

Interlaken Interoperability Recommendations

19 May 2016

Revision 1.9

Terms and Conditions

This document has been developed with input from a variety of companies, including members of the Interlaken Alliance, all of which have released their respective rights to information contained in this document for the sole and express purpose of implementation of this document and to encourage others to adopt this interface as an industry standard.

Any company or individual wishing to use all or a part of this document may do so only if they relinquish their proprietary rights to information contained or referenced herein. You agree not to enforce those intellectual property rights against any other party implementing or advocating the implementation of this document.

YOU ACKNOWLEDGE THAT THIS DOCUMENT IS PROVIDED “AS IS” AND WITHOUT WARRANTY OF ANY KIND. FOR THE SAKE OF CLARITY, ALL PARTIES USING OR ADVOCATING THE USE OF THIS DOCUMENT DISCLAIM ANY AND ALL WARRANTIES, EXPRESS OR IMPLIED, WITH RESPECT TO THIS DOCUMENT, INCLUDING WITHOUT LIMITATION ANY AND ALL WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE AND NON-INFRINGEMENT.

IN NO EVENT WILL ANY PARTY THAT UTILIZES OR ADVOCATES THE USE OF THIS DOCUMENT BE LIABLE TO ANY OTHER PARTY FOR ANY INDIRECT, INCIDENTAL OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH THIS DOCUMENT, AND IN NO EVENT WILL THE CUMULATIVE LIABILITY OF ANY SINGLE PARTY FOR ANY AND ALL CLAIMS IN CONNECTION WITH THIS DOCUMENT EXCEED, IN THE AGGREGATE, ONE HUNDRED DOLLARS (\$100.00). YOU AGREE THAT THIS PARAGRAPH IS AN ESSENTIAL BASIS OF THE USE OF THIS DOCUMENT.

Contents

Contents	3
Revision History	5
1 Overview	6
2 Logical Interoperability Recommendations	7
2.1 Definitions of Interoperability Parameters	7
2.2 Recommendations for All Transfer Rates	8
2.3 10 Gb/s Packet Transfers	8
2.3.1 Target Applications	8
2.3.2 Recommendations	8
2.4 25 Gb/s Packet Transfers	9
2.4.1 Target Applications	9
2.4.2 Recommendations	9
2.5 50 Gb/s Packet Transfers	10
2.5.1 Target Applications	10
2.5.2 Recommendations	10
2.6 100 Gb/s Packet Transfers	11
2.6.1 Target Applications	11
2.6.2 Recommendations	11
2.7 120 Gb/s Packet Transfers	12
2.7.1 Target Applications	12
2.7.2 Recommendation #1	12
2.8 400 Gb/s Packet Transfers	13
2.8.1 Target Applications	13
2.8.2 Recommendation #1	13
2.8.3 Recommendation #2	13
2.9 Summary of Lane Recommendations	13
2.10 Multirate Interoperability	14
2.11 Latency Requirements	15
3 Low Speed Physical Interoperability Recommendations	18
3.1 Connector Type	18
3.2 Connector Pin Assignment	18
3.3 Connector Placement	20
3.3.1 Board Height Requirements	21
3.4 Physical Interoperability Examples	21
3.4.1 Connecting 10 Gb/s to 10 Gb/s	21
3.4.2 Connecting 25 Gb/s to 25 Gb/s	22
3.4.3 Connecting 25 Gb/s to 10 Gb/s	22
3.4.4 Connecting 50 Gb/s to 50 Gb/s	22
3.4.5 Connecting 50 Gb/s to 25 Gb/s	23
3.4.6 Connecting 50 Gb/s to 10 Gb/s	23
3.4.7 Connecting 100 Gb/s to 100 Gb/s	24
3.4.8 Connecting 100 Gb/s to 50 Gb/s	25
3.4.9 Connecting to Multiple Interlaken Interfaces	25
3.5 PCB Intra-Lane Skew Budget	26
3.6 AC Coupling	27
3.7 Signal Integrity Considerations	27
4 High Speed Physical Interoperability Recommendations	28
4.1 Connector Type	28

4.2	Connector Pin Assignment	28
	4.2.1 Board Height and Spacing Requirements	32
4.3	Lane Assignment and Connections	32
4.4	Connectors and Cable Assembly Skew Budget	33
4.5	Connectors and Cable Assembly Insertion Loss.....	35
4.6	AC Coupling	36
4.7	Signal Integrity Considerations and PCB Items	36

Revision History

<p style="text-align: center;">Revision 1.9 Revision Date: 19 May 2016</p>
<ul style="list-style-type: none">• Connector pin assignment update in Figure 9: Changed TX to RX (pin locations: J8, J9)
<p style="text-align: center;">Revision 1.8 Revision Date: 26 March 2015</p>
<ul style="list-style-type: none">• Added High Speed (> 12.5G) Physical Interoperability Recommendations
<p style="text-align: center;">Revision 1.7 Revision Date: 4 March 2013</p>
<ul style="list-style-type: none">• Added 400G target applications and 100G version for 25G and 12.5G SerDes
<p style="text-align: center;">Revision 1.6 Revision Date: 11 October 2011</p>
<ul style="list-style-type: none">• Added 120G target applications• Added 16L Lmax-xoff entry in Table 1• Added skew budget sections 3.5-3.7
<p style="text-align: center;">Revision 1.5 Revision Date: 8 December 2010</p>
<ul style="list-style-type: none">• Added retransmission information to section 2.1 and 2.2.
<p style="text-align: center;">Revision 1.4 Revision Date: 04 June 2010</p>
<ul style="list-style-type: none">• Added Section 3.3.1, Board Height Requirements on page 21
<p style="text-align: center;">Revision 1.3 Revision Date: 05 January 2009</p>
<ul style="list-style-type: none">• Added 10.3125 Gb/s SerDes rate options• Added references to OIF's electrical specifications• Added legal disclaimer page
<p style="text-align: center;">Revision 1.2 Revision Date: 10 September 2008</p>
<ul style="list-style-type: none">• Added section to define Lmax-xoff.
<p style="text-align: center;">Revision 1.1 Revision Date: 11 June 2008</p>
<ul style="list-style-type: none">• Changed naming from "40G", "20G" interfaces to "50G", "25G" interfaces. The "50G" interface has exactly half the throughput of a "100G" interface.• Figure added showing interoperability with 10G narrow and 10G wide interfaces• Added Low Speed Physical Interoperability Recommendations section
<p style="text-align: center;">Revision 1.0 Revision Date: 31 October 2007</p>
<ul style="list-style-type: none">• First release.

1 Overview

The Interlaken protocol specification addresses a wide range of chip-to-chip packet transfer needs. It is scalable in bandwidth and has options to optimize the interface towards specific requirements as, for example, low latency burst transfers, unidirectional transfers, etc.

In order to facilitate interoperability between devices from different vendors, the Interlaken Alliance has defined a set of recommended baseline configurations. These recommendations target typical applications at data transfer rates of 10 Gb/s, 25 Gb/s, 50 Gb/s, 100 Gb/s and 400 Gb/s.

Physical interoperability recommendations are also included to facilitate interoperability between reference boards and hardware platforms.

2 Logical Interoperability Recommendations

2.1 Definitions of Interoperability Parameters

The parameters specified in the next sections are defined in the table below.

Property	Definition	Comments
Number of lanes	One lane means one differential transmit signal pair and one differential receive pair.	
Lane bit rate	Transceiver bit rate. All transceivers in a logical interface are frequency-locked and run at the same bit rate.	10.3125 Gb/s, 6.25 Gb/s and 3.125 Gb/s were selected to match common industry rates.
Backpressure method	Method to signal flow control on a specific logical channel or for the interface link as a whole.	Backpressure signals can either use separate pins, out of band, or be carried in bits in the data path control word, in band.
Packet transfer method	Packets can either be transferred one full packet at a time or as segments, where partial packets on different channels are interleaved.	The two modes are called packet mode and segment mode.
Stop boundary	Where a packet transfer stops as a response to backpressure.	Transmission can either stop at packet boundary, which require larger buffers, or stop at any burst control word within a packet.
BurstMax / BurstMin / BurstShort	<i>BurstMax</i> – maximum size of a data burst <i>BurstMin</i> – smallest size of an end-of-packet burst <i>BurstShort</i> – minimum interval between burst control words	
MetaFrameLength	The quantity of data sent on each lane including one Synchronization Word, one Scrambler State Word, one Diagnostic Word, one or more Skip Words, and the data payload	
Multiple use field	Field is part of the burst control word as defined in the Interlaken protocol specification.	Typical use of this field is to extend the channel ID field or to extend the flow control field.
Rate matching	Granularity of interface rate shaper.	This shaper prevents the transmitter from overflowing the receiver.
Retransmission	Interlaken supports an optional retransmission extension which allows bit errors to be corrected by retransmitting data.	
Comments	Additional comments and observations outside the interoperability recommendation.	

