



OpenCAPI 3.0

25 Gbps PHY Mechanical Specification

Version 1.0
24 February 2020

Approved

Approved for Distribution to OpenCAPI Members
Approved for Distribution to Non-Members for Learning Purposes Only

OpenCAPI 3.0: 25 Gbps PHY Mechanical Specification

PHY Signaling Work Group
OpenCAPI Consortium

Version 1.0 (24 February 2020)

Copyright © OpenCAPI Consortium 2020

Use of this document is controlled by the OpenCAPI Consortium License Agreement, which is available at <https://opencapi.org/license/>.

All capitalized terms in the following text have the meanings assigned to them in the OpenCAPI Intellectual Property Rights Policy (the "OpenCAPI IPR Policy"). The full Policy may be found at the OpenCAPI Consortium website.

THE SPECIFICATION IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, ACCURACY, COMPLETENESS AND NONINFRINGEMENT OF THIRD PARTY RIGHTS. IN NO EVENT SHALL LICENSOR, ITS MEMBERS OR ITS CONTRIBUTORS BE LIABLE FOR ANY CLAIM, OR ANY DIRECT, SPECIAL, INDIRECT OR CONSEQUENTIAL DAMAGES, OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THE SPECIFICATION.

OpenCAPI and the OpenCAPI logo design are trademarks of the OpenCAPI Consortium.

Other company, product, and service names may be trademarks or service marks of others.

Abstract

This document describes technical details for the design of mezzanine OpenCAPI add-in cards and 25 Gbps cable implementations. It is the work product of the OpenCAPI Consortium PHY Signaling Work Group.

This document is handled in compliance with the requirements outlined in the OpenCAPI Consortium Work Group (WG) process document. Comments, questions, etc. can be submitted to membership@opencapi.org.

Contents

List of figures.....	5
List of tables	6
Revision log	7
Preface	8
Conventions.....	8
Notes	8
Engineering notes	8
Developer notes	9
Editor notes	9
Terms	10
1. Overview	12
2. Mezzanine card mechanical detail	13
2.1 OpenCAPI mezzanine card connectors	13
2.1.1 Example of typical mezzanine connector.....	13
2.2 OpenCAPI mezzanine card outline	14
2.2.1 OpenCAPI mezzanine physical card diagram.....	14
2.2.2 Top view of the mezzanine card	15
2.2.3 Bottom view of the mezzanine card	16
2.3 Mechanical drawing for the co-docking connector pads.....	17
2.3.1 Thermal guidance	18
2.4 OpenCAPI mezzanine card mechanical and attach requirements	18
2.5 Overall mechanical stack.....	18
3. Mezzanine card electrical references	20
3.1 Mezzanine loss budget.....	20
3.2 Mezzanine card power-delivery capability.....	20
3.3 Mezzanine pinout	21
3.3.1 Connector 1 pinout.....	21
3.3.2 Connector 2 pinout.....	22
3.4 Mezzanine card device wiring	23
3.5 Mezzanine sideband signals	24
3.6 Mezzanine PCIe interface	24
4. Carrier card and advanced accelerator cable	25
4.1 Carrier card size	25
4.2 Carrier card connector.....	25
4.3 Carrier card power	25
4.4 Carrier card cooling	25
4.5 Advance accelerator cable	26

Approved

4.5.1	Accelerator cable circuit schematic.....	26
4.5.2	Accelerator cable signal definition.....	26
4.5.3	Advanced accelerator cable loss budget	27
4.5.4	Accelerator-cable pin definition	28
4.5.5	Internal cable pin ordering and mapping.....	29
4.5.6	AAC connector escape routing example.....	30
4.6	Cabled CPU to carrier card interconnect and loss budget	30
4.6.1	CPU-to-QSFP carrier card configuration.....	30
4.6.2	CPU-to-QSFP carrier card loss budget.....	31
4.6.3	CPU-to-accelerator configuration.....	32
4.6.4	CPU-to-accelerator loss budget.....	32
4.7	Accelerator card design supplement	33
4.7.1	Topology 1	33
4.7.2	Topology 2	33
4.8	Carrier card stack-up	34

List of figures

Figure 2-1. Mezzanine card	13
Figure 2-2. Example of typical mezzanine connector	13
Figure 2-3. Front-side view of mezzanine card layout	14
Figure 2-4. Top view of mezzanine card outline	15
Figure 2-5. Bottom view of mezzanine card outline	16
Figure 2-6. Reference dimension scheme for co-docking connector pads.....	17
Figure 2-7. Overall mechanical stack.....	18
Figure 3-1. Detailed pinout of connector 1	21
Figure 3-2. Detailed pinout of connector 2	22
Figure 3-3. OpenCAPI mezzanine card lane identification	23
Figure 4-1. Carrier card connecting to CPU board	25
Figure 4-2. Accelerator cable circuit schematic	26
Figure 4-3. Accelerator cable pin definition.....	28
Figure 4-4. GRD board pin ordering	29
Figure 4-5. AAC connector escape routing example	30
Figure 4-6. Typical 25G configuration.....	31
Figure 4-7. Example of a CPU-to-accelerator configuration	32
Figure 4-8. Example of a direct-attach accelerator card	33
Figure 4-9. Example of a cable-attached accelerator carrier card	33
Figure 4-10. Example of a carrier card stack up	34

