



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



Features...

- High-performance programmable logic device (PLD) family (see [Table 1](#))
 - Integrated high-speed transceivers with support for clock data recovery (CDR) at up to 1.25 gigabits per second (Gbps)
 - Look-up table (LUT)-based architecture optimized for high speed
 - Advanced interconnect structure for fast routing of critical paths
 - Enhanced I/O structure for versatile standards and interface support
 - Up to 14,400 logic elements (LEs)
- System-level features
 - Up to four general-purpose phase-locked loops (PLLs) with programmable multiplication and delay shifting
 - Up to 12 PLL output ports
 - Dedicated multiplier circuitry for high-speed implementation of signed or unsigned multiplication up to 16×16
 - Embedded system blocks (ESBs) used to implement memory functions including quad-port RAM, true dual-port RAM, first-in first-out (FIFO) buffers, and content-addressable memory (CAM)
 - Each ESB contains 4,096 bits and can be split and used as two 2,048-bit unidirectional dual-port RAM blocks

Table 1. Mercury Device Features

Feature	EP1M120	EP1M350
Typical gates	120,000	350,000
HSDI channels	8	18
LEs	4,800	14,400
ESBs (1)	12	28
Maximum RAM bits	49,152	114,688
Maximum user I/O pins	303	486

Note to Table 1:

(1) Each ESB can be used for two dual- or single-port RAM blocks.

...and More Features

- Advanced high-speed I/O features
 - Robust I/O standard support, including LVTTTL, PCI up to 66 MHz, 3.3-V AGP in 1× and 2× modes, 3.3-V SSTL-3 and 2.5-V SSTL-2, GTL+, HSTL, CTT, LVDS, LVPECL, and 3.3-V PCML
 - High-speed differential interface (HSDI) with dedicated circuitry for CDR at up to 1.25 Gbps for LVDS, LVPECL, and 3.3-V PCML
 - Support for source-synchronous True-LVDS™ circuitry up to 840 megabits per second (Mbps) for LVDS, LVPECL, and 3.3-V PCML
 - Up to 18 input and 18 output dedicated differential channels of high-speed LVDS, LVPECL, or 3.3-V PCML
 - Built-in 100-Ω termination resistor on HSDI data and clock differential pairs
 - Flexible-LVDS™ circuitry provides 624-Mbps support on up to 100 channels with the EP1M350 device
 - Versatile three-register I/O element (IOE) supporting double data rate I/O (DDRIO), double data-rate (DDR) SDRAM, zero bus turnaround (ZBT) SRAM, and quad data rate (QDR) SRAM
- Designed for low-power operation
 - 1.8-V internal supply voltage (V_{CCINT})
 - MultiVolt™ I/O interface voltage levels (V_{CCIO}) compatible with 1.5-V, 1.8-V, 2.5-V, and 3.3-V devices
 - 5.0-V tolerant with external resistor
- Advanced interconnect structure
 - Multi-level FastTrack® Interconnect structure providing fast, predictable interconnect delays
 - Optimized high-speed Priority FastTrack Interconnect for routing critical paths in a design
 - Dedicated carry chain that implements arithmetic functions such as fast adders, counters, and comparators (automatically used by software tools and megafunctions)
 - FastLUT™ connection allowing high speed direct connection between LEs in the same logic array block (LAB)
 - Leap lines allowing a single LAB to directly drive LEs in adjacent rows
 - The RapidLAB interconnect providing a high-speed connection to a 10-LAB-wide region
 - Dedicated clock and control signal resources, including four dedicated clocks, six dedicated fast global signals, and additional row-global signals

Tables 2 and 3 show the Mercury™ FineLine BGA™ device package sizes, options, and I/O pin counts.

Table 2. Mercury Package Sizes

Feature	484-Pin FineLine BGA	780-Pin FineLine BGA
Pitch (mm)	1.00	1.00
Area (mm ²)	529	841
Length × width (mm × mm)	23 × 23	29 × 29

Table 3. Mercury Package Options & I/O Count

Device	484-Pin FineLine BGA	780-Pin FineLine BGA
EP1M120	303	
EP1M350		486

General Description

Mercury devices integrate high-speed differential transceivers and support for CDR with a speed-optimized PLD architecture. These transceivers are implemented through the dedicated serializer, deserializer, and clock recovery circuitry in the HSDI and incorporate support for the LVDS, LVPECL, and 3.3-V PCML I/O standards. This circuitry, together with enhanced I/O elements (IOEs) and support for numerous I/O standards, allows Mercury devices to meet high-speed interface requirements.

Mercury devices are the first PLDs optimized for core performance. These LUT-based, enhanced memory devices use a network of fast routing resources to achieve optimal performance. These resources are ideal for data-path, register-intensive, mathematical, digital signal processing (DSP), or communications designs.

Mercury devices include other features for performance such as quad-port RAM, CAM, general purpose PLLs, and dedicated circuitry for implementing multiplier circuits. Table 4 shows Mercury performance.

Application	Resources Used		Performance			
	LEs	ESBs	-5 Speed Grade	-6 Speed Grade	-7 Speed Grade	Units
16-bit loadable counter (1)	16	0	400	400	400	MHz
32-bit loadable counter (1)	32	0	400	400	400	MHz
32-bit accumulator (1)	32	0	400	400	400	MHz
32-to-1 multiplexer	27	0	1.864	2.466	2.723	ns
32 × 64 asynchronous FIFO	103	2	290	258	242	MHz
8-bit, 37-tap FIR filter	251	1	290	240	205	MSPS

Note to Table 4:

- (1) The clock tree supports up to 400 MHz. Although the registered performance for these designs exceed 400 MHz, they are limited by the clock tree limit.

Configuration

The logic, circuitry, and interconnects in the Mercury architecture are configured with CMOS SRAM elements. Mercury devices are reconfigurable and are 100% tested prior to shipment. As a result, test vectors do not have to be generated for fault coverage purposes. Instead, the designer can focus on simulation and design verification. In addition, the designer does not need to manage inventories of different ASIC designs; Mercury devices can be configured on the board for the specific functionality required.

Mercury devices are configured at system power-up with data stored in an Altera® serial configuration device or provided by a system controller. Altera offers in-system programmability (ISP)-capable configuration devices, which configure Mercury devices via a serial data stream. Mercury devices can be configured in under 70 ms. Moreover, Mercury devices contain an optimized interface that permits microprocessors to configure Mercury devices serially or in parallel, synchronously or asynchronously. This interface also enables microprocessors to treat Mercury devices as memory and to configure the device by writing to a virtual memory location, simplifying reconfiguration.