2.2 Recommendations for All Transfer Rates

Property	Recommendation
Backpressure method	In-band
Channel count	Not specified, application dependent
Packet transfer method	Not specified, application dependent
Packet Mode Stop Boundary	For link level backpressure: Burst end For channel backpressure: Packet end
Burst Mode Stop Boundary	Burst
BurstMax / BurstMin / BurstShort	256 bytes / 64 bytes / 32 bytes
MetaFrameLength	2,048 words
Multiple use field	Not used
Rate matching	Yes, 1 Gb/s steps
Status Messaging	Not required
Retransmission	Optional, default is disabled
Comments	For unidirectional data paths, out-of-band backpressure may be useful. LVCMOS electrical levels are then preferred.

2.3 10 Gb/s Packet Transfers

2.3.1 Target Applications

The recommendations were created to provide sufficient performance for typical implementations of these applications:

Application	Frame Sizes
1 port 10 Gigabit Ethernet	64 bytes and up
12 ports Gigabit Ethernet	64 bytes and up
OC192 / STM64 POS framer	40 bytes and up

2.3.2 Recommendations

There are two different recommendations at the 10 Gb/s transfer rate to take advantage of widely used serial transceiver rates. For interoperability, an implementation would need to specify if it adheres to the narrow or wide interface option.

2.3.2.1 Recommendation #1 – Narrow Interface Option

Property	Recommendation
Number of lanes	3
Lane bit rate	6.25 Gb/s
Electrical Specification	CEI-6G-SR
Comments	10 GE can be optimized into just 2 lanes

2.3.2.2 Recommendation #2 – Wide Interface Option

Property	Recommendation
Number of lanes	5
Lane bit rate	3.125 Gb/s

2.4 25 Gb/s Packet Transfers

2.4.1 Target Applications

The recommendations were created to provide sufficient performance for typical implementations of these applications:

Application	Frame Sizes
2 ports 10 Gigabit Ethernet	64 bytes and up
25 ports Gigabit Ethernet	64 bytes and up
2x OC192 POS framer	40 bytes and up

2.4.2 Recommendations

Property	Recommendation
Number of lanes	5
Lane bit rate	6.25 Gb/s
Electrical Specification	CEI-6G-SR

2.5 50 Gb/s Packet Transfers

2.5.1 Target Applications

The recommendations were created to provide sufficient performance for typical implementations of these applications:

Application	Frame Sizes
5 ports 10 Gigabit Ethernet	64 bytes and up
50 ports Gigabit Ethernet	64 bytes and up
OC768 / STM256 POS framer	40 bytes and up
4x OC192 POS framer	40 bytes and up

2.5.2 Recommendations

There are two different recommendations at the 50 Gb/s transfer rate to take advantage of widely used serial transceiver rates. For interoperability, an implementation would need to specify if it adheres to the narrow or wide interface option.

2.5.2.1 Recommendation #1 – Wide Interface Option

Property	Recommendation
Number of lanes	10
Lane bit rate	6.25 Gb/s
Electrical Specification	CEI-6G-SR

2.5.2.2 Recommendation #2 – Narrow Interface Option

Property	Recommendation
Number of lanes	6
Lane bit rate	10.3125 Gb/s
Electrical Specification	CEI-11G-SR
Comment	It is recommended to allow double the skew tolerance (214 UI) for this rate

2.6 100 Gb/s Packet Transfers

2.6.1 Target Applications

The recommendations were created to provide sufficient performance for typical implementations of these applications:

Application	Frame Sizes
10 ports 10 Gigabit Ethernet	64 bytes and up
100 ports Gigabit Ethernet	64 bytes and up
2x OC768 POS framer	40 bytes and up
8x OC192 POS framer	40 bytes and up

2.6.2 Recommendations

There are four different recommendations at the 100 Gb/s transfer rate to take advantage of widely used serial transceiver rates. For interoperability, an implementation would need to specify which option it adheres to.

2.6.2.1 Recommendation #1 – 6G Interface Option

Property	Recommendation
Number of lanes	20
Lane bit rate	6.25 Gb/s
Electrical Specification	CEI-6G-SR

2.6.2.2 Recommendation #2 – 10G Interface Option

Property	Recommendation
Number of lanes	12
Lane bit rate	10.3125 Gb/s
Electrical Specification	CEI-11G-SR
Comment	It is recommended to allow double the skew tolerance (214 UI) for this rate

2.6.2.3 Recommendation #3 – 12.5G Interface Option

Property	Recommendation
Number of lanes	10
Lane bit rate	12.5 Gb/s
Electrical Specification	CEI-11G-SR+

Comment	It is recommended to allow the skew tolerance of 214 UI for this rate
---------	-----------------------------------------------------------------------

2.6.2.4 Recommendation #4 – 25G Interface Option

Property	Recommendation
Number of lanes	5
Lane bit rate	25.78125 Gb/s
Electrical Specification	CEI-28G-SR/MR
Comment	It is recommended to allow the skew tolerance of 214 UI for this rate

2.7 120 Gb/s Packet Transfers

2.7.1 Target Applications

The recommendations were created to provide sufficient performance for typical implementations of these applications:

Application	Frame Sizes
12 ports 10 Gigabit Ethernet	64 bytes and up
3 ports 40 Gigabit Ethernet	64 bytes and up
OTU4 or 3x OTU3/OTU3e	40 bytes and up
3x OC-768 POS Framer	40 bytes and up
12x OC-192 POS Framer	40 bytes and up

2.7.2 Recommendation #1

Property	Recommendation
Number of lanes	16
Lane bit rate	10.3125 Gb/s
Electrical Specification	CEI-11G-SR
Comment	It is recommended to allow double the skew tolerance (214 UI) for this rate

2.8 400 Gb/s Packet Transfers

2.8.1 Target Applications

The recommendations were created to provide sufficient performance for typical implementations of these applications:

Application	Frame Sizes
4 ports of 100Gigabit Ethernet	64 bytes and up
1 Port of 400Gigabit Ethernet	64 bytes and up

2.8.2 Recommendation #1

Property	Recommendation
Number of lanes	40
Lane bit rate	12.5 Gb/s
Electrical Specification	CEI-11G-SR+
Comment	It is recommended to allow the skew tolerance of 214 UI for this rate

2.8.3 Recommendation #2

Property	Recommendation
Number of lanes	20
Lane bit rate	25.78125 Gb/s
Electrical Specification	CEI-28G-SR/MR
Comment	It is recommended to allow the skew tolerance of 214 UI for this rate

2.9 Summary of Lane Recommendations

The following table has a summary of the lane recommendations.

Transfer Rate	Recommendations		
	Lanes	Rate [Gb/s]	Electrical Spec
10 Gb/s	5	3.125	
	3	6.25	CEI-6G-SR
25 Gb/s	5	6.25	CEI-6G-SR
50 Gb/s	10	6.25	CEI-6G-SR
	6	10.3125	CEI-11G-SR
100 Gb/s	20	6.25	CEI-6G-SR
	12	10.3125	CEI-11G-SR
	10	12.5	CEI-11G-SR+
	5	25.78125	CEI-28G-SR/MR
120 Gb/s	16	10.3125	CEI-11G-SR
400 Gb/s	40	12.5	CEI-28G-SR/MR
	20	25.78125	CEI-28G-SR/MR

2.10 Multirate Interoperability

The per-rate recommendations are for interoperability between components of the same capacity; e.g. a 50 Gb/s framer connected to a 50 Gb/s packet processor. Additionally, there is also a possibility to interoperate with multiple lower-capacity components if the lower speed recommendations are also implemented.




