

Jitter-, Signal Integrity-, Power-, and Process-Optimized Transceivers

1. Introduction	2	6. End-to-End Equalization	15
		6.1 Motivation	15
2. Trends and Requirements for High-Speed Links	3	6.2 Equalization in a Nutshell	17
2.1 Technology Trends and Challenges	3	6.3 Types of Signal Conditioning	18
2.2 I/O Protocol Standards Supported	4	6.3.1 Transmit Pre-Emphasis/De-Emphasis	18
		6.3.2 Continuous Time Linear Equalizer	20
3. 40-nm Process Node and Transceiver	6	6.3.3 Transversal FIR (TFIR)	21
3.1 Process Technology Roadmap	6	6.3.4 Decision Feedback Equalization	21
3.2 Analog Challenges vs. Process-Node Shrinkage	6	6.4 Backplanes and Their Impact on Equalization Se-	
3.2.1 Headroom	6	lection	22
3.2.2 Gain	7	6.4.1 Backplanes and Equalization Selection	23
3.2.3 Leakage	7	6.4.2 The Need for Adaptive Equalization	26
3.2.4 DSM Effects and Models	7	6.4.3 System Development Stages	27
3.2.5 Mismatch	7	6.5 A Historical Overview of Altera's Equalization	
3.3 Solutions and Trends	7	Solutions	28
3.3.1 High Supply Voltage for Analog Blocks	8	6.6 Built-In Oscilloscope	28
3.3.2 NML and Asymmetric Devices	8		
3.3.3 Analog Devices for FPGAs	8	7. Advanced Clock and Timing Generation	29
4. Architectures	9	7.1 Ring Oscillator	29
4.1 Data-Driven Architecture	9	7.2 LC Oscillator	29
4.2 Common Clock Architecture	9	8. Power and Jitter	31
4.3 Altera's Hybrid Architecture	10	8.1 Power Integrity	31
5. Mixed Signal Clock Recovery	12	8.2 Jitter and Noise Basics	33
5.1 Hybrid (Digitally Assisted) Clock Recovery	12	8.3 Jitter and Noise Generation	34
5.2 Capability Comparisons	12	8.4 Jitter and Noise Tolerance	35
5.2.1 Jitter Tolerance	12	8.5 Overall System BER	36
5.2.2 Flexibility and Versatility	13		
5.2.3 Transition Density and Maximum Run		9. Conclusion	38
Length (MRL)	14		
5.2.4 Lock Time	14	References	39
		Further Information	39
		Acknowledgements	39

1. Introduction

Technology advancement for the semiconductor industry is largely driven by Moore's Law in that the number of transistors in an integrated circuit doubles roughly every two years, as well as demanding a higher data rate for communication links between devices or systems. Moore's Law is facilitated by the feature size or process node shrinkage. Smaller features enable more functionalities, higher operation speed, logic density, integration, and lower power consumption per logic function. A higher data rate is often achieved through using advanced design methodologies and process technologies, and enables wireline and wireless communication, computer, storage, military and broadcast electronic systems to send and receive large amounts of data to meet the ever-increasing data transfer or bandwidth demand.

The 65-nm process technology is found in leading-edge products such as microprocessors and FPGAs. The next generation of these products will use the 45-nm or 40-nm process, available this year. The smaller feature size implies a smaller channel length for a transistor and shorter interconnects for a logic gate, resulting in faster switch time and shorter interconnect transport delay. The results of this process node shrinkage are favorable for logical operation, high density, and high-speed data transmission, as they are optimized for power consumption efficiency.

Today, the data rates for most of the advanced transceivers are in the range of 5–6 Gbps for communication and input/output (I/O) standards. A few examples of this lengthy list include CEI/OIF 6G, 2X XAUI (6.25 Gbps) for network communication, PCIe 2.0 (5 Gbps) for computer I/O buses, and SATA III/SAS II (6 Gbps) for storage area networks.

Altera® Stratix® IV GX FPGAs are based on 40-nm technology. The core logic fabric has up to 570K logic elements (LEs), enabling large SOC (system on chip) FPGA designs and applications. The high-speed transceiver has a quad topology with an up-to-48-channel capacity, which supports lane data rates up to 8.5 G Gbps. In summary, Altera's Stratix IV GX FPGAs deliver the highest density, the highest performance, and the lowest power. Leveraging 40-nm benefits, a proven transceiver, and memory interface technology, Stratix IV GX FPGAs provide an unprecedented level of system bandwidth with superior signal integrity. Stratix IV GX FPGAs, when coupled with HardCopy® IV ASICs, provide the benefits of FPGAs and ASICs using seamless prototyping.

This paper provides technical details of the Stratix IV GX FPGAs' performance, capabilities, and targeted applications. "2. Trends and Requirements for High-Speed Links" covers the technology, market, and application trends, as well as requirements for high-speed transceivers including the new high-speed I/O interface standards (PCI Express Generation 2 (PCIe 2.0), Hyper Transport 3.0 (HT 3.0), Interlaken, Common Public Radio Interface (CPRI), and SERDES Frame Interface Level 5 (SFI-5)). "3. 40-nm Process Node and Transceiver" covers Stratix IV GX FPGAs' capabilities and performance, and how they meet or exceed the technology and standards requirements. "4. Architectures" covers important and unique capabilities, technology superiority, and performance metrics, including the high-speed link and transceiver process node, and architectures. "5. Mixed Signal Clock Recovery" discusses the clock recovery circuit (CRC), and "6. End-to-End Equalization" covers both transmitter and receiver equalizations. "7. Advanced Clock and Timing Generation" covers various types of oscillators. "8. Power and Jitter" discusses built-in-self-test-jitter (BIST), noise, signal integrity, and bit error rate (BER), power management and power integrity, precision timing generation circuits, and supportability of high-speed standards. "9. Conclusion" summarizes and concludes the paper.

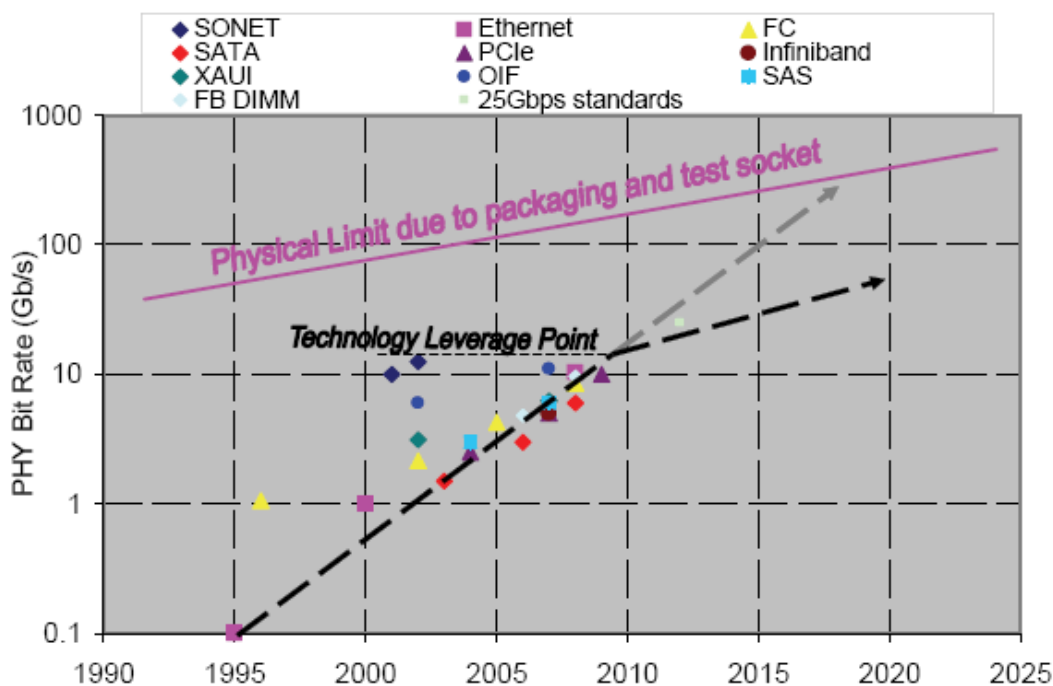
2. Trends and Requirements for High-Speed Links

This section covers the general technology trends and requirements for high-speed I/O links, as well as focusing on the corresponding markets and applications.

2.1 Technology Trends and Challenges

As shown in Figure 1, the International Technology Roadmap (ITRS) states that both the single-lane data rate and the pin counts for I/Os in system-on-chip (SOC)-type integrated devices have increased since 2000, due to the ever-increasing demand for more bandwidth and higher data rates. According to the ITRS 2007 revision, most high-speed I/O interfaces data rates double every two to three years, and are 5–6 Gbps in 2008. The predictions accurately reflect the reality of high-speed I/O devices available on the market and historical data. These multiple-Gbps, high-speed I/O interfaces are used for chip-to-chip, board-to-board, and system-to-system links, for data communication and telecommunication networks (e.g., Gigabit Ethernet (GbE), Optical Internetworking Forum (OIF) (e.g., SFI, Interlaken), computing I/Os (e.g., PCIe), storage area networking (e.g., Serial Advanced Technology Attachment (SATA), Fibre Channel (FC)) wireless networking (e.g., CPRI), and embedded processing (e.g., Serial RapidIO® (SRIO)). The next-generation I/O devices for most high-speed standards are expected to maintain similar data-rate growth to that of previous generations⁽¹⁾.

Figure 1. Mean Data Rate for High-Speed I/O Links (1)



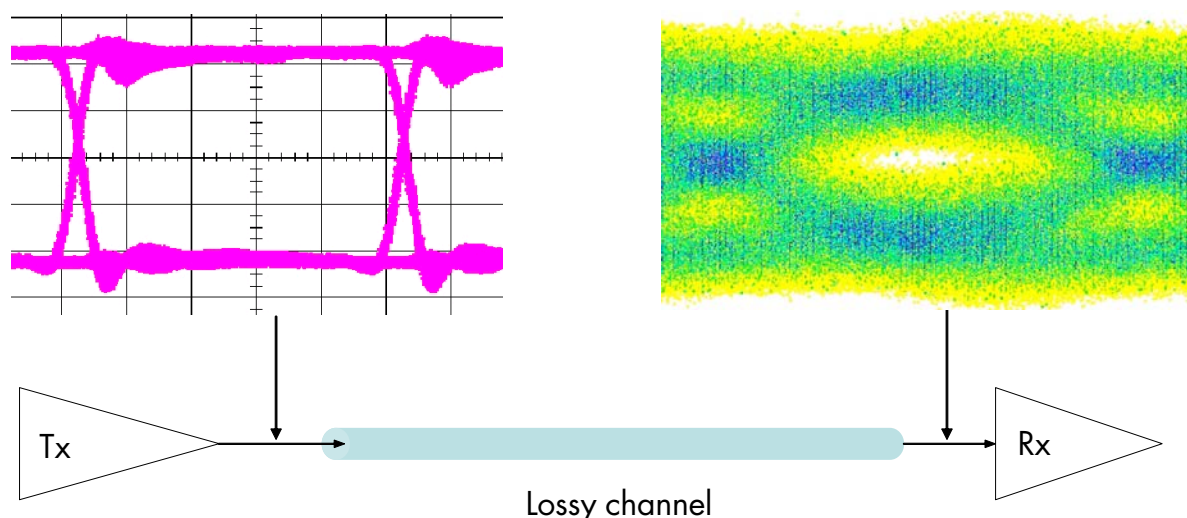
Note:

(1) From ITRS 2007.

Most multiple-Gbps link transceivers use a serializer/deserializer (SERDES) to format/deformat and transfer/receive data bits. At the transmitter side, many parallel and low-speed (commonly at a few hundred Mbps) coded data lines are multiplexed to a single lane at a much higher data rate (several Gbps). At the receiver side, data is received and processed in a reversal of the transmitted data. The received single-lane data first is recovered and then de-multiplexed to many parallel data lines, re-assembling the original data bits sent to the transmitter. Because of the much higher single-lane data rate, the high-speed SERDES enables much fewer I/O pins, channels, and associated channel materials, compared with low-speed synchronous parallel I/Os. Furthermore, the lane-to-lane skew is no longer an issue for SERDES since each lane is self-timed and asynchronous to each other. The self-timing is achieved by using a CRC that not only recovers a clock, but also tracks low-frequency jitter.

As the data rate keeps increasing, the unit interval (UI) for the data bit decreases, due to their reciprocal relationship. UI is the total timing budget available for the link, consisting of a minimum three components budget corresponding to transmitter, receiver, and channel subsystems. Most multiple-Gbps standards require a system BER be at 10^{-12} or lower. If the UI decreases and the BER is unchanged, the jitter generated by the subsystems must be reduced accordingly. Meanwhile, for cost consideration, printed circuit board (PCB) and backplane materials likely are kept unchanged when data rate increases. The same material means the same lossy physical characteristics, so the intersymbol interference (ISI) gets worse when a higher data rate signal propagates through the same channel material. ISI manifests as a more closed eye diagram, corresponding to increases in time jitter, amplitude noise, and a higher BER. Since UI is fixed for a given data rate to maintain the same BER metric, either the ISI (caused by the channel) must be equalized or compensated, or the jitter (from transmitter and receiver) must be reduced. Figure 2 shows what happens to the eye diagram when a high-speed signal propagates through a lossy channel.

Figure 2. Signal and Eye-Diagram Properties Before and After a Lossy Channel at Multiple Gbps



There are many different types of equalization. From the topology point view, equalization may be implemented either at the transmitter, the receiver, or both. From a signal and system point of view, equalization is digital (e.g., finite impulse response (FIR) or infinite impulse response (IIR)) or analog (e.g., continuous time linear equalization (CTLE)), linear or adaptive (e.g., decision-feedback equalization (DFE)). Each type of equalization has certain capabilities, advantages, and limitations. Using a better and intrinsically low noise/jitter component, such as a better oscillator, both the transmitter and receiver jitter can be reduced. Stratix IV GX equalization technology and its associated capabilities and advantages, as well as CRC and its jitter tracking, are covered in more detail in the following sections.

2.2 I/O Protocol Standards Supported

Stratix IV GX FPGAs have dedicated hardware and intellectual property (IP) to support new protocols—PCIe 2.0, HT 3.0, Interlaken, CPRI, and SFI-5—broadening its overall capability. In addition, Stratix IV GX FPGAs support many other high-speed I/O standards, covering a wide range of applications in areas of wireline, wireless, computer and storage, broadcast, military, test and measurement, medical, automotive, and consumer. Table 1 summarizes all the high-speed standards supported by Stratix IV GX FPGAs, with the newly added standards highlighted.

Table 1. High-Speed Protocol Standards Supported by Stratix IV GX FPGAs

Existing Protocols	Data Rate
PCIe Gen1	2.5 Gbps
GbE/TSE	10/100/1000 Mbps
GPON	1.244 Gbps upstream, 2.488 Gbps downstream
Interlaken	4.976 Gbps to 6.375 Gbps
CEI 6G (SR/LR)	4.976 Gbps to 6.375 Gbps
SFI-5.1	2.488 Gbps to 3.125 Gbps
PCIe Gen2	5 Gbps
SRIO	1.25 Gbps, 2.5 Gbps, 3.125 Gbps
HiGig2	4.0625 Gbps
SONET (OC3/OC12/OC48)	155.52 Mbps, 622.08 Mbps, 2.48832 Gbps
XAUI/HiGig/HiGig+	3.125 Gbps, 3.75 Gbps
CPRI v3.0+	614.4 Mbps, 1.2288 Gbps, 2.4576 Gbps, 3.072 Gbps
FibreChannel (FC1/FC2/FC4/FC8)	1.0625 Gbps, 2.125 Gbps, 4.25 Gbps, 8.5 Gbps
SDI (SD/HD/3G)	270 Mbps, 1.485 Gbps, 2.97 Gbps
OBSAI v3.1+	768 Mbps, 1.536 Gbps, 3.072 Gbps
PCIe Cable	2.5 Gbps
ASI	270 Mbps
Serial Data Converter (JESD204)	312.5 Mbps to 3.125 Gbps
3-Gbps Basic (proprietary)	Up to 3.125 Gbps
HT 3.0	2.4 Gbps, 2.8 Gbps, 3.2 Gbps
SerialLite II	See 3-Gbps Basic and 6-Gbps Basic
6-Gbps Basic (proprietary)	Up to 6.375 Gbps

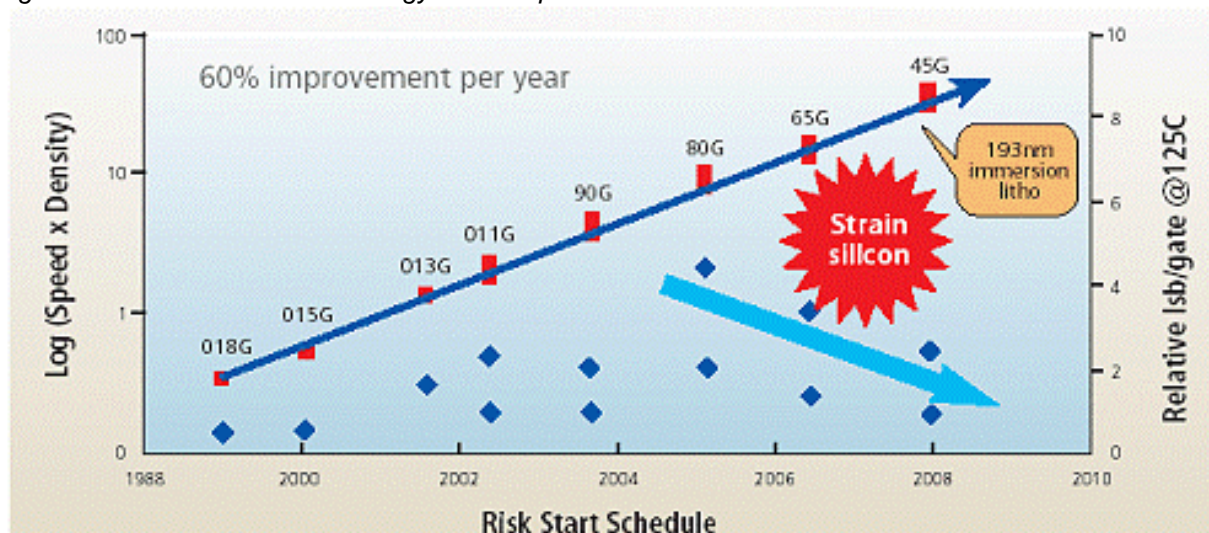
3. 40-nm Process Node and Transceiver

This section begins with the process technology roadmap, moves on to challenges and solutions associated with deep-sub-micron (DSM) effects, and ends with a discussion of transceiver circuits for FPGAs at 40 nm.

