

Interlaken Interoperability Recommendations

5 January 2009

Revision 1.3

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.....	4
1 Overview.....	5
2 Logical Interoperability Recommendations.....	6
2.1 Definitions of Interoperability Parameters.....	6
2.2 Recommendations for All Transfer Rates	7
2.3 10Gb/s Packet Transfers	7
2.3.1 Target Applications.....	7
2.3.2 Recommendations.....	7
2.4 25Gb/s Packet Transfers	8
2.4.1 Target Applications.....	8
2.4.2 Recommendations.....	8
2.5 50Gb/s Packet Transfers	9
2.5.1 Target Applications.....	9
2.5.2 Recommendations.....	9
2.6 100Gb/s Packet Transfers	10
2.6.1 Target Applications.....	10
2.6.2 Recommendations.....	10
2.7 Summary of Lane Recommendations	11
2.8 Multirate Interoperability	11
2.9 Latency Requirements	12
3 Physical Interop Recommendations	14
3.1 Connector Type	14
3.2 Connector Pin Assignment.....	14
3.3 Connector Placement.....	16
3.4 Physical Interop examples.....	16
3.4.1 Connecting 10Gb/s to 10Gb/s	17
3.4.2 Connecting 25Gb/s to 25Gb/s	17
3.4.3 Connecting 25Gb/s to 10Gb/s	17
3.4.4 Connecting 50Gb/s to 50Gb/s	18
3.4.5 Connecting 50Gb/s to 25Gb/s	18
3.4.6 Connecting 50Gb/s to 10Gb/s	18
3.4.7 Connecting 100Gb/s to 100Gb/s.....	19
3.4.8 Connecting 100Gb/s to 50Gb/s	20
3.4.9 Connecting to multiple Interlaken interfaces	20

Revision History

<p style="text-align: center;">Revision 1.3 Revision Date: 5 January 2009</p>
<ul style="list-style-type: none">• Added 10.3125Gb/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 Physical Interop 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 10Gb/s, 25Gb/s, 50Gb/s and 100Gb/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.	10.3125Gb/s, 6.25Gb/s and 3.125Gb/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 uses 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.
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, 1Gb/s steps
Status Messaging	Not required
Comments	For unidirectional data paths, out-of-band backpressure may be useful. LVCMOS electrical levels are then preferred.

2.3 10Gb/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 10Gb/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	6.25Gb/s
Electrical Specification	CEI-6G-SR
Comments	10GE 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	3.125Gb/s

2.4 25Gb/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	6.25Gb/s
Electrical Specification	CEI-6G-SR

2.5 50Gb/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 50Gb/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	6.25Gb/s
Electrical Specification	CEI-6G-SR

2.5.2.2 Recommendation #2 – Narrow Interface Option

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

2.6 100Gb/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 two different recommendations at the 100Gb/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.

2.6.2.1 Recommendation #1 – Wide Interface Option

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

2.6.2.2 Recommendation #2 – Narrow Interface Option

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

2.7 Summary of Lane Recommendations

Transfer Rate	Recommendation #1 Wide Options			Recommendation #2 Narrow Options		
	Lanes	Rate [Gb/s]	Electrical	Lanes	Rate [Gb/s]	Electrical
10Gb/s	5	3.125		3	6.25	CEI-6G-SR
25Gb/s	5	6.25	CEI-6G-SR			
50Gb/s	10	6.25	CEI-6G-SR	6	10.3125	CEI-11G-SR
100Gb/s	20	6.25	CEI-6G-SR	12	10.3125	CEI-11G-SR

