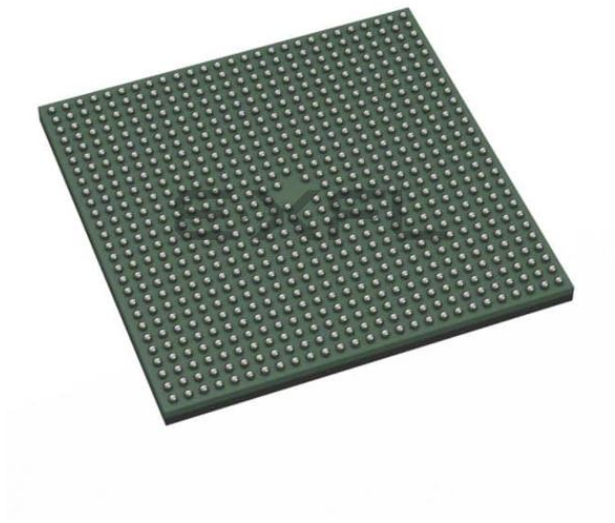

Welcome to E-XFL.COM

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	1664
Number of Logic Elements/Cells	16640
Total RAM Bits	425984
Number of I/O	492
Number of Gates	1900000
Voltage - Supply	1.425V ~ 1.575V
Mounting Type	Surface Mount
Operating Temperature	0°C ~ 85°C (TJ)
Package / Case	724-BBGA, FCBGA
Supplier Device Package	724-BGA (35x35)
Purchase URL	https://www.e-xfl.com/product-detail/intel/ep2a15b724c9

LE Operating Modes

The APEX II LE can operate in one of the following three modes:

- Normal mode
- Arithmetic mode
- Counter mode

Each mode uses LE resources differently. In each mode, seven available inputs to the LE—the four data inputs from the LAB local interconnect, the feedback from the programmable register, and the carry-in and cascade-in 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 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. Figure 8 shows the LE operating modes.

Normal Mode

The normal mode is suitable for general logic applications, combinatorial functions, or wide decoding functions that can take advantage of a cascade chain. In normal mode, four data inputs from the LAB local interconnect and the 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 output can be combined with the cascade-in signal to form a cascade chain through the cascade-out signal. LEs in normal mode support packed registers.

Arithmetic Mode

The arithmetic mode is ideal for implementing adders, accumulators, and comparators. An LE in arithmetic mode uses two 3-input LUTs. One LUT computes a three-input function; the other generates a carry output. As shown in Figure 8, the first LUT uses the carry-in signal and two data inputs from the LAB local interconnect to generate a combinatorial or registered output. For example, when implementing an adder, this output is the sum of three signals: DATA1, DATA2, and carry-in. The second LUT uses the same three signals to generate a carry-out signal, thereby creating a carry chain. The arithmetic mode also supports simultaneous use of the cascade chain. LEs in arithmetic mode can drive out registered and unregistered versions of the LUT output.

The Quartus II software implements parameterized functions that use the arithmetic mode automatically where appropriate; the designer does not need to specify how the carry chain will be used.

Counter Mode

The counter mode 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 an 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.

The counter mode uses two three-input LUTs: one generates the counter data, and the other generates the fast carry bit. A 2-to-1 multiplexer provides synchronous loading, and another AND gate provides synchronous clearing. If the cascade function is used by an LE in counter mode, the synchronous clear or load overrides any signal carried on the cascade chain. The synchronous clear overrides the synchronous load. LEs in arithmetic mode can drive out registered and unregistered versions of the LUT output.

Clear & Preset Logic Control

Logic for the register's clear and preset signals is controlled by LAB-wide signals. The LE directly supports an asynchronous clear function. The Quartus II Compiler can use a NOT-gate push-back technique to emulate an asynchronous preset. Moreover, the Quartus II Compiler can use a programmable NOT-gate push-back technique to emulate simultaneous preset and clear or asynchronous load. However, this technique uses three additional LEs per register. All emulation is performed automatically when the design is compiled. Registers that emulate simultaneous preset and load will enter an unknown state upon power-up or when the chip-wide reset is asserted.

