



OPEN
Compute Project

Bunch of Wires PHY Specification

The Open Domain-Specific Architecture BoW
Workstream

Release Candidate
WG Approved - 20 May 2022

Authors:

Elad Alon: Blue Cheetah Analog Design

Ken Poulton: Keysight

Shahab Ardalan: Luminous

Boris Murmann: d-Matrix

Namhoon Kim: Google

Greg Favor: Ventana Micro

Carrie Cox: IBM

Mike Bartley: Tessolve

Contents

Glossary of Terms

Language

1. License Agreement

- 1.1. Open Web Foundation (OWF) CLA

2. OCP Tenets Compliance

- 2.1. Openness
- 2.2. Efficiency
- 2.3. Scale
- 2.4. Impact

3. Revision Table

4. Scope

5. BoW Overview

- 5.1. Key Features and Conformance
- 5.2. BoW Slice
- 5.3. BoW Wires
- 5.4. BoW Modes
- 5.5. Logic Interface

6. Signal Definitions

- 6.1. Die-to-die Signals (Wires)
 - 6.1.1. DBI on the AUX wire
- 6.2. Slice Logic Interface
 - 6.2.1. Slice Logic Interface: Data Signals
 - 6.2.2. Slice Logic Interface: Control Signals
 - 6.2.3. Programming
 - 6.2.4. Link Controller
- 7. BoW Modes and Reach**
- 8. BoW Physical Configuration**
 - 8.1. Dead-Bug Views
 - 8.2. BoW Components
 - 8.3. Example Link
 - 8.4. Die-to-Die Signals
 - 8.5. Signal Ordering
 - 8.6. Bump Arrangements
 - 8.6.1. Alternate Bump Arrangements
 - 8.7. Cross Section
 - 8.8. Staggered Slices
 - 8.9. Slice Numbering
 - 8.10. Slice Stacking Pattern for Symmetric Links
- 9. BoW PHY Electrical Specifications**
 - 9.1. Voltages and Termination Resistance
 - 9.2. PHY protection
 - 9.3. ESD
 - 9.4. Return Loss and Parasitic Capacitance
 - 9.5. Receiver Bandwidth
- 10. BoW PHY Timing Specifications**
 - 10.1. Bit Ordering
 - 10.2. Clocking
 - 10.3. Clock and Data Specifications
 - 10.3.1. Transmitter Maximum Rise-Time
 - 10.3.2. Transmitter Correlated Jitter Filtering
 - 10.3.3. Transmitter Deterministic Timing Error
 - 10.3.4. Transmitter Total Timing Error
 - 10.3.5. CLK Receiver Sensitivity to Common-Mode Variations
 - 10.3.6. Receiver Sensitivity and Timing Margin
 - 10.3.7. Voltage Overshoot
 - 10.3.8. Slice-to-Slice CLK Skew
- 11. Chip-to-Chip Channel Specifications**
 - 11.1. Channel-Induced CLK Edge to D Transition Skew at RX
 - 11.2. Channel Impedance
 - 11.3. Example BoW Channel for Doubly-Terminated Links
 - 11.3.1. Channel Loss
 - 11.3.2. Crosstalk
- 12. Reset and Initialization**
 - 12.1. External Facilities
 - 12.2. Initialization Sequence
 - 12.2.1. A TX-RX Link should be brought up as follows:
 - 12.3. Unspecified Items
- 13. Configuration**
 - 13.1. Link Training
- 14. Control Register Mapping**
- 15. Testability**
 - 15.1. Test Patterns
 - 15.2. Loopback Test
- 16. BoW in an ODSA Design**
 - 16.1. BoW for Common Transaction Protocols

Appendix A - Requirements for IC Approval (to be completed Contributor(s) of this Spec)
Appendix B

Glossary of Terms

This section provides glossary used in this specification.

Term	Abbreviation	Definition
Bunch of Wires	BoW	The name for the PHY specification defined in this document
Die-to-die	D2D	Generic term used to refer to on-package interconnect
BoW Mode	N/A	A specific defined mode of operation for a BoW interface
picojoules per bit	pJ/bit	Energy required to transport a bit of data over a D2D interface
PHY		The set of circuitry physically communicating bits from one die to another
Core Logic		Digital logic transmitting data to/from the PHY
Control Logic		Logic used to manage the operation of the PHY
Tera/Giga bits per second	Tbps/Gbps	Measures of the speed of data transmission on the PHY
Beachfront		The length of die edge required by a PHY implementation
Clock		A signal that regulates the speed of data transmission
Bump		Solder balls grown on a die to allow connection to off-die wires
Channel		A term for the physical connection between a transmitter and a receiver
Test		The process of verifying functional correctness of a circuit
Initialization		The process of preparing an interface for data transmission

Table 1. Glossary of Terms

Language

This document uses the following terms as defined below.

- “Shall” or “must” indicates a requirement. Failure to meet the requirement results in non-conformance
- “Should” indicates a recommendation, but not a requirement. Failure to implement the recommendation does not result in non-conformance.
- “May” indicates an implementation option.
- The lack of one of the above verbs indicates the material is informative.

- “Reference” indicates a reference design that is provided as example for explanation, but is not a requirement.

1. License Agreement

1.1. Open Web Foundation (OWF) CLA

Contributions to this Specification are made under the terms and conditions set forth in the modified Open Web Foundation Contributor License Agreement (“OWF CLA 1.0”) (“Contribution License”) by:

BLUE CHEETAH ANALOG DESIGN, D-MATRIX, IBM, KEYSIGHT, TESSOLVE, VENTANA MICRO

You can review the signed copies of the applicable Contributor License(s) for this Specification on the OCP website at <http://www.opencompute.org/products/specsanddesign>

Usage of this Specification is governed by the terms and conditions set forth in the modified Open Web Foundation Final Specification Agreement (“OWFa 1.0”) (“Specification License”).

Notes:

1. The above license does not apply to the Appendix or Appendices. The information in the Appendix or Appendices is for reference only and non-normative in nature.

NOTWITHSTANDING THE FOREGOING LICENSES, THIS SPECIFICATION IS PROVIDED BY OCP “AS IS” AND OCP EXPRESSLY DISCLAIMS ANY WARRANTIES (EXPRESS, IMPLIED, OR OTHERWISE), INCLUDING IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, FITNESS FOR A PARTICULAR PURPOSE, OR TITLE, RELATED TO THE SPECIFICATION. NOTICE IS HEREBY GIVEN, THAT OTHER RIGHTS NOT GRANTED AS SET FORTH ABOVE, INCLUDING WITHOUT LIMITATION, RIGHTS OF THIRD PARTIES WHO DID NOT EXECUTE THE ABOVE LICENSES, MAY BE IMPLICATED BY THE IMPLEMENTATION OF OR COMPLIANCE WITH THIS SPECIFICATION. OCP IS NOT RESPONSIBLE FOR IDENTIFYING RIGHTS FOR WHICH A LICENSE MAY BE REQUIRED IN ORDER TO IMPLEMENT THIS SPECIFICATION. THE ENTIRE RISK AS TO IMPLEMENTING OR OTHERWISE USING THE SPECIFICATION IS ASSUMED BY YOU. IN NO EVENT WILL OCP BE LIABLE TO YOU FOR ANY MONETARY DAMAGES WITH RESPECT TO ANY CLAIMS RELATED TO, OR ARISING OUT OF YOUR USE OF THIS SPECIFICATION, INCLUDING BUT NOT LIMITED TO ANY LIABILITY FOR LOST PROFITS OR ANY CONSEQUENTIAL, INCIDENTAL, INDIRECT, SPECIAL OR PUNITIVE DAMAGES OF ANY CHARACTER FROM ANY CAUSES OF ACTION OF ANY KIND WITH RESPECT TO THIS SPECIFICATION, WHETHER BASED ON BREACH OF CONTRACT, TORT (INCLUDING NEGLIGENCE), OR OTHERWISE, AND EVEN IF OCP HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

2. OCP Tenets Compliance

The Bunch of Wires (BoW) is a simple, open and interoperable physical interface between two chiplets or chip-scale-packages (CSP) in a common package. This document specifies the BoW interface PHY layer. The BoW interface is a set of die-to-die parallel interfaces that provides the flexibility to trade off throughput/chipedge for design complexity, cost, and packaging technology. The use of BoW is expected to be confined to connect die placed close to one another within the same package. In this environment, signal attenuation is small and the interface can be simple. The definition of the BoW interface aims to meet the following OCP tenets as follows:

2.1. Openness

- Unencumbered by technology license costs

2.2. Efficiency

- Inexpensive to implement
- Very low power (< 0.5 - 1 pJ/bit) as defined by TX IO Pad, wire and RX IO Pad.
- Very low latency (<5 ns without FEC, <15 ns with FEC from logic interface to logic interface)
- High throughput density (100-1000+ Gbps/mm-chip-edge)
- Backwards compatible (across at least two major specification versions)

2.3. Scale

- Flexible to support both laminate and advancing packaging technologies
- Portable across multiple bump pitches
- Portable across IC process nodes ranging from 65 nm to 5 nm and beyond

2.4. Impact

The Bunch of Wires interface provides several key advantages for chiplet-based systems:

- Can operate at higher data rates per pin than existing parallel standards
 - or at lower data rates for compatibility with existing parallel standards
- Can be implemented in legacy technologies (process nodes) with generally available IP
- Can be implemented in low-cost laminates or higher-density silicon-based interconnect
- Can be implemented with much less design effort than a traditional SerDes
- Is not constrained to a specific bump pitch
 - interfaces with somewhat different bump pitches can be connected

Compared to SerDes, BoW uses a lower data rate/wire so it requires more wires. However the lower data rates allow use of single-ended signaling and denser wire packing. In addition, in laminates, BoW can take advantage of multiple wiring layers and in advanced packaging it can take advantage of the much-increased wire density.