2.8 Multirate Interoperability

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

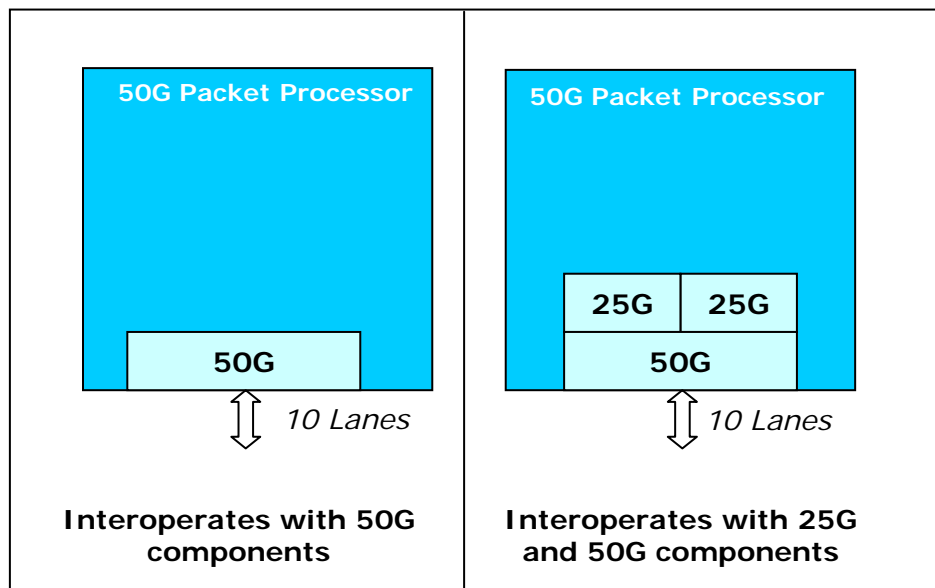


Figure 1: Example of interoperability with lower capacity components

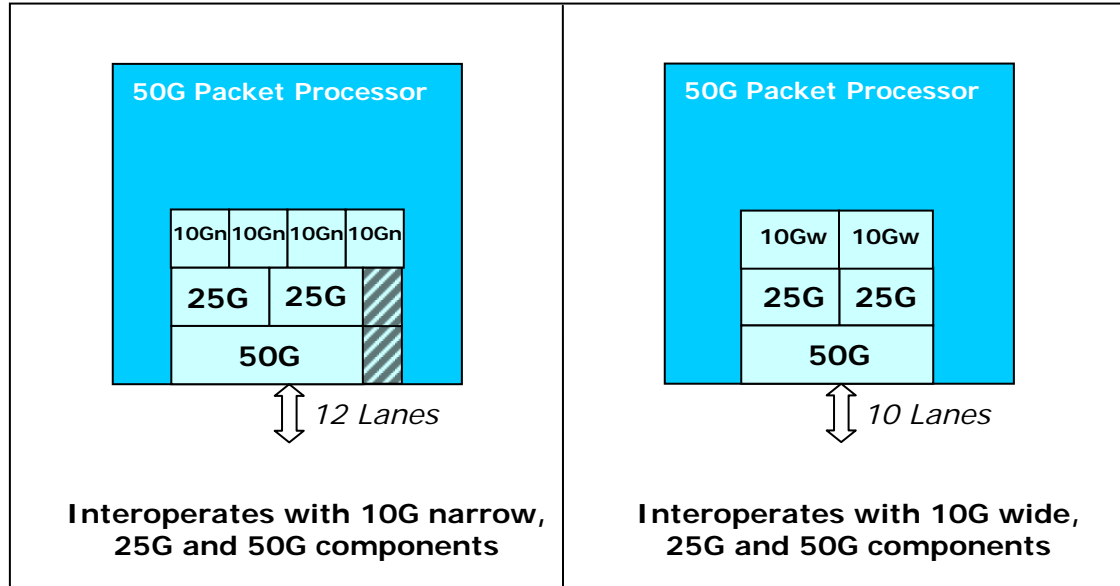


Figure 2: Example of interoperability with 10G components

2.9 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 $L_{max-xoff}$ bytes go through point D after XOFF is seen at point A. The various contributors to $L_{max-xoff}$ 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).