Figure 1: Example of Interoperability with Lower Capacity Components

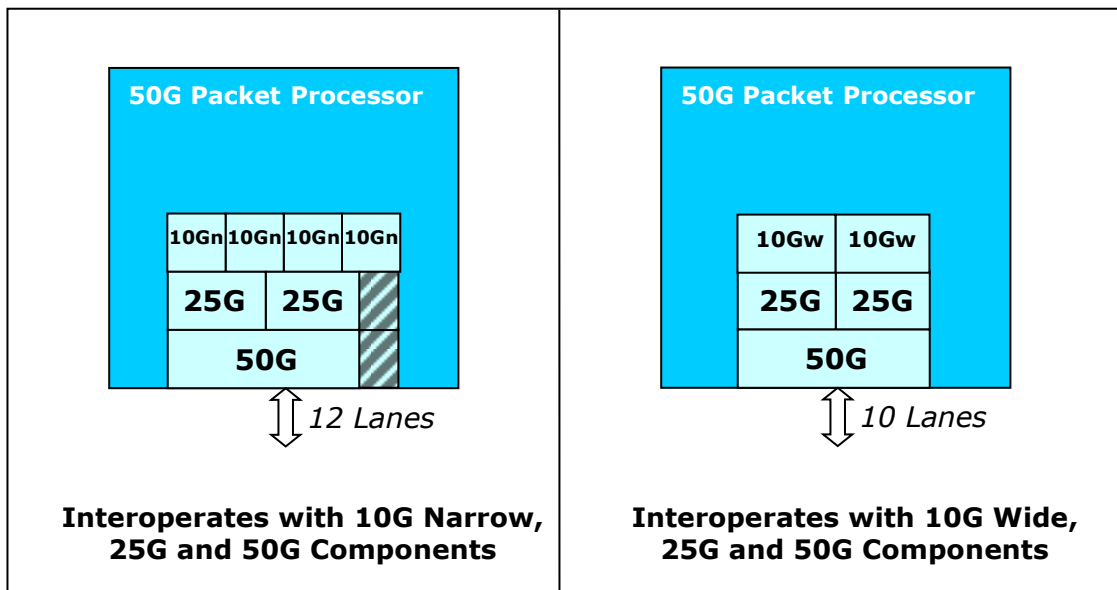


Figure 2: Example of Interoperability with 10G Components

2.11 Latency Requirements

In order to improve interoperability and to avoid over-design of buffers in a receiving device, the interoperability guideline defines the maximum time a transmitting device should respond after it receives XOFF information. This parameter is termed $L_{max-xoff}$ and is defined as the maximum number of bytes the transmitter is allowed to transmit after XOFF information appears at the input pins of the transmitting device.

Figure 3 shows a high level block diagram of a typical transmitter. The transmitter must ensure no more than Lmax-xoff bytes go through point D after XOFF is seen at point A. The various contributors to Lmax-xoff include the delay through the Interlaken RX block (point A->B), the time it takes for the TX Scheduler to stop transmitting data (point B->C), and how much data is already in flight through the Interlaken TX block (point C->D).

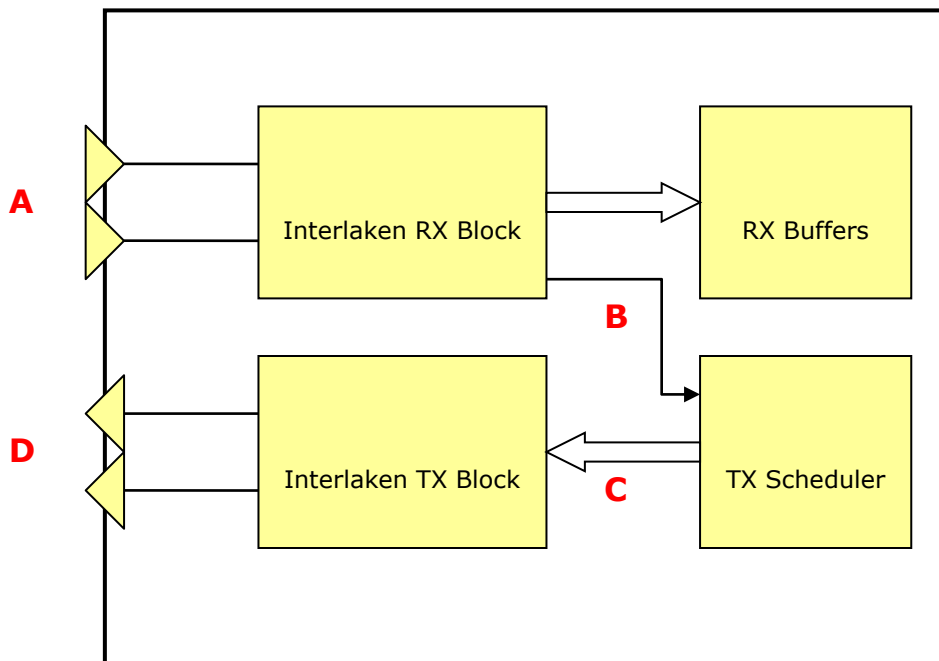


Figure 3 Transmitter Block Diagram

Table 1 defines Lmax-xoff for different interface configurations. The numbers apply to burst mode operation. For packet mode, one MTU must be added to the Lmax-xoff values. Additionally, the numbers are based on in-band flow control, since out-of-band flow control is typically faster and so the numbers would be less.

The numbers are independent from the number of channels, though valid for a small number of channels.

Profile	Number of Lanes	Serial Rate	Lmax-xoff
		(Gb/s)	(Bytes)
10G	5	3.125	1,536
	3	6.25	1,280
25G	5	6.25	1,792
50G	6	10.3125	6,400
	10	6.25	6,400
100G	20	6.25	12,484
	12	10.3125	12,484

	10	12.5	12,484
	5	25.78125	12,484
120G	16	10.3125	16,645
400G	40	12.5	TBD
	20	25.78125	TBD

Table 1 Lmax-xoff Definitions

3 Low Speed Physical Interoperability Recommendations

In order to facilitate the physical interoperability between devices, the Interlaken Alliance has defined a set of recommended connector types, connector pin assignment, and connector board placements. The connector chosen is for interoperability testing up to 12.5G. Higher speed interoperability should use the high speed connector. Please refer to the High Speed Physical Interoperability Recommendations section.

3.1 Connector Type

The recommended common connector is FCI AirMax VS[®] High Speed 4-Pair 120-pos 3-mm 10 IMLA Right Angle Header / Receptacle connectors for coplanar connections.

The part numbers for the connectors are 10035515 (header) and 10045722 (receptacle).

The header connector transmits data across the interface while the receptacle connector receives data from the interface.

3.2 Connector Pin Assignment

A common pin assignment for the connectors is shown in the following diagrams. Figure 4 below shows the signal names used in the pin assignment. Figure 5 below shows the pin assignment of the differential pairs for the receiver. Since the pin labeling of the header and receptacle connectors are mirror image of each other, the transmit and receive differential pairs have the same pin assignment (e.g. TX0_P at A7 of header will mate with RX0_P at A7 of receptacle).

