



Common Electrical I/O

Building for the New Future

Executive Summary

Recognizing industry's need for the entire component level infrastructure to be upgraded to support higher system capacity, the OIF authorized the PLL Working Group to begin a new project, referred to as the CEI-25G. The goal of the project is to define electrical specifications for 28 Gb/s signaling for chip-to-chip and chip-to-module applications and 25 Gb/s signaling for backplane applications. These specifications will provide the basis for the electrical interfaces that will be used to enable the higher capacities needed for the next generation of systems.

Introduction

The mission of the Optical Internetworking Forum (OIF) is to foster the development and deployment of interoperable products and services for data switching and routing using optical networking technologies. This requires addressing multiple issues related to optical internetworking at the carrier, system vendor, and component vendor levels. This integrated approach strengthens the OIF's ability to fulfill its mission.

The Physical and Link Layer (PLL) Working Group of the OIF specifies implementation agreements (IA) that not only enhance the interoperability of communication links at the optical level but also at the electrical interface, which includes the electrical properties and protocol, between the different components in a typical optical networking application.

In 2002 the OIF PLL Working Group undertook the creation of the Common Electrical I/O (CEI) specification, which defined electrical operation for serial 6+ and 11+ Gb/s data rates. Operation was defined to support both short reach ("SR": 0 to 200 mm or 8 inches of printed circuit board trace with up to one connector) and long reach ("LR": 0 to 1 meter or 39 inches of printed circuit board trace with up to two connectors, i.e. backplane) applications.

The influence of this effort permeated into other bodies that developed new standards aimed at these data rates. Ultimately, the CEI project provided the interconnect building blocks that are now being used to build the equipment that will satisfy the industry's immediate need for bandwidth.

The thirst for increased bandwidth continues, as is evident by the IEEE's formation of the IEEE P802.3ba 40 Gb/s and 100 Gb/s Ethernet Task Force. Furthermore, the ITU-T SG15 is currently working to define ODU4 to provide transparent backhauling of 100 GbE on a single wavelength. Before the formation of the IEEE P802.3ba Task Force, while it was still a study group, it was demonstrated, as illustrated in Figure 1, that networking applications are doubling approximately every 18 months while computing

applications are doubling approximately every 24 months. It is anticipated by many that 40GbE will leverage off of existing 10G serial technologies, optically and electrically. 100 GbE will also leverage 10 Gb/s serial technologies in initial deployments. As with the introduction of any new speeds, these initial deployments will consist of line cards consisting of one or maybe two ports of 100 GbE. Ultimately, however, there will be a need for switches and routers to provide a greater number of 100 GbE ports per given system. This need will drive the need for electrical interconnects beyond the current capabilities of CEI: chip-to-chip, chip-to-module, and backplane interconnect.