3.1 Process Technology Roadmap

The CMOS technology continues to evolve at the pace predicted by Moore's Law, as demonstrated by TSMC's technology roadmap(2) in Figure 3.

Figure 3. TSMC Process Technology Roadmap



CMOS scaling is driven mainly by digital applications such as memories and microprocessors. Memories place high density and low power constraints on CMOS technology, while microprocessors demand high performance and speed. Analog must use a digital process for two main reasons. First, most ICs are, in fact, SOCs, which integrate digital core, memories, microprocessors or digital signal processors, I/Os and analog interface circuits such as A/D and D/A converters, phase-locked loops (PLLs) and transceivers. However, for most of the SOCs, digital circuits dominate the die area; hence the choice of technology. Second, special RF and analog processes cost more and normally are offered later than the general digital process.

3.2 Analog Challenges vs. Process-Node Shrinkage

Analog design faces many new challenges as silicon continues to scale down, including headroom, gain, leakage, DSM effects and associated modeling, variations, and mismatches.

3.2.1 Headroom

Two forces drive supply voltage down. The feature size is getting smaller, as the electrostatic discharge (ESD) and electron migration (EM) considerations require low supply voltage, while other effects, like hot carrier injection (HCI), negative bias temperature instability (NBTI), and time-dependent dielectric breakdown (TDDB), limit supply voltage to the lower levels. In addition, a lower supply voltage cuts down power consumption for digital circuitry quadratically, based on the well-known equation, $P = CV^2f$.

Although low threshold voltage, V_t , means a high switching speed for digital applications, the dramatic leakage increase adds substantial static power consumption. That is why V_t does not decrease when the supply voltage is reduced. For example, the supply voltage was 1.2V in 65 nm but is 1.1V in 40 nm, so the V_t is only lowered by about 10~20mV.

The lower supply voltage and the stable V_t reduce the headroom, which creates a problem in some analog designs. The lower headroom makes it difficult to maintain signal integrity at an acceptable power consumption cost, and cascode structures in bias circuits and amplifiers do not work due to transistors in the triode region. The resulting consequences are:

- Current mirror errors add more variation in bias, speed, and aging
- Small signal swings lead to worse signal-to-noise ratio (SNR) and lower operational amplifier gain
- Variations in other current based analog parameters such as output drive voltage (V_{OD}) swings

3.2.2 Gain

Amplifier gain is expressed as $A_v = G_m/G_{ds}$. Scaling from 65-nm to 40-nm G_{ds} shows an increase while the normalized G_m stays at about the same level. This results in a worse G_m/G_{ds} , or gain. G_{ds} is a very critical parameter in analog design. Aside from gain (G_m/G_{ds}), a low G_{ds} is necessary for proper functioning of current sources and mirrors. Analog devices must be designed specially to maintain low G_{ds} by balancing pocket and channel implants. Without proper balancing, the G_{ds} suffers, too. For a generic digital process, it is customary simply to tune the pocket implant (practiced by most foundries) without introducing adverse effects to G_{ds} .

3.2.3 Leakage

There are two types of leakage in a metal-oxide semiconductor field-effect transistor (MOSFET). Source-drain current leakage has hit more on digital than analog due to static power consumption increase in logic. Transistors in analog blocks normally are biased in saturation, and source-drain leakage is not a concern. One exception is dynamic gates in high-speed analog circuits. Leakage prevents these circuits from working at low frequencies. Some remedies push down this low-speed boundary, but they also reduce the speed limit at the high end.

Gate-tunneling current leakage has a more severe impact on analog. Gate leakage increases drastically when migrating to a smaller process node. When the gate oxide thickness is reduced with the equivalent of one atomic layer, the gate current increases by approximately one order of magnitude. Despite technological remedies, gate leakage is becoming a critical consideration in analog design, especially for long transistors. Due to gate leakage, gate noise is increases, leading to increased jitter, and increases the mismatch for large devices (large $W \times L$). Normally the mismatch decreases as a function of the device area; however, the gate leakage offsets this common behavior. The digital designer normally does not need to worry about these effects.

3.2.4 DSM Effects and Models

SiGe and other strain technologies have been used in 65-nm and 40-nm technology nodes and are expected to be deployed for future technologies. This results in device characteristics becoming a function of many layout dependencies, such as length of defusing (LOD), well proximity effect (WPE), and poly spacing.

DSM effects lead to new challenges for technology nodes. For transistor modeling to have an accurate DSM feature size, the geometry extraction must take into account each individual device in relation to its surroundings. Such an extraction covers all DSM effects, with each device adjusting its performance accordingly. Since an individual device performance deviates from the ideal (pre-layout) model, complex layout rules must be developed to guide layout designers to reduce the number of post-layout iterations. Finally, complex geometry extractions lead to complicated models used in post-layout simulations, resulting in longer simulation times and analog design cycles.

3.2.5 Mismatch

Mismatch must be accounted for in analog design. Normally, increasing the area of the device reduces mismatch. However, as was noted in “3.2.3 Leakage”, gate leakage increases cause less improvement in mismatch by increasing area. Another dimension to be considered in respect to mismatch is detailed device engineering involved in DSM processes. Balancing the pocket and channel implants is critical for minimizing anomalous mismatch behavior. Normally, increasing channel length decreases mismatch. However, mismatch reduction due to channel length increasing is offset by mismatch contribution due to pocket implant.

3.3 Solutions and Trends

This section discusses solution and trends for analog device at DSM process node (e.g., 40 nm).

3.3.1 High Supply Voltage for Analog Blocks

Analog design can be subdivided into two general categories, precision analog and high-speed analog. In general, for analog design, the ITRS has proposed using different voltages and oxides than those used in digital logic. Higher voltages and thicker oxide are best suited for the precision analog category, while lower voltages and thin oxides are best for high-speed analog. Altera has utilized this dual-technology approach in analog design since 130-nm-technology node. Now at 40 nm, it is essential to move more precision analog circuits to a higher voltage supply to solve the headroom issue.

3.3.2 NML and Asymmetric Devices

Analog circuits use non-nominal length (NML) transistors, which have more balanced V_t , G_m/G_{ds} , linearity, and other relevant metrics. Transistors that have 2X to 5X of minimum channel length are commonly used in the mixed-signal design, a trend that was recognized in the 2005 edition of ITRS. As semiconductor companies move to more of a mixed-signal format, transistors are optimized for analog as well. Analog designs in deep sub-micron technology tend to be low voltage. A combination of high V_t and low V_t transistors throughout a complex analog block is found to be beneficial to the overall performance.

3.3.3 Analog Devices for FPGAs

To meet the increasing challenges of mixed-signal/analog design inside FPGA, Altera has designed analog transistors using the TSMC standard logic process. These analog transistors with specially tuned V_t levels, G_m/G_{ds} , and mismatch characteristics successfully meet the analog design requirements. They are designed to ensure adequate V_t headroom for the analog operation at power-supply voltage determined by logic operations in FPGA. The analog V_t headroom is designed to be a ratio of the power-supply voltage, which is needed to operate stacks of transistors typical of an analog design, such as the CMOS cascode current supply, and to maintain SNR integrity with low power.

During the engineering of these transistors, special attention was paid to ensure high gain and good mismatch behavior. Special channel implants are typically employed in an advanced logic process to control the short channel effects (SCE). These SCE control implants normally are optimized for the logic operation, not for analog. The analog transistors use optimized implants and analog parameters for the high-speed SERDES operation.

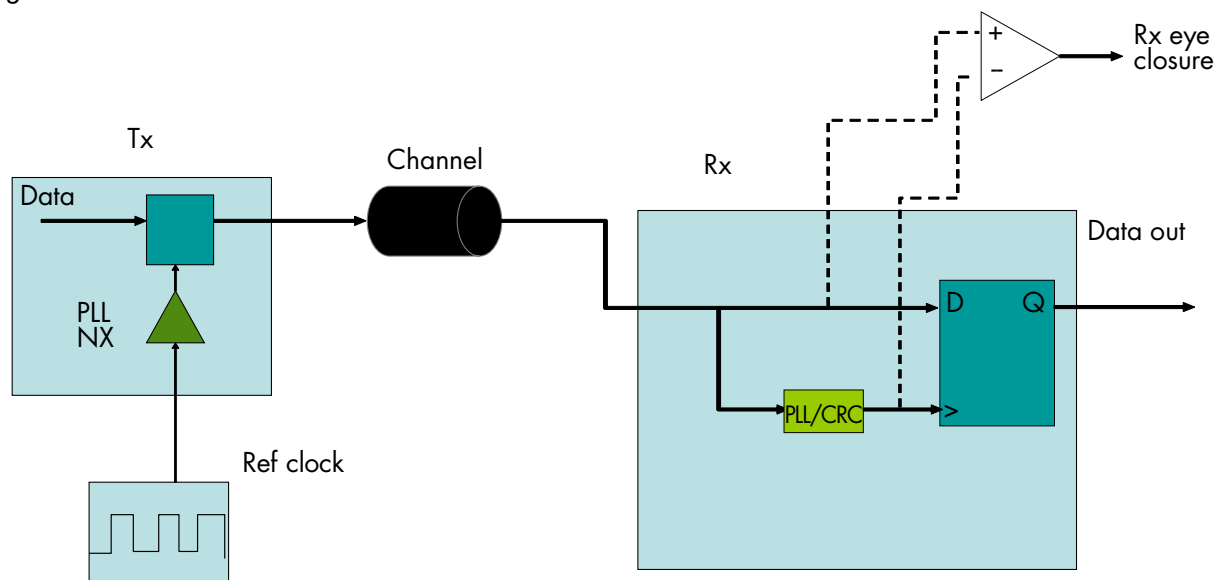
4. Architectures

This section reviews the mainstream multiple Gbps link architectures, starting with the conventional data-driven, then the common clock, and in the end, Altera's hybrid architecture. The advantages and limitations of each architecture also are discussed.

4.1 Data-Driven Architecture

The SERDES transceiver has emerged as the dominant transceiver for data rates above 1 Gbps. Earlier serial link architecture embedded the bit clock in the transmitting bitstream and relied on a receiver to recover the bit clock solely based on the incoming bitstream. Within the frame of this architecture, the CRC must recover both the frequency and phase of the clock from the received data and a PLL-type CRC fits well to meet the requirements. A PLL can be used as a second- or higher-order system that is capable of recovering both frequency and phase simultaneously for the bit clock. This type of architecture is often called data-driven architecture⁽³⁾⁽⁴⁾ and is illustrated in Figure 4.

Figure 4. Conventional Data-Driven Serial Link Architecture



There are many advantages for data-driven architecture. The two most important are: 1) it is simple, with no need to send a clock to the receiver along with the data; 2) it has good jitter tracking and tolerance capability since the PLL CRC can be used as a second- or third-order system, corresponding to a 40-dB/decade or a 60-dB/decade jitter tracking/tolerance. The implications of the first advantage are multiple folders, including easy testability, and a wide range of channel length reach, from chip-to-chip (~1 m), to board-to-board (~ 10 m), and system-to-system (~ or >1 km). A wide-range channel reach corresponds to a wide range of transceiver applications. The second advantage implies a relatively larger low-frequency jitter budget for transmitter, enabling a cost-effective transmitter design and the use of low-cost reference clocks. The drawbacks for the data-driven architectures with an analog PLL CRC include a relatively larger silicon area and a longer lock-in time. However, some limitations may be removed by using the hybrid digitally assisted PLL CRC, which is covered in “4.3 Altera's Hybrid Architecture”.

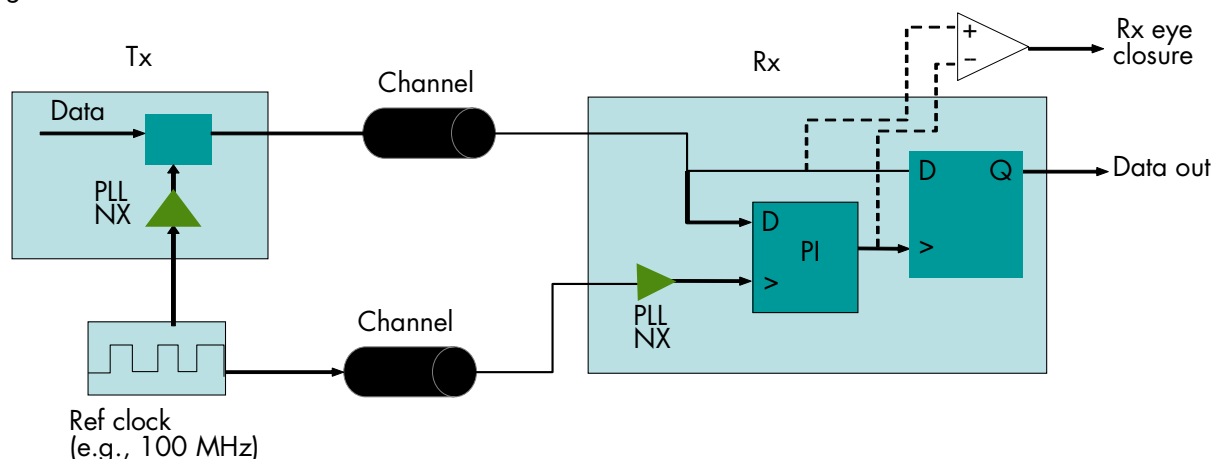
4.2 Common Clock Architecture

The SERDES transceiver was first developed and used in network communications (e.g., SONET, GbE), and soon adopted by computer I/O communication interfaces (e.g., PCIe). Due to the short distance reach requirement for computer I/Os (typically within 10 m), it is possible to send a low-frequency reference clock to both the transmitter and receiver, and then convert it to an in-rate clock for transmitting and receiving data through a multiplication PLL. With an in-rate clock at the receiver, the clock recovery only must recover the phase, making the use of a simple CRC

possible, such as phase interpolator (PI) that is essentially a “bang-bang” phase detector. The PI determines whether the phase between data and clock is earlier or later and aligns them over a certain time period. Thus, PI CRC is approximately equivalent to a first-order FIR.

The major advantages for a common clock architecture(5)(6), illustrated in Figure 5, include 1) maintenance of the phase correlation for the reference clocks used in the transmitter and receiver, enabling the use of a low-cost, relatively cheaper reference clock; and 2) a shorter lock time since it is digital based. The drawbacks include much complex architecture requiring both clock and data inputs for CRC, extra and complex test requirements, and less (20 dB/decade) jitter tracking and tolerance capability for its receiver CRC. Moreover, common clock architecture is limited to short-reach computer I/O application, thus it has limited applications.

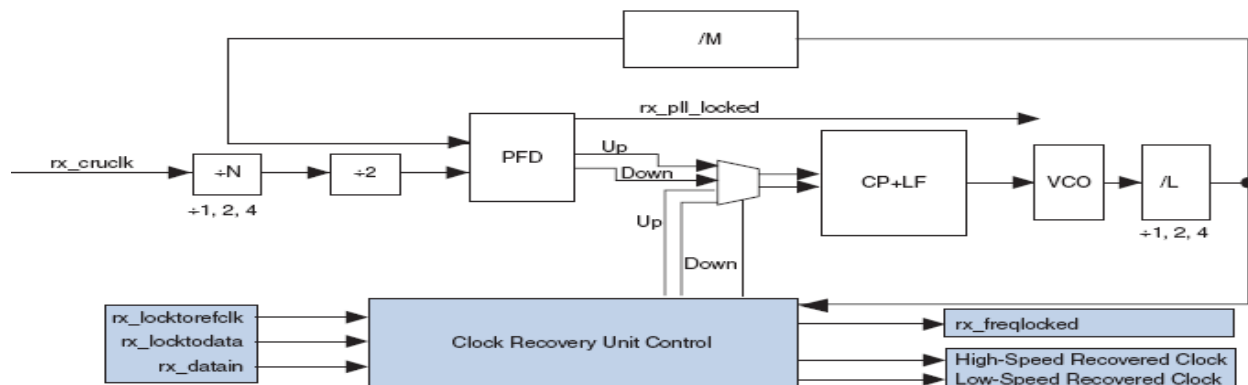
Figure 5. Conventional Common Clock Serial Link Architecture



4.3 Altera's Hybrid Architecture

Stratix IV GX architecture advances the conventional data-driven architecture by allowing two operational modes, as well as digital controllability for its CRC(7). The two operational modes are lock-to-data and lock-to-clock, and can be used automatically or manually. A common usage mode is to use the reference clock as the input and lock to the desired frequency, then switch the input to data signal to lock to the phase. In the end, both frequency- and phase-aligned bit clocks are recovered. In addition to maintaining most of the advantages of the data-driven architecture, this hybrid architecture offers better lock time, power consumption optimization, and tolerance of jitter and transition density. A schematic drawing for this architecture is shown in Figure 6.