In addition to the two clear and preset modes, APEX II devices provide a chip-wide reset pin (DEV_CLRn) that resets all registers in the device. Use of this pin is controlled through an option in the Quartus II software that is set before compilation. The chip-wide reset overrides all other control signals. Registers using an asynchronous preset are preset when the chip-wide reset is asserted; this effect results from the inversion technique used to implement the asynchronous preset.

FastTrack Interconnect

In the APEX II architecture, connections between LEs, ESBs, and I/O pins are provided by the FastTrack interconnect. The FastTrack interconnect is a series of continuous horizontal and vertical routing channels that traverse the device. This global routing structure provides predictable performance, even in complex designs. In contrast, the segmented routing in FPGAs requires switch matrices to connect a variable number of routing paths, increasing the delays between logic resources and reducing performance.

Figure 12. APEX II FastRow Interconnect

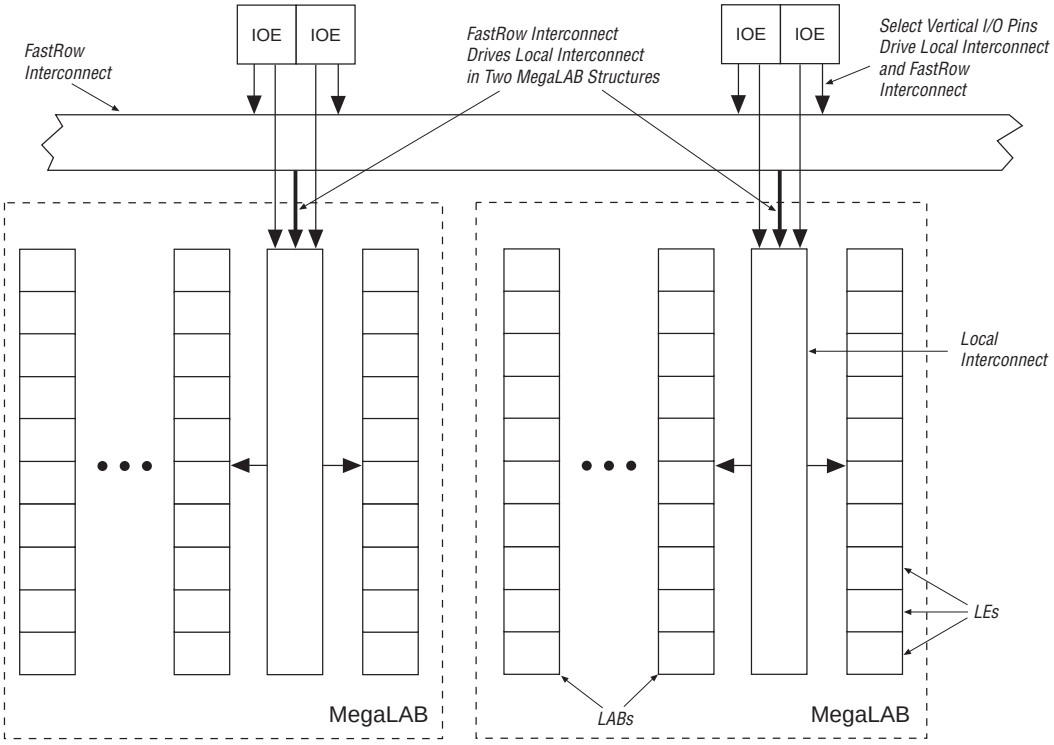

Figure 14. APEX II Macrocell

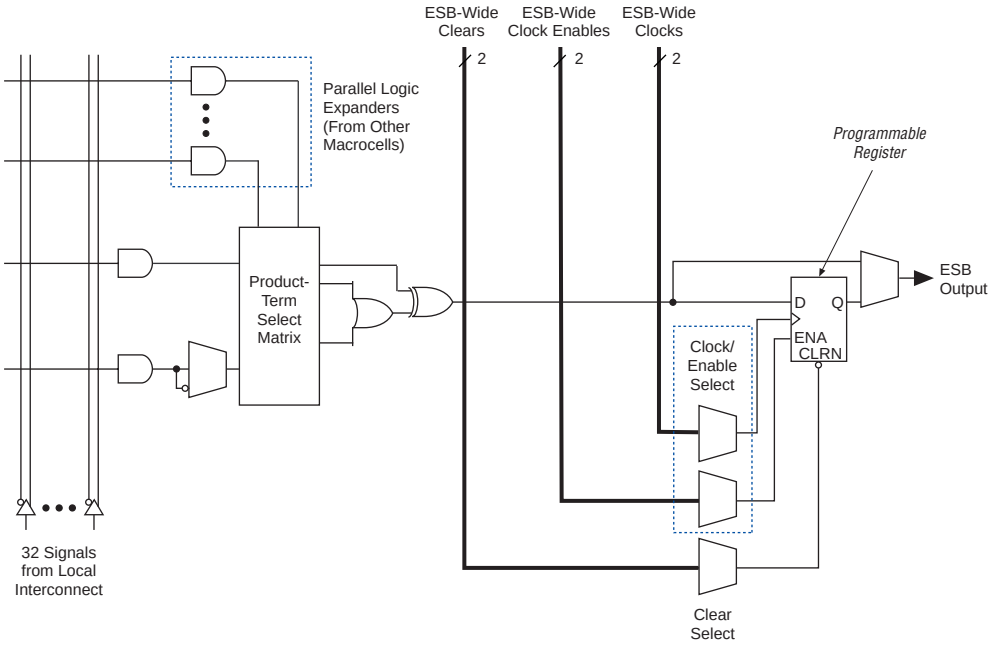

For registered functions, each macrocell register can be programmed individually to implement D, T, JK, or SR operation with programmable clock control. The register can be bypassed for combinatorial operation. During design entry, the designer specifies the desired register