



Chipsmall Limited consists of a professional team with an average of over 10 year of expertise in the distribution of electronic components. Based in Hongkong, we have already established firm and mutual-benefit business relationships with customers from,Europe,America and south Asia,supplying obsolete and hard-to-find components to meet their specific needs.

With the principle of “Quality Parts,Customers Priority,Honest Operation,and Considerate Service”,our business mainly focus on the distribution of electronic components. Line cards we deal with include Microchip,ALPS,ROHM,Xilinx,Pulse,ON,Everlight and Freescale. Main products comprise IC,Modules,Potentiometer,IC Socket,Relay,Connector.Our parts cover such applications as commercial,industrial, and automotives areas.

We are looking forward to setting up business relationship with you and hope to provide you with the best service and solution. Let us make a better world for our industry!



Contact us

Tel: +86-755-8981 8866 Fax: +86-755-8427 6832

Email & Skype: info@chipsmall.com Web: www.chipsmall.com

Address: A1208, Overseas Decoration Building, #122 Zhenhua RD., Futian, Shenzhen, China





LatticeXP2™ Family Data Sheet

DS1009 Version 2.2, September 2014

Features

■ flexiFLASH™ Architecture

- Instant-on
- Infinitely reconfigurable
- Single chip
- FlashBAK™ technology
- Serial TAG memory
- Design security

■ Live Update Technology

- TransFR™ technology
- Secure updates with 128 bit AES encryption
- Dual-boot with external SPI

■ sysDSP™ Block

- Three to eight blocks for high performance Multiply and Accumulate
- 12 to 32 18x18 multipliers
- Each block supports one 36x36 multiplier or four 18x18 or eight 9x9 multipliers

■ Embedded and Distributed Memory

- Up to 885 Kbits sysMEM™ EBR
- Up to 83 Kbits Distributed RAM

■ sysCLOCK™ PLLs

- Up to four analog PLLs per device
- Clock multiply, divide and phase shifting

■ Flexible I/O Buffer

- sysIO™ buffer supports:
 - LVCMOS 33/25/18/15/12; LVTTTL
 - SSTL 33/25/18 class I, II
 - HSTL15 class I; HSTL18 class I, II
 - PCI
 - LVDS, Bus-LVDS, MLVDS, LVPECL, RSDS

■ Pre-engineered Source Synchronous Interfaces

- DDR / DDR2 interfaces up to 200 MHz
- 7:1 LVDS interfaces support display applications
- XGMII

■ Density And Package Options

- 5k to 40k LUT4s, 86 to 540 I/Os
- csBGA, TQFP, PQFP, ftBGA and fpBGA packages
- Density migration supported

■ Flexible Device Configuration

- SPI (master and slave) Boot Flash Interface
- Dual Boot Image supported
- Soft Error Detect (SED) macro embedded

■ System Level Support

- IEEE 1149.1 and IEEE 1532 Compliant
- On-chip oscillator for initialization & general use
- Devices operate with 1.2V power supply

Table 1-1. LatticeXP2 Family Selection Guide

Device	XP2-5	XP2-8	XP2-17	XP2-30	XP2-40
LUTs (K)	5	8	17	29	40
Distributed RAM (KBits)	10	18	35	56	83
EBR SRAM (KBits)	166	221	276	387	885
EBR SRAM Blocks	9	12	15	21	48
sysDSP Blocks	3	4	5	7	8
18 x 18 Multipliers	12	16	20	28	32
V _{CC} Voltage	1.2	1.2	1.2	1.2	1.2
GPLL	2	2	4	4	4
Max Available I/O	172	201	358	472	540
Packages and I/O Combinations					
132-Ball csBGA (8 x 8 mm)	86	86			
144-Pin TQFP (20 x 20 mm)	100	100			
208-Pin PQFP (28 x 28 mm)	146	146	146		
256-Ball ftBGA (17 x 17 mm)	172	201	201	201	
484-Ball fpBGA (23 x 23 mm)			358	363	363
672-Ball fpBGA (27 x 27 mm)				472	540

Introduction

LatticeXP2 devices combine a Look-up Table (LUT) based FPGA fabric with non-volatile Flash cells in an architecture referred to as flexiFLASH.

The flexiFLASH approach provides benefits including instant-on, infinite reconfigurability, on chip storage with FlashBAK embedded block memory and Serial TAG memory and design security. The parts also support Live Update technology with TransFR, 128-bit AES Encryption and Dual-boot technologies.

The LatticeXP2 FPGA fabric was optimized for the new technology from the outset with high performance and low cost in mind. LatticeXP2 devices include LUT-based logic, distributed and embedded memory, Phase Locked Loops (PLLs), pre-engineered source synchronous I/O support and enhanced sysDSP blocks.

Lattice Diamond[®] design software allows large and complex designs to be efficiently implemented using the LatticeXP2 family of FPGA devices. Synthesis library support for LatticeXP2 is available for popular logic synthesis tools. The Diamond software uses the synthesis tool output along with the constraints from its floor planning tools to place and route the design in the LatticeXP2 device. The Diamond tool extracts the timing from the routing and back-annotates it into the design for timing verification.

Lattice provides many pre-designed Intellectual Property (IP) LatticeCORE[™] modules for the LatticeXP2 family. By using these IPs as standardized blocks, designers are free to concentrate on the unique aspects of their design, increasing their productivity.

Architecture Overview

Each LatticeXP2 device contains an array of logic blocks surrounded by Programmable I/O Cells (PIC). Interspersed between the rows of logic blocks are rows of sysMEM™ Embedded Block RAM (EBR) and a row of sys-DSP™ Digital Signal Processing blocks as shown in Figure 2-1.

On the left and right sides of the Programmable Functional Unit (PFU) array, there are Non-volatile Memory Blocks. In configuration mode the nonvolatile memory is programmed via the IEEE 1149.1 TAP port or the sysCONFIG™ peripheral port. On power up, the configuration data is transferred from the Non-volatile Memory Blocks to the configuration SRAM. With this technology, expensive external configuration memory is not required, and designs are secured from unauthorized read-back. This transfer of data from non-volatile memory to configuration SRAM via wide busses happens in microseconds, providing an “instant-on” capability that allows easy interfacing in many applications. LatticeXP2 devices can also transfer data from the sysMEM EBR blocks to the Non-volatile Memory Blocks at user request.

There are two kinds of logic blocks, the PFU and the PFU without RAM (PFF). The PFU contains the building blocks for logic, arithmetic, RAM and ROM functions. The PFF block contains building blocks for logic, arithmetic and ROM functions. Both PFU and PFF blocks are optimized for flexibility allowing complex designs to be implemented quickly and efficiently. Logic Blocks are arranged in a two-dimensional array. Only one type of block is used per row.

LatticeXP2 devices contain one or more rows of sysMEM EBR blocks. sysMEM EBRs are large dedicated 18Kbit memory blocks. Each sysMEM block can be configured in a variety of depths and widths of RAM or ROM. In addition, LatticeXP2 devices contain up to two rows of DSP Blocks. Each DSP block has multipliers and adder/accumulators, which are the building blocks for complex signal processing capabilities.

Each PIC block encompasses two PIOs (PIO pairs) with their respective sysIO buffers. The sysIO buffers of the LatticeXP2 devices are arranged into eight banks, allowing the implementation of a wide variety of I/O standards. PIO pairs on the left and right edges of the device can be configured as LVDS transmit/receive pairs. The PIC logic also includes pre-engineered support to aid in the implementation of high speed source synchronous standards such as 7:1 LVDS interfaces, found in many display applications, and memory interfaces including DDR and DDR2.

