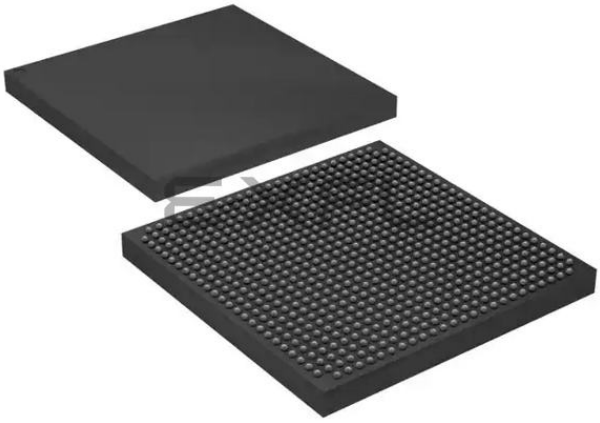

Welcome to E-XFL.COM

Understanding Embedded - FPGAs (Field Programmable Gate Array)

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

Applications of Embedded - FPGAs

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

Details	
Product Status	Obsolete
Number of LABs/CLBs	780
Number of Logic Elements/Cells	15600
Total RAM Bits	419328
Number of I/O	366
Number of Gates	-
Voltage - Supply	1.15V ~ 1.25V
Mounting Type	Surface Mount
Operating Temperature	0°C ~ 85°C (TJ)
Package / Case	672-BBGA
Supplier Device Package	672-FBGA (27x27)
Purchase URL	https://www.e-xfl.com/product-detail/intel/ep2s15f672c5n

Chapter 6. Reference & Ordering Information

1. Introduction

SII51001-4.2

Introduction

The Stratix® II FPGA family is based on a 1.2-V, 90-nm, all-layer copper SRAM process and features a new logic structure that maximizes performance, and enables device densities approaching 180,000 equivalent logic elements (LEs). Stratix II devices offer up to 9 Mbits of on-chip, TriMatrix™ memory for demanding, memory intensive applications and has up to 96 DSP blocks with up to 384 (18-bit × 18-bit) multipliers for efficient implementation of high performance filters and other DSP functions. Various high-speed external memory interfaces are supported, including double data rate (DDR) SDRAM and DDR2 SDRAM, RLDRAM II, quad data rate (QDR) II SRAM, and single data rate (SDR) SDRAM. Stratix II devices support various I/O standards along with support for 1-gigabit per second (Gbps) source synchronous signaling with DPA circuitry. Stratix II devices offer a complete clock management solution with internal clock frequency of up to 550 MHz and up to 12 phase-locked loops (PLLs). Stratix II devices are also the industry's first FPGAs with the ability to decrypt a configuration bitstream using the Advanced Encryption Standard (AES) algorithm to protect designs.

Features

The Stratix II family offers the following features:

- 15,600 to 179,400 equivalent LEs; see Table 1–1
- New and innovative adaptive logic module (ALM), the basic building block of the Stratix II architecture, maximizes performance and resource usage efficiency
- Up to 9,383,040 RAM bits (1,172,880 bytes) available without reducing logic resources
- TriMatrix memory consisting of three RAM block sizes to implement true dual-port memory and first-in first-out (FIFO) buffers
- High-speed DSP blocks provide dedicated implementation of multipliers (at up to 450 MHz), multiply-accumulate functions, and finite impulse response (FIR) filters
- Up to 16 global clocks with 24 clocking resources per device region
- Clock control blocks support dynamic clock network enable/disable, which allows clock networks to power down to reduce power consumption in user mode
- Up to 12 PLLs (four enhanced PLLs and eight fast PLLs) per device provide spread spectrum, programmable bandwidth, clock switch-over, real-time PLL reconfiguration, and advanced multiplication and phase shifting