type; the Quartus II software then selects the most efficient register operation for each registered function to optimize resource utilization. The Quartus II software or other synthesis tools can also select the most efficient register operation automatically when synthesizing HDL designs.

Each programmable register can be clocked by one of two ESB-wide clocks. The ESB-wide clocks can be generated from device dedicated clock pins, global signals, or local interconnect. Each clock also has an associated clock enable, generated from the local interconnect. The clock and clock enable signals are related for a particular ESB; any macrocell using a clock also uses the associated clock enable.

If both the rising and falling edges of a clock are used in an ESB, both ESB-wide clock signals are used.

Figure 17. Bidirectional Dual-Port Memory Configuration

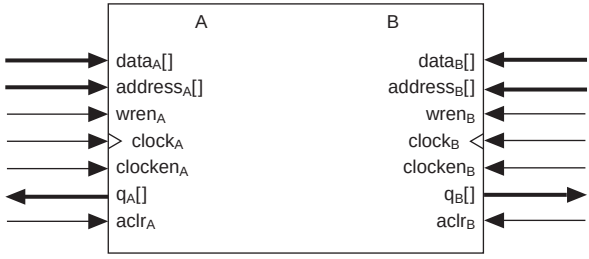

In addition to bidirectional dual-port memory, the ESB also supports dual-port, and single-port RAM. Dual-port memory supports a simultaneous read and write. Single-port memory supports independent read and write. Figure 18 shows these different RAM memory port configurations for an ESB.

Figure 18. Dual- & Single-Port Memory Configurations

Dual-Port Memory

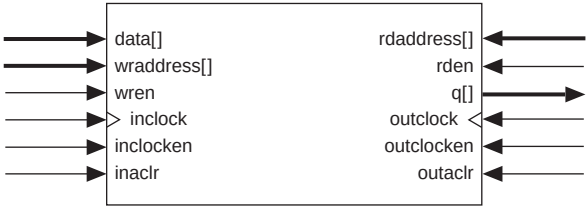

Single-Port Memory *(1)*

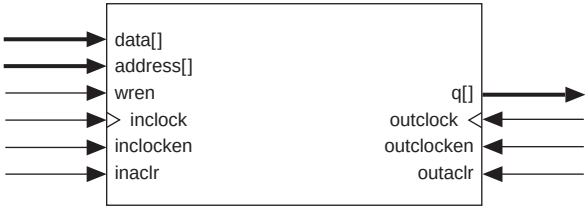

Note to Figure 18:

(1) Two single-port memory blocks can be implemented in a single ESB.