List of tables

Table 3-1. Mezzanine loss budget	20
Table 3-2. System power delivery capability per add-in card (set of two connectors)	20
Table 4-1. Advanced accelerator cable loss budget	27
Table 4-2. Example of CPU-to-QSFP carrier card loss budget	31
Table 4-3. Example of CPU-to-accelerator loss budget (OpenCAPI module)	32

Approved

Revision log

Each release of this document supersedes all previously released versions. The revision log lists all significant changes made to the document since its initial release.

Revision date	Summary of changes
24 February 2020	Version 1.0. Initial release to OpenCAPI Consortium.

Preface

This specification describes technical details for the design of mezzanine OpenCAPI add-in cards and 25 Gbps cable implementations.

This specification is intended for designers who plan to develop products that use OpenCAPI. This document provides a technical overview of the mechanical component guidance for 25 Gbps OpenCAPI system design.

Conventions

The following typographical conventions are used in this document.

Convention	Description
Hyperlink	Web-based <u>URLs</u> are displayed in blue text to denote a virtual link to an external document.
Note: This is note text.	The note text denotes information that emphasizes a concept or provides critical information.
Footnote reference. ¹ 1. Descriptive footnote text.	A footnote is an explanatory note or reference inserted at the foot of the page or under a table that explains or expands upon a point within the text or indicates the source of a citation or peripheral information.

Notes

This section describes Engineering and Developer notes.

Engineering notes

Engineering notes provide additional implementation details and recommendations not found elsewhere. The notes might include architectural compliance requirements. That is, the text might include Architecture compliance terminology. These notes should be read by all implementation and verification teams to ensure architectural compliance.

Engineering note:

This is an example of an Engineering note. Lorem ipsum dolor sit amet, consectetur adipiscing elit. Proin cursus hendrerit enim, vel tempus nibh ornare ut. Quisque ac augue eu augue convallis hendrerit. Mauris iaculis viverra ipsum nec dapibus. Nunc at porta libero. Curabitur luctus ultrices augue non pulvinar. Vestibulum mattis non ipsum at venenatis. Suspendisse euismod, neque et suscipit luctus, odio metus semper lectus, quis volutpat est libero quis nunc. Vivamus rutrum mauris sed tristique malesuada.

Approved

Developer notes

Developer notes are used to document the reasoning and discussions that led to the current version of the architecture. These notes might also include recommended changes for future versions of the architecture, or warnings of approaches that have failed in the past. These notes should be read by verification teams and contributors to the architecture.

Developer note:

This is an example of a Developer note. Lorem ipsum dolor sit amet, consectetur adipiscing elit. Proin cursus hendrerit enim, vel tempus nibh ornare ut. Quisque ac augue eu augue convallis hendrerit. Mauris iaculis viverra ipsum nec dapibus. Nunc at porta libero. Curabitur luctus ultrices augue non pulvinar. Vestibulum mattis non ipsum at venenatis. Suspendisse euismod, neque et suscipit luctus, odio metus semper lectus, quis volutpat est libero quis nunc. Vivamus rutrum mauris sed tristique malesuada.

Editor notes

Editor notes are reminders to the editors and other contributors of additional work that is required. Editor notes may appear as red text in the body of any section or may include an entire section. This is to indicate text that is either out of date or is speculative (unapproved) in nature. The red text might also include directions to the editor for changes to be applied to a future revision of the document. Approved versions of a document are not expected to be released with red text unless approved by the owners of the document.

Approved

Terms

The following terms are used in this document.

A	Amperes.
AAC	Advance accelerator cable.
AC	Alternating current.
ASIC	Application-specific integrated circuit.
BGA	Ball grid array.
CAPI	Coherent Accelerator Processor Interface.
CPU	Central processing unit.
C4	Controlled Collapse Chip Connection.
FPGA	Field-programmable gate array.
Gbps	Gigabits per second.
GND	Ground.
HLGA	Hybrid land grid array.
I2C	Inter-integrated circuit.
lbf	Pound-force.
LGA	Land grid array.
LSB	Least-significant byte.
OD	Open drain.

Approved

PCB	Printed circuit board.
PCIe	Peripheral Component Interconnect Express.
QSFP	Quad small form-factor pluggable.
UI	Unit interval (for example, 25Gbps = $1/25 = 40\text{ps}$).
V	Volt.