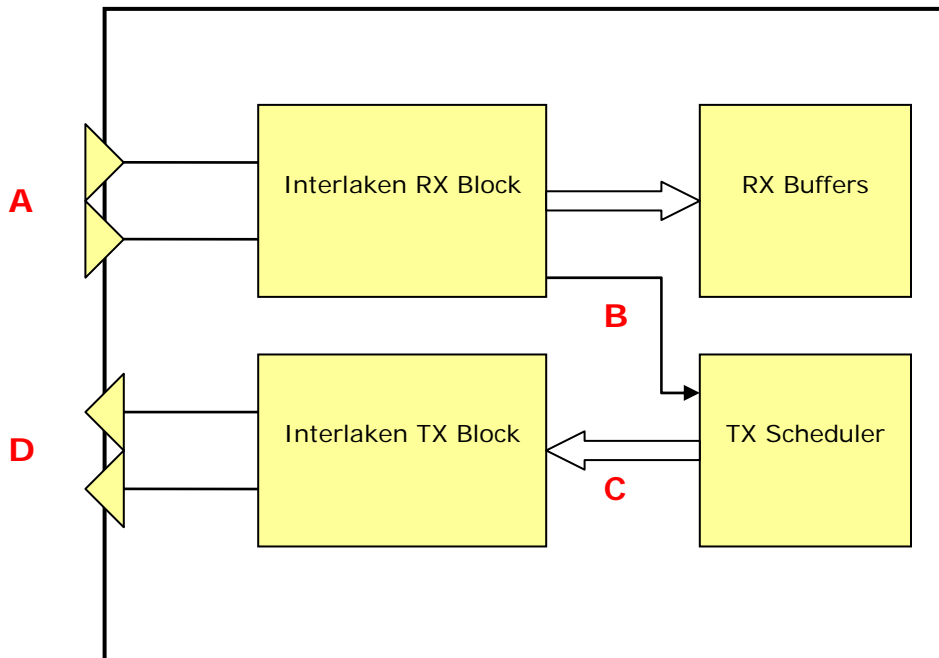


Figure 3 Transmitter Block Diagram

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.

Number of Lanes	Serial Rate	Lmax-xoff
	(Gbps)	(Bytes)
5	3.125	1,536
3	6.25	1,280
5	6.25	1,792
6	10.3125	6,400
10	6.25	6,400
12	10.3125	12,484
20	6.25	12,484

Table 1 Lmax-xoff Definitions

3 Physical Interop 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.

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).

