



OPEN
Compute Project

ELECTRICAL INTERFACES PERFORMANCE METRICS

Shahab Ardalan (Enosemi)

Kevin Donnelly (Eliyan Corporation)

James Wong (Palo Alto Electron, Inc.)

Marc Greenberg (Blue Cheetah Analog Design, Inc.)

Boon Chong Ang (Intel)

Majid Foodeei (Kandou)

Ahmad Tavakoli (Yorchip)

Anu Ramamurthy (Microchip)

Armin Tajalli (Kandou)

Bapi Vinnakota (OCP)

Cliff Grossner (OCP)

Elad Alon (Blue Cheetah Analog Design, Inc.)

Keivan Khasraghi (Synopsys)

Kevin Yee (Samsung)

Madhumita Sanyal (Synopsys)

Quinn Jacobson (CMU)

Trent Uehling (NXP)

Greg Taylor (Greg Taylor Consulting)

Executive Summary

The Die-to-Die interface integrates the Link Layer and the Physical Layer to ensure smooth data transmission across networks, each operating at different levels of the communication stack. The Link Layer handles reliable data transmission between directly connected devices by organizing data into frames, managing error detection and correction, and controlling data flow to prevent congestion. It also coordinates access to the physical transmission medium, enabling efficient sharing of network resources.

The Physical Layer, as the base layer of the communication stack, converts digital information from the Link Layer into physical signals for transmission over various media. It manages the electrical, optical, or radio frequency aspects of transmission and defines hardware specifications like connectors, pin configurations, and signal types.

This white paper focuses on standardizing the performance metrics of the Physical Layer, particularly in the context of Electrical Interface Physical Layer Intellectual Property (PHY IP). PHY IP is crucial for enabling seamless data exchange in modern electronic systems across interfaces like BoW, UCIe, PCIe, and Ethernet. The growing demand for high-speed, low-power, and versatile PHY IPs highlights their importance in developing advanced electronic systems for a range of applications.

The document aims to establish a unified framework for measuring and standardizing the performance and functionality parameters of PHY IP. Standardization is essential for ensuring interoperability, simplifying integration, enabling compatibility, promoting quality assurance, reducing development time, fostering industry collaboration, and ensuring regulatory compliance. By addressing these factors, the paper supports the development of more sophisticated and reliable electronic systems through the advancement and standardization of PHY IP practices.



Table of Contents

Executive Summary	2
Table of Contents	3
Introduction	5
Objective of this Document	5
1. General Information	6
1.1. Type (Serial or Parallel)	6
1.2. Clocking/Data Rate: (SDR/DDR/QDR)	7
1.3. Logical Link (Unidirectional Simplex / Bidirectional Simplex / Duplex)	8
1.4. Single-Ended, Differential, and Multi-Wire Signaling	10
1.5. Signaling type (NRZ/PAM2, PAMn, Multi-Wire PAMn)	12
1.6. IO swing type (rail2rail, partial)	13
1.7. Clocking (forwarded/recovered)	14
1.8. Termination	15
1.9. Specification and IP Maturity	16
1.10. Publicly announced customers / Reference users	17
1.11. Access to Specification (Open Access, Registration, Membership)	17
1.12. IP License (open design, fee-free / fee)	18
1.13. Organization (if open)	18
1.14. Reference (to standard, must if open)	18
1.15. Interoperability	18
1.16. Packaging technology (organic, bridge, interposer, fanout, 3D)	19
2. Unit Bidirectional Link Parameters	20
2.1. Number of data lanes	20
2.2. Number of data wires	20
2.3. Number of clock wires	21
2.4. Number of other wires	21
2.5. Mandatory Side Band (number of lanes, if mandatory)	22
2.6. Targeted speed per data lane	22
2.7. Example table of Link Parameters	23
3. Reference Implementation Related Parameters	24
3.1. Process Technology	24
3.2. Number of power supplies and values	24
3.3. IO swing value Vpp, single-ended	25
3.4. Interface Physical Dimensions	25
3.5. Min Data Pad spacing (X1, Y1, Pdiag, P)	27
3.6. Unit Beachfront per die	28
3.7. Unit Depth per die	28
3.8. NS/EW Orientation	29
3.9. Total Area (X & Y)	29
3.10. Transmitter Receiver, and clocking Area	29
3.11. How many TX/RX macros can share one clocking block	30
3.12. Max Stacked units	30
3.13. Number of substrate layers for reported edge density	31



3.14. Latency and BER	31
3.15. Reach	35
4. Calculations and Figures of Merit (FOMs)	36
4.1. Unit Bandwidth	36
4.2. Unit Power	36
4.3. Technology Power Efficiency	38
4.4. Edge Bandwidth Density	38
4.5. Area Bandwidth Density	39
4.6. Technology FOM1	39
4.7. Technology FOM2	40
5. Recommended Minimum List of Parameters	41
6. Conclusion and Next Steps	41
7. License	42
8. About Open Compute Foundation	43
9. Compliance with Open Compute Project Tenets	43



Introduction

The Electrical Interface Physical Layer Intellectual Property (PHY IP) serves as a crucial component in modern electronic systems, facilitating the seamless communication and data exchange between different devices. This specialized IP module operates at the physical layer of the communication stack, defining the electrical characteristics, signaling, and connectivity protocols necessary for reliable data transmission. It plays a pivotal role in ensuring compatibility and efficiency in various interfaces, such as BoW, UCIe, PCIe, Ethernet, and more. As technology evolves, the demand for high-speed, low-power, and versatile PHY IPs continues to grow, making them essential building blocks for the development of advanced electronic systems across a wide range of applications, from consumer electronics to industrial automation and beyond.

Objective of this Document

The primary objective of this white paper within the realm of Electrical Interface Physical Layer Intellectual Property (PHY IP) is to meticulously identify a comprehensive list of pertinent parameters associated with the performance and functionality of such IP. Concurrently, the paper aims to propose a standardized and widely accepted method of measurement tailored specifically to these identified parameters. By undertaking this dual mission, the white paper seeks to establish a unified framework that not only provides a nuanced understanding of the intricacies inherent in Electrical PHY IP but also advocates for a consistent language and methodology in the measurement and assessment of these critical parameters. This concerted effort is motivated by the overarching goal of fostering greater coherence, comparability, and interoperability within the field, thereby contributing to the advancement and standardization of practices specific to Electrical PHY IP.

A common method of parameter measurement and performance reporting in electronics Intellectual Property (IP) is essential for several reasons:

1. **Interoperability:** A standardized approach ensures that different IP components from various vendors or sources can seamlessly work together within a larger system. This interoperability is crucial for the successful integration of diverse IP blocks in complex electronic designs.
2. **Ease of Integration:** Common measurement and reporting methods simplify the integration process for designers. When parameters and performance metrics are consistently defined and reported, engineers can more easily assess the compatibility of different IP components, reducing the risk of integration challenges.
3. **Comparability:** Standardized measurement methods enable direct comparisons between different IP offerings. Engineers can evaluate the performance of various IP solutions objectively, aiding in the selection of components that best meet the requirements of a particular application.
4. **Quality Assurance:** A standardized approach to parameter measurement and reporting establishes a baseline for quality standards. It allows IP developers to communicate and adhere to industry-accepted benchmarks, promoting consistency and reliability in the performance of electronic components.
5. **Reduced Development Time:** With a common method in place, engineers can accelerate the development process by avoiding the need to decipher and adapt to varying



This work is licensed under a [Creative Commons Attribution-ShareAlike 4.0 International License](https://creativecommons.org/licenses/by-sa/4.0/).

measurement and reporting practices. This, in turn, leads to more efficient design cycles and faster time-to-market for electronic products.

6. **Industry Collaboration:** Common standards foster collaboration within the electronics industry. By providing a shared framework for parameter measurement and reporting, different entities can collaborate more effectively, spending less time agreeing on measurements and more time driving innovation and advancing the state of the art in electronic design.
7. **Regulatory Compliance:** In certain industries, adherence to standardized measurement and reporting methods may be necessary for regulatory compliance. Common standards help ensure that electronic products meet regulatory requirements and industry specifications.

In summary, a common method for parameter measurement and performance reporting in electronics IP is fundamental for promoting interoperability, facilitating integration, supporting comparability, ensuring quality, reducing development time, encouraging collaboration, and meeting regulatory standards within the electronics industry.

1. General Information

This section contains descriptions of the general information entered for die interfaces.

1.1. Type (Serial or Parallel)

Parameter	Valid Values	Units
Type (Serial or Parallel)	Serial, Parallel	n/a

The Electrical Interface PHY IP differentiates significantly between serial and parallel interfaces, each with distinct technical features concerning clocking, error correction, equalization, and data management.

Serial Interface: Serial interfaces transmit data sequentially over a single channel, notably reducing wiring needs. This is partly due to the elimination of a dedicated clocking wire, as serial systems embed clock information within the data signal. Typically, these are differential channels, enhancing signal integrity, especially against electromagnetic interference. In specific scenarios, such as in cost-sensitive applications, single-ended implementations are used. The absence of a separate clock line necessitates Clock Data Recovery (CDR) mechanisms at the receiver end, adding to design complexity.

