

Defining a New High-Speed, Multi-Protocol SerDes Architecture for Advanced Nodes

By Eric Naviasky, Cadence Design Systems

Designing a high-speed, high-performance serializer/deserializer (SerDes) for advanced process nodes can be challenging on many levels. The speed (16Gbps) stresses the capabilities of even the most modern process with limited gain available and without area-intensive peaking inductors. The high channel loss (-30dB) requires a decision feedback equalization (DFE) architecture, and the low-power constraints preclude the use of speculative designs. With multiple protocols and different processes in play, it's difficult to optimize for a single set of requirements; this situation instead calls for more general solutions. For example, consider the emergence of FinFET processes. The FinFET process presents characteristics that are very different from those of prior semiconductor devices, such as high Rout (low gds), low static current densities, and very restrictive device sizing. In response to these criteria, this paper defines a new high-speed, multi-protocol SerDes architecture that is ideal for meeting the challenges of advanced-node designs.

Contents

Introduction.....	1
Design Approach	2
Design Architecture	2
Clock and Data Recovery: Why Separate, Optimized Paths are Advantageous.....	5
Addressing Traditional Receiver Issues	6
Improving Transmitter Performance	6
Resolving Clock Generation Issues	7
Power Advantages of Multi-Protocol SerDes Architecture	8
New SerDes Architecture Available as Cadence PHY IP	8
Summary	9
For Further Information.....	9

Introduction

Next-generation peripherals, tablets, servers, and other applications are demanding greater bandwidth at lower cost and power. To meet these demands, communications protocols like PCI Express® (PCIe®) have gotten substantially faster—PCIe Gen4 calls for signal transmission speeds of 16Gbps, as well as transmission with margin through channel losses as high as 27dB at Nyquist. As with previous versions of the protocol, PCIe Gen4 is also suited for a variety of cost- and power-sensitive applications.

To minimize the number of I/O pins and interconnects in what are already complex designs, design engineers look to SerDes physical-layer (PHY) intellectual property (IP) to provide the data transmission over a single or differential line. Traditional low-power SerDes architectures either provide good channel loss performance at lower speeds or good speed for lower loss channels. However, these architectures have limitations in addressing the need for low-power, high-speed performance on the higher loss channels of next-generation applications. Therefore, supporting the latest communications protocol specifications, including PCIe, calls for a new type of SerDes architecture that addresses the following needs with minimal power dissipation:

- Data and clock recovery requirements in high-speed, high-dB-loss, and high-crosstalk channels
- Critical loop timing specifications for the DFE
- Environmental and process variations
- Transmitter performance under low-supply conditions
- High-speed clock distribution

A new multi-protocol, high-speed SerDes architecture, designed for advanced nodes, addresses all of these challenges while offering the following characteristics:

- Support for data rates of 1Gbps up to 16Gbps, with a continuous frequency range
- Compliance with:
 - PCIe Gen4: 2.5Gbps, 5Gbps, 8Gbps, 16Gbps
 - 10G-KR: 10.3125Gbps, 12.5Gbps
 - XAUI: 3.125Gbps
 - RXAUI: 6.25Gbps
 - Gigabit Ethernet/SGMII: 1.25Gbps
 - SATA: 1.5, 3, 6Gbps
 - HMC-SR: 10Gbps, 12.5Gbps, 15Gbps
- Equalization up to 30dB channel loss (in the presence of -48dB crosstalk)
- Low power
- Flexibility and robustness

Let's take a closer look at this new multi-protocol SerDes architecture and how it was developed to address the limitations of traditional SerDes architectures.