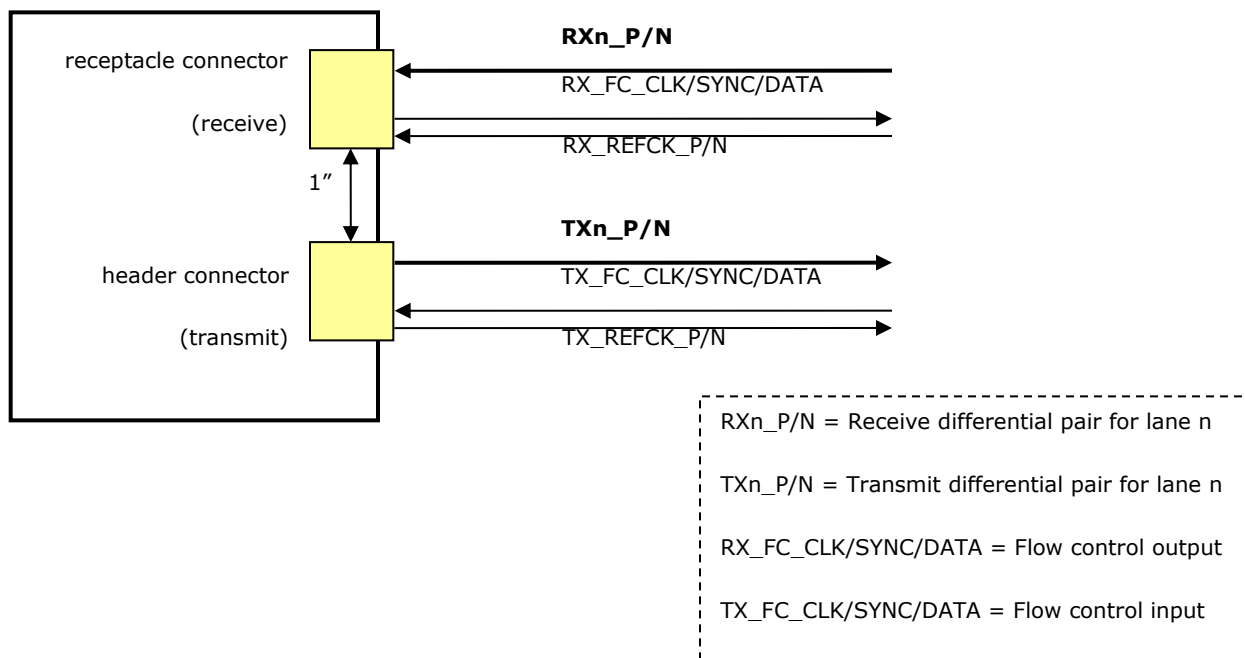


Figure 4: Signal Names for Connector Pin Assignment

	10	9	8	7	6	5	4	3	2	1
A	GND	RX3_P	GND	RX0_P	GND	RX7_P	GND	RX4_P	GND	RX_REFCK_P
B	GND	RX3_N	GND	RX0_N	GND	RX7_N	GND	RX4_N	GND	RX_REFCK_N
C	GND	GND	GND	GND	GND	GND	GND	GND	GND	GND
D	GND	GND	RX2_P	GND	RX1_P	GND	RX6_P	GND	RX5_P	GND
E	RX_FC_CK	GND	RX2_N	GND	RX1_N	GND	RX6_N	GND	RX5_N	GND
F	GND	GND	GND	GND	GND	GND	GND	GND	GND	GND
G	GND	GND	GND	GND	GND	RX8_P	GND	RX9_P	GND	RX11_P
H	GND	RX_FC_SYNC	GND	RX_FC_DATA	GND	RX8_N	GND	RX9_N	GND	RX11_N
I	GND	GND	GND	GND	GND	GND	GND	GND	GND	GND
J	RX13_P	GND	RX15_P	GND	RX12_P	GND	RX10_P	GND	RX14_P	GND
K	RX13_N	GND	RX15_N	GND	RX12_N	GND	RX10_N	GND	RX14_N	GND
L	GND	GND	GND	GND	GND	GND	GND	GND	GND	GND

Figure 5: Connector Pin Assignment for Differential Pairs (receiver shown)

The lanes needed to support a certain data transfer rate shall occupy the lower-numbered differential pairs (e.g. 25 Gb/s interface shall use RX0_P/N through RX4_P/N).

The upper-numbered pairs are treated as spare pins for user-defined purposes. If unused, these pins shall be left unconnected.

The flow control signals (FC_CK, FC_SYNC and FC_DATA) are used only when out-of-band backpressure is supported. If unused, these pins shall be left unconnected.

The REFCK_P/N pair can be optionally used to pass a reference clock across the interface. If unused, these pins shall be left unconnected.

The 100 Gb/s interface shall employ 2 sets of connectors with the same pin assignment for 50 Gb/s interface. The placement of these connectors is described in Section 0 below.

3.3 Connector Placement

For all interface rates, the header and receptacle connectors shall be placed 1" apart (1155 mil between pins A10 of 2 connectors) as shown in Figure 6 below. When viewed from the top with the connectors along the right edge of the board, the header shall be placed below the receptacle.

For 100 Gb/s interface, two sets of connectors are needed. The two connector sets shall be placed 5" apart as shown below. For 100 Gb/s interface to interoperate with 50 Gb/s (or lower) interface, channel mapping will be needed on the 100 Gb/s interface, as shown in Section 3.4.8 on page 25.

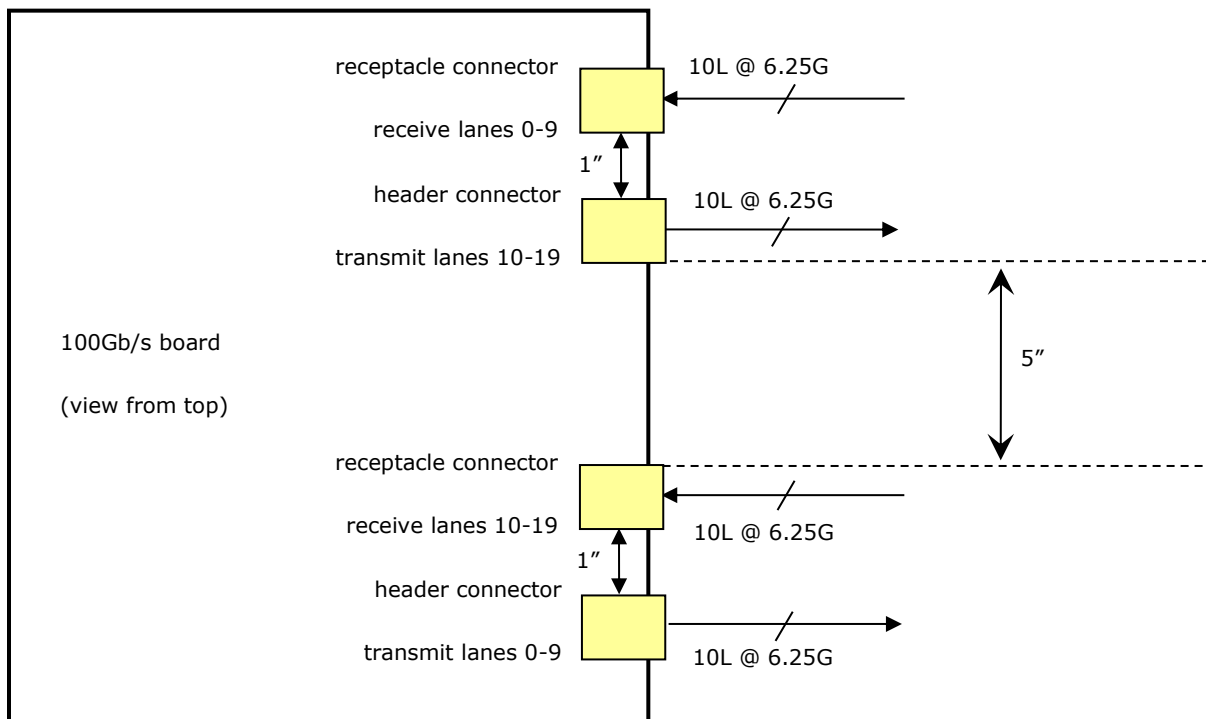


Figure 6: Connector Placement for 100 Gb/s Interface

3.3.1 Board Height Requirements

For ease of interoperability testing, a standard board height of 1.37 inches to the top of the board is recommended as shown in Figure below.





Figure 7: Standard Interlaken Board Height


3.4 Physical Interoperability Examples

The following sections show examples of how two evaluation boards can be connected together based on the recommendations above.