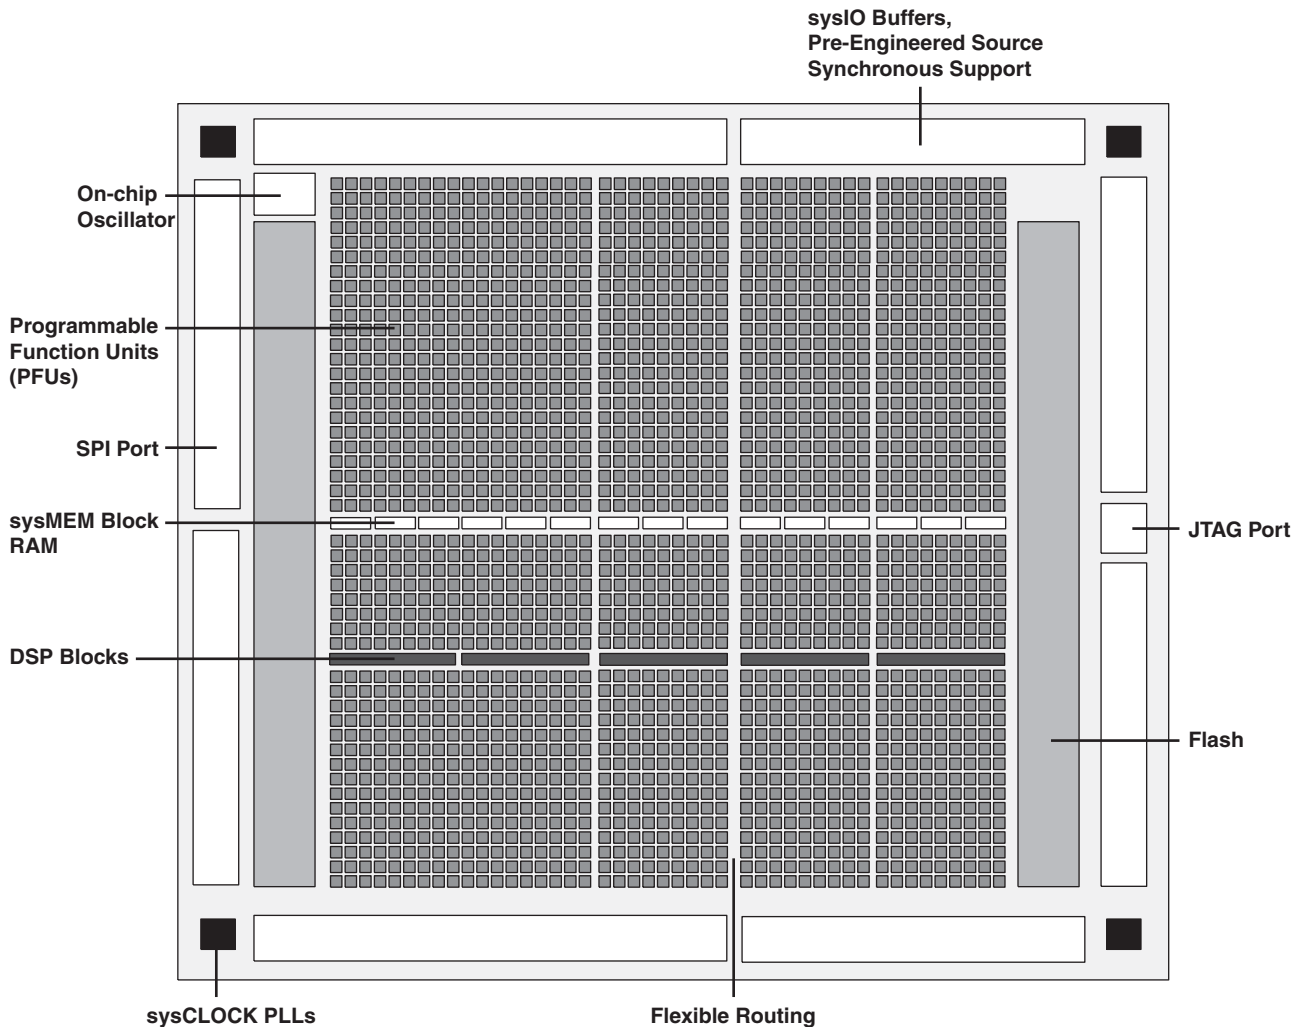
The LatticeXP2 registers in PFU and sysI/O can be configured to be SET or RESET. After power up and device is configured, the device enters into user mode with these registers SET/RESET according to the configuration setting, allowing device entering to a known state for predictable system function.

Other blocks provided include PLLs and configuration functions. The LatticeXP2 architecture provides up to four General Purpose PLLs (GPLL) per device. The GPLL blocks are located in the corners of the device.

The configuration block that supports features such as configuration bit-stream de-encryption, transparent updates and dual boot support is located between banks two and three. Every device in the LatticeXP2 family supports a sysCONFIG port, muxed with bank seven I/Os, which supports serial device configuration. A JTAG port is provided between banks two and three.

This family also provides an on-chip oscillator. LatticeXP2 devices use 1.2V as their core voltage.

Figure 2-1. Simplified Block Diagram, LatticeXP2-17 Device (Top Level)

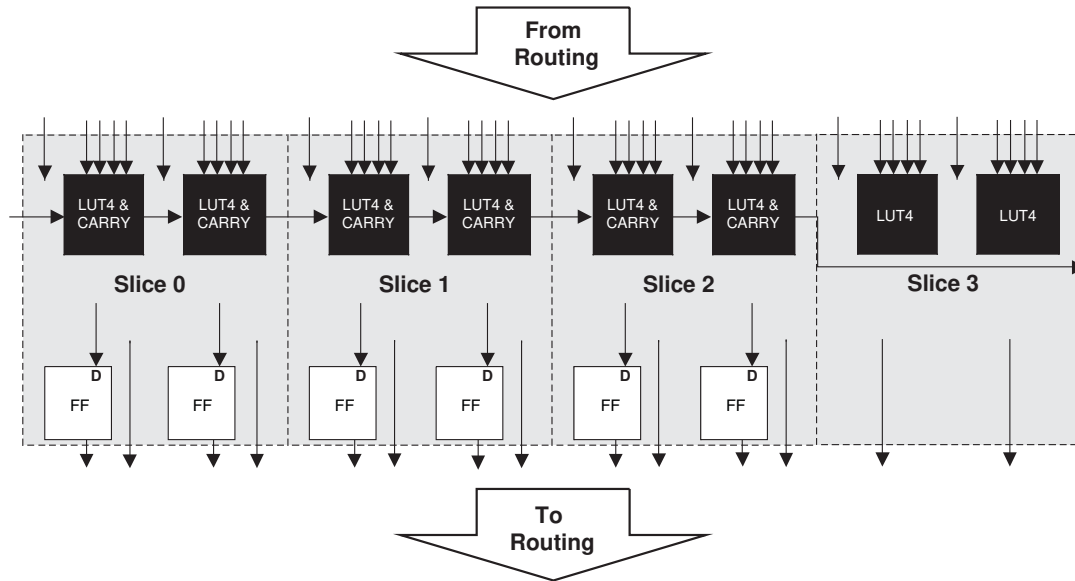


PFU Blocks

The core of the LatticeXP2 device is made up of logic blocks in two forms, PFUs and PFFs. PFUs can be programmed to perform logic, arithmetic, distributed RAM and distributed ROM functions. PFF blocks can be programmed to perform logic, arithmetic and ROM functions. Except where necessary, the remainder of this data sheet will use the term PFU to refer to both PFU and PFF blocks.

Each PFU block consists of four interconnected slices, numbered Slice 0 through Slice 3, as shown in Figure 2-2. All the interconnections to and from PFU blocks are from routing. There are 50 inputs and 23 outputs associated with each PFU block.

Figure 2-2. PFU Diagram



Slice

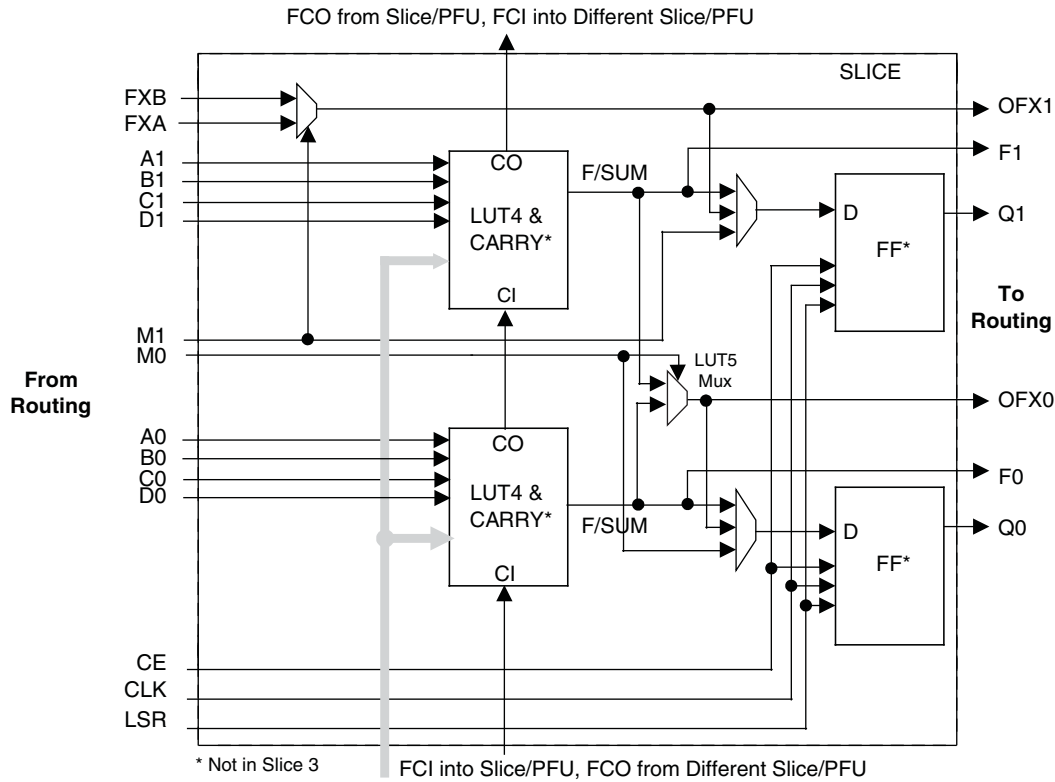
Slice 0 through Slice 2 contain two 4-input combinatorial Look-Up Tables (LUT4), which feed two registers. Slice 3 contains two LUT4s and no registers. For PFUs, Slice 0 and Slice 2 can also be configured as distributed memory, a capability not available in PFF blocks. Table 2-1 shows the capability of the slices in both PFF and PFU blocks along with the operation modes they enable. In addition, each PFU contains logic that allows the LUTs to be combined to perform functions such as LUT5, LUT6, LUT7 and LUT8. There is control logic to perform set/reset functions (programmable as synchronous/asynchronous), clock select, chip-select and wider RAM/ROM functions. Figure 2-3 shows an overview of the internal logic of the slice. The registers in the slice can be configured as positive/negative edge triggered or level sensitive clocks.