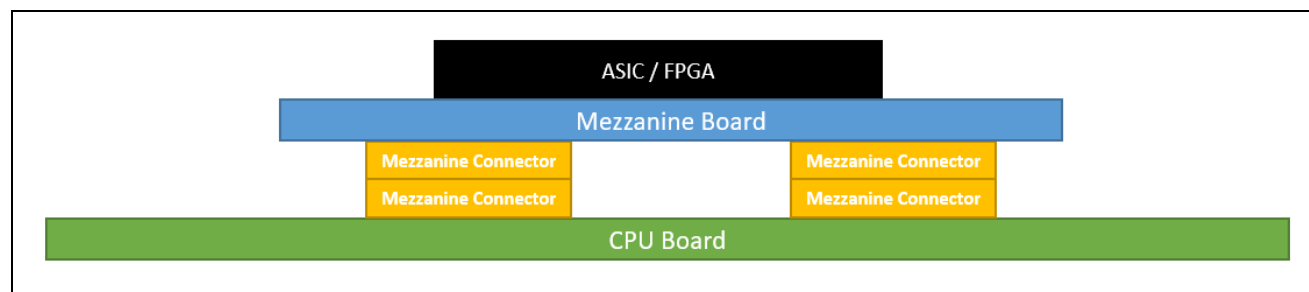
1. Overview

This specification describes the technical details for the design of OpenCAPI mezzanine cards, carrier cards, and 25 Gbps cable implementations. This specification is constrained to the use of the tested and verified connector technology and I/O assignments. The scope of this document is to provide a technical overview of the mechanical requirements for 25 Gbps OpenCAPI system design.

2. Mezzanine card mechanical detail

This section describes the mechanical constraints and details of designing an OpenCAPI add-in card.

Figure 2-1. Mezzanine card



2.1 OpenCAPI mezzanine card connectors


OpenCAPI add-in cards use mezzanine connectors as the electrical interface to the system planar. The connector is a 400 pin, 4 mm stack-height, high-speed interconnect. The OpenCAPI add-in card requires the use of connectors that meet the electrical and mechanical performance requirements. Examples of these connectors are shown in *Figure 2-2. Example of typical mezzanine connector.*

Each add-in card co-docks two connectors. One connector is primarily for high-speed signals and the other connector is for power and lower-speed signals. Co-docking these connectors drives the specific PCB tolerances described in *Section 2.2. OpenCAPI mezzanine card outline.* The low nominal pin wipe in the connectors drives specific mechanical tolerances on the bottom-side stiffener. See *Section 2.4. OpenCAPI mezzanine card mechanical and attach requirements.*

2.1.1 Example of typical mezzanine connector

Figure 2-2 shows an example of a typical mezzanine connector.

Figure 2-2. Example of typical mezzanine connector



2.2 OpenCAPI mezzanine card outline


The top view of the mezzanine card outline is shown in *Figure 2-4. Top view of mezzanine card outline*. The maximum add-in card planar dimensions are 78 mm × 140 mm. For orientation in the system, the north side of the card is identified on the drawing. Two alignment pins are used to orient the card within the system. The origin pin (south side) drives the alignment and docking of the card to the planar. This feature is critical for proper alignment of the add-in card and the system planar. The north-side alignment pin provides angular alignment and can be slotted on the card for hole positional tolerance.

The mounting hole dimensions, alignment holes, and pins A1 and K40 are depicted in *Figure 2-3*. The orientation and connector identification (1 and 2) are also defined in this view, as well as the location of two copper grounding pads located between the connector pair.

2.2.1 OpenCAPI mezzanine physical card diagram

Figure 2-3 shows the OpenCAPI mezzanine physical card diagram with pin numbering details.

Figure 2-3. Front-side view of mezzanine card layout




Approved

2.2.2 Top view of the mezzanine card

Figure 2-4 shows the top view of the mezzanine card outlines with a North direction indicator. Figure 2-4 includes datum identification, connector placement, and pin identification. Dimensions are in millimeters.


Figure 2-4. Top view of mezzanine card outline



2.2.3 Bottom view of the mezzanine card

Figure 2-5 shows the bottom view of the mezzanine card outlines with a North direction indicator. The figure includes datum identification, connector placement, and pin identification. Dimensions are in millimeters.

Figure 2-5. Bottom view of mezzanine card outline




For a 2U chassis, the maximum component height on the top side of the card is approximately 66 mm. The total height of the OpenCAPI cards must be less than 71.75 mm. It is preferred that the card extend the full height (including the heat sink) or provide airflow blocking features.

2.3 Mechanical drawing for the co-docking connector pads

Figure 2-6 shows the mechanical drawing for the co-docking connector pads.

Figure 2-6. Reference dimension scheme for co-docking connector pads



Approved

2.3.1 Thermal guidance

Unless the mezzanine card design is known to be significantly lower power, the thermal solution should be capable of dissipating a maximum of 305 W, which is the limit of the input power for the mezzanine card. Because the mezzanine card is intended to be mounted inside a host chassis, any cooling airflow must be assumed to be approximately 5 °C above ambient inlet air temperature to account for pre-heating by the host chassis and motherboard component.