3.4.1 Connecting 10 Gb/s to 10 Gb/s



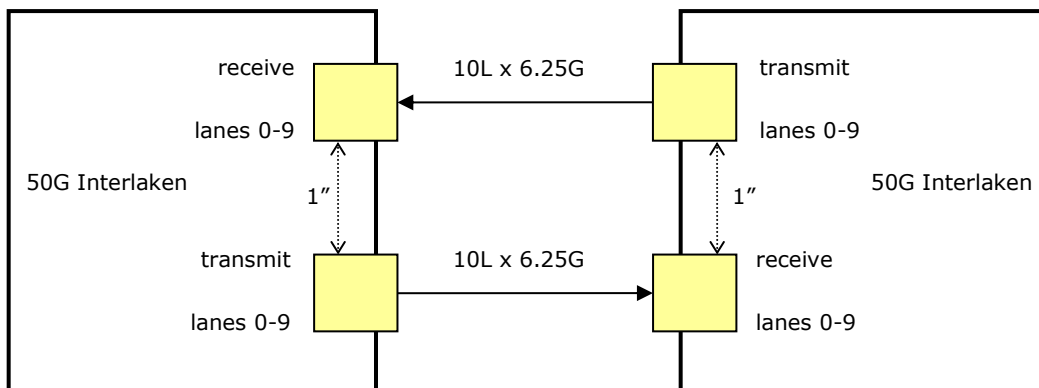
3.4.2 Connecting 25 Gb/s to 25 Gb/s




3.4.3 Connecting 25 Gb/s to 10 Gb/s




3.4.4 Connecting 50 Gb/s to 50 Gb/s



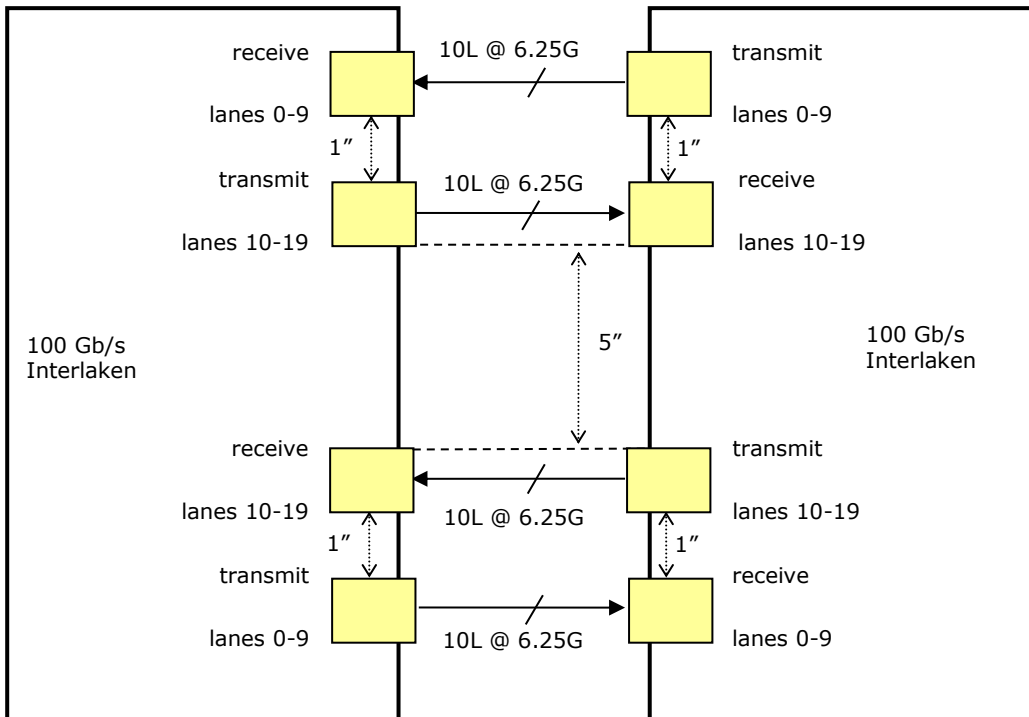
3.4.5 Connecting 50 Gb/s to 25 Gb/s



3.4.6 Connecting 50 Gb/s to 10 Gb/s



3.4.7 Connecting 100 Gb/s to 100 Gb/s

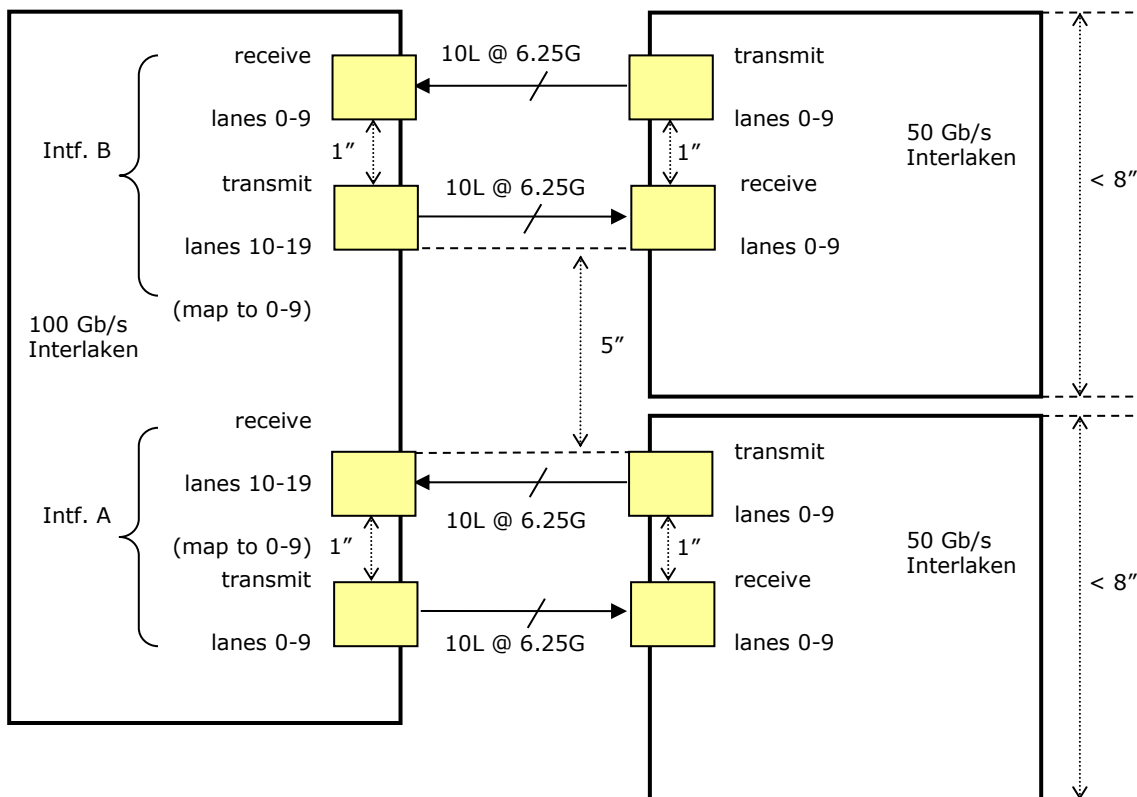


3.4.8 Connecting 100 Gb/s to 50 Gb/s

For 100 Gb/s to interop with lower rate interfaces, channel mapping is required as shown below:

- Map RX lanes 10-19 to RX lanes 0-9 of Interlaken Interface A
- Map TX lanes 10-19 to TX lanes 0-9 of Interlaken Interface B

Also, for connecting 2 evaluations boards of lower interface rate, there should be sufficient space for the boards to be connected side-by-side. Each board should not be more than 8" wide.



3.4.9 Connecting to Multiple Interlaken Interfaces

As shown in Section 0, interoperability with multiple lower-capacity components is also possible. This section describes the general recommendations of how this can be achieved. However, not all combinations can be covered by these recommendations. As such, for more complex interoperability cases, they will serve only as a set of guidelines.

3.4.9.1 Lanes Assignment

For connection between a higher-capacity component with multiple lower-capacity components (e.g. 50 Gbps component connecting to two 25 Gbps components), these recommendations apply:

- Components supporting multiple Interlaken interfaces of rates of 50 Gbps and lower should use consecutive lanes starting with lane 0 in the following order:
 1. for all 50G interfaces (if any)
 2. for all 25G interfaces (if any)
 3. for all 10G-wide interfaces (if any)

4. for all 10G-narrow interfaces (if any)

These lanes should use only 1 set of connectors. If necessary, the spare lanes can be used, giving a maximum of 16 lanes per connector.

For example, a 50G component that also supports interoperability with four 10G-narrow components, uses 12 lanes (lanes 0-2, 3-5, 6-8 and 9-11) all from the same set of connectors. In this case, spare lanes 10 and 11 are used.

- Components supporting 100 Gbps shall use 2 sets of connectors as shown in Figure 5 above. Each set of connectors should be treated as a separate interface, connecting to a single component, or connecting to multiple lower-rate components.

For example, a 100G component that also supports 2 x 25G interfaces can either route each 25G interface to lanes 0-4 of each of the 2 connector sets, or route both interfaces to 1 connector set, using lanes 0-4 and 5-9.

- Lane mapping should be used to support multirate interoperability.
- Since in-band flow control is part of the interoperability recommendations for all rates, out-of-band flow control need not be supported.
- For any cases where extra lanes are required, the spare lanes should be used.

