Extending Transceiver Leadership at 28 nm

WP-01130-2.1 White Paper

High-speed serial protocols with increased data rates and expanded capabilities are addressing the demand for more network bandwidth. Efficiently addressing the subsequent increase in system bandwidth by attaining higher data rates and achieving greater integration is becoming an ever-greater challenge. This challenge includes targeting lower bit error ratios (BERs) and ensuring signal and power integrity while maintaining power efficiency and optimizing design productivity. This white paper explores transceiver architecture in Altera® 28 nm FPGAs for applications at 10 to 28 Gbps, and highlights the architectural advantages for making high performance systems with low BER.

Introduction

The current 28 nm process technology found in leading-edge ASICs, microprocessors, and FPGAs provides great capabilities for fast and power-efficient transceivers at 10 and 28 Gbps. The smaller feature sizes allow for shorter transistor channel lengths and shorter interconnects for logic gates, resulting in faster switching times and shorter interconnect transport delays. Smaller feature sizes allow greater hardware functions, faster operation speeds, higher logic density, tighter integration, and lower power consumption per logic function. The results of process node shrinkage are favorable for logic operations, higher density, higher data transmission, and other advanced features while optimized for power efficiency. These advances along with higher transceiver data rates at the 28 nm node enable wireline and wireless communication, computer, storage, military, and broadcast electronic systems to send and receive large amounts of data efficiently.

The data rates for most advanced transceivers used for communication and I/O standards today are either within the range of 5 to 6 Gbps for relatively established markets, or within the range of 8 to 14 Gbps for emerging markets. Some examples of applications using the 5 to 6 Gbps data rate include:

- *Network communications*—CEI/OIF 6G, 6G Interlaken, and CPRI 4.0
- *Computer I/O buses*—PCI Express® (PCIe®) Gen2 (5 Gbps), QPI, and HyperTransport™
- *Storage area networks*—Fibre Channel 4G (4.25 Gbps) and SATA III/SAS II (6 Gbps) Some examples of applications that use the 8 to 12 Gbps data rate include:
- *Network communications*—IEEE 10G Ethernet, IEEE 40G/100G Ethernet (802.3ba, 4X/10X 10.3125 Gbps, chip-to-chip, and chip-to-optical module over short, medium, and long channels, such as backplane and cable assembly), and CEI/OIF 11G (9.95 – 11.1 Gbps)

101 Innovation Drive
San Jose, CA 95134
www.altera.com

- *Computer I/O buses*—PCIe Gen3 (8.0 Gbps) and QPI
- *Storage area networks*—16G Fibre Channel (14.1 Gbps) and SATA IV/SAS III (12 Gbps)

Optical modules are used by the host receiver in communication and computer systems to convert electrical signals to optical signals before driving them to optical fiber channels. Similarly, optical modules are used by the host transmitter side to convert the optical signals to electrical signals before driving them to electrical copper channels. The evolution of optical modules focuses on increasing data rates, lowering power consumption, and reducing form factor by removing components off the module. A few examples of optical modules are:

- 10G form-factor pluggable (XFP)
- Small form-factor pluggable (SFP) and SFP+
- 40G (4 × 10G) quad small form-factor pluggable (QSFP)
- 100G (10 × 10G and 4 × 25G) 100G form-factor pluggable (CFP)

Altera's 28 nm Stratix® V FPGAs deliver the highest bandwidth and levels of system integration, and flexibility for high-end applications on a single chip. This device family goes beyond Moore's law with ground-breaking innovations in its core architecture with integrated transceivers up to 28 Gbps and a unique array of integrated hard intellectual property (IP) blocks. With these innovations, Stratix V FPGAs deliver a new class of application-targeted devices optimized for:

- High-performance, bandwidth-centric applications including PCIe Gen3
- Data-intensive applications for 40G/100G and beyond, such as 400G Ethernet
- Ultra-high bandwidth backplanes and switches
- High-performance, high-precision digital signal processing (DSP) applications

High-Speed Serial I/O Protocol Support

Stratix V FPGAs support data rates from 600 Mbps to 28 Gbps with the lowest bit error ratio (BER) and the most power-efficient (200 mW/channel at 28 Gbps) transceivers. Stratix V FPGAs have dedicated hardware and IP to support a wide range of high-speed I/O standards, covering applications in the wireline, wireless, computer storage, broadcast, military, test and measurement, and medical fields. Table 1, Table 2, and Table 3 highlight the 8 Gbps, 10 Gbps, and 28 Gbps high-speed standards supported by Stratix V FPGAs.

Table 1. 8G High-Speed Standards/Protocols Supported by Stratix V FPGAs (Part 1 of 2)

Standard	Electrical Serial Line Rate	Link	Lanes
PCIe 3.0	8 Gbps	Chip-to-chip and backplane	1, 2, 4, 8
PCIe 2.0	5 Gbps	Chip-to-chip and backplane	1, 2, 4, 8
Interlaken	4.976 Gbps to 6.375 Gbps	Chip-to-chip and backplane	1 to 24
Serial RapidIO® 2.0+	1.25, 2.5, 3.125, 5 to 6.25 Gbps	Chip-to-chip and backplane	1, 2, 4
CPRI 4.0+	0.6144, 1.2288, 2.4576, 3.072, 4.9152, 6.144 Gbps	Chip-to-chip	1 to *N*
OBSAI 4.0+ (RP3)	0.768, 1.536, 3.072, 6.144 Gbps	Chip-to-chip and backplane	1 to *N*

Table 1. 8G High-Speed Standards/Protocols Supported by Stratix V FPGAs (Part 2 of 2)

Standard	Electrical Serial Line Rate	Link	Lanes
SATA 3.0	6 Gbps	Chip-to-chip and backplane	1 to *N*
SAS 2.0	6 Gbps	Chip-to-chip and backplane	1 to *N*
SPAUI	6.375 Gbps	Chip-to-chip and backplane	4 or 6
DDR-XAUI	6.25 Gbps	Chip-to-chip and backplane	4
QPI	4, 4.8, 6.4, 8 Gbps	Chip-to-chip	(5, 10, 20)+1
HyperTransport 3.0+	0.4, 2.4, 2.8, 3.2 Gbps	Chip-to-chip and backplane	(2, 4, 8)+2, (16)+4
HighGig+, HighGig2+	3.75, 6.25 Gbps	Chip-to-chip and backplane	4
8G FC	8.5 Gbps	Chip-to-chip and chip-to-module	1 to *N*
OIF/CEI 6G-SR	4.976 to 6.375 Gbps	Chip-to-chip	I/O technology
OIF/CEI 6G-LR	4.976 to 6.375 Gbps	Backplane	I/O technology
4G FC	4.25 Gbps	Chip-to-chip and chip-to-module	1 to *N*

Table 2. 10G High-Speed Standards/Protocols Supported by Stratix V FPGAs

Standard	Electrical Serial Lane Rate	Link	Lanes
IEEE 802.3ba 40G	10.3125 Gbps	Chip-to-chip, backplane	4
IEEE 802.3ba 100G	10.3125 Gbps	Chip-to-module	10
IEEE 802.3ba 10GBASE-R	10.3125 Gbps	Chip-to-module	1 to *N*
IEEE 802.3ba 10GBASE-KR	10.3125 Gbps	Backplane	1 to *N*
10G GPON/EPON	10 Gbps	Chip-to-chip, chip-to-module	1 to *N*
OIF SFI-S	9.95 to 11.1 Gbps	Chip-to-module	(8, 10, 12, 16)+1
OIF SFI-5.2 (40G)	9.95 to 11.1 Gbps	Chip-to-module	5
10G Interlaken	10.6921 Gbps	Chip-to-chip, cable	1 to *N*
SONET/SDH OC-192 (10G)	9.95 Gbps	Chip-to-chip	1 to *N*
SONET/SDH OC-192 (40G)	9.95 Gbps	Chip-to-chip	4
SFP+	8.5 to 11.32 Gbps	Optical module std	1 to *N*
XFP	9.95328 to 1 1/32 Gbps	Optical module std	1 to *N*
OIF/CEI 11G-SR	9.95 to 11.1 Gbps	Chip-to-chip	I/O technology
OIF/CEI 11G-LR	9.95 to 11.1 Gbps	Backplane	I/O technology
OTU2	10.709 Gbps	Chip-to-chip	See SFI-S
OTU3	10.7545 Gbps	SFI-S	See SFI-S
OTU4	11.2 Gbps	SFI-S	See SFI-S
10G SDI	10.6921 Gbps	Chip-to-chip, cable	1 to *N*
10G InfiniBand	10 Gbps	Chip-to-chip, chip-to-module, cable, backplane	1 to *N*