Table 2-1. Resources and Modes Available per Slice

Slice	PFU BBlock		PFF Block	
	Resources	Modes	Resources	Modes
Slice 0	2 LUT4s and 2 Registers	Logic, Ripple, RAM, ROM	2 LUT4s and 2 Registers	Logic, Ripple, ROM
Slice 1	2 LUT4s and 2 Registers	Logic, Ripple, ROM	2 LUT4s and 2 Registers	Logic, Ripple, ROM
Slice 2	2 LUT4s and 2 Registers	Logic, Ripple, RAM, ROM	2 LUT4s and 2 Registers	Logic, Ripple, ROM
Slice 3	2 LUT4s	Logic, ROM	2 LUT4s	Logic, ROM

Slice 0 through Slice 2 have 14 input signals: 13 signals from routing and one from the carry-chain (from the adjacent slice or PFU). There are seven outputs: six to routing and one to carry-chain (to the adjacent PFU). Slice 3 has 13 input signals from routing and four signals to routing. Table 2-2 lists the signals associated with Slice 0 to Slice 2.

Figure 2-3. Slice Diagram



For Slices 0 and 2, memory control signals are generated from Slice 1 as follows:
WCK is CLK
WRE is from LSR
DI[3:2] for Slice 2 and DI[1:0] for Slice 0 data
WAD [A:D] is a 4bit address from slice 1 LUT input

Table 2-2. Slice Signal Descriptions

Function	Type	Signal Names	Description
Input	Data signal	A0, B0, C0, D0	Inputs to LUT4
Input	Data signal	A1, B1, C1, D1	Inputs to LUT4
Input	Multi-purpose	M0	Multipurpose Input
Input	Multi-purpose	M1	Multipurpose Input
Input	Control signal	CE	Clock Enable
Input	Control signal	LSR	Local Set/Reset
Input	Control signal	CLK	System Clock
Input	Inter-PFU signal	FCI	Fast Carry-In ¹
Input	Inter-slice signal	FXA	Intermediate signal to generate LUT6 and LUT7
Input	Inter-slice signal	FXB	Intermediate signal to generate LUT6 and LUT7
Output	Data signals	F0, F1	LUT4 output register bypass signals
Output	Data signals	Q0, Q1	Register outputs
Output	Data signals	OFX0	Output of a LUT5 MUX
Output	Data signals	OFX1	Output of a LUT6, LUT7, LUT8 ² MUX depending on the slice
Output	Inter-PFU signal	FCO	Slice 2 of each PFU is the fast carry chain output ¹

1. See Figure 2-3 for connection details.

2. Requires two PFUs.

Modes of Operation

Each slice has up to four potential modes of operation: Logic, Ripple, RAM and ROM.

Logic Mode

In this mode, the LUTs in each slice are configured as LUT4s. A LUT4 has 16 possible input combinations. Four-input logic functions are generated by programming the LUT4. Since there are two LUT4s per slice, a LUT5 can be constructed within one slice. Larger LUTs such as LUT6, LUT7 and LUT8, can be constructed by concatenating two or more slices. Note that a LUT8 requires more than four slices.

Ripple Mode

Ripple mode allows efficient implementation of small arithmetic functions. In ripple mode, the following functions can be implemented by each slice:

- Addition 2-bit
- Subtraction 2-bit
- Add/Subtract 2-bit using dynamic control
- Up counter 2-bit
- Down counter 2-bit
- Up/Down counter with async clear
- Up/Down counter with preload (sync)
- Ripple mode multiplier building block
- Multiplier support
- Comparator functions of A and B inputs
 - A greater-than-or-equal-to B
 - A not-equal-to B
 - A less-than-or-equal-to B

Two carry signals, FCI and FCO, are generated per slice in this mode, allowing fast arithmetic functions to be constructed by concatenating slices.

RAM Mode

In this mode, a 16x4-bit distributed Single Port RAM (SPR) can be constructed using each LUT block in Slice 0 and Slice 2 as a 16x1-bit memory. Slice 1 is used to provide memory address and control signals. A 16x2-bit Pseudo Dual Port RAM (PDPR) memory is created by using one slice as the read-write port and the other companion slice as the read-only port.

The Lattice design tools support the creation of a variety of different size memories. Where appropriate, the software will construct these using distributed memory primitives that represent the capabilities of the PFU. Table 2-3 shows the number of slices required to implement different distributed RAM primitives. For more information on using RAM in LatticeXP2 devices, please see TN1137, [LatticeXP2 Memory Usage Guide](#).

Table 2-3. Number of Slices Required For Implementing Distributed RAM

	SPR 16X4	PDPR 16X4
Number of slices	3	3

Note: SPR = Single Port RAM, PDPR = Pseudo Dual Port RAM

ROM Mode

ROM mode uses the LUT logic; hence, Slices 0 through 3 can be used in the ROM mode. Preloading is accomplished through the programming interface during PFU configuration.

Routing

There are many resources provided in the LatticeXP2 devices to route signals individually or as busses with related control signals. The routing resources consist of switching circuitry, buffers and metal interconnect (routing) segments.

The inter-PFU connections are made with x1 (spans two PFU), x2 (spans three PFU) or x6 (spans seven PFU) connections. The x1 and x2 connections provide fast and efficient connections in horizontal and vertical directions. The x2 and x6 resources are buffered to allow both short and long connections routing between PFUs.

The LatticeXP2 family has an enhanced routing architecture to produce a compact design. The Diamond design tool takes the output of the synthesis tool and places and routes the design. Generally, the place and route tool is completely automatic, although an interactive routing editor is available to optimize the design.

sysCLOCK Phase Locked Loops (PLL)

The sysCLOCK PLLs provide the ability to synthesize clock frequencies. The LatticeXP2 family supports between two and four full featured General Purpose PLLs (GPLL). The architecture of the GPLL is shown in Figure 2-4.

CLKI, the PLL reference frequency, is provided either from the pin or from routing; it feeds into the Input Clock Divider block. CLKFB, the feedback signal, is generated from CLKOP (the primary clock output) or from a user clock pin/logic. CLKFB feeds into the Feedback Divider and is used to multiply the reference frequency.

Both the input path and feedback signals enter the Voltage Controlled Oscillator (VCO) block. The phase and frequency of the VCO are determined from the input path and feedback signals. A LOCK signal is generated by the VCO to indicate that the VCO is locked with the input clock signal.

The output of the VCO feeds into the CLKOP Divider, a post-scalar divider. The duty cycle of the CLKOP Divider output can be fine tuned using the Duty Trim block, which creates the CLKOP signal. By allowing the VCO to operate at higher frequencies than CLKOP, the frequency range of the GPLL is expanded. The output of the CLKOP Divider is passed through the CLKOK Divider, a secondary clock divider, to generate lower frequencies for the CLKOK output. For applications that require even lower frequencies, the CLKOP signal is passed through a divide-by-three divider to produce the CLKOK2 output. The CLKOK2 output is provided for applications that use source synchronous logic. The Phase/Duty Cycle/Duty Trim block is used to adjust the phase and duty cycle of the CLKOP Divider output to generate the CLKOS signal. The phase/duty cycle setting can be pre-programmed or dynamically adjusted.

The clock outputs from the GPLL; CLKOP, CLKOK, CLKOK2 and CLKOS, are fed to the clock distribution network.

For further information on the GPLL please see TN1126, [LatticeXP2 sysCLOCK PLL Design and Usage Guide](#).

Figure 2-4. General Purpose PLL (GPLL) Diagram

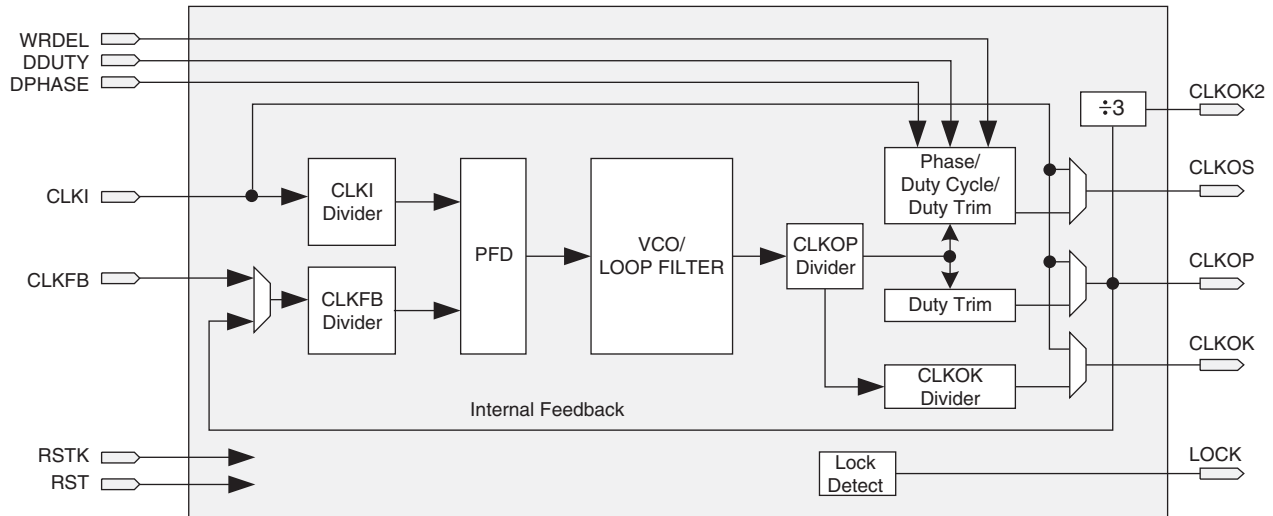


Table 2-4 provides a description of the signals in the GPLL blocks.

