

Emerging Standards at ~10 Gbps for Wireline Communications and Associated Integrated Circuit Design and Validation

(An Invited Paper for CICC)

Mike Peng Li and Sergey Shumarayev
Altera Corporation
101 Innovation Road
San Jose, CA 95134

Abstract-We first review the signaling and jitter requirements for emerging high-speed wireline communication standards at ~10 Gbps, including CEI 11G, XLAUI/CAUI, XFI, and SFP+. We then present a FPGA transceiver architecture and subsystem/circuit blocks for clocking and timing generation, a transmitter buffer, and receiver CDR and DFE, all designed and manufactured with 40-nm process node. Lastly, we present the signal/jitter transmitter output and receiver-tolerance measurement results at 10.3125 Gbps, with an ultra-low random jitter at ~550 fs.

I. INTRODUCTION

Semiconductor technology advancements largely follow Moore's Law in that the number of transistors in integrated circuits doubles roughly every two years. Moore's Law is facilitated by the feature size or process node shrinkage. Smaller features enable more functionality and integration, higher operation speed and logic density, and lower power consumption per logic function. As the device functionality and logic density increase, the I/O bandwidth must increase proportionally, which is achieved by increasing the data rate for each I/O lane.

A higher data rate often is achieved through using advanced design methodologies and process technologies. 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. 40 nm is the leading-edge process node used for some advanced, high-density, and high-performance devices, including microprocessors, FPGAs, and graphic processors.

At the present, the data rates for most high-speed transceivers are in the range of 5–6 Gbps for communication and I/O standards. A few examples 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. The leading-edge and/or next-generation standards—IEEE 10G Ethernet (802.3ae), IEEE 10G Ethernet for backplane (802.3ap), IEEE 40G/100G Ethernet (802.3ba, 4x10.3 Gbps/10x10.3125 Gbps), CEI/OIF 11G for network communication, PCIe 3.0 (8.0 Gbps) for computer I/O buses, and Fibre Channel 8x (8.5 Gbps) for storage area networks—have data rates in the neighborhood of 8–11 Gbps per lane with aggregations.

To meet the data-rate increase requirements imposed by high-speed standards, transmitter (Tx) and receiver (Rx) of the link need to be advanced so that they can send and receive data at a target bit error rate (BER), commonly set at 10^{-12} or lower. However, as the data rate increases, the signal distortion, or inter-symbol interference (ISI) caused by the lossy channel gets worse, which degrades the BER of the link if it is not reduced or compensated. There are two basic ways to address the copper channel caused ISI at higher data rates: one is to add the equalization circuit at Tx, Rx, or both, the other is to design or choose a better channel with less loss and crosstalk (if the link has

multiple channels) properties. For cost reasons, the industry and standards have chosen the first approach to deliver higher data rates with the same or slightly better channel material in recent years. This approach will be the case for links with data rates up to 10 Gbps and higher in the near future.

At data rates greater than 1 Gbps, most high-speed links use serial asynchronous architecture to send only data from the Tx and recover the clock and data at the Rx side via a clock and data recovery (CDR) unit. The clock recovery circuit (CRC) commonly is implemented with a phase-locked loop (PLL), although phase interpolator (PI) clock recovery is also used in computer I/Os such as PCIe. In addition to recover clock and data, CDR also tracks or reduces the low-frequency jitter that otherwise degrades the BER performance of the link. As data rates increase, the CDR bandwidth must increase proportionally.

As the data rate increases, the unit interval (UI) of the data bit stream decreases proportionally. To maintain the same BER, the Tx and Rx jitter must be reduced proportionally, assuming that the equalization can compensate ISI increases for the same channel. Circuit functionality, performance, and power requirements pose new and outstanding design and verification challenges for higher data rates (e.g., 10 Gbps). Designing manufacturable and general-purpose transceiver ICs at 10 Gbps and higher with multiple channels, while supporting and complying multiple standards, is an even harder and more challenging task, and this focus of this paper.