Figure 6. Digitally Assisted Hybrid Serial Link Architecture in Stratix IV GX FPGAs

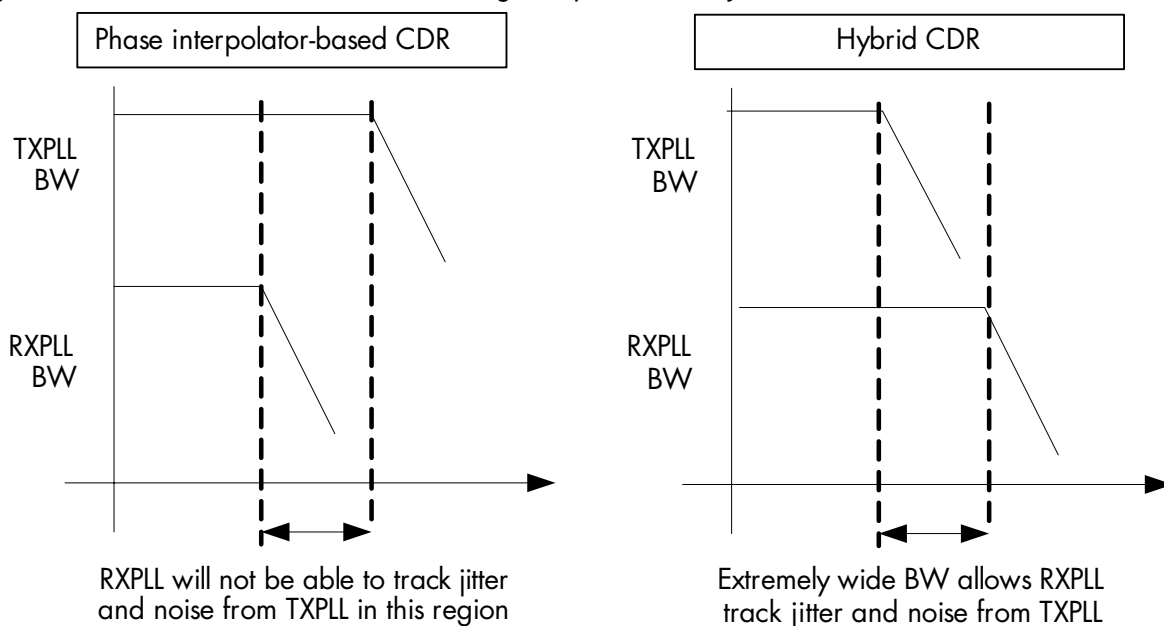


When compared to other link architectures, the hybrid architecture of Stratix IV GX FPGAs has advanced capabilities in filtering and reducing jitter from transmitter and reference clock, therefore enabling best possible BER performance for the link system. In the hybrid architecture, the reference clock is only used in the initial training phase for CRC, and is not used in the actual data recovery phase. Therefore, reference clock jitter does not contribute to the system BER at all, leaving system designers a significant jitter margin for their system design. In contrast, the common-clock architecture uses the reference clock directly for both clock and data recovery, therefore the reference jitter contributes to the system BER and consumes system jitter budget. As for the conventional data-driven architecture, the CRC may never recover the bit clock or take a long time to lock on the incoming data when the receiver data input has excessive jitter, while hybrid architecture is largely immune from such problems since its frequency-locked bit clock is already in place when recovering the data.

In addition, because the hybrid architecture does not use reference clock for data recovery, there is no need for the reference compliance test, simplifying the system design and reducing the associated cost, while the reference-clock compliance test is required and necessary for common-clock architecture.

Finally, due to its digitally assisted design, the hybrid architecture sets the bandwidths of clock generation PLL at the transmitter and the clock recovery PLL at the receiver independently, so the receiver PLL CRC bandwidth is much wider than that of the transmitter PLL (as shown in Figure 7). Consequently, the transmitter's high-frequency jitter is attenuated by its PLL, while the transmitter's low-frequency jitter is attenuated by the receiver CRC, leaving a minimum jitter contribution in the system BER. In contrast, common-clock architecture does not have the freedom to set the receiver PLL bandwidth to be wider than that of its transmitter PLL (8)(9), so jitter from transmitter contributes to the system BER.

Figure 7. Stratix IV GX PLL Bandwidth Setting Comparison for Hybrid and Common Clock Architectures



5. Mixed Signal Clock Recovery

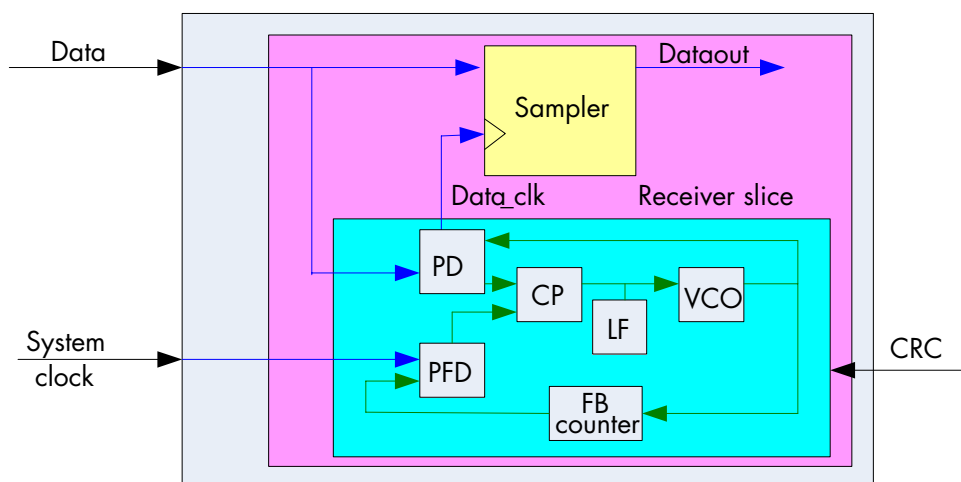
This section covers clock recovery (10), an importance unit for a serial receiver, starting with detailed review of the Stratix IV GX clock recovery subsystem architecture and working mechanisms, then moving to capabilities and performance, as well as comparison with other relevant clock recovery technologies.

5.1 Hybrid (Digitally Assisted) Clock Recovery

Stratix IV GX FPGAs use a hybrid clock-recovery architecture, which is a mixture between an analog PLL-based CRC and a phase-interpolator-based CRC, as shown in Figure 8. In this architecture, the CRC receives both data and reference clocks. The CRC itself includes two fully integrated loops, which are PFD (lock-to-clock) and PD (lock-to-data) loops. They share a common charge pump (CP), loop filter (LF), and voltage control oscillator (VCO). The PFD loop is equivalent to normal 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 PFD mode). Once the CDR reaches the desired output frequency, the CRC automatically switches from the PFD loop to the PD loop so that it can track incoming data and generate the recovered clock quickly since frequency is recovered at this time. Within the CRC, there is a PPM detection circuit that constantly checks the parts-per-million (PPM) difference between recover 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.

Figure 8. Stratix IV GX Clock Recovery



The Stratix IVGX CRC is also a design for power optimization. In order to reduce the power consumption, the CRC is implemented by using half-rate design in which the VCO frequency is half of the data transfer rate and both data and clock edges are used to sample incoming data.

5.2 Capability Comparisons

Each of the link architectures and associated CRCs has its own advantages and drawbacks, including jitter tolerance, design flexibility and area, transition density, run-length tolerance, and lock time performance. It is important for system designers to understand the capabilities and limitations for the integrated devices so that appropriate design and design trade-offs are made to meet their system requirements.

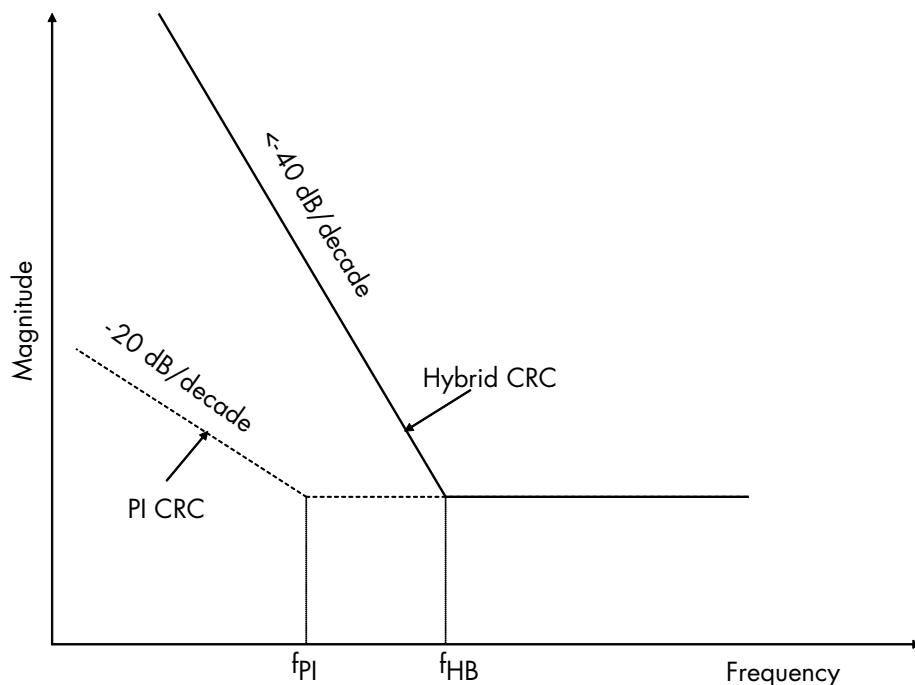
5.2.1 Jitter Tolerance

Jitter tolerance measures how much jitter in the incoming data the receiver can tolerate without losing data. Jitter tolerance is closely related to jitter transfer function (JTF), which is governed by CRC properties, and is called for as a compliance test by many high-speed communication I/O standards. It is important for system designers to choose a

receiver that meets or exceeds the standard requirement. For example, for the hybrid PLL CRC discussed in “4. Architectures”, the JTF is a second-order high-pass minimum, with a 3-dB bandwidth at up to several 10 MHz. The hybrid PLL CRC tolerates a significant amount of jitter at low frequencies.

In contrast, for PI-based CRC, its bandwidth depends on the number of integration steps (phases), up and down counter settings, and a controller update rate. It is desirable to be able to generate as many steps as possible to achieve the necessary phase resolution, but it reduces the CDR bandwidth. Also discussed in “4. Architectures”, the PI-based CRC has a first-order high-pass JTF, with its 3-dB bandwidth typically at around MHz range. Due to the limited jitter tolerance, a low-jitter and high-performance transmitter must design the link system that uses a PI CRC. Figure 9 shows the jitter tolerance comparison between a PI CRC and a hybrid CRC.

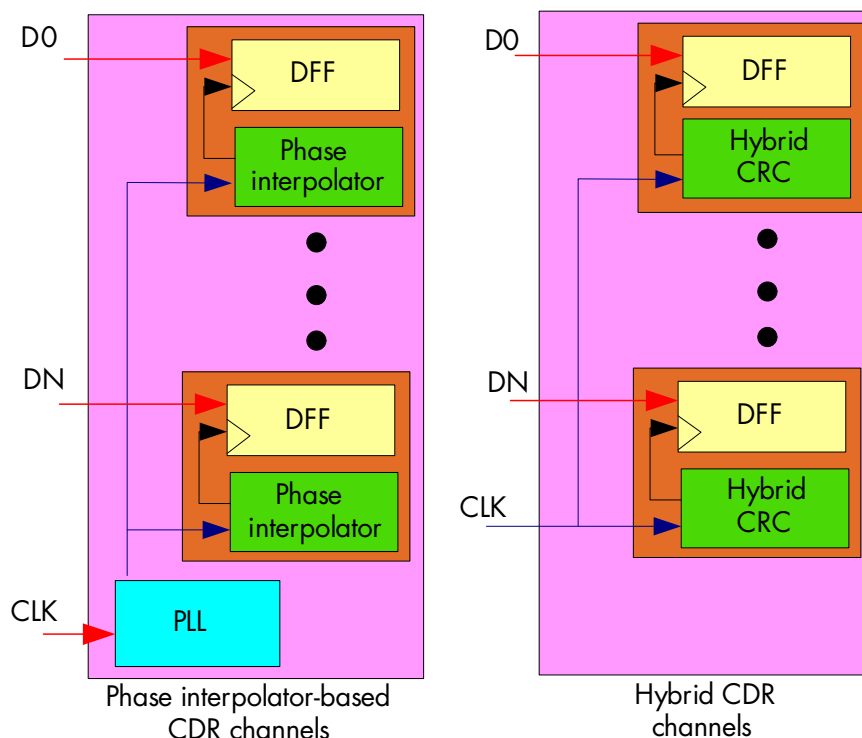
Figure 9. Jitter Tolerance Comparison



5.2.2 Flexibility and Versatility

Most of the serial I/O link systems must support multiple lanes, which affects the clock and data signal input structure for receiver. Figure 10 shows multiple channel signal and clock topologies for PI and hybrid architectures.

Figure 10. Multiple-Lane Signal Structure Comparison



A PI CRC-based receiver normally uses a single PLL to support a multiple-channel operation, but this limits the system to support only one standard and one data rate, unless more PLLs or a new clock distribution network is added. In contrast, in a hybrid CRC-based receiver, each channel has its own CRC and they are very independent of each other. Since each hybrid CRC in the multiple lane receiver is independent, each supports a different data rate in a different lane. This architecture provides system designers with good flexibility to support different standards and/or data rates simultaneously.

5.2.3 Transition Density and Maximum Run Length (MRL)

All the CRCs rely on some data transition to recover and maintain bit clock from the incoming data. A PI CRC maintains its recovered clock even if it does not receive any data transition for many UIs. That is because its clock is derived from the input clock phase, which is independent of the data signal. In general, analog CRCs have a lower tolerance with MRL, as such it normally require designers to either encode the data stream (e.g., 8b/10b encoding, MRL = 5) or scrambled data (e.g., PRBS 2⁷-1, MRL = 7) such that it has a reasonable MRL and transition density to work without drifting away from the desired frequency. On average, the transition density is inversely proportional to the MRL. The hybrid CRC MRL capability is in the range of 500 to 1000 UIs, between analog PLL and PI CRCs and much longer than most of the MRLs defined by I/O standards.

5.2.4 Lock Time

Analog PLL CRC has a much longer lock time compared with the digital PI CRC. Longer lock time causes latency and lower efficiency for the link system. A hybrid CRC, due to its two-mode operation capability, however, is able to lock quickly to the desired frequency and phase, so its lock time is shorter than the conventional analog PLL-based CRC.

6. End-to-End Equalization

This section covers the need and motivation for equalization, different equalization methods and their trade-offs, and how the backplane influences the type of equalization chosen. The described architecture satisfies a continuous data rate from 0.622 to 6.5 Gbps, but proper equalization requires fine-tuning for each specific backplane. Supporting a host of customer backplanes requires thousands of settings, so adaptive equalization is necessary to tune for each unique backplane. Altera has a long history of best-in-class equalization solutions.

6.1 Motivation

As introduced in “2.1 Technology Trends and Challenges”, increases in the data rate lead to higher attenuation, reflections, and coupling. Chips are now so fast that when communicating through a backplane, it hits the bandwidth limitations of the wires(11). Figure 11 shows a diagram of a typical backplane, where the Tx device transmits a signal to Rx device. The goal is to deliver an electrical signal from Tx to Rx at as high a data rate as possible. As the transmitted signal leaves the Tx driver, it must go through the Tx I/O card traces, followed by an I/O card connector, then the backplane traces, another connector for the Rx I/O card, plus another set of I/O card traces at Rx side. The bandwidth limitations of this system are a result of the I/O cards, connectors, and the backplane itself, which is over 40" of FR-4 material. The skin effect, dielectric loss in the wires as well as reflections caused by the vias, significantly distorts the signal from Tx to Rx. The bandwidth of the backplane wires is typically in a range between several hundred MHz up to 2 GHz.