Table 3. 28G High-Speed Standards/Protocols Supported by Stratix V FPGAs

Standard	Electrical Serial Lane Rate	Link	Lanes
OIF/CEI 28G-SR	19.9 to 28 Gbps	Chip-to-chip	1 to *N*
OIF/CEI 28G-VSR	19.9 to 28 Gbps	Chip-to-module	1 to *N*
IEEE 802.3ba 100G	25 Gbps	Chip-to-chip, chip-to-module, cable	4
32G FibreChannel	28 Gbps	Chip-to-module	1 to *N*
25G InfiniBand	25 Gbps	Chip-to-chip, chip-to-module, cable	1 to *N*

28 nm Transceiver Block Architecture

Altera's 28 nm transceivers offer an advanced architecture with scalability and flexibility through a continuous bank of up to 66 transceiver channels. A schematic overview of the block architecture is shown in Figure 1. Each channel consists of a transmit (TX) and receive (RX) pair with full physical media attachment (PMA) and physical coding sublayer (PCS). Each channel can operate both independently or in a bonded mode with a common clock source. In addition, the TX and RX pair for each channel can operate at independent data rates, further increasing flexibility and port density.

Figure 1. 28 nm Transceiver Block Architecture

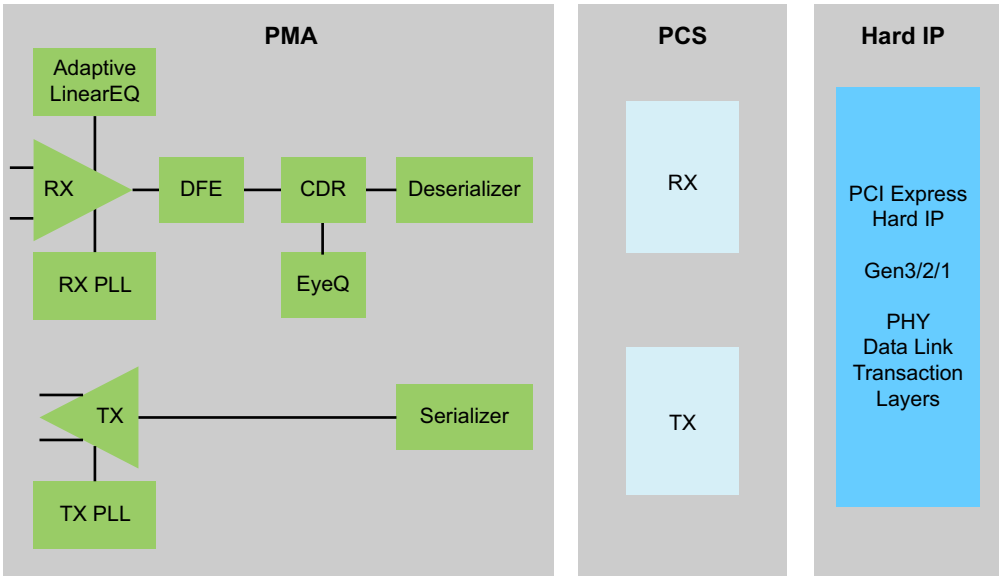

Advanced Clock and Timing Generation

The TX clock can come from three different sources:

- *PLL with LC oscillator (LC PLL)*—Provides the best jitter performance (sub-ps random jitter) with good tuning range
- *PLL with ring oscillator (CMU PLL)*—A ring oscillator (RO) derived from a clock multiplication unit (CMU) channel, which can also be used as a transceiver. Provides good jitter performance with the widest tuning range
- *Fractional PLL (fPLL)*—Offers both integer and fractional multipliers for a wide range of frequency outputs and can also be used for transmit clocking for transceivers to help reduce off-chip clock sources

Clocking and timing generation play important roles in high-speed transceivers. Jitter is an important metric used to determine the quality of a clock. As clock jitter affects both the transmitter and the receiver jitter performance, the BER for the link system is affected as well. On the TX side, the clock jitter degrades the eye opening at its output. The clock jitter on the RX side prevents the receiver from latching data correctly, reduces the setup and hold time of the data latch, and consumes a portion of the total available jitter budget of the link, leaving less jitter budget for the transmitter and channel.

All clock generation and distribution circuits in a transceiver produce a certain amount of jitter. The clock generator normally uses PLL circuits. The key component of a PLL is the oscillator, which is also the major source of jitter. Currently, the two main types of CMOS-based oscillators used in multiple-GHz PLLs are ROs and LC tanks (LCs), each of which has advantages and disadvantages.

The voltage control ring oscillator (VCRO) normally has a wide frequency tuning range, from 10 – 100 MHz to1 – 10 GHz. This feature enables transceivers to accommodate many different data rates, but it also results in very large gains. A high VCRO gain makes the PLL more sensitive to front-end noises and spurs. In addition, the RO can be sensitive to power supply and substrate noises. The VCRO phase noise or jitter is typically dominated by the supply noise injection if there is a lack of a high power supply rejection ratio (PSRR) voltage regulator. Good substrate isolation is also able to improve the phase noise or jitter of the RO.

System reliability is associated with the number of errors transmitted per bits transmitted, which is also known as the BER. A key component in maintaining system reliability across process, voltage and temperature variation (PVT) is to maintain the random jitter (non-bounded) component of a system. LCs offer a superior phase-noise performance due to its highly selective and high-Q LC tank. Discrete LCs have existed in RF applications for a long time, but high-Q LCs in mixed-signal integrated circuits are only a recent development. There are two main factors driving LCs to be used in integrated transceiver designs. The first is to reduce jitter in order to increase system reliability. The second is due to process feature size shrinkage that makes inductors small enough to be integrated on a die as the LC's frequency increases. In Stratix V FPGAs, the LC continues to offer a few hundred femtoseconds (RJ-rms) of transmit jitter, and a wider data rate range addressing transceiver applications from 3.25 Gbps to 14.1 Gbps and 20 Gbps to 28 Gbps.

Hybrid CDR Architecture

The innovative hybrid clock-data recovery (CDR) architecture of Stratix V FPGAs enhances the conventional data-driven phase interpolator-based CDR by allowing two operational modes, as well as digital controllability for its clock recovery circuit (CRC). The two operational modes are "lock to data" (LTD) and "lock to reference" (LTR), both of which can be used automatically or manually. A common usage model is to use the reference clock as the input, lock to the desired frequency, and then switch the input to the data signal to lock to the phase. This model allows both frequency and phase-aligned bit clocks to be recovered. In addition to retaining most of the advantages of a data-driven architecture, the hybrid architecture offers lower lock time, optimized power consumption, and greater tolerance of jitter and transition density. A schematic drawing for this architecture is shown in Figure 2.