3. Revision Table

Date	Revision	Author	Description
05/24/22	1.0	Elad Alon	Feedback from community
03/15/22	0.9	Elad Alon	Channel simulations and models
12/01/21	0.8	Ken Poulton	Figures following convention
07/13/21	0.7	Ken Poulton	Control model
03/30/21	0.5	Ken Poulton	Final (current) definition of the basic BoW interface

4. Scope

The scope of this document has several levels.

1. The specification of the BoW interface includes these requirements:
 - a. Operating modes
 - b. Chip-to-chip wire signals
 - c. Wire ordering

- d. Timing and electrical specifications on the chip-to-chip interface
 - e. Signals at the logic (Link Layer) interface
 - f. Configuration, initialization, calibration
 - g. Functions that must be supported at the Link Layer or above
2. The specification includes recommendations for these elements:
- a. Bump patterns
 - b. Arrangement of multiple slices in a link
 - c. Arrangement of wires in laminate and advanced packaging
 - d. signal integrity of the wire channel
 - e. Configuration and management programming
 - f. Design for test and test methods
 - g. Performance estimates
 - h. Conformance verification
3. The following activities are outside the scope of this document:
- a. Specific implementations of the interface
 - b. Integration of the interface with system-level data flow e.g. interface to a PHY-layer abstraction such as PIPE/PCle interface to the BoW
 - c. The use of this interface outside of a package or entirely inside a chip
 - d. Definition of protocols for logical data transfer
4. The following aspects may be addressed in subsequent versions of this specification:
- a. Simultaneous bidirectional data (full duplex on each wire)
 - b. Security

5. BoW Overview

This section provides an overview of the BoW physical interface (PHY) and its use in a multi-chiplet design.

5.1. Key Features and Conformance

The specifications must be met over process variation, supply voltage range and temperature range (PVT). Each implementation must document its supported I/O voltage range, supply voltage range and temperature range.

Table 2 summarize the conformance points that shall be met in order to comply with the BoW specification. Each of the conformance points is discussed in the specification.

Description	Section	Detail
BoW Modes	5.4	
Die-to-die Signals (Wires)	6.1	
Slice Logic Interface	6.2	
BoW Modes and Reach	7	Table 7
Wire and Slice Ordering	8	
Voltages and Termination Resistance	9.1	
PHY Protection	9.2	
ESD	9.3	
Return Loss and Parasitic Capacitance	9.4	
Clocking	10.2	
Clock and Data Specs	7.2	
Channel Skew	8.1	

Description	Section	Detail
External Facilities	12.1	
Initialization	12.2	
Control Register Mapping	14	

Table 2. BoW Conformance Summary

5.2. BoW Slice

BoW is an energy-efficient, easy-to-use PHY interface between a pair of die inside a single package as shown in Figure 1. The BoW PHY is defined as a single unidirectional slice. Multiple slices are combined to create links of the desired throughput. A link may be symmetric, asymmetric or unidirectional. The BoW PHYs between two die are physically connected through wires on a substrate or interposer. A BoW PHY does not have enough drive strength for off-package interfaces, nor is it designed for buses that are entirely on die.

This document specifies the protocol for a BoW PHY slice. The aggregation of multiple PHYs into a link is beyond the scope of this document.

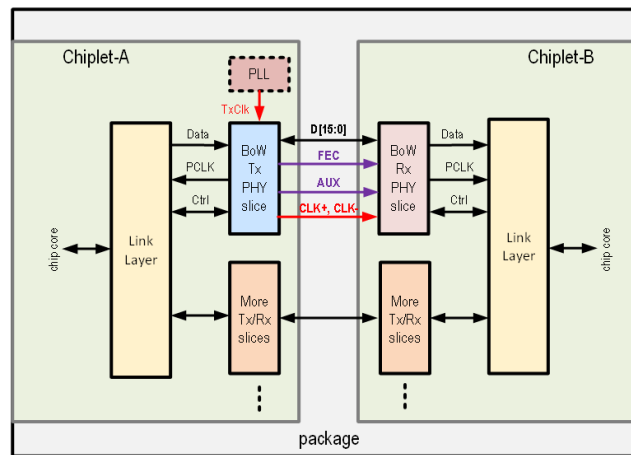


Figure 1. BoW Overview

A BoW PHY slice either transmits or receives 16 bits of data between die. The BoW is a source-synchronous PHY and each transmitting PHY slice transmits a complementary clock signal CLK+ and CLK- with the data. A BoW PHY optionally has two additional wires designated FEC (for Forward Error Correction) and AUX, for other optional functions such as Data Bus Inversion (DBI).

5.3. BoW Wires

Within the package, the BoW datapath is transported on physical passive wires between the pair of connected die. The specifics of the wires, such as their density, maximum length, impedance characteristics and how they are realized vary with the packaging technology. In order to minimize power, unterminated and source-terminated links will have short reaches requiring chips to be adjacent.

5.4. BoW Modes

A BoW PHY must be operable in one of the BoW Modes listed in ascending order in Table 3. A BoW Mode defines the speed of clock and data of the PHY on the die-to-die wires. In all

modes, the data must be clocked DDR: the chip-to-chip data wire bit rate is double the clock wire frequency. All BoW interfaces faster than BoW-64 should also be able to support BoW-64. Supporting rates other than the defined four modes is an implementation choice. There is more detail on BoW Modes in section 7.

BoW Mode	Slice Data Rate Gbps	Wire Bit Rate Gbps/wire	TxClock GHz
BoW-32	32	2	1
BoW-64	64	4	2
BoW-128	128	8	4
BoW-256	256	16	8

Table 3. BoW Modes




Figure 2. BoW Data Rate vs. Reach tradeoff

Figure 2 shows the tradeoff between package, data rate, termination, and reach. Source-terminated BoW on laminate allows a longer reach than advanced packaging, but the wider design rules in laminate means that both of these cases are barely able to reach 8 Gbps/wire. A doubly-terminated link offers longer distances and higher rates, but requires a more complicated receiver design.

5.5. Logic Interface

Figure 3 shows the logic interface between a BoW slice and the digital Link Layer logic in a chip. The speed at the logic interface (Figure 1) is implementation-dependent. Typically, PCLK will be the TxClock frequency divided by a power of 2, so 250, 500 and 1000 MHz are common rates. The data at the logic interface is SDR (bit rate equal to PCLK frequency).

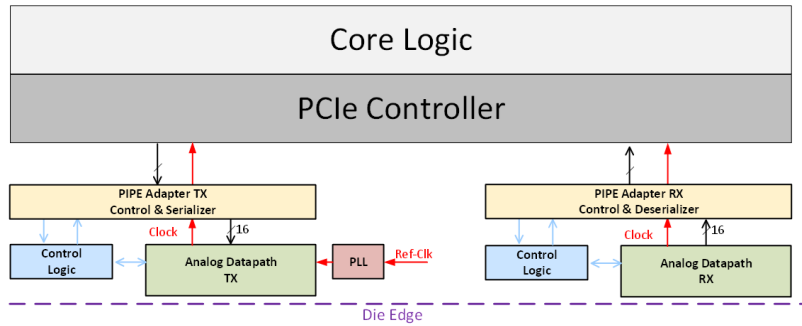


Figure 3. BoW slice logic interface

6. Signal Definitions

This section specifies the control data signals into and out of device logic and package for BoW RX and TX slices.

6.1. Die-to-die Signals (Wires)

As shown in Figure 1, each BoW slice consists of a differential clock pair, 16 single-ended data wires, and an optional pair of wires FEC and AUX.

Each BoW slice is unidirectional when in operation. A chiplet may be designed with RX-only and TX-only slices, or each slice may have both TX and RX capability which is configured at runtime. A bidirectional link is composed of some number of slices configured for RX and some for TX.

FEC (Forward Error Correction) is an optional signal that allows using FEC to improve the bit error rate (BER). By using an additional wire when FEC is enabled, the payload data rate is not affected and the wire data rate is unaffected. This allows $F(\text{PCLK}) = F(\text{TxClk}) / 2^n$ with FEC off or on, which simplifies the clock generation and serialization functions. If used, FEC is implemented in the Link Layer, and the PHY treats the FEC bit the same as the other data bits.

AUX is an optional signal that may be used for purposes such as Data Bus Inversion (DBI), flow control, redundancy for defect repair, etc. The Link Layers of Chiplets A and B will need to agree on the details on FEC and AUX usage. An implementation may choose to support the FEC and AUX wires, or to omit both of them. If FEC and AUX are included in a PHY implementation, the PHY carries them in the same way as the data bits without acting on the content.

Table 4 summarizes these signals.