Design Approach

Developing the multi-protocol SerDes architecture involved these techniques:

- Optimizing the block design for one function and the system design for overall performance. The paths and small blocks were optimized for speed and low power. Accuracy was layered on with background digital calibration. The design also used simple analog blocks that do not rely on precise knowledge of device characteristics.
- Integrating adaption loops that measure at the highest level possible, while controlling at the lowest level possible
- Tracking inter-stage common mode voltage for internal stages

Design Architecture

As shown in Figures 1 and 2, the new SerDes architecture has a common module and a module with 1 to 16 lanes. The architecture features a dual-path receiver (RX) and a hybrid transmitter (TX). The receiver is configured in a dual path for better jitter tolerance, crosstalk rejection, and power dissipation. The hybrid transmitter allows the ability to perform real emphasis, provides larger swings than would be available from a conventional voltage mode driver, results in less wasted power in de-emphasis, supports broadband matching, and also allows use of larger and better protecting electrostatic discharge diodes.

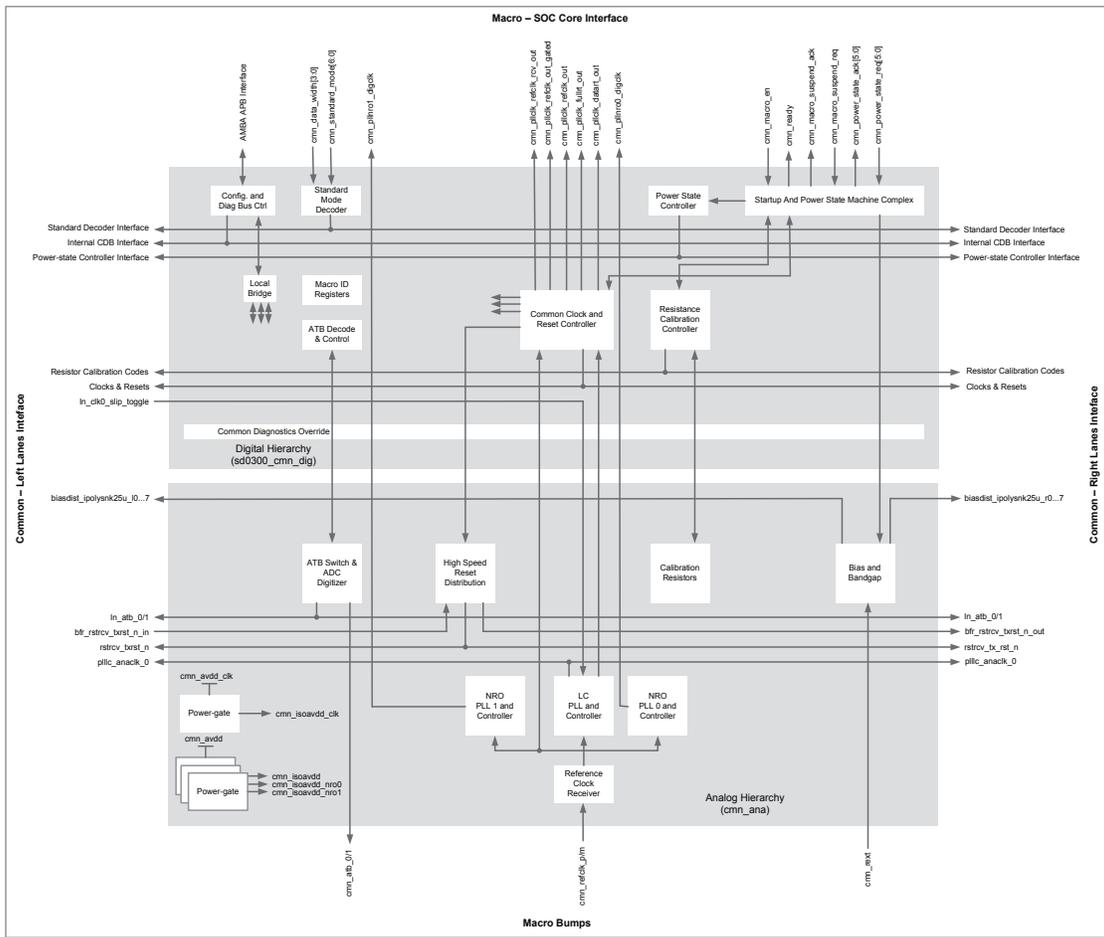


Figure 1: Common Module Block Diagram

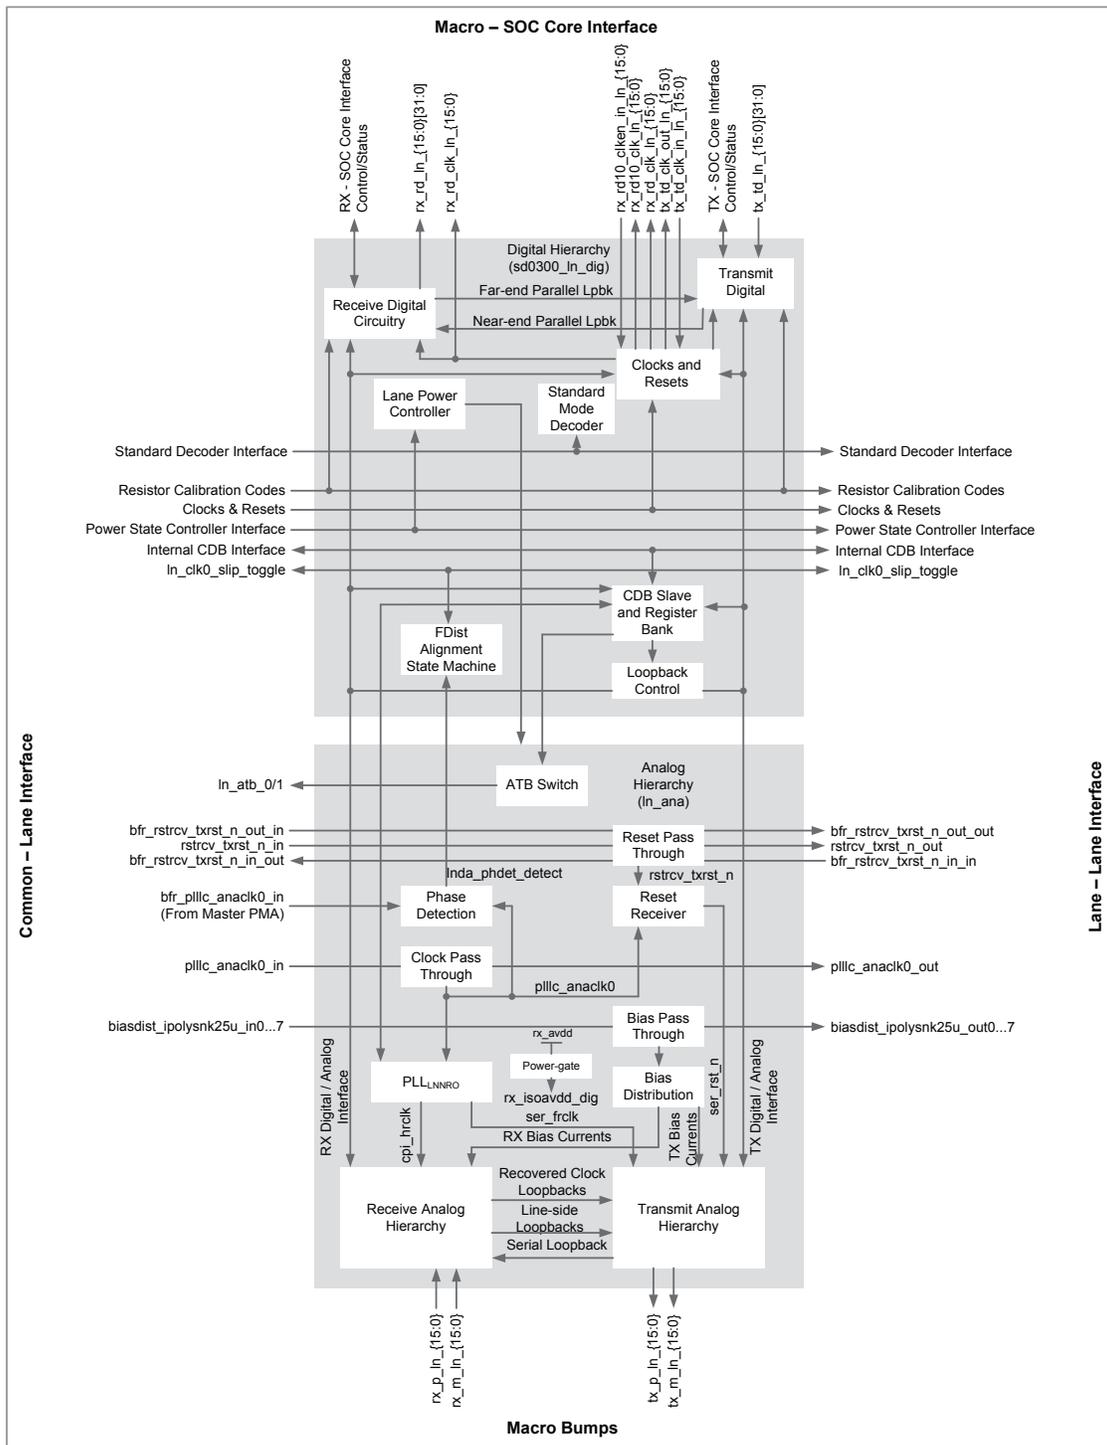


Figure 2: Lane Module Block Diagram

For clock generation and distribution, a low-jitter clean-up phase-locked loop (PLL) in the common area allows the use of a cost-effective reference. Unbuffered clock distribution on high-level metal avoids jitter that's induced by power supply noise. The architecture also includes an in-lane local PLL operating at the TX baud rate.