Advanced techniques like error correction and signal equalization become crucial in long-distance serial links. Error correction addresses data corruption risks over long distances, while equalization combats signal attenuation and inter-symbol interference, vital for high-speed data integrity.

Parallel Interface: Parallel interfaces handle multiple data bits simultaneously across several channels, optimizing throughput for short distances. While generally single-ended, differential



implementations exist for noise-sensitive applications. A forwarded clock signal in parallel interfaces aids synchronization, circumventing the need for CDR. However, this necessitates precise phase matching and deskewing to ensure data bits' simultaneous arrival at the receiver. Techniques such as adjustable delay lines or initialization calibration (“training”) routines are employed to manage timing discrepancies due to varied electrical characteristics in parallel lines.

The clock signal's overhead in parallel interfaces is distributed across multiple lines, reducing its impact. However, issues like crosstalk and signal integrity degradation become significant with increased distances, particularly for single-ended signaling, making parallel interfaces less ideal for long-distance communication. Single-ended signaling also exacerbates EMI issues for longer-distance connections.

Selecting between serial and parallel interfaces in Electrical Interface PHY IP involves evaluating distance, data rate, complexity, and clock synchronization requirements. Serial interfaces are often preferred for their scalability in high-speed, long-distance applications (where the cost of the channel itself may be significant), while parallel interfaces are relevant for high-bandwidth, short-distance scenarios and lower power communications. This choice reflects a balance between technical feasibility and specific application needs.

1.2. Clocking/Date Rate: (SDR/DDR/QDR)

Parameter	Valid Values	Units
Clocking/Date Rate	SDR, DDR, QDR	n/a

SDR (Single Data Rate), DDR (Double Data Rate), and QDR (Quad Data Rate) are terms associated with the clocking mechanisms used in computer memory technologies.

SDR employs a straightforward approach where data is transferred once per clock cycle. This means that one edge (rising or falling) of the clock signal corresponds to one piece of data being transmitted.

DDR, on the other hand, doubles the data transfer rate by transmitting data on both the rising and falling edges of the clock signal. This effectively allows for twice the amount of data to be transferred per clock cycle compared to SDR. DDR also means that the clock and the data have approximately the same bandwidth and the same needs for equalization.

QDR takes the concept of increased data transfer rates even further. As the name suggests, it operates at four times the rate of the base clock, transmitting data like DDR on both rising and falling edges of the clock signal (at 0o and 180o) but also introducing additional clock phases (at 90o and 270o) to enable the transmission of data in four distinct steps within a single clock cycle. This results in a more efficient use of the clock signal and delivers higher performance compared to SDR and DDR data transfer methods, although with the tradeoff of added complexity required to generate the additional clock phases needed.




Figure 1: SDR, DDR, and QDR Clocking/Data Rate

1.3. Logical Link (Unidirectional Simplex / Bidirectional Simplex / Duplex)

Parameter	Valid Values	Units
Logical Link	Unidirectional, Bidirectional, Dynamic Bidirectional, Simultaneous Bidirectional	n/a

In the field of Electrical Interface PHY IP, logical link configurations are integral, encompassing unidirectional (simplex), bidirectional (dual simplex), dynamic bidirectional (half-duplex), and simultaneous bidirectional (full-duplex) communication modes. Each presents distinct technical features, suited to specific operational needs.

Unidirectional (Simplex): Unidirectional (simplex) communication is characterized by one-way data transmission. It employs a single transmitter on one chip and a single receiver on another chip, and lacks mechanisms for receiver feedback to the transmitting chip. Typically utilized in control interfaces and older display interfaces where commands are issued without need for acknowledgment, unidirectional simplex is prized for its simplicity. However, it's limited by the unavailability of two-way communication.

Bidirectional (Dual-Simplex): Bidirectional (dual-simplex) enables simultaneous two-way data transmission through two distinct channels, one dedicated to transmission and the other to reception, allowing concurrent operations. Bidirectional links are the most commonly used in die-to-die and chip-to-chip communications, facilitating continuous and seamless two-way communication and providing symmetric data bandwidth in both directions, essential for applications where high-speed data exchange and minimal latency are crucial.

Dynamic Bidirectional (Half-Duplex): Dynamic Bidirectional, or half-duplex, permits two-way communication, but not concurrently. A shared channel alternates between transmitting and receiving states, governed by protocols that manage this switch. Such

half-duplex signaling is also called dynamic bidirectional signaling, and is widely adopted in memory and USB interfaces that alternate between Read and Write data traffic. The primary technical challenge in half-duplex systems is the management of turnaround time and maintaining synchronization between transmission and reception. While more complex than unidirectional simplex, half-duplex systems offer improved bandwidth efficiency.

Simultaneous Bidirectional (Full-Duplex): Simultaneous Bidirectional, or full-duplex communication, enables simultaneous two-way data transmission over the same channel. Simultaneous bidirectional systems are the most sophisticated, requiring the application of echo-cancellation and other techniques to enable simultaneous two-way communication over a single channel, but result in improved bandwidth efficiency over standard bidirectional communication using separate channels dedicated to transmission and reception.

The selection of communication mode in Electrical Interface PHY IP hinges on factors like required data throughput, latency, system complexity, and specific application needs. Unidirectional (simplex) signaling can be used for simpler, one-way applications; bidirectional (dual-simplex) signaling is used for scenarios requiring cost-effective two-way communication with moderate performance needs; dynamic bidirectional (half-duplex) signaling is often used for memory applications, and simultaneous bidirectional (full-duplex) signaling can be used for high-end applications demanding robust and continuous bidirectional communication.

These logical link types are pivotal in defining the operational scope and capabilities of Electrical Interface PHY IPs, influencing system design and application suitability.




Figure 2: Logical Link Terminology

1.4. Single-Ended, Differential, and Multi-Wire Signaling

Parameter	Valid Values	Units
Signaling	Single-ended, Differential, Multi-Wire	n/a

Single-ended signaling is the most basic way of transferring data between two points using copper as the communication media. In this approach, a single wire is required to transfer a stream of binary bits. There is a continuous demand to increase the overall throughput in serial data communication systems. As a result, there is a tendency to increase the number of bits transferred over every symbol (defined by the modulation scheme), and also maximize the usage of the available number of wires. Pin-efficiency or wire-efficiency is defined to be:

$$\text{PE (or WE)} = \text{Number of Bits-per-UI} / \text{Number of Pins (or Wires)}$$

With the fewest number of wires per signal, single-ended signaling offers a higher pin-efficiency compared to differential signaling. The presence of strong inter-wire crosstalk (XTALK), common-mode noise (CMN), and supply noise (SN) can limit the usefulness of single-ended links to applications in which all those limiting factor (XTALK, CMN, and SN) are well controlled. Applications such as chiplets, very short-reach die-to-die and memory links, are among them.

Differential links are used in applications where the influence of XTALK, CMN, and SN cannot be neglected. The cost of differential signaling, however, is reducing the overall throughput of the system by a factor of two compared to single-ended links at the same data rates. In long-reach applications, where channel loss is high, and the received signal is weak and its SNR is very low, differential links are pervasive due to their inherent resistance against any type of common-mode noise. Example application areas are Ethernet, USB, and JESD, and PCIe links.

Multi-wire links have been proposed as an alternative to the differential links to enhance the throughput while maintaining a certain level of resistance against different types of common-mode noise. In such approaches, some precoding techniques will be employed to keep the signal transitions balanced (or partially balanced). This way, link sensitivity to XTALK and CMN can be reduced while the overall pin-efficiency of the link improves.

Channel Response

Figure 3 shows the frequency domain response of a typical wireline channel. The most important parameters to characterize such channels are frequency dependent channel loss, crosstalk (XTALK) and reflection (Return Loss). To determine the right signaling scheme for an application, it is important to take into account channel response. Generally, proper equalization is required to address the channel frequency dependent loss, also called insertion loss (IL). Adequate matching network is essential to keep the channel reflection under control. To control crosstalk, physical design of the channel is critically important. For example, increasing the spacing in between two adjacent conductors helps to reduce the crosstalk. Balanced signaling schemes, such as differential signaling, can be employed to suppress the crosstalk.






Figure 3: (a) Differential Signal and (b) Channel Response

Single-ended signaling (SE): Single-ended signaling is the simplest signaling scheme and uses a single wire ($n=1$) to transmit the information data. The voltage D is measured against the reference common voltage (global “ground”) VCM.

Differential signaling (DS): Differential signaling is shown in Figure 2, uses two complementary signals ($D+$ and $D-$) over two wires ($n=2$) for each information link. This signaling scheme sends the same data information as a differential pair of signals, each on its own wire. In DS, $D+$ reference is $D-$, the negative of the signal which is also transmitted. This differential routing both generates less emitted noise and can reduce the sensitivity of the pair to noise from external sources.

Multi-Wire Balanced Signaling: Multi-wire balanced signaling has been developed to maintain many of the advantages of differential signaling (i.e., rejecting common-mode sources of noise) with fewer wires. There are different ways to realize multi-wire balanced signaling schemes. Chord signaling is a multi-wire balanced signaling scheme, which uses orthogonal coding and can be considered as generalized DS but with pin-efficiency that approaches that of SE. Offering a higher pin-efficiency compared to DS, Chord signaling suppresses CMN and XTALK without sacrificing ISI sensitivity, like DS. Chord signaling has been standardized (OIF CEI-5) for both long-reach (called Ensemble NRZ, ENRZ) and short-reach (Correlated NRZ, CNRZ) applications.