After a Mercury device has been configured, it can be reconfigured in-circuit by resetting the device and loading new data. Real-time changes can be made during system operation, enabling innovative reconfigurable computing applications.

Software

Mercury devices are supported by the Altera Quartus™ II development system, a single, integrated package that offers HDL and schematic design entry, compilation and logic synthesis, full simulation and worst-case timing analysis, SignalTap™ logic analysis, and device configuration. The Quartus II software also ships with Altera-specific HDL synthesis tools from Exemplar Logic and Synopsys, and Altera-specific Register Transfer Level (RTL) and timing simulation tools from Model Technology. The Quartus II software supports PCs running Windows 98, Windows NT 4.0, and Windows 2000; UNIX workstations running Solaris 2.6, 7, or 8, or HP-UX 10.2 or 11.0; and PCs running Red Hat Linux 7.1.

The Quartus II software provides NativeLink™ interfaces to other industry-standard PC- and UNIX-workstation-based EDA tools. For example, designers can invoke the Quartus II software from within the Mentor Graphics LeonardoSpectrum software, Synplify's Synplify software, and the Synopsys FPGA *Express* software. The Quartus II software also contains built-in optimized synthesis libraries; synthesis tools can use these libraries to optimize designs for Mercury devices. For example, the Synopsys Design Compiler library, supplied with the Quartus II development system, includes DesignWare functions optimized for the Mercury architecture.

For more information on the Quartus II development system, see the *Quartus II Programmable Logic Development System & Software Data Sheet*.

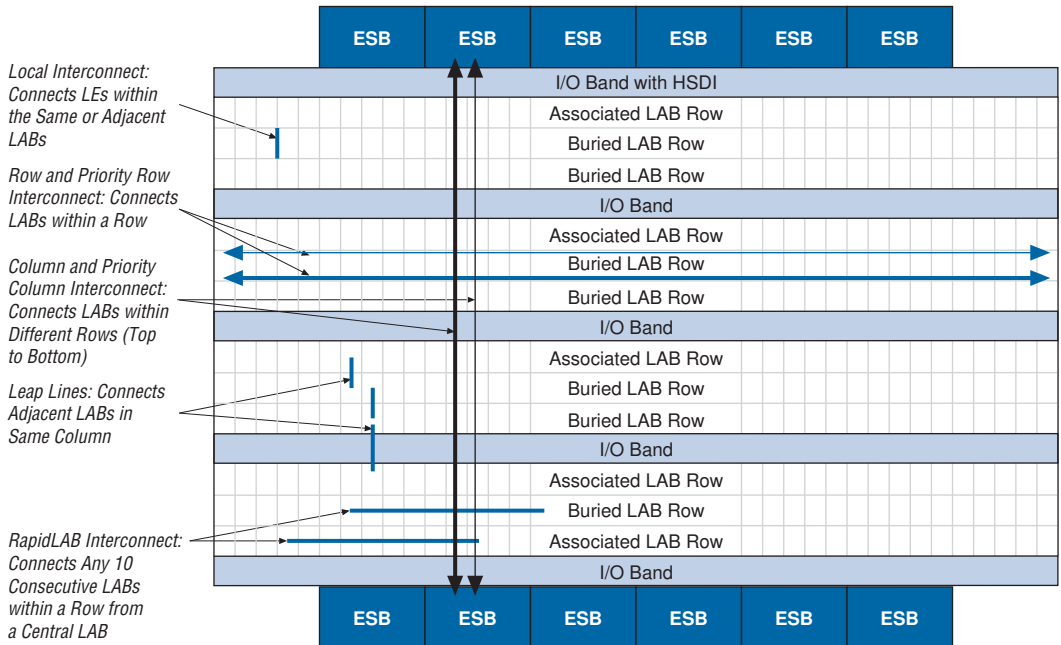
Functional Description

The Mercury architecture contains a row-based logic array to implement general logic and a row-based embedded system array to implement memory and specialized logic functions. Signal interconnections within Mercury devices are provided by a series of row and column interconnects with varying lengths and speeds. The priority FastTrack Interconnect structure is faster than other interconnects; the Quartus II Compiler places design-critical paths on these faster lines to improve design performance.

Mercury device I/O pins are evenly distributed across the entire device area; other Altera device families have I/O pins placed on the device periphery. Mercury device I/O pin placement allows for higher I/O count at a given die size; pad size is no longer a limiting issue. Each I/O pin is fed by an IOE. IOEs are grouped in IOE row bands from the top to the bottom of the device. IOE row bands are separated by several LAB rows. LABs from the associated LAB row closest to the I/O row band drive IOEs through the local interconnect. This feature allows fast clock-to-output times when a pin is driven by any of the 10 LEs in the adjacent associated LAB. Each IOE contains a bidirectional buffer along with an input register, output register, output enable (OE) register, and input latch for DDR. When used with a global clock, these dedicated registers provide exceptional bidirectional I/O performance.

IOEs provide a variety of features, such as 3.3-V, 64-bit, 66-MHz PCI compliance; 3.3-V, 64-bit, 133-MHz PCI-X compliance; Joint Test Action Group (JTAG) boundary-scan test (BST) support; output drive strength control; slew-rate control; tri-state buffers; bus-hold circuitry; programmable pull-up resistors; programmable input and output delays; and open-drain outputs. Mercury devices offer enhanced I/O support, including support for 1.8-V I/O, 2.5-V I/O, LVCMOS, LVTTTL, HSTL, LVPECL, 3.3-V PCML, 3.3-V PCI, PCI-X, LVDS, GTL+, SSTL-2, SSTL-3, CTT, and 3.3-V AGP I/O standards. CDR (up to 1.25 Gbps) and source-synchronous (up to 840 Mbps) transfers are supported with HSDI circuitry for LVDS, LVPECL, and 3.3-V PCML I/O standards.

The ESB can implement a variety of memory functions, including CAM, quad-port RAM, true dual-port RAM, dual- and single-port RAM, ROM, and FIFO functions. ESBs are grouped into two rows: one at the top and one at the bottom of the device. Embedding the memory directly into the die improves performance and reduces die area compared to distributed-RAM implementations. Moreover, the abundance of cascadable ESBs, in conjunction with the ability for one ESB to implement two separate memory blocks, ensures that the Mercury device can implement multiple wide memory blocks for high-density designs. The ESB's high speed ensures the implementation of small memory blocks without any speed penalty. The abundance of ESBs ensures that designers can create as many different-sized memory blocks as the system requires. [Figure 1](#) shows an overview of the Mercury device.