We structure this paper as follows: In Section 2, we start with reviewing the leading-edge ~10-Gbps wireline communication standards with a focus on link architecture and physical layer signaling/jitter. Electrical requirements, such as output voltages, rise/fall times, compliance eye mask/jitter, for Tx; voltage sensitivity, worst-case eye/jitter tolerance, and CDR jitter tracking for Rx are also covered. As we are unable to cover all of the 10-Gbps specifications, we focus on the common and emerging standards at ~10 Gbps or higher, including electrical specifications of link standards of CEI/OIF (11.1 G) (chip-to-chip)[1], XLAUI/CAUI (10.3125 G) (chip-to-chip and chip-to-module (optical))[2], and XFI (11G)[3] and SFP+ (to 11G) [4] (chip-to-module interfaces).

In Section 3, we review an FPGA general-purpose transceiver architecture and IC technologies that meet, or in some cases exceed, the standard requirements, with a focus on the latest and advanced 40-nm process node. Transceiver and link architectures, ICs for equalizations, voltage control oscillators (VCOs), clock generation, and CDR are reviewed. In Section 4, transceiver IC performance/compliance measurement results, including signaling/jitter output for Tx and signaling/jitter tolerance for Rx, are presented. We summarize and conclude in Section 5, along with some forward-looking discussion.

II. HIGH-SPEED I/O DATA RATE TREND and STANDARD REQUIREMENTS

In this section, we first review the high-speed I/O data rate and link bandwidth trend. Then we move on to discuss the signaling/jitter and electrical requirements of high-speed I/O standards at ~10 Gbps.

A. Data rate trends for high-speed I/Os

As shown in Fig. 1, the International Technology Roadmap for Semiconductors (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 steadily 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 interface data rates double every two to three years, and the leading edge rates are at ~10 Gbps in 2009. The trend line prediction accurately reflects both the reality of high-speed I/O devices currently 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. The next-generation I/O devices for most high-speed standards are expected to maintain similar data-rate growth to that of previous generations. Specifically, PCIe 3.0 (8.0 Gbps), FC 8x (8.5 Gbps), IEEE Ethernet 40G/100G, and ITU OTN 40G/100G will emerge in the 2009–2010 time frame

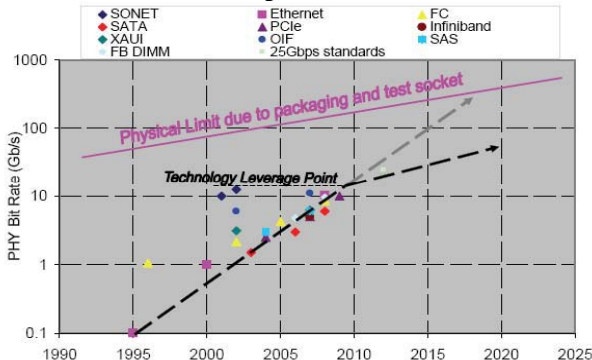


Fig. 1. Mean Data Rate for High-Speed I/O Links From ITRS 2007 [5]

The data rate increases pose challenges for the designing and manufacturing high-speed transceivers such as signaling and jitter. In the next subsection, we discuss the requirement advancements of high-speed standards in response to these challenges.

B. Signaling/jitter for ~10 Gbps standards

We first review a typical link architecture at 10 Gbps. Then we move on to the signaling/jitter for Tx and Rx.

B.1 Link system overview




Fig. 2. A High-Speed Link System

Fig. 2 shows a typical link architecture at 10 Gbps. Within Tx, a parallel-in-serial-out (PISO) or serializer converts multiple parallel lower data-rate lanes to a single higher data-rate lane. A linear equalization (LEQ), such as a finite impulse response (FIR) is commonly used for Tx equalization. A clock generator, often implemented with PLL circuits, provides the PISO and equalization timing, while a driver provides the necessary voltage level and associated rise and fall times. Tx jitter largely depends on the clock and PLL jitter, especially random jitter (RJ). In addition to the bandwidth of the Tx driver, Tx package and parasitic are also sources of ISI, due to their lossy properties. Reflection-induced jitter between silicon driver and package, and between package and channel, can also occur if the impedance between them is not matched. Due to capacitive or inductive coupling, crosstalk can also occur on integrated circuits on-die or on the package. For the channel, its insertion loss causes ISI for the signal at the channel output or Rx input. Moreover, crosstalk due to the channel coupling, as well as reflection due to the impedance mismatch between the channel and Tx/Rx can further degrade the signal quality at the Rx input. Within the Rx, the incoming signal is first detected by its buffer. At around 10 Gbps, the loss of the channel at Nyquist can be as high as -35 dB at certain longer channels (e.g., backplanes), and the Tx LEQs are unable to compensate for all of the ISI caused by the channel loss. Therefore, Rx equalization is also required. A LEQ, followed by decision feedback equalization (DFE), is the common equalization topology for Rx equalization. Rx DFE can be adaptive without needing to create a back channel, so is suitable for compensating the variation or drifting of channel loss properties. However, DFE can only compensate the postcursor ISI. As such, a LEQ, such as continuous time linear equalization (CTLE) or feed-forward equalization (FFE), is needed to work with DFE to compensate both precursor and postcursor ISIs. Clock recovery commonly takes the data signal from the LEQ and recovers the clock from it. The recovered clock then is used for DFE and data-sample timing.