3.4.9.2 More Multirate Interoperability Examples

- 50G connecting to 2 x 25G: lanes 0-4 connect to one 25G component while lanes 5-9 connect to the other.
- 50G connecting to 4 x 10G-narrow: lanes 0-2, 3-5, 6-8, 9-11 connect to each of the 10G-narrow components. In this example, lanes 10 and 11 are spare lanes.
- 50G connecting to 1 x 25G and 2 x 10G-narrow: lanes 0-4 connect to the 25G component, while lanes 5-7, 8-10 connect to each of the 10G-narrow components. Lane 11 is a spare lane.
- 100G connecting to 2 x 50G: this is shown in Section 3.4.8 on page 25.
- 100G connecting to 4 x 10G-narrow: lanes 0-2, 3-5, 6-8, 9-11 (from one set of connectors) connect to each of the 10G-narrow components. Lanes 10 and 11 are spare lanes. Alternatively, 4 x 10G-narrow components can use 2 sets of connectors. Two of the 10G-narrow interfaces can go through 1 set of connectors using lanes 0-2 and 3-5, while the other two 10G-narrow interfaces go to the 2nd set of connectors.

3.5 PCB Intra-Lane Skew Budget

Skew between P/N signals in each differential pair must be managed for each transmit and receive lane and include compensation for skew within the connector. Maximum skew between P/N shall be 3ps after subtracting connector delay compensation shown in the table below from the P lane (e.g. If RX0_P is 710ps then RX0_N should be 700ps +/-3ps or 697 to 703 ps). The connector skew compensation should be located as close to the connector as possible.

Lane	Delay(ps)
0, 3, 4, 7	10
1, 2, 5, 6	2.5
8, 9, 11	8
10, 12, 13, 14, 15	7.5

3.6 AC Coupling

Each Lane connected to the receive receptacle connector should include AC coupling capacitors on each lane using a value of 100nF.

3.7 Signal Integrity Considerations

Lanes 0, 3, 4 and 7 have higher insertion loss through the connector. For interface rates higher than 6.5 Gb/s these lanes should be closely examined to minimize any additional PCB losses. Suggested layout options include:

- Backdrill these lanes at the connector and use the lowest available routing layer at the connector via to minimize via stubs
- Include an outer GND via for the P signal to reference on row As

4 High Speed Physical Interoperability Recommendations

In order to facilitate the physical interoperability between devices, the Interlaken Alliance has defined a set of recommended items: a connector type, a connector pin assignment, and a cable assembly to simplify board connector placement and interop usage. The connector and cable assembly chosen are for interoperability testing greater than 12.5G.

4.1 Connector Type

The recommended common connector is a TE Connectivity 3.9 mm STRADA Whisper connector. It is a connector that can be used for cabled connection.

The part number for the connector is 2187194-1. The base part number for the cable assembly is 2274955 with a -x extension, where x refers to the length of the cable assembly.

The connector handles both transmit and receive signals; the cable assembly crosses the signals to provide connectivity from transmit to receive and receive to transmit. This cabled connector assembly provides the necessary scalability to 40 lanes of connectivity at 25.7125 Gbps and beyond simply by having an additional connector placed on the board. Figure 8 shows a pictorial representation of how a single connector or two connectors can be setup and demonstrates its respective capability.

Component	Part Number
3.9 mm STRADA Whisper Connector (R/A Receptacle)	2187194-1
20 inch cable assembly	2274955-1
9 inch cable assembly	2274955-2

Table 2: Part Numbers for Interoperation Connector and Cable Assembly

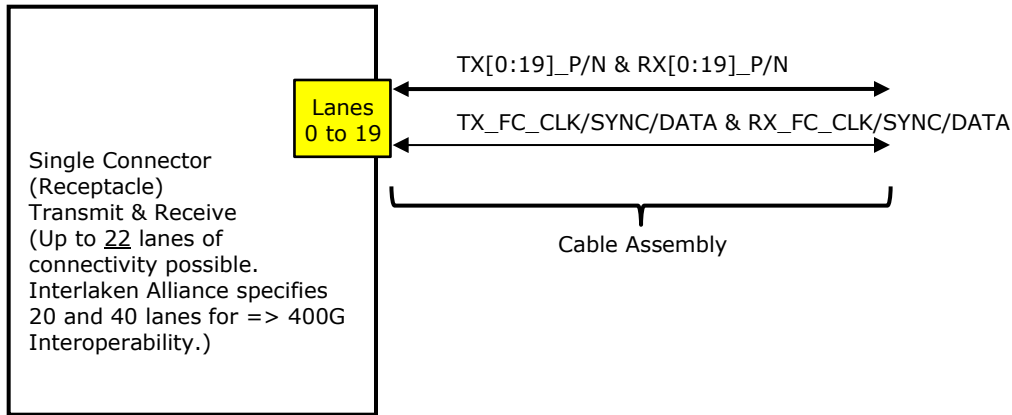
4.2 Connector Pin Assignment

Figure 8 shows the two options on how your board can be setup for interoperability. The top option is a single connector which allows up to 22 lanes of high-speed differential connectivity to another board for interoperability. Note, the Interlaken Alliance currently specifies only 20 lanes are needed for interoperability. The bottom option consists of two connectors which allows up to 44 lanes of high-speed differential connectivity to another board. In this case, the Interlaken Alliance currently specifies the need for only 40 lanes for interoperability. The connector and cable assembly are rated to support data rates up to 30 Gbps with non-return to zero (NRZ) encoding.

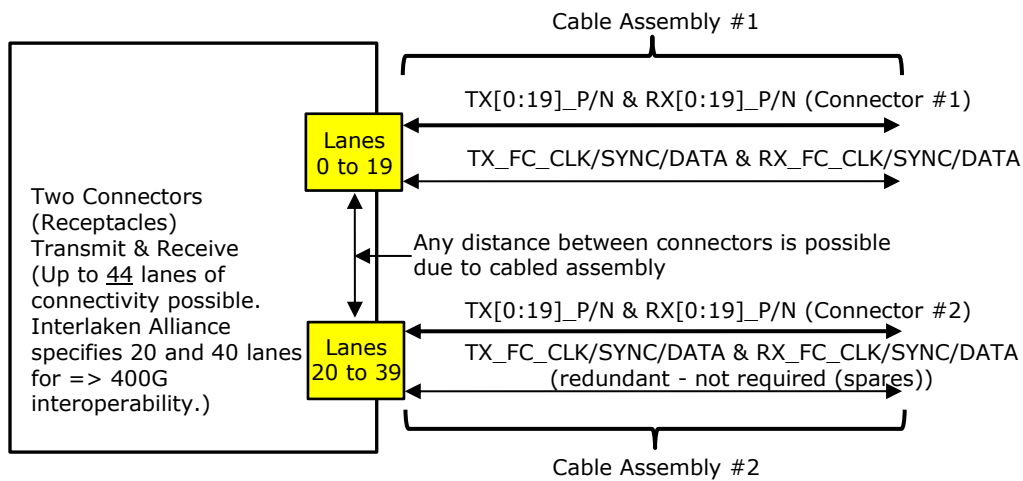
Note: TX[20:21] and RX[20:21] are spare lanes that can be used for other cases when extra lanes may be required.

Figure 9 shows the pin assignment of the high-speed differential pairs and out-of-band flow control signals for the connector. Signals labeled TX are outputs from the board; RX is an input to the board. The cable assembly that connects each PCB's on-board connector(s) (STRADA Whisper Connectors – R/A Receptacles) takes care of swapping TX to RX, therefore the connector pinout is identical on either interoperating board. Note: TX_REFCK_P/N and RX_REFCK_PN are not provided for high speed physical interoperability.