There are many adaption loops in the architecture beyond the traditional DFE training, including clock offset, clock equalization, and sampling point adjust. Datapath gain, offset, TX-FFE, and DFE tap weights are calculated in a correlation calculation at the data sampling point using data history and a pair of calibrated samplers. Clock

path equalization, offset, and gain are calculated in a statistics-based engine from the clock sampler results. The sampling point adjustment is based on a statistical measurement of the eye opening at the data sampling point. The phase interpolator calibration is based on a foreground calibration at startup and a delta-based background calibration that occurs periodically after startup. Data word alignment is only performed at startup.

The next few sections of this paper will examine how the new architecture addresses typical SerDes design challenges where traditional architectures fall short.

Clock and Data Recovery: Why Separate, Optimized Paths are Advantageous

Equalization of high-dB-loss channels

Traditional SerDes architectures have a limit in maximum achievable equalization. Typically, the DFE can't begin opening the eye until clock recovery has occurred. Clock recovery, in turn, can't start until the eye is slightly open. Also at odds are clock recovery and data recovery continuous time linear equalizer (CTLE) frequency response. High-frequency peaking CTLE for clock recovery amplifies crosstalk and reduces signal/noise. Data recovery prefers low-frequency equalization to extend to areas that the DFE cannot reach. High-frequency equalization for data recovery is best provided by the DFE to avoid high-frequency concentration of crosstalk.

High-loss DFE systems