Figure 1. Mercury Architecture Block Diagram *Note (1)***Note to Figure 1:**

- (1) **Figure 1** shows an EP1M120 device. Mercury devices have a varying number of rows, columns, and ESBs, as shown in **Table 5**.

Table 5 lists the resources available in Mercury devices.

Device	LAB Rows	LAB Columns	I/O Row Bands	ESBs
EP1M120	12	40	5	12
EP1M350	18	80	4	28

Mercury devices provide four dedicated clock input pins and six dedicated fast I/O pins that globally drive register control inputs, including clocks. These signals ensure efficient distribution of high-speed, low-skew control signals. The control signals use dedicated routing channels to provide short delays and low skew. The dedicated fast signals can also be driven by internal logic, providing an ideal solution for a clock divider or internally generated asynchronous control signal with high fan-out. The dedicated clock and fast I/O pins on Mercury devices can also feed logic. Dedicated clocks can also be used with the Mercury general purpose PLLs for clock management.

Each I/O row band also provides two additional I/O pins that can drive two row-global signals. Row-global signals can drive register control inputs for the LAB row associated with that particular I/O row band.

High-Speed Differential Interface

The top I/O or HSDI band in Mercury devices contains dedicated circuitry for supporting differential standards at speeds up to 1.25 Gbps. Mercury devices have dedicated differential buffers and circuitry to support LVDS, LVPECL, and 3.3-V PCML I/O standards. Two dedicated high-speed PLLs (separate from the general purpose PLLs) multiply reference clocks and drive high-speed differential serializer/deserializer channels. In addition, clock recovery units (CRUs) at each receiver channel enable CDR. EP1M120 devices support eight input channels, eight output channels, and two dedicated clock inputs for feeding the receiver and/or transmitter PLLs. EP1M350 devices support 18 input channels, 18 output channels, and two dedicated clock inputs.

Mercury devices have optional built-in 100- Ω termination resistors on HSDI differential receiver data pins and the HSDI_CLK1 and HSDI_CLK2 pins.

Designers can use the HSDI circuitry for the following applications:

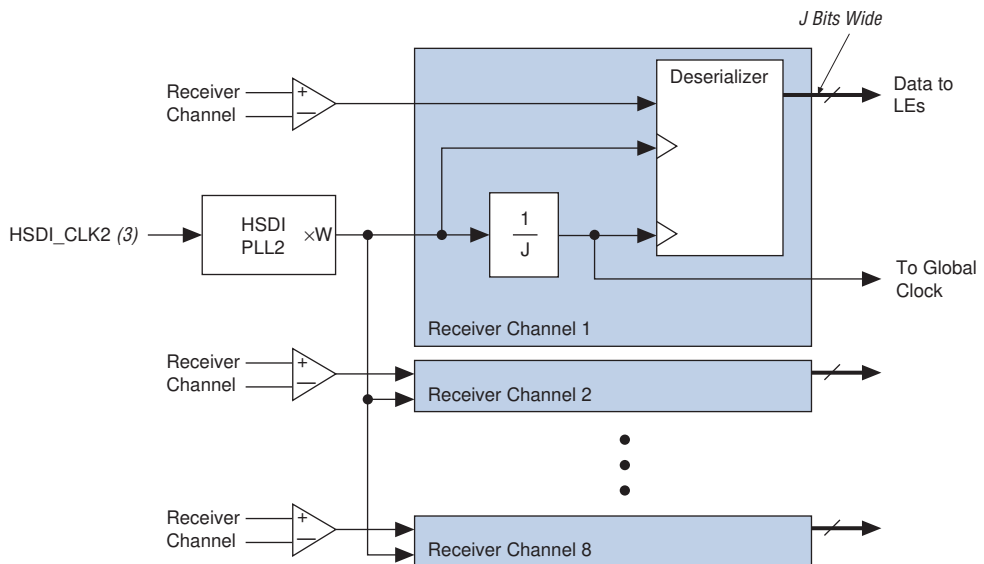
- Gigabit Ethernet backplanes
- ATM, SONET
- RapidIO
- POS-PHY Level 4
- Fibre Channel
- SDTV

The HSDI band supports one of two possible modes:

- Source-synchronous mode
- Clock data recovery (CDR) mode

In source-synchronous mode, source synchronous interfacing is supported at up to 840 Mbps. Serial channels are transmitted and received along with a low speed clock. The receiving device then multiplies the clock by a factor of 1 to 12, 14, 16, 18, or 20. The serialization/deserialization rate can be any number from 4, 7, 8, 9 to 12, 14, 16, 18, or 20 and does not have to equal the clock multiplication value. For example, an 840-Mbps LVDS channel can be received along with a 84-MHz clock. The 84-MHz clock is multiplied by 10 to drive the serial shift register, but the register can be clocked out in parallel at 7-, 8-, 9- to 12-, 14-, 16-, 18-, or 20-bits wide at 42 to 120 MHz. See Figures 2 and 3.

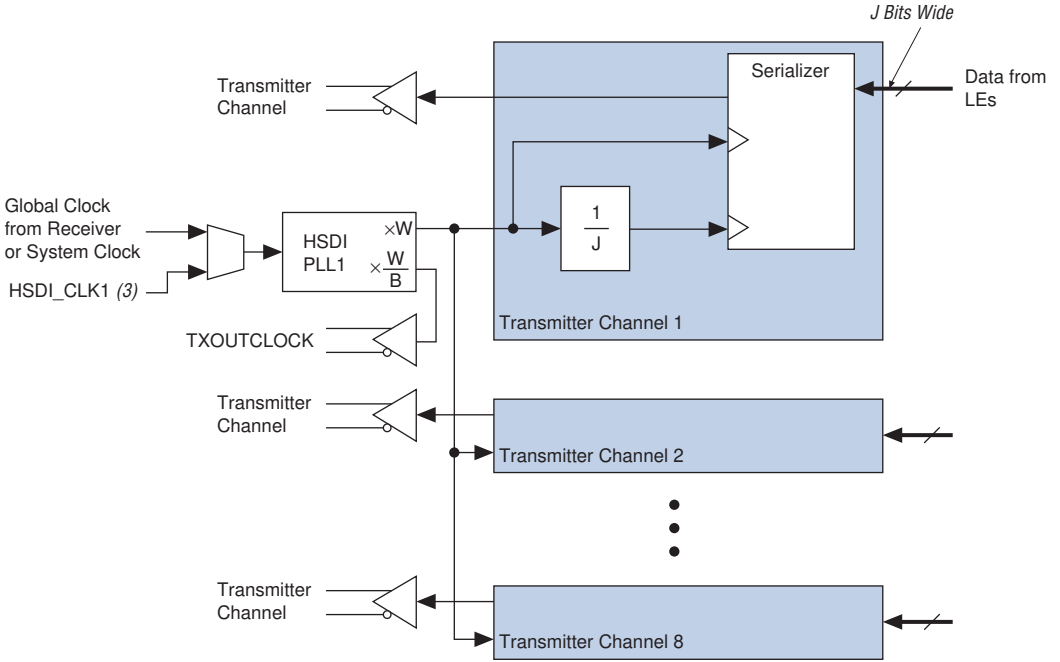
Figure 2. Receiver Diagram for Source Synchronous Mode Notes (1), (2)



Notes to Figure 2:

- (1) EP1M350 devices have 18 individual receiver channels. EP1M120 devices have 8 individual receiver channels.
- (2) $W = 1$ to 12, 14, 16, 18, or 20
 $J = 4, 7, 8, 9$ to 12, 14, 16, 18, or 20
 W does not have to equal J .
- (3) This clock pin drives an HSDI PLL only. It does not drive to the core.