To insure the interoperability and $BER=10^{-12}$, signaling and jitter must be specified at the output of the Tx and the input of the Rx, with the assumptions on the channel electrical properties. For the channel, the reference Tx and Rx with the minimum equalization and clock recovery capabilities must be assumed to quantify the time-domain ISI and jitter contribution and to assess whether it is complied or not. This is the general trend in specifying the signal and jitter, as well as the electrical properties for high-speed I/O standards at ~10 Gbps.

B.2 Tx signaling and jitter requirements

Tx signaling and jitter commonly is specified via eye mask, as shown in Fig. 3.




Fig. 3. A Common Tx Signaling/Jitter Compliance Mask

For the y-axis, A1 and -A1 define the minimum pk-pk differential voltage, while A2 and -A2 define the maximum differential voltage. For the x-axis, and for a given hexagon compliance mask, a smaller X1 and/or X2 imply a relaxed eye mask. The mask is intended to correspond to a target BER (e.g., 10^{-12} , or 10^{-15}) boundary. Furthermore, 2x1 represents the total jitter (TJ) allowed for the transmitter, corresponding to the same mask BER. The mask also sets the maximum values for the rise/fall time for the transmitter.

In addition to TJ, deterministic jitter (DJ) and/or RJ also are specified. DJ commonly is assumed to have a dual Dirac probability density function (PDF), and RJ is assumed to have a Gaussian PDF (see [6] for a recent review on jitter). With the DJ and RJ components, the link jitter budget gets some statistical relief (e.g., via root-square sum (RSS)), leaving a maximum jitter margin for the link subsystems (e.g., Tx, Rx, and channel).

In Table I, we list the eye-mask and jitter parameters for the ~10 Gbps high-speed standards discussed in Section 1.

Table I: Tx Signaling and Jitter Requirements

	CEI 11GSR	CEI 11G MR/LR	XLAUI/CAUI	XFI	SFP+
BER	10^{-15}	10^{-15}	10^{-12}	10^{-12}	10^{-12}
Data Rate (Gbps)	9.95–11.1	9.95–11.1	10.3125	9.95–11.1	up to 11.1
A1 (mv)	180	400	200	180	95
A2 (mv)	385	600	380	385	350
X1 (UI)	0.15	0.15	0.16	0.15	0.12
X2 (UI)	0.40	0.40	0.38	0.4	0.33
TJ (UI)	0.30	0.30	0.32	0.30	0.28
RJ (UI)	0.15	0.15	-	-	-
DJ (UI)	-	-	0.17	0.15	-
DDJ(UI)	-	-	-	-	0.1
UBHPJ (UI)	0.15	0.15	-	-	0.023
PWJ	-	-	-	-	0.055
DCD (UI)	-	0.05	-	-	-

DDJ: data dependent jitter
 UBHPJ: uncorrelated bounded high-probability jitter
 PWJ: pulse width jitter
 DCD: duty-cycle distortion

B.3 Rx signaling and jitter requirements

Due to channel loss, the signal reaching the receiver side is distorted. A square wave may become a sinusoidal wave. To prevent this, a diamond-shaped eye mask is used for Rx worst case signaling/jitter compliance for many standards. In some applications, such as backplane, the significant loss due to the channel causes the signal eye to be closed completely at the Rx input. Fig. 4 shows the eye mask at the receiver input for most of the 10-Gbps I/O interfaces, except for backplanes or very lossy channels.




Fig. 4. A Common Tx Signaling/Jitter Compliance Mask

For the y-axis, A1 and -A1 define the minimum pk-pk differential voltage needed for the Rx to comply. For the x-axis, 2x1 represents the TJ that the receiver must tolerate. The mask corresponds to a target BER (e.g., 10^{-12} , or 10^{-15}) boundary.

