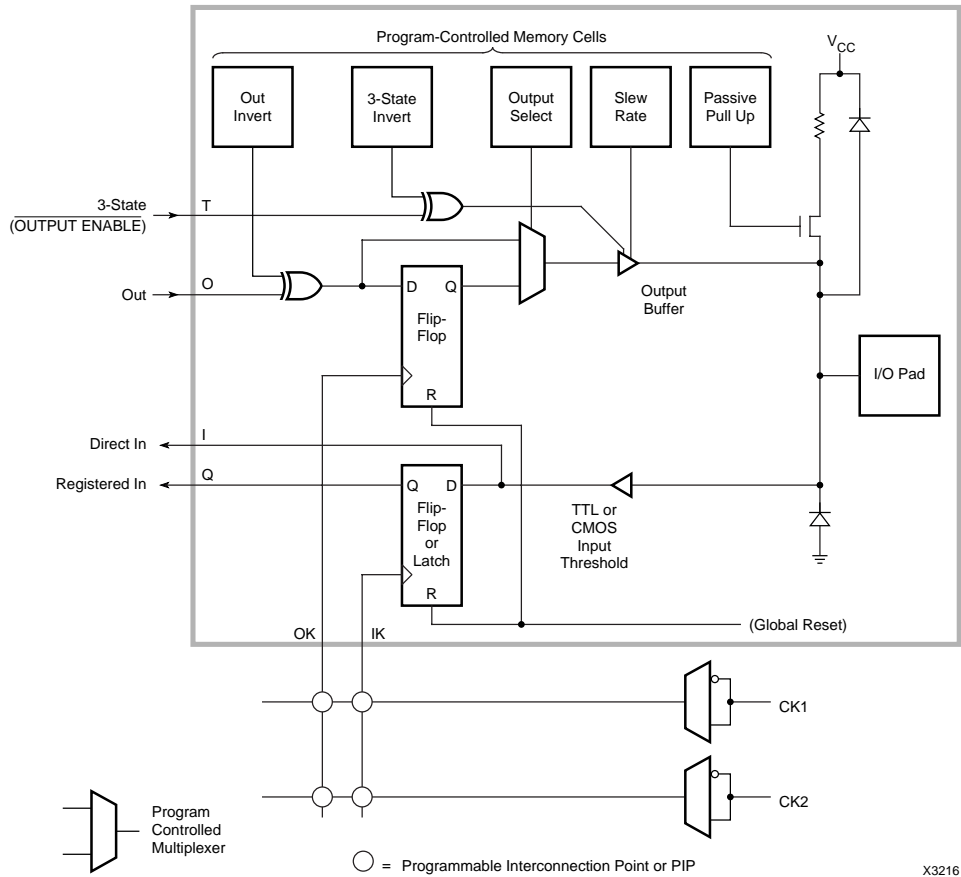


Figure 3: Input/Output Block (IOB)

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.

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 FPGA 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. Figure 4 and Figure 5 show output current/voltage curves for typical XC3000 and XC3100 devices.

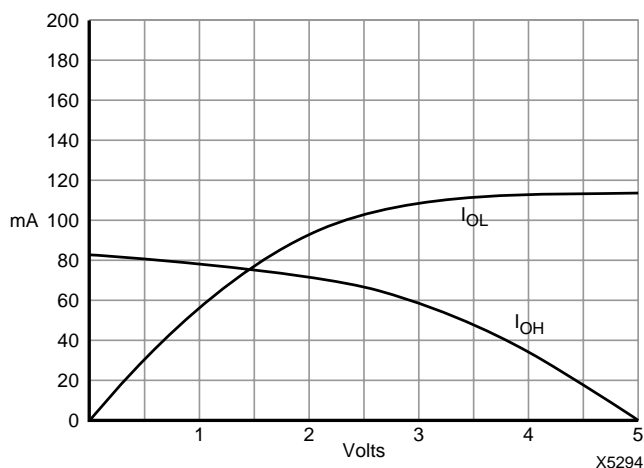


Figure 4: Output Current/Voltage Characteristics for XC3000, XC3000A, XC3100 and XC3100A Devices

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 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 FPGA 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.