Figure 2–3. Direct Link Connection

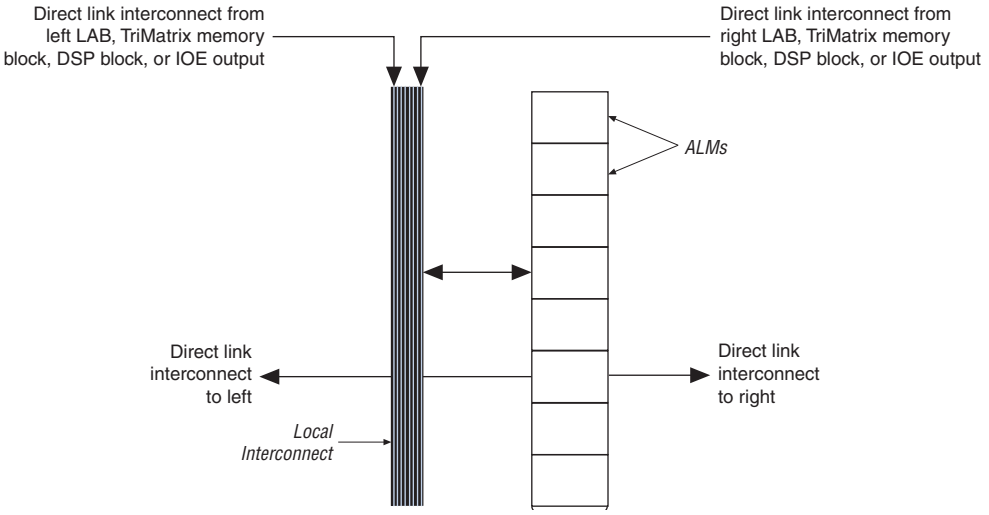

LAB Control Signals

Each LAB contains dedicated logic for driving control signals to its ALMs. The control signals include three clocks, three clock enables, two asynchronous clears, synchronous clear, asynchronous preset/load, and synchronous load control signals. This gives a maximum of 11 control signals at a time. Although synchronous load and clear signals are generally used when implementing counters, they can also be used with other functions.

Each LAB can use three clocks and three clock enable signals. However, there can only be up to two unique clocks per LAB, as shown in the LAB control signal generation circuit in Figure 2–4. Each LAB's clock and clock enable signals are linked. For example, any ALM in a particular LAB using the `labclk1` signal also uses `labclkena1`. If the LAB uses both the rising and falling edges of a clock, it also uses two LAB-wide clock signals. De-asserting the clock enable signal turns off the corresponding LAB-wide clock.

Each LAB can use two asynchronous clear signals and an asynchronous load/preset signal. By default, the Quartus II software uses a `NOT` gate push-back technique to achieve preset. If you disable the `NOT` gate push-up option or assign a given register to power up high using the Quartus II software, the preset is achieved using the asynchronous load

Clear & Preset Logic Control

LAB-wide signals control the logic for the register's clear and load/preset signals. The ALM directly supports an asynchronous clear and preset function. The register preset is achieved through the asynchronous load of a logic high. The direct asynchronous preset does not require a NOT-gate push-back technique. Stratix II devices support simultaneous asynchronous load/preset, and clear signals. An asynchronous clear signal takes precedence if both signals are asserted simultaneously. Each LAB supports up to two clears and one load/preset signal.

In addition to the clear and load/preset ports, Stratix II devices provide a device-wide reset pin (DEV_CLRn) that resets all registers in the device. An option set before compilation in the Quartus II software controls this pin. This device-wide reset overrides all other control signals.

MultiTrack Interconnect

In the Stratix II architecture, connections between ALMs, TriMatrix memory, DSP blocks, and device I/O pins are provided by the MultiTrack interconnect structure with DirectDrive™ technology. The MultiTrack interconnect consists of continuous, performance-optimized routing lines of different lengths and speeds used for inter- and intra-design block connectivity. The Quartus II Compiler automatically places critical design paths on faster interconnects to improve design performance.

DirectDrive technology is a deterministic routing technology that ensures identical routing resource usage for any function regardless of placement in the device. The MultiTrack interconnect and DirectDrive technology simplify the integration stage of block-based designing by eliminating the re-optimization cycles that typically follow design changes and additions.

The MultiTrack interconnect consists of row and column interconnects that span fixed distances. A routing structure with fixed length resources for all devices allows predictable and repeatable performance when migrating through different device densities. Dedicated row interconnects route signals to and from LABs, DSP blocks, and TriMatrix memory in the same row. These row resources include:

- Direct link interconnects between LABs and adjacent blocks
- R4 interconnects traversing four blocks to the right or left
- R24 row interconnects for high-speed access across the length of the device