The ESB also enables variable width data ports for reading and writing to the RAM ports in dual-port RAM configuration. For example, the ESB can be written in 1× mode at port A while being read in 16× mode from port B. Table 6 lists the supported variable width configurations for an ESB in dual-port mode.

Table 6. Variable Width Configurations for Dual-Port RAM

Read Port Width	Write Port Width
1 bit	2 bits, 4 bits, 8 bits, or 16 bits
2 bits, 4 bits, 8 bits, or 16 bits	1 bit

ESBs can implement synchronous RAM, which is easier to use than asynchronous RAM. A circuit using asynchronous RAM must generate the RAM write enable (WE) signal while ensuring that its data and address signals meet setup and hold time specifications relative to the WE signal. In contrast, the ESB's synchronous RAM generates its own WE signal and is self-timed with respect to the global clock. Circuits using the ESB's self-timed RAM only need to meet the setup and hold time specifications of the global clock.

ESB inputs are driven by the adjacent local interconnect, which in turn can be driven by the MegaLAB or FastTrack interconnects. Because the ESB can be driven by the local interconnect, an adjacent LE can drive it directly for fast memory access. ESB outputs drive the MegaLAB and FastTrack interconnects and the local interconnect for fast connection to adjacent LEs or for fast feedback product-term logic.

When implementing memory, each ESB can be configured in any of the following sizes: 512 × 8, 1,024 × 4, 2,048 × 2, or 4,096 × 1. For dual-port and single-port modes, the ESB can be configured for 256 × 16 in addition to the list above.

The ESB can also be split in half and used for two independent 2,048-bit single-port RAM blocks. The two independent RAM blocks must have identical configurations with a maximum width of 256 × 8. For example, one half of the ESB can be used as a 256 × 8 single-port memory while the other half is also used for a 256 × 8 single-port memory. This effectively doubles the number of RAM blocks an APEX II device can implement for its given number of ESBs. The Quartus II software automatically merges two logical memory functions in a design into an ESB; the designer does not need to merge the functions manually.