2.4 OpenCAPI mezzanine card mechanical and attach requirements

A stiffener is required on the bottom side of the add-in cards to ensure that the connectors are loaded appropriately. The stiffener requires a minimum thickness of 3.93 ± 0.1 mm with a minimum insulator thickness of 0.125 mm on the bottom of the stiffener. The total thickness of stiffener and insulators (bottom and top insulators with adhesive) is 4.18 ± 0.12 mm.

2.5 Overall mechanical stack

Figure 2-7 illustrates the overall mechanical stack.

Figure 2-7. Overall mechanical stack



The stiffener is required to be full thickness around the mounting locations. However, the stiffener can be less than full thickness in other locations, which allows the placement of components on the bottom side of the mezzanine card. The design of the stiffener should allow full-force insertion of the connectors (approximately 31 lbf connector) and maintain integrity of the BGA grid on the connector. Care should be taken when designing a bottom-side stiffener (as well as a top-side stiffener if required). Minimize board strain in the area around the connector BGA to reduce damage to the connector and solder joints during card insertion and extraction.

The stiffener requires including two conductive-fabric-over-foam gaskets to touch the copper grounding pads. (The gasket-attach plane is nominally inset 0.35 mm from the bottom-stiffener plane to achieve this)

Approved

compression.) This fabric-over-foam gasket must also connect electrically to the OpenCAPI bottom stiffener and provide a ground path for the card. This is a chassis ground.

The system planar should have a package keep-out for the full 78 mm × 140 mm on the top side on the main planar. However, top-side traces are allowed in this region. Add-in cards must ensure that areas of the OpenCAPI card that are in contact with the main planar are electrically isolated (except for the conductive foam gaskets). It is recommended that a thin insulator (0.125 mm) be used for this purpose. Nonconductive coatings can be used but might not be considered the primary method of isolation and damage mitigation to the planar card.

Attachment to the planar is by eight threaded fasteners at the mounting locations. These eight mounting locations are defined as M3 × 0.5 and the threaded depth is 5 mm. Accounting for a planar thickness of 3.01 mm, the maximum threaded depth beyond the bottom surface for the OpenCAPI card is 8 mm. Use an exposed threaded fastener with a length between 6.5 - 7.5 mm from the bottom of the OpenCAPI card.

Removing the mezzanine card from the system planar requires an upward force to un-mate the connectors.

3. Mezzanine card electrical references

This section details the power delivery capability and pinout of the mezzanine connectors labeled 1 and 2. Detailed requirements are found in the *OpenCAPI 25 Gbps Physical Signaling Specification*.

3.1 Mezzanine loss budget

Table 3-1 lists the mezzanine loss budget for the following:
CPU ↔ Host board ↔ Mezzanine Card CONNECTOR ↔ Mezzanine card trace ↔ ASIC or FPGA

Table 3-1. Mezzanine loss budget

Component	Loss Budget (dB)	Notes
CPU socket	4.5	34 mm trace, LGA 7-2-7 laminate package
Reference mother board trace	8.5	
Mezzanine card connector	0.3	
Mezzanine card trace	6.7	
Total	20	Receiver specification

3.2 Mezzanine card power-delivery capability

Table 3-2 details typical power delivery capability of the system as a function of voltage domain per add-in card. All amperages are assumed to be equally distributed amongst all pins on that domain.

Table 3-2. System power delivery capability per add-in card (set of two connectors)

Voltage domain	Amperage	Allowable excursion
12 V	25 A	30% for < 100 ms
5 V	1 A	30% for < 100 ms

Approved

3.3 Mezzanine pinout

3.3.1 Connector 1 pinout

Figure 3-1 shows a detailed pinout of connector 1. The top view of the mezzanine card outlines with a North direction indicator for placement definition.