Figure 2–17. Shared Arithmetic Chain, Carry Chain & Register Chain Interconnects

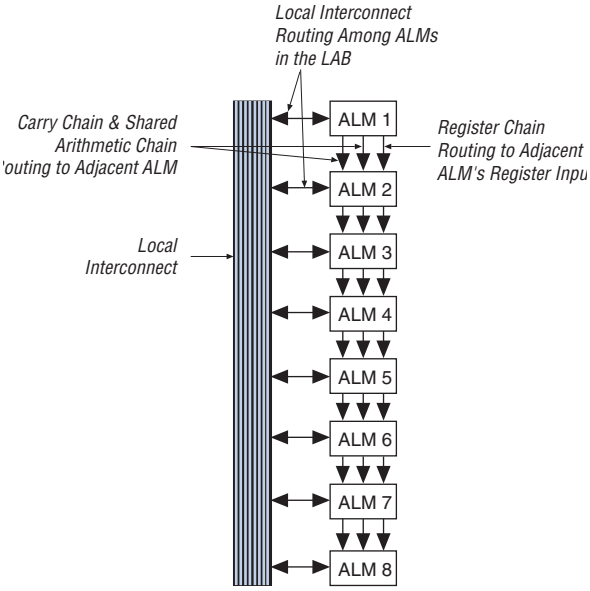

The C4 interconnects span four LABs, M512, or M4K blocks up or down from a source LAB. Every LAB has its own set of C4 interconnects to drive either up or down. Figure 2–18 shows the C4 interconnect connections from an LAB in a column. The C4 interconnects can drive and be driven by all types of architecture blocks, including DSP blocks, TriMatrix memory blocks, and column and row IOEs. For LAB interconnection, a primary LAB or its LAB neighbor can drive a given C4 interconnect. C4 interconnects can drive each other to extend their range as well as drive row interconnects for column-to-column connections.

Figure 2–25. M-RAM Block LAB Row Interface *Note (1)*

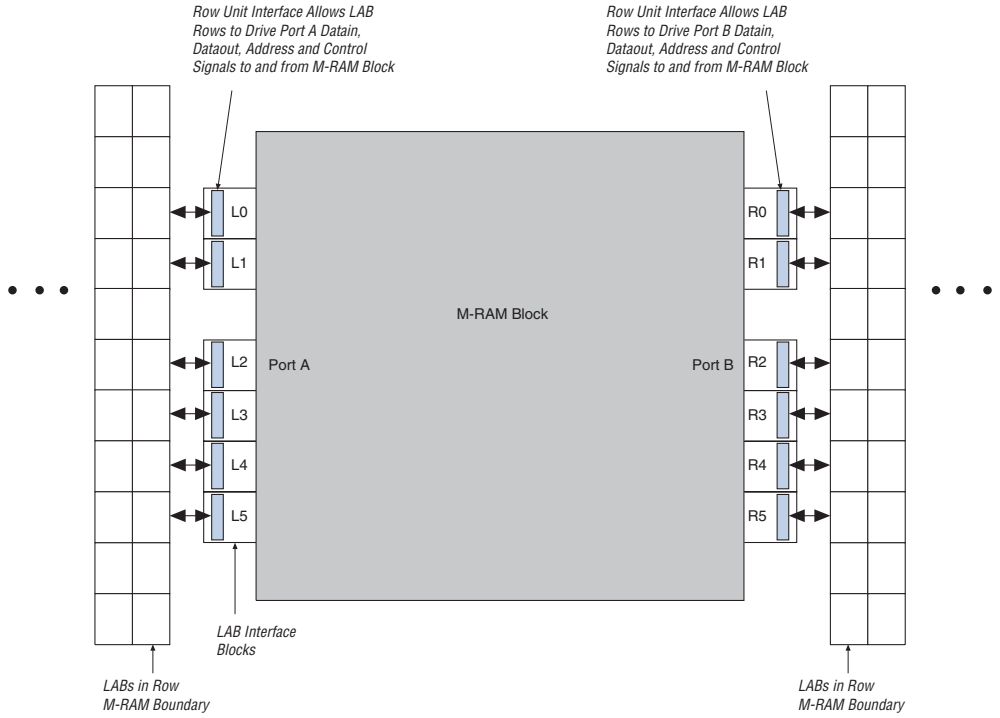

Note to Figure 2–25:
(1) Only R24 and C16 interconnects cross the M-RAM block boundaries.

Figure 2–27 shows one of the columns with surrounding LAB rows.