Figure 19. ESB in Input/Output Clock Mode *Note (1)*

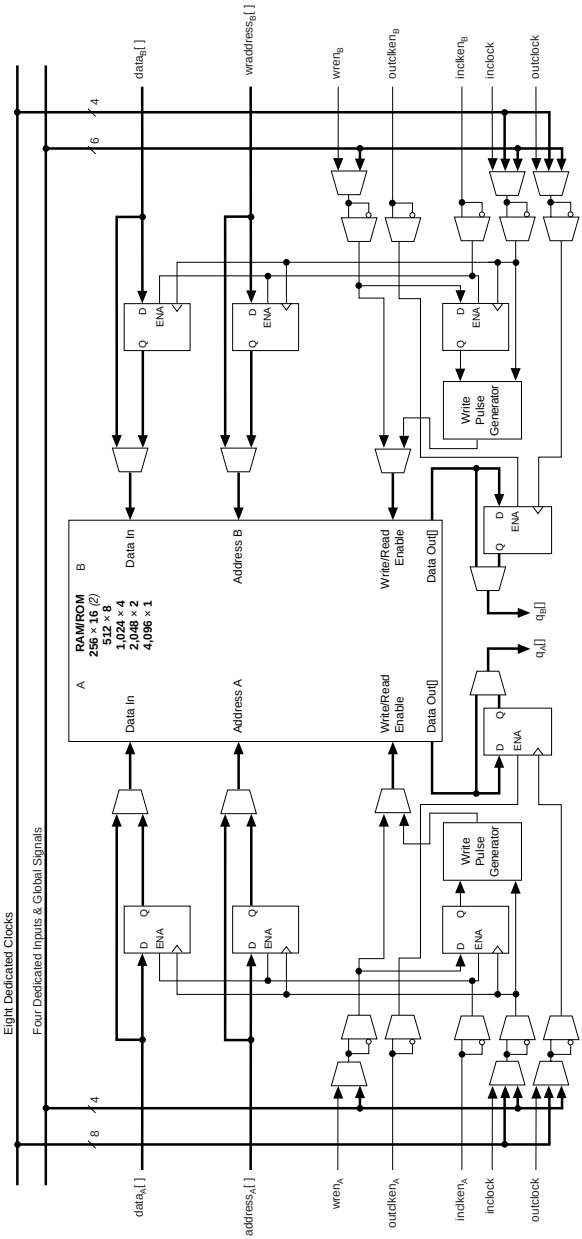

Notes to Figure 19:

(1) All registers can be cleared asynchronously by ESB local interconnect signals, global signals, or the chip-wide reset.

(2) This configuration is not supported for bidirectional dual-port configuration.

If the same data is written into multiple locations in the memory, a CAM block can be used in multiple-match or fast multiple-match modes. The ESB outputs the matched data's locations as an encoded or unencoded address. In multiple-match mode, it takes two clock cycles to write into a CAM block. For reading, there are 16 outputs from each ESB at each clock cycle. Therefore, it takes two clock cycles to represent the 32 words from a single ESB port. In this mode, encoded and unencoded outputs are available. To implement the encoded version, the Quartus II software adds a priority encoder with LEs. Fast multiple-match is identical to the multiple match mode, however, it only takes one clock cycle to read from a CAM block and generate valid outputs. To do this, the entire ESB is used to represent 16 outputs. In fast multiple-match mode, the ESB can implement a maximum CAM block size of 16 words.

A CAM block can be pre-loaded with data during configuration, or it can be written during system operation. In most cases, two clock cycles are required to write each word into CAM. When don't-care bits are used, a third clock cycle is required.

f For more information on CAM, see *Application Note 119 (Implementing High-Speed Search Applications with APEX CAM)*.

Driving Signals to the ESB

ESBs provide flexible options for driving control signals. Different clocks can be used for the ESB inputs and outputs. Registers can be inserted independently on the data input, data output, read address, write address, WE, and RE signals. The global signals and the local interconnect can drive the WE and RE signals. The global signals, dedicated clock pins, and local interconnects can drive the ESB clock signals. Because the LEs drive the local interconnect, the LEs can control the WE and RE signals and the ESB clock, clock enable, and synchronous clear signals. Figure 24 shows the ESB control signal generation logic.

Figure 26. Row IOE Connection to the Interconnect

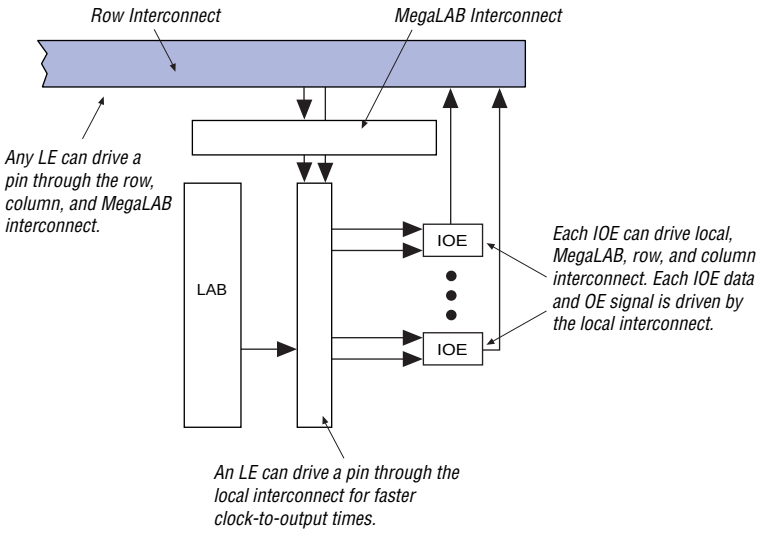

Figure 27 shows how a column IOE connects to the interconnect.