Figure 3. Transmitter Diagram for Source Synchronous Mode Notes (1), (2)



Notes to Figure 3:

- (1) EP1M350 devices have 18 individual transmitter channels. EP1M120 devices have 8 individual transmitter channels.
- (2) $W = 1$ to 12, 14, 16, 18, or 20
 $B = 1$ to 12, 14, 16, 18, or 20
 $J = 4, 7, 8, 9$ to 12, 14, 16, 18, or 20
 $W, B,$ and J do not have to be equal.
- (3) This clock pin drives an HSDI PLL only. It does not drive to the logic array.

The Mercury device’s source-synchronous mode also supports the RapidIO interface protocol at up to 500 Mbps using the LVDS I/O standard.



For more information on source synchronous interfacing see [AN 159: Using HSDI in Source-Synchronous Mode in Mercury Devices](#).

Table 6 defines the support for source-synchronous mode applications.

Table 6. Source-Synchronous Mode			
Data Rate	I/O Standard		
	LVDS	LVPECL	3.3-V PCML
≤ 840 Mbps	(1)	✓	✓

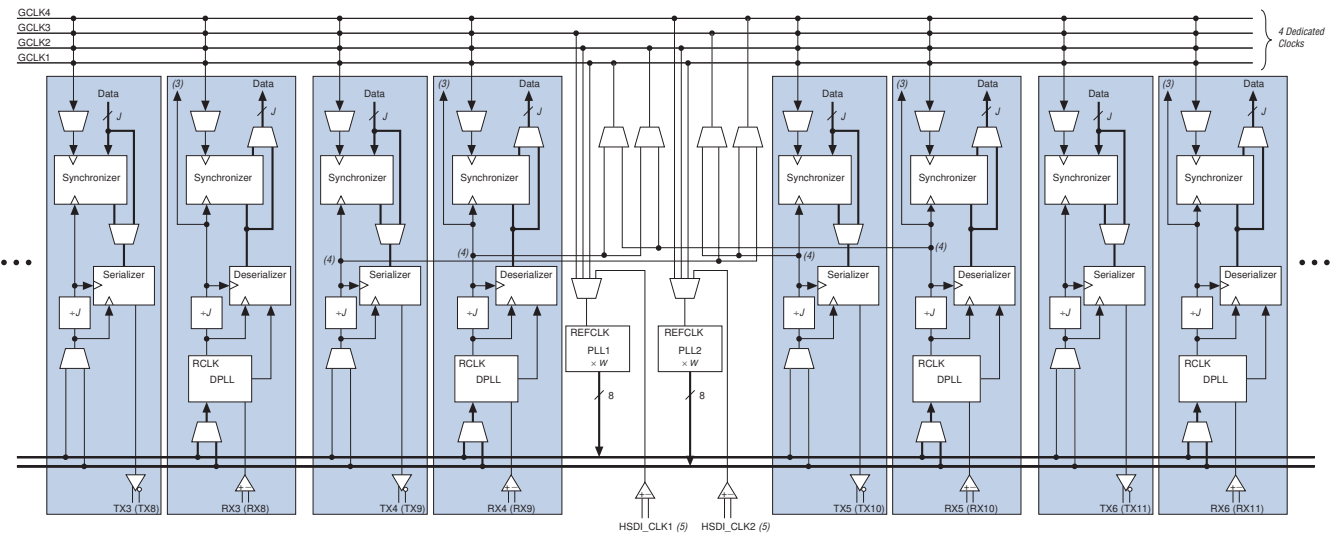
Note to Table 6:

- (1) You can use the CDR circuit to achieve data rates for DC coupled LVDS applications. You must AC-couple the clock to a 2.2-V common mode voltage (V_{CM}) using the AC-coupling schemes in *AN 134: Using Programmable I/O Standards in Mercury Devices*. The data channels should be DC-coupled. The byte alignment relative to the clock is lost when using the CDR circuit. Therefore, a byte-alignment circuit is required. Most Mercury source-synchronous designs already include byte-alignment logic since they usually use DDR or SDR clocks. The CDR run length requirement is waived if the reference clock and the receiver data come from the same source and have the same frequency.

In CDR mode, serial data is supported up to 1.25 Gbps per channel. The system provides a reference clock which is multiplied by the receiver or transmitter PLL to the same rate as the data is provided. For the receiver, this multiplied reference clock is used by a CRU on each receiver channel to generate a recovered clock in-phase with the received data. That recovered clock drives the programmable deserializer and synchronizer. The synchronizer is a FIFO for data transfer between the recovered clock domain and the global clock domain. The dedicated synchronizers can be bypassed if necessary. For every receiver channel in the EP1M350 and EP1M120 devices, the $\div J$ recovered clock can drive a priority column line for use as a clock. See [Figure 4](#).

Figure 4. Receiver & Transmitter Diagrams for CDR Mode

Notes (1), (2)



Notes to Figure 4:

- (1) EP1M350 devices have 18 individual receiver and transmitter channels. EP1M120 devices have 8 individual receiver and transmitter channels. Receiver and transmitter channel numbers in parenthesis are for EP1M350 devices.
- (2) $W = 1$ to 12, 14, 16, 18, or 20
 $J = 3$ to 12, 14, 16, 18, or 20
 W does not have to equal J .
- (3) For every receiver channel in EP1M350 and EP1M120 devices, the $\div J$ recovered clock can drive the priority column interconnect for use as a clock.
- (4) The two center channels adjacent to the HSDI PLLs (channels 4 and 5 for EP1M120 devices, channels 9 and 10 for EP1M350 devices) can drive the Mercury device's global clocks.
- (5) HSDI_CLK1 and HSDI_CLK2 pins must be differential. These clock pins drive HSDI PLLs only. They do not drive to the logic array.