Figure 11. Diagram of Typical Backplane Communication Between Two Chips

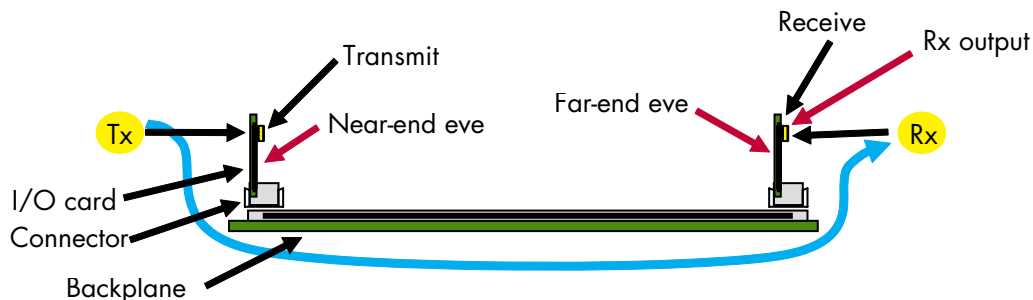
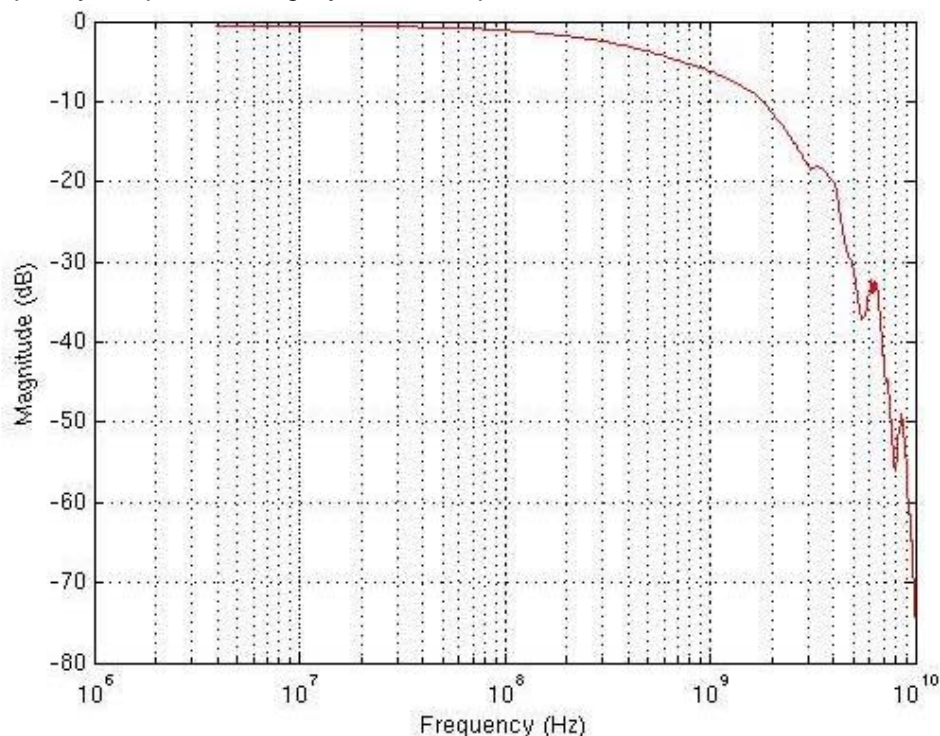


Figure 12 shows the frequency response of legacy XAUI, a subset of I/O standards for 10G GbE backplane built for 3.125-Gbps signaling. Approximately 15 dB of attenuation occurs at three-quarters of the baud rate, or roughly 2.3 GHz. However, if the same legacy backplane is used for a 6.5-Gbps operation, over 30 dB of attenuation occurs.

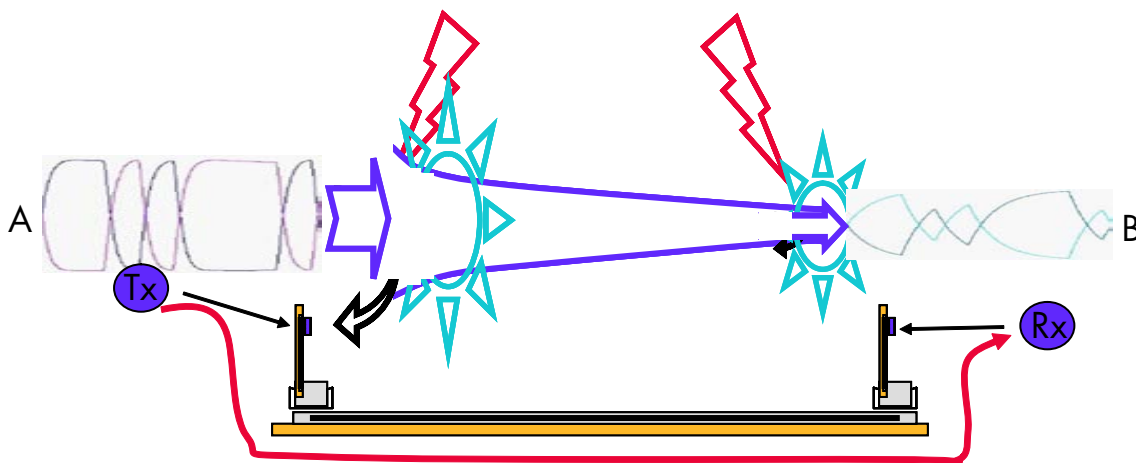
Figure 12. Frequency Response of Legacy XAUI Backplane



Assuming the legacy XAUI backplane is used to communicate between Tx and Rx, Figure 13 illustrates various effects the Tx signal experiences. Initially, there is a good quality differential signal observed at point A, the near end in the link. 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 in the link, a degraded version of the signal is launched from point A. In extreme cases, the signal at point B may be so attenuated that the two differential signals do not even cross. This is typically true in long backplanes, where the incident signal degrades further, and for higher data rates in the same backplane, which would experience higher degradation. These highly distorted signals still must be processed somehow.

Figure 13. Transmitted Signal Affected by a Transmission Medium

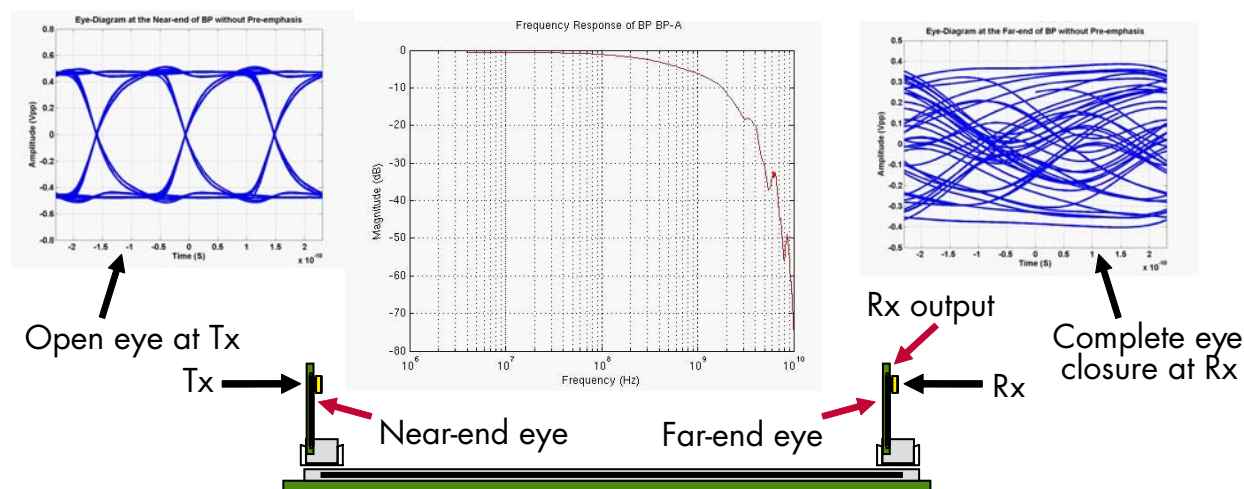
Incident => Attenuation + reflection + radiation + coupling



Support of legacy backplanes at increased data rates (beyond what they were originally intended for) is a cost-effective way to upgrade present systems. All that is required for a system upgrade is an I/O card replacement with new-generation high-data-rate transceivers capable of signal conditioning in legacy systems. Exotic backplane materials, which support increased data rates, are more expensive than conventional FR-4 material. Furthermore, replacement of the backplane may imply a costly overall system replacement, which is why system designers try to push as much bandwidth as possible with new silicon transceivers operating with legacy backplanes.

Figure 14 shows a 6.5-Gbps signal sent over the example XAUI backplane. The transmitted signal is so attenuated that no recognizable eye is seen at the far end. The most practical solution to this dilemma is signal conditioning to open the eye before receiver samples the data, and on-chip equalization is the most practical way to compensate for the backplane attenuation.

Figure 14. Example With XAUI Backplane at 6.5 Gbps



6.2 Equalization in a Nutshell

The transmission medium is a linear system, which creates an inverse transfer function. When this inverse transfer function is added to the transfer function of the link, the goal is to have a resultant transfer function that is relatively “flat” up to the required frequency. Figure 15 shows this simplified concept in frequency domain.

Figure 15. Equalization Simplified

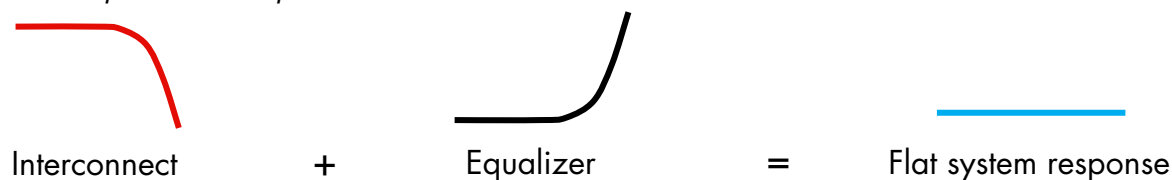


Figure 16 shows eye diagrams at three key points in the link. The first (left) is a launched signal from Tx. As this signal goes through the backplane, the resulting signal is attenuated (center). After the attenuated signal is sent through the equalizer, the original eye is restored (right).

Figure 16. Resulting Eye Diagram After Equalization



Since linear system is assumed, it is important to note that signal conditioning can be applied before or after the interconnect. In the previous example, the equalizer was placed at the far end, at the receiver. Similarly, the signal at Tx is pre-distorted so that after it goes through the interconnect, the resulting signal is clean for recovery at Rx. This type of signal conditioning is usually called “emphasis,” of which there are two types of emphasis, pre-emphasis and de-emphasis. However, a complicating factor is that the link may not have a simple attenuation curve. There may be multiple poles in the transfer function, as well as reflections, crosstalk, and resonances at some frequencies. Therefore, it is important to choose the right equalization selection from large toolbox of available signal conditioning techniques.

6.3 Types of Signal Conditioning

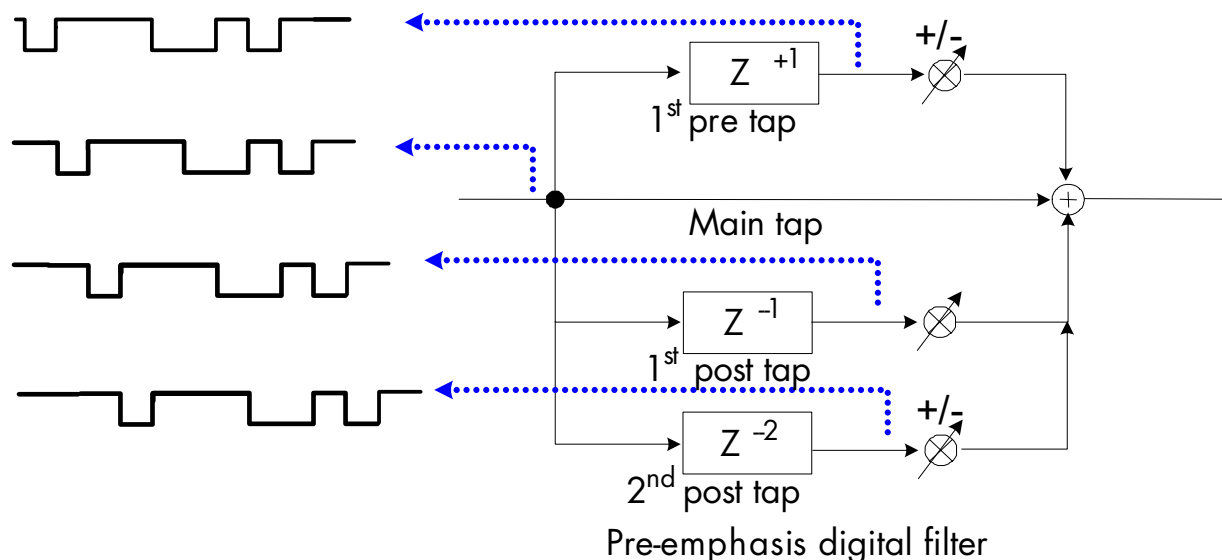
There are many types of signal conditioning, or compensation/equalization, each of which has benefits and disadvantages. This section reviews some of the more popular implementations.

6.3.1 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 so that the signal high-frequency contents are amplified (pre-emphasis), or the signal low-frequency contents are reduced. The benefit of this method is relative simplicity and low power. All the sampled data is readily available at transmitting devices. Delayed versions of transmitted data easily are created to be one UI apart by placing a bank of registers that to hold both prior and upcoming serial data bits. Similarly, fractional sampled data (1/2 UI) is easily available in Tx by taking information from the intermediate latch stages that are commonly used to create registers. Since both prior and upcoming (future) data bits are available, this signal conditioning technique addresses both pre- and post-cursor ISI.

Figure 17 demonstrates a block diagram for a possible pre-emphasis implementation. In general, transmit pre-emphasis is relatively simple to implement since the design has “control” of the delays and clocking, in contrast to the receive side. The causal taps (post-taps) remove the post-cursor ISI, while the anti-causal taps (pre-taps) address the pre-cursor ISI.

Figure 17. Block Diagram of Pre-Emphasis Implementation



The single pulse response in Figure 18 shows the time domain pre-emphasis response. It has pre-tap, followed by main pulse and two post taps.

Figure 18. Examples of Pre-, Main, and Two Post-Taps

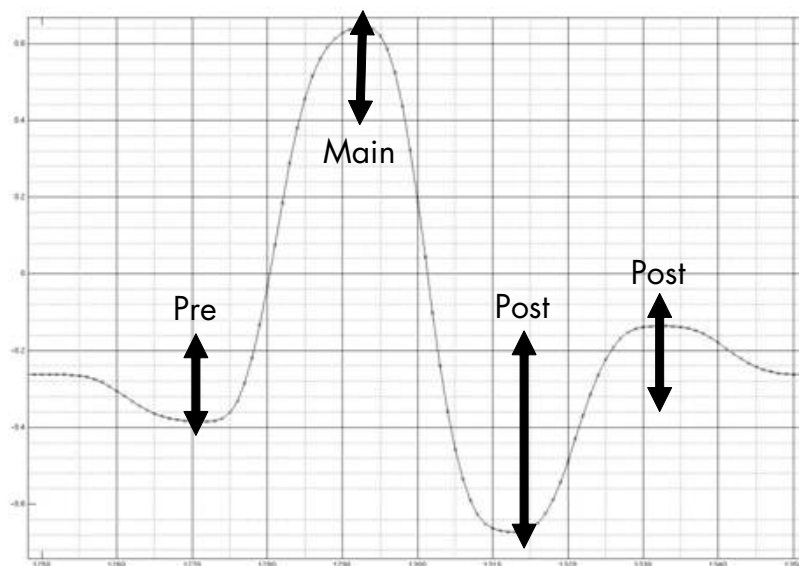
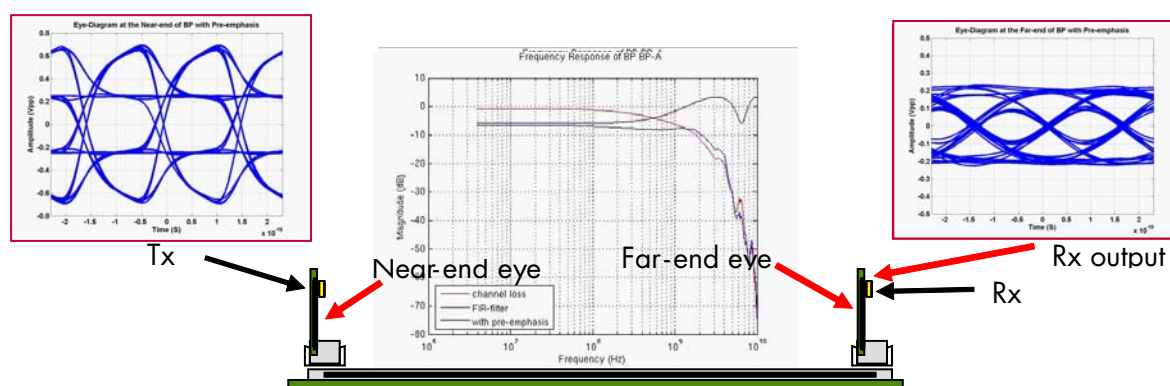


Figure 19 shows the same XAUI backplane used in Figure 14, but with a pre-distorted Tx signal. The open eye is seen clearly at the far end now, even though the same 6.5-Gbps data rate is used.

Figure 19. Transmitter Pre-Emphasis at Near End and Far End



The main disadvantages of signal conditioning on the Tx side is that they do not address crosstalk and are very hard to make adaptive in real time. Regarding crosstalk, Tx pre-emphasis may even increase the amount of crosstalk in the system. This is because the signal is pre-conditioned on the Tx side to have an excessive amount of high-frequency content, so after it is lost to the lossy channel, the resulting signal has balanced high and low frequency. The consequence is that in multiple serial links systems, the increased high frequency content tends to escape into adjacent links, thus causing crosstalk.

Real-time adaptation is a critical parameter to consider, specifically for link operations above 3 Gbps. The actual system signals are different even for identical launched signals at Tx, because at the Rx side, two adjacent links would receive different eye diagrams. For a system that relies on Tx for signal conditioning, a back channel must be created between each Rx and corresponding Tx to change the pre- and post taps of the driver to tailor it to the individual link.

6.3.2 Continuous Time Linear Equalizer

CTLE is implemented on the receive side, and based on the equalization theory, the linear equalizer is well suited here. This simplifies the design, since non-sampled (i.e., continuous time) implementation suffices. As a result, CTLE-based signal conditioning is usually the lowest power choice. Similar to transmit pre-emphasis, CTLE addresses pre- and post-cursor ISI but in continuous time versus being limited to a pre-set number of Tx taps.