Function	# Wires	Signal Name	Notes
Clock	2	CLK+, CLK-	Differential
Data	16	D0-15	
Forward Error Correction	0/1	FEC	Optional
Auxiliary	0/1	AUX	Optional

Table 4. BoW Signals at the Die To Die Interface

6.1.1. DBI on the AUX wire

Data Bus Inversion (DBI) may be used to mitigate simultaneous switching output (SSO) noise or to optimize energy of a BoW PHY by reducing the number of BoW data wires that switch between adjacent data transfer cycles. DBI functionality is optional; it one of several possible uses of the AUX wire. If implemented, DBI is in the Link Layer and must be implemented on both RX and TX.

6.2. Slice Logic Interface

Figure 3 shows the data and control signals in the interfaces to the logic in the die in each BoW transmit and receive slice. The data at the slice logic interface must be SDR (Single Data Rate - bit rate equal to the PCLK frequency).

6.2.1. Slice Logic Interface: Data Signals

The signals in Table 5 shall constitute the data and clocks in the logic interface of the PHY. N is the ratio of the chip-to-chip per-wire data rate to the logic interface per-wire data rate.

Signal	# Bits	TX Slice	RX Slice	Description
Data	16*N	In	Out	Data
FEC	N or 0	In	Out	Forward Error Correction (optional)
AUX	N or 0	In	Out	Auxiliary uses (optional)
PCLK	1	Out	Out	
TxCk	1	In	NA	Comes from a PLL or other clock source, not the Link Layer. The TxCk source is usually shared among many TX slices. May be differential
RxCk	1 or 0	NA	Out	May be differential

Table 5. Logic Interface Signals

6.2.2. Slice Logic Interface: Control Signals

A BoW interface must provide the following control and status signals:

- PHY Ready
- PHY Reset

The signals in Table 6 shall constitute the control portion of the logic interface of the PHY.

Signal	# Bits	TX Slice	RX Slice	Description
PHYResetB	1	In	In	Resets the BoW slice. 0 causes a reset
PHYReady	1	Out	Out	Indicates that the PHY is ready to transmit/receive mission mode data. 1 indicates ready

Table 6. Logic Interface Control Signals

6.2.2.1. PHYResetB TX and RX

The PHYResetB pin shall be asserted by the link controller to initialize the PHY. While the PHYResetB signal is asserted, the PHY shall stay in its reset state. When the PHYResetB signal is de-asserted, the PHY shall perform any necessary self-alignment. The reset states are otherwise implementation-dependent and shall be documented in the datasheet of a particular implementation.

6.2.2.2. PHYReady TX

On a TX slice, the PHY shall assert PHYReady to indicate it is transmitting appropriate CLK and PCLK signals, and that it is ready to transmit data.

6.2.2.3. PHYReady RX

On an RX slice, when PHYResetB is deasserted, the PHY assumes that the corresponding TX slice is sending CLK and that the TX Link Layer is sending training data on the data wires.

After the RX slice clock self-alignments are complete, each RX PHY slice shall assert its PHYReady pin. How an RX PHY slice determines completion of the self-alignment is implementation-dependent. For instance, it may be determined by observing the settling of the DLL or by a simple timer. PHYReady asserted indicates that any data received will be captured correctly.

6.2.3. Programming

There shall be an AMBA APB programming interface to control internal registers for control and status readout of the PHY.

The internal registers are implementation-dependent. The internal registers shall be fully documented in the PHY datasheet.

6.2.4. Link Controller

There shall be a Link Controller (LC) outside the PHY. This will manage initialization of the Link. It may reside on one of the chiplets of the link, in a third chiplet in the package or outside the package.

Communication from the Link Controller across chiplets shall be by a transport mechanism outside the BoW link. This could be a serial link like SPI or I2C, but this is not specified at this time.

Link initialization is described in Section 12. Clocks are described in 10.2.

7. BoW Modes and Reach

A BoW PHY slice must conform to at least one of the BoW Modes seen in Table 3. The recommended maximum wire reach for different packaging types and terminations is seen in Table 7. Exceeding these reach values may degrade the voltage margins at the receiver. See section 11 for how TX, RX and channels are qualified.

“Laminate” is intended to include organic laminate packages (a.k.a. “buildup”) and similar technologies with approximately 25 μm line and space rules. The minimum wire length for closely spaced chips in these technologies is around 3 mm for the slice closest to the chip edge.

“Advanced” is intended to include silicon interposer and similar technologies. These have much finer line and space dimensions, but traces are usually much more resistive than in organic laminate packages and will be limited to much shorter trace lengths. Due to these short traces, termination is not expected to be useful for implementations targeting Advanced packaging. The minimum wire length in these technologies may be less than 1 mm.

Package			Laminate	Laminate	Laminate	Advanced
Termination			None	Source	Double	None
BoW Mode	Wire Bit Rate	TxClock	Reach	Reach	Reach	Reach
	(Gbps/wire)	(GHz)	(mm)	(mm)	(mm)	(mm)
BoW-32	2	1	10	20	25+	4
BoW-64	4	2	NA	10	25+	2
BoW-128	8	4	NA	5	25+	1
BoW-256	16	8	NA	NA	25+	NA

Table 7. Recommended BoW Wire Reaches

Adding termination increases the speed and/or reach, at the expense of greater design complexity and power.

8. BoW Physical Configuration

8.1. Dead-Bug Views

The physical diagrams and descriptions in this document must be interpreted as looking down at the top layer of the unpackaged chiplets. Since these are flip-chip packages, these views are equivalent to looking through the bottom of the package with the balls up (dead bug view). For the view as seen looking down on a package as mounted on a PCB (live bug view), these views must be mirrored.

8.2. BoW Components




Figure 4. BoW Link Components

A BoW link between two chiplets is made up of wires, slices, and stacks as seen in Figure 4.

- The signal traces in the package between chiplets are called **wires**.
- A **slice** is the basic unit of a BoW PHY. It must have 18 or 20 signal bumps. It must have 2 bumps for the differential clock and 16 single-ended data bumps. It may also have the optional single-ended signals AUX and FEC. The long edge of a slice must be parallel to the chip edge.
- A **stack** is composed of one or more slices stacked from the chip edge towards the center. The slice positions are designated A, B, C, etc, starting with the slice closest to the edge of the chip.

- The signals for a TX slice are in the following order at the chip edge, going clockwise around the chiplet in a dead-bug view: AUX, D0, D1, D2, D3, D4, D5, D6, D7, CLK+, CLK-, D8, D9, D10, D11, D12, D13, D14, D15, FEC
- The signals for an RX slice are in the reversed order (ascending goes counter-clockwise)
- The same clockwise/counter-clockwise ordering is used on all four sides of a chiplet
- The AUX and FEC signals may be omitted

8.6. Bump Arrangements

Note that bump patterns are not specified by BoW; only the signal *ordering* at the chip edge is specified for interoperability.

The reference example in Figure 5 uses hexagonal closest packing for the bumps: two rows for signal bumps and one row for power and ground bumps. In this pattern, the wire pitch is half the bump pitch.

8.6.1. Alternate Bump Arrangements

Alternate bump arrangements may include:

- 90-degree rotation of the hexagonal packing direction (to decrease the wire pitch 14%)
- square bump arrays instead of hexagonal (for regularity of layout)
- more than two rows of signal bumps (to decrease the wire pitch without changing the bump pitch)
- different ordering of power and ground bumps
- multiple power and ground rows

Somewhat different wire pitches between two chiplets may be accommodated with fan-out in the chip-to-chip wires. This is limited by the maximum skew due to different wire lengths - see section 11.1.

8.7. Cross Section

An example cross section for an organic laminate (a.k.a. “buildup”) package is shown in Figure 6.

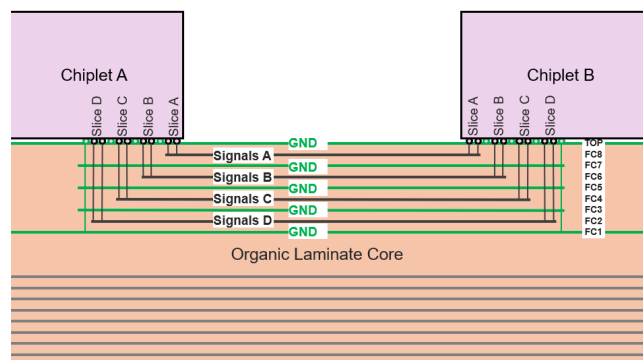


Figure 6. Cross section of a BoW Link in an Organic Laminate Package

In an organic laminate package, signal layers should be alternated with ground layers in order to maintain a controlled impedance of 50 Ω . Each slice position (A, B, C, D) should be

associated with one signal layer and there should be no mixing of signals from multiple slices.

In any technology, the position-A slice on chiplet A must be connected to the position-A slice on chiplet B (one must be configured for TX and one for RX). The position-B slices are connected together, and so on.

There is no specified limit to the number of slices in a stack. In organic laminate, the practical limit in 2020 is an 8-2-8 laminate which supports 4 slices as shown in Figure 6. A 7-2-7 laminate may support 4 slices by omitting the top GND layer, but with reduced signal integrity. Layers on the bottom side of the package typically cannot be used for BoW signals due to low via density passing through the thick central core layer.

In advanced packaging technologies, the shorter wire lengths and higher wire resistance suggests the use of non-controlled-impedance wires and unterminated transmitters and receivers. The smaller wire and space dimensions may allow the wires for multiple slices to be interleaved on a single wiring layer. The wire order within each slice must be maintained, even if interleaving with other slices is used.