Figure 3-1. Detailed pinout of connector 1

	A	B	C	D	E	F	G	H	J	K
1	OP3_RX7_N	GND	OP3_RX5_N	GND	GND	GND	GND	OP3_TX6_N	GND	OP3_TX5_N
2	OP3_RX7_P	OP3_RX6_N	OP3_RX5_P	OP3_RX4_N	GND	GND	OP3_TX7_N	OP3_TX6_P	OP3_TX4_N	OP3_TX5_P
3	GND	OP3_RX6_P	GND	OP3_RX4_P	GND	GND	OP3_TX7_P	GND	OP3_TX4_P	GND
4	OP3_RX3_N	GND	OP3_RX2_N	GND	GND	GND	GND	OP3_TX3_N	GND	OP3_TX2_N
5	OP3_RX3_P	OP3_RX0_N	OP3_RX2_P	OP3_RX1_N	GND	GND	OP3_TX0_N	OP3_TX3_P	OP3_TX1_N	OP3_TX2_P
6	GND	OP3_RX0_P	GND	OP3_RX1_P	GND	GND	OP3_TX0_P	GND	OP3_TX1_P	GND
7	OP2_RX7_N	GND	OP2_RX5_N	GND	GND	GND	GND	OP2_TX6_N	GND	OP2_TX5_N
8	OP2_RX7_P	OP2_RX6_N	OP2_RX5_P	OP2_RX4_N	GND	GND	OP2_TX7_N	OP2_TX6_P	OP2_TX4_N	OP2_TX5_P
9	GND	OP2_RX6_P	GND	OP2_RX4_P	GND	GND	OP2_TX7_P	GND	OP2_TX4_P	GND
10	OP2_RX3_N	GND	OP2_RX2_N	GND	GND	GND	GND	OP2_TX3_N	GND	OP2_TX2_N
11	OP2_RX3_P	OP2_RX0_N	OP2_RX2_P	OP2_RX1_N	GND	GND	OP2_TX0_N	OP2_TX3_P	OP2_TX1_N	OP2_TX2_P
12	GND	OP2_RX0_P	GND	OP2_RX1_P	GND	GND	OP2_TX0_P	GND	OP2_TX1_P	GND
13	GND	GND	PRSENT1A_N	GND	GND	GND	GND	GND	GND	GND
14	GND	OP5_RX6_N	GND	OP5_RX5_N	GND	GND	OP5_TX7_N	GND	OP5_TX5_N	GND
15	OP5_RX7_N	OP5_RX6_P	OP5_RX4_N	OP5_RX5_P	GND	GND	OP5_TX7_P	OP5_TX6_N	OP5_TX5_P	OP5_TX4_N
16	OP5_RX7_P	GND	OP5_RX4_P	GND	GND	GND	GND	OP5_TX6_P	GND	OP5_TX4_P
17	GND	OP5_RX3_N	GND	OP5_RX2_N	GND	GND	OP5_TX3_N	GND	OP5_TX2_N	GND
18	OP5_RX0_N	OP5_RX3_P	OP5_RX1_N	OP5_RX2_P	GND	GND	OP5_TX3_P	OP5_TX0_N	OP5_TX2_P	OP5_TX1_N
19	OP5_RX0_P	GND	OP5_RX1_P	GND	GND	GND	GND	OP5_TX0_P	GND	OP5_TX1_P
20	GND	GND	GND	GND	REFCLK_156M_N	REFCLK_156M_P	GND	GND	GND	GND
21	OP4_RX7_N	GND	OP4_RX5_N	GND	GND	GND	GND	OP4_TX6_N	GND	OP4_TX5_N
22	OP4_RX7_P	OP4_RX6_N	OP4_RX5_P	OP4_RX4_N	GND	GND	OP4_TX7_N	OP4_TX6_P	OP4_TX4_N	OP4_TX5_P
23	GND	OP4_RX6_P	GND	OP4_RX4_P	GND	GND	OP4_TX7_P	GND	OP4_TX4_P	GND
24	OP4_RX3_N	GND	OP4_RX2_N	GND	GND	GND	GND	OP4_TX3_N	GND	OP4_TX2_N
25	OP4_RX3_P	OP4_RX0_N	OP4_RX2_P	OP4_RX1_N	GND	GND	OP4_TX0_N	OP4_TX3_P	OP4_TX1_N	OP4_TX2_P
26	GND	OP4_RX0_P	GND	OP4_RX1_P	GND	GND	OP4_TX0_P	GND	OP4_TX1_P	GND
27	GND	GND	GND	GND	GND	GND	GND	GND	GND	GND
28	GND	OP1_RX6_N	GND	OP1_RX5_N	GND	GND	OP1_TX7_N	GND	OP1_TX5_N	GND
29	OP1_RX7_N	OP1_RX6_P	OP1_RX4_N	OP1_RX5_P	GND	GND	OP1_TX7_P	OP1_TX6_N	OP1_TX5_P	OP1_TX4_N
30	OP1_RX7_P	GND	OP1_RX4_P	GND	GND	GND	GND	OP1_TX6_P	GND	OP1_TX4_P
31	GND	OP1_RX3_N	GND	OP1_RX2_N	GND	GND	OP1_TX3_N	GND	OP1_TX2_N	GND
32	OP1_RX0_N	OP1_RX3_P	OP1_RX1_N	OP1_RX2_P	GND	GND	OP1_TX3_P	OP1_TX0_N	OP1_TX2_P	OP1_TX1_N
33	OP1_RX0_P	GND	OP1_RX1_P	GND	GND	GND	GND	OP1_TX0_P	GND	OP1_TX1_P
34	GND	PRSENT1B_N	GND	GND	SMB_DAT	SMB_CLK	GND	GND	GND	GND
35	OP0_RX7_N	GND	OP0_RX5_N	GND	GND	GND	GND	OP0_TX6_N	GND	OP0_TX5_N
36	OP0_RX7_P	OP0_RX6_N	OP0_RX5_P	OP0_RX4_N	GND	GND	OP0_TX7_N	OP0_TX6_P	OP0_TX4_N	OP0_TX5_P
37	GND	OP0_RX6_P	GND	OP0_RX4_P	GND	GND	OP0_TX7_P	GND	OP0_TX4_P	GND
38	OP0_RX3_N	GND	OP0_RX2_N	GND	GND	GND	GND	OP0_TX3_N	GND	OP0_TX2_N
39	OP0_RX3_P	OP0_RX0_N	OP0_RX2_P	OP0_RX1_N	GND	GND	OP0_TX0_N	OP0_TX3_P	OP0_TX1_N	OP0_TX2_P
40	GND	OP0_RX0_P	GND	OP0_RX1_P	GND	GND	OP0_TX0_P	GND	OP0_TX1_P	GND