Figure 2. Digitally-Assisted Hybrid Clock Architecture in Stratix V FPGAs

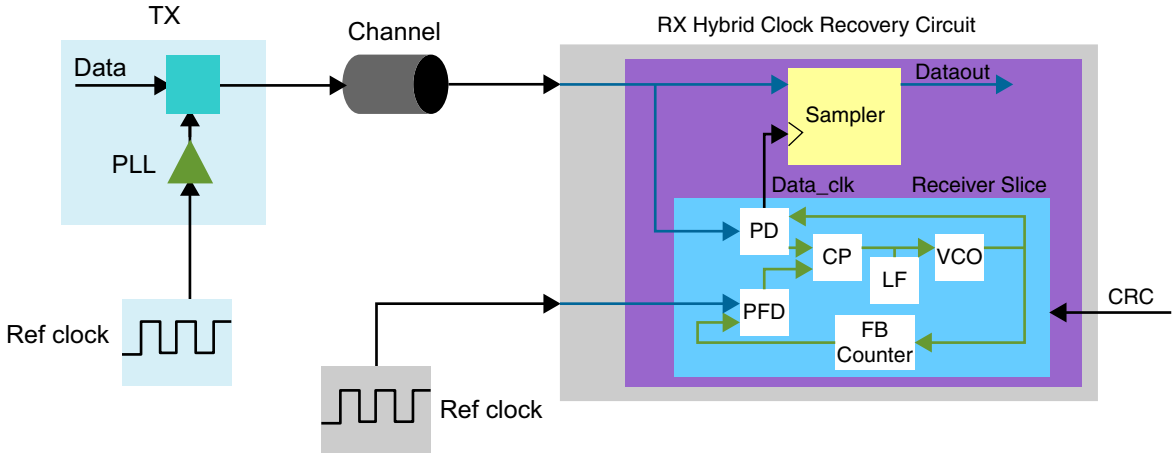

In this architecture, the CRC receives both data and reference clocks. The CRC itself includes two fully integrated loops, which are the phase frequency detector (PFD) (i.e., lock-to-reference) loop, and the phase detector (PD) (i.e., lock-to-data) loop. They share a common charge pump (CP), loop filter (LF), and voltage-controlled oscillator (VCO). The PFD loop is equivalent to an analog PLL structure, and is used to train the PLL to a desired frequency by using the correct feedback counter value. The PD loop is used to align the PLL output clock frequency with incoming data.

The CRC always starts up using the PFD loop (or LTR mode). Once the CDR reaches the desired output frequency, the CRC automatically switches from the PFD loop to the PD in LTD mode so that it can track incoming data and generate the recovered clock quickly since frequency is recovered and locked at this time. Within the CRC, there is a parts-per-million (ppm) detection circuit that constantly checks the difference between the recovered clock and reference clock. If the CRC clock "drifts" too far from the desired frequency (as in case of excessive spread spectrum), the ppm detection circuit switches the CRC from the PD loop to the PFD loop so that the CDR output frequency is maintained.

The hybrid clocking architecture of Stratix V FPGAs has advanced capabilities in filtering and reducing jitter from the transmitter and reference clocks. This capability allows the best possible BER performance for the link system. In the hybrid clocking architecture, the reference clock can come from the same reference clock as the transmitter, or it can be a separate clock with relaxed ppm requirements. As the RX

reference clock is only used in the initial lock-to-clock phase, its jitter will not impact the receiver CDR. This feature offers the user flexibility in providing a reference clock source without any performance degradation. In the hybrid clocking architecture, the reference clock jitter does not contribute to the system BER at all, providing significant jitter margin for designs.

In contrast, a phase interpolator, such as that used in competing FPGAs, uses the reference clock directly for both clock and data recovery. Therefore, the reference clock jitter in this architecture does contribute to the system BER, consuming system jitter budget. For the conventional data-driven architecture, the CRC either may never recover the bit clock or requires a long time to lock incoming data when the receiver data input has excessive jitter. The hybrid architecture is unaffected by such problems as its frequency-locked bit clock is already in place when recovering data.

28 nm Transceiver Power Efficiency

One of the key advantages of moving to higher data rates is the power efficiency achieved per Gbps. At 12.5 Gbps, Stratix V transceivers consume less than 170 mW per channel (full duplex), while the 28 Gbps transceiver consumes around 200 mW. Normalized, this power efficiency is about 7 mW/Gbps at 28 Gbps (Figure 3).

Figure 3. 28 nm Stratix V Transceiver Power

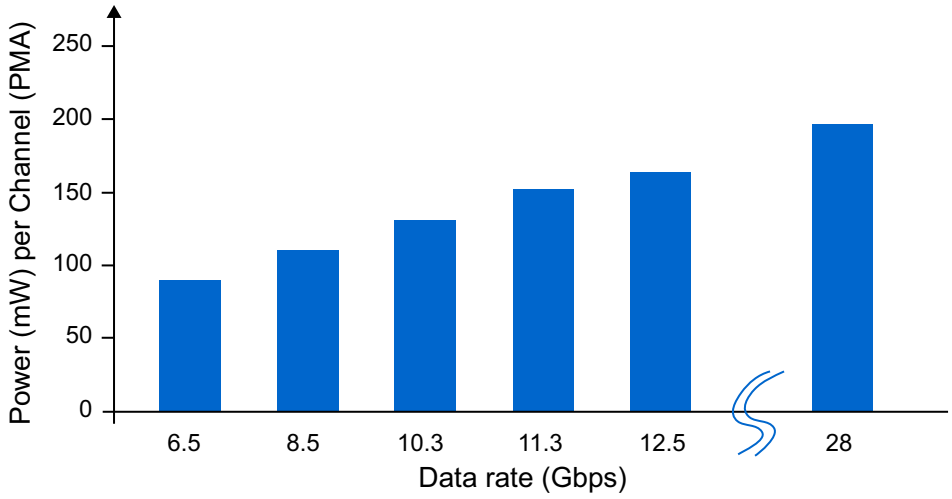

The 28 nm transceiver categorizes the power supply voltage options into three basic configurations (Figure 4) to address low-power and high-performance needs:

- *Low power (0.85 V)*—This configuration is used for data rates of equal to or less than 6 Gbps in short reach, chip-to-chip, and chip-to-module applications with simple linear equalization techniques, such as transmit pre-emphasis/de-emphasis and continuous time linear equalizer (CLTE).
- *High performance (1.0 V)*—This configuration is used for noisy and lossy channels, such as long-reach and backplane applications at data rates of more than 6 Gbps, with more advanced equalization circuit blocks, such as decision feedback equalization (DFE) and the automatic dispersion compensation engine (ADCE), in addition to simple linear equalization circuits.
- *Off (0 V)*—The option to turn off the transceiver saves power on unused channels.

Figure 4. Three Power Supply Voltage Options

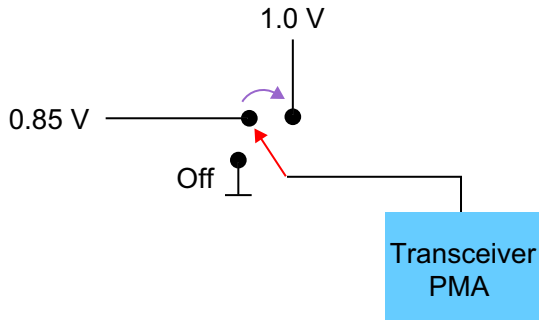

One architectural advantage that provides Altera's power advantage is the H-tree output driver. This architecture uses half the current of the more power-hungry CML driver. Figure 5 shows the diagram for the H-tree link versus a CML link. The current for the H-tree circuit flows from the positive arm into the negative arm across the load. The CML output is generated by only one arm, while the other arm is tri-stated. Thus, twice the amount of current is required in order to generate the same amount of swing.