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.

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.

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.

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

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.

Internal Bus Contention

XC3000 and XC4000 Series devices have internal 3-state bus drivers (TBUFs). As in any other bus design, such bus drivers must be enabled carefully in order to avoid, or at least minimize, bus contention. (Bus contention means that one driver tries to drive the bus High while a second driver tries to drive it Low).

Since the potential overlap of the enable signals is lay-out dependent, bus contention is the responsibility of the FPGA user. We can only supply the following information:

While two internal buffers drive conflicting data, they create a current path of typically 6 mA. This current is tolerable, but should not last indefinitely, since it exceeds our (conser-

vative) current density rules. A continuous contention could, after thousands of hours, lead to metal migration problems.

In a typical system, 10 ns of internal bus contention at 5 MHz would just result in a slight increase in I_{cc} .

$16 \text{ bits} \times 6 \text{ mA} \times 10 \text{ ns} \times 5 \text{ MHz} \times 50\% \text{ probability} = 2.5 \text{ mA}$.

There is a special use of the 3-state control input: When it is directly driven by the same signal that drives the data input of the buffer, i.e. when D and T are effectively tied together, the 3-state buffer becomes an “open collector” driver. Multiple drivers of this type can be used to implement the “wired-AND” function, using resistive pull-up.

In this situation there cannot be any contention, since the 3-state control input is designed to be slow in activating and fast in deactivating the driver. Connecting D to ground is an obvious alternative, but may be more difficult to route.

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.

Power Dissipation

As in most CMOS ICs, almost all FPGA 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%. Undoubtedly, 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 5, comprises a high-speed, high-gain inverting amplifier with its input connected to the dedicated XTAL2 pin, and its output connected to the XTAL1 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 series form the load on the crystal. This load is specified by the crystal manufacturer, and is typically 20 pF. The capacitors should be approximately equal: 40 pF each for a 20 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.

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.

Crystal operation below 1 MHz is not supported. Low-frequency crystals have a high resonant impedance and require more gain than provided by the single stage inverter in the XC3000 devices. Low-frequency applications are usually also more power-conscious and would not accept the power consumption of the fast general-purpose Xilinx oscillator circuit. Inexpensive complete oscillator packages are often a better choice.




Figure 5: Crystal Oscillator

Table 6: Third-Harmonic Crystal Oscillator Tank-Circuit

Frequency (MHz)	LC Tank			R2 (Ω)	C1 (pF)
	L (μ H)	C2 (pF)	Freq (MHz)		
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

Crystal-Oscillator Considerations

There is nothing Xilinx-specific about the oscillator circuit. It's a wide-band inverting amplifier, as used in all popular microcontrollers. When a crystal and some passive components close the feedback path, this circuit becomes a reliable and stable clock source.

The path from XTAL2 to XTAL1 inside the LCA device is a single-stage inverting amplifier, which means it has a low-frequency phase response of 180°, increasing by 45° at the 3-dB frequency.

Input impedance is 10-15 pF, input threshold is CMOS, but dc bias must be supplied externally through a megohm resistor from XTAL1 to XTAL2.

Low-frequency gain is about 10, rolling off 3dB at 125 MHz.

Output impedance is between 50 and 100 Ω and the capacitance on the output pin is 10 to 15 pF.

Pulse response is a delay of about 1.5 ns and a rise/fall time of about 1.5 ns.

For stable oscillation,

- the loop gain must be exactly one, i.e., the internal gain must be matched by external attenuation, and
- the phase shift around the loop must be 360° or an integer multiple thereof. The external network must, therefore, provide 180° of phase shift.