Figure 2–27. DSP Blocks Arranged in Columns

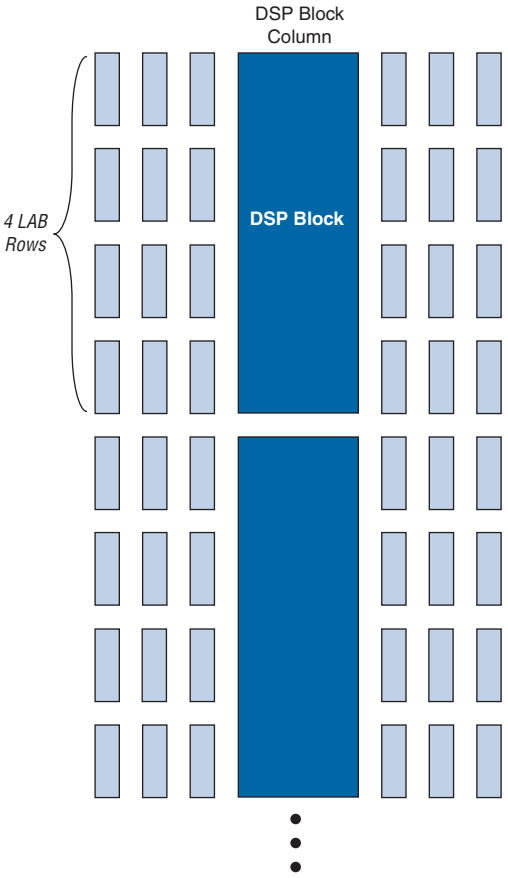

Figure 2–28. DSP Block Diagram for 18 × 18-Bit Configuration

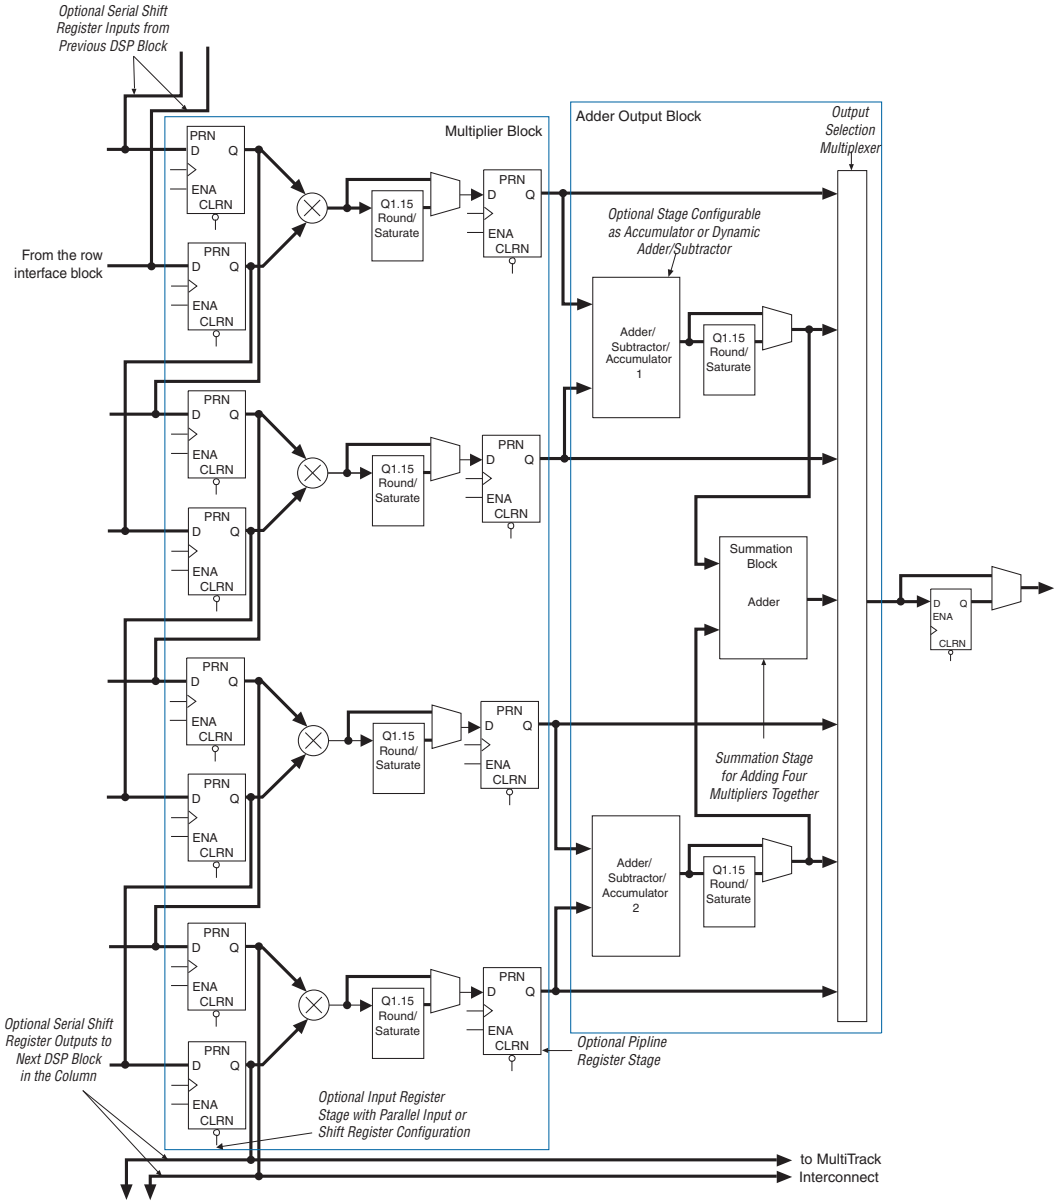

The LAB row source for control signals, data inputs, and outputs is shown in Table 2–7.