Single Connector Interoperation Setup



Two Connector Interoperation Setup



TX_n_P/N = Transmit differential pair for lane n
 RX_n_P/N = Receive differential pair for lane n
 TX_FC_CLK/SYNC/DATA = Flow control input
 RX_FC_CLK/SYNC/DATA = Flow control output

Figure 8: Signal Names for Connector Interoperability Options

Red line denotes PCB edge

	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1
A	GND	TX20_N	TX20_P	GND	TX18_N	TX18_P	GND	TX16_N	TX16_P	GND	TX4_N	TX4_P	GND	TX2_N	TX2_P	GND	TX0_N	TX0_P	GND
B	GND	RX20_N	RX20_P	GND	RX18_N	RX18_P	GND	RX16_N	RX16_P	GND	RX4_N	RX4_P	GND	RX2_N	RX2_P	GND	RX0_N	RX0_P	GND
C	GND	TX21_N	TX21_P	GND	TX19_N	TX19_P	GND	TX17_N	TX17_P	GND	TX5_N	TX5_P	GND	TX3_N	TX3_P	GND	TX1_N	TX1_P	GND
D	GND	RX21_N	RX21_P	GND	RX19_N	RX19_P	GND	RX17_N	RX17_P	GND	RX5_N	RX5_P	GND	RX3_N	RX3_P	GND	RX1_N	RX1_P	GND
E	UNPOPULATED ROW (No Connect)																		
F	GND	GND	RX_FC_CLK	GND	TX14_N	TX14_P	GND	TX12_N	TX12_P	GND	TX10_N	TX10_P	GND	TX8_N	TX8_P	GND	TX6_N	TX6_P	GND
G	GND	GND	TX_FC_CLK	GND	RX14_N	RX14_P	GND	RX12_N	RX12_P	GND	RX10_N	RX10_P	GND	RX8_N	RX8_P	GND	RX6_N	RX6_P	GND
H	GND	RX_FC_DATA	TX_FC_DATA	GND	TX15_N	TX15_P	GND	TX13_N	TX13_P	GND	TX11_N	TX11_P	GND	TX9_N	TX9_P	GND	TX7_N	TX7_P	GND
J	GND	RX_FC_SYNC	TX_FC_SYNC	GND	RX15_N	RX15_P	GND	RX13_N	RX13_P	GND	RX11_N	RX11_P	GND	RX9_N	RX9_P	GND	RX7_N	RX7_P	GND

Note: TX[20:21] and RX[20:21] are spare lanes that can be used for other cases when extra lanes may be required.

Figure 9: Connector Pin Assignment for Differential Pairs and Out-of-Band Flow Control Signals (Bottom View of Connector)

Figure 10 displays a sample routing configuration and provides fixtures orientations to help with proper connector placement and ideas on how to limit board layer count.

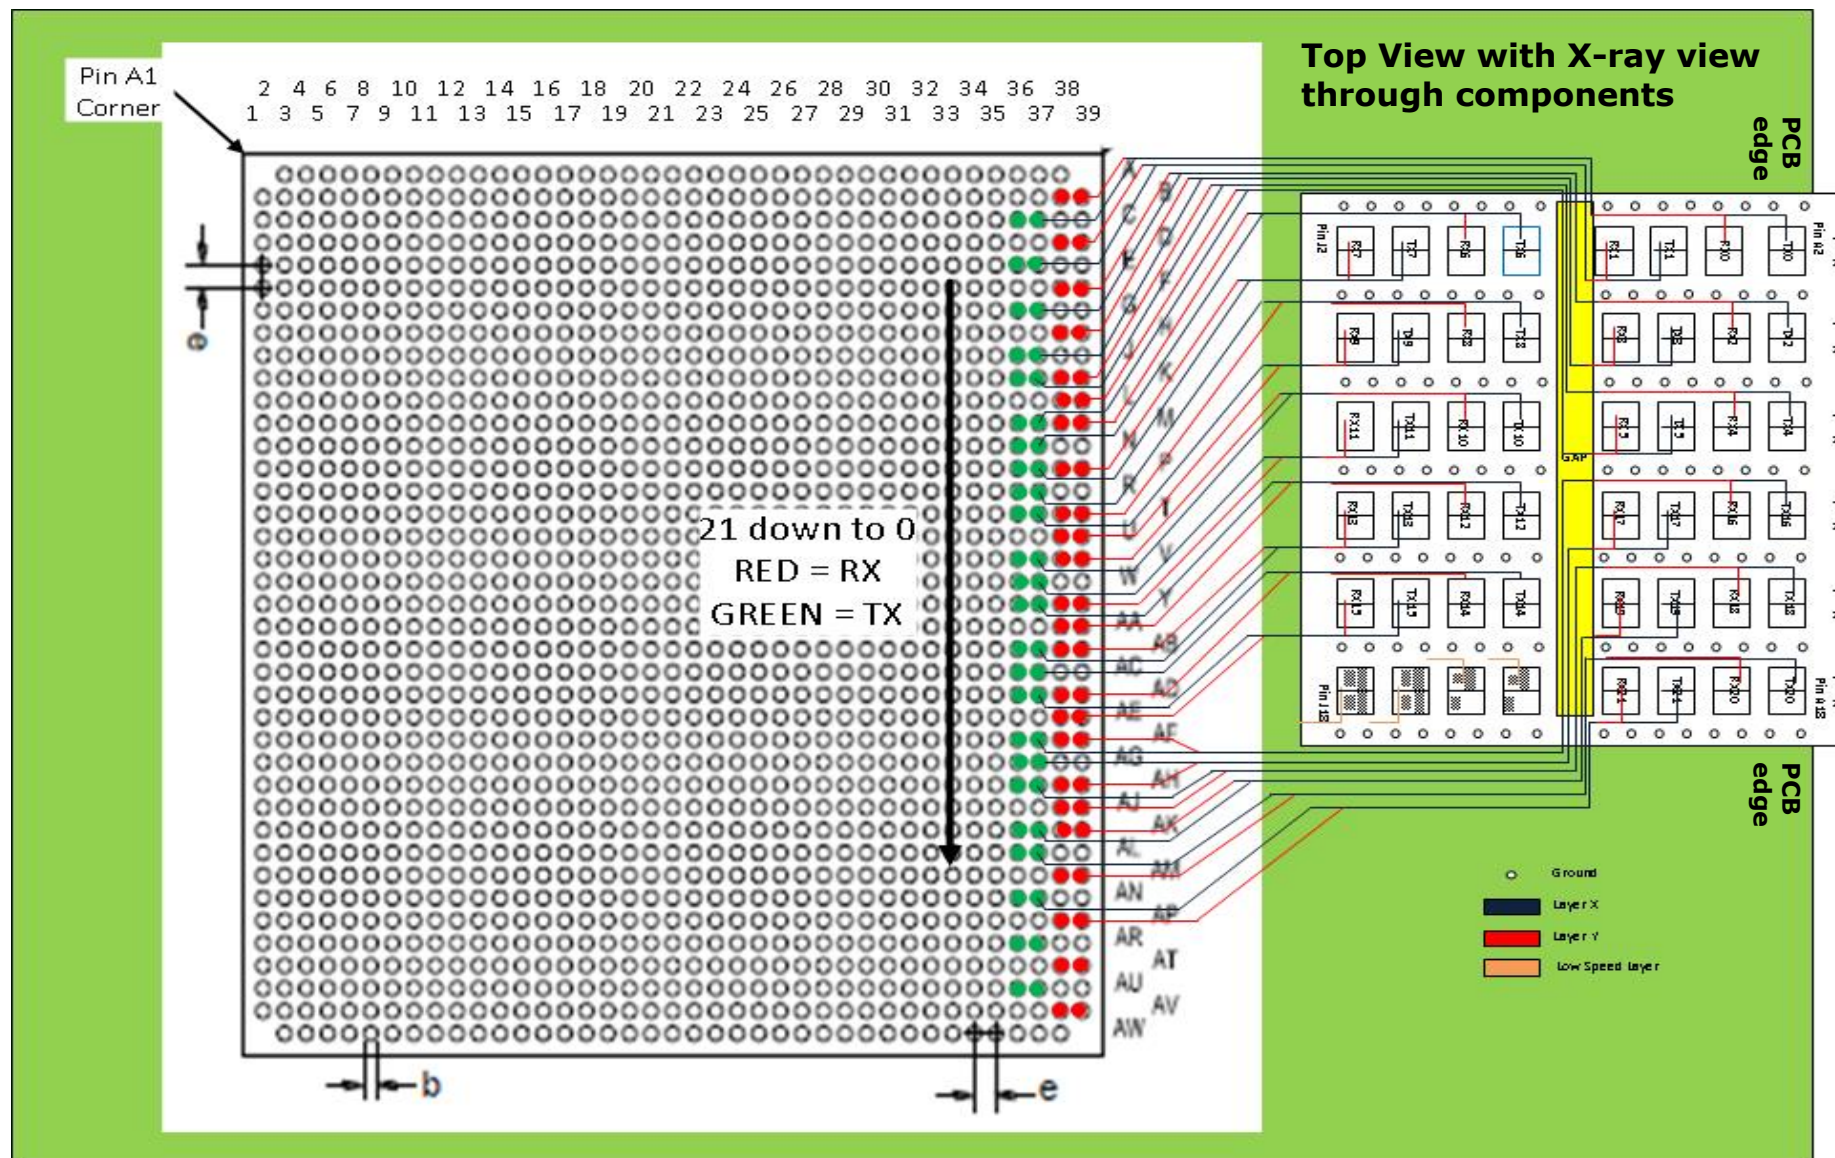


Figure 10: Sample Routing Configuration

4.2.1 Board Height and Spacing Requirements

For interoperability testing, the board height should be greater than or equal to 7 mm – from a flat surface to the top of the board. An example is shown in Figure 11, where the height is greater than 7mm (.28 inches). For the example case shown, the board height is 1.37 inches. Also, note the spacing between the two connectors is 1.85 inches. The spacing between a two-connector solution (using up to 40 Interlaken lanes) is up to the board developer as other system or board constraints may play a factor. The connector & cable assembly solution offers dimensional flexibility to board designers that alleviate such constraints.