8.8. Staggered Slices

To optimize the density of hexagonal bump arrays, slices in positions B and D may be offset horizontally by one half the bump pitch as seen in Figure 7. This necessitates a one-bump-pitch horizontal jog in the wires for slices B and D. The practical effect of this 130-um jog across a 2.5+ mm wire between chiplets is very small.

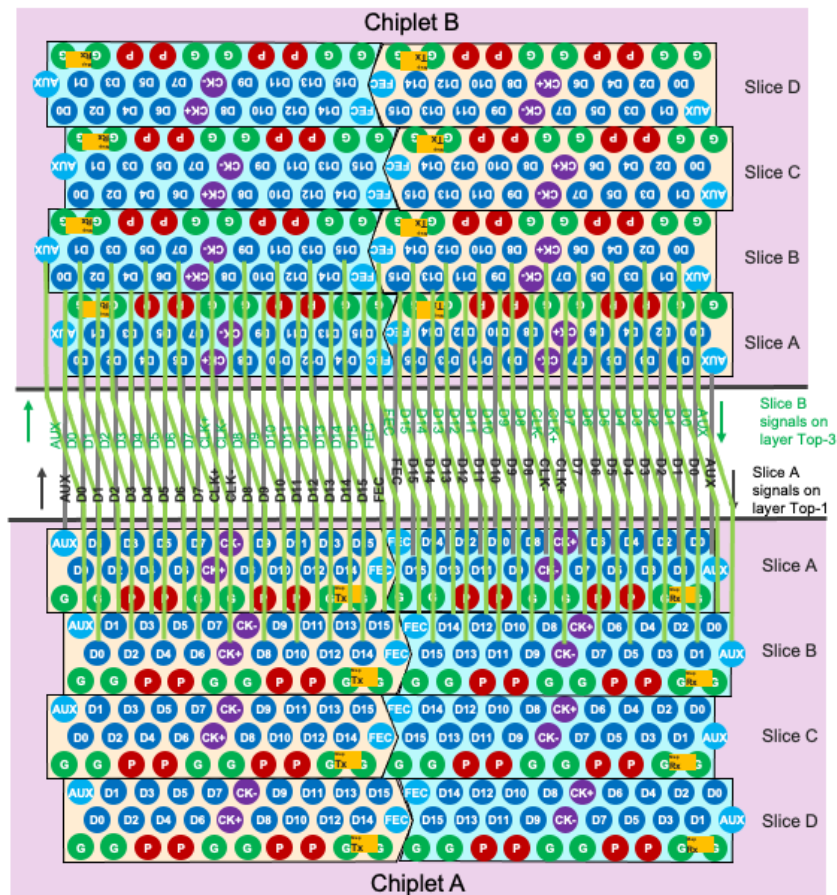


Figure 7. Staggered slices for the densest bump packing

An alternative arrangement is to keep the slices aligned vertically. This requires adding a small extra vertical space between the slices, for an overall increase of 4% of the slice area.

8.9. Slice Numbering

A BoW interface must conform to these slice numbering rules:

- The TX slices in a link are numbered from 0 at the upper left edge of the link (facing from the chip center to the edge in a dead-bug view) and ascending through the TX slices in a stack, then from stack to stack clockwise.
- The RX slices in a link are numbered from 0 at the upper right edge, through the RX slices in a stack, then stack to stack counterclockwise.

An example of this numbering is shown in Figure 8.

The signal ordering and slice numbering rules allow BoW chiplets to be connected without signal reordering regardless of chiplet rotations.

8.10. Slice Stacking Pattern for Symmetric Links

For bidirectional links, a pattern of alternating TX and RX stacks should be used. Figure 8 shows an example bidirectional link with 4 stacks of 4 slices each, for 8 TX and 8 RX slices on each chiplet. The first TX stack should be at the left edge of the link.

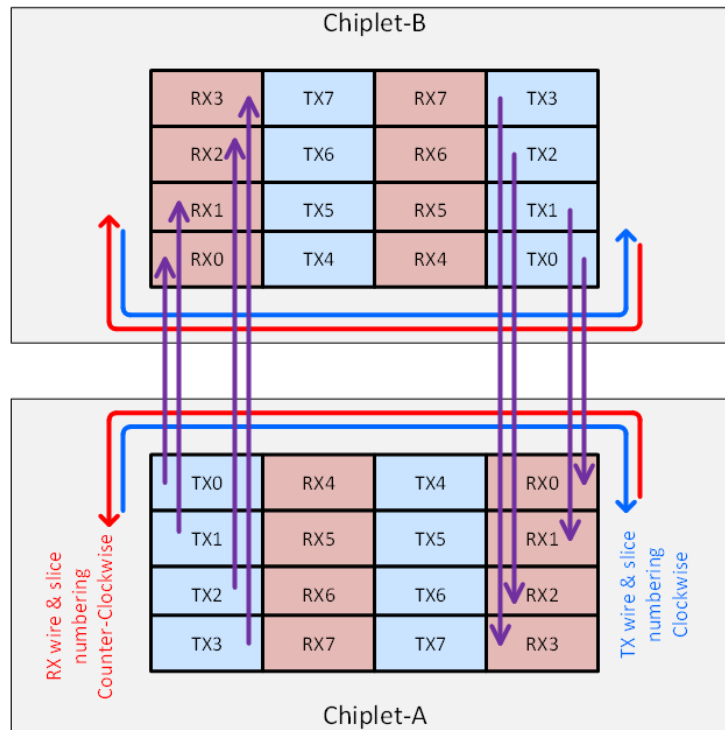


Figure 8. Alternating-Stacks Pattern of TX and RX Slices in a Link

Asymmetric and unidirectional links may use any slice pattern, but the slice numbering rules must be observed.

An alternate approach with more flexibility is to design every slice to operate as either RX or TX, to be configured after assembly or upon powerup. This allows complete flexibility in link configuration and interoperability and also provides an opportunity for wafer-level loopback testing. In this case, number the slices as if they are all TX slices.

In BoW-256 at 16 Gbps/wire, the link in Figure 8 provides a total of 2.0 Tb/s in each direction. In an organic substrate using the hexagonal bump pattern of Figure 5 with a bump pitch is 130 μm , the total edge width is 5.2 mm (4.16 mm without AUX and FEC); the depth from the edge is 1.35 mm. In an interposer, if the bump pitch is 40 μm , the edge width is 1.60 mm (or 1.28 mm) and the depth is 0.42 mm.

9. BoW PHY Electrical Specifications

In order to ensure interoperability between differing BoW PHY implementations, this chapter provides a set of electrical specifications that all such BoW PHY implementations must meet.

9.1. Voltages and Termination Resistance

All BoW implementations must support signaling based on a 0.75 V “I/O voltage”. BoW PHYs may also support higher or lower signaling voltages, but must support 0.75 V based signaling for interoperability.

Note that the simplest implementation is to provide a 0.75 supply voltage to the BoW VDD bumps, but the supply voltage may be different from the I/O voltage as long as the signal voltages meet the specification.

In doubly terminated modes of operation, the RX termination resistance must be connected to 0V, the I/O voltage, or mid-rail of the I/O voltage (e.g., 375 mV with a 0.75 V I/O voltage). The selection of termination voltage is expected to be static (hardwired) in the RX, and must be specified in the receiver's datasheet. It is expected that Source-Series-Terminated (SST) Transmitters will be largely agnostic to the choice of termination voltage on the receiver.

Regardless of the value selected for the I/O supply voltage, BoW transmitters and receivers must meet the DC termination resistance requirements defined in Table 9. Note that TX/RX termination (output/input) resistance values are skewed low/high compared to the channel impedance in order to ensure that the DC single-ended voltage swing at the RX is never reduced to less than half of the I/O voltage (i.e., 375 mV for a 0.75 V I/O voltage). Note that these termination resistance values must be met with all combinations of data inputs (logical 1 and logical 0), termination voltage selections (ground terminated, supply terminated, or mid-rail terminated), and termination resistance values. For example, a BoW TX must achieve between 36 Ω and 50 Ω resistance when driving a load resistance (modeling the RX termination) of 50-69 Ω with any of the three termination options.

	Unterminated	Source-Terminated	Doubly Terminated
TX DC Term.	As required to meet TX rise-time	36 Ω - 50 Ω (0.72 - 1.0 Z_{chan})	36 Ω - 50 Ω (0.72 - 1.0 Z_{chan})
RX DC Term.	-	-	50 Ω - 69 Ω (1.0 - 1.38 Z_{chan})
Within-Slice DC Term. Matching	-	$\sigma = 1.333\%$ (8% over 6 σ)	$\sigma = 0.667\%$ (4% over 6 σ)

Table 9. TX and RX Termination Resistance Requirements vs. Mode

Especially in doubly terminated modes, within-slice variations of termination resistance would directly result in varying swing levels at each pin. Thus, in order to reduce or

eliminate the need for per-pin voltage reference adjustment at the RX, Table 9 also specifies requirements on DC termination resistance matching across all I/O's within a given BoW slice. The σ for this variation in the table must be interpreted as capturing within-slice manufacturing variability across worst-case voltage/temperature operating conditions, and is expected to be primarily influenced by some combination of transistor and explicit resistor matching (with the mix depending on the circuit implementation).