Table 2–7. DSP Block Signal Sources & Destinations

LAB Row at Interface	Control Signals Generated	Data Inputs	Data Outputs
0	clock0 aclr0 ena0 mult01_saturate addnsub1_round/ accum_round addnsub1 signa sourcea sourceb	A1[17..0] B1[17..0]	OA[17..0] OB[17..0]
1	clock1 aclr1 ena1 accum_saturate mult01_round accum_sload sourcea sourceb mode0	A2[17..0] B2[17..0]	OC[17..0] OD[17..0]
2	clock2 aclr2 ena2 mult23_saturate addnsub3_round/ accum_round addnsub3 sign_b sourcea sourceb	A3[17..0] B3[17..0]	OE[17..0] OF[17..0]
3	clock3 aclr3 ena3 accum_saturate mult23_round accum_sload sourcea sourceb mode1	A4[17..0] B4[17..0]	OG[17..0] OH[17..0]

See the *DSP Blocks in Stratix II & Stratix II GX Devices* chapter in volume 2 of the *Stratix II Device Handbook* or the *Stratix II GX Device Handbook,* for more information on DSP blocks.

Figure 2–32. Regional Clocks

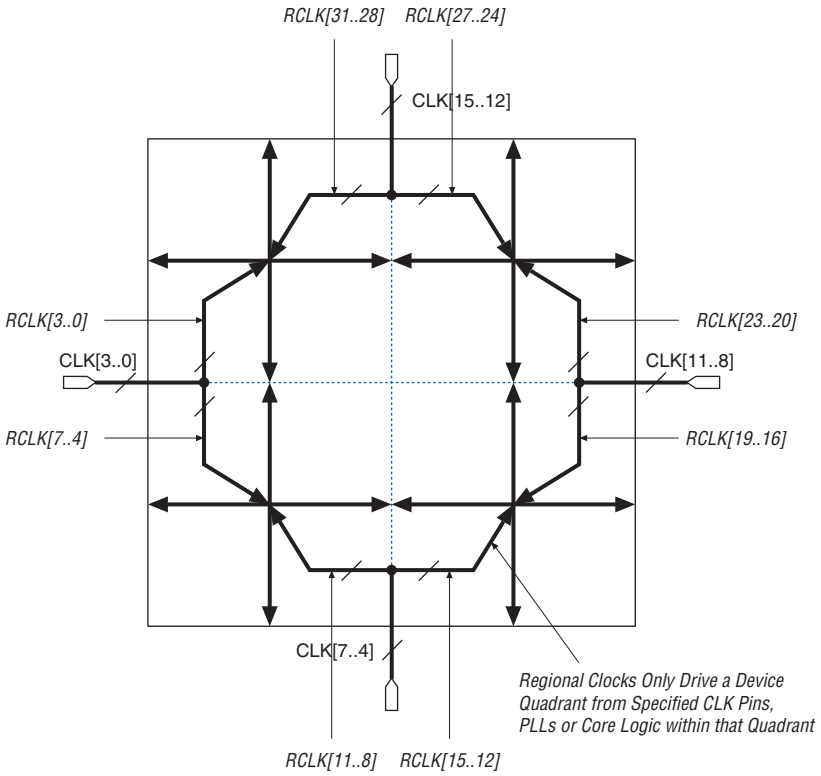

Dual-Regional Clock Network

A single source (CLK pin or PLL output) can generate a dual-regional clock by driving two regional clock network lines in adjacent quadrants (one from each quadrant). This allows logic that spans multiple quadrants to utilize the same low skew clock. The routing of this clock signal on an entire side has approximately the same speed but slightly higher clock skew when compared with a clock signal that drives a single quadrant. Internal logic-array routing can also drive a dual-regional clock. Clock pins and enhanced PLL outputs on the top and bottom can drive horizontal dual-regional clocks. Clock pins and fast PLL outputs on the left and right can drive vertical dual-regional clocks, as shown in Figure 2–33. Corner PLLs cannot drive dual-regional clocks.

The Stratix II clock networks can be disabled (powered down) by both static and dynamic approaches. When a clock net is powered down, all the logic fed by the clock net is in an off-state thereby reducing the overall power consumption of the device.

The global and regional clock networks can be powered down statically through a setting in the configuration (**.sof** or **.pof**) file. Clock networks that are not used are automatically powered down through configuration bit settings in the configuration file generated by the Quartus II software.

The dynamic clock enable/disable feature allows the internal logic to control power up/down synchronously on GCLK and RCLK nets and PLL_OUT pins. This function is independent of the PLL and is applied directly on the clock network or PLL_OUT pin, as shown in Figures 2–37 through 2–39.