1.5. Signaling type (NRZ/PAM2, PAMn, Multi-Wire PAMn)

Parameter	Valid Values	Units
Signaling Type	NRZ, PAMn, Multi-Wire PAMn, ...	n/a

Non-Return-to-Zero (NRZ), is a family of signaling schemes that symbols carry information (data) without going back to zero or reset state in between two consecutive symbol intervals. Unlike NRZ, in RZ signaling, at the end of each symbol period, data goes back to zero for a certain amount of time. The term NRZ has been misused in industry for Pulse Amplitude Modulation 2-level (PAM2). However, all different types of PAM signaling, including PAM3 and PAM4 are part of the NRZ signaling family. If we use the term NRZ for PAM2, then it refers to a binary code using 2 levels (high and low) to represent 1 and 0 information content. An NRZ based system transmits 1 bit per symbol period (1 bit/ Symbol). While differential NRZ are widely used in long-reach applications, ultra short-reach die-to-die links use single-ended NRZ to carry data to enhance the overall throughput. The maximum energy efficiency can usually be achieved using single-ended NRZ running at the minimum voltage level that a technology can support.




Figure 4: RZ vs NRZ Signaling

Pulse Amplitude Modulation (PAMn) is a family of multilevel signal modulation schemes. 4-level (PAM4) is a multilevel (4) signal modulation signaling technique where each signal level represents 2 bits of information encoded in their amplitude. As seen in Figure 4, a PAM4 signal includes three distinct eyes, corresponding to the four expected voltage levels (00, 01, 11, 10). It is obvious that the available eye height margin in PAM4 is three times smaller than an NRZ (PAM2) link, while the number of bits transferred over the same baud rate will be two times more:

$$\text{Bit/ Second} = (\text{Bit/ Symbol}) * (\text{Symbol/ Second})$$

3-level (PAM3) is another common amplitude modulation scheme that uses three values (-1, 0, +1) to encode 1.5 bits of information per cycle. Where PAM4 has three distinct eyes, PAM3 has two distinct eyes.

As an example, PCIe G1 to G5 use differential NRZ signaling, while PCIe Gen6 employs differential PAM4 and GDDR7 uses single-ended PAM3.

PAM signaling (PAM2, PAM3, PAM4, or any other order of PAM) can be combined with multi-wire signaling to enhance link performance.

Multi-Wire PAMn: PAMn Signaling can be done over a single wire, two wires, or more than two wires (multi-wire PAMn). Example of single wire PAM2 is SE Signaling in BoW 2.0 Die-to-die specification. Example of 2-wire PAM4 is PCIe 64 GT/s in PCI-SIG Specification. Example of 4-wire PAM2 Chord signaling is ENRZ, specified in OIF CEI-5 Specification, and carries data over four wires with three PAM2 differential subchannels (3 bits over 4 wires).

1.6. IO swing type (rail2rail, partial)

Parameter	Valid Values	Units
IO Swing Type	Rail2Rail, Partial	n/a

IO swing is an important parameter in the design of serial communication systems. The IO swing can span the entire voltage available (supply rail-to-rail) as with standard CMOS drivers, or can be a partial swing as is common with current-steering driver IOs. A higher I/O swing results in more eye margin and thus a better available SNR. Moreover, a higher I/O swing can relax the design of receiver analog front-end as a lower voltage gain will be required. Proper circuit design techniques, such as a combined current-mode and voltage-mode transmitter output driver circuit, can be employed to increase the I/O swing. The other possibility is to choose a proper termination resistance value for the receiver and transmitter to increase the I/O swing. One possibility, for example, is to reduce the termination resistor on the transmitter side and increase it on the receiver end. Of course, such techniques can be applied when the effect of channel reflection remains under control and does not degrade the link quality.

While adequate channel matching is essential in high-speed links operating over high-loss channels, the requirements can be relaxed in very short-reach systems, such as in chiplet and die-to-die links. Different aspects of imbalance termination structures have been explored for different applications, with the goal to find a proper balance between link quality (controlled reflections), link power, and circuit design complexities. Some applications employ a higher termination resistor on the transmitter end (compared to the receiver termination), mainly to save power in the output driver. However, the drawback is their smaller vertical opening. On the other side, one can reduce the termination resistance on the transmitter end to increase the eye size at the cost of higher line power.

An optimal design point shall be devised based on application and the tradeoffs between circuit power, I/O swing, and the effect of reflections. When silicon is designed to operate across a range of channels (length, loss, and mixed vendors) training may be needed to choose the best set of operating parameters for each individual system; offering higher data



rates, lower energy per bit, and possibly both when compared to a design with fixed parameters.

1.7. Clocking (forwarded/recovered)

Parameter	Valid Values	Units
Clocking	CCK, FCK, CR	n/a

IO Clocking Architectures can be divided into the following categories

- Common Clock (Synchronous)
- Forward Clock (Source Synchronous)
- Embedded Clock (Clock Recovery)




Figure 5: Common Clock, Forwarded Clock, and Embedded Clock

Common-Clock (CCK) Scheme

In the simplest synchronous clocking scenario, a common reference clock is provided to two devices (Device A and Device B). Utilizing identical multiplication factors can ensure that the interfaces of the two devices run at the same frequency, but the phase relationship between the clocks on Device A and Device B will be unknown. For such cases, typically elastic buffers or FIFOs are used to cross between the clock domains with unknown phases, which add to latency. Also, there is limited jitter tracking between the two sides.

Forwarded-Clock (FCK)

In a forwarded clock scenario, a clock is forwarded from one device to another (from Device A to Device B) to provide the clocks on both devices. Having device A provide a clock to device B can improve jitter tracking as both clocks have the same source, but unless the data paths and clock paths are identical the phase relationship on the two devices will not be identical, resulting in the need for clock crossing logic again and proper clock/data alignment circuits. The overall jitter tracking of FCK is better than CCK, so the settling time is lower.

Embedded-Clock (Clock Recovery, CR)

In longer-distance links (for example, between boards in a system), a common clock cannot be provided to the two devices (Device A and Device B), so the clock frequencies may be different. As long as the two devices use similar clock sources, the frequency offset may be small (usually measured in ppm), and embedded clock recovery logic can be utilized to extract the clock information from the data. This approach is widely used in serial links such

as PCIe. While this clocking approach enables long-reach links, the jitter tracking BW is low compared to the forwarded clock schemes, which results in slow lock times.

1.8. Termination

Parameter	Valid Values	Units
Termination	Unterminated, Source terminated, Load terminated, Doubly terminated	n/a

PHY (Physical Layer) I/O termination is a fundamental aspect of the physical layer interface in a communication system, designed to ensure the integrity of signals as they are transmitted across a medium. Proper termination is pivotal for maintaining signal quality by minimizing reflections, and thereby preventing signal degradation or distortion that could compromise the performance of the communication system.

Unterminated Links: This approach does not employ any termination resistors at either end of the transmission line. While this simplifies the design and reduces the components required, it makes the system susceptible to signal reflections, especially in high-speed applications, leading to potential signal integrity issues.

Source Terminated: In source termination (also known as series termination), a resistor is added at the transmitting end of the line. The resistor's value is chosen to match the characteristic impedance of the transmission line, effectively reducing reflections towards the source. This technique is particularly beneficial in scenarios where the transmission line drives signals to multiple receivers, as it helps preserve signal integrity while minimizing power consumption.

Load Terminated: In load termination (also known as parallel termination), a termination resistor is placed at the receiving end, matching the line's characteristic impedance to absorb the transmitted signal energy fully and prevent reflections. Load termination is essential in systems where the accuracy and quality of the received signal are critical, even though it may result in increased power dissipation at the receiver.

Doubly Terminated: Combining both source and load termination, doubly terminated links feature resistors at both ends of the transmission line. This comprehensive approach ensures maximal signal integrity by mitigating reflections from both directions. Although it incurs either reduced swing or higher power consumption, it is indispensable in complex systems where maintaining optimal signal quality is paramount.

The choice of termination method depends on various factors, including the signaling standards, the nature of the transmission medium, system requirements, and the specific design goals of the communication interface. Each termination strategy offers a different set of trade-offs in terms of complexity, cost, signal integrity, and power efficiency. Designers



must carefully consider these factors to select the most appropriate termination technique for their specific application.

PHY I/O termination techniques are critical for achieving high-performance communication, ensuring that the interface meets the requisite standards for signal integrity and reliability. By judiciously applying these termination methods, designers can optimize the performance of their communication systems, ensuring efficient, accurate, and high-speed data transfer.

1.9. Specification and IP Maturity

The following table lists the stages for IP maturity, starting with the definition for least mature:

Status	Impact
Maturity	
Specification under development	The specification hasn't been finalized and it is work in progress.
Specification available, but no implementations announced	The specification has been finalized but there is no IP (silicon) based on the specification.
IP available	When IP is available, no changes would be expected unless the specification changes.
Many Vendors	When IP is available from multiple vendors, the specification has higher maturity. However, a vendor may make implementation choices that impact interoperability.
Silicon Status	
No implementation	The IP has not been implemented in silicon in any process node or in the target process node for the product.
Soft IP	The design is available as a verified soft IP, but has not been implemented in any silicon.
Test chip	The IP has been implemented in a test chip Even better if the test chip is in the process node targeted by a design.
Product sampling	A product incorporating the IP is currently sampling. Even better if a product incorporating the IP in the target process node is currently sampling.