A crystal is a piezoelectric mechanical resonator that can be modeled by a very high-Q series LC circuit with a small resistor representing the energy loss. In parallel with this series-resonant circuit is unavoidable parasitic capacitance inside and outside the crystal package, and usually also discrete capacitors on the board.

The impedance as a function of frequency of this whole array starts as a small capacitor at low frequencies (Figure 6). As the frequency increases, this capacitive reactance decreases rapidly, until it reaches zero at the series resonant frequency.




Figure 6: Reactance as a Function of Frequency

At slightly higher frequencies, the reactance is inductive, starting with a zero at series resonance, and increasing very rapidly with frequency. It reaches infinity when the effective inductive impedance of the series LC circuit equals the reactance of the parallel capacitor. The parallel resonance frequency is a fraction of a percent above the series-resonance frequency.

Over this very narrow frequency range between series and parallel resonance, the crystal impedance is inductive and changes all the way from zero to infinity. The energy loss represented by the series resistor prevents the impedance from actually reaching zero and infinity, but it comes very close.

Microprocessor- and FPGA-based crystal oscillators all operate in this narrow frequency band, where the crystal impedance can be any inductive value. The circuit oscillates at a frequency where the attenuation in the external

circuit equals the gain in the FPGA device, and where the total phase shift, internal plus external, equals 360°.

Figure 7 explains the function. At the frequency of oscillation, the series-resonant circuit is effectively an inductor, and the two capacitors act as a capacitive voltage divider, with the center-point grounded. This puts a virtual ground somewhere along the inductor and causes the non-driven end of the crystal to be 180° out of phase with the driven end, which is the external phase shift required for oscillation. This circuit is commonly known as a Pierce oscillator.




Figure 7: Pierce Oscillator

Practical Considerations

- The series resonance resistor is a critical parameter. To assure reliable operation with worst-case crystals, the user should experiment with a discrete series resistor roughly equal to the max internal resistance specified by the crystal vendor. If the circuit tolerates this additional loss, it should operate reliably with a worst-case crystal without the additional resistor.
- The two capacitors affect the frequency of oscillation and the start-up conditions. The series connection of the two capacitors is the effective capacitive load seen by the crystal, usually specified by the crystal vendor.
- The two capacitors also determine the minimum gain required for oscillation. If the capacitors are too small, more gain is needed, and the oscillator may be unstable. If the capacitors are too large, oscillation is stable but the required gain may again be higher. There is an optimum capacitor value, where oscillation is stable, and the required gain is at a minimum. For most crystals, this capacitive load is around 20 pF, i.e., each of the two capacitors should be around 40 pF.
- Crystal dissipation is usually around 1 mW, and thus of no concern. Beware of crystals with “drive-level dependence” of the series resistor. They may not start up. Proper drive level can be checked by varying Vcc. The frequency should increase slightly with an increase in Vcc. A decreasing frequency or unstable amplitude indicate an over-driven crystal. Excessive swing at the

XTAL2 input results in clipping near V_{CC} and ground. An additional 1 to 2 k Ω series resistor at the XTAL1 output usually cures that distortion problem. It increases the amplifier output impedance and assures additional phase margin, but results in slower start-up.

- Be especially careful when designing an oscillator that must operate near the specified max frequency. The circuit needs excess gain at small signal amplitudes to supply enough energy into the crystal for rapid start-up. High-frequency gain may be marginal, and start-up may be impaired.
- Keep the whole oscillator circuit physically as compact as possible, and provide a single ground connection. Grounding the crystal can is not mandatory but may improve stability.

Series Resonant or Parallel Resonant?

Crystal manufacturers label some crystals as series-resonant, others as parallel-resonant, but there really is no difference between these two types of crystals, they all operate in the same way. Every crystal has a series resonance, where the impedance of the crystal is extremely low, much lower than at any other frequency. At a slightly higher frequency, the crystal is inductive and in parallel resonance with the unavoidable stray capacitance or the deliberate capacitance between its pins.