Connector directivity example: OP3_RX7_N means that the receiver is on the OpenCAPI card and the transmitter is on the system planar, while OP3_TX7_N means that the transmitter is on the OpenCAPI card and the receiver is on the system planar.

Approved

3.3.2 Connector 2 pinout

Figure 3-2 shows a detailed pinout of connector 2. The top view of the mezzanine card outlines with a North direction indicator for placement definition.


Figure 3-2. Detailed pinout of connector 2

	A	B	C	D	E	F	G	H	J	K
1	GND	PE1_RX5_N	GND	PE1_RX7_N	GND	GND	PE1_TX5_N	GND	PE1_TX6_N	GND
2	PE1_RX4_N	PE1_RX5_P	PE1_RX6_N	PE1_RX7_P	GND	GND	PE1_TX5_P	PE1_TX4_N	PE1_TX6_P	PE1_TX7_N
3	PE1_RX4_P	GND	PE1_RX6_P	GND	GND	GND	GND	PE1_TX4_P	GND	PE1_TX7_P
4	GND	PE1_RX2_N	GND	PE1_RX3_N	GND	GND	PE1_TX2_P	GND	PE1_TX3_P	GND
5	PE1_RX1_N	PE1_RX2_P	PE1_RX0_N	PE1_RX3_P	GND	GND	PE1_TX2_N	PE1_TX1_P	PE1_TX3_N	PE1_TX0_P
6	PE1_RX1_P	GND	PE1_RX0_P	GND	GND	GND	GND	PE1_TX1_N	GND	PE1_TX0_N
7	GND	INT_RST_N	GND	GND	REFCLK_100M_N	REFCLK_100M_P	GND	GND	PRSENT2A_N	GND
8	PE0_RX5_P	GND	PE0_RX6_P	GND	GND	GND	GND	PE0_TX5_P	GND	PE0_TX7_N
9	PE0_RX5_N	PE0_RX4_N	PE0_RX6_N	PE0_RX7_P	GND	GND	PE0_TX4_N	PE0_TX5_N	PE0_TX6_P	PE0_TX7_P
10	GND	PE0_RX4_P	GND	PE0_RX7_N	GND	GND	PE0_TX4_P	GND	PE0_TX6_N	GND
11	PE0_RX2_P	GND	PE0_RX3_P	GND	GND	GND	GND	PE0_TX2_P	GND	PE0_TX3_P
12	PE0_RX2_N	PE0_RX1_P	PE0_RX3_N	PE0_RX0_P	GND	GND	PE0_TX1_P	PE0_TX2_N	PE0_TX0_P	PE0_TX3_N
13	GND	PE0_RX1_N	GND	PE0_RX0_N	GND	GND	PE0_TX1_N	GND	PE0_TX0_N	GND
14	GND	GND	GND	GND	GND	GND	GND	GND	GND	GND
15	RSVD	GND	RSVD	GND	GND	GND	GND	RSVD	GND	RSVD
16	RSVD	RSVD	RSVD	RSVD	GND	GND	RSVD	RSVD	RSVD	RSVD
17	GND	RSVD	GND	RSVD	GND	GND	RSVD	GND	RSVD	GND
18	PWR_EN	GND	I2C_ADDR_ID1	GND	PWR_BRAKE_N	PERST_N	GND	I2C_ADDR_ID0	GND	PWR_GOOD
19	TH_OVERT_N	GND	JTAG_SEL	GND	GND	GND	GND	GND	GND	GND
20	JTAG_TRST_N	JTAG_TMS	JTAG_TCK	JTAG_TDO	JTAG_TDI	+5.0V	+5.0V	+5.0V	+5.0V	+5.0V
21	GND	GND	GND	GND	GND	GND	GND	GND	GND	GND
22	+12.0V	+12.0V	+12.0V	+12.0V	+12.0V	+12.0V	+12.0V	+12.0V	+12.0V	+12.0V
23	+12.0V	+12.0V	+12.0V	+12.0V	+12.0V	+12.0V	+12.0V	+12.0V	+12.0V	+12.0V
24	GND	GND	GND	GND	GND	GND	GND	GND	GND	GND
25	+12.0V	+12.0V	+12.0V	+12.0V	+12.0V	+12.0V	+12.0V	+12.0V	+12.0V	+12.0V
26	+12.0V	+12.0V	+12.0V	+12.0V	+12.0V	+12.0V	+12.0V	+12.0V	+12.0V	+12.0V
27	GND	GND	GND	GND	GND	GND	GND	GND	GND	GND
28	+12.0V	+12.0V	+12.0V	+12.0V	+12.0V	+12.0V	+12.0V	+12.0V	+12.0V	+12.0V
29	+12.0V	+12.0V	+12.0V	+12.0V	+12.0V	+12.0V	+12.0V	+12.0V	+12.0V	+12.0V
30	GND	GND	GND	GND	GND	GND	GND	GND	GND	GND
31	+12.0V	+12.0V	+12.0V	+12.0V	+12.0V	+12.0V	+12.0V	+12.0V	+12.0V	+12.0V
32	+12.0V	+12.0V	+12.0V	+12.0V	+12.0V	+12.0V	+12.0V	+12.0V	+12.0V	+12.0V
33	GND	GND	GND	GND	GND	GND	GND	GND	GND	GND
34	+12.0V	+12.0V	+12.0V	+12.0V	+12.0V	+12.0V	+12.0V	+12.0V	+12.0V	+12.0V
35	+12.0V	+12.0V	+12.0V	+12.0V	+12.0V	+12.0V	+12.0V	+12.0V	+12.0V	+12.0V
36	GND	GND	GND	GND	GND	GND	GND	GND	GND	GND
37	+12.0V	+12.0V	+12.0V	+12.0V	+12.0V	+12.0V	+12.0V	+12.0V	+12.0V	+12.0V
38	+12.0V	+12.0V	+12.0V	+12.0V	+12.0V	+12.0V	+12.0V	+12.0V	+12.0V	+12.0V
39	GND	GND	GND	GND	GND	GND	GND	GND	GND	PRSENT2B_N
40	+12.0V	+12.0V	+12.0V	+12.0V	+12.0V	+12.0V	+12.0V	+12.0V	+12.0V	+12.0V