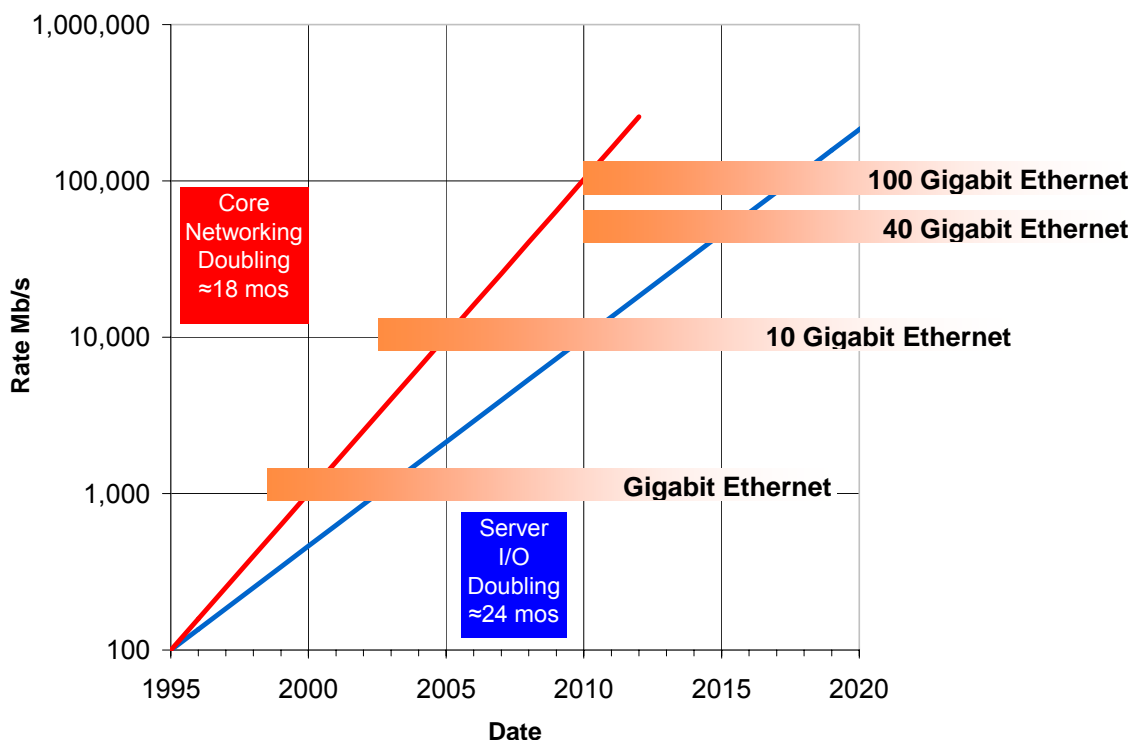


Figure 1- Bandwidth Growth Forecasts¹

Introducing the CEI Project

Recognizing the industry's need for the entire component level infrastructure to be upgraded to support higher system capacity, the OIF authorized the PLL Working Group to begin the CEI-25 project. This project will define electrical specifications for 28 Gb/s signaling for chip-to-chip and chip-to-module applications and 25 Gb/s signaling for backplane applications. This will allow the development of narrower interfaces that support 100 Gb/s payloads, and provide smaller package sizes, lower pin count components, connectors and optical modules, lower power dissipation, and clockless interfaces.

OIF System Reference Model

The OIF system reference model is used as a basis for development of OIF implementation agreements and is illustrated in Figure 2.

Interfaces in telecom applications fall into two generalized categories within this model: System Packet Interfaces (SPI) connect Link Layer devices to Physical Layer devices; and Serdes-Framer Interfaces

¹ http://www.ieee802.org/3/hssg/public/nov07/HSSG_Tutorial_1107.zip
www.oiforum.com

(SFI) connect Physical Layer devices to other Physical Layer devices. Physical Layer devices are generally located on the same card and therefore the electrical requirements for Serdes-Framer Interfaces support relatively short chip-to-chip or chip-to-module connections. On the other hand, Link Layer devices are frequently not on the same card as the Physical Layer devices, and therefore System Packet Interfaces must traverse a backplane. CEI defines electrical requirements for both “Short Reach” and “Long Reach” interfaces, supporting both SFI chip-to-chip and chip-to-module applications, and SPI backplane applications.