Figure 5. H-Tree Link and CML Link

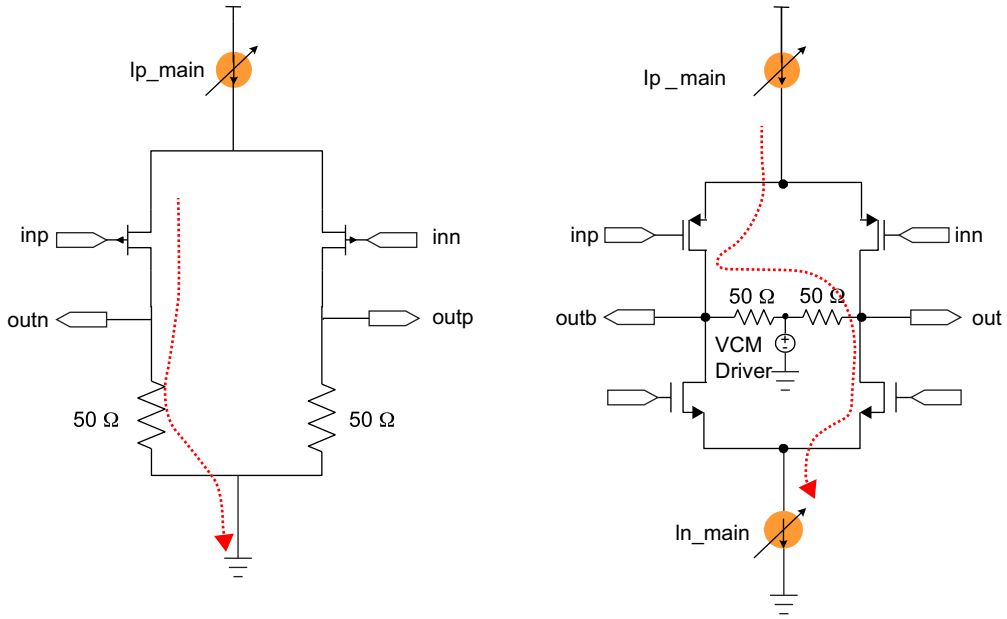

CML Driver **H Tree Driver**

End-to-End Equalization

When communication is done through backplanes and lossy channels at higher data rates, the increase in attenuation, reflections, and coupling causes limitations in bandwidth on the medium in which the signals travel. Figure 6 shows a diagram of a typical backplane used by the TX device to transmit a signal to the RX device. As the signal leaves the TX driver, it must go through the TX I/O card, the I/O card connector, the backplane traces, the RX I/O card connector, and finally to the RX I/O card. The bandwidth limitations of this system result from the I/O cards, connectors,

and the backplane itself, which can be over 40 inches of FR-4 material. The skin effect, the dielectric loss in the medium, and the reflections caused by the vias significantly distort the signal going from TX to RX. To allow the applications of 10 to 12.5 Gbps backplanes, standards have evolved to overcome the limitations of lossy backplanes, namely the IEEE 802.3ap 10GBASE-KR and CEI-11G-LR.

Figure 6. Typical Backplane Communication Between Two Chips

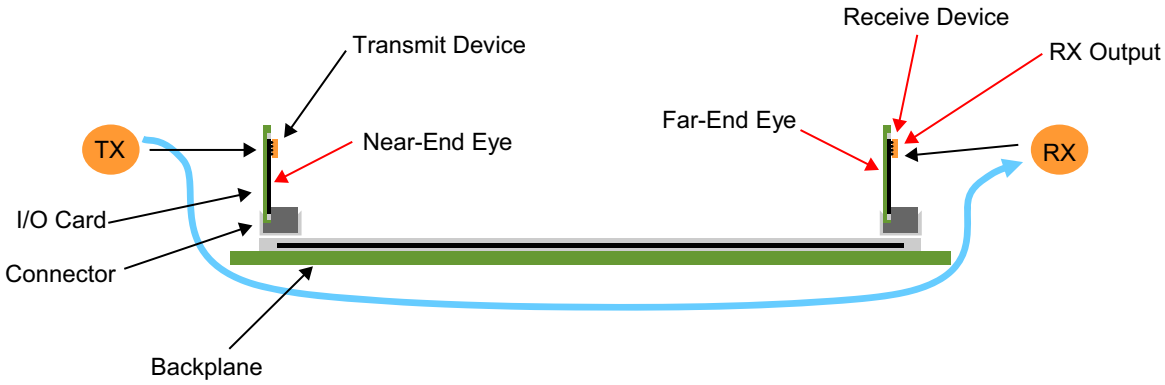

Assuming the 10GBASE-KR backplane is used to communicate between TX and RX, Figure 7 illustrates various effects the TX signal experiences. Initially, the differential signal observed at point A, the near end of the link, is of good quality. As the signal propagates through the I/O cards, connectors, and backplanes, it is distorted by attenuation, reflections, radiation, and coupling. At point B, the far end of the link, the signal has become degraded. In extreme cases, the signal may be so attenuated that the two differential signals do not even cross. This extreme attenuation is most likely to occur in long backplanes, where the incident signal degrades further. Higher data rates in the same backplane will experience an even higher amount of degradation. Yet, these highly distorted signals must still be processed somehow, usually using signal conditioning or compensation/equalization.

Figure 7. Transmitted Signal Affected by a Transmission Medium

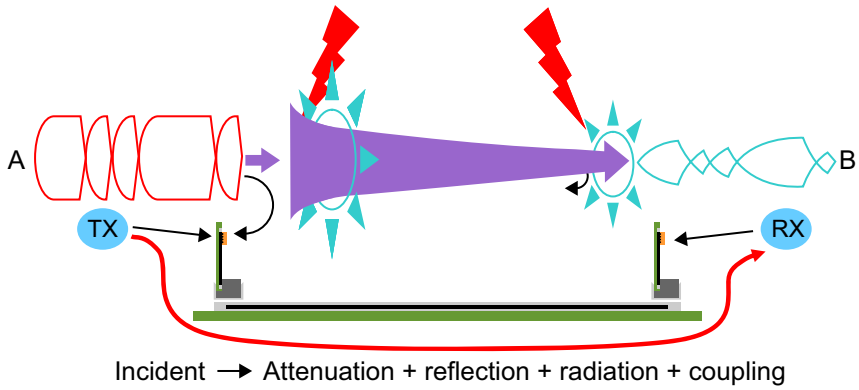

Signal Conditioning

There are many types of signal conditioning, or compensation/equalization, each of which has advantages and disadvantages. This section reviews some of the implementations available in Stratix V FPGAs.

Transmit Pre-Emphasis/De-Emphasis

Transmit pre-emphasis/de-emphasis is implemented at the TX driver by pre-conditioning the signal before it is launched into the channel. The purpose is to amplify (pre-emphasize) the high-frequency contents of the signal and to reduce the low-frequency contents of the signal. The benefits of this method are relative simplicity and low power. All sampled data is readily available at transmitting devices. Delayed versions of transmitted data are easily created to be one unit interval (UI) apart by placing a bank of registers to hold both prior and upcoming serial data bits. Similarly, fractional sampled data (1/2 UI) is easily available at the TX driver from the intermediate latch stages that are commonly used to create registers. This signal-conditioning technique addresses both pre-cursor and post-cursor intersymbol interference (ISI) as both prior and upcoming bits are available.

CTLE

A CTLE is implemented on the RX side. This linear equalizer is well suited here as non-sampled or continuous time implementation suffices. As a result, CTLE-based signal conditioning is usually the best choice for low power. Similar to transmit pre-emphasis, a CTLE addresses pre-cursor and post-cursor ISI, but in continuous time instead of being limited to a pre-set number of TX taps.

Figure 8 shows an example of a first-order CTLE transfer function. A zero value is inserted to compensate the poles in the channel transfer function. This implementation is simple and consumes low power. More equalizer stages can be added as given links are required for a selected data rate. Multiple equalizer stages not only increase the order of the resulting equalizer, but also increase the maximum boost achieved in a given frequency interval. However, it is important to note that parasitic poles and their locations should be carefully considered when designing a CTLE.