3.4 Mezzanine card device wiring

Figure 3-3 shows the OpenCAPI mezzanine card lane identification.

Figure 3-3. OpenCAPI mezzanine card lane identification



The OpenCAPI mezzanine card pinout is pre-defined and contains six, 8-lane ports (OP0 - OP5). The implementation and assignment of the available ports is determined by the developer's requirements.

Potential usage cases are as follows:

- Develop the mezzanine card for a predefined back plane. Care must be taken to ensure that the chosen ports and pin assignments of the mezzanine card match the back-plane wiring. The chosen lanes must be attached to valid OpenCAPI lanes.
- Develop the mezzanine card and the back plane. The developer has more freedom in this case to maximize the lane use and back-plane wiring. The port use is determined by the developer's application, which would potentially factor in system planar component layout, number of CPUs, and wiring plane restrictions. The chosen lanes must attach to the valid OpenCAPI lanes from the processor module.

3.5 Mezzanine sideband signals

The following sideband signals are present:

- REFCLK_100MHz – PCIe reference clock (from system logic).
- REFCLK_133MHz – 156.25 MHz reference clock (from the CPU). Can also be configured to be 133 MHz.
- I2C_ADDR0/1 – I²C bus address LSBs. Pulled high (1.8 V only) or pulled low on motherboard. These are used to set the physical address of the OPEN_CAPI device. These pins must be at a steady state voltage before the OPEN_CAPI device is taken out of reset.
- SMB_CLK/DAT – SMBUS. Only supports operation at 1.8 V. Care must be taken that this signal does not activate before PWR_EN and PWR_GOOD is valid. It is recommended that this signal be ANDED with PWR_EN at minimum because PWR_GOOD is optional.
- JTAG_SEL – Control signal into OPEN_CAPI device to enable JTAG debug interface. 1.8 V signal. Pull down on motherboard for normal operation. Pull high on mezzanine card for JTAG operation.
- JTAG_TCK/TMS/TDI/TDO/TRST_N – JTAG interface for debug. Scan rings can span multiple riser sites. Ring accessible through debug connector on motherboard. 1.8 V operation only.
- PERST_N – “PERST” signal to OPEN_CAPI device. Effectively is the master reset. 1.8 V level signal.
- PWR_EN – Master “on/off” switch to OPEN_CAPI device power subsystem. 1.8 V level signal only (OD, pull up is on motherboard). Signal is generated by logic on the system planar and is controlled by the power sequencer. This signal indicates to the OPEN_CAPI card that the OPEN_CAPI card power is enabled. The system asserts this signal to power on the module and may be asserted only after all input rails are stable.
- PWR_BRAKE_N – “throttle” signal to OPEN_CAPI device. 1.8 V signal. OD. Pull up on motherboard. Motherboard logic contribution from CPU and power supply throttle signal.
- PWR_GOOD – OPEN_CAPI device PGOOD signal-to-system. 1.8 V OD. Pull up with 10 KΩ on motherboard. Indicates to the power sequencer that OPEN_CAPI regulators are good. Module power good. The module will assert this signal when all of its internal power regulators are stable. Use by system is optional.
- TH_OVERT_N – emergency “panic” signal from an OPEN_CAPI device to the system that a catastrophic thermal condition is imminent and the device is going to shut down. 1.8 V level signal to motherboard. Signals for each OPEN_CAPI device gathered together and fed to the power sequencer. Care must be taken that this signal does not activate before PWR_EN and PWR_GOOD is valid. It is recommended that this signal be ANDED with PWR_EN at a minimum because PWR_GOOD is optional and pulled up when not used.
- INT_RST_N – Reset to OpenCAPI FPGA controlled by the processor I2C bus. This signal must have a pull up to 3.3 V on the mezzanine card.
- PRSNT_1A/2A/1B/2B – presence detect loop. Short PRSNT_1A to PRSNT_1B and short PRSNT_2A to PRSNT_2B on OPEN_CAPI device. Motherboard uses these to form a presence detect loop that is only ‘true’ when the card is fully and correctly seated. Signal is used by JTAG scan ring bypass logic and PCIe device prsnt detect logic.