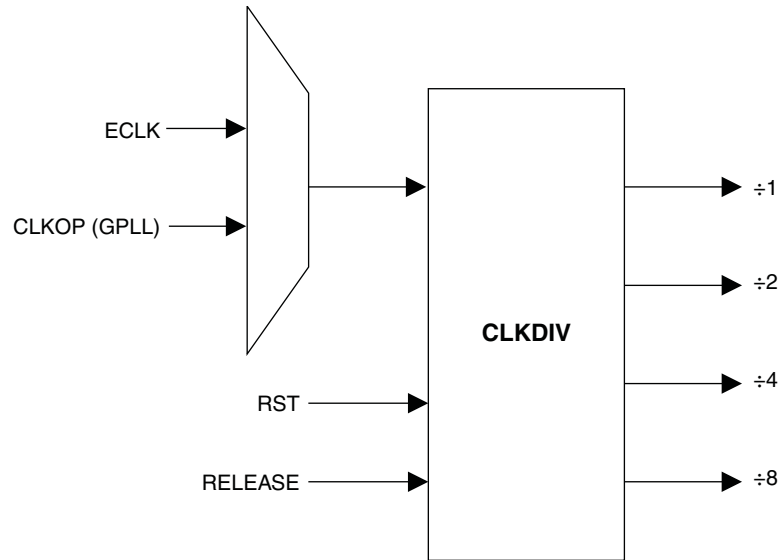
Table 2-4. GPLL Block Signal Descriptions

Signal	I/O	Description
CLKI	I	Clock input from external pin or routing
CLKFB	I	PLL feedback input from CLKOP (PLL internal), from clock net (CLKOP) or from a user clock (PIN or logic)
RST	I	"1" to reset PLL counters, VCO, charge pumps and M-dividers
RSTK	I	"1" to reset K-divider
DPHASE [3:0]	I	DPA Phase Adjust input
DDUTY [3:0]	I	DPA Duty Cycle Select input
WRDEL	I	DPA Fine Delay Adjust input
CLKOS	O	PLL output clock to clock tree (phase shifted/duty cycle changed)
CLKOP	O	PLL output clock to clock tree (no phase shift)
CLKOK	O	PLL output to clock tree through secondary clock divider
CLKOK2	O	PLL output to clock tree (CLKOP divided by 3)
LOCK	O	"1" indicates PLL LOCK to CLKI

Clock Dividers

LatticeXP2 devices have two clock dividers, one on the left side and one on the right side of the device. These are intended to generate a slower-speed system clock from a high-speed edge clock. The block operates in a $\div 2$, $\div 4$ or $\div 8$ mode and maintains a known phase relationship between the divided down clock and the high-speed clock based on the release of its reset signal. The clock dividers can be fed from the CLKOP output from the GPLLs or from the Edge Clocks (ECLK). The clock divider outputs serve as primary clock sources and feed into the clock distribution network. The Reset (RST) control signal resets the input and forces all outputs to low. The RELEASE signal releases outputs to the input clock. For further information on clock dividers, please see TN1126, [LatticeXP2 sysCLOCK PLL Design and Usage Guide](#). Figure 2-5 shows the clock divider connections.

Figure 2-5. Clock Divider Connections



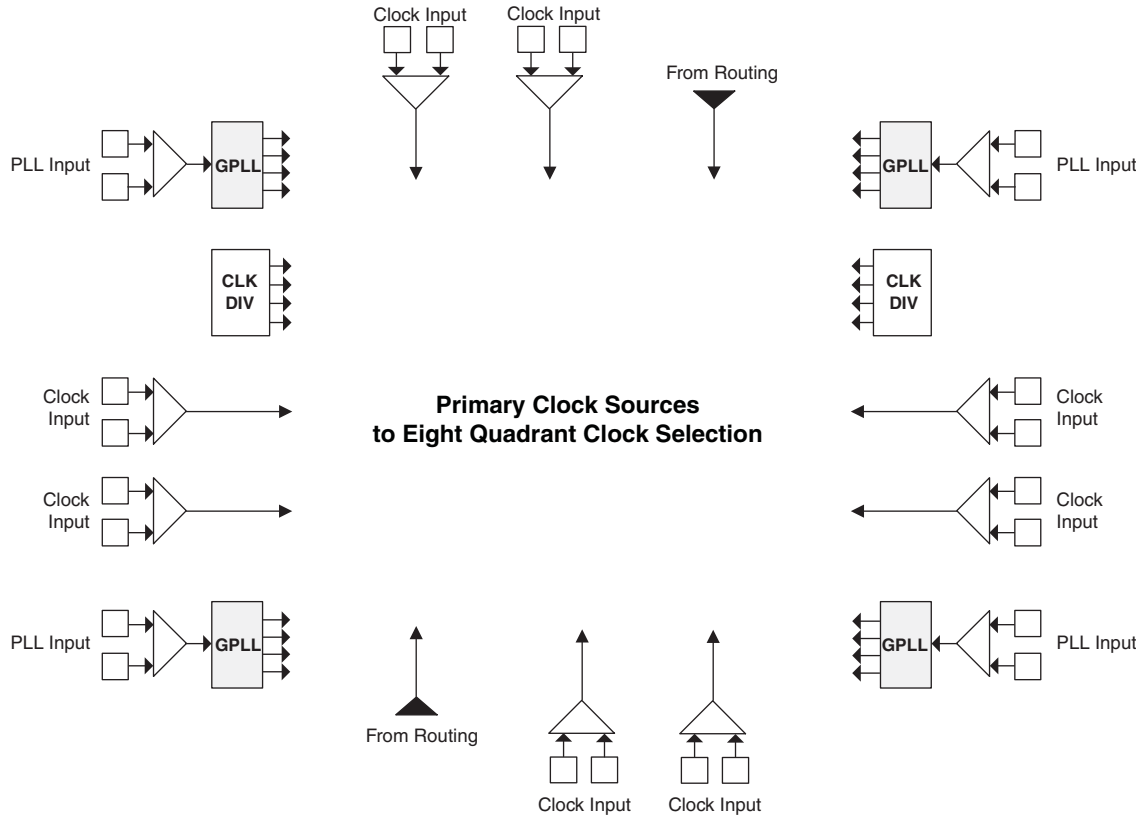
Clock Distribution Network

LatticeXP2 devices have eight quadrant-based primary clocks and between six and eight flexible region-based secondary clocks/control signals. Two high performance edge clocks are available on each edge of the device to support high speed interfaces. The clock inputs are selected from external I/Os, the sysCLOCK PLLs, or routing. Clock inputs are fed throughout the chip via the primary, secondary and edge clock networks.

Primary Clock Sources

LatticeXP2 devices derive primary clocks from four sources: PLL outputs, CLKDIV outputs, dedicated clock inputs and routing. LatticeXP2 devices have two to four sysCLOCK PLLs, located in the four corners of the device. There are eight dedicated clock inputs, two on each side of the device. Figure 2-6 shows the primary clock sources.