9.2. PHY protection

A BoW PHY should not draw excessive current nor be damaged under the following conditions:

- The bumps are open-circuited, e.g., at wafer test, or if not connected in a package assembly during power-up.
- In the un-powered state when connected to a matching PHY which may be powered or un-powered.
- In the reset state when connected to a matching PHY which may be powered or un-powered.
- In the operational state.

In particular, RX slices should avoid “crowbar” states. This may be done by disabling the receiver circuit or with a weak pulldown at the bumps (except CLK-, which may get a pullup or may be left open). Pullups should not be used on bumps other than CLK-.

A PHY slice which can be configured as TX or RX must have the TX circuit disabled when PHYResetB is asserted and an explicit APB command must be performed to turn on the TX circuit.

9.3. ESD

BoW I/O shall be designed to withstand **50 V CDM** (Charged Device Model) and **250 V HBM** (Human Body Model) at the bumps. This requirement is deemed sufficient for intra-package signaling, similar to other die-to-die interface standards.

9.4. Return Loss and Parasitic Capacitance

Since BoW PHYs are targeted for relatively dense and simple realizations, it is expected that the primary frequency-dependent parasitics seen at a PHY's I/Os will be capacitive in nature. Table 10 provides limits on the maximum “equivalent” capacitance allowed on each side of each BoW I/O pin. (E.g., a BoW-128 TX is allowed to have up to 500 fF of equivalent capacitance.) Note that while the maximum capacitance specification does increase at lower data-rates, it is recommended that BoW PHY implementations retain as low of a capacitance as practical in order to reduce power consumption and improve signal integrity.

	BoW-32 or BoW-64	BoW-128	BoW-256
Maximum Equivalent Capacitance (TX or RX)	800 fF	400 fF	200 fF

Table 10. Maximum Parasitic Capacitance at a BoW I/O vs. Mode

Since the actual frequency-dependent impedance profile of any given implementation may be comprised of a complex electrical network, conformance with the “equivalent” capacitance metric is formally defined by requiring that the magnitude of the return loss of any BoW I/O must be lower than the maximum limits shown in Figures 9 and 10 below. (Note that the return loss requirements are different for TX and RX because of the differences in DC termination between the two sides.) Similarly to DC termination

resistance, the maximum s_{11} magnitude in the figure must be met with all combinations of data inputs (logical 1 and logical 0), termination voltage selections (ground terminated, supply terminated, or mid-rail terminated), and termination resistance values.

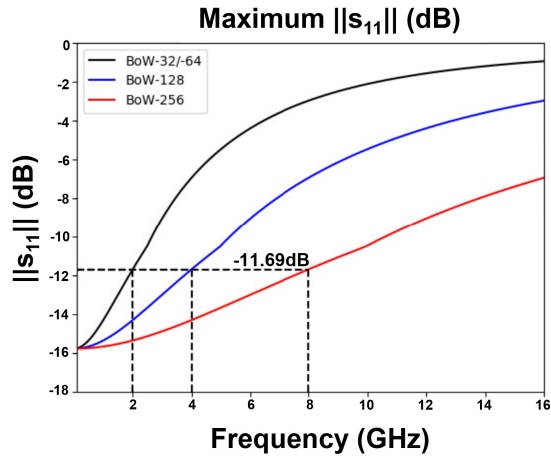


Figure 9. BoW TX Termination Maximum Return Loss




Figure 10. BoW RX Termination Maximum Return Loss

9.5. Receiver Bandwidth

While this specification does not place a direct requirement on the bandwidth of a BoW receiver implementation, such receivers should maintain an effective 3 dB bandwidth of at least $(0.667/T_{bit})$ Hz. For example, for a BoW-256 PHY, the receiver 3 dB bandwidth is recommended to be at least 10.667 GHz.

10. BoW PHY Timing Specifications

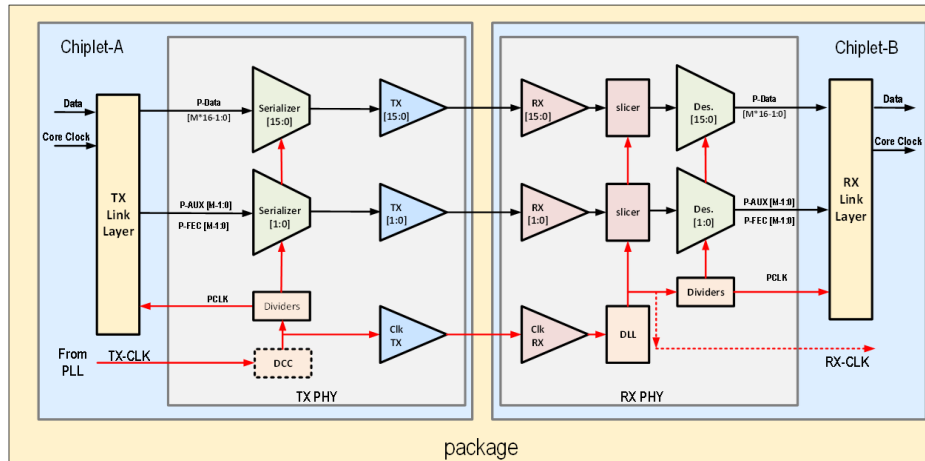


Figure 11. BoW Clock and Data Block Diagram - One TX Slice, One RX Slice

10.1. Bit Ordering

The PHY TX serializer shall order data this way (referring to Figure 11):

- On the first CLK edge (CLK+ rising) bits P_D[0:15] are sent on wires D[0:15]
- On the second CLK edge (CLK+ falling) bits P_D[16:31] are sent on wires D[0:15]
- and so on to bits P_D[M*16-16:M*16-1].
- Then the cycle repeats.

The RX PHY shall order bits in the same fashion. However, the bits at the RX PHY logic interface P_D[*] may be offset by a multiple of 16 bits from the TX order if the TX and RX PCLK dividers are not aligned. A PHY implementation may provide a way to align the TX and RX dividers, or it may rely on the Link Layer to rotate the RX P_D[*] bits to provide that alignment as part of the training of the Link Layer.

10.2. Clocking

Figure 11 shows the clock and data flow for a single TX slice and a single RX slice. On the TX side, data bits (and optional FEC and AUX bits) come in a wide word from the Link Layer, and are serialized to the line rate. At the RX side, they are sampled with a common slicer clock in most BoW implementations. BoW PHYs may optionally implement per-bit delay adjust or per-bit slicer clock adjust.

BoW PHYs shall be DDR (Double Data Rate) at the chip-to-chip interface: the data bit rate is twice the clock frequency, so data is clocked in on both edges of the clock in the RX slice. BoW PHYs shall be SDR (Single Data Rate) at the logic interface.

Table 11 provides recommended clock and data rates for each BoW mode. The ratio M should be limited to integers, preferably powers of two, and any other ratios should be implemented outside the PHY.

Note that higher PCLK rates (and lower M ratios) help reduce gate count and Link Layer latency, but lower rates are often more power efficient. The best PCLK rate(s) to implement for a particular chiplet will tend to be a function of its process node. For implementations in process nodes at 16 nm and below, supporting 1000 MHz is recommended.

	Data	PCLK	Mux	Logic
--	-------------	-------------	------------	--------------

Mode	Data Rate (Gbps)	PCLK (MHz)	Ratio M	Logic Width
BoW-2	2	250	2	2x18
BoW-64	4	250	16	16x18
		500	8	8x18
		1000	4	4x18
		2000	2	2x18
BoW-128	8	500	16	16x18
		1000	8	8x18
		2000	4	4x18
BoW-256	16	500	32	32x18
		1000	16	16x18
		2000	8	8x18

Table 11. Recommended PCLK and Logic Data Rates for Figure 11

Table 12 provides clock and data rates for an example with 4 Gbps wire data rate and M=4 to support a 1 Gbps data rate at the Link-PHY interface.

Signal	Rate	SDR/DDR
TxCk	2 GHz	
CLK+,CLK- D[15:0],AUX,FEC	2 GHz 4 Gbps	DDR
PCLK	1 GHz	
P_D[63:0],P_AUX[3:0],P_FEC[3:0]	1 Gbps	SDR

Table 12. Example Clock and Data Rates for Figure 11 with 4 Gbps, M=4

The DDR clock TxClk is provided to the TX PHY from elsewhere on Chiplet-A. This may come for example from an on-chip PLL (typically shared across multiple slices) or routed from the RxClk of an RX slice on Chiplet-A. In order to meet duty cycle requirements, a Duty Cycle Corrector (DCC) may be needed in the TX slice. TxClk is used to drive the serializers and provide the output CLK+, CLK- to Chiplet-B.

On the RX side, the PHY must align the slicer clock to sample the data correctly. This may be done with a DLL, adjustable delays, or other methods. If the PHY includes control logic to self-align the slicer clock for correct sampling of the data, the PHYReady signal must be asserted after the logic has determined that such alignment is complete. The RX PHY may output the received CLK as RxClk to the logic interface.

All BoW interfaces shall be source synchronous at the die-to-die interface within a slice. No modes of BoW require per-wire or per-slicer delay adjustments, but such capability may be optionally included.