High volume manufacturing	One or more products incorporating the IP are currently being produced at high volume. One or more products incorporating the IP in the target process node are currently produced at high volume.
---------------------------	--

For companies considering licensing PHY IP cores, having many IP vendors with the same PHY IP proven on the desired target process node gives the highest degree of flexibility, especially when performance is not a priority. If proven PHY IP is not available on the desired process node, the next best would be to have PHY IP available at test chip on another process node, followed by PHY IP being under development. For companies where PHY performance is the highest priority, customization of existing PHY IP will often be necessary.

1.10. Publicly announced customers / Reference users

Parameter	Valid Values	Units
Publicly announced customers Reference users	Number of names	n/a

This information highlights whether any customers or reference users have been publicly announced for the chiplet. Having publicly announced customers or users can serve as a validation of the chiplet's capabilities and market acceptance.

1.11. Access to Specification (Open Access, Registration, Membership)

Parameter	Valid Values	Units
Access to Specification	Open Access, Registration, Membership	n/a

This detail specifies whether the interface specification is openly accessible to the public and potentially collaborative, or if access requires registration or membership, which may be paid or unpaid.

1.12. IP License (open design, fee-free / fee)

Parameter	Valid Values	Units
IP License (fee-free / fee)	Open design, Fee-free, Fee	n/a

This indicates whether the intellectual property (IP) associated with the interface specification is available through open source (open design), or under a fee-free (no licensing) or fee-based license. It's important to note that IP licensing is distinct from access to the specification.

1.13. Organization (if open)

Parameter	Valid Values	Units
Specification Organization	Organization name	n/a

If the chiplet technology or specifications are open, this detail specifies the organization or entity responsible for overseeing the open aspects. This organization may manage community contributions, maintain the specifications, and provide support to users.

1.14. Reference (to standard, must if open)

Parameter	Valid Values	Units
Specification Reference	Standard version number	n/a

If the chiplet follows open standards or specifications, this detail specifies the reference or version to those standards. Compliance with open standards ensures interoperability and compatibility with other chiplets and systems following the same standards.

1.15. Interoperability

Parameter	Valid Values	Units
Interoperability	Yes / No Compliant to other PHYs	n/a



This section highlights if the PHY/D2D IP is compliant with a specification and can interoperate with different vendor implementations. Expected values are Yes/No and compliant to other PHYs if any. Any spec limitations/unimplemented features should be documented. For example, if the vendor will provide the chiplet phy simulation model and validate for spec compliance, and address the interoperability method use.

1.16. Packaging technology (organic, bridge, interposer, fanout, 3D)

Parameter	Valid Values	Units
Packaging technology	Organic, Bridge, Interposer, Fanout, 3D	n/a

For Electrical Interface PHY IP, the PHY interconnect physical design has broad implications on the package technologies available for integration of chips and chiplets. Package technology, in turn, drives the cost of chiplet integration. To aid in specification of applicable package technologies for a given die-to-die physical interface, four categories are defined as shown in Figure 6.

Level	Package Technology	Attributes	Example Illustration
P1	Organic Substrate	Chiplets in 2D side-by-side arrangement. Build-up or laminate organic dielectric. Interconnect by solder bump or Cu pillar.	
P2	Fan-Out Package	Chiplets in 2D side-by-side arrangement. RDL interconnect. Die may or may not be bumped.	
P3	Si Bridge Interconnect	Chiplets in 2D side-by-side arrangement. Si or glass bridge embedded in substrate. Interconnect by solder bump or Cu pillar.	
	Si or Glass Interposer	Chiplets in 2D side-by-side arrangement. Interposer interconnect. Interposer on package substrate.	
P4	3D Stack	Chiplets stacked in 3D configuration. μ -bump or hybrid bond interconnect.	

Figure 6: Four Chiplet Package Categories

Each of these four package categories has specific characteristics that align to chiplet interface physical design attributes. While these characteristics serve as general guidelines, there are no hard lines between package technology capabilities. For instance, an interface with a 60 μ m pitch bump pattern may be suitable for interposer, bridge, fan-out and organic substrate technologies. However, the package routing layers required for organic substrate technology are likely higher count than the other package types. For this reason, other considerations must be taken into account to accurately specify the applicable package technologies for a given interface. These additional considerations are discussed in Sections 3.4 through 3.14.



2. Unit Bidirectional Link Parameters

Multiple terms are used to describe the number and types of connections used in Electrical Interface PHY IP links. These terms are defined for bidirectional (dual simplex) links.

2.1. Number of data lanes

Parameter	Valid Values	Units
Number of data lanes	Number	n/a

The "number of data lanes for chiplet input/output (I/O)" refers to the specific count of electrical channels or paths used for transmitting data between a chiplet and its host or other interconnected components within a semiconductor device or system. These data lanes are part of the chiplet's I/O interface and are primarily responsible for carrying data signals.

Each data lane typically consists of one or multiple wires or signal lines used to transmit data in parallel. The exact number of data lanes for chiplet I/O can vary depending on the chiplet's design and intended functionality. It is a critical parameter in determining the data transfer bandwidth between the chiplet and the rest of the system.

The number of data lanes is entered based on bidirectional links. For example, a PHY with 16 Tx + 16 Rx lanes would be entered as 32 data lanes.

2.2. Number of data wires

Parameter	Valid Values	Units
Number of data wires	Number	n/a

The term "number of wires for chiplet I/O" refers to the total count of individual electrical conductors or signal lines that are used to transmit data, commands, and control signals between a chiplet and its host or other interconnected components within a semiconductor device or system. These wires are a critical part of the I/O interface of a chiplet and are responsible for carrying information to and from the chiplet.

The actual number of wires for chiplet I/O can vary widely depending on the chiplet's design, functionality, and the specific communication requirements. It includes wires for data transmission (data lanes), control and configuration signals, and other necessary connections. The total wire count is a crucial aspect of chiplet design, as it influences factors like data transfer speed, power consumption, and the overall complexity of the chiplet and its interface.



The number of data wires is entered based on bidirectional links. For example, a PHY using single-ended signaling that has 16 Tx + 16 Rx lanes would be entered as 32 data wires.

2.3. Number of clock wires

Parameter	Valid Values	Units
Number of clock wires	Number	n/a

The "number of clock wires for chiplet I/O" refers to the count of electrical signal lines or wires used to transmit clock signals between a chiplet and its host or other interconnected components within a semiconductor device or system. Clock wires are crucial for synchronizing the operation of different parts of the chiplet and ensuring that data is transferred and processed at the correct timing intervals.

Clock wires are essential for maintaining the timing and synchronization of various operations within a chiplet. They help coordinate activities such as data sampling, signal processing, and data transmission. The specific number of clock wires needed for chiplet I/O can vary depending on the chiplet's design, functionality, and communication requirements.

In many cases, chiplets may use one or more dedicated clock wires for different purposes, such as data transmission, control signals, and synchronization with the system's clock domain. The exact number of clock wires for chiplet I/O is typically determined by the chiplet's specifications and its interface standards.

The number of clock wires is entered based on bidirectional links. For example, a PHY that has 2 Tx clocks + 2 Rx clocks would be entered as 4 clock wires.

2.4. Number of other wires

Parameter	Valid Values	Units
Number of other wires	Number	n/a

The term "number of other wires for chiplet I/O" implies I/O other than data, clock, and sideband I/O. Example of other wires include the FEC, valid signals for miscellaneous functionalities such as control and status within the context of semiconductor chiplet interfaces or I/O communication. These "other wires" could include dedicated signal lines for specific functions or custom communication needs unique to the chiplet's application, or for the repair of lanes that fail to work after assembly.

The number of other wires is entered based on bidirectional links. For example, a PHY that has 2 Tx FEC + 2 Rx FEC wires would be entered as 4 other wires.



2.5. Mandatory Side Band (number of lanes, if mandatory)

Parameter	Valid Values	Units
Number of mandatory Sideband wires	Number	n/a

The term "number of sideband signals for chiplet I/O" refers to the count of additional electrical signal lines or channels that are used for transmitting auxiliary or non-data-related information between a chiplet and its host or other interconnected components within a semiconductor device or system.

Sideband signals are typically used for various purposes, such as control, configuration, monitoring, status reporting, or any other communication that is separate from the main data transfer. These signals are called "sideband" because they are separate from the primary data and clock signals used for data communication.

The specific number of sideband signals for chiplet I/O can vary widely depending on the chiplet's design, functionality, and the requirements of the system it is integrated into. The sideband signals provide a means for the chiplet and the system to exchange information beyond just data, enabling the chiplet to communicate its state, receive commands, or provide status updates.

The number of sideband wires is entered based on bidirectional links. For example, a PHY that has 2 Tx sideband + 2 Rx sideband wires would be entered as 4 sideband wires. If the sideband signals are not mandatory, this field is left blank.