The only difference between the two types of crystal is the manufacturer's choice of specifying either of the two frequencies. If series resonance is specified, the actual frequency of oscillation is a little higher than the specified value. If parallel resonance is specified, the frequency of oscillation is a little lower. In most cases, these small deviations are irrelevant.

CCLK Frequency Variation

The on-chip R-C oscillator that is brought out as CCLK also performs several other internal functions. 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 -50%. 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: Typical CCLK Frequency Variation

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

CCLK Low-Time Restriction

When used as an input in Slave Serial and Readback modes, CCLK does not tolerate a Low time in excess of 5 μ s. For very low speed operation, the CCLK High time can be stretched to any value, but the Low time must be kept short. XC4000 and XC5200 devices do not have this restriction.

Battery Back-up

Since SRAM-based FPGAs 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, XC4000 and XC5200 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, XC2000L, XC3000, XC3000A and XC3000L devices.

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

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

Important considerations include the following.

- Insure that \overline{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. \overline{PWRDWN} edges must not rise or fall slowly.
- Insure "glitch-free" switching of the power connections to the FPGA device from the primary V_{CC} to the battery and back.
- Insure that, during normal operation, the FPGA V_{CC} is maintained at an acceptable level, 5.0 V \pm 5% (\pm 10% for Industrial and Military).

Figure 8 shows a power-down circuit developed by Shel Epstein of Epstein Associates, Wilmette, IL. Two Schottky diodes power the FPGA 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 $\overline{\text{PWRDWN}}$ Low whenever V_{CC} falls below 4 V.

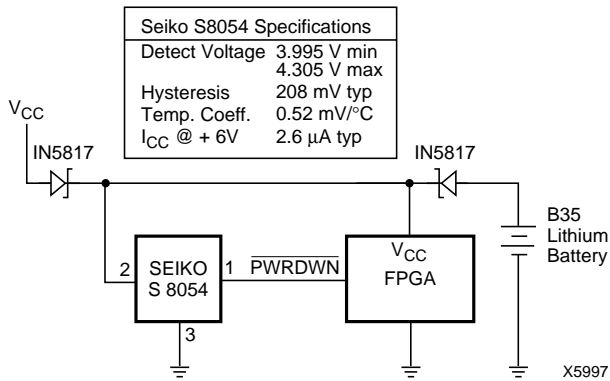


Figure 8: Battery Back-up Circuit

Powerdown Operation

A Low level on the $\overline{\text{PWRDWN}}$ input, while V_{CC} remains higher than 2.3 V, stops all internal activity, thus reducing I_{CC} to a very low level:

- All internal pull-ups (on Long lines as well as on the I/O pads) are turned off.
- The crystal oscillator is turned off
- All package outputs are three-stated.
- All package inputs ignore the actual input level, and present a High to the internal logic.
- All internal flip-flops or latches are permanently reset.
- The internal configuration is retained.
- When $\overline{\text{PWRDWN}}$ is returned High, after V_{CC} is at its nominal value, the device returns to operation with the same sequence of buffer enable and $\overline{\text{D/P}}$ as at the completion of configuration.

Things to Remember

Powerdown retains the configuration, but loses all data stored in the device. Powerdown three-states all outputs and ignores all inputs. No clock signal will be recognized, and the crystal oscillator is stopped. All internal flip-flops and latches are permanently reset and all inputs are interpreted as High, but the internal combinatorial logic is fully functional.

Things to Watch Out For

Make sure that the combination of all inputs High and all internal flip-flop outputs Low in your design will not generate internal oscillations or create permanent bus contention

by activating internal bus drivers with conflicting data onto the same Longline. These two situations are farfetched, but they are possible and will result in considerable power consumption. It is quite easy to simulate these conditions since all inputs are stable and the internal logic is entirely combinatorial, unless latches have been made out of function generators.

During powerdown, the V_{CC} monitoring circuit is disabled. It is then up to the user to prevent V_{CC} dips below 2.3 V, which would corrupt the stored configuration.