Figure 8. Example of a CTLE Transfer Function

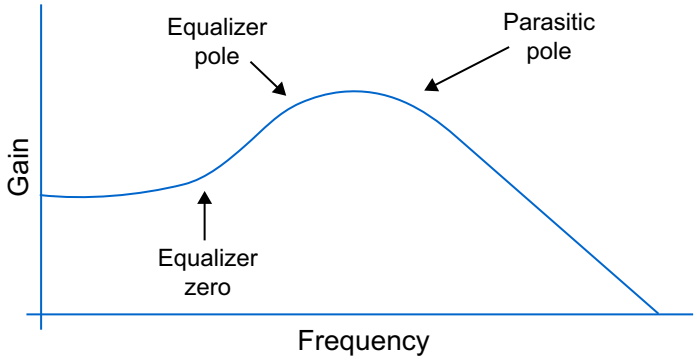

For the same near-end TX signal, the far end shows an open eye when an internal CTLE is enabled. The amount of crosstalk does not increase in a system with a CTLE versus a transmit pre-emphasis/de-emphasis system due to the reduced high-frequency content at the TX driver. Lastly, it is important to note that a system that has a CTLE is well suited for real-time adaptation, as signal information after the channel is readily available for processing and reconditioning at RX. In Stratix V FPGAs, the CTLE bandwidth is programmable and covers frequency components up to 6.25 GHz (corresponding to data rates of 12.5 Gbps).

DFE

Unlike prior equalization schemes, a DFE is a non-linear system, shown in Figure 9. DFE not only requires the sampling of data at the correct instant in time within a unit interval (UI), but also computing the DFE tap weight prior to the next sample. This complexity makes timing closure a challenge. A DFE requires a proper data sample, leading to the design complexity of the combined equalization and recovery sections of the transceiver.

Figure 9. DFE Scheme Diagram

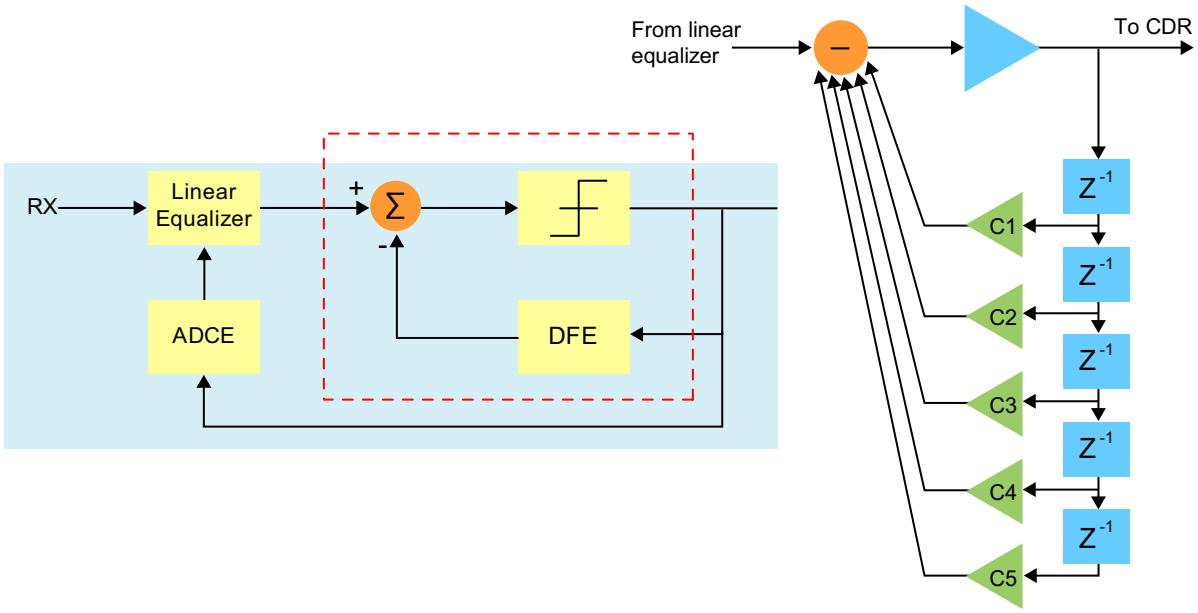

The major benefit of a DFE is improved handling of crosstalk, especially when it is assumed to be additive white Gaussian noise (AWGN). The benefit of the DFE in an AWGN system lies in the signal-to-noise ratio (SNR). However, for correlated crosstalk, a DFE is not as effective as a CTLE. The system has mostly post-cursor ISI as the DFE does not address pre-cursor ISI. The SNR is computed as the ratio of signal power to noise power. As the CTLE is continuous in time and does not really "know" or "need to know" the spectral density of incoming signals, it boosts both signal and noise by an equal amount. In effect, the CTLE preserves the SNR of the original link.

The assumed ability of the DFE engine to compute coefficients for incoming bits based on the history of received bits makes it attractive. In fact, for data rates above 3 Gbps, the adaptive equalization feature is worth looking out for in transceiver vendor offerings. However, it is important to note that there is a clear difference between generic DFE and adaptive equalizer systems, such as electronic dispersion compensation (EDC) and adaptive dispersion compensation engine (ADCE), and "programmable" DFE systems. The key difference is a real-time adaptation. Programmable DFE systems presume prior knowledge of each given channel that must be programmed into the DFE coefficients. All other systems, like the ADCE, compute coefficients in real time without any prior knowledge of channel or data patterns and can address a large number of programmable signal conditioning settings. Stratix V FPGAs offer a 5 tap DFE with auto-adaptation to address a wide range of standards for backplane applications including CEI-11G long reach.

Power Integrity, Jitter, and BER

Although supporting increasing data rates is the key to address emerging standards and applications, it is important to assure that power supply noise does not affect the performance of the circuitry. In Stratix V FPGA transceivers, the high-speed analog sections of the transmit and receive paths are separated because the transceiver allows completely independent frequency selections on the RX path from that chosen for the TX path. Separated power prevents the injection of uncorrelated noise sources. Precision analog blocks, such as bang gap, current bias, and on-chip voltage regulators, receive power from one dedicated power supply. On-chip voltage regulators are added to isolate the sensitive PLL circuitry of the TX and RX, such as VCO, CP, and LF shown in Figure 10. The TX driver has its own power supply, VccEH, which provides different power levels. Both on-chip and on-package decoupling is used in the transceivers to provide necessary noise filtering from external power supplies while reducing the number of decoupling capacitors needed on board.

Figure 10. Regulator for VCO/CP/LF in the TX PLL and CDR

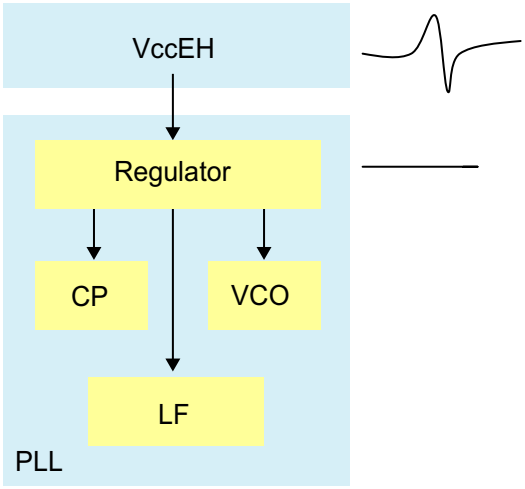

Jitter and Noise Generation

A good transmitter is expected to produce a minimum amount of jitter and noise, which means that the eye diagram measured at the transmitter output should be wide open. Many high-speed I/O standards, such as PCIe, CEI/OIF, and FC, define an eye mask to determine whether the jitter, noise, and signaling of the transmitter output comply with the requirement. The eye mask should typically correspond to a BER of 10^{-12} or lower. Unretimed optical interfaces such as SFP+ have particularly stringent total jitter (Tj) requirements. A component spec of 0.28 UI at a BER of 10^{-12} at 10.3125 Gbps is often required. Figure 11 shows the jitter distribution over PVT for Stratix V FPGAs, with a significant margin to the SFP+ requirement allowing for robust and reliable system operation.