2.6. Targeted speed per data lane

Parameter	Valid Values	Units
Target speed per data lane	Number	Gbps

The term "target speed per data lane for chiplet I/O" refers to the desired or intended data transfer rate or bandwidth for each individual data lane within the I/O interface of a chiplet. It represents the rate at which data can be transmitted or received over a single data lane, and it is typically measured in terms of data bits per second (bps) or gigabits per second (Gbps).

The target speed per data lane is a critical parameter in chiplet design and communication protocols, as it determines the overall data throughput and performance of the chiplet's I/O interface. It is important to ensure that the data lanes can meet the required speed to support the chiplet's functionality and the system's performance goals.



2.7. Example table of Link Parameters

The table below shows Link Parameters for two example D2D PHYs in standard packaging.

	Bunch of Wires (BoW) Two 16-bit slices 16 TX + 16 RX	Universal Chiplet Interface (UCIe) x16 Module 16 Tx + 16Rx
Data Lanes	32	32
Data Wires	32 (16 Rx + 16 Tx)	32 (16 Rx + 16 Tx)
Clock Wires	4 2 Rx and 2 Tx	4 2 Rx and 2 Tx
Other wires	4 TX: FEC, AUX RX: FEC, AUX (Optional)	4 Valid Rx, Track Rx, Valid Tx, Track Tx
Side Band	0	4 Rx SB_Clock, SB_Data, Tx SB_Clock, SB_Data
Targeted Speed	2~32 Gbps	4~32 Gbps

3. Reference Implementation Related Parameters

Multiple terms are used to describe physical design parameters used in Electrical Interface PHY IP links.

3.1. Process Technology

Parameter	Valid Values	Units
Process Technology	Node number	nm

Process technology refers to the manufacturing processes used to fabricate individual semiconductor devices. It includes the techniques and methodologies involved in fabricating the intricate structures and components of electronic circuits on a semiconductor substrate, typically made of silicon.

From heterogeneous chiplet integration perspective, advanced package technologies have made it possible to disaggregate SoC into various chip functions and each chip can be implemented in different technology nodes. That means we can take advantage of using advanced nodes for computing, dedicated processes for DRAM, flash, or analog, and older technology nodes for I/O interfaces to increase savings and decrease time to market. That's why there is a need to have Die-to-Die interfaces developed in multiple technology nodes from the most advanced to the mature ones. It will eventually enable the robust open chiplet ecosystem to fulfill the motive of interoperability, ease of integration and standard measurement methods of electrical interface PHY IP's.

3.2. Number of power supplies and values

Parameter	Valid Values	Units
Number of supplies & values	Number of supplies Supply voltage	n/a V

The number and type of power supplies in a PHY can vary based on the specific implementation, technology node, and the performance target requirements of the application. Vendors should follow the recommendations to minimize the number of power supplies, but they might have additional power supply requirements in their specific circuitry like PLL and clock generation in their PHY IPs.

Power supply operating conditions should be specified as nominal, min and max range, and AC and DC ripple noise tolerated by the power supply.



3.3. IO swing value V_{pp} , single-ended

Parameter	Valid Values	Units
IO swing value V_{pp} , single-ended	Voltage peak-to-peak, single-ended	V pkpk

Because IO swing is an important parameter for both TX and RX design in an electronic system, achieving the higher SNR or wide eye margin is the key. Every TX and RX design has minimum voltage swing requirements set to make sure the functionality and SNR are not compromised. For example, a D2D interface defines TX electrical parameters with a min voltage swing. Maximum channel reach for a specific data rate is also dependent on TX swing. For higher than minimum transmitter swing, an unterminated receiver can be extended to longer channels and higher data rate.

3.4. Interface Physical Dimensions

The die-to-die interface is defined by physical dimensions as shown in figure XX. Definition of these dimensions is included in the table below.




Figure 7: Illustration of die-to-die physical interface and dimensions

Dimension	Description
X_1	Spacing (center to center) of adjacent bump columns in 'X' direction
Y_1	Spacing (center to center) of adjacent bump rows in 'Y' direction
X_2	Spacing (center to center) of bumps in same row in 'X' direction
Y_2	Spacing (center to center) of bumps in same column in 'Y' direction
P	Minimum bump spacing (center to center)
Beachfront	Total width of PHY along chip edge
Depth	Distance from chip edge to furthest extent of PHY interior to chip

The physical dimensions of an interface refer to its layout footprint on the die or chiplet. This includes key linear measurements that determine the size and spacing requirements of the interface. Some important physical dimensions to specify are:

- Pitch - The center-to-center distance between adjacent signal pads or bumps. A lower pitch allows for higher pin density but raises package routing and manufacturing challenges.
- Pad size - The surface area occupied by each individual signal or I/O pad. Bigger pads simplify bonding or bumping but use up more space.
- Footprint size - The total linear extent of the interface occupied on the die/chiplet perimeter. Defined separately for the x and y dimensions.
- Alignment targets - Measurement points and markers provided for precision alignment between tiers when stacking chiplets. Their size and spacing impacts the ability to finely register chiplets.
- Exclusion zones - Regions around the interface footprint where no additional devices or routing is allowed. Ensure reliable connections by spacing the interface away from obstructions.

Orientation refers to the directional positioning of components or channels. Common orientation designations include:

N - North

S - South

E - East

W - West



These orientation values are used to define the positioning of PHY channels or other physical interconnection components when their geographical placement is important to characterize, such as in high-speed serializer/deserializer links or multi-wire bus architectures. Standardized orientation naming helps ensure consistent interpretation of channel configurations between design teams or documentation.

Specifying these key physical dimensions early helps validate manufacturability and guides tradeoffs in the interface design space. It also facilitates integration by communicating interface constraints to potential partners and foundries.

3.5. Min Data Pad spacing (X1, Y1, Pdiag, P)

Parameter	Valid Values	Units
Min Data Bump Spacing - X1	Spacing center-to-center of adjacent bump columns along Beachfront	um
Min Data Bump Spacing - Y1	Spacing center-to-center of adjacent bump rows into Depth	um
Min Diagonal pitch, Pdiag	If linear bump array - NA. If staggered bump array, minimum center-to-center between adjacent bumps on diagonal	um
Minimum pitch, P	Smallest repeat distance between pad centers.	um

The pad spacing in the physical interface will drive package routing conductor width and space as well as layer count. The min pad spacing constraints determine how closely signal pads can be packed within the interface footprint. Separate criteria often apply in the X and Y directions:

- X1, Y1 - Minimum center-to-center spacing between adjacent pads moving in the positive X or Y direction. Ensures no overlaps during bonding.
- X2, Y2 - Minimum center-to-center spacing moving in the linear X or Y direction. Prevents accidental bonds if pads are misaligned during assembly.
- Pdiag - If staggered bump array as shown in the figure, Pdiag is the minimum center-to-center spacing between bumps on the diagonal. If linear bump array, then N/A.
- P - Minimum pitch, the smallest repeat distance between pad centers. Defines the theoretical maximum achievable I/O density.

Closer spacing enables higher I/O counts but challenges assembly yields. Optimal values balance density with manufacturability. Design rules for pad spacing typically come from assembly house design kits based on proven yield models for a process. Communicating



these constraints early prevents inefficient interface layouts and yields robust, manufacturable chiplet designs.

3.6. Unit Beachfront per die

Parameter	Valid Values	Units
Unit Beachfront per die	Total length of the IO area across the beachfront	um

The unit beachfront refers to the total linear extents of the I/O areas occupied on the perimeter of a single die or chiplet. Referring to the diagram in Section 3.4, the measurement is determined as follows:

- Beachfront X - Total length in micrometers of the I/O areas stretching along the top and bottom edges of the die. Defines the minimum chiplet length along the chip edge.

Specifying the beachfront allows potential partners to assess the silicon footprint needed for a chiplet and its I/O requirements. It also ensures the physical dimensions of chiplets using the same interface can be stacked in standard tiles or units. The beachfront must account for not just I/O pads but also associated routing, ESD cells, alignment markers and other peripheral circuitry bound to the die edges. Communication upfront simplifies physical design efforts when using shared interfaces.

3.7. Unit Depth per die

Parameter	Valid Values	Units
Unit Depth per die	Total depth of the IO area away from chip edge	um

The unit depth refers to the total linear extents of the I/O areas occupied going into the single die or chiplet away from the chip edge or beachfront. It is to be measured as:

- Depth Y - Total length in micrometers of the I/O areas along the left and right sides. Determines minimum chiplet width (depth into ASIC away from chip edge or beachfront).

Specifying the depth allows potential partners to assess the silicon footprint needed for a chiplet and its I/O requirements. It also ensures the physical dimensions of chiplets using the same interface can be stacked in standard tiles or units. The depth must account for not just I/O pads but also associated routing, ESD cells, alignment markers and other peripheral circuitry bound to the die edges. Communication upfront simplifies physical design efforts when using shared interfaces.



3.8. NS/EW Orientation

Parameter	Valid Values	Units
NS/EW Orientation	Is the PHY poly orientation North-South or East-West	n/a

Is the PHY placed on the North-South (NS) or East-West (EW) edge of the die? Specifying the interface I/O region orientation relative to the die provides important context for multi-chiplet integration designs. The two common options are:

- NS orientation - I/O regions placed along the top and bottom edges of the die.
- EW orientation - I/O regions along the left and right side edges.