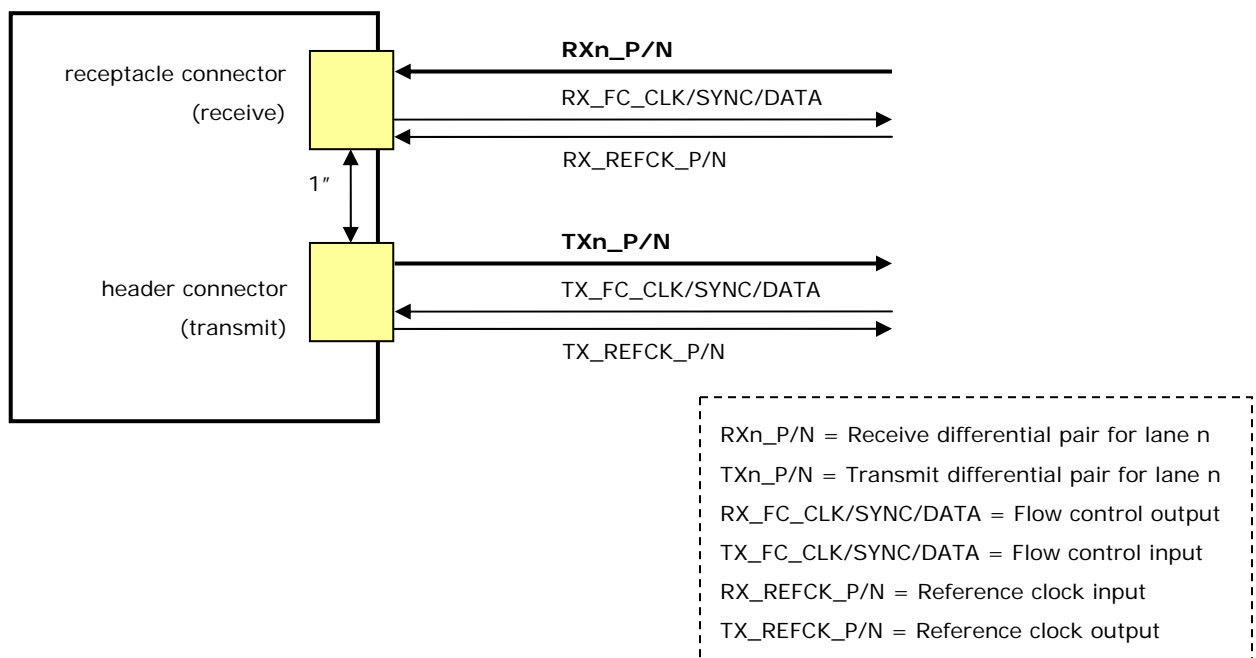


Figure 4: Signal names for connector pin assignment

	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. 25Gb/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 100Gb/s interface shall employ 2 sets of connectors with the same pin assignment for 50Gb/s interface. The placement of these connectors is described in Section 3.3 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 100Gb/s interface, two sets of connectors are needed. The two connectors shall be placed 5" apart as shown below. For 100Gb/s interface to interoperate with 50Gb/s (or lower) interface, channel mapping will be needed on the 100Gb/s interface, as shown in Section 3.4.8 below.