Figure 2-6. Primary Clock Sources for XP2-17

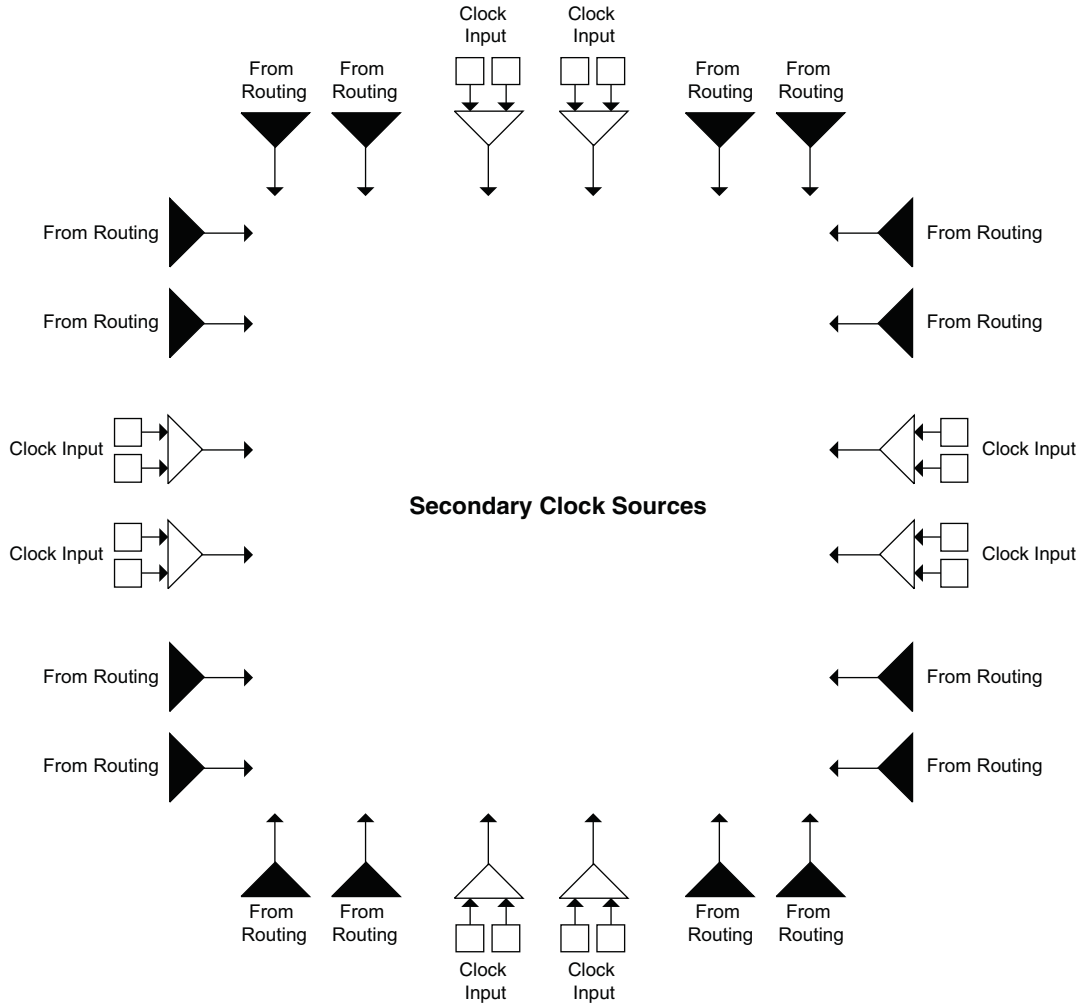


Note: This diagram shows sources for the XP2-17 device. Smaller LatticeXP2 devices have two GPLLs.

Secondary Clock/Control Sources

LatticeXP2 devices derive secondary clocks (SC0 through SC7) from eight dedicated clock input pads and the rest from routing. Figure 2-7 shows the secondary clock sources.

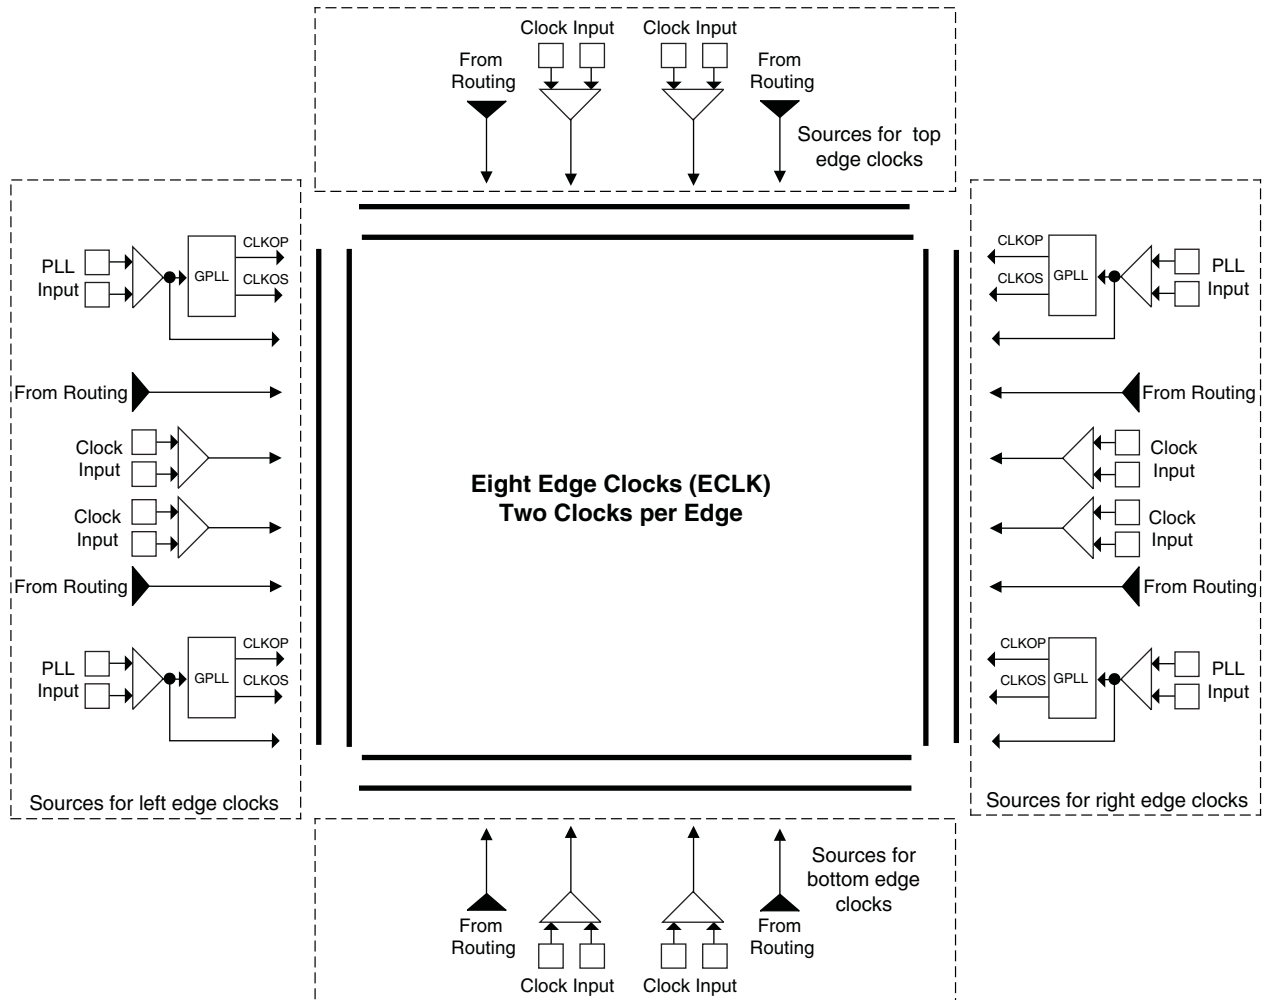
Figure 2-7. Secondary Clock Sources



Edge Clock Sources

Edge clock resources can be driven from a variety of sources at the same edge. Edge clock resources can be driven from adjacent edge clock PIOs, primary clock PIOs, PLLs and clock dividers as shown in Figure 2-8.

Figure 2-8. Edge Clock Sources

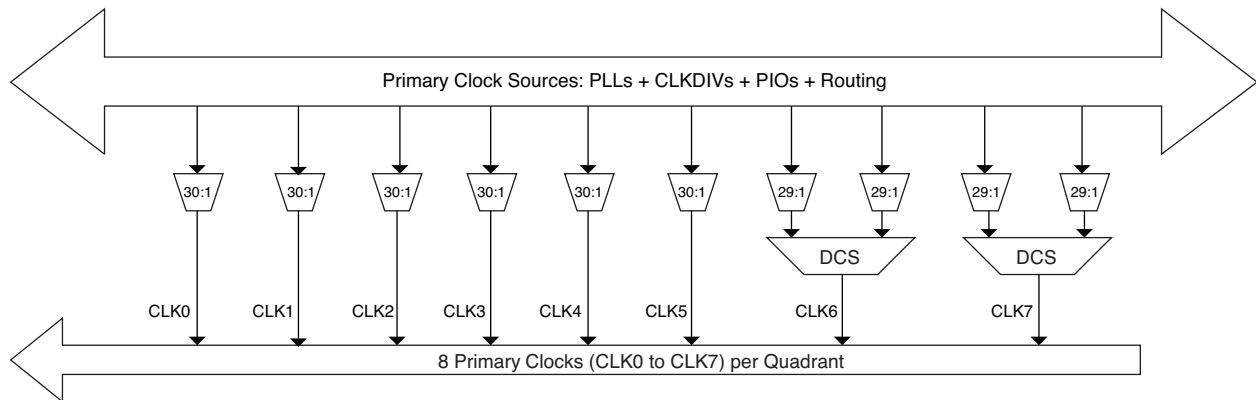


Note: This diagram shows sources for the XP2-17 device. Smaller LatticeXP2 devices have two GPLLs.

Primary Clock Routing

The clock routing structure in LatticeXP2 devices consists of a network of eight primary clock lines (CLK0 through CLK7) per quadrant. The primary clocks of each quadrant are generated from muxes located in the center of the device. All the clock sources are connected to these muxes. Figure 2-9 shows the clock routing for one quadrant. Each quadrant mux is identical. If desired, any clock can be routed globally.

Figure 2-9. Per Quadrant Primary Clock Selection




Dynamic Clock Select (DCS)

The DCS is a smart multiplexer function available in the primary clock routing. It switches between two independent input clock sources without any glitches or runt pulses. This is achieved irrespective of when the select signal is toggled. There are two DCS blocks per quadrant; in total, eight DCS blocks per device. The inputs to the DCS block come from the center muxes. The output of the DCS is connected to primary clocks CLK6 and CLK7 (see Figure 2-9).

Figure 2-10 shows the timing waveforms of the default DCS operating mode. The DCS block can be programmed to other modes. For more information on the DCS, please see TN1126, [LatticeXP2 sysCLOCK PLL Design and Usage Guide](#).