Figure 27. Column IOE Connection to the Interconnect

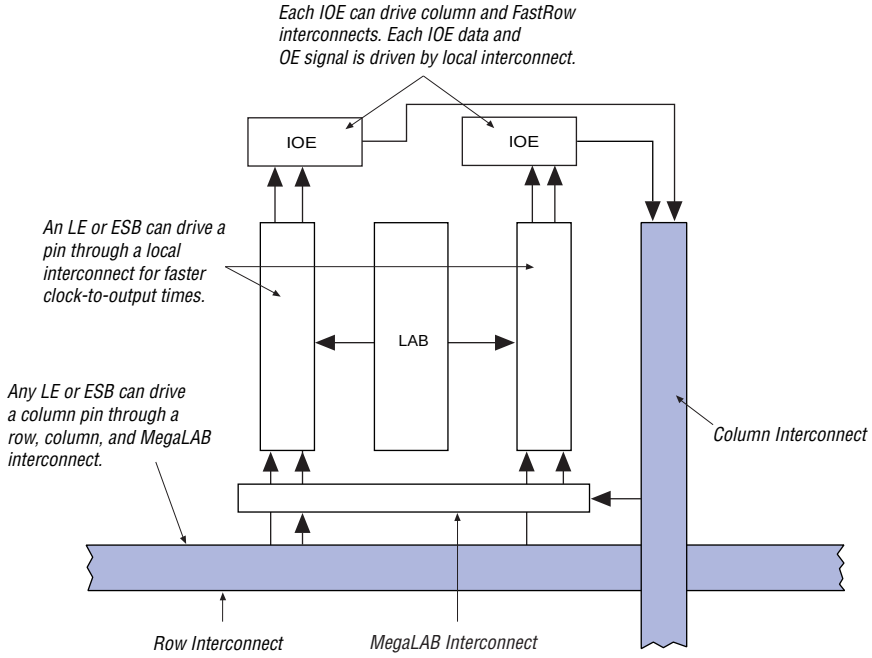

FastRow interconnects connect a column I/O pin directly to the LAB local interconnect within two MegaLAB structures. This feature provides fast setup times for pins that drive high fan-outs with complex logic, such as PCI designs. For fast bidirectional I/O timing, LE registers using local routing can improve setup times and OE timing.

APEX II devices have a peripheral control bus made up of 12 signals that drive the IOE control signals. The peripheral bus is composed of six output enables, `OE[5:0]` and six clock enables, `CE[5:0]`. These twelve signals can be driven from internal logic or from the Fast I/O signals. Table 7 lists the peripheral control signal destinations.

Bus Hold

Each APEX II device I/O pin provides an optional bus-hold feature. When this feature is enabled for an I/O pin, the bus-hold circuitry weakly holds the signal at its last driven state. By holding the last driven state of the pin until the next input signal is present, the bus-hold feature eliminates the need to add external pull-up or pull-down resistors to hold a signal level when the bus is tri-stated. The bus-hold circuitry also pulls undriven pins away from the input threshold voltage where noise can cause unintended high-frequency switching. This feature can be selected individually for each I/O pin. The bus-hold output will drive no higher than V_{CCIO} to prevent overdriving signals. If the bus-hold feature is enabled, the programmable pull-up option cannot be used. The bus-hold feature should also be disabled if open-drain outputs are used with the GTL+ I/O standard.

The bus-hold circuitry weakly pulls the signal level to the last driven state through a resistor with a nominal resistance (R_{BH}) of approximately 7 k . Table 41 on page 74 gives specific sustaining current that will be driven through this resistor and overdrive current that will identify the next driven input level. This information is provided for each V_{CCIO} voltage level.