Figure 20 shows an example of a first-order CTLE transfer function. Zero is inserted to compensate the poles in channel transfer function (see Figure 12). This implementation is relatively simple and has a very low power consumption rate. Many equalizer stages can be added as given links required for a selected data rate. Multiple stages not only increase the order of the resulting equalizer but also increase the maximum boost achieved in a given frequency interval. (It is worthwhile to note that parasitic poles and their location should be carefully considered when designing CTLE, but a detailed technical discussion of pole and zero placements is outside of the scope of this paper.)

Figure 20. Example of a CTLE Transfer Function

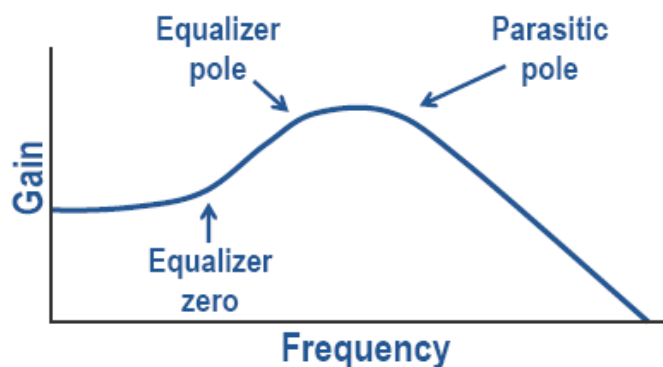
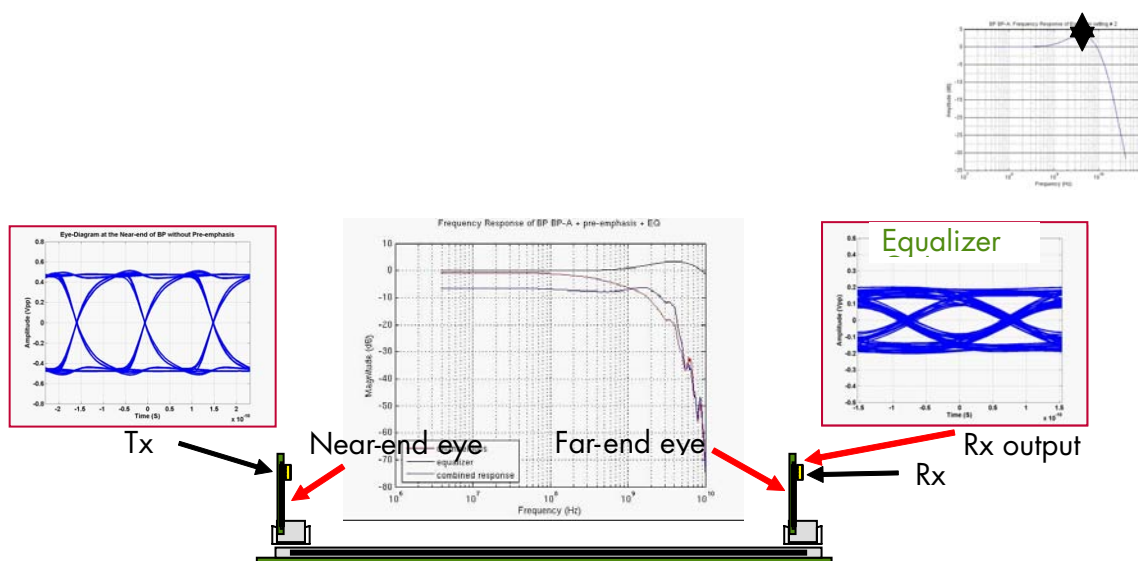


Figure 21 uses the example from Figure 19, but with a CTLE. For the same near-end Tx signal, the far end shows an open eye when internal CTLE is enabled. It is clear that the amount of crosstalk is not increased in a system with CTLE versus Tx emphasis, because of the reduced high frequency content at the Tx drive. Lastly, it is important to note that a system that has a CTLE is well suited for real time adaptation, since signal information after the channel is readily available for processing and reconditioning at Rx.

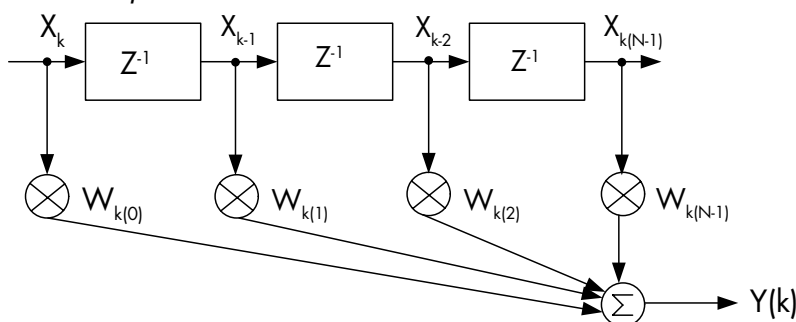
Figure 21. FCTLE at Near End and Far End



6.3.3 Transversal FIR (TFIR)

TFIR is also a non-sampled-based discrete equalization, as shown in Figure 22. It is a linear system that addresses post-cursor ISI, but it cannot deal with pre-cursor signal distortion unless additional delay is introduced to hold the prior data bit, as this information is not available at the time of signal processing. The requirement of a very accurately matched delay element is another major disadvantage for this approach. This disadvantage is two-fold. Firstly, creation of absolute delays that are accurate is a challenging technical problem considering device variation over large production quantities and normal voltage and temperature variation. Secondly, creation of accurate delay elements restricts a design to a fixed frequency of operation that is inversely proportional to chosen delay. This is a main reason this approach is not found in multi-rate transceivers.

Figure 22. Diagram of Unsamplerd Rx FIR

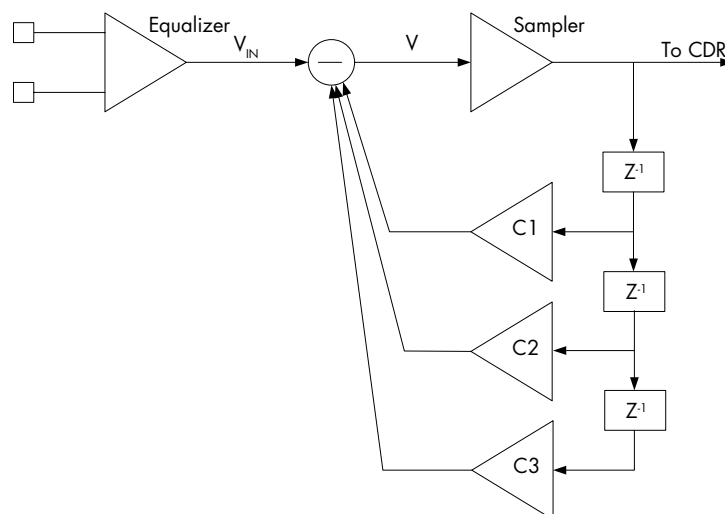


6.3.4 Decision Feedback Equalization

Unlike prior equalization schemes, decision feedback equalization (DFE) is a non-linear system (see Figure 23). The generic DFE system requires not only sampling the data but also computing the new coefficient prior to the next sample. This makes timing closure extremely challenging. The DFE system requires a proper data sample, leading to increased design complexity of combined equalization and recover sections of the transceiver.

Furthermore, an issue of error propagation arises. The error propagation phenomenon is due to the fact that present DFE decision is based on prior samples. An error in that would lead to erroneous coefficient computation for present data sample, hence a single wrongfully captured bit would propagate out to a few consecutive bits until correct samples are obtained again. Since DFE system bases its decisions on prior bits, it only addresses post-cursor ISI. This leaves the pre-cursor ISI uncompensated. As a result, a CTLE type of equalization is still required in a DFE system to accommodate the pre-cursor ISI.

Figure 23. Diagram of DFE Scheme(12)

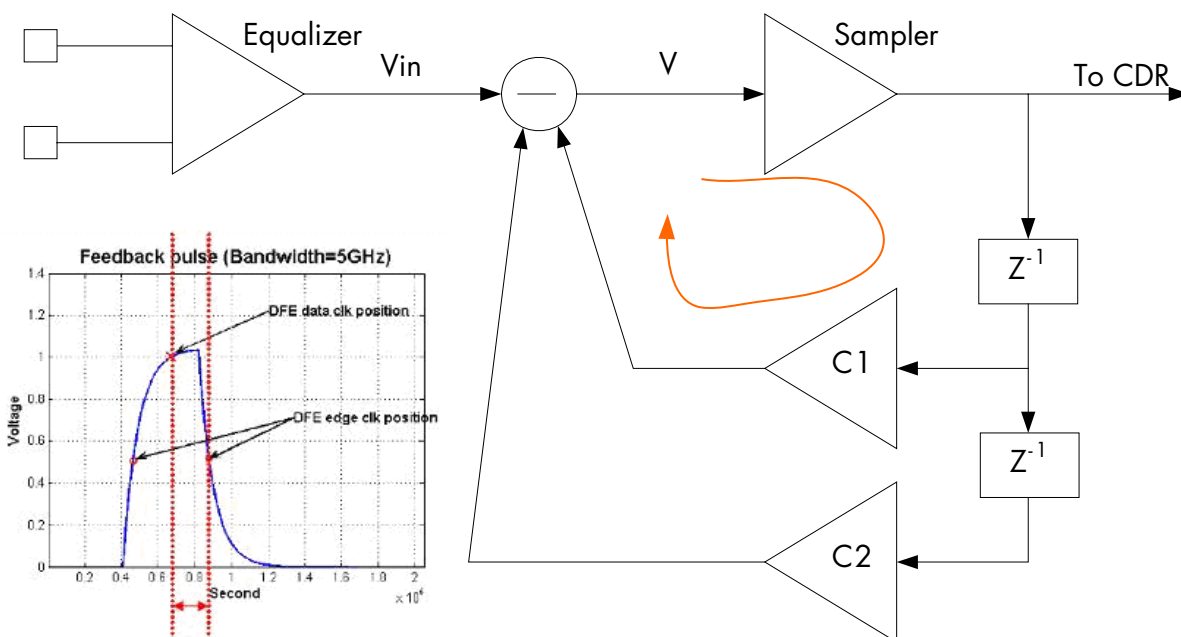


The major benefit of DFE is improved immunity in the case of crosstalk, especially when it is assumed to be additive white Gaussian noise (AWGN). In the case of correlated crosstalk, DFE is not effective over CTLE. The benefit of DFE in an AWGN system is understood by considering the signal-to-noise ratio (SNR) for cases of CTLE and DFE. The system has predominately post-cursor ISI (since DFE does not address pre-cursor ISI). The SNR is computed as ratio of power of signal to power of noise. Since the CTLE is continuous in time and CTLE does not really “know” or “need to know” the incoming signal spectral density, it boosts both signal and noise by equal amount. This, in effect, preserves the original link SNR.

Signal, unlike AWGN, is limited in spectrum to the data rate and harmonics. In the case when the DFE is used to equalize the signal, SNR improves. This is because the DFE operates on sampled data, and not on the whole spectrum. Even though both signal and noise are present in frequency spectrum of data and are boosted by DFE, boosting the fixed frequency spectrum improves the overall SNR because the power of noise is equally distributed over the whole spectrum, a key AWGN assumption. Lastly, it is important to underscore the fairness of the AWGN assumption. The first limit theorem of probability theory (aka Central Limit Theorem) simply states that with a large number of independent variables any process approaches normal distribution.

The DFE effectiveness in improving crosstalk is reduced when implementation difficulties are considered. As pointed out earlier, the timing closure is extremely difficult to meet. In many published high-speed applications, the first tap is not even available due to timing. As seen [Figure 24](#), the feedback loop timing is extremely tight. It is necessary to slice the bit, multiply by the coefficient, and sum it in 0.5 UI. At 10 Gbps, this leaves 50 ps for the first successful tap operation. Additionally, the clock must be recovered from the serial datastream. It has its own jitter, resulting in a further reduction in the available timing margin, and further complicating the design.

Figure 24. Loop Timing



A DFE system is rather complex and hence requires a large amount of power. A general industry rule is 5 mW per Gbps of data per each DFE tap. For example, for 10 Gbps with five taps, DFE alone would require 250 mW/ch.

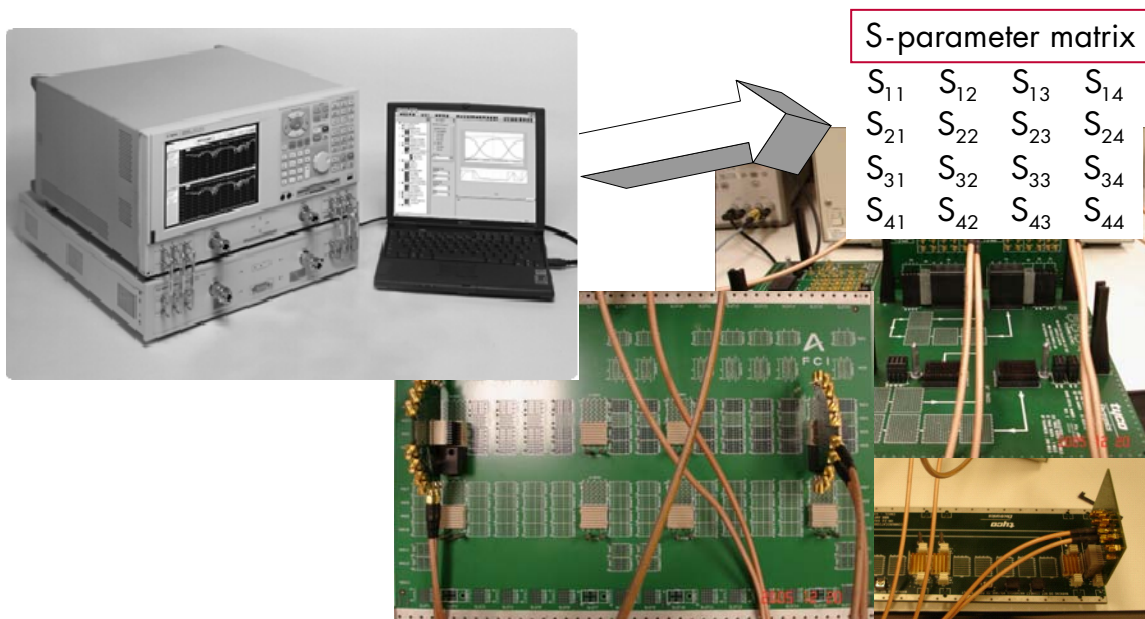
6.4 Backplanes and Their Impact on Equalization Selection

Each link in a system has different characteristics due to many factors, such as the layer it is routed on, length of the route or neighboring signals, board material, etc. How does the designer decide which equalization scheme is good enough? What topology is required and how much boost is needed?

6.4.1 Backplanes and Equalization Selection

In order to design the correct equalization scheme successfully, it is required to build a database of customer backplanes. **Figure 25** shows an overview of this procedure. The critical characteristics of a backplane are extracted with a vector network analyzer (VNA). The extracted S-parameters then are put into a database that can be accessed by the various simulation tools, such as MATLAB, ADS, or hspicrf. Over the years, Altera has created an extensive database of customer backplanes, including those from popular connector vendors and large board manufacturers.

Figure 25. Overview of Characterizing a Backplane System



Backplane studies can be performed on this internal database, with the attenuation curves analyzed by curve fitting and the different equalization schemes modeled. Simulations are run on the database of backplanes to determine the equalization order and architecture required for the best fit.

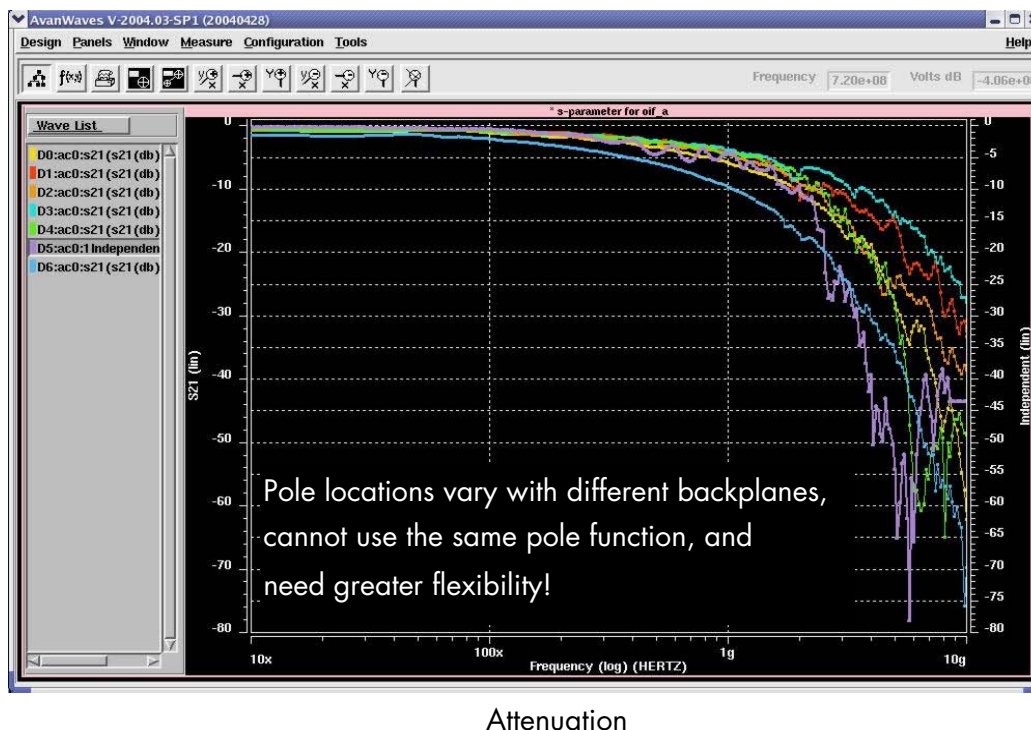
Figure 26 shows the attenuation of a 40" XAUI legacy backplane used throughout this section. Third- and fourth-order pole functions are also plotted. At 6.5 Gbps, a third-order function does a relatively good job of fitting the attenuation curve, while at 10 Gbps it requires a fourth-order function. Curve fitting with the number of poles tells you the number of zeroes required in your equalization to undo the poles of the backplane. (This is a first stab at the problem; high-level modeling is required to model the design truly.) Once proven, the design is converted to a transistor-level implementation and simulated in the frequency and time domain against the database of extracted backplanes.