Figure 2-10. DCS Waveforms



Secondary Clock/Control Routing

Secondary clocks in the LatticeXP2 devices are region-based resources. The benefit of region-based resources is the relatively low injection delay and skew within the region, as compared to primary clocks. EBR rows, DSP rows and a special vertical routing channel bound the secondary clock regions. This special vertical routing channel aligns with either the left edge of the center DSP block in the DSP row or the center of the DSP row. Figure 2-11 shows this special vertical routing channel and the eight secondary clock regions for the LatticeXP2-40.

LatticeXP2-30 and smaller devices have six secondary clock regions. All devices in the LatticeXP2 family have four secondary clocks (SC0 to SC3) which are distributed to every region.

The secondary clock muxes are located in the center of the device. Figure 2-12 shows the mux structure of the secondary clock routing. Secondary clocks SC0 to SC3 are used for clock and control and SC4 to SC7 are used for high fan-out signals.

Figure 2-11. Secondary Clock Regions XP2-40

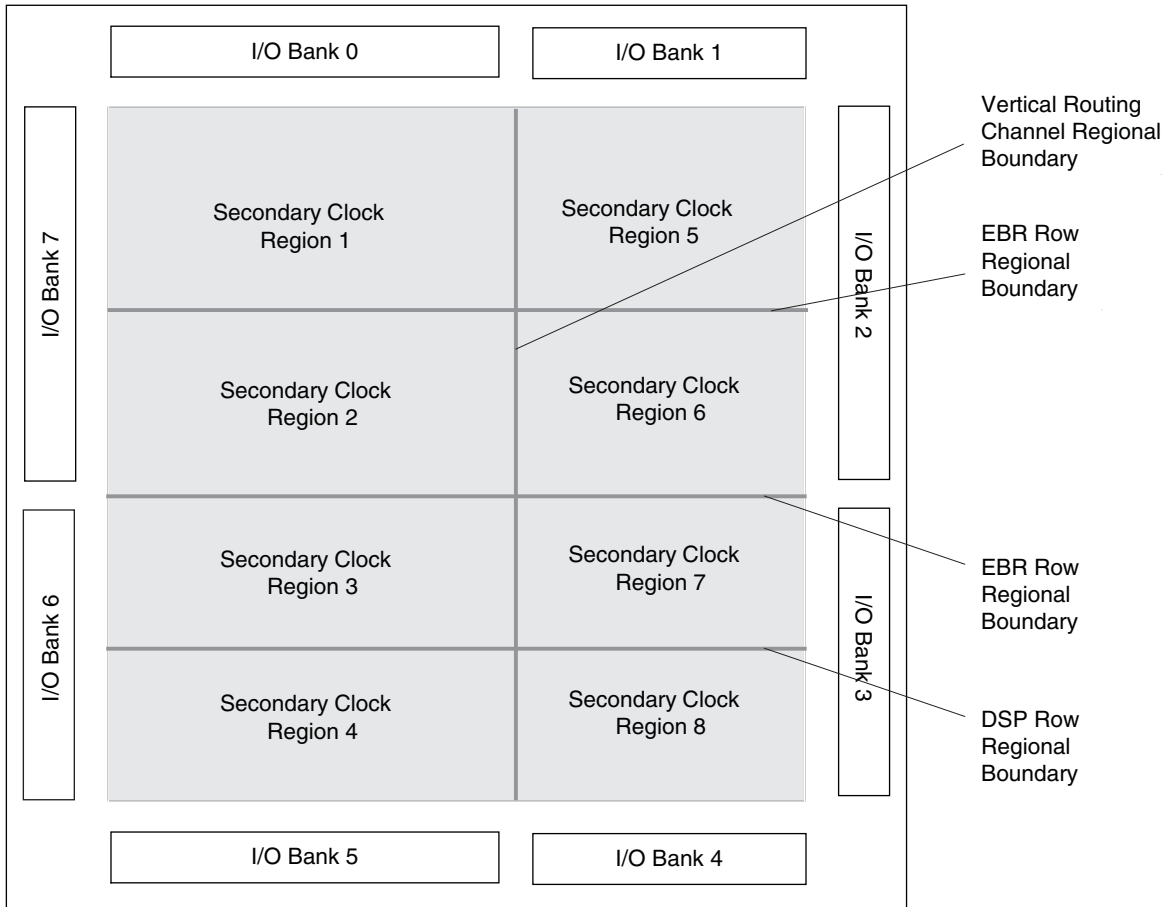
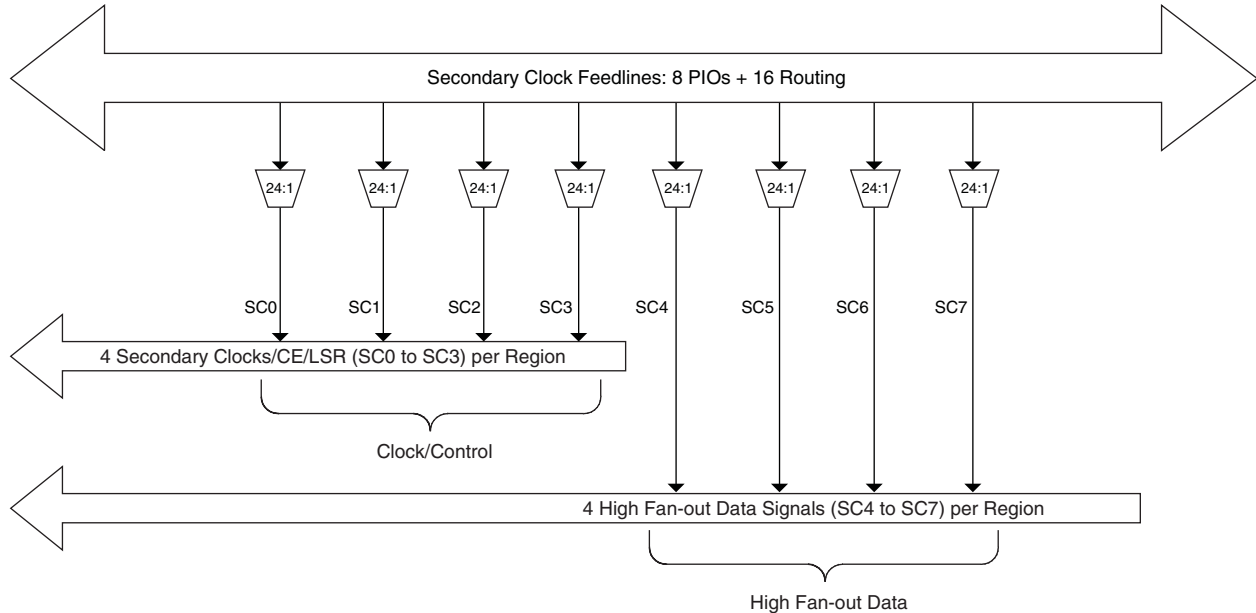


Figure 2-12. Secondary Clock Selection



Slice Clock Selection

Figure 2-13 shows the clock selections and Figure 2-14 shows the control selections for Slice0 through Slice2. All the primary clocks and the four secondary clocks are routed to this clock selection mux. Other signals, via routing, can be used as clock inputs to the slices. Slice controls are generated from the secondary clocks or other signals connected via routing.

If none of the signals are selected for both clock and control, then the default value of the mux output is 1. Slice 3 does not have any registers; therefore it does not have the clock or control muxes.

Figure 2-13. Slice0 through Slice2 Clock Selection


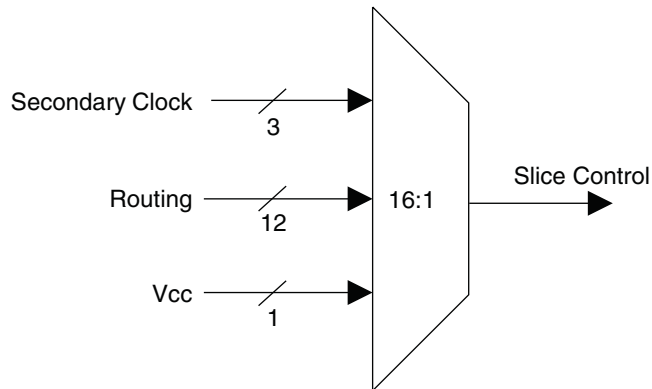


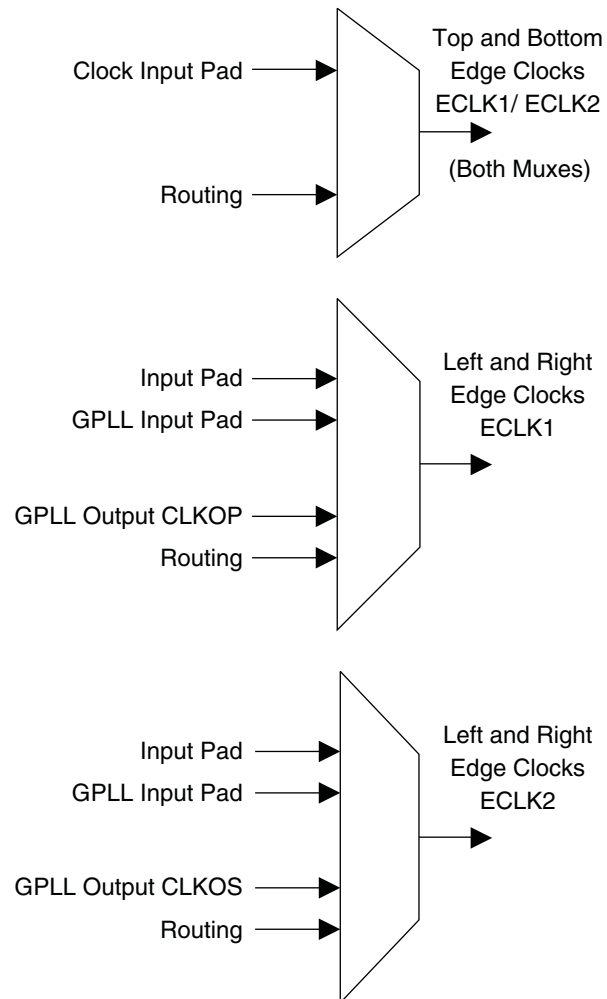
Figure 2-14. Slice0 through Slice2 Control Selection



Edge Clock Routing

LatticeXP2 devices have eight high-speed edge clocks that are intended for use with the PIOs in the implementation of high-speed interfaces. Each device has two edge clocks per edge. Figure 2-15 shows the selection muxes for these clocks.