The orientation choice affects how chiplets can be packed layout and stacked physically. NS lends itself better to horizontally stacking columns of chiplets while EW favors vertical stacking in rows. This impacts how chiplets may be assembled onto interposers or substrates. Alignment and routing are also affected. Communicating orientation conventions early in the design process ensures smooth IP integration between providers and product designers.

3.9. Total Area (X & Y)

Parameter	Valid Values	Units
Total area dimension X & Y	X & Y	um

The total area for single channel transmitter and receiver including clocking reported separately for the two dimensions:

- X-dim indicates the cumulative length in μm along the top and bottom edges occupied by transmitter, receiver and clocking circuitry.
- Y-dim provides the equivalent length along the left and right sides.

Specifying the total area guides physical planning when such circuitry has to be shared.

3.10. Transmitter Receiver, and clocking Area

Parameter	Valid Values	Units
Total amortized physical area dimension for transmitter, receiver and clocking block separately	TX_X & TX_Y RX_X & RX_Y CK_X & CK_Y	um



The total area for single Transmitter / Receiver / Clocking reported separately for the two dimensions:

- X-dim indicates the cumulative length in μm along the top and bottom edges occupied by transmitter / receiver / clocking circuitry.
- Y-dim provides the equivalent length along the left and right sides.

Specifying the total area Transmitter, Receiver, and Clocking separately helps guide physical floor planning.

3.11. How many TX/RX macros can share one clocking block

Parameter	Valid Values	Units
Number of Tx/Rx macros that share one clocking block	Number of macros	n/a

The number of TX/RX macros that can share one clocking block parameter specifies the maximum number of Tx-Rx signal path macros that can share a single reference clock distribution block/network. This number depends on factors such as:

- Fanout capability of clock buffers/drivers - Larger fanout circuits can drive more loads, enabling greater sharing.
- Clock frequency/Edge rate - Higher speeds require tighter skew budgets, limiting sensible fanout.
- Jitter/timing closure - More shared loads increases non-determinism, challenging robust timing coordination between data and clock domains.
- Layout constraints - Space/routability limitations impact how many macros can physically cluster together.

Understanding the fanout capability of clocking circuits provides insight into physical planning - like how many I/O pairs can be placed together while reusing common synchronization circuitry to optimize area usage.

3.12. Max Stacked units

Parameter	Valid Values	Units
Max stacked units	Number of PHY units that can be stacked (away from chip edge, into ASIC)	n/a

The max stacked unit parameter refers to the maximum number of identical PHY blocks that can be stacked perpendicular to the die edge using the given interface specification.



It provides an important capacity metric for applications requiring 2D/2.5D multi-die integration. The limit depends on factors like:

- Underlying substrate design rules and routing capabilities
- Interconnect bandwidth/power delivery

It also aligns expectations for physical integration in multi-chiplet systems.

3.13. Number of substrate layers for reported edge density

Parameter	Valid Values	Units
Number of substrate layers Conductor width/spacing	Number of substrate layers, Line width/spacing	n/a, um/um

The number of substrate layers, conductor width and spacing parameter communicates key routing constraints for the substrate/interposer that provides physical support and interconnects between die/chiplets assembled using the standard interface. Specifically, it defines:

- Number of metal layers available for routing on the substrate
- Minimum line width of conductors on each layer
- Spacing between adjacent conductors to prevent shorts

Understanding these design rules let interface adopters assess the routing resources and density limitations for bridging I/O links between chiplets. This ensures reported interface edge densities along chiplet boundaries can actually be realized with a given interposer technology. It also eases planning 2D/2.5D assembly for a target capacity.

For standard packaging PHYs for organic/laminate substrates, the entry for Number of Substrate Layers should be for the package stackup (for example, 4-2-4, 8-2-8, etc.).

For advanced packaging PHYs using InFO or silicon interposers/bridges, the entry for Number of Substrate Layers should be the total number of layers used (for example, 4, 6, etc.)

3.14. Latency and BER

Parameter	Valid Values	Units
Latency for PHY-only without overhead for Packet/FEC/CRC	(number, ex: 2.0)	ns
BER for PHY-only without overhead for Packet/FEC/CRC	(number, ex: 1e-15)	errors



Parameter	Valid Values	Units
Latency for PHY together with Packet/FEC/CRC	(number, ex: 3.5)	ns
BER for PHY together with Packet/FEC/CRC	(number, ex: 1e-20)	errors
BER including CRC with Retry mechanisms	(number, ex: 1e-20)	errors

BER stands for Bit Error Rate in communication systems. It is a metric used to quantify the accuracy of data transmission by measuring the percentage of bits that are received incorrectly compared to the total number of bits transmitted. In other words, BER assesses the probability of errors occurring in the transmission of digital data.

The Bit Error Rate is typically expressed as a ratio or percentage, representing the number of erroneous bits divided by the total number of transmitted bits. A lower BER indicates a more reliable communication system, while a higher BER suggests a higher likelihood of errors.

CRC (Cyclic Redundancy Check) and FEC (Forward Error Correction) are two error detection and correction techniques used in data communication to ensure the accuracy of transmitted information.

All error handling schemes (CRC, FEC, Retry) will result in some degree of added latency and overhead (e.g. FEC or CRC Bytes).

A. CRC (Cyclic Redundancy Check):

CRC is an error-detection technique that involves appending a fixed-size check value to the data being transmitted. The check value, known as the CRC checksum, is generated by dividing the data and appended with a polynomial divisor. At the receiving end, the same polynomial divisor is used to divide the received data along with the CRC checksum. If the remainder is zero, the data is considered error-free. If not, errors are detected, and the data may need to be retransmitted. CRC is commonly used in network communication protocols, such as Ethernet, to detect errors in data frames.

B. FEC (Forward Error Correction):

FEC is an error-correction technique that involves adding redundant information (error-correcting codes) to the transmitted data. This redundant information allows the receiver to detect and correct errors without requiring retransmission of the data. FEC is particularly useful in scenarios where retransmission is impractical or introduces unacceptable delays, such as in real-time audio or video streaming. There are various FEC codes, such as Reed-Solomon codes and Turbo codes, each with its own properties and applications. While FEC introduces some overhead due to the redundant information, it can significantly improve the reliability of communication in the presence of errors. FEC often permits higher data rates at the same BER that more than compensate for the overhead in channel width.

C. "Retry" refers to the process of retransmitting data that was not successfully received or acknowledged by the intended recipient. When data is transmitted from a sender to a receiver, various factors such as noise, interference, or channel conditions may lead to errors or loss of data during transmission. To address this, communication protocols often incorporate mechanisms for error detection and recovery, and one common approach is through the use of retries.

Latency refers to the time delay between the initiation of a signal or data transmission and its reception. In the context of electrical PHY interfaces, latency can be influenced by various factors, including the transmission medium, signal processing, error correction and the complexity of the communication protocol.

While FEC enhances the reliability of communication by reducing the probability of data errors, it does introduce some level of latency both on the transmit side for the calculation of the redundant information and on the receive side for the detection and correction of any errors. CRC also enhances the reliability of communication, and it also introduces some latency on the transmit side for the checksum calculation but not less latency on the receive-side since it provides detection but not correction of errors. In addition to FEC and CRC, latency can also be introduced if the data is formatted into packets such as FLITS (used in PCIe and CXL). For example, if 256B packets are used, then all 256B must be accumulated before error detection and correction can be performed.

A Clock Domain Crossing (CDC) buffer, also known as a synchronizer or FIFO (First-In-First-Out) buffer, is a digital circuit component used to safely transfer data between different clock domains within a system. Clock domains have independent clock signals, and crossing between them requires careful handling to avoid data integrity issues. The elastic buffer absorbs variations in data arrival rates between the two clock domains, allowing for a more flexible and adaptable data transfer. The term "elastic" refers to its ability to stretch or compress data as needed to align with the different clock frequencies. The effect of Clock Domain Crossing buffers on system latency is the result of synchronization processes and pipeline stages introduced to ensure reliable data transfer between different clock domains.

For accurate performance evaluation of both Bit Error Rate (BER) and Latency, measurements should be conducted in specific conditions. The paper recommends a three-tiered approach to BER and Latency measurement: with Packets/FEC/CRC, without Packets/FEC/CRC, and targeting a specific BER with retry mechanisms. This strategy provides a comprehensive understanding of system performance across various scenarios.

This paper recommends a three-tiered approach to describing BER and Latency performance for a given PHY:

- BL-0 BER & Latency for PHY-only without overhead for Packet/FEC/CRC,
- BL-1 BER & Latency for PHY together with Packet/FEC/CRC, and
- BL-2 BER including CRC with Retry mechanisms.

This strategy provides a comprehensive understanding of system performance across various scenarios.

For BL-0 BER & Latency for the PHY without Packets/FEC/CRC, this paper recommends capturing the latency and BER results when the input data rate to the serializer is 1GT/s and



again at the maximum speed per data lane, in order to capture both analog and UI-based delays.