The multiplied reference clock is also used to synchronize and serialize at the transmitter side.

Up to two different serial data rates are supported for input channels or output channels. Received data must be non-return-to-zero (NRZ).

Table 7 defines the support for CDR-mode applications. Table 8 shows the supported data rates for each speed grade.

Table 7. CDR-Mode Applications						
Data Rate	CDR Mode					
	DC-Coupled LVDS	DC-Coupled LVPECL	DC-Coupled 3.3-V PCML	AC-Coupled LVDS (1)	AC-Coupled LVPECL (1)	AC-Coupled 3.3-V PCML (1)
1.0 to 1.25 Gbps	(2)	✓	✓	✓	✓	✓
≤ 1.0 Gbps	✓	✓	✓	✓	✓	✓

Notes to Table 7:

- (1) The V_{CM} operating range for AC-coupled applications is from 0 to 0.7 V and from 1.8 to 2.4 V.
- (2) Use AC-coupled LVDS or another I/O standard. The DC-coupled LVDS I/O standard provides performance up to 1.0 Gbps.



For more information on CDR, see [AN 130: CDR in Mercury Devices](#).



Mercury device HSDI performance is finalized for certain speed grades. Also, the industrial-grade CDR specification is the same as the -6 speed grade for commercial-grade CDR specification. See [Table 8](#).

Table 8. CDR & Source-Synchronous Data Rates

Device	Speed Grade	Number of Channels	Maximum CDR Data Rate (Gbps)	Maximum Source-Synchronous Data Rate (Mbps)
EP1M120	-5	8	1.25	840
	-6 (1)	8	1.25	840
	-7	8	1.0	840
EP1M350	-5	18	1.25	840
	-6 (1)	8 (2)	1.25	840
		10 (2)	1.0	840
	-7	18	1.0	840

Notes to Table 8:

- (1) The -6 speed grade specifications apply for both commercial and industrial devices.
- (2) EP1M350 devices can support any 8 channels at 1.25 Gbps. The other 10 channels must run at 1.0 Gbps or less.


Logic & Interconnect

Mercury device logic is implemented in LEs. LE resources are used differently according to specific operating modes and the type of logic function being implemented. LEs are grouped into LABs in a row-based architecture. The multi-level FastTrack Interconnect structure provides the routing connection between LEs, ESBs, and IOEs.

Logic Array Block

Each LAB consists of 10 LEs, LE carry chains, multiplier circuitry, LAB control signals, local interconnect, and FastLUT connection lines. The local interconnect transfers signals between LEs within the same or adjacent LABs. FastLUT connections transfer the output of one LE to the adjacent LE for ultra-fast sequential LE connections within the same LAB. The Quartus II Compiler places associated logic within a LAB or adjacent LABs, allowing the use of fast local and FastLUT connections for high performance. [Figure 5](#) shows the Mercury LAB structure.

Figure 5. Mercury LAB Structure

**Notes to Figure 5:**

- (1) Priority column lines drive priority row lines, but not other row lines.
- (2) The RapidLAB interconnect can be driven by priority column lines, but not other column lines.
- (3) In multiplier mode, the RapidLAB interconnect drives LEs directly.

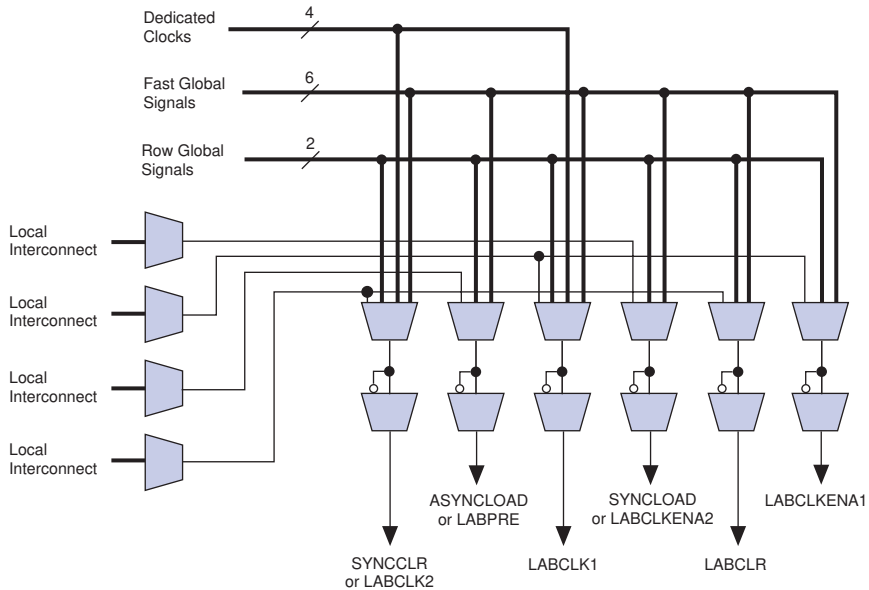
Mercury devices use an interleaved LAB structure, which allows each LAB to drive two local interconnect areas. Every other LE drives to either the left or right local interconnect area, alternating by LE. The local interconnect can drive LEs within the same LAB or adjacent LABs. This feature minimizes use of the row and column interconnects, providing higher performance and flexibility. Each LAB structure can drive 30 LEs through fast local interconnects.

Each LAB contains dedicated logic for driving control signals to its LEs. The control signals include clock, clock enable, asynchronous clear, asynchronous preset, asynchronous load, synchronous clear, and synchronous load signals. A maximum of six control signals can be used 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 two clocks and two clock enable signals. Each LAB's clock and clock enable signals are linked (e.g., any LE in a particular LAB using `LABCLK1` will also use `LABCLKENA1`). In addition to LAB-wide control of clock enables, Mercury devices can also control clock enable signals on individual LEs, allowing more than two clock enables in a given LAB. The Quartus II software automatically chooses whether a clock enable is LAB-wide for individual LEs. If both the rising and falling edges of a clock are used in a LAB, both LAB-wide clock signals are used.

The LAB local interconnect, fast global signals, row-global signals, and dedicated clock pins can generate the LAB-wide control signals. The multi-level FastTrack Interconnect's inherent low skew allows it to be used for clock distribution. [Figure 6](#) shows the LAB control signal generation circuit.