The following restrictions for the input clock pins apply:

- CLK0 pin -> inclk[0] of CLKCTRL
- CLK1 pin -> inclk[1] of CLKCTRL
- CLK2 pin -> inclk[0] of CLKCTRL
- CLK3 pin -> inclk[1] of CLKCTRL

In general, even CLK numbers connect to the inclk[0] port of CLKCTRL, and odd CLK numbers connect to the inclk[1] port of CLKCTRL.

Failure to comply with these restrictions will result in a no-fit error.

Enhanced & Fast PLLs

Stratix II devices provide robust clock management and synthesis using up to four enhanced PLLs and eight fast PLLs. These PLLs increase performance and provide advanced clock interfacing and clock-frequency synthesis. With features such as clock switchover, spread-spectrum clocking, reconfigurable bandwidth, phase control, and reconfigurable phase shifting, the Stratix II device's enhanced PLLs provide you with complete control of clocks and system timing. The fast PLLs provide general purpose clocking with multiplication and phase shifting as well as high-speed outputs for high-speed differential I/O support. Enhanced and fast PLLs work together with the Stratix II high-speed I/O and advanced clock architecture to provide significant improvements in system performance and bandwidth.

Figure 2–43 shows the global and regional clocking from enhanced PLL outputs and top and bottom CLK pins. The connections to the global and regional clocks from the top clock pins and enhanced PLL outputs is shown in Table 2–11. The connections to the clocks from the bottom clock pins is shown in Table 2–12.

Table 2–11. Global & Regional Clock Connections from Top Clock Pins & Enhanced PLL Outputs (Part 1 of 2)

Top Side Global & Regional Clock Network Connectivity	DLLCLK	CLK12	CLK13	CLK14	CLK15	RCLK24	RCLK25	RCLK26	RCLK27	RCLK28	RCLK29	RCLK30	RCLK31
Clock pins													
CLK12p	✓	✓	✓			✓				✓			
CLK13p	✓	✓	✓				✓						✓
CLK14p	✓			✓	✓			✓				✓	
CLK15p	✓			✓	✓				✓		✓		
CLK12n		✓				✓				✓			
CLK13n			✓				✓						✓
CLK14n				✓				✓				✓	
CLK15n					✓				✓		✓		
Drivers from internal logic													
GCLKDRV0		✓											
GCLKDRV1			✓										
GCLKDRV2				✓									
GCLKDRV3					✓								
RCLKDRV0						✓				✓			
RCLKDRV1							✓				✓		
RCLKDRV2								✓				✓	
RCLKDRV3									✓				✓
RCLKDRV4						✓				✓			
RCLKDRV5							✓				✓		
RCLKDRV6								✓				✓	
RCLKDRV7									✓				✓
Enhanced PLL 5 outputs													
c0	✓	✓	✓			✓				✓			
c1	✓	✓	✓				✓				✓		
c2	✓			✓	✓			✓				✓	
c3	✓			✓	✓				✓				✓

There are 32 control and data signals that feed each row or column I/O block. These control and data signals are driven from the logic array. The row or column IOE clocks, `io_clk[7..0]`, provide a dedicated routing resource for low-skew, high-speed clocks. I/O clocks are generated from global or regional clocks (see the "PLLs & Clock Networks" section). Figure 2–49 illustrates the signal paths through the I/O block.