The Rx must tolerate the worst-case signal distortion/jitter due to both Tx and channel. Worst-case jitter may include various types of jitter. The Rx responses to various jitter inputs are different, depending on the kind of subsystems included. For example, equalization only compensates correlated jitter, thus the receiver's response to correlated and uncorrelated jitter is different. Correlated jitter stresses the equalization, while uncorrelated jitter stresses the setup and hold timing margins of the Rx sampler. In Table II, we list the eye-mask and jitter parameters for Rx.

Table II: Rx Signaling and Jitter Requirements

	CEI 11G SR	CEI 11G MR	CEI 11G LR	XLAUI/CAUI	XFI	SFP+
BER	10^{-15}	10^{-15}	10^{-12}	10^{-12}	10^{-12}	10^{-12}
Data Rate (Gbps)	9.95–11.1	9.95–11.1	-	10.3125	9.95–11.1	up to 11.1
A1 (mv)	55	55	-	42.5	55	150
A2 (mv)	525	525	600	425	525	425
X1 (UI)	0.35	0.35	-	0.31	0.325	0.35
X1' (UI)	0.25	0.25	-	-	-	-
TJ (UI)	0.7	0.7	-	0.62	0.65	0.70
TJ' (UI)	0.5	0.5	-	-	-	-
RJ (UI)	0.20	0.20	-	-	-	-
DJ (UI)	-	-	-	0.42	-	-
UBHPJ (UI)	0.25	0.25	-	-	-	-
CBHPJ (UI)	0.20	0.20	-	-	-	-
NEQJ	-	-	-	-	0.45	-
J2	-	-	-	-	-	0.42
PWJ	-	-	-	-	-	0.3
SJ Max (UI)	5	5	5	-	-	-
SJ HF (UI)	0.05	0.05	0.05	-	-	-

UBHPJ: uncorrelated bounded high-probability jitter, CBHPJ: correlated bounded high-probability jitter, SJ Max: maximum sinusoidal jitter (SJ), SJ HF: high-frequency SJ, X1': exclude CBHPJ, TJ': exclude CBHPJ, NEQJ: non-equalizable jitter

In addition to signal and jitter tolerance, a receiver also must have a compliance jitter-tolerance frequency mask designed to stress the receiver clock recovery jitter tracking capability using SJ frequency sweeping. Fig. 5 shows a typical jitter-tolerance frequency mask.




Fig. 5. A Typical SJ Frequency Mask for Rx Clock Recovery Jitter Tracking Stressing

The corner frequency f_c is commonly set within the range of data rate/1667. This frequency corresponds to the minimum Rx clock-recovery corner frequency. For the 10-Gbps I/O interfaces list in Table I, f_c is commonly within the range of 4–10 MHz. In addition, f_h is commonly set at $10f_c$, and f_l is commonly set at $f_c/100$. JM_2 is the SJ magnitude at frequencies below f_l , JM_1 is the magnitude at above f_c , and between f_l and f_c , the mask follows a -20 dB/decade slope. If a receiver-clock recovery corner frequency is at f_c , then a SJ with frequency above f_c is not trackable.

Having reviewed the signaling and jitter requirements for these ~10 Gbps standards, we discuss how they are met by the general-purpose FPGA transceivers manufactured at the 40-nm process node. Note that for the purpose of this paper, the Tx and Rx signaling and jitter requirements reviewed are representative ones from authors view. We refer readers to the respective standards documents for more details and complete coverage.

III. 40-nm TRANSCEIVER UP TO 11.3 Gbps

A. Architecture Overview

A FPGA 40-nm transceiver architecture is shown in Fig. 6. It consists of four full-duplex channels and two dual-function central clock management units (CMUs). Each channel independently operates at its own speed. It can select either of two CMUs to become an individual clock source or can listen to the global clock from CMUs from other quads so that groups of multiple Tx paths can be synchronized to operate at the same data rate. Each channel consists of one Tx path and one Rx path. The Tx path transmits the data to the outside world through a programmable Tx driver, and the Rx path receives the incoming data through a Rx receiver. To maximize customer flexibility, the Tx path and the Rx path of the same channel can operate at different data rates. Furthermore, each central CMU can be converted to a full duplex channel, allowing for additional channels instead of CMU functionality. In such a configuration, REFCLK pins become Rx path signal pairs, while CLKOUT pins are configured as the Tx path. The Tx path inside each CMU can be configured as center- or edge-aligned clock output for those standards that require forwarded clocks.