Figure 2-15. Edge Clock Mux Connections



sysMEM Memory

LatticeXP2 devices contains a number of sysMEM Embedded Block RAM (EBR). The EBR consists of 18 Kbit RAM with dedicated input and output registers.

sysMEM Memory Block

The sysMEM block can implement single port, dual port or pseudo dual port memories. Each block can be used in a variety of depths and widths as shown in Table 2-5. FIFOs can be implemented in sysMEM EBR blocks by using support logic with PFUs. The EBR block supports an optional parity bit for each data byte to facilitate parity checking. EBR blocks provide byte-enable support for configurations with 18-bit and 36-bit data widths.

Table 2-5. sysMEM Block Configurations

Memory Mode	Configurations
Single Port	16,384 x 1
	8,192 x 2
	4,096 x 4
	2,048 x 9
	1,024 x 18
True Dual Port	512 x 36
	16,384 x 1
	8,192 x 2
	4,096 x 4
	2,048 x 9
Pseudo Dual Port	1,024 x 18
	16,384 x 1
	8,192 x 2
	4,096 x 4
	2,048 x 9
	1,024 x 18
	512 x 36

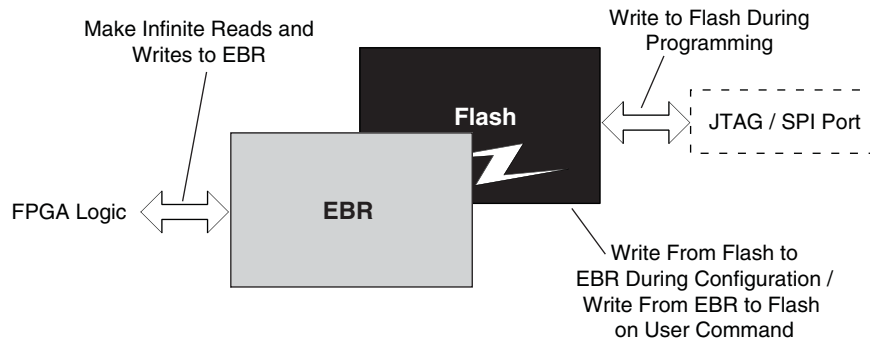
Bus Size Matching

All of the multi-port memory modes support different widths on each of the ports. The RAM bits are mapped LSB word 0 to MSB word 0, LSB word 1 to MSB word 1, and so on. Although the word size and number of words for each port varies, this mapping scheme applies to each port.

FlashBAK EBR Content Storage

All the EBR memory in the LatticeXP2 is shadowed by Flash memory. Optionally, initialization values for the memory blocks can be defined using the Lattice Diamond design tools. The initialization values are loaded into the Flash memory during device programming and into the SRAM at power up or whenever the device is reconfigured. This feature is ideal for the storage of a variety of information such as look-up tables and microprocessor code. It is also possible to write the current contents of the EBR memory back to Flash memory. This capability is useful for the storage of data such as error codes and calibration information. For additional information on the FlashBAK capability see TN1137, [LatticeXP2 Memory Usage Guide](#).

Figure 2-16. FlashBAK Technology



Memory Cascading

Larger and deeper blocks of RAMs can be created using EBR sysMEM Blocks. Typically, the Lattice design tools cascade memory transparently, based on specific design inputs.

Single, Dual and Pseudo-Dual Port Modes

In all the sysMEM RAM modes the input data and address for the ports are registered at the input of the memory array. The output data of the memory is optionally registered at the output.

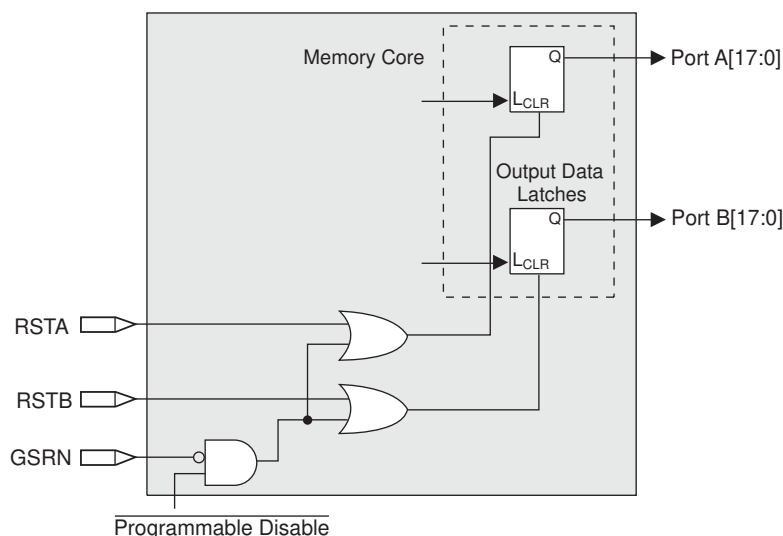
EBR memory supports two forms of write behavior for single port or dual port operation:

1. Normal – Data on the output appears only during a read cycle. During a write cycle, the data (at the current address) does not appear on the output. This mode is supported for all data widths.
2. Write Through – A copy of the input data appears at the output of the same port during a write cycle. This mode is supported for all data widths.

Memory Core Reset

The memory array in the EBR utilizes latches at the A and B output ports. These latches can be reset asynchronously or synchronously. RSTA and RSTB are local signals, which reset the output latches associated with Port A and Port B respectively. GSRN, the global reset signal, resets both ports. The output data latches and associated resets for both ports are as shown in Figure 2-17.

Figure 2-17. Memory Core Reset




For further information on the sysMEM EBR block, please see TN1137, [LatticeXP2 Memory Usage Guide](#).

EBR Asynchronous Reset

EBR asynchronous reset or GSR (if used) can only be applied if all clock enables are low for a clock cycle before the reset is applied and released a clock cycle after the low-to-high transition of the reset signal, as shown in Figure 2-18. The GSR input to the EBR is always asynchronous.

Figure 2-18. EBR Asynchronous Reset (Including GSR) Timing Diagram



If all clock enables remain enabled, the EBR asynchronous reset or GSR may only be applied and released after the EBR read and write clock inputs are in a steady state condition for a minimum of $1/f_{MAX}$ (EBR clock). The reset release must adhere to the EBR synchronous reset setup time before the next active read or write clock edge.

If an EBR is pre-loaded during configuration, the GSR input must be disabled or the release of the GSR during device Wake Up must occur before the release of the device I/Os becoming active.

These instructions apply to all EBR RAM and ROM implementations.

Note that there are no reset restrictions if the EBR synchronous reset is used and the EBR GSR input is disabled.

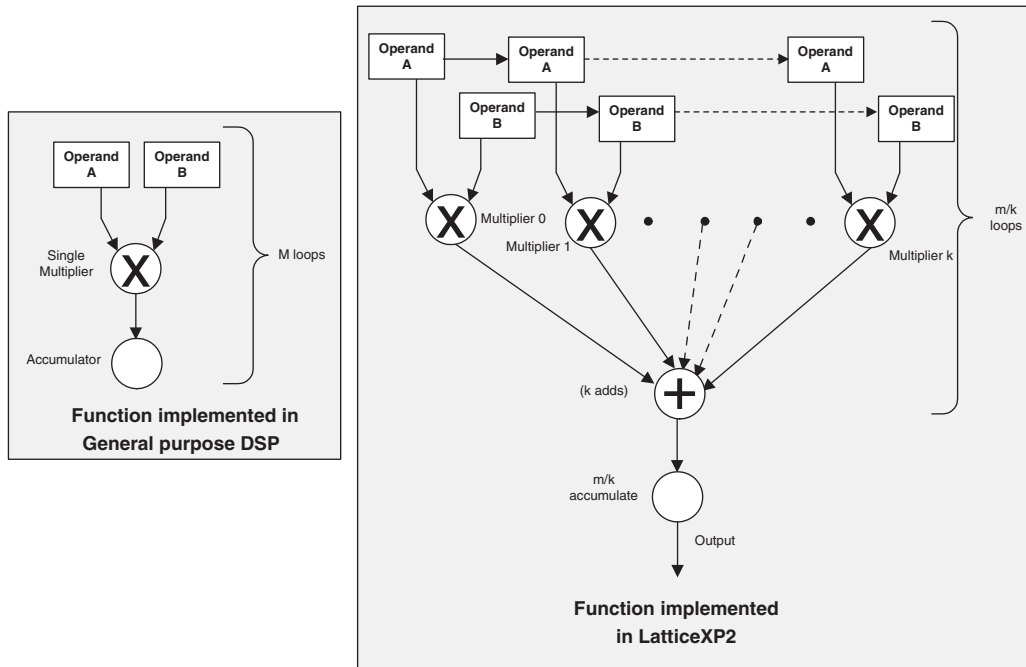
sysDSP™ Block

The LatticeXP2 family provides a sysDSP block making it ideally suited for low cost, high performance Digital Signal Processing (DSP) applications. Typical functions used in these applications include Bit Correlators, Fast Fourier Transform (FFT) functions, Finite Impulse Response (FIR) Filter, Reed-Solomon Encoder/Decoder, Turbo Encoder/Decoder and Convolutional Encoder/Decoder. These complex signal processing functions use similar building blocks such as multiply-adders and multiply-accumulators.