Figure 11. Stratix V Jitter Distribution at 10.3125 Gbps for SFP+ Requirements

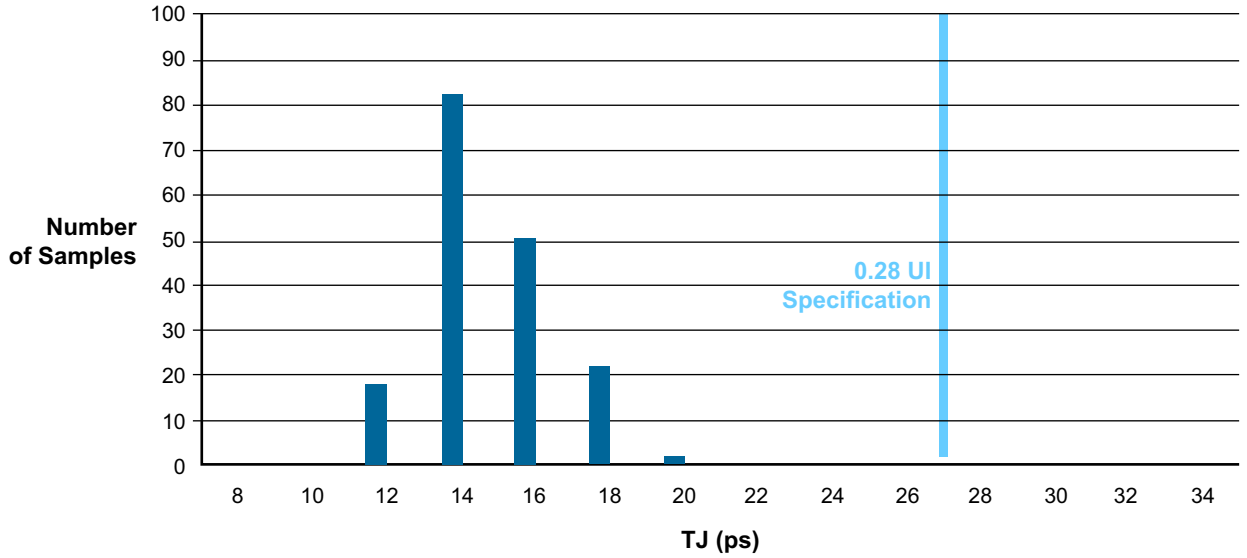

On the other hand, a good receiver should be able to tolerate a minimum amount of jitter and noise. The two important subsystems to test for a receiver are the CRC and the equalization subsystems, such as CLTE and DFE. To verify whether a receiver CRC has the required jitter-tracking capability, a jitter frequency mask or jitter tolerance mask is often defined by protocol standards. A jitter tolerance mask curve is the reciprocal of the jitter transfer function of the receiver. A receiver with a good CRC tolerates more jitter than is required by a standard. Figure 12 shows the transceiver jitter tolerance performance for Stratix V FPGAs at 28 nm FPGA transceiver. Due to its enhanced CRC, CTLE, and DFE capabilities, jitter and noise tolerance is further advanced for the RX of 28 nm Stratix V FPGAs. From these results, the jitter tracking capability of the hybrid CDR architecture should be evident from the wide jitter tolerance margin.

Figure 12. CEI-11G Receiver Jitter Tolerance Results for Stratix V FPGAs

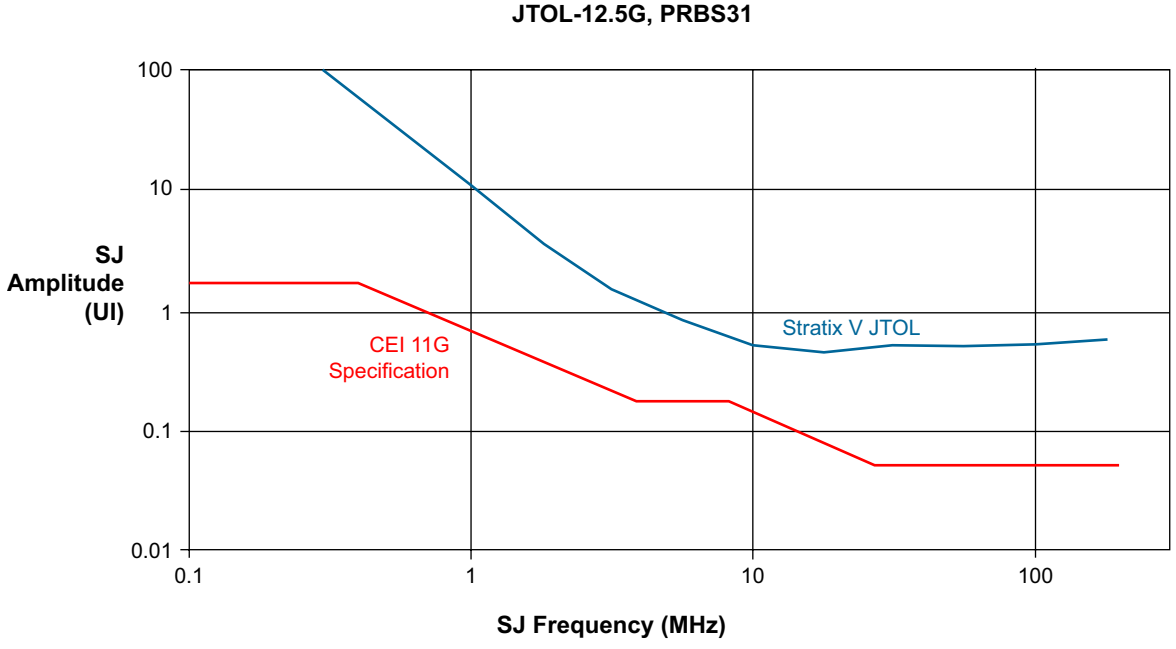

System Performance and BER

The BER is a system metric for a link, and thus is tightly coupled with the link system architecture and subsystem performances. A quantitative and accurate system model and understanding is essential for estimating or measuring the BER of the system. The BER for a subsystem may also be defined if the rest of the characteristics of the link system are defined.

The BER can be caused by jitter, noise, or both. By definition, the BER is the sum of integrations of PDFs of jitter and noise, manifesting two-dimensional characteristics for such performance metrics. Accordingly, the BER is a cumulative distribution function (CDF). The BER is often estimated and viewed as a function of sampling time at a given reference voltage, such as at zero-crossing or at 50 percent swing levels, or as a function of sampling voltage at a given reference time location, such as at the center of the UI data cell. Figure 13 shows an eye diagram with the jitter PDF at the zero-crossing level and the BER as a function of sampling time, often called a bathtub curve.

Figure 13. Integrated and Correlated View for the Eye Diagram, Jitter PDF, and BER CDF