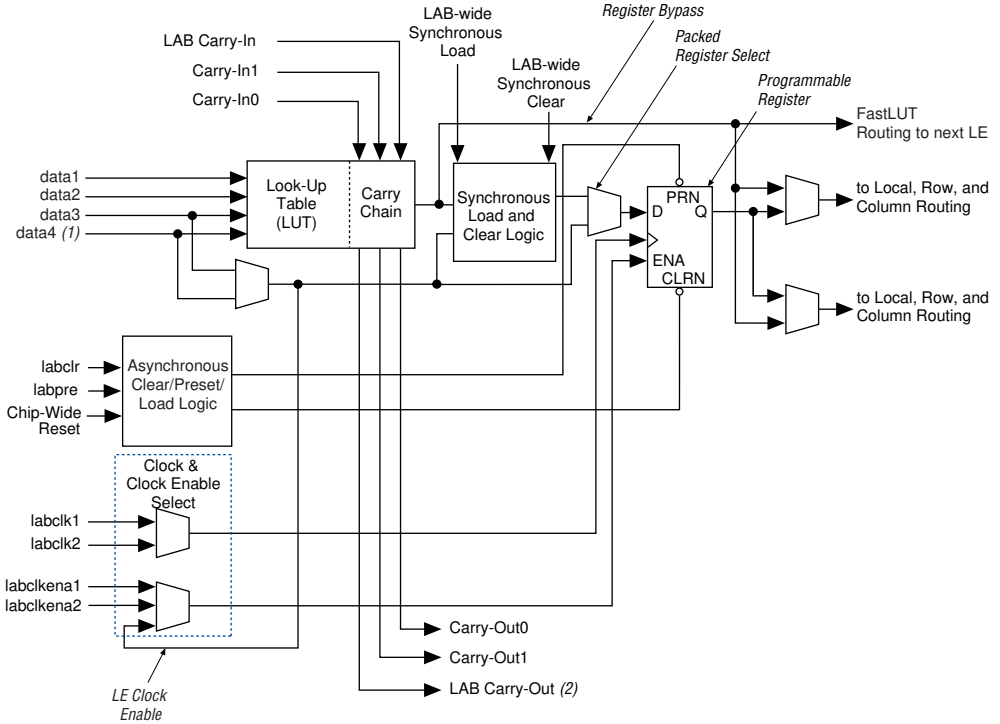
Figure 6. LAB-Wide Control Signals



Logic Element

The LE, the smallest unit of logic in the Mercury architecture, is compact and provides efficient logic usage. Each LE contains a four-input LUT, which is a function generator that can quickly implement any function of four variables. In addition, each LE contains a programmable register and carry chain with carry select look ahead capability. Each LE drives all interconnect types: local interconnect, row and priority row interconnect, column and priority column interconnect, leap lines, and RapidLAB interconnect. Each LE also has the ability to drive its combinatorial output directly to the next LE in the LAB using FastLUT connections. See [Figure 7](#).

Figure 7. Mercury LE



Notes to Figure 7:

- (1) FastLUT interconnect uses the data4 input.
- (2) LAB carry-out can only be generated by LE 4 and/or LE 10.

Each LE’s programmable register can be configured for D, T, JK, or SR operation. The register’s clock, clock enable, and clear control signals can be driven by global signals, general-purpose I/O pins, or any internal logic. For combinational functions, the register is bypassed and the output of the LUT drives directly to the outputs of the LE.

Each LE has four data inputs that can drive the internal LUT. One of these inputs has a shorter delay than the others, improving overall LE performance. This input is chosen automatically by the Quartus II software as appropriate.

Each LE has two outputs that drive the local, row, and column routing resources. Each output can be driven independently by the LUT's or register's output. For example, the LUT can drive one output, while the register drives the other output. This feature, called register packing, improves device utilization because the register and the LUT can be used for unrelated functions. The LE can also drive out registered and unregistered versions of the LUT output.

LE Operating Modes

The Mercury LE can operate in one of the following modes:

- Normal
- Arithmetic
- Multiplier

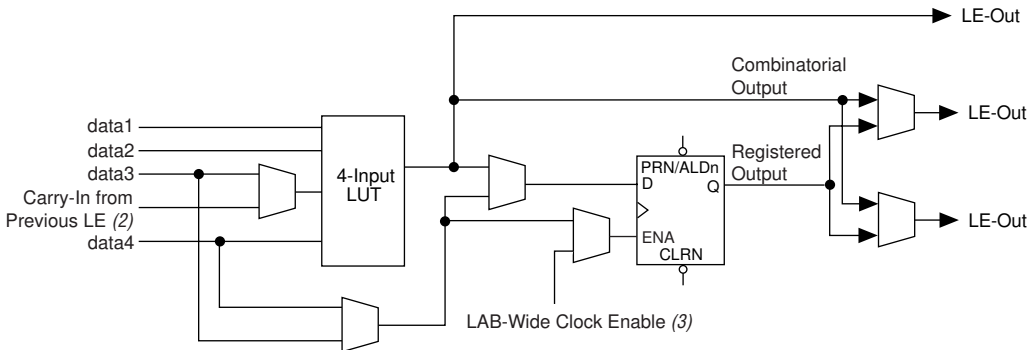
Each operating mode uses LE resources differently. In each operating mode, eight available inputs to the LE—the four data inputs from the LAB local interconnect; `carry-in0`, `carry-in1` from the previous LE; the LAB carry-in from the previous carry-chain generation; and the FastLUT Connection input from the previous LE—are directed to different destinations to implement the desired logic function. LAB-wide signals provide clock, asynchronous clear, asynchronous preset, asynchronous load, synchronous clear, synchronous load, and clock enable control for the register. These LAB-wide signals are available in all normal and arithmetic LE modes.

The Quartus II software, in conjunction with parameterized functions such as LPM and DesignWare functions, automatically chooses the appropriate mode for common functions, such as counters, adders, and multipliers. If required, the designer can also create special-purpose functions that specify which LE operating mode to use for optimal performance.

Normal Mode

The normal mode is suitable for general logic applications and combinatorial functions. In normal mode, four data inputs from the LAB local interconnect and a single carry-in are inputs to a four-input LUT. The Quartus II Compiler automatically selects the carry-in or the `data3` signal as one of the inputs to the LUT. The LUT (combinatorial) output can be driven to the FastLUT connection to the next LE in the LAB. LEs in normal mode support packed registers. [Figure 8](#) shows an LE in normal mode.