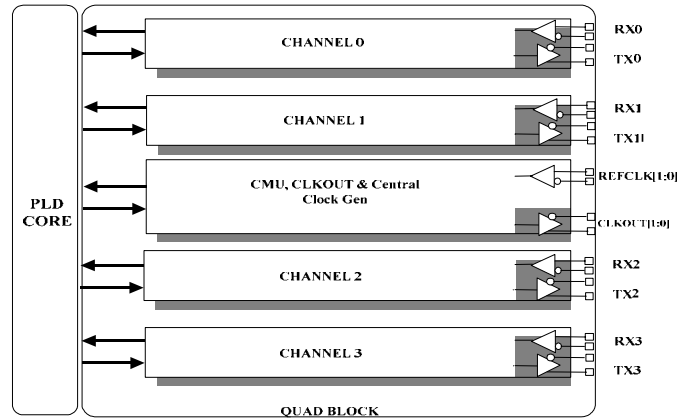


Fig. 6. A 40-nm FPGA Overall Transceiver Chip Architecture

In next two subsections, we discuss clocking and timing generation, and Tx, and Rx architectures.

B. Advanced Clock and Timing Generation

Clocking and timing generation play important roles in high-speed transceivers. An important metric that reflects the quality of a clock is jitter, because clock jitter affects both Tx and Rx jitter performance, in turn causing BER to increase for the link system.

On the Tx side, the clock jitter limits the eye opening at its output. The clock jitter on the Rx side affects the Rx to latch the receiving data correctly, and to consume 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 the 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 main source of jitter. Currently, there are two main types of oscillator used in multiple GHz PLLs—ring oscillators and LC tank (LC) oscillators—both of which have advantages and shortfalls. We focus on details of LC oscillators in the next subsection.

B.1 LC Oscillator

Many papers, including [7] and [8], describe the design of LC oscillators. A LC oscillator offers superior phase noise performance due to its highly selective and high Q LC tank. Discrete LC oscillators have been used in RF applications for long time, but the integration of LC oscillators into mixed-signal IC occurred only in recent years. Two main factors drive the use of LC oscillators in integrated transceiver designs. The first is that ring oscillator phase noise is unable to meet the Tx jitter requirements at the multi-GHz frequency range. The second is that, due to the process-feature size shrinkage, the inductor is now small enough to be integrated on a die.

The cross-coupled LC oscillator is the most widely used architecture (see Fig. 7). The LC tank oscillates in either current mode or voltage mode. In current mode, the signal amplitude is determined by a biased tail current. Although the waveform is immune to the supply variation, the tail current's 1/f noise upconversion degrades its phase noise. In voltage mode, there is no tail current and the amplitude is limited by supply voltage. A voltage regulator is needed to prevent supply noise and spur injection.

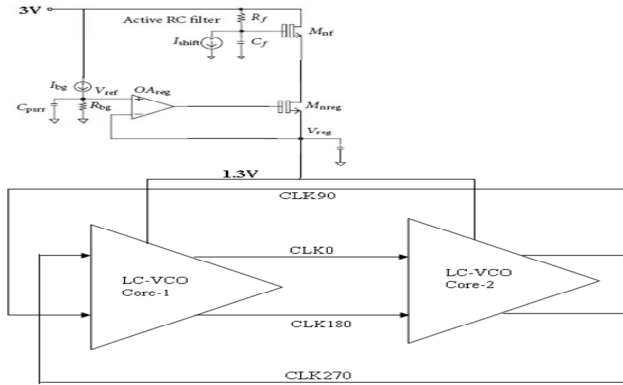


Fig. 7. LC VCO With Regulator and Filter

Compared to a ring oscillator, a LC oscillator occupies more die area due to its bulky inductors. Unlike a ring oscillator, an integrated LC oscillator has a limited frequency tuning range. The only tunable element of an on-chip LC oscillator is the varactor, and the tuning range of a either diode or MOS varactor is limited due to the current process technologies. Wide-range LC VCOs have been designed[9], but require a large number of bands, each covering a small range of frequencies. Frequency overlap of the bands across process, voltage, and temperature (PVT) is difficult to guarantee and the tuning range still is much less than a ring oscillator. The continuous frequency tuning range of a LC oscillator is typically around 20%. However, a LC oscillator normally offers much better phase noise/jitter performance than a ring oscillator.