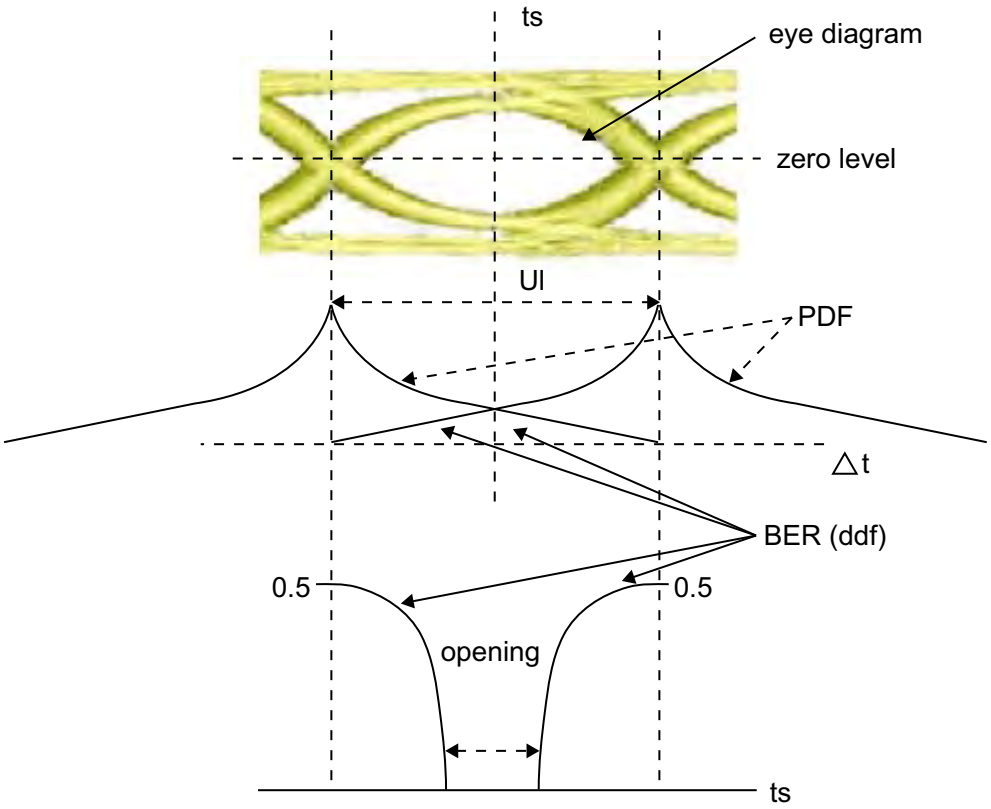

The BER of the overall link system depends on the jitter, noise, and signaling performances of its subsystems, such as the transmitter, channel, and receiver. Altera's innovative design enables the transceiver to meet and even exceed the requirements posed by various high-speed I/O standards. Due to the excess margins in jitter and noise, it also enables users to design systems with lower costs and better BER system performance.

28 Gbps Advanced Transceiver Technology

Closely following Moore's law, the data rate per lane doubles approximately every two years. For a link system, such as 100G Ethernet and OTN, the number of pins on the system-on-chip (SoC) device and the width of the link are reduced by half when the data rate per lane doubles. For example, the next-generation signal rate for 100G Ethernet will be 25 Gbps, an approximately 2.5× increase compared with the current 10.3125 Gbps lane signaling rate. The number of transceiver pins and physical channels of the PCB board will be reduced from 10 to 4, thus reducing the width of the link. Due to power and form-factor constraints, the 25 Gbps chip-to-module interface will likely be asymmetrical. CDR and equalization circuit blocks are likely to be removed from the optical module, with their functionality delivered by host SoC devices such as 25 to 28 Gbps FPGA transceivers.

At 28 Gbps, the targeted channel loss can be 15 dB for short reach and 25 dB for long reach. To cover both short-reach and long-reach applications, TX equalization alone will not be enough. To compensate for the channel insertion loss as well as uncorrelated distortions such as crosstalk and reflection, both TX and RX equalizations are needed. Figure 14 shows the 28 Gbps TX architecture of the Stratix V FPGA.

Figure 14. 28 Gbps TX Architecture in Stratix V FPGAs

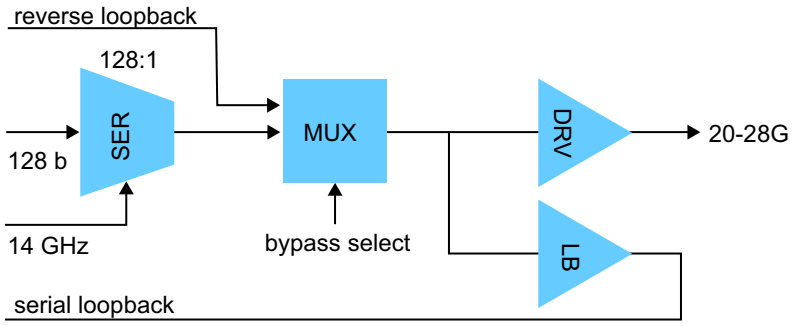

The clocking scheme for the 28 Gbps serializer uses a similar architecture as the one used in the 14.1 Gbps channel for 28 nm Stratix V FPGAs. A shared LC PLL is used in Stratix V FPGAs to generate a raw 14 GHz clock.

Sub-phases are generated to drive the 128 b to 1 b data for the TX driver. The integrity and load balancing of the clock is maintained through carefully designed buffers and dividers. Layout matching and delays are critical for the final two-to-one multiplexer clock and data timing. Figure 15 shows the Stratix V FPGA's transceiver clocking architecture.

Figure 15. 28 Gbps Transceiver Clocking in Stratix V FPGAs

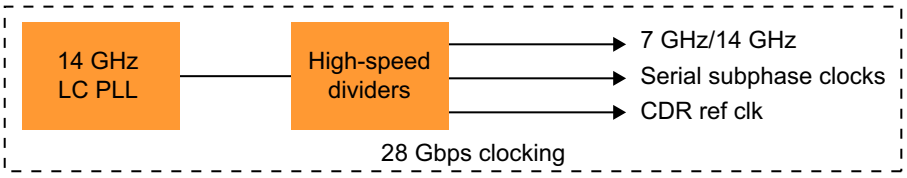

The TX driver uses a similar architecture as the 12 Gbps channel. The design is simplified to reduce the number of supported output swings to conform to CEI-25/28 specifications. The output driver structure is an H-bridge design that offers the advantage of flexible output common mode and lower power consumption compared to CML-type designs. The pre-emphasis/de-emphasis features are easily integrated into the TX driver, which is important for longer trace lengths and optical module interconnects. To achieve low jitter TX characteristics, output capacitance reduction is important. Deterministic jitter is also a challenge at these high data rates due to settling issues in current sources, which are mitigated by various design techniques.

Stratix V FPGAs use a hybrid clock-recovery architecture, a combination of an analog PLL-based CRC and a phase-interpolator-based CRC for both 14.1 Gbps and 28 Gbps receiver data rates. Similar jitter tolerance and tracking, run-length and transition-density tolerance, and lock time are expected. Figure 16 shows the Stratix V FPGA's 28 Gbps RX architecture.

Figure 16. 28 Gbps RX Architecture in Stratix V FPGAs

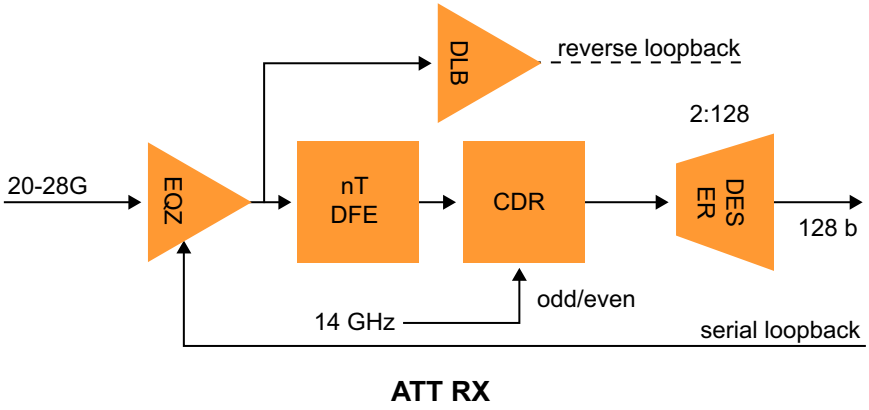

On-Die Instrumentation (ODI)

With the integration of equalization at the receiver, it is challenging to understand the true eye opening before the signal is latched in the logic domain. Altera's on-die instrumentation tool, EyeQ, provides in-situ signal and jitter generation and measurement capability at various signal nodes within the transceivers. This information cannot be easily generated or measured by external equipment, but it is very important for the TX and RX equalization and clock-recovery diagnostics and debug as well as link characterization and verification. The functionality and performance of EyeQ are quite similar to those of an external instrument or a tester, but at no additional cost. In addition, it naturally solves the high-speed interface problems faced by external instruments as the ODI resides inside the transceiver and accesses the internal signal nodes directly. Figure 17 shows the Stratix V FPGA's ODI architecture.