Figure 8. Normal-Mode LE *Note (1)*



Notes to Figure 8:

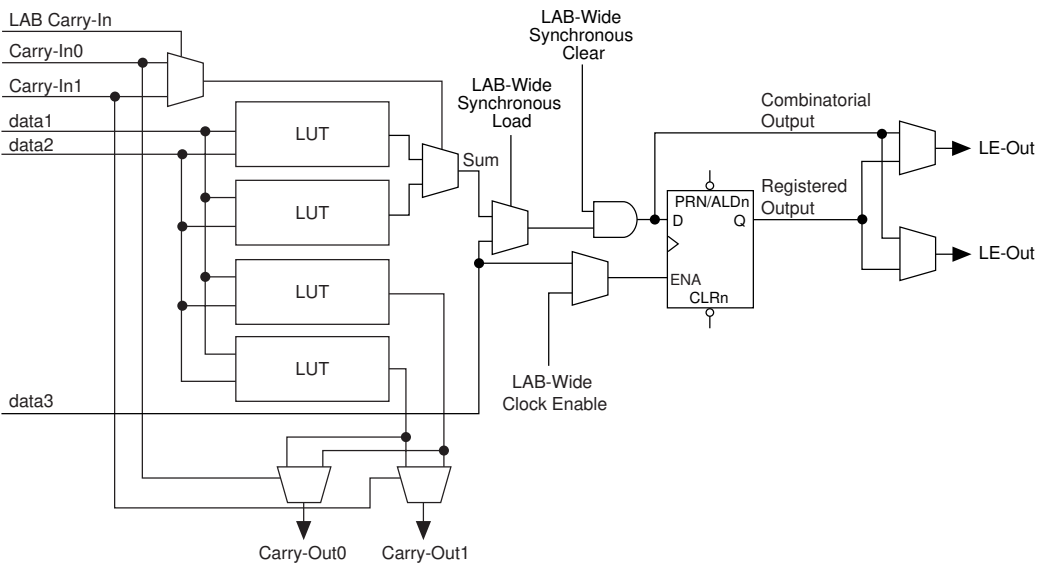
- (1) LEs in normal mode support register packing.
- (2) When using the carry-in in normal mode, the packed register feature is unavailable.
- (3) There are two LAB-wide clock enables per LAB in addition to LE-specific clock enables.

Arithmetic Mode

The arithmetic mode is ideal for implementing adders, accumulators, and comparators. A LE in arithmetic mode contains four 2-input LUTs. The first two 2-input LUTs compute two summations based on a possible carry of 1 or 0; the other two LUTs generate carry outputs for the two possible chains of the carry-select look-ahead (CSLA) circuitry. As shown in Figure 9, the LAB carry-in signal selects the appropriate carry-in chain (either carry-in0 or carry-in1). The logic level of the chain selected in turn selects which parallel sum is generated as a combinatorial or registered output. For example, when implementing an adder, this output is the signal comprised of the sum $data1 + data2 + carry$, where *carry* is 0 or 1. The other two LUTs use the *data1* and *data2* signals to generate two possible carry-out signals—one for a carry of 1 and the other for a carry of 0. The *carry-in0* signal acts as the carry select for the *carry-out0* output; *carry-in1* acts as the carry select for the *carry-out1* output. LEs in arithmetic mode can drive out registered and unregistered versions of the LUT output. Figure 9 shows a Mercury LE in arithmetic mode.

The arithmetic mode also offers clock enable, counter enable, synchronous up/down control, synchronous clear, and synchronous load options. The counter enable and synchronous up/down control signals are generated from the data inputs of the LAB local interconnect. The synchronous clear and synchronous load options are LAB-wide signals that affect all registers in the LAB. Consequently, if any of the LEs in a LAB use the counter mode, other LEs in that LAB must be used as part of the same counter or be used for a combinatorial function. The Quartus II software automatically places any registers that are not used by the counter into other LABs.

Figure 9. Arithmetic Mode LE



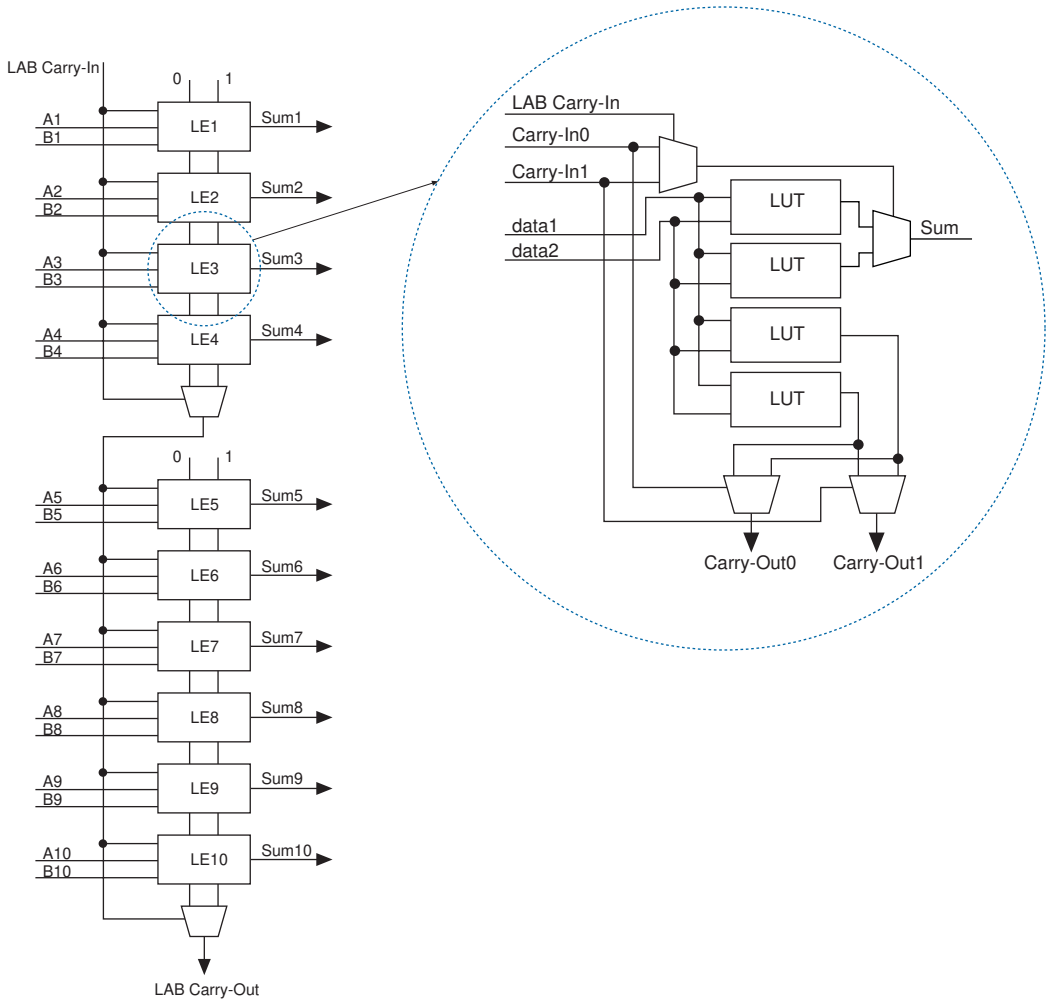
Carry-Select Look-Ahead Chain

The CSLA chain provides a very fast carry-forward function between LEs in arithmetic mode or multiplier mode. The CSLA chain uses the redundant carry calculation to increase the speed of carry functions. The LE can calculate sum and carry values for a possible carry-in of 1 and carry-in of 0 in parallel. The carry-in0 and carry-in1 signals from a lower-order bit drive forward into the higher-order bit via the parallel carry chain and feed into both the LUT and the next portion of the CSLA chain. CSLA chains can begin in any LE within a LAB.