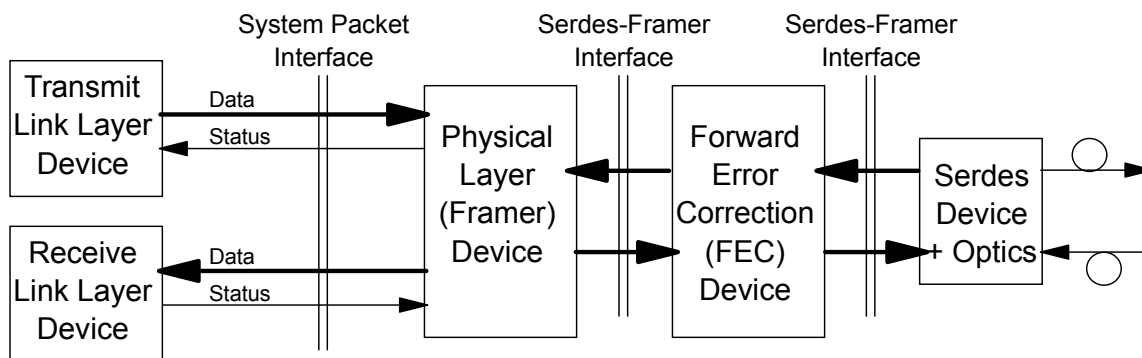


Figure 2 - OIF System Reference Model

Baud rate requirements also differ for SPI and SFI applications. While applications with 100 Gb/s payloads are on the horizon, implementation of serial 100 Gb/s electrical signals across backplanes will not be realizable within the necessary timeframe. Conversely, using existing 10 Gb/s serial link technology to transport a 100 Gb/s payload does not meet density requirements discussed previously. Past experience with 40 Gb/s technology shows that the technology sweet spot is achieved when four serial electrical links are used to transport the payload through the system to/from the optics. This implies that CEI electrical requirements must support baud rates up to 25 Gb/s in order to support transport a 100 Gb/s payloads across four links.

Higher baud rates are necessary for interfaces between Physical Layer devices. As shown in Figure 2, SFI is used to connect Framer, FEC, and Optics devices. Consistent with prior generations, it is expected that telecom applications will transport the 100 Gb/s payload within an OTU-n protocol as defined by ITU-T. Such protocols insert additional overhead for link management and Forward Error Correction. Current industry expectations are that 12% overhead is required to support a strong FEC code, and therefore CEI Short Reach interfaces need to support baud rates up to 28 Gb/s on the FEC to optics path.


Constraints and Trade-offs

Specification development requires trade-offs between channel requirements and Serdes requirements in order to achieve a cost-effective solution. Channel performance must improve in order to support baud rates required by the CEI-25 project, but forcing the solution to be entirely dependent on such improvements increases system costs unreasonably. More complex signal processing in the Serdes device can also solve the problem, but this increased Serdes complexity also increases chip power and area, and thus is also unacceptable.


Channel Considerations

Channel requirements must meet the following constraints in order to have a viable specification:

- Short Reach chip-to-chip and chip-to-module interfaces must support insertion loss consistent with channel length up to 12 inches and up to one connector. Long Reach interfaces must support insertion loss consistent with channel length up to 27 inches and up to two connectors in order to support the physical system configuration. Proposed loss limits are shown in Figure 3.



“SR” Channel Limits



“LR” Channel Limits

Figure 3 - Proposed Loss Limits

- Channel requirements must assume circuit board designs can be built with reasonable manufacturing yields. Requirements cannot be developed based on assumptions that exotic board materials, connectors, or manufacturing techniques will be used if those elements result in lower manufacturing yields and substantially increased costs.
- Next generation systems will target higher capacities. The underlying infrastructure will have to be overhauled to support higher aggregate bandwidths to provide a better performance to cost ratio, while simultaneously reducing overall space, power, and cost. Figure 4 demonstrates the impact of the serial line rate on the number of electrical pairs that would be needed at a switch card as the capacity of the system is varied from 400 Gb/s to 4 Tb/s. Signal routing densities, especially on backplanes, cannot be substantially reduced – otherwise density benefits of higher baud rates would be negated. This implies that the solution must budget for crosstalk, and that this crosstalk may come from arbitrary sources.




Figure 4 - Impact of Increased Backplane Capacity

- Channel requirements *can* assume circuit boards are designed by competent designers employing state-of-the-art techniques to minimize signal integrity impairments.