sysDSP Block Approach Compare to General DSP

Conventional general-purpose DSP chips typically contain one to four (Multiply and Accumulate) MAC units with fixed data-width multipliers; this leads to limited parallelism and limited throughput. Their throughput is increased by higher clock speeds. The LatticeXP2 family, on the other hand, has many DSP blocks that support different data-widths. This allows the designer to use highly parallel implementations of DSP functions. The designer can optimize the DSP performance vs. area by choosing appropriate levels of parallelism. Figure 2-19 compares the fully serial and the mixed parallel and serial implementations.

Figure 2-19. Comparison of General DSP and LatticeXP2 Approaches



sysDSP Block Capabilities

The sysDSP block in the LatticeXP2 family supports four functional elements in three 9, 18 and 36 data path widths. The user selects a function element for a DSP block and then selects the width and type (signed/unsigned) of its operands. The operands in the LatticeXP2 family sysDSP Blocks can be either signed or unsigned but not mixed within a function element. Similarly, the operand widths cannot be mixed within a block. DSP elements can be concatenated.

The resources in each sysDSP block can be configured to support the following four elements:

- MULT (Multiply)
- MAC (Multiply, Accumulate)
- MULTADDSUB (Multiply, Addition/Subtraction)
- MULTADDSUBSUM (Multiply, Addition/Subtraction, Accumulate)

The number of elements available in each block depends on the width selected from the three available options: x9, x18, and x36. A number of these elements are concatenated for highly parallel implementations of DSP functions. Table 2-6 shows the capabilities of the block.

Table 2-6. Maximum Number of Elements in a Block

Width of Multiply	x9	x18	x36
MULT	8	4	1
MAC	2	2	—
MULTADDSUB	4	2	—
MULTADDSUBSUM	2	1	—

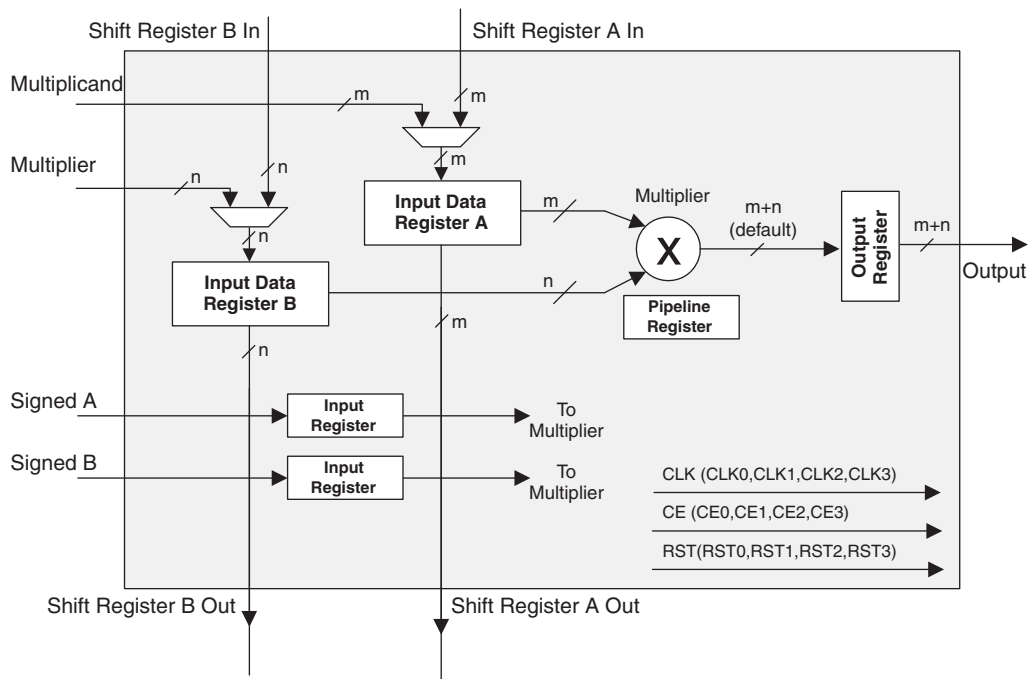
Some options are available in four elements. The input register in all the elements can be directly loaded or can be loaded as shift register from previous operand registers. By selecting 'dynamic operation' the following operations are possible:

- In the 'Signed/Unsigned' options the operands can be switched between signed and unsigned on every cycle.
- In the 'Add/Sub' option the Accumulator can be switched between addition and subtraction on every cycle.
- The loading of operands can switch between parallel and serial operations.

MULT sysDSP Element

This multiplier element implements a multiply with no addition or accumulator nodes. The two operands, A and B, are multiplied and the result is available at the output. The user can enable the input/output and pipeline registers. Figure 2-20 shows the MULT sysDSP element.

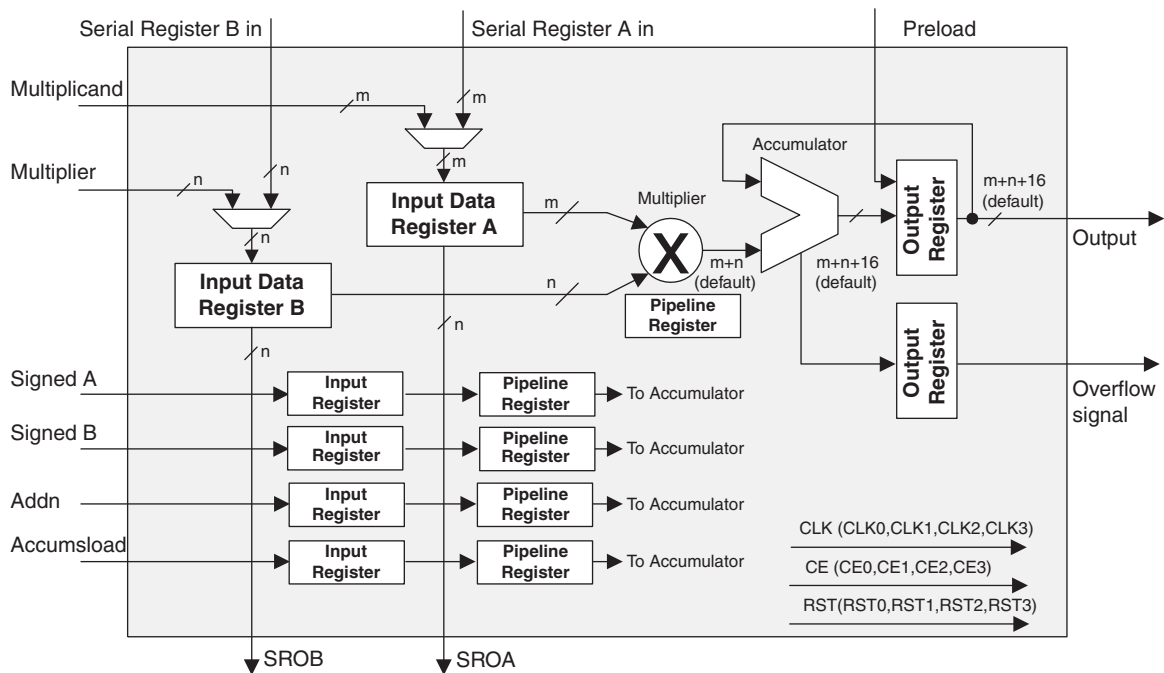
Figure 2-20. MULT sysDSP Element



MAC sysDSP Element

In this case, the two operands, A and B, are multiplied and the result is added with the previous accumulated value. This accumulated value is available at the output. The user can enable the input and pipeline registers but the output register is always enabled. The output register is used to store the accumulated value. The Accumulators in the DSP blocks in LatticeXP2 family can be initialized dynamically. A registered overflow signal is also available. The overflow conditions are provided later in this document. Figure 2-21 shows the MAC sysDSP element.

Figure 2-21. MAC sysDSP



MULTADDSUB sysDSP Element

In this case, the operands A0 and B0 are multiplied and the result is added/subtracted with the result of the multiplier operation of operands A1 and B1. The user can enable the input, output and pipeline registers. Figure 2-22 shows the MULTADDSUB sysDSP element.

Figure 2-22. MULTADDSUB

