

Supporting Advanced-Node FinFET SoCs with 16Gbps Multi-Protocol SerDes PHY IP

By William Chen and Osman Javed, Cadence Design Systems

Applications such as the Internet of Things, cloud computing, and high-definition video streaming are all taxing enterprise/server computing performance and bandwidth. To support these demands, the IT infrastructure calls for higher data rates, and communications protocols such as PCI Express® are responding by increasing the serial data rate to 16Gbps per channel. On the process side, the advance-node FinFET process offers better speed/gain transistor with lower power leakage, compared to its planar predecessors (32nm/28nm). This benefit, along with its given geometric advantage—approximate 0.53X logic area scaling from generation to generation—make the FinFET process ideal for high-speed/high-performance system-on-chip (SoC) design. To help assuage what are already complex SoC designs on advanced-node processes, many SoC architects look to SerDes PHY intellectual property (IP) vs. in-house development. The market trend of incorporating multiple protocol interfaces into a single SoC design offers a great value proposition for the concept of adopting a single multi-protocol PHY. This paper takes a look at how a new 16Gbps multi-protocol SerDes PHY IP addresses the unique challenges of advanced-node FinFET design. The paper also discusses the benefits offered by its multi-protocol capability to support SoCs for server computing applications.

Contents

Introduction	1
Why a Multi-Protocol PHY?	2
How the Multi-Protocol PHY Works	2
Overview: 16Gbps Multi-Protocol PHY IP	2
Conclusion	4
For Further Information	5
Sources.....	5

Introduction

We’re living in a high-speed, high-bandwidth era, where we are sharing and transacting on an increasingly voluminous amount of digital data. The IT backbone supporting these connected applications must meet this continuous demand for faster processing performance and higher I/O bandwidth.

Applications—particularly those in areas such as consumer, networking, and storage—are transitioning as a result, moving from typical 5-8Gbps to 10-16Gbps (per physical channel) speeds. See Table 1 for a quick snapshot of the protocol standards that offer data rate support of greater than 10Gbps.

Standard	Application	Max. Data Rate (Gbps per channel)	Remark
10G-KR (IEEE 802.3 clause 49)	Networking	10.3125 12.5	Mature products in the market
HMC-SR	Storage	15	April 2013 v1.0 spec release
USB 3.1	Consumer	10	July 2013 v1.0 spec release
PCIe Gen 4.0	Consumer	16	Completed specification anticipated in Q4’ 2015

Table 1: Major 10Gbps+ Standards

PCI Express (PCIe®) 4.0 is the next generation of the ubiquitous and general-purpose PCI Express I/O specification. With this latest generation, the interconnect bandwidth has doubled over the PCIe 3.0 specification (8Gbps). Meanwhile, compatibility with software and mechanical interfaces is maintained.

When used along with modern advanced processes, PCIe 4.0 enables SoC designs to integrate more digital logic and run at faster speeds with lower active power consumption. For example, you'd be able to take advantage of L1 sub-states engineering change notice (ECN). Designing a high-speed, high-performance SerDes for such advanced technology, however, can be challenging on many levels. As an example, consider the 16nm FinFET process—this technology brings advantages including lower leakage and power and higher drive currents. But this type of design also presents characteristics that are very different from recent semiconductor devices, such as high output impedance (because of low gds), low static current densities, and very restrictive device sizing. In addition, because device width is quantized, place and route and custom design tools need to be updated. In addition, the 3D structure requires 3D RC extraction. There's an increase in gate capacitance, which impacts static timing analysis (STA). Stronger Miller Effects call for certification of electro-migration, IR, and STA tools.¹

In the following sections, let's discuss how a silicon-proven high-performance 16Gbps PHY can support not only PCIe 4.0 applications at the 16nm FinFET process but also deliver on the market's needs for IP cores that support multiple protocols.

Why a Multi-Protocol PHY?

Modern SoC designs tend to integrate more and more interface protocols in response to market demands for:

- Flexibility to bridge different systems, such as PCIe, SATA, USB, and Ethernet, in an all-in-one unit
- Alleviated engineering burden for handling the increased technical complexity on each protocol
- Competitive costs along with fast time to market, through the sharing the common components and by reducing the extra time required to bring up each individual protocol PHY

How the Multi-Protocol PHY Works

Typically, a PHY consist of two modules: the physical media attachment (PMA) and the physical coding sublayer (PCS). The PMA is the analog part of the SerDes implementation, while the PCS is a digital block establishing the functional link to the protocol supported. A multi-protocol PHY implements support for multiple protocols by optimizing the usage of signal-path blocks needed for each protocol. In the case of the PMA, this support consists of a common high-speed SerDes signal path across multiple protocols. For the digital PCS, it entails a configurable design that uses the same signal-path components multiplexed as needed for different protocols. This approach leads to optimized die area and, as a result, lower SoC cost and an accelerated design schedule.