Serdes Considerations

Similarly, Serdes requirements must meet the following constraints in order to have a viable specification:

- Serdes power-per-link for 25 Gb/s must not exceed 1.5 times typical power-per-link for 10 Gb/s Serdes technology, since power distribution and heat dissipation are a very significant concern in most systems. Substantial increases to power-per-link would overly challenge the system design, and lead to substantially increased system cost and/or limit density. However, adherence to this power requirement imposes limits on the complexity of signal processing which may be considered as part of the solution.
- Electrical specifications can compensate for insertion loss by increasing signal launch amplitudes. However, Serdes launch amplitudes must not exceed levels which can be driven by CMOS circuits using a 1.2V power supply. Sub-micron CMOS process technologies are used to implement Link Layer and Physical Layer chips (with the exception of devices on the optics module). CMOS chips use a 1.2V power supply to power most of the circuits on the chip. Serdes transmitters use this supply to drive serial signals; using a higher voltage would increase power consumption and dictate additional mask layers (increasing cost).

Signaling Explorations

Most serial link standards for 10 Gb/s legacy baud rates are based on equalized NRZ signaling. NRZ signaling uses two signal levels to encode '0' and '1' bit values. Equalized NRZ signaling additionally uses equalizer circuits in both the transmitter and receiver to correct for various signal impairments induced by the channel. Feed Forward Equalizer (FFE) circuits are typically included in the transmitter, and Decision Feedback Equalizer (DFE) circuits are typically included in the receiver. It should be noted

that the FFE in effect produces a multi-level signaling code where the transmitted signal amplitude is not only based on the current bit value, but also values of adjacent bits. In addition, the DFE varies the slicing threshold of the receiver based on the history of prior bits received.

OIF member companies have contributed extensive analysis of Duobinary, PR-3, and PAM-4 multi-level signaling schemes for the CEI-25 project. In general, equalized NRZ has been found to perform either equivalently or better than other signaling schemes.² In order to understand this result, it is necessary to understand that equalized NRZ signaling using an n -tap FFE is in effect generating a multi-level signal with 2^n provisionable signal levels. PR- n multi-level signaling (where a PR- n code has $n+1$ signal levels) can be generated by an n -tap FFE.³ PR- n codes are therefore special cases within the optimization search space for an equalized NRZ signaling system. It can be shown that when presented with special cases of channel characteristics where multi-level signaling is expected to have an advantage, an equalized NRZ signaling system produces equalizer settings which effectively result in PR- n signaling.⁴

Based on the results of the above analysis, OIF has chosen to proceed with equalized NRZ signaling as the basis of an electrical interface specification for 28 Gb/s Short Reach applications. This signaling approach is sufficient to meet system needs given the proposed insertion loss curve for short reach channels in Figure 3.

As of the time of this writing, OIF is still investigating signaling for 25 Gb/s Long Reach applications. The higher insertion loss associated with long reach channels remains a challenge for equalized NRZ signaling. However, based on analysis performed to date, it is clear that solutions for this application space must draw on approaches which are not inherently a subset of the equalized NRZ system capabilities. Multi-level signaling schemes analyzed to date have not proven to offer any advantage; innovation is required.

Summary

The development of 100 GbE and its future transport over the optical transport network is driving the need for the entire component level infrastructure to be upgraded. The OIF authorized its PLL Working Group to begin the CEI-25 project, which will define electrical specifications for 28 Gb/s signaling for chip-to-chip and chip-to-module applications and 25 Gb/s signaling for backplane applications.

The work being undertaken in the OIF will provide the future basis for the next level in electrical signaling. Companies interested in participating in the CEI-25 project are invited to contact the OIF about membership.

This paper is the collaborative effort of the following members of the OIF:

Beth Donnay, Force10 Networks

David R. Stauffer, IBM Microelectronics

² “Multi-level Signaling in High-density, High-speed Electrical Links”, D.G.Kam, et al., DesignCon 2008.

³ Oif2008.099, “Extended NRZ Coding Technique”, J.Goergen, et al.

⁴ oif2008.114, “Link Simulations for Partial Response Signaling”, D.G. Kam and D.R.Stauffer.