BL-0A: Measure latency from the input of the serializer to the output of the deserializer at a data rate of 1 GT/s with a recorded bit error rate (BER-0A) and excluding latency of clock domain crossing elastic buffer.

BL-0B: Measuring latency from the input of the serializer to the output of the deserializer at the maximum supported speed per data lane (SP-0B) by the interface PHY with a recorded bit error rate (BER-0B) and recorded input data rate to the serializer (DR-0B) and excluding latency of clock domain crossing elastic buffer.

For BL-1 BER & Latency for the PHY including Packets/FEC/CRC overhead, this paper recommends capturing the latency and BER results when the input data rate to the serializer is 1GT/s and again at the maximum speed per data lane, in order to capture both analog and UI-based delays. For cases with FEC, the BER should reflect the improvement from error correction, while for cases with CRC, the BER reported should be the uncorrected bit error rate.

BL-1A: Measuring latency from the input of Packets/FEC/CRC block to the output of Packets/FEC/CRC at the receiver at a data rate of 1 GT/s with a recorded bit error rate (BER) while excluding latency of the clock domain crossing elastic buffer.

BL-1B: Measuring latency from the input of Packets/FEC/CRC block to the output of Packets/FEC/CRC at the receiver at the maximum supported speed per data lane (SP-1B) by the interface PHY with a recorded bit error rate (BER) and recorded input data rate to the serializer (DR-1B) while excluding latency of the clock domain crossing elastic buffer.

For B-2 BER for the PHY including CRC and Retry, this paper recommends capturing the bit error rate results (BL2) when CRC calculation and error detection are combined with Retry mechanisms to improve BER.




Figure 8: Latency and BER Measurement Points

	Data Rate	Latency	BER
Ser to Des	1GTps	T0-A	BL0-A
Ser to Des	Max Supported Data Rate (DR0)	T0-B	BL0-B
Packet to Packet	1GTps	T1-A	BL1-A
Packet to Packet	Max Supported Data Rate (DR1)	T1-B	BL1-B
With Retry	Max Supported Data Rate (DR2)	N/A	BL2

3.15. Reach

Parameter	Valid Values	Units
Reach	Signaling distance	mm

In the context of electrical PHY (Physical Layer) interfaces, "reach" refers to the maximum distance in millimeter [mm] over which a signal can be transmitted while still maintaining target BER.

In certain implementations, the PHY interface is arranged in multiple rows, necessitating the definition of the MAX reach based on the maximum number of stackable units illustrated in the diagram below.

4. Calculations and Figures of Merit (FOMs)

Multiple figures of merit (FOMs) are used to describe the performance of Electrical Interface PHY IP links. The figures of merit are calculated for bidirectional links, including both Tx and Rx lanes.

4.1. Unit Bandwidth

Parameter	Valid Values	Units
Unit Bandwidth	Speed/lane * num_lanes	Gbps

Unit Bandwidth provides the effective bandwidth contribution in a single direction (chip A to chip B) for the given number of lanes, wires and targeted line rate per lane. When combined with the other direction (chip B to chip A), it results in the total full-duplex unit bandwidth.

Unit bandwidth in each direction (unit_bw) = speed / lane x number of data lanes

The unit bandwidth is calculated based on bidirectional links. For example, a PHY that has 16 Tx + 16 Rx lanes would have bandwidth calculated based on 32 lanes. A PHY that has 16 lanes that are selectable to be Tx OR Rx should report bandwidth based on two PHYs to give bidirectional data of 16 Tx + 16 Rx lanes.

Where:

speed/lane is the targeted speed per data lane as described in Section 2.6

num_lanes is the number of data lanes as described in Section 2.1

Unit is Gbps to align with the typical bandwidth specifications

4.2. Unit Power

Parameter	Valid Values	Units
Unit Power	Total power for unit including Tx + Rx + clocks	mW

Unit Power refers to the total power consumption of the transmitter (Tx), receiver (Rx) and clocking circuitry within a single computational/memory unit. Unit power is reported as the actual measured or calculated total power consumption of the Tx, Rx and clocking blocks/circuitry in milliwatts (mW).

Unit power (unit_power) = Tx + Rx + Clock power

The unit power is calculated based on bidirectional links. For example, a PHY that has 16 Tx + 16 Rx wires would have unit power calculated based on all Tx + Rx + Clock power supporting all 32 lanes of bandwidth, while a PHY that has 16 lanes that are selectable to be



This work is licensed under a [Creative Commons Attribution-ShareAlike 4.0 International License](https://creativecommons.org/licenses/by-sa/4.0/).

Tx OR Rx should report bandwidth based on two PHYs to give bidirectional data of 16 Tx + 16 Rx lanes + Clock power supporting 32 lanes of bandwidth.

Where:

Tx is transmitter power (mW)

Rx is receiver power (mW)

Clock is clock power (mW)

Using a pseudorandom pattern such as Pseudo-Random Bit Sequence (PRBS) is generally recommended for measuring power for the following reasons:

- Randomness - A truly random bit pattern provides the most realistic approximation of real data traffic.
- Repeating nature - Pseudorandom patterns repeat deterministically to enable multiple measurements and average/compare them easily.
- Error detection - Patterns like PRBS allow use of BER tests to ensure integrity of power measurements.
- Standardized - Patterns like PRBS are widely used so results are comparable to specs and other equipment/testing.
- Full bandwidth - The rapid transitions of pseudorandom patterns fully stresses the interface and utilizes its bandwidth similar to real data.
- Worst-case power - Random patterns can generate worst-case switching scenarios that produce maximum highs and lows in power consumption/dissipation.
- Consistency - Pseudorandom patterns repeat identically each time so you get consistent readings without variation from real data traffic patterns.

Typical steps to force a PRBS pattern and measure power on a physical interface:

- Configure the interface for PRBS mode via setting various registers to enable PRBS generation/checking on the transmitter and receiver.
- Select the PRBS pattern type such as PRBS13, etc.
 - For Multilevel Signaling (PAM-X), patterns such as PRBS13 should be used for each level with different seeds
- Start pattern generation on the transmitter side.
- Configure power measurement hardware/software such as connecting an electrical power meter to the interface and setting it up to take measurements.
- Optionally loopback or connect the interface to itself so the transmitted PRBS is received back.
- Start pattern transmission and take unit power measurements.
- Compare measurements to specifications/requirements to check for compliance.
- Analyze and record test results. Adjusts output drive levels, equalization, or other tuning as necessary and repeat the steps above.

The unit is in milliwatts (mW)



4.3. Technology Power Efficiency

Parameter	Valid Values	Units
Power Efficiency	Unit power / BW	pJ/bit

This provides a measure of the energy efficiency of the technology, in terms of the power required to transmit a single bit. Lower values indicate better power efficiency. This is a key figure of merit for evaluating and comparing the energy efficiency of various interconnect technologies.

$$\text{Technology power efficiency} = \text{unit\_power} / \text{unit\_bw}$$

The power efficiency is calculated based on bidirectional links. For example, a PHY that has 16 Tx + 16 Rx wires would have power efficiency calculated based on all Tx + Rx + Clock power divided by all 32 lanes of bandwidth, while a PHY that has 16 lanes that are selectable to be Tx OR Rx should report bandwidth based on two PHYs to give bidirectional data of 16 Tx + 16 Rx lanes + Clock power supporting 32 lanes of bandwidth.

Where:

unit_power = Total Tx + Rx + Clock power consumption from Section 4.2 above

unit_bw = Unit bandwidth from Section 4.1 above

The unit of the Technology power efficiency parameter is pJ/bit.

Note that the units pJ/bit = mW/Gbps, so the two units are interchangeable.

4.4. Edge Bandwidth Density

Parameter	Valid Values	Units
Edge BW density	Unit BW / Beachfront	Tbps/mm

This parameter provides a measure of the edge density or bandwidth that can be achieved per unit of beachfront/interconnect distance. It quantifies how much aggregate bandwidth can be delivered across a given interconnect beachfront resource. This helps evaluate the effective density achievable with different technology solutions.

For PHY IP cores designed to implement standard interfaces such as UCIe or BoW, the technology edge density reported here should match the specification requirements defined by those standards. Adhering to the edge density targets ensures interoperability when integrating the PHY IP into designs using those interface standards.

$$\text{Technology edge density for technology} = \text{unit\_bw} * \text{stack\_units} / \text{beachfront}$$

Where:

unit_bw is the total unit bandwidth from Section 4.1



This work is licensed under a [Creative Commons Attribution-ShareAlike 4.0 International License](https://creativecommons.org/licenses/by-sa/4.0/).

stack_units is the max number of units that can be stacked from Section 3.12
beachfront is the unit beachfront per die in mm from Section 3.6

The unit for this metric is Tbps/mm.

4.5. Area Bandwidth Density

Parameter	Valid Values	Units
Areal BW Efficiency	Unit BW / Area	Tbps/mm ²

This parameter provides a measure of the area density or bandwidth that can be achieved within the footprint area required by the interconnect unit. It quantifies how much bandwidth can be delivered per unit area by dividing the total unit bandwidth by the physical die/chip area consumed by the Tx, Rx and clock circuitry.