During product development (e.g., 40-nm Stratix IV GX FPGAs), clock generation has both ring oscillator and LC oscillator technologies implemented and optimized for power, area, and jitter, as well as a variety of data rates, to support virtually all of the multiple-Gbps communication standards.

B.3 Noise Isolation and Reduction

Noise isolation/reduction is a key focus area when integrating analog circuitry into a noisy FPGA core. The magnitude and frequency of the noise is unknown, due to the programmable nature of the FPGA, which can affect the performance of the analog circuitry adversely. This especially is critical for the PLL and bias circuitry when one considers that the UI at 10 Gbps is 100 ps. Any sensitivity to noise has a large effect on the allowable timing/jitter budget in properly sampling and transmitting the serial data. Power-supply partitioning is critical. To enable acceptable noise performance, each channel has separate power and ground buses on the chip. Furthermore, voltage regulators with active filtering techniques are employed to reject out-of-band noise (see Fig. 7). A typical power-supply rejection ratio (PSRR) for this regulator is in the range of -40 to -60dB. Integrated regulators and filters greatly reduce the effort and cost of providing regulators and filtering on the board[10]. In some cases, over 128 ferrite beads and bypass capacitors may be used, as well as external linear regulators, to achieve the jitter performance goals.

C. Tx Architecture

The high-level Tx architecture is similar to that shown in Fig. 2. Because clock and timing is covered in sub-section B, here we focus on the Tx driver and give some circuit details.

C.1 Tx Driver


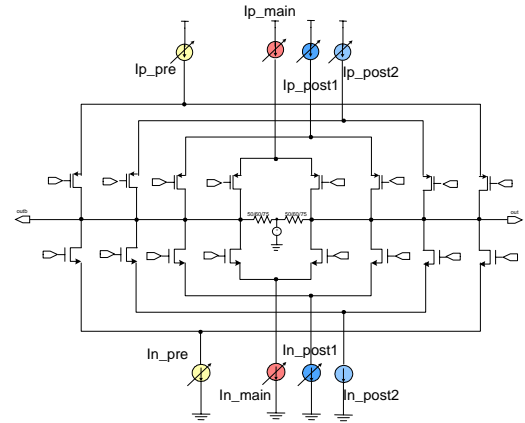


Fig. 8. H-Tree Driver FIR Block Diagram

The upper diagram in Fig.8 shows the H-tree Tx driver, also known as an AB-class driver. The H-tree driver has dynamic performance on par with a CML (A-class) driver with the additional benefits of inherent symmetry, better power rejection, and much lower power consumption. The symmetry of the H-tree driver reduces common mode noise and reduces electro-magnetic interference (EMI). The symmetry is due to the use of a current source in both the charge and discharge paths. The current source in the charge path also leads to better power supply rejection since the current source provides a high impedance path of over 1K Ω to the supply compared to the 50 Ω load used by a CML driver.

To illustrate power consumption savings between a CML and a H-tree buffer, Fig. 9 shows an H-tree link and a CML link with corresponding equations for peak-to-peak differential output voltages.




Fig. 9. H-Tree and CML Link Analyses

The current for the H-tree circuit flows from the positive arm into the negative arm across the load. The CML output is generated by only one arm, while the other arm is tri-stated. Thus, twice the current is required in order to generate the same swing.

The dynamic performance of the H-tree driver is on par with the CML driver because the dynamic current is the same as the CML driver. In both links, the termination impedance seen by each output is the same (for the H-tree, the common mode node is a virtual ground). The main speed difference comes from the differences in loading. In both designs, the loading is due to on-chip routing metallization and the on-chip bump, as well as package parasitic. In the CML layout, the routing loading to the pad is similar to the H-tree because electro-migration rules require that the routing must be twice as wide to

support twice the static current. Additionally, the transistor or device loading is roughly the same. The CML driver requires an explicit electro-static discharge (ESD) protection device to protect against positive ESD strikes, while the H-tree design uses the existing PMOS switch transistor and associated NW diode to clamp the output against positive strikes.