Figure 17. ODI: The EyeQ Architecture

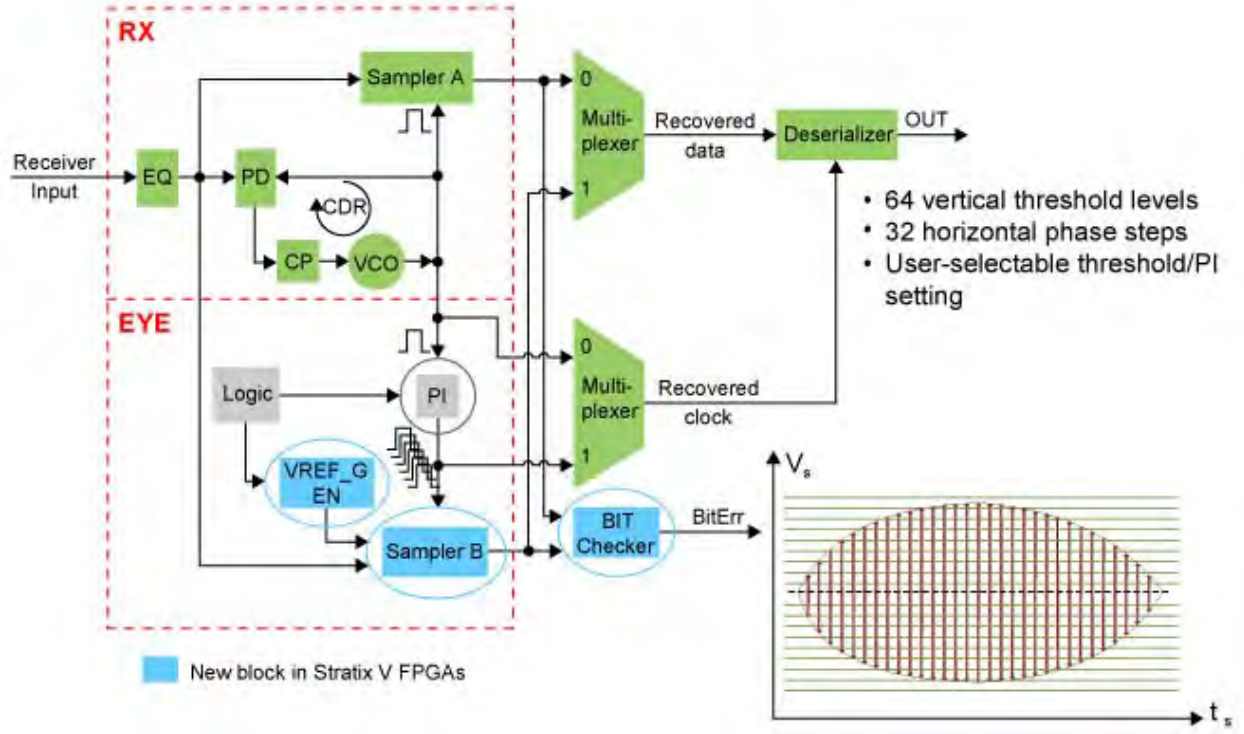

Notes:

(1) EQ: equalizer, PI: phase interpolator

An important capability of the EyeQ is to measure and show the effective eye-width of the equalized RX data. Not only can this information be used to monitor RX signal conditions, but also to adjust the equalization amount and settings, a task that cannot be achieved with external equipment. The EyeQ is also useful for system-level debugging in field and traffic watching as it does not require any special or fixed data patterns. At 28 nm, the Stratix V FPGA's ODI goes beyond its predecessor by providing both timing and voltage sampling up to 28 Gbps, enabling complete eye-diagram reconstruction with EyeQ. Furthermore, ODI also provides serial BER measurement capability before the deserializer to help solve most test, verification, characterization and debug challenges for high-speed transceivers through a non-intrusive interface.

Conclusion

To address the growing technology trends and challenges of high-speed serial links and transceivers of the next generation, a careful balance in transceiver designs is needed. Stringent performance requirements for jitter, noise, power integrity, and BER must be met, all while minimizing power consumption and increasing design productivity. Altera's Stratix V FPGAs meet and even exceed those requirements. The 28 nm process technology allows Stratix V FPGAs to achieve the best possible logic density, memory speed, and capacity. In addition, transceiver innovations enable superior jitter, noise, signal integrity, and BER performances with the highest data rates of 28 Gbps in production today, along with 12.5 Gbps backplane capability when optimized for power.

In addition to the Quartus® II development software, Altera's unique system tools, including PELE and HST, provide fast and accurate link design simulations. These simulations enable a quick and reliable design of Stratix V FPGA transceivers as well as transceiver jitter margin for a cost-optimized and high-performance link channel design. Stratix V FPGAs, enhanced with 28 Gbps transceivers, are the best in class in terms of speed, performance, power consumption, and built-in-measurement capability, with a state-of-art design in the 28 nm process node that supports a wide range of high-speed standards.

Further Information

- Stratix V FPGAs: Built for Bandwidth: www.altera.com/products/devices/stratix-fpgas/stratix-v/stxv-index.jsp
- Stratix V FPGA Transceivers: www.altera.com/products/devices/stratix-fpgas/stratix-v/transceivers/stxv-transceivers.html
- Webcast: "Extending Transceiver Leadership at 28 nm": www.altera.com/education/webcasts/all/wc-2010-transceiver-leadership-28nm.html
- Documentation: Stratix V Devices: www.altera.com/literature/lit-stratix-v.jsp
- M. Li, A. Martwick, G. Talbot, J. Wilstrup, "Transfer Functions for the Reference Clock Jitter in a Serial Link: Theory and Applications," ITC/IEEE, 2004.
- M. Li, "Jitter and Signaling Test for High-Speed Links," an Invited Paper, CICC/IEEE, 2006.
- B. Razavi (Editor), *Monolithic Phase-Locked Loops and Clock Recovery Circuits: Theory and Design*, Wiley-IEEE Press, 1996.
- V. Stojanovic, M. Horowitz, "Modeling and Analysis of High-Speed Links," Stanford University lecture.
- B. Casper, P. Pupalaikis, J. Zerbe, "Serial Data Equalization," DesignCon 2007.
- W. Wong et al., "Digitally Assisted Adaptive Equalizers in 90-nm With Wide-Range Support From 2.5 Gbps to 6.5 Gbps," DesignCon, 2007.
- M. Li, *Jitter, Noise, and Signal Integrity at High-Speed*, Prentice Hall, ISBN 0132429616, 2007.
- M. Li, "Statistical and System Approaches for Jitter, Noise, and Bit Error Rate (BER) Tests for High-Speed Serial Links and Devices," ITC/IEEE, 2005.

Acknowledgements

- Weichi Ding, Design Manager, Analog Design Group, Altera Corporation
- Tim Hoang, IC Design Manager, Analog Design Group, Altera Corporation
- Salman Jiva, Product Marketing Manager, Altera Corporation
- Mike Peng Li, Ph.D., Principle Architect/Distinguished Engineer, Product Engineering, Altera Corporation

- Sriram Narayan, Principle Design Engineer, Analog Design Group, Altera Corporation
- Sergey Shumareyev, Director of Engineering, Analog Design Group, Altera Corporation
- Wilson Wong, Principle Design Engineer, Analog Design Group, Altera Corporation
- Lux Joshi, Product Marketing Manager, Altera Corporation

Document Revision History

Table 4 shows the revision history for this document.

Table 4. Document Revision History

Date	Version	Changes
October 2012	2.1	■ Updated Figure 1.
October 2012	2.0	■ Text edits and reformatting. ■ Added Figure 1, Figure 5, Figure 11, and Figure 12. ■ Added "Hybrid CDR Architecture".
May 2010	1.0	Initial release.