By tapping into a single, multi-protocol PHY IP, an SoC designer can reuse many of the common design building blocks, enjoy a much simpler design verification process, benefit from a faster bring-up and characterization process, and perform only a single test chip tapeout. Other advantages of a single PHY IP:

- Create a single design that supports multiple applications. The advanced transceiver and clocking architecture minimizes power and the number of logic components required.
- Change protocol during the design process via a soft code change, not a full chip redesign and without the headache of readjusting the chip floorplan
- Relieve engineering burden and reduce overall design costs by getting multiple uses from a single IP

Overview: 16Gbps Multi-Protocol PHY IP

Cadence offers a 16Gbps multi-protocol PHY IP that is a hard PHY macro consisting of a PMA layer and a soft PCS. It provides continuous frequency range support from 1.25Gbps to 16Gbps, with the capability to equalize channel loss up to 30dB at Nyquist, and complies with the PCIe 4.0, 10G-KR, RXAUI, XAUI, GbE/SGMII, SATA3, and HMC-SR specifications. The IP supports low-power mode (such as L1 sub-states for PCIe) for power-efficient applications. Implemented on TSMC's 16nm FinFET process, the PHY IP can help control costs and power for high-performance designs.

Specification	Data Rate
PCIe 4.0	2.5Gbps, 5Gbps, 8Gbps, 16Gbps
10G-KR	10.3125Gbps, 12.5Gbps
XAUI	3.125Gbps
RXAUI	6.25Gbps
Gigabit Ethernet/SGMII	1.25Gbps
SATA	1.5Gbps, 3Gbps, 6Gbps
HMC-SR	10Gbps, 12.5Gbps, 15Gbps

Table 2: 16Gbps Multi-Protocol PHY IP – Compliance

The PHY also offers two different options of multi-protocol configuration: single link and multi link (Figure 1). Single link offers conventional multi-protocol function with all the lanes running at one protocol at a time. Thanks to the novel dual phase-locked loop (PLL) clocking topology employed in this PHY design, the multi-link option offers greater flexibility for SoCs to exploit different protocols via different lanes concurrently.

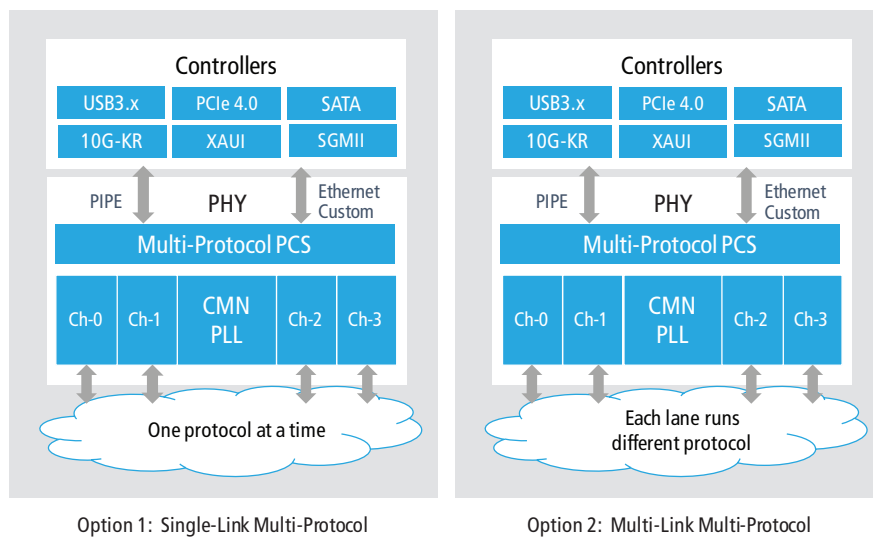


Figure 1: Single-link and multi-link configuration options in multi-protocol PHY.

The following eye diagram demonstrates the superior PHY performance at 16Gbps.




Figure 2: Eye diagram of 16Gbps PHY performance.

Given the aforementioned side effects introduced by the advanced FinFET process, this PHY IP is architected and implemented in a way to overcome the following challenges:

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

Its block design is optimized for one function and its system design for overall performance (speed vs. power trade-off). To reduce dependency on highly accurate process characteristics, simple analog blocks that don't rely on precise device characteristic knowledge are used. In order to gain the maximum margin for IR drops and supply noise interference, the supply voltage is set as low as possible. Adaptation loops measure at the highest level possible and control at the lowest level possible.