Clock skew between the slices in each direction of a link likely depends on the implementation of the TxClk distribution to all the TX slices. That is, for the data flow from Chiplet A to Chiplet B, the TxClk distribution on Chiplet A probably dominates the clock skew of the TX slices on Chiplet A and the clock skew of the RX slices on Chiplet B, and vice versa for flow from B to A. The skew between Tx CLK signals within one direction of a link should be no more than 150 ps/stack along the chip edge. There is no specification of the skew between TxClk on Chiplet A vs. TxClk on Chiplet B nor between different links.

Note that the dividers creating PCLK in each PHY slice are not required to be aligned. This implies that they will tend to have random starting states, leading to additional PCLK misalignment between slices of up to one PCLK period. PHY implementations may optionally include methods to align these dividers.

On both the TX and RX sides, the Link Layer will usually need to include a Clock Domain Crossing (CDC) to align the data between CoreClk and PCLK. The Link Layer must be able to absorb the slice-to-slice clock skew and core clock distribution skew across a whole BoW link. Word alignment across a link need not be supported by the PHY; if required, it should be done in the Link Layer.

10.3. Clock and Data Specifications

In order to not introduce excess pessimism into the link budgets implied by the BoW specification and avoid unnecessary over-design of BoW PHY circuitry, note that both the TX and RX voltage and timing error component requirements account for deterministic (bounded) terms separately from random (unbounded) terms. However, in order to retain some degree of design flexibility on each of the TX and RX, a bound is always placed on the maximum deterministic error and on the maximum total error budget at the target error rate of $1e-15$. Thus, if a given BoW PHY design achieves deterministic error performance better than that requirement set by the deterministic component, the random errors introduced by that design may be increased as long as the total error requirement at $1e-15$ probability is still met.

Note that the error rate of $1e-15$ is at the level of any individual wire within a slice. In other words, in a conformant BoW interface, no individual wire within the interface would have an error-rate exceeding $1e-15$.

10.3.1. Transmitter Maximum Rise-Time

The maximum 20% - 80% rise-time at the output of BoW TX shall not exceed **23% of a UI**. For example, for a BoW-128 Transmitter, the 20% - 80% rise-time shall not exceed 31.25 ps. This rise-time shall be simulated with the TX (including all of its parasitics) driving an ideal load of 50Ω (Z_{chan}).

10.3.2. Transmitter Correlated Jitter Filtering

For timing-error specifications provided in the following sections that are impacted by transmitter jitter, this jitter must be evaluated for CLK edges that are up to **3 UI** earlier than the CLK edge that launched the data bit being captured at the receiver. This is due to the fact that even though jitter on the data edges may be correlated with the CLK jitter, the slicer in the RX side is likely to use a different CLK edge due to delays in the RX-side clock alignment circuit (usually a DLL and clock distribution).

In order to properly account for the jitter filtering/peaking that will occur due to the difference in delay between the data launching edge at the TX and the data capturing edge at the RX, when evaluating the transmitter's jitter (and whether it meets the requirements described in this document), the jitter at the TX output that is **correlated** between the CLK and D lines shall be filtered by the following frequency-dependent transfer function:

$$H_{tx_jit}(j\omega) = 1 - e^{-j\omega t_{clk-d}}$$

where t_{clk-d} is the delay between the CLK edge that launched the data bit and the CLK edge used to capture it. Note that jitter that is not correlated between the CLK and D signals shall **not** be filtered by this transfer function. (I.e., if the CLK signal and a given D signal have completely independent sources of jitter added to them such as non-shared portions of the clock distribution network, those jitter sources shall not be filtered by $H_{tx_jit}(j\omega)$). Since the total TX jitter after filtering by this transfer function might not be monotonic with t_{clk-d} , and since receiver implementations may realize varying values of t_{clk-d} , a transmitter must meet all related specifications for $t_{clk-d} = T_{bit}$, for $t_{clk-d} = 2T_{bit}$, for $t_{clk-d} = 3T_{bit}$, and for $t_{clk-d} = 4T_{bit}$

10.3.3. Transmitter Deterministic Timing Error

The total deterministic (bounded) timing errors introduced by the TX shall not exceed **14% of a UI peak-to-peak**. The evaluation of these timing errors must include all possible deterministic contributors, such as reference clock, clock distribution networks, duty cycle error (i.e., deviation from 50% duty cycle), skew between CLK and any D line, and power supply variation induced jitter or skew. Note that any such time-dependent error terms (i.e., jitter) that are correlated between the CLK and D lines must be filtered as described in Section 10.3.2.

This specification is a peak-to-peak requirement, so if a given design has e.g. +/-5% UI of duty cycle error, this would imply that it can achieve a TX deterministic timing error of no better than 10% UI.

10.3.4. Transmitter Total Timing Error

The total timing error introduced between the CLK and any data (D) line at the output of the TX shall not exceed **25% of a UI peak-to-peak at an error rate of 1e-15**. The evaluation of errors must encompass all possible deterministic as well as random timing error contributors, including all sources of random jitter in addition to the representative deterministic error sources described in Section 10.3.3.

Assuming a Gaussian distribution for the random jitter, then in order to account for the 1e-15 error rate and peak-to-peak requirement, the total timing error $t_{err,tot}$ may be computed as:

$$t_{err,tot} = t_{err,deterministic} + 15.9 \sigma_{tj,random}$$

10.3.5. CLK Receiver Sensitivity to Common-Mode Variations

The differential receiver for the CLK signal within a BoW receiver must achieve an input-referred common-mode to differential conversion gain of less than **0.2 V / V**. This requirement must be met across any common-mode input frequency less than or equal to $1/T_{bit}$. For example, with a conformant BoW RX, 20mV of common-mode variation on the CLK+/CLK- lines must impact the effective differential input by less than 4mV.

Note that common-mode variations on the CLK+/CLK- lines of ~10-15% of the signal swing may be expected on typical in BoW channels.

10.3.6. Receiver Sensitivity and Timing Margin

A BoW receiver must meet a set of requirements on the following sets of timing and voltage error components:

- **Maximum RX Deterministic Voltage Error:** This term ($V_{err,det,RX}$) must include all deterministic voltage errors that would shift the receiver's voltage threshold relative to its ideal position in the middle of the signal swing. For example, any deterministic voltage errors due to residual offset, reference level error, and supply noise must be included. This specification accounts for the double-sided loss in margin, so if a given design has e.g. a residual threshold error of 0mV to +10mV, this would imply that the design can achieve a $V_{err,det,RX}$ of no better than 20 mV.
- **Maximum Total Required RX Voltage Margin:** This term must include all possible voltage errors at the RX at a probability of 1e-15 or higher. In addition to the deterministic RX voltage error sources mentioned above, error sources such as receiver thermal/flicker noise must therefore be included in this term. For Gaussian random voltage noise, the total required voltage margin ($V_{err,tot,RX}$) may be computed as:

$$V_{err,tot,RX} = V_{err,det,RX} + 15.9 \sigma_{Verr,random,RX}$$

- **Maximum RX Deterministic Timing Error:** This term ($t_{err,det,RX}$) must include all deterministic timing errors that would shift the receiver's sampling timing relative to its ideal position for any data line. This term must therefore include errors due to e.g. residual sample timing position error, DLL dither, and power supply induced jitter. This specification accounts for the double-sided loss in margin, so if a given design has a mismatch-induced shift of the sampling position (relative to the ideal) of 0% to 5%, this would imply that the design can achieve a $t_{err,det,RX}$ of no better than 10% UI.
- **Maximum Total RX Timing Error:** This term must include all possible timing errors at the RX at a probability of 1e-15 or higher. In addition to the deterministic RX timing error sources mention above, error sources such as RX clock receiver or clock distribution random jitter must be included in the total timing error. For Gaussian random jitter, the maximum total required timing error ($t_{err,tot,RX}$) may be computed as:

$$t_{err,tot,RX} = t_{err,det,RX} + 15.9 \sigma_{terr,random,RX}$$

Since swing and signal integrity are expected to vary with termination as well as data-rate, the RX voltage and timing requirements are termination- and rate-dependent, as outlined in Table 13.

	BoW-256	BoW-128	BoW-128	BoW-64 or BoW-32	BoW-64 or BoW-32
	Any Termination	Doubly Terminated	Source- or Unterminated	Doubly Terminated	Source- or Unterminated
$V_{err,det,RX}$	40 mV	40 mV	100 mV	65 mV	150 mV
$V_{err,tot,RX}$	75 mV	75 mV	150 mV	100 mV	200 mV
$t_{err,det,RX}$	32% T_{bit}	32% T_{bit}	32% T_{bit}	28% T_{bit}	28% T_{bit}
$t_{err,tot,RX}$	40.5% T_{bit}	40.5% T_{bit}	40.5% T_{bit}	36.5% T_{bit}	36.5% T_{bit}

Table 13. Receiver Voltage and Timing Requirements for BoW

10.3.7. Voltage Overshoot

This specification does not place a specific requirement on the overshoot observed by the RX, but it is expected that the overshoot should have magnitude of less than 300 mV for 750 mV I/O supply. Since overshoot is most likely to create potential reliability and other issues in unterminated operating modes, BoW RX's are allowed to turn on termination resistors to reduce the overshoot they observe. The value of the RX termination resistance in this case must be larger than 50 Ω , but is otherwise unconstrained as long as the receiver is able to meet its timing margin requirements with the resulting swing/channel.

BoW TX's therefore shall be designed to achieve their lifetime, reliability, and other requirements regardless of whether the BoW RX selects to operate with or without termination.