In high-loss DFE systems, there are a number of issues around clock data recovery to address:

- Because DFE corrections are nearly as large as the signal, they often fool the clock data recovery. The corrections shift the apparent center of the eye for clock recovery and obscure the real data transitions, creating a blind spot. Such actions eventually cause runaway clock data recovery or result in the creation of zones of reduced CDR gain.
- DFE equalization for data is incorrect for clock recovery. DFE corrections are converged for data samples. Since DFE is a discreet time equalization, the correction value at the edge samples is fundamentally wrong.
- Clock recovery requires high-frequency equalization, but it is only minimally impacted by crosstalk amplification
- At high speeds, such as the 16Gbps required by PCIe Gen4, closing the timing around the DFE loop becomes even more difficult

The multi-protocol SerDes architecture was designed to address the above limitations. The architecture consists of separate optimized paths for clock recovery and data recovery. For each path, the relative timing is adjusted by an adaptive loop, which saves the power that would have been required to maintain good alignment of phases in the clock distribution. The DFE path samplers are only used for sampling open eye signals and can, therefore, be designed for speed rather than metastability. In turn, this metastability improves the critical loop timing of the DFE so that no speculation is needed. Decoupled clock recovery also allows for much better jitter tolerance because the CTLE and edge samplers are optimized for the clock path. Unlike many SerDes architectures, this new architecture allows use of every edge in the datastream for clock recovery.