During configuration, the $\overline{\text{PWRDWN}}$ pin must be High, since configuration uses the internal oscillator. Whenever V_{CC} goes below 4 V, $\overline{\text{PWRDWN}}$ must already be Low in order to prevent automatic reconfiguration at low V_{CC} . For the same reason, V_{CC} must first be restored to 4 V or more, before $\overline{\text{PWRDWN}}$ can be made High.

$\overline{\text{PWRDWN}}$ has no pull-up resistor. A pull-up resistor would draw supply current when the pin is Low, which would defeat the idea of powerdown, where I_{CC} is only microamperes.

Configuration and Start-up

Start-Up

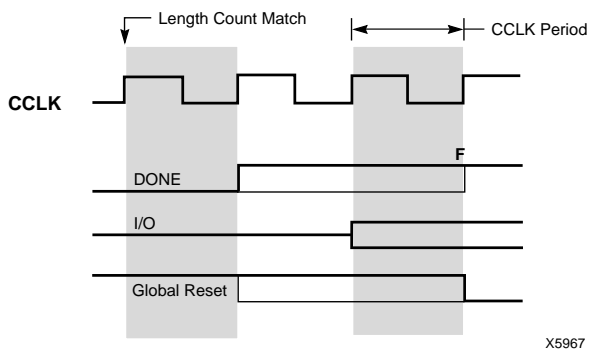
Start-up is the transition from the configuration process to the intended user operation. This means a change from one clock source to another, and a change from interfacing parallel or serial configuration data where most outputs are 3-stated, to normal operation with I/O pins active in the user-system. 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 10 describes Start-up timing for the XC3000 families in detail.

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 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 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 11, but the designer can modify it to meet particular requirements.

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 FPGA devices hooked up in a daisy chain will all go active simultaneously



X5967

Figure 9: Start-up Timing

on the same CCLK edge. This is well documented in the data sheets.

Not documented, however, is how the internal combinational 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 FPGA 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 10 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, \overline{LDC} is asserted Low and holds the D-input of the flip-flop High, while Q is held Low by the internal reset, and \overline{RESET} is kept High by internal and external pull-up resistors. At the end of configuration, the \overline{LDC} pin is

unasserted, but D remains High since the function generator acts as an R-S latch; Q stays Low, and \overline{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 \overline{RESET} Low. This resets the whole chip until the Low on Q permits \overline{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 \overline{LDC} prevents the R-S latch from becoming set.


Beware of a Slow-Rising XC3000 Series \overline{RESET} Input

It is a wide-spread habit to drive asynchronous \overline{RESET} inputs with a resistor-capacitor network to lengthen the reset time after power-on. This can also be done with Xilinx FPGAs, but the user should question the need, and should beware of certain avoidable problems.

Xilinx FPGAs contain an internal voltage-monitoring circuit, and start their internal housekeeping operation only after V_{CC} has reached ~ 3.5 V. The internal housekeeping and configuration memory clearing operation then takes between about 10 and 100 ms, depending on configuration mode and processing variations. Any RC delay shorter than 40 ms for a device in master configuration mode, or shorter than 10 ms for a device in slave configuration mode, is clearly redundant.

A significantly longer RC delay can be used to hold off configuration. Without the use of an external Schmitt trigger circuit, the rise time on the \overline{RESET} input will be very slow, and is likely to cross the threshold of ~ 1.4 V several times, due to external or internal noise. This can cause the FPGA to start configuration, then immediately abort it, then start it again, after having automatically cleared the configuration memory once more.

This is no problem for the FPGA, but it requires that the source of configuration data, especially an XC1700 serial PROM, be reset accordingly. This is another reason to use the \overline{INIT} output of the lead FPGA, instead of \overline{LDC} , to drive the \overline{RESET} input of the XC1700 serial PROMs.



X3222

Figure 10: Synchronous Reset