10.3.8. Slice-to-Slice CLK Skew

The slice to slice clock skew t_{skew} across the width of a BoW link (along the chip edge) should be less than 150 ps/stack. (E.g., for a 4-stack interface, the skew from end to end must be less than 600 ps.) This skew includes only analog delays and specifically does not include any clock-related timing skew due to flip-flops/latches or varying reset states.

This skew is expected to be dominated by the TxClk distribution network.

11. Chip-to-Chip Channel Specifications

BoW does not place any direct requirements on characteristics such as channel loss or crosstalk. Instead, BoW channels are considered conforming if they are able to achieve the required error rate of $1e-15$ in conjunction with reference transmitters and receivers that meet all of the requirements provided in Section 9 and Section 10.3.

To assist with evaluating conformance of a given channel, open-source software evaluating signal integrity and the overall link budgets with the reference transmitters and receivers will be provided at a future date.

11.1. Channel-Induced CLK Edge to D Transition Skew at RX

Within each slice, a BoW channel should not introduce more than **2% UI** of skew between any D lines and the CLK lines. For BoW-256, this corresponds to $\sim 187.5 \mu\text{m}$ of length mismatch on a substrate with an ϵ_r of 4.

Note that the skew recommendation above is based on achieving sufficient timing margin on representative channels; channels with better signal integrity may allow for larger skew between the D and CLK lines as they meet the overall timing margin requirements.

Note further that if the BoW PHYs used on a given channel include per-bit delay adjustment, channels with larger skew can be supported. Note however that all of the timing requirements described in Section 10.3 must then be met with the residual skew and its variation over time taken into account.

11.2. Channel Impedance

In laminate packages, the channel characteristic impedance should be between 45 and 55 Ω .

11.3. Example BoW Channel for Doubly-Terminated Links

To provide guidance on the types of channels that are expected to meet the requirements for conformance with the BoW reference receiver and transmitter, this section provides examples of typical loss and crosstalk profiles for doubly-terminated channels supporting 16 Gbps operation (which are most sensitive to channel signal integrity). Note that when operating at lower rates, the frequency axes in the figures below should be scaled with the data-rate relative to 16 Gbps.

11.3.1. Channel Loss

To avoid the need for equalization, a BoW-256 channel should typically have lower loss than shown in Figure 12.




Figure 12. BoW Doubly-Terminated Wire Channel Loss Limit

11.3.2. Crosstalk

The total crosstalk observed on an individual signal within a BoW-256 channel should typically be less than ~35% of the signal swing.

12. Reset and Initialization

12.1. External Facilities

These facilities must be provided outside the PHY:

- A Link Controller (LC) which will manage initialization of the Link. It may reside on one of the chiplets of the Link, in a third chiplet in the package or outside the package.
- A communication path from the Link Controller to the PHY slices outside the BoW link. This could be a serial link like SPI or I2C, but this is not specified at this time.
- A source of training pattern data outside the PHY, assumed to be the Link Layer here. This must be able to repetitively transmit an arbitrary 16 bit per wire pattern (256 or 288 bits pattern depending on inclusion of FEC+AUX) required by the RX slice for clock alignment as specified in the datasheet for the RX PHY.
- The PHYResetB input to each PHY shall be asserted upon powerup. It may also be asserted by commands from the LC.

An example topology is shown in Figure 13. The BoW interface communicates to interface and core logic (I&C) blocks.

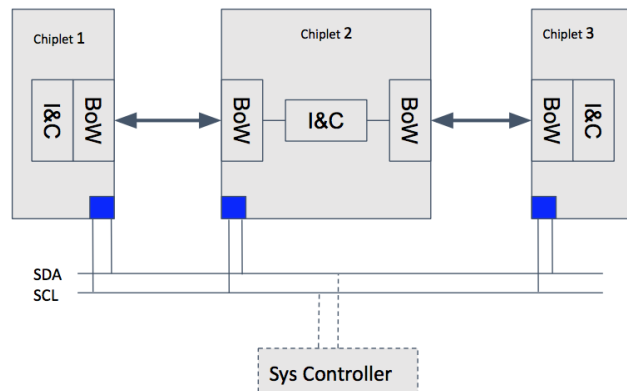


Figure 13. Example BoW System Configuration

[Add another figure with the LC embedded rather than external](#)

12.2. Initialization Sequence

12.2.1. A TX-RX Link should be brought up as follows:

1. The Link Controller (LC) asserts PHYResetB to the PHY slices on both ends of the link.
2. The Link Controller (LC) **de-asserts PHYResetB to the TX** PHY slices
3. The LC performs any needed configuration of the TX PHY slices via the APB interface. This is implementation dependent. If the TX needs to be enabled (see section 9.2), it happens here.

4. Once its outgoing CLK and PCLK stabilize, each TX PHY slice asserts PHYReady to the LC
5. When all TX PHY slices are ready, the LC signals the TX Link Layer to send the training pattern (repeating) specified by the RX PHY
6. The Link Controller (LC) **de-asserts PHYResetB to the RX** PHY slices
7. The LC performs any needed configuration of the RX PHY slices via the APB interface. This is implementation dependent. If the RX needs to be enabled, it happens here.
8. Each RX PHY slice performs clock and data alignment and signals PHYReady to the LC when done
9. When all RX PHY slices are ready, the LC signals the TX and RX Link Layers to proceed with channel bonding

Note that BoW PHY implementations that do not adopt the order recommended above are not guaranteed to interoperate with other BoW PHY implementations.

Implementation dependent:

- Whether the up and down links are initialized one at a time or in parallel
- How the signals from the the LC get from and to the PHYReady and PHYResetB pins of the PHY
- How the Link Layer performs channel bonding or start of data transmission
- Any PHY registers required to implement this process

12.3. Unspecified Items

- Whether the APB registers are separate for each slice or shared among slices in a link
- There is no low-power standby mode defined.
- There is no specification for when a PHY should de-assert its PHYReady pin. PLL or DLL losing lock are possible causes.
- There is no definition of what occurs if the PHY does de-assert PHYReady
- There is no definition of what should be done with unused PHYs (that are on the chip but have no partner on another chiplet)
- There is no definition of logical addressing of chiplets, Links or slices
- Possible use of PRBS patterns for training

13. Configuration

PHY configuration is implementation dependent. It may include:

- TX vs RX for configurable slices
- PLL, DLL, DCC or similar circuit configuration

PHY configuration may be hardwired in the chiplet implementation, or it may be programmable.

13.1. Link Training

Link training will be addressed in a future revision of the specification.

14. Control Register Mapping

The interface control registers are implementation dependent. The registers shall be fully documented in the PHY datasheet.

15. Testability

15.1. Test Patterns

In order to support die-to-die (in package) testing, within a BoW implementation, either the link layer, the PHY, or both must be able to support the generation (on the TX side) or checking (on the RX side) of repeating data patterns.

Users of BoW systems should check that one or more of the test patterns supported on the TX is also supported on the RX. Such pattern generators / checkers should therefore support the following patterns:

- PRBS-9 Pattern, defined by polynomial of $X^9 + X^5 + 1$
- PRBS-31 Pattern, defined by polynomial of $X^{31} + X^{28} + 1$
- Isolated 1 and 0 pattern to test DC wander and single bit response:
 - $['0'] \times 10 + ['1'] + ['0'] \times 10 + ['1'] \times 10 + ['0'] + ['1'] \times 10 + ['0'] \times 10$
 - This may be prepended to a PRBS pattern as seen in Figure 14




Figure 14. Stress Test Pattern

15.2. Loopback Test

A BoW interface may implement loopback testing for several use cases: at chiplet wafer-sort test, post-assembly package test, and debug/validation.

Wafer sort tests are currently only practical for the BoW interface with regular bump pitches ($\sim 130 \mu\text{m}$), where ATE (automatic testing equipment) probe boards with matching pin pitches are available. Microbump probes will require additional effort.

Unidirectional links should support open-loop testing. In TX open loop testing, shown in Figure 15, Chiplet-A transmits a known test pattern (PRBS9 or PRBS31) to a golden reference receiver through the ATE load board. The received pattern should be verified in the ATE load board.

RX open loop testing, shown in Figure 16, is used for a link where the DUT is only a receiver. A golden reference TX transmits a known pattern (PRBS9 or PRBS31) through the channel to the chiplet. The received pattern should be analyzed for quality and functional tests.

The logic for generating and testing the PRBS sequences is outside the PHY, e.g., in the Link Layer.