Ideal for advanced-node designs, the new SerDes architecture offers the following advantages:

- Extends the theoretical limit of equalization by removing the limitation that the eye must be opened before recovery can begin
- Optimizes data recovery for accuracy in the presence of crosstalk; clock recovery takes advantage of inherent crosstalk insensitivity
- Eliminates obscuring of the clock recovery loop by DFE corrections. Instead, clock recovery equalization is adapted at clock sample times and all edges contribute information, so there are no blind patterns or dead spots.
- Offers an additional degree of freedom, with optimal data sampling point set by an adaptive loop

Addressing Traditional Receiver Issues

The RX in the new SerDes architecture, shown in Figure 3, is configured with dual paths. The dual-path RX allows for better jitter tolerance, better crosstalk rejection, and comparable or better power dissipation than the single-path RX. Jitter tolerance is improved because a lean clock recovery path design minimizes latency and because no edges are lost to pattern filtering or DFE artifacts. A benchmark test against a competitive solution demonstrated some advantages of this new architecture: up to 100% better jitter tolerance at low frequency, up to 50% better jitter tolerance at high frequency, and 2X higher clock recovery bandwidth. The power savings result from a singular optimized function of the blocks.

The optimized clock path gets more signal and less noise than in a single-path design, due to a number of factors. For one, a separate CTLE for clock recovery, shown in the lower portion of Figure 3, allows high-frequency peaking optimization for clock recovery. The equalizer is converged at the clock sample time, without having to rely on incorrect discrete equalization converged at the data sample time. And, all data patterns can contribute to clock recovery.

In Figure 3, the red blocks show the adaptive loops in the receiver. A digital controller manages all of the loops. Some of the adaptive loops are for start-up only. Others run in the background, so if there are changes in, for instance, humidity or temperature, then the backplane automatically adjusts to accommodate the changes. This approach allows continuous uptime, as the background adaptive loops do not interrupt the flow of data through the system.