The CSLA chain's speed advantage results from the parallel pre-computation of carry chains. Instead of including every LUT in the critical path, only the propagation delays between LAB carry-in generation circuits (LE 4 and LE 10) make up the critical path. This feature allows the Mercury architecture to implement high-speed counters, adders, multipliers, parity functions, and comparators of arbitrary width.

Figure 10 shows the CSLA circuitry in a LAB for a 10-bit full adder. One portion of the LUT generates the sum of two bits using the input signals and the appropriate carry-in bit; the sum is routed to the output of the LE. The register can be bypassed for simple adders or used for accumulator functions. Another portion of the LUT generates carry-out bits. A lab-wide carry-in bit selects which chain is used for the addition of given inputs. The actual carry-in signal for that selected chain, `carry-in0` or `carry-in1`, selects the carry-out to carry forward, which is routed to the carry-in signal of the next-higher-order bit. The final carry-out signal is routed to an LE, where it is driven to local, row, or column interconnects.

Figure 10. CSLA Details



The Quartus II Compiler can create CSLA logic automatically during design processing. Alternatively, the designer can create CSLA logic manually during design entry. Parameterized functions such as library of parameterized modules (LPM) and DesignWare functions automatically take advantage of carry chains for the appropriate functions.

The Quartus II Compiler creates carry chains longer than ten LEs by linking LABs together automatically. For enhanced fitting, a long carry chain skips intermediate LABs in a row structure. A carry chain longer than one LAB skips either from an even-numbered LAB to the next even-numbered LAB, or from an odd-numbered LAB to the next odd-numbered LAB. For example, the last LE of the first LAB in a LAB row carries to the first LE of the third LAB in the same LAB row.

Multiplier Mode

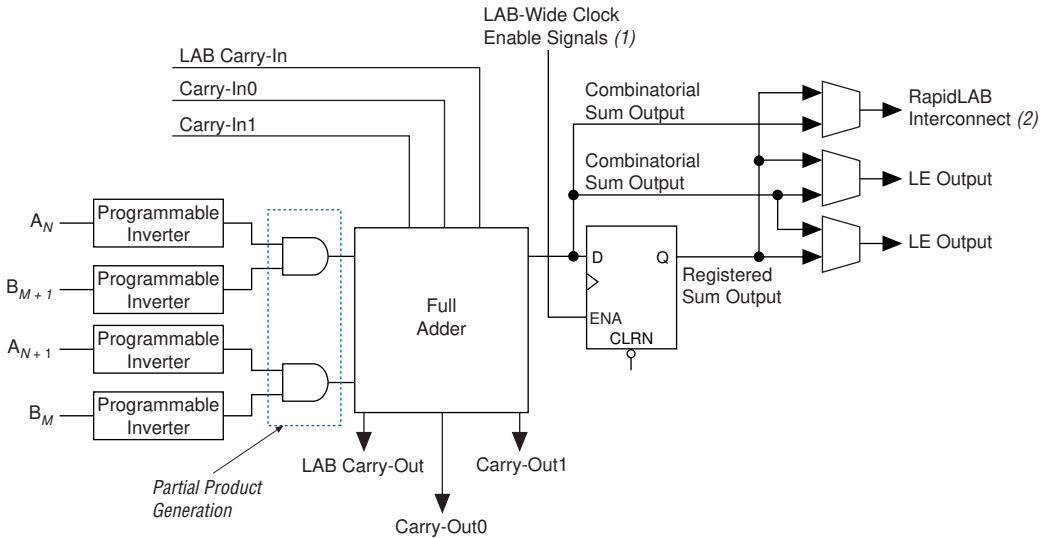
Multiplier mode is used for implementing high-speed multipliers up to 16×16 in size. The LUT implements the partial product formation and summation in a single stage for a $N \times M$ -bit multiply operation. A single LE can implement the summation of $A_N B_{M+1} + A_{N+1} B_M$ for the multiplier and multiplicand inputs. To increase the speed of the multiplication, LAB wide signals are used to control the partial product sum generation. These multiplier LAB-wide signals use the `LABCLKENA1` and `PRESET/ASYNLOAD` resources. The multiplier mode takes advantage of the CSLA circuitry for optimized sum and carry generation in the partial product sum. There is a special CSLA circuitry mode used for the multiplier where the carry chain runs vertically between LABs in the same column. The Quartus II Compiler automatically uses this special mode for dedicated multiplier implementation only. The summation of the multiplier and multiplicand bits is driven out along with the `carry-out0` and `carry-out1` bits. The combinatorial or registered versions of the sum can be driven out, allowing the multiplier to be pipelined.

The RapidLAB interconnect has dedicated fast connections to the LE inputs in multiplier mode, further increasing the speed of the multiplier. These dedicated connections allow RapidLAB lines to avoid delay incurred by driving onto local interconnects and then into the LE.

The Quartus II software implements parameterized functions that use the multiplier mode automatically when multiply operators are used.

Figure 11 shows a Mercury device LE in multiplier mode.

Figure 11. Multiplier Mode LE

**Notes to Figure 11:**

- (1) LABCLKENA1 cannot be used in multiplier mode.
- (2) When the RapidLAB output is used, local interconnect outputs are unavailable.

The basis for the high-speed 16×16 -bit multiplier in a Mercury device is the binary tree multiplier. In the first stage of the binary tree, the multiplicand bits, $a[15:0]$, and the multiplier bits, $b[15:0]$, are multiplied together. The results of the first stage are sixteen 16-bit partial products, $a[15:0]b[15]$, $a[15:0]b[14]$, \dots , $a[15:0]b[0]$. The partial products are then grouped into pairs and added together in the second stage. In a similar fashion, the results of the previous stage are grouped in pairs and then added forming the binary tree structure seen in Figure 12.