Figure 26. Third- and Fourth-Order Equalizers Plotted Against Attenuation Curve of XAUI Legacy Backplane



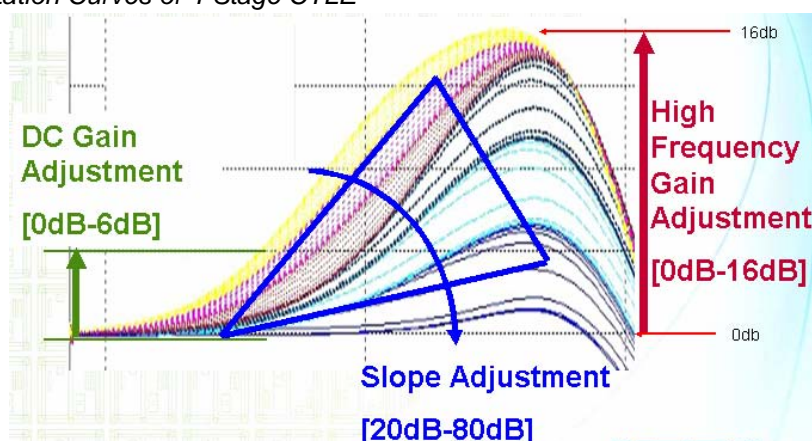
Even the few customer backplanes shown in Figure 27 have quite different attenuation characteristics, so the same equalization transfer curve will not fit all of these designs perfectly. An equalization scheme flexible enough to fit a majority of these curves must be developed.

Figure 27. Attenuation Curves of Sample Customer Backplanes



Analysis of the backplane database of up to 6.5 Gbps shows that all equalizer needs are supported by a maximum slope of 80 dB/dec. Hence, the receiver CTLE has a minimum requirement of four equalizer stages, each with 20 dB/dec slope and independent control over zero placement, to achieve the overall required curve fit. Each stage contributes a zero in the transfer function. Additionally, each stage is controlled independently to allow tuning the zero location as well as the slope of the boost. Lower EQ settings are available to prevent over-equalization for short-reach or chip-to-chip applications. Figure 28 shows some of possible ranges of boost and possible slopes using this architecture. Over 1200 unique settings are possible.

Figure 28. Equalization Curves of 4-Stage CTLE



At the transmit side, a similar backplane study has revealed the need for one pre-tap to address pre-cursor ISI and two post-taps to compensate for post-cursor ISI. In addition, coefficients on all four taps are individually programmable, as well as signs of pre- and second-post-taps. Table 2 shows the available coefficients for taps. The main tap in this table is shown as V_{OD} selection. It allows for over 2700 unique settings for the one V_{OD} setting shown. As a result, a very competitive boost between 15 and 500 percent is achieved.

Table 2. Available Pre-Emphasis Settings

Level	V_{OD} (mv)	pre_tap	1st post	2nd post
1	1000	0.025	0.05	0.025
2	1000	0.05	0.11	0.05
3	1000	0.075	0.16	0.075
4	1000	0.1	0.22	0.1
5	1000	0.125	0.27	0.125
6	1000	0.15	0.33	0.15
7	1000	0.2	0.37	0.2
8	1000		0.42	
9	1000		0.47	
10	1000		0.52	
11	1000		0.56	
12	1000		0.61	

Additionally, Altera Tx and Rx drivers and receivers are designed to support auto-calibration of the termination resistors. Nominally, the process variation of the resistor in deep sub-micron technologies is on the order of ± 20 percent, but this does not even take into account the additional variation of active devices normally used to build termination cells. As a result, an uncalibrated termination cell could easily have over 30 percent variation due to process, voltage, and temperature (PVT). Accurate termination resistors reduce reflected back energy. This improves overall system margins.

An additional Tx feature worth mentioning is programmable slew rate control. The need for agile slew rate is clear from “6.3.1 Transmit Pre-Emphasis/De-Emphasis”. Reducing the high-frequency content of a transmitted signal edge till it most resembles sine-like edges for a given data rate goes a long way toward reducing the overall system crosstalk level and EMI.

6.4.2 The Need for Adaptive Equalization

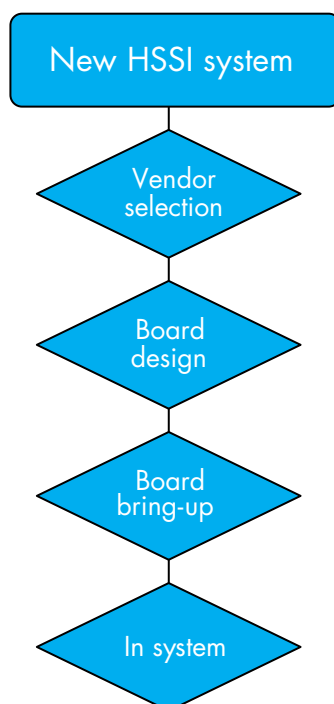
One reason DFEs are attractive is the assumed ability of the DFE engine to compute coefficients for incoming bits based on the history of received bits. In fact, for data rates above 3 Gbps, the adaptive equalization feature is a key differentiation worth looking for in transceiver vendor offerings. However, it is important to note that there is a clear difference between generic DFEs and adaptive equalizers—a.k.a. electronic dispersion compensation (EDC) and adaptive dispersion compensation engine (ADCE)—and “programmable” DFEs. The key difference is a real-time adaptation. Programmable DFE presumes prior knowledge of each given channel that must be programmed into the DFE coefficients. All others, like ADCE, compute coefficients in real time without any prior channel or data pattern knowledge, and as mentioned previously, above the 3-Gbps data rate, the number of programmable signal conditioning settings is very large.

6.4.3 System Development Stages

For every new system, four separate development stages (shown in Figure 29) can be clearly identified:

1. *Vendor-selection stage*: Does the selected transceiver vendor offer enough signal conditioning for all links in the system?
2. *Board-design stage*: The SI engineer must select the best setting for each link, often as the board layout is initiated, when new settings must be found rapidly.
3. *Board bring-up stage*: Each link must be activated in such a way so as to improve the overall system-level BER. Trade-offs are made between the ability to drive an individual link perfectly versus overall system goodness, once again iterating priority settings.
4. *In-system stage*: Look at the system from a long-term perspective. Specifically, will the link characteristics change with time and environment?

Figure 29. System Development Flowchart



The first key observation is that throughout these four development stages, the designer must select the best setting out of the thousands available. As mentioned previously, one Tx V_{OD} level has over 2700 unique settings and the Rx equalizer has over 1200 unique settings, so the task of choosing the optimal setting can be daunting, if not impractical. For the vendor selection and board design stages, Altera offers a fast-simulation, pre-emphasis equalization link estimator (PELE) that, given the link layout or extracted S-parameters, rapidly (within minutes) converges to a proper solution. However, the “board bring-up” and “in system” stages cannot be addressed via simulation, so hardware is required to automate this process.

To do so, Altera developed the adaptive dispersion compensation engine (ADCE). This hardware automatically and uniquely adapts to each link in the customer’s system. Regardless of the Tx pre-emphasis setting, ADCE selects one of the available Rx CTLE settings. The adaptive engine has both one-time- and continuous-adaptation selections. One-time adaptation is intended for additional power savings, and is enabled during initial link training. However, the customer may enable it periodically either over predetermined time intervals or by observing degradation in system BER over time. In this setting, power is used only during the adaptation process, as once convergence is detected, the CTLE settings are frozen. Continuous adaptation constantly monitors the link in mission mode without disturbing normal traffic. Equalization settings adapt to track not only link aging but also voltage and temperature variations on both side of the link.

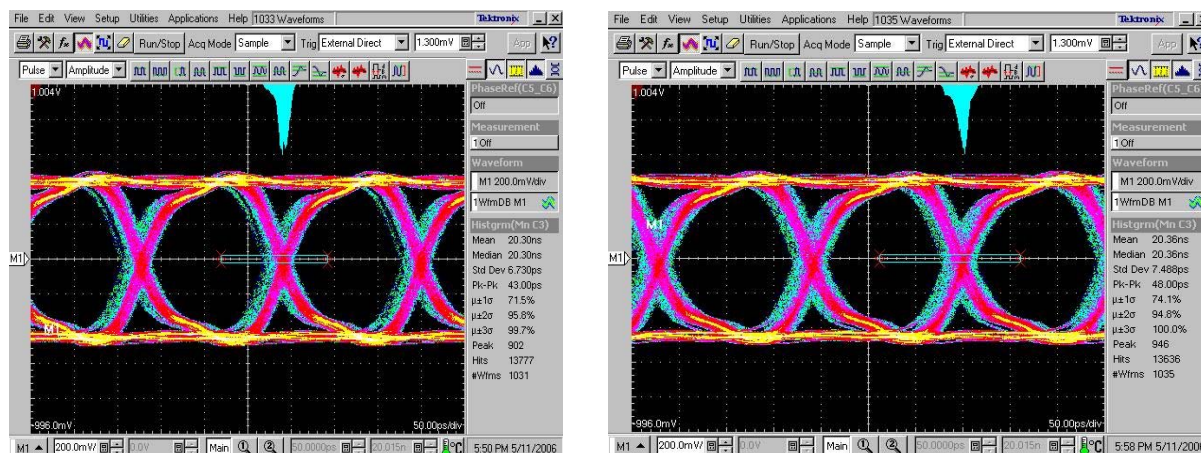
The goal of the ADCE is to make the system easy to use. Each link in the customer’s design has different characteristics. The amount of equalization required varies with backplane length, type, data rate, aging, process, voltage, temperature, etc. The ADCE examines the output of all equalizer stages and adjusts the equalizer to increase or decrease the amount of equalization. Setting and forgetting may not be adequate for changing environments or aging(13).

Blind equalization, which means that no specific training pattern is required, is also one of the goals. However, the data must be DC balanced, and have reasonable transition density and run lengths. The ideal patterns would be similar to PRBS2⁷-1 or PRBS2¹⁰-1 patterns. This usually matches most customer applications very well. For those customers that have sub-optimal data patterns, a training pattern may still be used. In situations where an optimal solution cannot be found, the ability to read out ADCE values is provided inside FPGA. These values are used as a starting point for manual adjustment.

The ADCE is controlled by a flexible state machine that allows changing the adaptation sequence as well as the maximum zero locations in the equalizer. The design also allows monitoring the link in real time during continuous adaptation or randomly enabled on the system level to gather link statistics in the field.

Figure 30 shows two eye diagrams after adaptively reaching a solution. Both eyes are open. The first diagram is for a 12" trace, while the second diagram is for a 33" trace. The second eye diagram would be closed without equalization. Note that both output solutions are roughly the same quality even though the trace length is very different. The eye diagrams are measured after a diagnostic loopback path known to have bandwidth limitations. The actual quality of the eyes inside Rx is even better once measurement bandwidth deterioration is accounted for.

Figure 30. Equalized Eye Diagrams After Adaptive Equalization



FCI N4000 - 13SI
12", Mod4, BGA, Slot4J5K5 to Slot10B5A5

FCI N4000 - 13SI
33", Mod4, BGA, Slot1J5K5 to Slot12B5A5

6.5 A Historical Overview of Altera's Equalization Solutions

AnalogZone says, "The two biggest areas that differentiate [Altera's] Stratix SERDES elements are in their ... transceiver equalization technology they use to achieve better signal integrity, reach, and overall performance under real world conditions."

Altera has demonstrated successful transceiver test chips at 65-nm and 40-nm process nodes, qualifying this newly developed TSMC process for advanced analog needs. The 65-nm test chip was demonstrated to perform above 10 Gbps at DesignCon 2007 and DesignCon 2008. Altera has nearly a decade of SI experience and five different technology nodes of successful transceiver designs.

6.6 Built-In Oscilloscope

In Stratix IV GX FPGAs, the built-in (BI) oscilloscope feature provides in situ signal and jitter measurement capability at various signal nodes within the transceivers. This signal and jitter information cannot be measured by external equipment, yet it is very important for receiver equalization and clock recovery diagnostics and debug, as well as link characterization and verification. The functionality and performance of BI oscilloscope are quite similar to those of an external instrument or a tester, without any additional cost. One of the BI oscilloscope's capabilities is to measure and show the effective eye width of the equalized Rx data. This information not only can 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. Since the BI oscilloscope does not require any special or fixed data patterns, this feature is also very useful for system-level debugging in field and live data traffic watching. In short, the BI oscilloscope solves some emerging test, verification, characterization, debug challenges on the Stratix IV GX FPGA transceiver that are not possible with external test and measurement equipment, and at a cost that is nearly zero or much lower.

7. Advanced Clock and Timing Generation

Clocking and timing generation play important roles in high-speed transceivers. Jitter is an important metric used to measure the quality of a clock, since clock jitter affects both transmitter and receiver jitter performance, in turn causing BER to increase for the link system. (See “8. Power and Jitter” for more details on jitter and noise.) On the transmitter side, the clock jitter limits the eye opening at its output. The clock jitter on the receiver side affects the receiver to latch the receiving data correctly, and consumes a portion of the total available jitter budget of the link (i.e., 1 UI), leaving less jitter budget available 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. A key component of a PLL is the oscillator, which is the major source of jitter. Currently, there are two main types of oscillator used in multiple-GHz PLLs, ring oscillators (ROs) and LC tanks (LCs), each of which has advantages and shortfalls.

7.1 Ring Oscillator

The RO is the most widely used VCO architecture due to its well-known design and performance in clock generation and clock recovery. To achieve high oscillating frequencies, the number of inverter stages must be minimized. Normal RO consists of at least three inverter stages, but some high frequency oscillators have only two stages, which require adding carefully designed coupling for the extra phase shift. These 2-stage and 4-stage ring oscillators are widely used in transceiver designs as they provide quadrature outputs.

The voltage control ring oscillator (VCRO) normally has a wide frequency tuning range (from 10–100 MHz to 1–10 GHz), which enables transceivers to accommodate many different data rates, but also results in very large gains. A high VCRO gain makes the PLL more sensitive to front-end noises and spurs. In addition, the ring oscillator is very sensitive to power supply and substrate noises. The VCRO phase noise or jitter is typically dominated by the supply noise injection if there is lack of a high power supply rejection ratio (PSRR) voltage regulator. Good substrate isolation also improves the phase noise/jitter of ring oscillator.

In addition, the noises coming from switching devices and the biasing circuitry are the main contributions to the phase noise/jitter of a RO in a clean supply and substrate environment. Generally, minimizing the active device count, boosting the current, and removing the biasing circuitry improve the VCRO phase noise. Maintaining a symmetric waveform prevents 1/f noise up-conversion and generally improves phase noise at the low end. However, the jitter improvement may not be due to the PLL shaping on the VCRO phase noise.

There are many different inverter structures—single-end, differential, pseudo-differential, etc.—used in ROs, each with its own advantages and disadvantages. Details such as oscillating frequency, power consumption, and area have to be taken into account in VCO designs. A well-designed RO enables a VCRO to achieve phase noise of $-70 \sim -100$ dBc/Hz at 1 MHz offset, depending on the oscillating frequencies. The phase noise/jitter of a ring oscillator degrades at higher frequencies. For example, a 3-GHz CMOS VCRO manufactured at the 65-nm node achieves a phase noise of -91 dBc/Hz at 1 MHz and a RMS jitter (1–80 MHz) at 1.1 ps. However, a 6-GHz VCRO at the same process node has a phase noise of -86 dBc/Hz at 1 MHz and a RMS jitter (1–80 MHz) at 1.24 ps.

7.2 LC Oscillator

LC oscillator offers a superior phase noise performance due to its highly selective and high Q LC tank. Discrete LC oscillators have been existed in RF applications for long time, but integrating LC oscillators in mixed-signal IC became common only in recent years. There are two main factors driving LC oscillators to be used in integrated transceiver designs. The first is that RO phase noise has difficulties meeting the transmitter jitter requirements at a multi-GHz frequency range. The second is that, due to the process's feature size shrinkage, inductors are getting small enough to be integrated on a die as LC oscillator frequency increases.

The cross-coupled LC oscillator is the most widely used architecture. The LC tank oscillates in either current mode or voltage mode. In the current mode, the signal amplitude is determined by a biased tail current. Although the waveform is immune to the supply variation, the tail current 1/f noise up-conversion degrades its phase noise. In the

voltage mode, there is no tail current and the amplitude is limited by supply voltage. A voltage regulator prevents supply noise and spur injection.