Figure 11: Interlaken Board Height

4.3 Lane Assignment and Connections

For a given interoperability, use lane 0 to lane x where x is the highest lane needed to support a given interoperability bandwidth (number of lanes x data rate).

Interlaken Alliance currently specifies 20 and 40 lanes of connectivity for greater than 400G interoperability. As shown on the pinout, two additional differential pairs (TX[20:21] and RX[20:21]) are spare lanes that can be used for any other cases where extra lanes are required.

Since in-band flow control is part of the interoperability recommendations for all rates, out-of-band flow control is an optional portion of the interoperability and signals are provided via the connector pinout if such interoperability is required between vendors.

For out-of-band flow control, the I/O pins (TX_FC_CLK, RX_FC_CLK, TX_FC_SYNC, RX_FC_SYNC, TX_FC_DATA, RX_FC_DATA) must be compatible to +1.8 V LVCMOS electrical levels to ensure compatibility between devices. These signals are only used when out-of-band backpressure is required to be supported. If unused, these pins shall be left unconnected.

4.4 Connectors and Cable Assembly Skew Budget

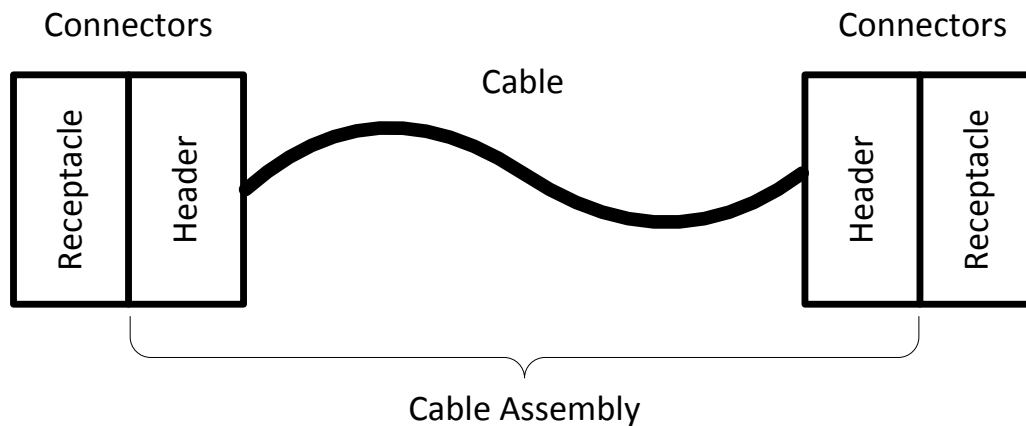


Figure 12: Connector & Cable Assembly Channel

RECEPTACLE and HEADER: The skew between the P/N signals for each differential pair and the skew between the differential pairs in regards to the STRADA Whisper connector are as follows:

- The maximum skew between any two rows of high-speed differential pairs for the connector is 189 ps (Row A to Row H = Row A to Row B + Row B to Row C + Row C to Row D + Row D to Row F + Row F to Row G + Row G to Row H + Row H to Row J). Note the goal is to match TX data signals and RX data signals independently – worst case skew amongst all TX data paths grouped and all RX data paths grouped. Table 3 outlines the various skews between the nine rows of the connector to the header portion of the cable assembly (see Figure 12 for reference). Recall, the center row of the connector is unpopulated and has no physical pins for connectivity to the printed circuit board.
 - The skews between connecting data path pairs across the channel: TX to RX and RX to TX cancel one another as the routing was reserved to adjacent rows only (recall, the cable assembly cables are swizzled to connect TX to RX to maintain the same connector pinouts on both boards). For example, adjacent row connectivity (TX to RX) is shown by routing row A (TX0_P/N) on the source connector through the cable assembly to row B (RX0_P/N) on the sink connector and its adjacent counterpart (RX to TX) is shown by routing row B (RX0_P/N) on the source connector through the cable assembly to row A (TX0_P/N) on the sink connector, thereby maintaining symmetry across the whole data path (connector footprint (source) to connector footprint (sink)). Figure 13 provides a diagram of the swizzling feature.

Skew Between Pairs	Skew¹ (ps)
Row A to Row B	22
Row B to Row C	21
Row C to Row D	20
Unpopulated Row (E)	
Row D to Row F	64
Row F to Row G	20
Row G to Row H	21
Row H to Row J	21

Note 1: Skew path is from connector footprints through cable header

Table 3: Skews for STRADA Whisper Connector and Header




Figure 13: Cable Assembly Connectivity Swizzle Example

CABLE: The skew the cable itself introduces is minimal and will vary slightly by cable length. The cable length is matched between the P/N signals of each differential pair. The worst case skew tolerances between any of the differential pairs for the various cable lengths are shown in Table 4.

<i>Cable Length</i>	<i>Skew Tolerance (ps)</i>
20 inches / 9 inches	+/-2

Table 4: Cable Skew Tolerances

All the before mentioned skews are provided to be taken into account for the board’s overall data skew budget. The Interlaken Alliance’s 400G+ interoperability specified skew tolerance is 214 UI and must be adhered to for compliant Interlaken interoperability at these higher bandwidths. The connector and cable assembly minimally impact the skew budget, therefore each PCB can use up to 100 UI of the skew budget.

The maximum intra-pair skew (between P/N) is 2 ps.

4.5 Connectors and Cable Assembly Insertion Loss

The differential insertion loss includes receptacles at both ends with the cable assembly (cable + header) completing the channel connectivity. The model measurement does not take PCB footprints in account. See Section 4.7 – Signal Integrity Considerations and PCB Items for additional information on footprint modeling. Various colors are shown in Figure 14 which is representative of the insertion losses for each channel. From the data shown below, the data paths are well matched and exhibit roughly the same degree of loss across the swept frequencies.




Figure 14: Differential Insertion Loss (9 Inch Cable Assembly)

To ensure end-to-end interoperability connectivity, a loss budget is presented in Table 5. The data shown in Table 5 provides the loss acceptances for the PCBs, connectors, and cable assembly that are to be

used in an interoperability activity. The FPGA-to-Connector (Source) allows up to 7 dB of loss to occur across the PCB when routing high-speed traces from the FPGA package balls through the STRADA Whisper connector. For the cable assembly loss graph shown in Figure 14, the loss attributed to the overall budget from this item that attaches to the connector is ~4 dB from 12 GHz to 15 GHz. As the cable assembly enables connection to another board with a STRADA Whisper connector, the FPGA-to-Connector's (Sink) loss is identical to the adjacent PCB's requirements of 7 dB of loss. By adhering to the various stages of the interoperability loss budget, end-to-end connectivity is achieved.

<i>FPGA-to-Connector (PCB) Loss Source</i>	<i>Cable Assembly Loss</i>	<i>FPGA-to-Connector (PCB) Loss Sink</i>	<i>Total End-to-End Link Loss</i>	<i>Frequency</i>
7 dB	4 dB	7 dB	18 dB	From 12.5 GHz to 15 GHz

Table 5: End-to-End Interoperability Loss Budget

4.6 AC Coupling

AC coupling capacitors must be placed on the receiver lanes between the device and the connector. Recommended AC coupling capacitor value for data rates in the range from 12.5 Gbps to 30 Gbps are 100 nF.

4.7 Signal Integrity Considerations and PCB Items

Please refer to TE Connectivity's collateral document for further details on signal integrity considerations, mechanical drawings, and Allegro board files to simplify the design and usage of the 400G+ Interlaken interoperability connector and cable assembly:

- STRADA Whisper 3.9 mm Connector Routing Guide (Report # 29GC001)
- Engineering Test Report: end-to-end channel: includes source board connector, cable assembly, and sink board connector. Also includes simulation plots and signal integrity measurements of the end-to-end channel
- Allegro board files with footprint model
- S-parameter models

Please email TE Connectivity for the above mentioned collateral items: modeling@te.com