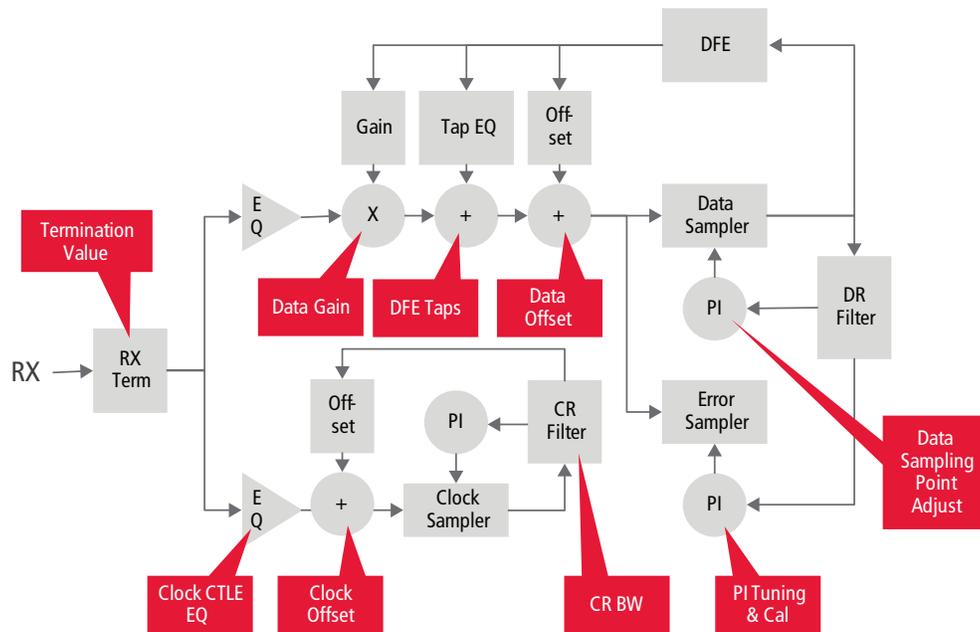


Figure 3. Dual-path RX in multi-protocol SerDes architecture

The clock recovery process is a band-limited process where the information (clock phase) is encoded on a much higher frequency carrier (the baud rate). The loss of the information contained in a few transitions is not fatal to the clock recovery process, as there is a good bit of redundancy. The problem occurs when too many transitions are “discarded” due to contamination from the DFE. Traditional clock and data recovery systems either allow those errors to contaminate the desired information, or some improved systems “blank” the recovery during which pattern-matching filters determine certain patterns to be risky. The loss of the subset of the data reduces the “gain” of the phase detection. While gain can be made up in other parts of the loop, the noise is not reduced, so the signal-to-noise ratio is degraded.

Improving Transmitter Performance

In many traditional SerDes architectures, transmitter performance is limited by a variety of issues:

- The transmit drive amplitude (V_{ppd}) of the H-bridge is theoretically limited to the supply voltage
- Regulators aren't desirable due to power dissipation and area
- Power dissipation of the H-bridge increases quickly as de-emphasis is applied
- More margin is needed, as higher frequency effects steal amplitude at transition bits and real-world boards have many imperfections and high loss due to economic concerns

In addition, there are a number of real-world effects that also hinder the transmitter:

- Reflections that are caused by stubs, vias, and discontinuities in package and board
- Ringing to LC resonances
- Stray capacitance in package
- Capacitance to power planes
- Imperfect trace impedance
- Ground return current issues

The hybrid TX path (Figure 4) in the multi-protocol, high-speed SerDes architecture is designed with a hybrid driver with true emphasis, not just de-emphasis. This path offers better rise times due to boost circuit, lower output cap than H-bridge, and less wasted power in de-emphasis. Its transmit amplitude is $>900mV_{ppd}$ with an internal worst-case 0.765V supply. By comparison, competing H-bridge designs are unable to reach $800mV_{ppd}$. Many traditional designs also tend to use regulators to boost amplitude, an approach that comes at a power disadvantage.

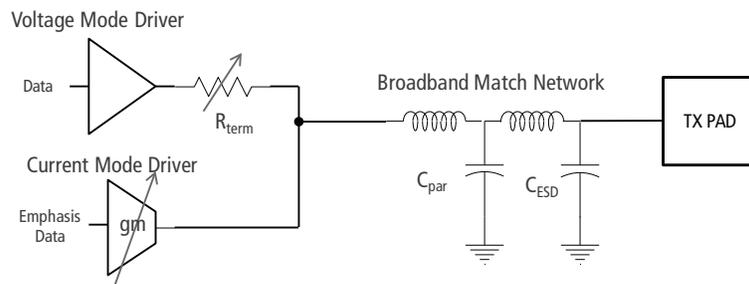


Figure 4. Hybrid TX path in new SerDes architecture

The hybrid TX path addresses transmitter effects by:

- Maintaining power advantages inherent to non-emphasized voltage mode
- Allowing additional amplitude in excess of what the voltage mode can produce
- Requiring less power in emphasis than conventional voltage mode or current mode driver
- Allowing, through the use of broadband matching, the use of larger and better protecting electrostatic discharge (ESD) diodes

Addressing transmit return loss, the broadband matched output stage of the transmitter achieves the required $-6dB$ at 8Gbps. The transmitter is more tolerant of output capacitance from larger ESD protection and, based on its distribution of load capacitance, provides better transmit rise time. The transmitter also has the flexibility to adjust for different packages and reduces demands on package design.