Comparing with VCRO, VCLCO occupies more die area due to its bulky inductors. Unlike VCRO, integrated VCLCO has a very limited frequency-tuning range. The only tunable element of an on-chip LC oscillator is the varactor. The tuning range of either diode or MOS varactor is very limited due to the current process technologies. The continuous frequency tuning range of a VCLCO is typically around 20 percent. However, a VCLCO normally offers much better phase noise/jitter performance than a VCRO. For example, in 65 nm, a 6-GHz VCLCO has a phase noise of -110 dBc/Hz at 1-MHz offset and a RMS jitter of 100 fs (1–80 MHz).

8. Power and Jitter

Throughout the literature, careful partition between digital and analog domains is emphasized to provide the isolation necessary to meet the stringent analog jitter requirements.

8.1 Power Integrity

In the Stratix IV GX FPGAs' transceiver, the high-speed analog sections of receive and transmit paths are separated. This is because the transceiver allows completely independent frequency selection on Rx path from that chosen for Tx path. Separated power prevents injection of uncorrelated noise sources. In Figure 31, these power supplies are identified as V_{CCET} and V_{CCER} . Furthermore, the clock path has its own power supply (V_{CCEL} in Figure 32), which is separated out to prevent noise injection from the data path into the transmitter clocking.

Figure 31. Power Domain Scheme in the Transceiver

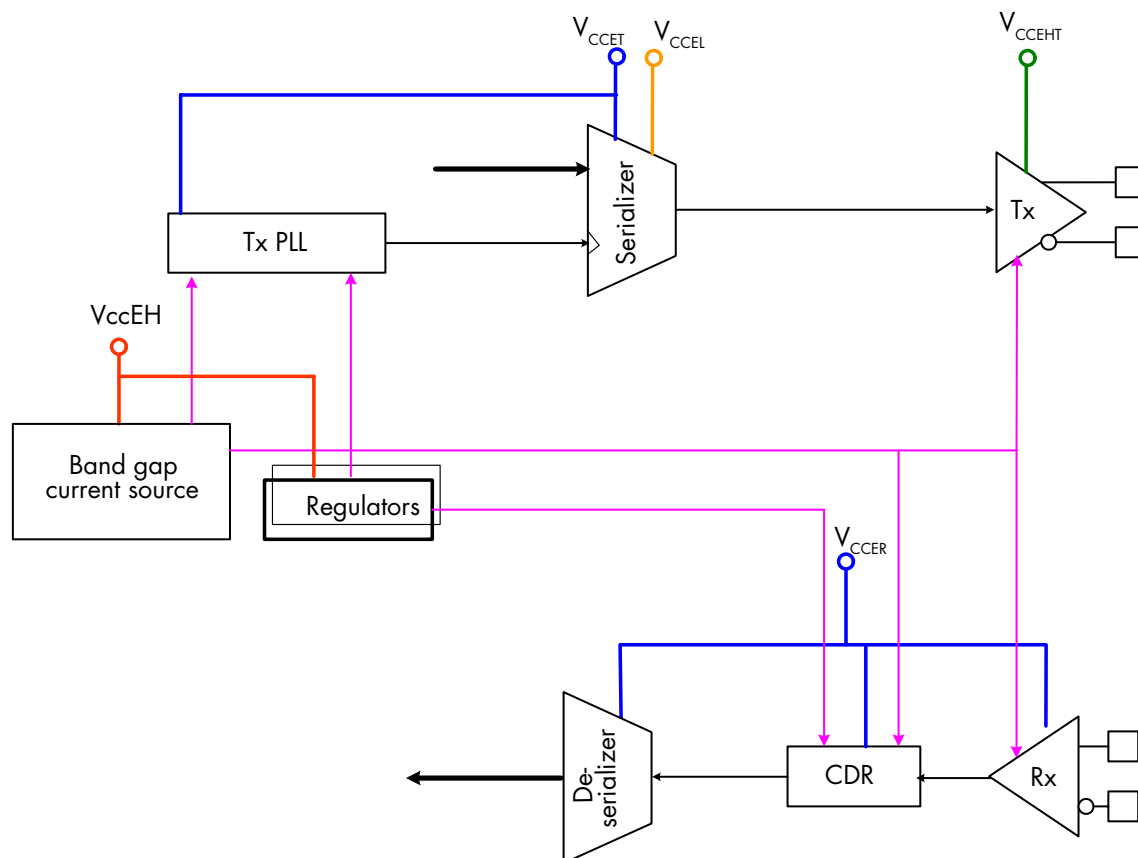
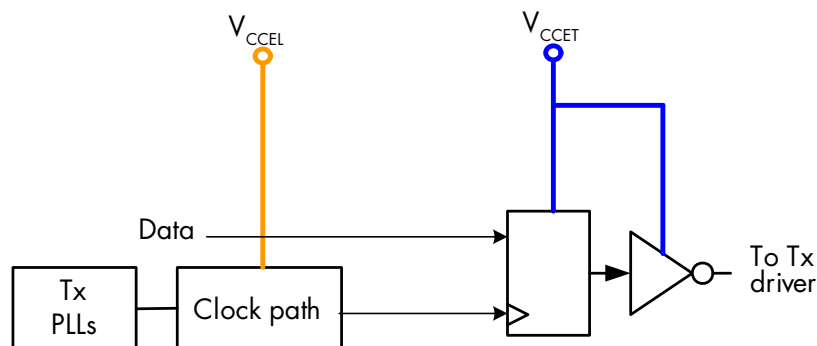


Figure 32. Power for Clock and Data Path in the Transmitter Path



Precision analog blocks, such as bang gap, current bias, and on-chip voltage regulators, receive power from one dedicated V_{CCEH} power supply. One chip voltage regulator is intended to isolate out sensitive circuitry of each Tx and Rx PLL, such as VCO, charge pump and loop filter, as shown in Figure 33. The Tx driver in Figure 31 has its own power supply, V_{CCEHTS} which provides different power levels. Both on-chip and on package de-coupling is used in the transceivers to provide necessary noise filtering from the external power supplies. Figure 34 shows the PSRR for the regulator-supporting CP and LP, and Figure 35 shows the regulator for VCO. Note that the PSRRs are below -50 dB in both cases for frequencies greater than 1 GHz.

Figure 33. Regulator for VCO/CP/LF in the Tx PLL and CDR

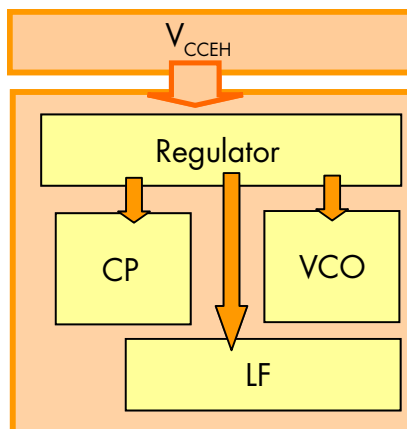


Figure 34. PSRR for Charge Pump Regulator and Loop Filter

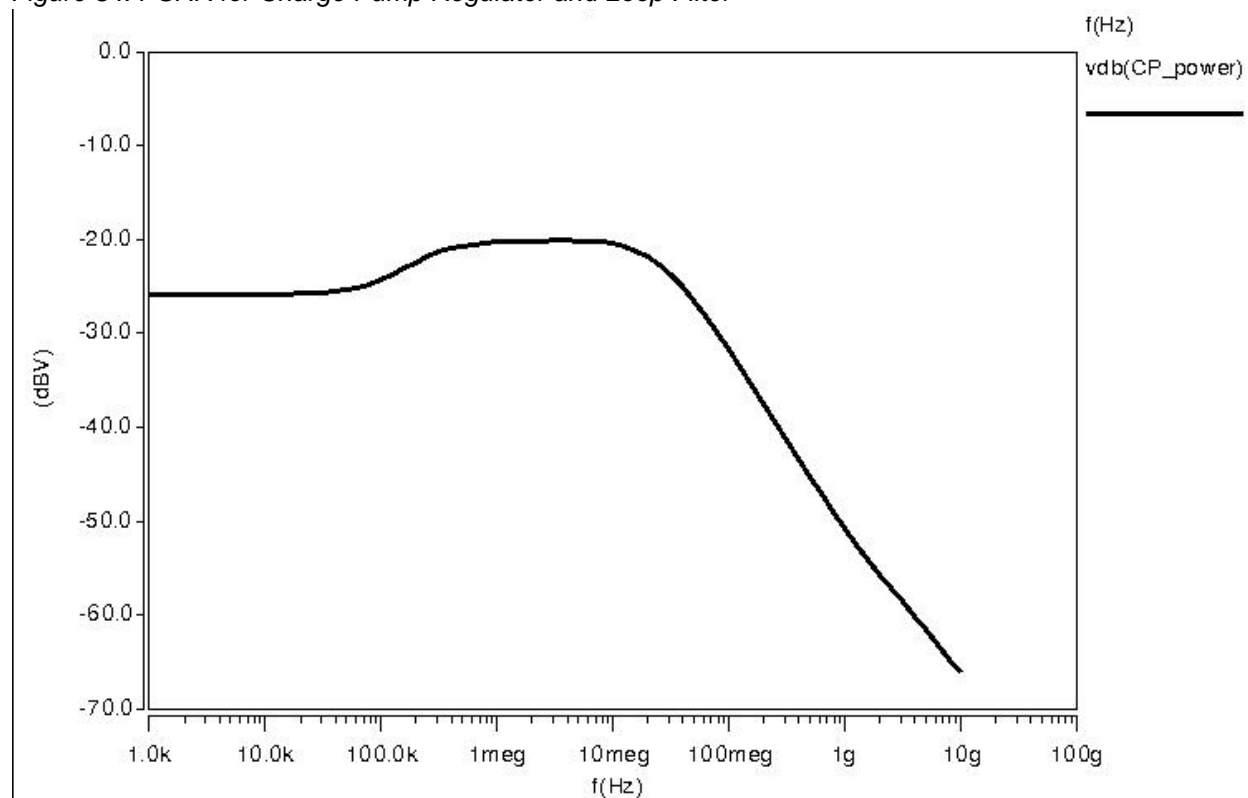
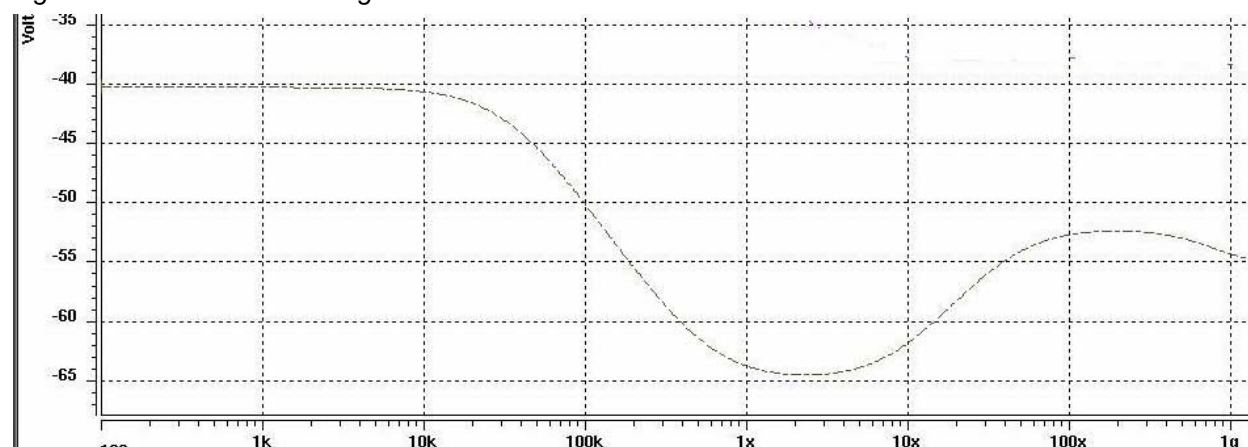


Figure 35. PSRR for VCO Regulator

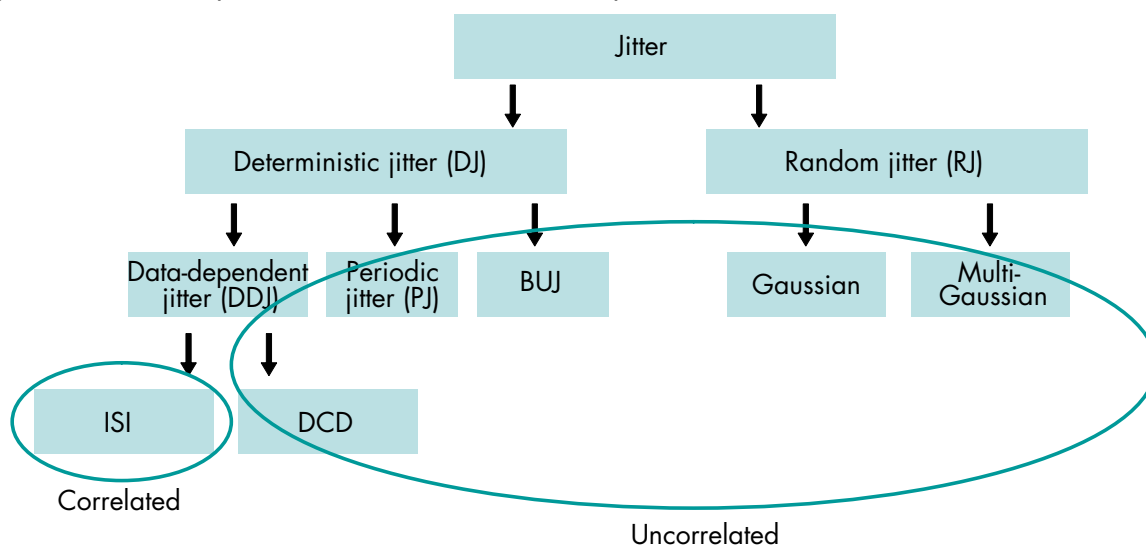


8.2 Jitter and Noise Basics

Two important metrics for quantifying the performance of a transceiver are jitter and noise. Jitter is commonly defined as any deviation from ideal timings such as an ideal bit clock; and noise is commonly defined as any deviation from a reference voltage or power. Both excessive jitter and/or excessive noise increase the BER for the link system.

Jitter can be separated into deterministic jitter (DJ), which is bounded, and random jitter (RJ), which is unbounded. A jitter component family tree hierarchy is shown in Figure 36.

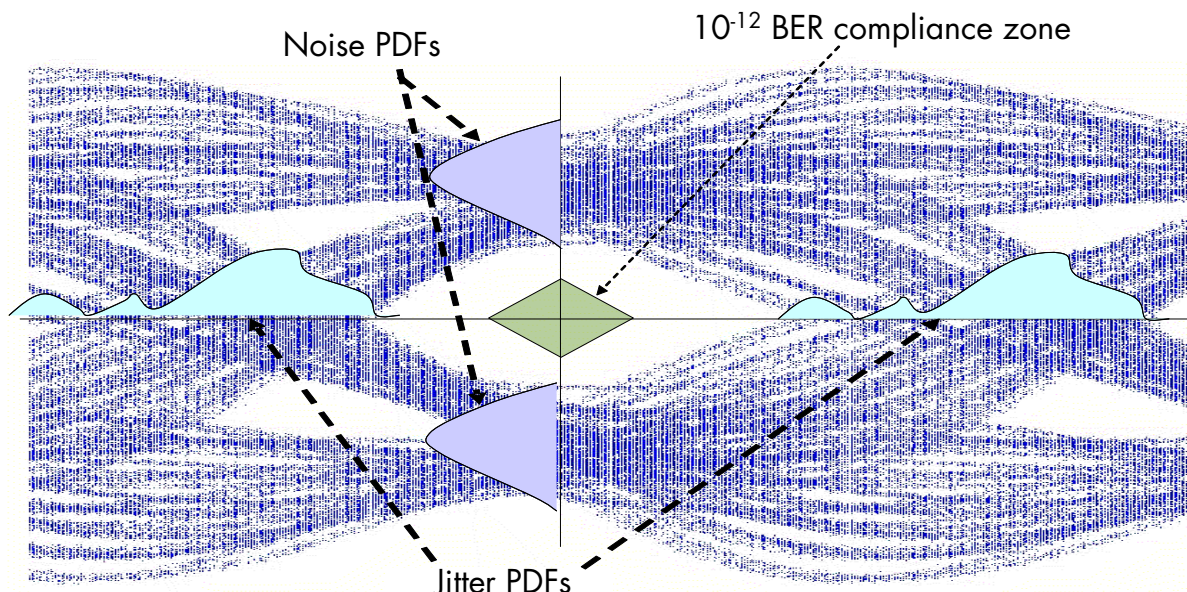
Figure 36. Jitter Components and Their Interrelationships



DJ may include components of data-dependent jitter (DDJ), periodic jitter (PJ), and bounded-uncorrelated jitter (BUJ). DDJ is typically caused by band limiting effects such as lossy channel. PJ is caused periodic modulation such as switch power-supply coupling, and BUJ is caused by crosstalk. DDJ may include components of ISI and duty cycle distortion (DCD). ISI and DCD are the consequences of band-limiting effect. However, DCD may also be caused by shifting of reference voltage. RJ is commonly caused by thermal noise and its distribution is best described by a Gaussian distribution.

A similar separation concept also can be applied to noise. The statistical properties for jitter and noise are described by a probability density function (PDF) best illustrated by an eye diagram, as shown in Figure 37.