Similarly, the area density reported for PHY IP cores supporting standard interfaces like UCIe or BoW should meet the area density specifications defined by those standards. Matching the standard area density targets allows such PHY IP cores to be drop-in compatible with designs using the standardized interface.

$$\text{Technology area density for technology} = \text{unit\_bw} / \text{unit\_area}$$

Where:

unit_bw is the total unit bandwidth from in Section 4.1

unit_area is the total Tx + Rx + clocking physical area (mm²) in Section 3.9

The unit for this metric is Tbps/mm².

4.6. Technology FOM1

Parameter	Valid Values	Units
Technology FOM1 Figure of Merit #1	Edge density / power efficiency	(Tbps/mm) / (pJ/bit)

This Figure of Merit allows a combined assessment of bandwidth density and energy efficiency to compare technologies in a single metric. A technology with a higher FOM1 provides better performance within the same power budget. FOM1 balances edge density versus power consumption, giving an indication of overall technological efficiency.

A higher FOM1 value would imply either:

- Higher edge density for the same power efficiency
- Same edge density but better power efficiency



This work is licensed under a [Creative Commons Attribution-ShareAlike 4.0 International License](https://creativecommons.org/licenses/by-sa/4.0/).

$$\text{Technology FOM1} = \text{max\_edge\_density} / \text{power\_eff}$$

Where:

max_edge_density is the Technology edge density (Tbps/mm) in Section 4.4

power_eff is the Technology power efficiency (pJ/bit) in Section 4.3

The unit for this metric is (Tbps/mm)/(pJ/bit).

4.7. Technology FOM2

Parameter	Valid Values	Units
Technology FOM2 Figure of Merit #2	Power efficiency / reach	(pJ/bit) / (mm)

This Figure of Merit allows an assessment of the tradeoff between energy consumption and physical transmission span. A technology with a lower FOM2 indicates it can transmit data either more efficiently or over a longer distance.

This evaluates:

- How efficient the technology is in terms of pJ/bit of power
- In relation to how far it can transmit data over a given interconnect distance

$$\text{Technology FOM2} = \text{power\_efficiency} / \text{reach}$$

Where:

power_efficiency is the Technology power efficiency (pJ/bit) in Section 4.3

reach is the unit reach specified (mm) in Section 3.15

The unit for this metric is (pJ/bit)/mm.



5. Recommended Minimum List of Parameters

Of the metrics listed in this whitepaper, a good subset list that provides a quick overview of any D2D interface should include the following:

Metric	Units
D2D Spec/Version	(N/A)
Data lane speed	Gbps
PHY Bandwidth (Rx + Tx)	Gbps
PHY Bandwidth (Rx + Tx) per shoreline	Tbps/mm
Reach	mm
Power efficiency	pJ/bit
Packaging type	Std Pkg or Adv Pkg
Min bump pitch	um
BER	Ex: 1E-15
Latency T0	ns

6. Conclusion and Next Steps

This document provides a comprehensive framework for understanding and standardizing the performance metrics of Electrical Interface Physical Layer Intellectual Property (PHY IP), a critical component in modern electronic systems. By focusing on the Die-to-Die interface, which integrates the Link Layer and Physical Layer, the paper emphasizes the importance of seamless data transmission and the need for reliable, high-speed communication across various network interfaces.

The standardization of performance metrics is crucial for ensuring interoperability, ease of integration, comparability, quality assurance, and reduced development time within the electronics industry. These standards foster industry collaboration, meet regulatory requirements, and advance the development of sophisticated electronic systems.

The document details specific technical aspects such as serial and parallel interfaces, clocking mechanisms, signaling types, and termination methods, which influence the design and implementation of PHY IPs. The recommended figures of merit, including bandwidth efficiency, power efficiency, edge density, and area density, provide a quantitative basis for evaluating and comparing different PHY technologies. Adopting these standardized metrics ensures that PHY IPs meet necessary performance benchmarks, leading to more reliable and efficient electronic systems.



Next Steps:

Data Collection on Standards (BoW and UCle): Begin gathering data on existing standards such as Bunch of Wires (BoW) and Universal Chiplet Interface Express (UCle) based on the guidelines established in this white paper. This data should be compiled and reported as a practical example of how the proposed metrics can be applied in real-world scenarios, helping to validate the framework and identify areas for improvement.

Consideration of Advanced Die-to-Die Interfaces: Explore more advanced die-to-die interface configurations, such as Leader and Follower links, to expand the applicability of the framework. This involves studying how these advanced interfaces operate and how the proposed metrics can be adapted or extended to capture their unique characteristics and performance requirements.

Incorporation of Additional Parameters for 3D-ICs: Extend the scope of the white paper to include additional parameters specific to 3D-IC (Three-Dimensional Integrated Circuit) interfaces. As 3D-IC technology becomes more prevalent, it is essential to consider how these parameters interact with the existing framework and what new metrics might be required to accurately assess performance in a 3D-IC environment.

Updating Performance Metrics Based on Feedback: Establish a process for regularly updating the performance metrics and framework based on customer feedback and technological advancements. This iterative approach will ensure that the framework remains relevant and aligned with industry needs, allowing it to evolve alongside emerging technologies and new use cases.

By addressing these steps, the paper supports the ongoing development and standardization of Electrical Interface PHY IP practices, contributing to the creation of more sophisticated, efficient, and reliable electronic systems.

7. License

OCP encourages participants to share their proposals, specifications and designs with the community. This is to promote openness and encourage continuous and open feedback. It is important to remember that by providing feedback for any such documents, whether in written or verbal form, that the contributor or the contributor's organization grants OCP and its members irrevocable right to use this feedback for any purpose without any further obligation.

It is acknowledged that any such documentation and any ancillary materials that are provided to OCP in connection with this document, including without limitation any white papers, articles, photographs, studies, diagrams, contact information (together, "Materials") are made available under the Creative Commons Attribution-ShareAlike 4.0 International License found here: <https://creativecommons.org/licenses/by-sa/4.0/>, or any later version, and without limiting the foregoing, OCP may make the Materials available under such terms.



This work is licensed under a [Creative Commons Attribution-ShareAlike 4.0 International License](https://creativecommons.org/licenses/by-sa/4.0/).

As a contributor to this document, all members represent that they have the authority to grant the rights and licenses herein. They further represent and warrant that the Materials do not and will not violate the copyrights or misappropriate the trade secret rights of any third party, including without limitation rights in intellectual property. The contributor(s) also represent that, to the extent the Materials include materials protected by copyright or trade secret rights that are owned or created by any third-party, they have obtained permission for its use consistent with the foregoing. They will provide OCP evidence of such permission upon OCP's request. This document and any "Materials" are published on the respective project's wiki page and are open to the public in accordance with OCP's Bylaws and IP Policy.

This can be found at <http://www.opencompute.org/participate/legal-documents/>. If you have any questions please contact OCP.

8. About Open Compute Foundation

At the core of the Open Compute Project (OCP) is its Community of hyperscale data center operators, joined by telecom and colocation providers and enterprise IT users, working with vendors to develop open innovations that, when embedded in products, are deployed from the cloud to the edge. The OCP Foundation is responsible for fostering and serving the OCP Community to meet the market and shape the future, taking hyperscale led innovations to everyone. Meeting the market is accomplished through open designs and best practices, and with data center facility and IT equipment embedding OCP Community-developed innovations for efficiency, at-scale operations and sustainability. Shaping the future includes investing in strategic initiatives that prepare the IT ecosystem for major changes, such as AI & ML, optics, advanced cooling techniques, and composable silicon. Learn more at www.opencompute.org

9. Compliance with Open Compute Project Tenets

This whitepaper, authored by the PHY metrics workstream of the Open Domains Specific Architecture Group, complies with the five tenets of the Open Compute Project as detailed in the paragraphs below.

Openness

Participation in the ODSA and the PHY metrics workstream is open to everyone. The meeting times are posted online and anyone can join, participate and contribute. The drafts for this whitepaper were hosted online and sent out regularly to anyone signed up for the ODSA mailing list to seek feedback.

Efficiency

The goal of this whitepaper is to address the challenges of documenting the performance of links in chiplet-based systems. Chiplets are a design methodology to build systems out of modular die (chiplets) with standardized interfaces. This allows for chiplets to be reused



This work is licensed under a [Creative Commons Attribution-ShareAlike 4.0 International License](https://creativecommons.org/licenses/by-sa/4.0/).

across different product designs which allows complex systems to be created faster, with less development effort and often with lower unit cost clearly embodying the tenet of efficiency.

Impact

Chiplet technology has a transformative impact on the semiconductor industry. Today, many big companies leverage chiplet technology to achieve the highest levels of performance at manageable design time and cost. We believe this whitepaper is impactful because it will help smaller companies understand the challenges of describing chiplet link metrics better and reap the benefits of chiplet-based designs.

Scale

Coming up with chiplet-based products is a design philosophy rather than a rule-based classification. As such, the insights described in this whitepaper can be easily scaled to many application areas within semiconductor product design. Additionally, chiplets enable building scalable product portfolios through defining base chiplets that can be defined in various ways to meet different performance, feature and cost targets.

Sustainability

The modular and reusable nature of chiplets makes this design approach more sustainable than repeated custom monolithic designs. As a result, we believe this paper is in-line with building a more sustainable semiconductor industry.