The bus-hold circuitry is active only after configuration. When going into user mode, the bus-hold circuit captures the value on the pin present at the end of configuration.

Programmable Pull-Up Resistor

Each APEX II device I/O pin provides an optional programmable pull-up resistor during user mode. When this feature is enabled for an I/O pin, the pull-up resistor (typically 25 k) weakly holds the output to the V_{CCIO} level of the bank that the output pin resides in.

Dedicated Fast I/O Pins

APEX II devices incorporate dedicated bidirectional pins for signals with high internal fanout, such as PCI control signals. These pins are called dedicated fast I/O pins (FAST1, FAST2, FAST3, and FAST4) and can drive the four global fast lines throughout the device, ideal for fast clock, clock enable, preset, clear, or high fanout logic signal distribution. The dedicated fast I/O pins have one output register and one OE register, but they do not have input registers. The dedicated fast lines can also be driven by a LE local interconnect to generate internal global signals.

Figure 33. True-LVDS Transmitter Diagram *Notes (1), (2)*

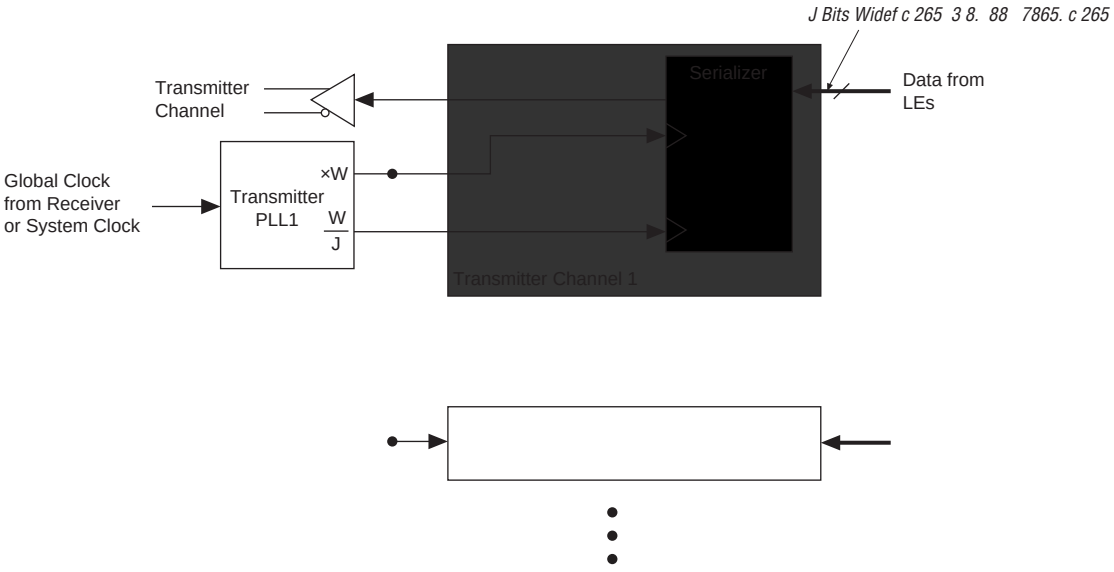

Notes to Figure 33:

(1) Two sets of 18 transmitter channels are located in each APEX II device. Each set of 18 channels has one transmitter PLL.

(2) *W* = 1, 2, 4 to 10
J = 1, 2, 4 to 10
W does not have to equal *J*. When *J* = 1 or 2, the deserializer is bypassed. When *J* = 2, DDR I / O registers are used.

Clock-Data Synchronization

In addition to dedicated serial-to-parallel converters, APEX II True-LVDS circuitry contains CDS circuitry in every receiver channel. The CDS feature can be turned on or off independently for each receiver channel. There are two modes for the CDS circuitry: single-bit mode, which corrects a fixed clock-to-data skew of up to ±50% of the data bit period, and multi-bit mode, which corrects any fixed clock-to-data skew.