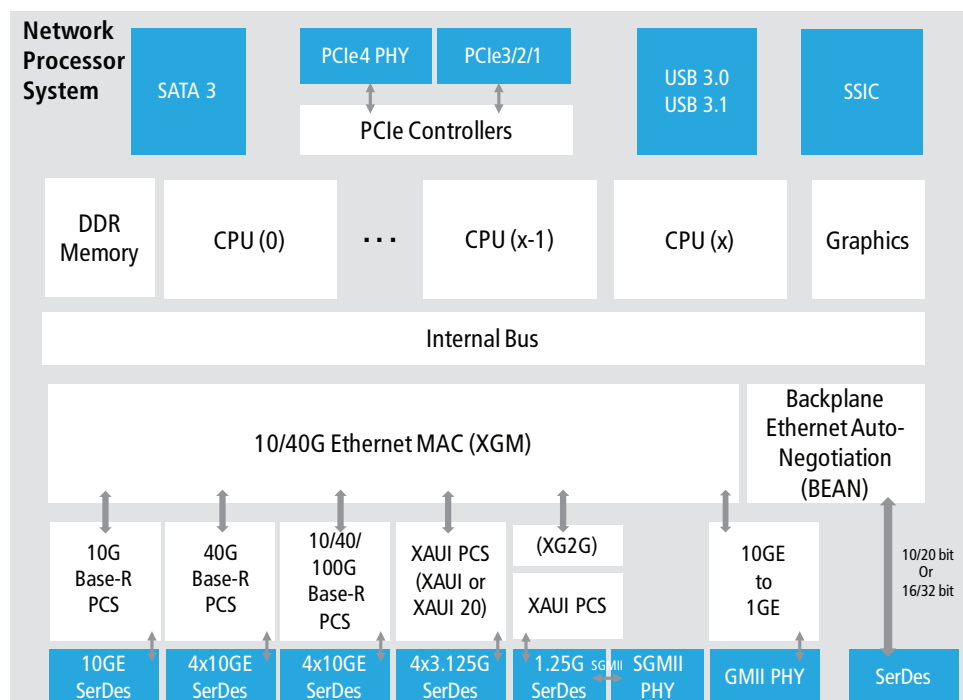


Figure 3: Example system-level block diagram. All of the light blue blocks are to be covered by this multi-protocol PHY design.

Conclusion

Increasing digital demands in our world mean that the data rate evolution we are experiencing will continue. While new industry protocols keep rolling out to offer system perspective solutions, technical challenges to engineering design will remain relentless.

Modern SoC designs incorporate more and more interface protocols in a single chip to meet demands for product flexibility, fast time to market, and cost competitiveness. The proposed multi-protocol PHY IP provides the market an easy-to-use unified PMA + configurable PCS approach in conjunction with cost benefits. The IP offers an optimum PHY solution to satisfy application needs ranging from data rates of 1.25Gbps to 16Gbps and covering most major and popular standards, such as PCIe, SATA, USB 3.x, SGMII/XAUI, 10G-KR, and others.

From the process technology point of view, in contrast to its predecessor process at 28nm, the 16nm FinFET process offers better speed performance and power-saving merits. However, the 16nm FinFET process also brings unfavorable side effects, such as much higher device rout requiring extra effort for loop stabilization, device current

density that is limited by metal before overdrive, and area-expensive decoupling capacitance. All of these design implications require experienced engineering and the persistence of quality commitment on iterations of design and layout optimizations with respect to the desired power, performance, and area.

FinFET technology facilitates higher speed SerDes design (from 8Gbps to 28Gbps) for addressing emerging standards such as PCIe 4.0. Given the degree of complexity on designing a high-performance multi-Gbps data rate SerDes, system houses and chip companies tend to out-source these types of sophisticated jobs to capable and qualified vendors, like Cadence. Cadence has committed resources for TSMC 16nm FinFET IP development across the board, building on its analog EDA experience and successful IP development in advanced process technology nodes (28nm and 16nm). Cadence's 16Gbps multi-protocol PHY IP can help designers meet cost and power targets for high-performance designs.

For Further Information

Learn more about the architecture of Cadence's multi-protocol SerDes PHY IP in this white paper (you will need to log in to your Cadence® account for access): http://ip.cadence.com/uploads/file/1021/652/Multi_protocol_SerDes_PHY_IP_WP_final.pdf

Sources

1. <https://www.semiwiki.com/forum/content/3165-designing-soc-16nm-finfet.html>



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