Tx signal conditioning is implemented at the Tx driver (see the second diagram in Fig. 8). The Tx signal is pre-conditioned before it is launched into the channel so the signal high-frequency contents get amplified (pre-emphasis), or the signal low-frequency contents get reduced. The benefit of this method is relative simplicity and low power. All of the sampled data information is readily available at the transmitting devices. Delayed versions of transmitted data can be easily created to be one UI apart by adjusting the bank of registers that holds both prior and upcoming serial data bits. Similarly, fractional sampled data (1/2 UI) can be made available in the Tx by taking information from the intermediate latch stages that commonly are used to create registers. Since both prior and upcoming (future) data bits are available, this signal conditioning technique can address both pre- and postcursor ISI well.

D. Rx Architecture

We first discuss the Rx buffer, followed by DFE, and then CDR.

D.1 Rx Buffer

Similar to pre-emphasis, equalization uses key signal conditioning techniques on the Rx side. The FPGA transceiver uses several Rx equalizers: CTLE, DFE, and adaptive dispersion compensation engine (ADCE). (Details of CTLE and ADCE designs are covered in [11] and not discussed here.) In this paper, we focus on DFE.

D.2 DFE

DFE is a non-linear system, see Fig. 10 for an example diagram of a 2-tap DFE. In general, a DFE can have n number of taps. A generic DFE system not only samples the data but also computes new coefficients prior to the next sample. This makes timing closure extremely challenging. Because DFE requires a proper data sample, this leads to increased design complexity of the combined equalization and recover sections of the transceiver. Furthermore, acute error propagation can arise, because the present DFE decision is based on prior samples. An error in the prior samples lead to an erroneous coefficient computation for the present data sample, therefore a single wrongfully captured bit can propagate out to a few consecutive bits until correct samples are obtained again. Since the DFE system bases its decisions on prior bits, it can only address postcursor ISI, leaving precursor ISI uncompensated. As a result, a LEQ such as CTLE (or FFE) still is required in a DFE system to accommodate the precursor ISI.




Fig. 1. Diagram of DFE Scheme

Even though both signal and noise are present in the frequency spectrum of data and are boosted by the DFE, boosting the fixed frequency spectrum improves the overall SNR because the power of noise is equally distributed over whole spectrum. (AWGN is assumed.)

D.3 Hybrid CDR Architecture

The CDR architecture advances the conventional data-driven architecture by allowing two operational modes, as well as digital controllability for its clock recovery. 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 the data signal to lock to the phase. In the end, both the frequency and the phase-aligned bit clock 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 Fig. 10.

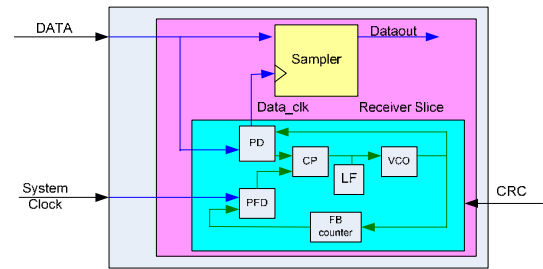


Fig. 10. A Digitally Assisted Hybrid CDR Architecture

The hybrid architecture has advanced capabilities in filtering and reducing jitter from the Tx and reference clock, thereby enabling the best possible BER performance for the link system. The reasons are as follows.

First, in the hybrid architecture, the reference clock only is used in the initial training phase for CRC, and is not used in the actual data recovery phase. Therefore, the reference clock jitter does not contribute to the system BER at all, leaving system designers with 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 may take a long time to lock on to the in-coming data when the receiver data input has excessive jitter, while the hybrid architecture largely is immune from such problems since its frequency-locked bit clock is already in place when recovering data.

Second, due to its digitally-assisted design, the hybrid architecture can set the bandwidths of a clock generation PLL at the Tx and a clock recovery PLL at the Rx independently, so the receiver PLL CRC bandwidth is much wider than that of the Tx PLL, as shown in Fig.11. Consequently, the Tx high-frequency jitter is attenuated by its PLL, while the transmitter low-frequency jitter is attenuated by the receiver CRU, leaving a minimum jitter contribution the system BER. In contrast, the common-clock architecture does not have the freedom to set the receiver PLL bandwidth to be wider than that of its transmitter PLL, since it is designed to match (cancel out) the




Fig. 13a. XLAUI/CAUI Receiver Worst-Case Stressed Eye Tolerance and Compliance Test Results From a Stratix IV GT FPGA

Fig. 13b shows the receiver CDR SJ sweeping tolerance test results. The test pattern is a PRBS $2^{31}-1$. The results not only meet but exceed the XLAUI/CAUI specifications.