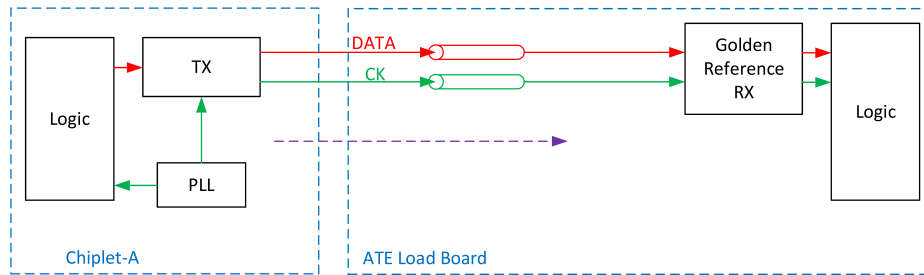


Figure 15. Open loop TX wafer test

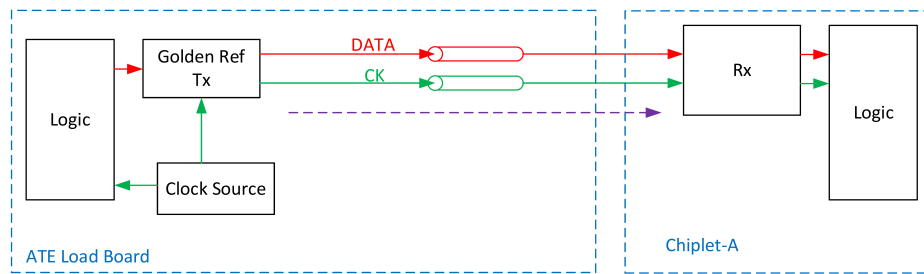


Figure 16. Open loop RX wafer test

In bidirectional links, loopback tests may be implemented in several modes:

- Slice-to-slice short loopback
 - Data is looped back within the chip from a TX slice to an RX slice using on-chip switching (shown in Figure 17). The short loopback path is configured by the ATE using implementation-dependent registers.
 - Loopback may be implemented before the PHY serializer; between the serializer and the output buffer; and/or at the bumps.
- Intra-slice short loopback
 - A single slice containing both RX and TX paths sharing the same bumps may perform on-chip loopback testing simply by turning on both the RX and TX paths at once. This has more on-chip circuitry, but allows loopback testing with no switches or extra lines connected to the bumps other than the TX driver tristate switches. Figure 17 applies, except there is only one shared set of bumps for a TX/RX slice.
- Long loopback
 - The PRBS pattern is generated by chiplet-A, sent over the replica channel on the ATE load board which loops it back (shown in Figure 18). The received pattern should be passed to a bit error rate tester (BERT) to analyze the performance of the link with off-chip data and clock wires.

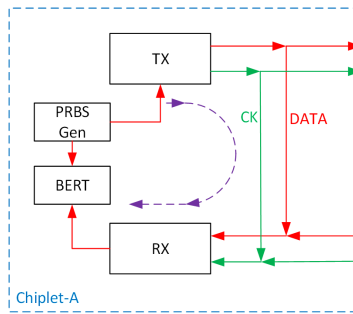


Figure 17. Short loopback testing

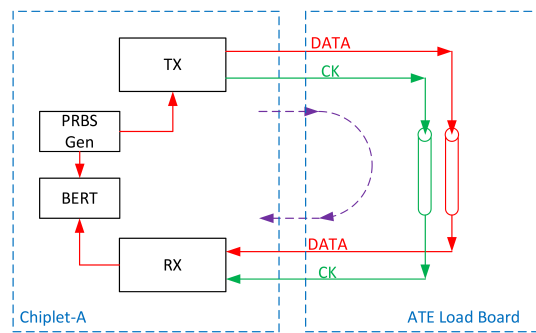


Figure 18. Long loopback testing

Both loopback modes may potentially be used for in-field validation bring-up and test. Cooperation across chiplets will be required to execute these tests in the field. Open-loop testing requires the use of a fixed test pattern recognized by both ends and is the only option for unidirectional links. Long loopback mode can be implemented on interposer or organic laminate for validation/verification purposes.

Figure 19 shows how a long loopback mode is executed across two chiplets for in-field validation and test where TX and RX are in different chiplets. Furthermore, this configuration may be expanded to loop back the data from the transmitter of chiplet-A to the receiver of chiplet-A.

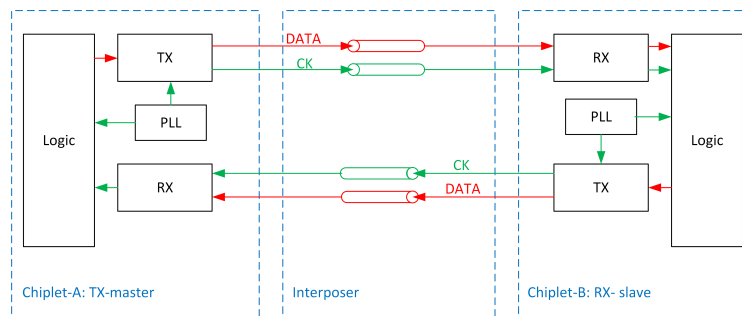


Figure 19. Chiplet-to-chiplet long loopback

16. BoW in an ODSA Design

Chiplet-based designs require logical connectivity between the die in a single package, in addition to physical connectivity. This section provides an overview of how the Open Domain-Specific Architecture stack may be used as an underlay for popular transaction protocols.

16.1. BoW for Common Transaction Protocols

Two connected die in a multi-chiplet device need to exchange logical information. The ODSA aims to define an open physical and logical interface for chiplets, as shown in Figure 20 to enable chiplets from multiple vendors to interoperate and be integrated in a multi-die package. The Bunch of Wires is an open D2D PHY option in the interface. The logical component of the ODSA interface aims to support protocols used for the two most common chiplet use cases, package aggregation and die disaggregation across a wide range of open and proprietary D2D PHYs such as PCIe, CXL, CCIX, AXI and proprietary streaming protocols.

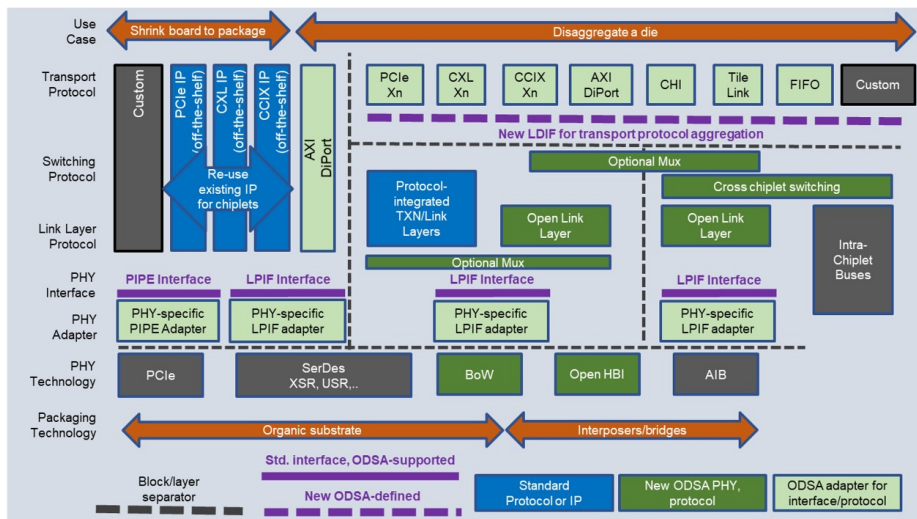


Figure 20. The BoW PHY in the ODSA Stack

The ODSA stack abstracts the PHY layer from the logical interface by using the well-defined abstraction interfaces PIPE and LPIF. Any logic transaction controller, such as a PCIe controller, that supports a PIPE or LPIF interface may use any D2D PHY that also supports that interface as its physical layer. As shown in Figure 20, the BoW interface may receive data through either the PIPE or LPIF interfaces to support common transaction protocols. For this use case, some BoW-specific adapter logic will be needed to support the requirements of PIPE or LPIF. The specifications for these adapters are outside the scope of this document. Figure 21 shows how the BoW with an PIPE adapter may be interfaced to a PCIe controller.

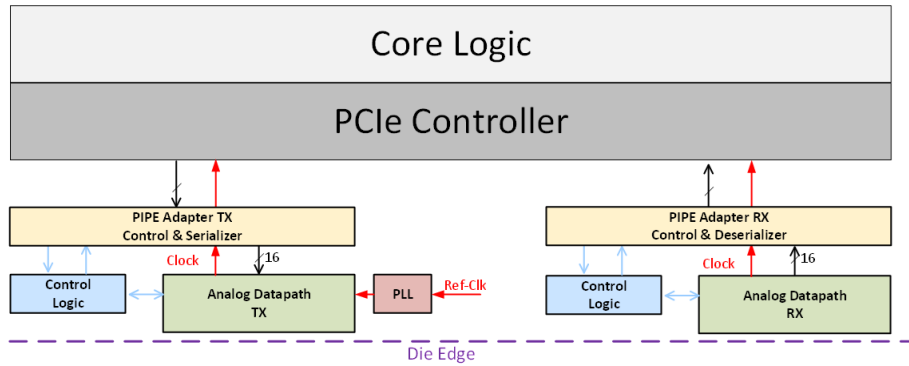


Figure 21. BoW with a PIPE adapter for PCIe transactions

Appendix A - Requirements for IC Approval (to be completed Contributor(s) of this Spec)

List all the requirements in one summary table with links from the sections.

Requirements	Details	Link to which section in the Spec
Contribution License Agreement	OWF CLA 1.0 (modified)	Please refer to Section 1
Are All Contributors listed in Sec 1: License?	Yes	
Did All the Contributors sign the appropriate license for this spec? Final Spec Agreement/HW License?	Yes	
Which 3 of the 4 OCP Tenets are supported by this Spec?	All four	
Is there a Supplier(s) that is building a product based on this Spec? (Supplier must be an OCP Solution Provider)		
Will Supplier(s) have the product available for GENERAL AVAILABILITY within 120 days?	Seeking exception to have extended period for silicon availability. Test chips expected in 2H'2022 by multiple vendors.	

Appendix B