Figure 37. Eye Diagram and its Relationship to Jitter PDFs and Noise PDFs



Excessive jitter and/or noise increases the BER, in turn enhancing the chance of data samples falling within the compliance BER zone, resulting in a link failure. Thus, maintaining good jitter and noise performance is the key to achieve good BER performance for a link.

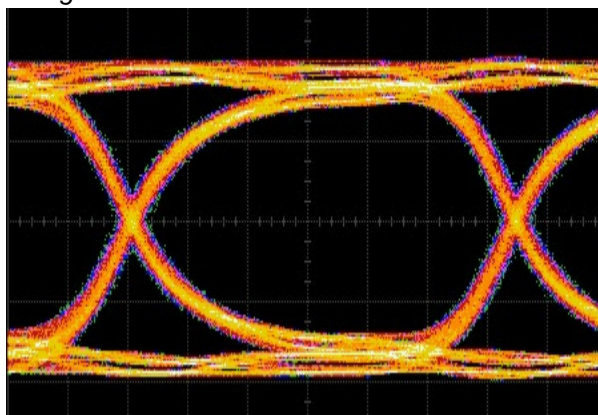
More information on jitter and noise fundamentals can be found in (14).

8.3 Jitter and Noise Generation

A good transmitter is expected to generate minimum jitter and noise, which means that the eye diagram measured at the transmitter output should be wide open. Many high-speed I/O standards (e.g., PCIe, CEI/OIF, 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. If no measurement data fall within the eye mask, then the transmitter jitter, noise, and signaling pass the requirements and the transmitter performance is assured.

Transmitter jitter is largely depends on the jitter of its PLL that uses a VCO. Due to the advanced design and the use of a LC-based VCO, the Stratix IV GX jitter and noise performance is much improved. Figure 38 shows an eye diagram and associated DJ, RJ, and TJ (at $\text{BER} = 10^{-12}$) for a Stratix IV GX test chip, running at 8.5 Gbps, with a PRBS²¹⁵⁻¹ pattern.

Figure 38. “Wide-Open” Eye Diagram Measured From a Stratix IV GX Transmitter

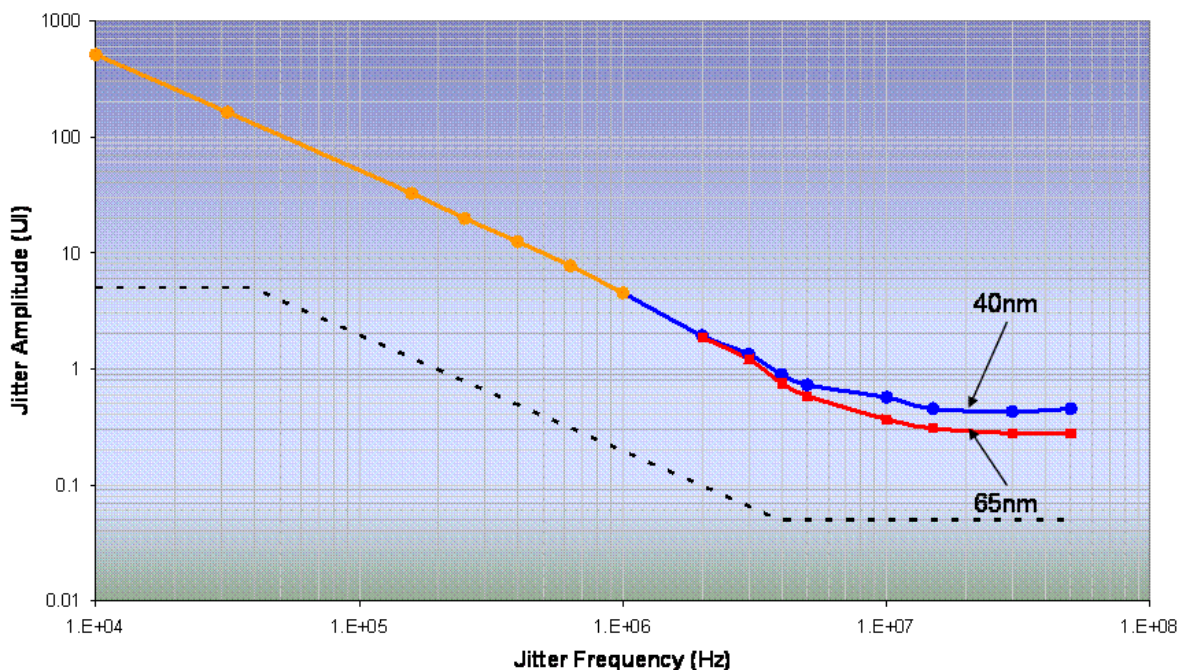


8.4 Jitter and Noise Tolerance

A good transmitter should generate a minimum amount of jitter and noise. In contrast, 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 clock recovery and the equalization, such as CLTE and/or DFE. To verify whether a receiver CRC has the required jitter-tracking capability, a jitter frequency mask or tolerance mask is often defined by a standard. 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 it is required by a standard.

Stratix IV GX FPGAs, due to their innovative hybrid CRC design, have superior jitter tolerance capability. Figure 39 shows the jitter tolerance performance for a Stratix IV GX receiver. The results should not only meet the receiver needs, but also exceeds the specifications. If other parts of the link subsystem also meet the specification requirements, a receiver better than the specification implies that links built with Stratix IV GX FPGAs have a better (lower) BER than 10^{-12} .

Figure 39. Jitter Frequency Tolerance (1) Measured From a Stratix IV GX Receiver

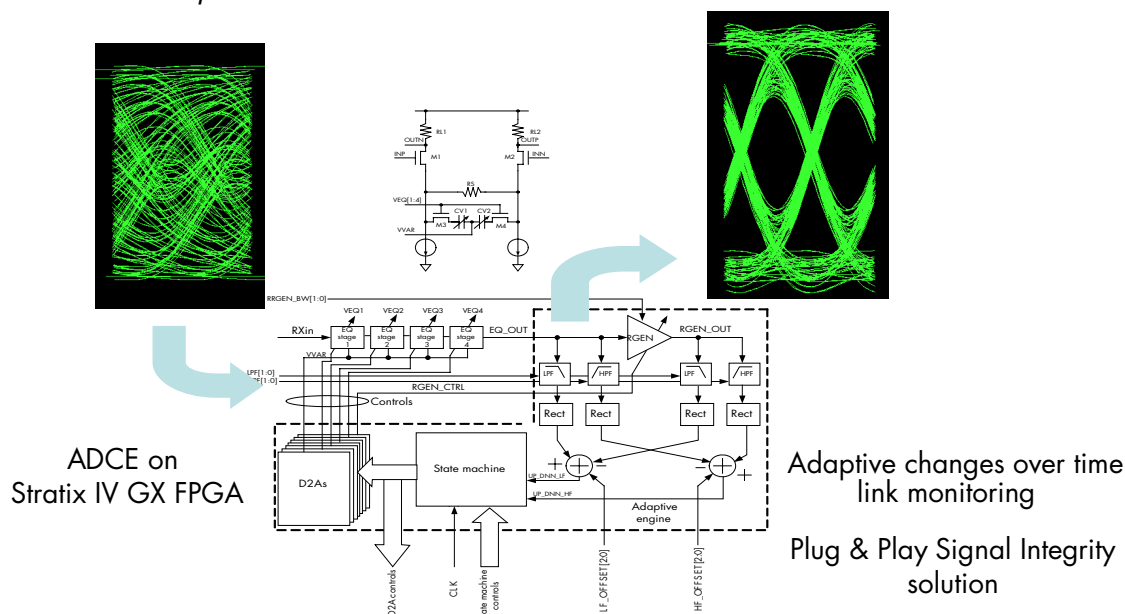


Note:

(1) The jitter tolerance level exceeds the standard requirement.

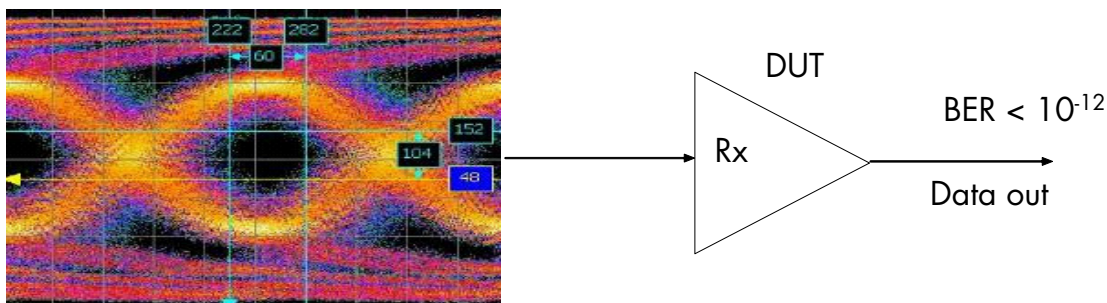
As shown in Figure 40, Stratix IV GX FPGAs have complete sets of end-to-end equalizations built-in. In addition to transmitter pre-emphasis equalization, the receiver has CTLE, ADCE, as well as DFE, with the most powerful equalization capabilities and covering a wide range of channel materials and reaches. An importance test aspect for receiver equalization is to verify that the receiver is capable of opening the closed eye caused by the lossy channel and meeting the required BER of 10^{-12} or better imposed by a specific standard.

Figure 40. End-to-End Equalization for Stratix IV GX FPGA Transceivers



The third aspect of the receiver verification is to insure that the receiver operates under the worst-case jitter, noise, and signal conditions and is still able to achieve BER of 10^{-12} or lower. This is a receiver system tolerance test, while jitter frequency mask and ISI mask are receiver component verifications. With both subsystem and system verifications, the receiver test fault coverage is 100%. Stratix IV GX FPGAs operate under the worst-case input jitter, noise, and signaling conditions, and achieve a BER of 10^{-12} or better results. Figure 41 shows the test example.

Figure 41. Worst-Case Jitter, Noise, and Signaling Test Setup for the Stratix IV GX Receiver

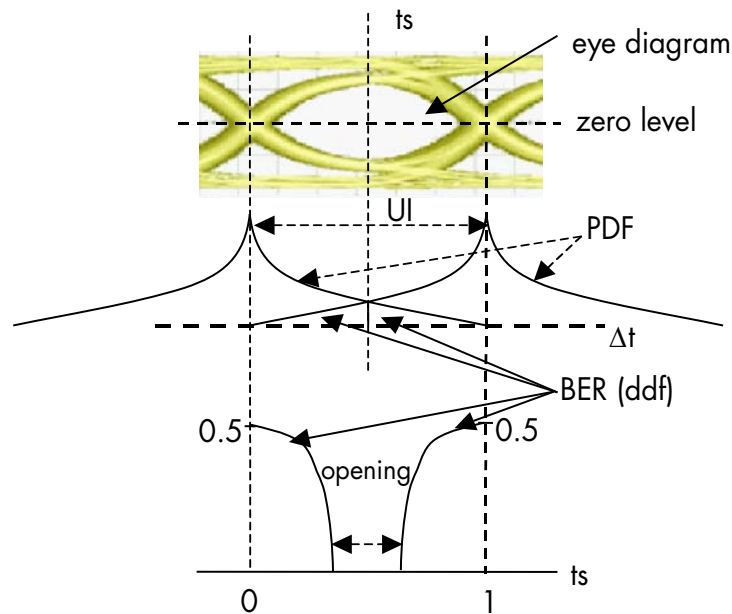


8.5 Overall System BER

BER is a system metric for a link. Thus, BER is tightly coupled with link system architecture and subsystem performances. A quantitative and accurate system model and understanding is essential for estimating or measuring the system BER. The BER for a subsystem may also be defined if the rest of the characteristics of the link system are defined or assumed (15).

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). For simplicity, the BER is often estimated and viewed as a function of sampling time at a given reference voltage (e.g., at zero crossing or 50 percent swing level), or as a function of sampling voltage at a given reference time location (e.g., at the center of the UI data cell). Figure 42 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 42. Integrated and Correlated View for Eye-Diagram, Jitter PDF, and BER CDF



The overall link system BER depends on the jitter, noise, and signaling performances of its subsystems, such as transmitter, channel, and receiver. Due to Altera's innovative designs in system and subsystem architecture, clock recovery, equalization, power and power integrity, the use of the latest 40-nm process technology for fabrication, and its jitter, noise, and power consumption, the transceiver meets or exceeds the requirements posed by various high-speed I/O standards. Due to the excess margins in jitter and noise, it also enables its users to design systems with lower cost and/or better BER system performance (e.g., 10^{-15}).

9. Conclusion

In reviewing the technology trends and associated challenging requirements for high-speed links and transceivers, the challenges range from data rate increase, process node decrease, and optimal power consumption, to stringent performance requirements for jitter, noise, power integrity, and BER, and supportability of various multiple-Gbps high-speed I/O standards. Altera's Stratix IV GX FPGAs enhanced with transceivers meet and/or exceed those challenges and requirements using 40-nm design and fabrication, innovative hybrid transceiver architecture, hybrid/mixed-signal clock recovery, complete end-to-end equalizations, and design and use of an ultra-low noise/jitter LC-based oscillator, and BI oscilloscope. The 40-nm process technology allows Stratix IV GX FPGAs to achieve the best possible logical density, memory speed, and capacity, while the transceiver innovations enable superior jitter, noise, signal integrity, and BER performances at the minimum or optimized power.

Stratix IV GX FPGAs deliver the highest density, the highest performance, and the lowest power. Leveraging 40-nm benefits and proven transceiver and memory interface technology, Stratix IV GX FPGAs provide an unprecedented level of system bandwidth with superior signal integrity. Stratix IV GX FPGAs, when coupled with HardCopy IV ASICs, provide the benefits of FPGAs and ASICs using seamless prototyping. Quartus® II design software enables success with the highest logic utilization and the shortest compile times in the industry. Take advantage of Altera's 40-nm portfolio and innovate your products without compromise.

New circuits and IP are added to enable Stratix IV GX FPGAs to support emerging standards of PCIe 2.0, HT 3.0, Interlaken, CPRI, and SFI-5. The general purpose, scalable, and programmable design permits Stratix IV GX FPGAs to support most of the high-speed link standards.

In summary, Stratix IV GX FPGAs enhanced with transceivers are the best in class in terms of speed, performance, capability, power consumption, state-of-art design, wide range of standard supportability, design and fabrication with the 40-nm process node, and built-in-self-test capability.

References

1. ITRS roadmap, 2007 Edition,
www.itrs.net/Links/2007ITRS/Home2007.htm
2. PCIe Standard: PCI-SIG,
www.pcisig.com/specifications/pciexpress
3. Hyper Transport Standard:
www.hypertransport.org/index.cfm
4. TSMC Advanced Technology Roadmap:
www.tsmc.com/download/english/a05_literature/2_Advanced_Technology_Overview_Brochure_2007.pdf
5. 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.
6. M. Li, "Jitter and Signaling Test for High-Speed Links," An Invited Paper, CICC/IEEE, 2006.
7. Stratix IV GX and HardCopy IV GX Transceiver Overview:
www.altera.com/products/devices/stratix-fpgas/stratix-iv/transceivers/stxiv-transceivers.html
8. PCI Express 1.1 Base Specification, 2005:
www.pcisig.com/specifications/pciexpress/base
9. PCI Express 2.0 Base Specification, 2007:
www.pcisig.com/specifications/pciexpress/base
10. B. Razavi (Editor), "Monolithic Phase-Locked Loops and Clock Recovery Circuits: Theory and Design," Wiley-IEEE Press, 1996
11. V. Stojanovic, M. Horowitz, "Modeling and Analysis of High-Speed Links," Stanford University lecture.
12. B. Casper, P. Pupalaiakis, J. Zerbe, "Serial Data Equalization," DesignCon 2007.
13. 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.
14. M. Li, *Jitter, Noise, and Signal Integrity at High-Speed*, Prentice Hall, ISBN-10: 0132429616, 2007.
15. 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.

Further Information

- *Leveraging the 40-nm Process Node to Deliver the World's Most Advanced Custom Logic Devices:*
www.altera.com/literature/wp/wp-01058-stratix-iv-40nm-process-node-custom-logic-devices.pdf
- *40-nm Power Management and Advantages:*
www.altera.com/literature/wp/wp-01059-stratix-iv-40nm-power-management.pdf
- *Increasing Productivity With Quartus II Incremental Compilation:*
www.altera.com/literature/wp/wp-01062-quartus-ii-increasing-productivity-incremental-compilation.pdf

Acknowledgements

- Weichi Ding, Design Manager, Analog Design Group, Altera Corporation
- Bernhard Friebe, Product Marketing Manager, High-End FPGA Products, Altera Corporation
- Tim Hoang, IC Design Manager, Analog Design Group, Altera Corporation
- Mike Peng Li, Ph.D., Principle Architect/Distinguished Engineer, Product Engineering, Altera Corporation
- Sergey Shumareyev, Director of Engineering, Analog Design Group, Altera Corporation
- Tina Tran, Senior Design Manager, Analog Design Group, Altera Corporation
- Wilson Wong, Principle Design Engineer, Analog Design Group, Altera Corporation
- Zhi Wong, Sr. Technical Marketing Manager, Product Marketing, Altera Corporation



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

Copyright © 2008 Altera Corporation. All rights reserved. Altera, The Programmable Solutions Company, the stylized Altera logo, specific device designations, and all other words and logos that are identified as trademarks and/or service marks are, unless noted otherwise, the trademarks and service marks of Altera Corporation in the U.S. and other countries. All other product or service names are the property of their respective holders. Altera products are protected under numerous U.S. and foreign patents and pending applications, maskwork rights, and copyrights. Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera Corporation. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.