Figure 2–49. Signal Path through the I/O Block

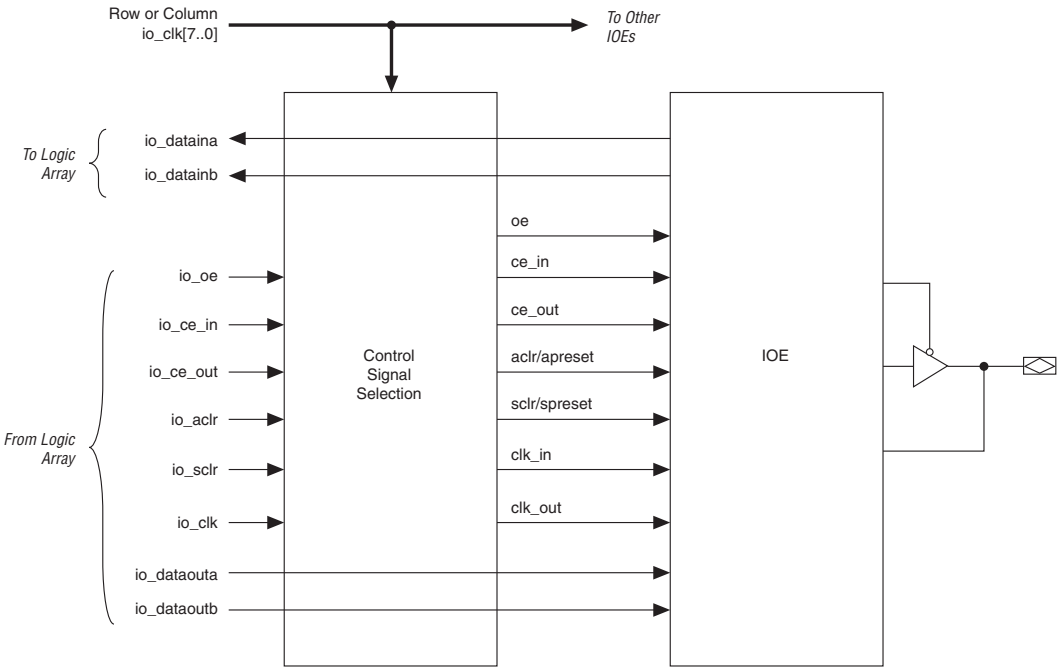

Each IOE contains its own control signal selection for the following control signals: `oe`, `ce_in`, `ce_out`, `aclr/apreset`, `sclr/spreset`, `clk_in`, and `clk_out`. Figure 2–50 illustrates the control signal selection.

Table 2–27. Document Revision History (Part 2 of 2)

Date and Document Version	Changes Made	Summary of Changes
January 2005, v2.0	• Updated the "MultiVolt I/O Interface" and "TriMatrix Memory" sections. • Updated Tables 2–3, 2–17, and 2–19.	—
October 2004, v1.2	• Updated Tables 2–9, 2–16, 2–26, and 2–27.	—
July 2004, v1.1	• Updated note to Tables 2–9 and 2–16. • Updated Tables 2–16, 2–17, 2–18, 2–19, and 2–20. • Updated Figures 2–41, 2–42, and 2–57. • Removed 3 from list of SERDES factor *J*. • Updated "High-Speed Differential I/O with DPA Support" section. • In "Dedicated Circuitry with DPA Support" section, removed XSBI and changed RapidIO to Parallel RapidIO.	—
February 2004, v1.0	Added document to the Stratix II Device Handbook.	—

The PLL_ENA pin and the configuration input pins (Table 3–4) have a dual buffer design: a 3.3-V/2.5-V input buffer and a 1.8-V/1.5-V input buffer. The VCCSEL input pin selects which input buffer is used. The 3.3-V/2.5-V input buffer is powered by $V_{CCPD,}$ while the 1.8-V/1.5-V input buffer is powered by V_{CCIO}. Table 3–4 shows the pins affected by VCCSEL.

Table 3–4. Pins Affected by the Voltage Level at VCCSEL

Pin	VCCSEL = LOW (connected to GND)	VCCSEL = HIGH (connected to V_{CCPD})
nSTATUS (when used as an input)	3.3/2.5-V input buffer is selected. Input buffer is powered by V_{CCPD}.	1.8/1.5-V input buffer is selected. Input buffer is powered by V_{CCIO} of the I/O bank.
nCONFIG		
CONF_DONE (when used as an input)		
DATA[7..0]		
nCE		
DCLK (when used as an input)		
CS		
nWS		
nRS		
nCS		
CLKUSR		
DEV_OE		
DEV_CLRn		
RUnLU		
PLL_ENA		

VCCSEL is sampled during power-up. Therefore, the VCCSEL setting cannot change on the fly or during a reconfiguration. The VCCSEL input buffer is powered by V_{CCINT} and must be hardwired to V_{CCPD} or ground. A logic high VCCSEL connection selects the 1.8-V/1.5-V input buffer, and a logic low selects the 3.3-V/2.5-V input buffer. VCCSEL should be set to comply with the logic levels driven out of the configuration device or MAX® II/microprocessor.

If you need to support configuration input voltages of 3.3 V/2.5 V, you should set the VCCSEL to a logic low; you can set the V_{CCIO} of the I/O bank that contains the configuration inputs to any supported voltage. If

4. Hot Socketing & Power-On Reset

SII51004-3.2

Stratix® II devices offer hot socketing, which is also known as hot plug-in or hot swap, and power sequencing support without the use of any external devices. You can insert or remove a Stratix II board in a system during system operation without causing undesirable effects to the running system bus or the board that was inserted into the system.

The hot socketing feature also removes some of the difficulty when you use Stratix II devices on printed circuit boards (PCBs) that also contain a mixture of 5.0-, 3.3-, 2.5-, 1.8-, 1.5- and 1.2-V devices. With the Stratix II hot socketing feature, you no longer need to ensure a proper power-up sequence for each device on the board.