Resolving Clock Generation Issues

In traditional SerDes architectures, clock generation faces some challenges:

- To achieve the lowest jitter, an LC tank PLL is needed; however, LC tank PLLs have a very narrow tuning range and consume a lot of area
- Ring oscillator PLLs have a wide range and consume less area and power, but, without a high reference frequency, they cannot achieve the required jitter performance
- Many applications call for simultaneous multi-rate lanes within the macro, which creates the need for a second, large LC tank PLL and clock distribution network
- To reduce costs, many applications are designed for low reference frequencies (~25MHz), but these reference frequencies require a high level of LC tank performance
- The clock must be distributed long distances from the common area to lanes with minimal jitter, skew, and duty cycle distortion

The multi-protocol, high-speed SerDes architecture features a clock distribution network that addresses these issues. A low-jitter clean-up PLL in the common area allows the use of a less expensive reference and lower reference frequencies. A lane-local digital PLL allows for flexible, multi-rate operation and improved duty cycle distortion (DCD). Finally, unbuffered clock distribution on high-level metal avoids the common ingress point for power supply noise-induced jitter.

Transmit jitter is minimized because the full-rate TX clock is locally generated and has minimal distribution to the TX. Full-rate operation avoids common sources of DCD. This configuration also eliminates the need for a 16Gb clock distribution network, resulting in a more flexible, tolerant design.

Power Advantages of Multi-Protocol SerDes Architecture

The new SerDes architecture was designed to minimize power consumption through architectural changes that allow the use of lower power circuit blocks. Power per Gbps (x8 configuration) is 7.5mW. Its split clock and data recovery RX paths allow for speed-optimized samplers for data and metastability and power-optimized samplers for edge detection. Since the data path is a full-speed DFE, it avoids the substantial increase in required circuitry and clock distribution that unrolling would need. In addition, the IP:

- Eliminates the need for a critical IQ phase-aligned clock distribution
- Uses a reduced area and power, wide-frequency-range phase interpolator
- Features lower frequency, top-level clock distribution

Leakage optimizations for PCIe, including L1 sub-state engineering change notice (ECN) and lane-up configure, also contribute to the power savings.

New SerDes Architecture Available as Cadence PHY IP

The multi-protocol, high-speed SerDes architecture described in this paper is available as PHY IP from Cadence, which has more than 15 years of experience developing SerDes architectures. Ideal for advanced-node processes, such as 16nm FinFET, the SerDes high-speed PHY IP is based on an analog five-tap DFE architecture with adaptive CTLE and offset correction. Given its multi-protocol support, a single instance of Cadence® SerDes PHY IP can be used for the different communication links in a design, making the IP lower cost, easier to integrate, and easier to validate.

The IP has a full suite of built-in test features and bypass modes for verification in place. BIST features include multiple loopbacks, standard and user-programmable pattern generators and checkers, analog test bus with available A/D for reading results, and full register access to all digital parameters. The user has the ability to remap CTLE and VGA tables or to use the standard ones programmed in. Available with a GUI, the IP includes an Eye Surf tool which provides internal eye plot capability that lets users visualize, on the fly, eye and bathtub curve diagrams depicting the speed of the internal signals via the output of the TX inside the SerDes.

The SerDes PHY IP is easy to integrate into a design and is part of a complete PCIe Gen4 solution including the controller and firmware, also available from Cadence. The IP supports inverse bifurcation in PCIe and 10G-KR modes, features configurable parallel data word widths of 16, 20, or 32, and offers selectable serial pin polarity reversal for transmit and receive paths. It is designed for channel loss of 24dB but can meet a target of 30dB.

Summary

Traditional SerDes architectures have limitations when it comes to addressing the demands of next-generation applications for high bandwidth and signal transmission, along with low cost and power. A new multi-protocol, high-speed SerDes architecture, available in PHY IP from Cadence, is designed to support the performance, cost, and power requirements of advanced-node designs. As an established supplier of proven, configurable IP, Cadence can provide engineers with the assurance that their advanced-node designs will meet the stringent requirements of a variety of next-generation applications.

For Further Information

Learn more about Cadence's multi-protocol, high-speed SerDes PHY IP at <http://ip.cadence.com/ipportfolio/ip-portfolio-overview/interface-ip/serdes-ip>.



Cadence Design Systems enables global electronic design innovation and plays an essential role in the creation of today's electronics. Customers use Cadence software, hardware, IP, and expertise to design and verify today's mobile, cloud and connectivity applications. www.cadence.com