Pre-programmed CDS may also be used to resolve clock-to-data skew greater than 50% of the bit period. However, internal logic must be used to implement the byte alignment circuitry for this operation.

Flexible-LVDS I/O Pins

A subset of pins in the top two I/O banks supports interfacing with Flexible-LVDS, LVPECL, and HyperTransport inputs. These Flexible-LVDS input pins include dedicated LVDS, LVPECL, and HyperTransport input buffers. A subset of pins in the bottom two I/O banks supports interfacing with Flexible-LVDS and HyperTransport outputs. These Flexible-LVDS output pins include dedicated LVDS and HyperTransport output buffers. The Flexible-LVDS pins do not require any external components except for 100- termination resistors on receiver channels. These pins do not contain dedicated serialization/deserialization circuitry; therefore, internal logic is used to perform serialization/deserialization functions.

The EP2A15 and EP2A25 devices support 56 input and 56 output Flexible-LVDS channels. The EP2A40 and larger devices support 88 input and 88 output Flexible-LVDS channels. All APEX II devices support the Flexible-LVDS interface up to 400 Mbps (DDR) per channel. Flexible-LVDS pins along with the True-LVDS pins provide up to 144-Gbps total device bandwidth. Table 13 shows the Flexible-LVDS timing specification.

Table 13. APEX II Flexible-LVDS Timing Specification

Symbol	Timing Parameter Definition	Speed Grade						Unit
		-7		-8		-9		
		Min	Max	Min	Max	Min	Max	
Data Rate	Maximum operating speed		400		311		311	Mbps
TCCS	Transmitter channel-to-channel skew		700		900		900	ps
SW	Receiver sampling window	1,100		1,400		1,400		ps

MultiVolt I/O Interface

The APEX II architecture supports the MultiVolt I/O interface feature, which allows APEX II devices in all packages to interface with systems of different supply voltages. The devices have one set of V_{CC} pins for internal operation and input buffers (VCCINT), and another set for I/O output drivers (VCCIO).

Note to Figure 35:

(1) *n* represents the prescale divider for the PLL input. *m* represents the multiplier. *k* and *v* represent the different post scale dividers for the two possible PLL outputs. *m* and *k* are integers that range from 1 to 160. *n* and *v* are integers that range from 1 to 16.

Advanced ClockBoost Multiplication & Division

APEX II PLLs include circuitry that provides clock synthesis for eight internal outputs and two external outputs using $m / (n \times \text{output divider})$ scaling. When a PLL is locked, the locked output clock aligns to the rising edge of the input clock. The closed loop equation for Figure 35 gives an output frequency $f_{clock0} = (m / (n \times k))f_{IN}$ and $f_{clock1} = (m / (n \times v))f_{IN}$. These equations allow the multiplication or division of clocks by a programmable number. The Quartus II software automatically chooses the appropriate scaling factors according to the frequency, multiplication, and division values entered.

A single PLL in an APEX II device allows for multiple user-defined multiplication and division ratios that are not possible even with multiple delay-locked loops (DLLs). For example, if a frequency scaling factor of 3.75 is needed for a given input clock, a multiplication factor of 15 and a division factor of 4 can be entered. This advanced multiplication scaling can be performed with a single PLL, making it unnecessary to cascade PLL outputs.

External Clock Outputs

APEX II devices have two low-jitter external clocks available for external clock sources. Other devices on the board can use these outputs as clock sources.

There are three modes for external clock outputs.

Zero Delay Buffer: The external clock output pin is phase aligned with the clock input pin for zero delay. Multiplication, programmable phase shift, and time delay shift are not allowed in this configuration. The MegaWizard interface for `altclklock` should be used to verify possible clock settings.
External Feedback: The external feedback input pin is phase aligned with clock input pin. By aligning these clocks, you can actively remove clock delay and skew between devices. This mode has the same restrictions as zero delay buffer mode.
Normal Mode: The external clock output pin will have phase delay relative to the clock input pin. If an internal clock is used in this mode, the IOE register clock will be phase aligned to the input clock pin. Multiplication is allowed with the normal mode.