The Stratix II hot socketing feature provides:

- Board or device insertion and removal without external components or board manipulation
- Support for any power-up sequence
- Non-intrusive I/O buffers to system buses during hot insertion

This chapter also discusses the power-on reset (POR) circuitry in Stratix II devices. The POR circuitry keeps the devices in the reset state until the V_{CC} is within operating range.

Stratix II Hot-Socketing Specifications

Stratix II devices offer hot socketing capability with all three features listed above without any external components or special design requirements. The hot socketing feature in Stratix II devices allows:

- The device can be driven before power-up without any damage to the device itself.
- I/O pins remain tri-stated during power-up. The device does not drive out before or during power-up, thereby affecting other buses in operation.
- Signal pins do not drive the V_{CCIO}, V_{CCPD}, or V_{CCINT} power supplies. External input signals to I/O pins of the device do not internally power the V_{CCIO} or V_{CCINT} power supplies of the device via internal paths within the device.

Clock Network Skew Adders

The Quartus II software models skew within dedicated clock networks such as global and regional clocks. Therefore, intra-clock network skew adder is not specified. Table 5–68 specifies the clock skew between any two clock networks driving registers in the IOE.

Table 5–68. Clock Network Specifications

Name	Description	Min	Typ	Max	Unit
Clock skew adder EP2S15, EP2S30, EP2S60 *(1)*	Inter-clock network, same side			±50	ps
	Inter-clock network, entire chip			±100	ps
Clock skew adder EP2S90 *(1)*	Inter-clock network, same side			±55	ps
	Inter-clock network, entire chip			±110	ps
Clock skew adder EP2S130 *(1)*	Inter-clock network, same side			±63	ps
	Inter-clock network, entire chip			±125	ps
Clock skew adder EP2S180 *(1)*	Inter-clock network, same side			±75	ps
	Inter-clock network, entire chip			±150	ps

Note to Table 5–68:
(1) This is in addition to intra-clock network skew, which is modeled in the Quartus II software.

Table 5–80. Maximum DCD for Non-DDIO Output on Row I/O Pins (Part 2 of 2) *Note (1)*

Row I/O Output Standard	Maximum DCD for Non-DDIO Output		
	-3 Devices	-4 & -5 Devices	Unit
1.8 V	180	180	ps
1.5-V LVCMOS	165	195	ps
SSTL-2 Class I	115	145	ps
SSTL-2 Class II	95	125	ps
SSTL-18 Class I	55	85	ps
1.8-V HSTL Class I	80	100	ps
1.5-V HSTL Class I	85	115	ps
LVDS/ HyperTransport technology	55	80	ps

Note to Table 5–80:
(1) The DCD specification is based on a no logic array noise condition.

Here is an example for calculating the DCD as a percentage for a non-DDIO output on a row I/O on a -3 device:

If the non-DDIO output I/O standard is SSTL-2 Class II, the maximum DCD is 95 ps (see Table 5–80). If the clock frequency is 267 MHz, the clock period T is:

T = 1/ f = 1 / 267 MHz = 3.745 ns = 3745 ps

To calculate the DCD as a percentage:

(T/2 – DCD) / T = (3745ps/2 – 95ps) / 3745ps = 47.5% (for low boundary)

(T/2 + DCD) / T = (3745ps/2 + 95ps) / 3745ps = 52.5% (for high boundary)

Table 5–92. Enhanced PLL Specifications (Part 2 of 2)

Name	Description	Min	Typ	Max	Unit
t_{LOCK}	Time required for the PLL to lock from the time it is enabled or the end of device configuration		0.03	1	ms
t_{DLOCK}	Time required for the PLL to lock dynamically after automatic clock switchover between two identical clock frequencies			1	ms
$f_{SWITCHOVER}$	Frequency range where the clock switchover performs properly	4		500	MHz
f_{CLBW}	PLL closed-loop bandwidth	0.13	1.20	16.90	MHz
f_{VCO}	PLL VCO operating range for –3 and –4 speed grade devices	300		1,040	MHz
	PLL VCO operating range for –5 speed grade devices	300		840	MHz
f_{SS}	Spread-spectrum modulation frequency	30		150	kHz
% spread	Percent down spread for a given clock frequency	0.4	0.5	0.6	%
t_{PLL_PSERR}	Accuracy of PLL phase shift			±15	ps
t_{ARESET}	Minimum pulse width on `areset` signal.	10			ns
$t_{ARESET\_RECONFIG}$	Minimum pulse width on the `areset` signal when using PLL reconfiguration. Reset the PLL after `scandone` goes high.	500			ns

Notes to Table 5–92:

(1) Limited by I/O f_{MAX}. See Table 5–78 on page 5–69 for the maximum. Cannot exceed f_{OUT} specification.

(2) If the counter cascading feature of the PLL is utilized, there is no minimum output clock frequency.