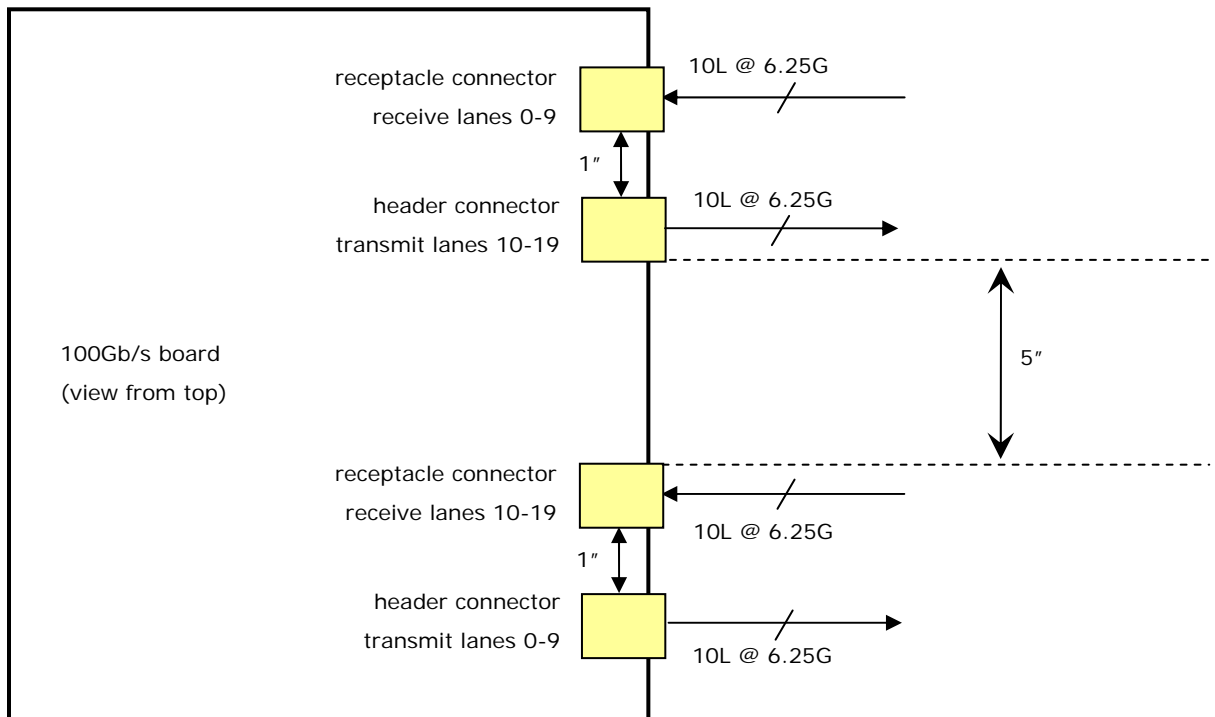



Figure 6: Connector placement for 100Gb/s interface


3.4 Physical Interop examples

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


3.4.1 Connecting 10Gb/s to 10Gb/s




3.4.2 Connecting 25Gb/s to 25Gb/s



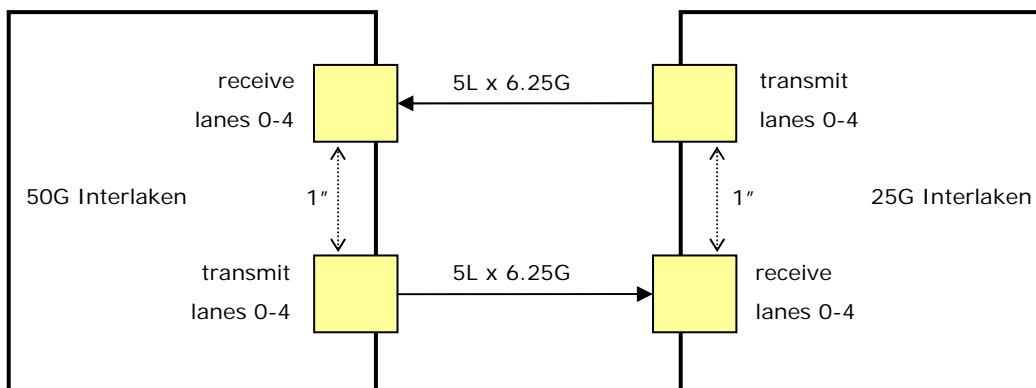
3.4.3 Connecting 25Gb/s to 10Gb/s




3.4.4 Connecting 50Gb/s to 50Gb/s



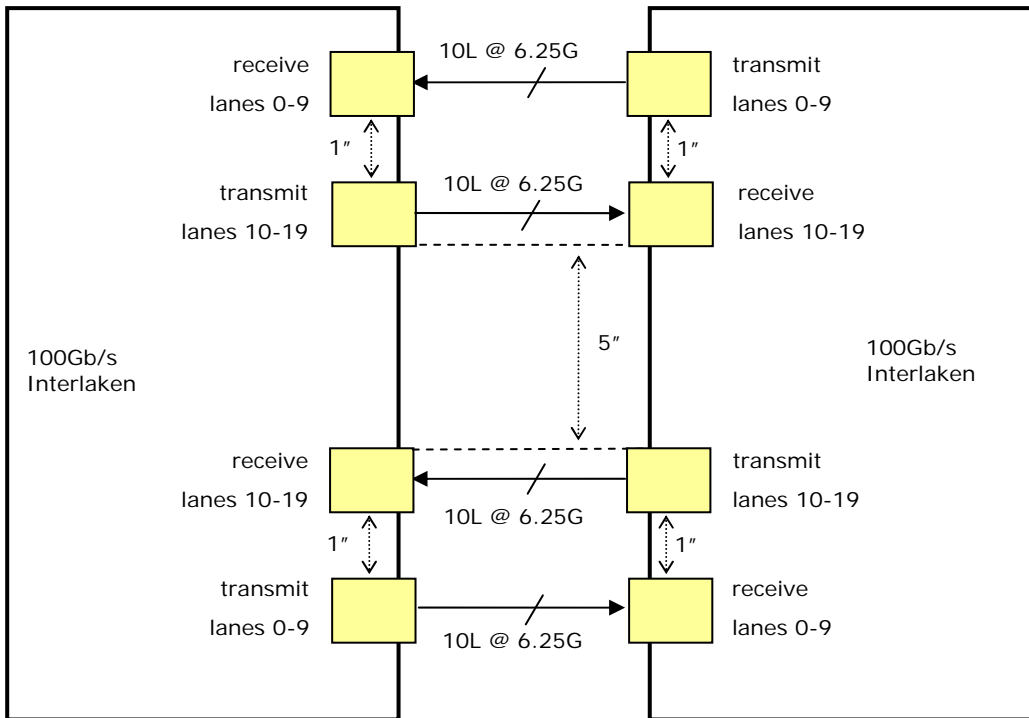
3.4.5 Connecting 50Gb/s to 25Gb/s



3.4.6 Connecting 50Gb/s to 10Gb/s



3.4.7 Connecting 100Gb/s to 100Gb/s

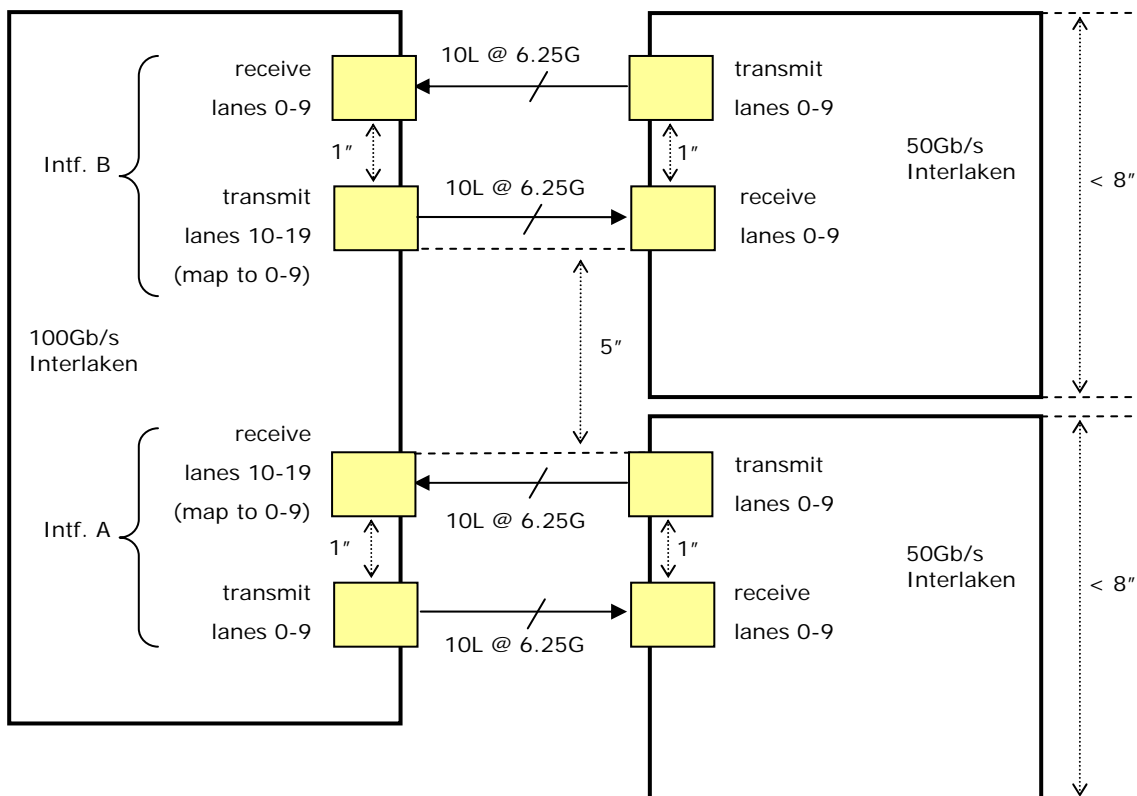


3.4.8 Connecting 100Gb/s to 50Gb/s

For 100Gb/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 2.7, 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. 50Gbps component connecting to two 25Gbps components), these recommendations apply:

- Components supporting multiple Interlaken interfaces of rates of 50Gbps 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 100Gbps 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 Interop 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 10 is a spare lane.
- 100G connecting to 2 x 50G: this is shown in Section 3.4.8.
- 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.