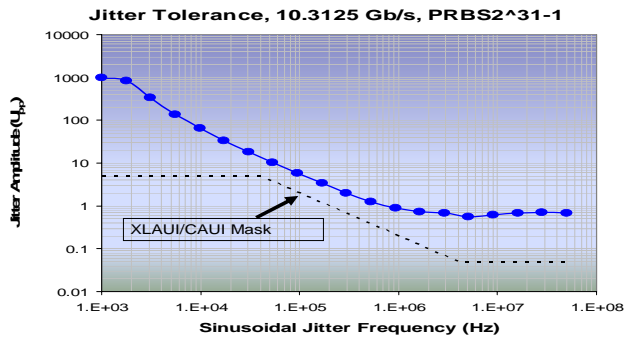


Fig.13b. XLAUI/CAUI Receiver Jitter Tolerance Compliance Test Results From a Stratix IV GT FPGA

V. SUMMARY

We reviewed the signaling and jitter requirements for emerging high-speed wireline communication standards at ~ 10 Gbps, including CEI 11G, XLAUI/CAUI, XFI, and SFP+, and covering transmitter output and receiver tolerance. We then presented a 40-nm FPGA general-purpose transceiver architecture and subsystem/circuit blocks aimed at meeting or exceeding the requirements for these emerging 10-Gbps standards. Critical functional/circuit blocks for clocking and timing generation with both ring- and LC-based VCOs, driver buffer and pre-emphasis/deemphasis, and receiver CDR and DFE were presented. In the end, we presented the signal/jitter output and tolerance measurement at 10.3125 Gbps, all of which exceed the XLAUI/CAUI standard requirements. In particular, we show that an ultra-low RJ of 550 fs can be achieved with the LC-based VCO.

Given the trend that high-speed I/O data rate doubles every 2–3 years, as suggested by the ITRS data, it is conceivable that the industry will design and manufacture transceivers at operating at 20–25 Gbps in two years. These transceivers will create many opportunities for innovative design and validation methods to solve the forthcoming challenges at these rates.

REFERENCES

- [1] “Common Electrical I/O (CEI)-Electrical and Jitter Interoperability agreements for 6G+ bps and 11G+ bps I/O,” IA # OIF-CEI-020, *OIF*, 2005, www.oiforum.com
- [2] “40 Gb/s Attachment Unit Interface (XLAUI) and 100 Gb/s Attachment Unit Interface (CAUI),” Annex 83A, IEEE P802.3™D2.0, *IEEE*, 2009, www.ieee802.org/3/ba
- [3] “10-Gigabit Small Form Factor Pluggable Module,” INF-8077i, Re4.5, *XFPMSA*, 2005, www.xfpmsa.org
- [4] “Enhanced Small Form Factor Pluggable Module ‘SFP+’,” SFF-8431, Rev 3.2, *SFF*, 2009, www.sffcommittee.com
- [5] ITRS roadmap, *ITRS*, 2007 Edition: www.itrs.net/Links/2007ITRS/Home2007.htm
- [6] M. Li, *Jitter, Noise, and Signal Integrity at High-Speed*, Prentice Hall, 2007.
- [7] J. Craninckx and M. Steyaert, “Low-Noise Voltage-Controlled Oscillators Using Enhanced LC-Tanks,” *IEEE Transactions on Circuits and Systems-II*, December 1995.
- [8] K. A. Kouznetsov and R. G. Meyer, “Phase Noise in LC Oscillators,” *JSSC*, Aug. 2000.
- [9] A. Berny, A. M. Niknejad, R. G. Meyer, “A Wideband Low-Phase-Noise CMOS VCO,” *IEEE CICC*, 2003, pp. 555-558.
- [10] S. Gondi et al, “A 10Gb/s CMOS Adaptive Equalizer for Backplane Applications,” *JSSCC*, 2005, session 18.1.
- [11] W. Wong, et al, “Digitally Assisted Adaptive Equalizer in 90 nm,” *DesignCon*, 2007, Track 9-TP1.



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

Copyright © 2009 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.