3.6 Mezzanine PCIe interface

Each mezzanine card has an X2 PCI Express (PCIe) interface that can be used in many ways, with one being a sideband support. These buses are labeled PE0 and PE1 in the connector table. This interface is compliant with the PCIe base specifications.

4. Carrier card and advanced accelerator cable

The internal cable connecting the CPU board to the PCIe carrier card is shown in *Figure 4-1*. There are two connector slots on the CPU board for 25 Gbps $\times 8$ connection. One $\times 8$ cable is assumed for this guideline. If a vertical connector is selected, select a right-angle connector and place it at the end of the PCIe card, toward the CPU side, to avoid interference with the next PCIe card.

Figure 4-1. Carrier card connecting to CPU board



4.1 Carrier card size

The carrier card design should conform to the size requirements for PCIe add-in cards defined by the PCI Express® Card Electromechanical Specification Revision 3.0 or later. It is recommended that the carrier card be kept to the half or three-quarter lengths per the specification to better accommodate the Advanced Accelerator connector routing in the chassis.

4.2 Carrier card connector

Only power is being removed from the PCIe slot for the carrier card, so the size of the PCIe edge connector can be determined by the amount of card power required and the mounting stability and flexibility required. The X1 connector size provides full power and is most universal but might not provide the best retention and mechanical stability compared to a $\times 8$ connector.

4.3 Carrier card power

The power available from the PCIe edge connector is limited to 75 W. Additional power requirements may be addressed by the Auxiliary Power Connector options identified in the PCI Express® Card Electromechanical Specification.

4.4 Carrier card cooling

Refer to the PCI Express® Card Electromechanical Specification for information on cooling requirements for

Approved

the carrier card.

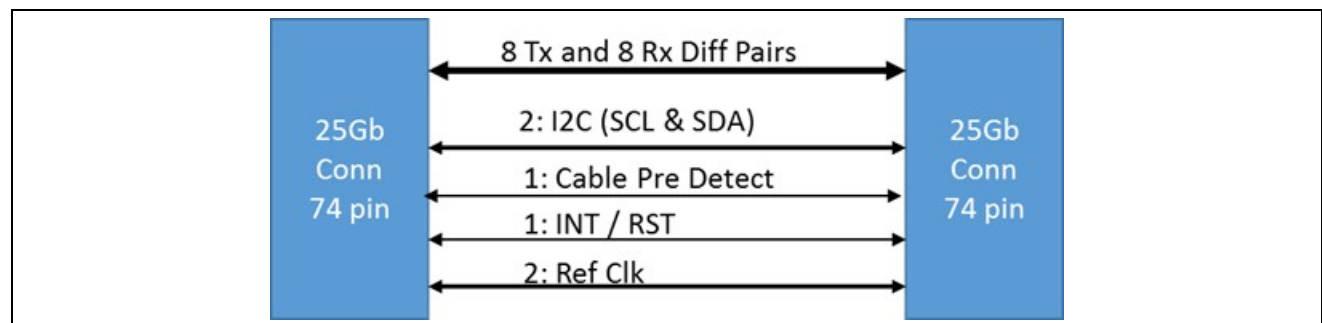
4.5 Advance accelerator cable

The OpenCAPI platform supports the optional 25 Gbps interface to the advance accelerator processor unit in a different drawer of the rack or the riser card plug-in to the PCIe slot in the same system. This section only contains information on topologies, connectivity, and routing guidelines for the advance accelerator cable (AAC) interface.

4.5.1 Accelerator cable circuit schematic

Figure 4-2 illustrates an accelerator cable circuit schematic.

Figure 4-2. Accelerator cable circuit schematic



4.5.2 Accelerator cable signal definition

The interconnect allows lane and polarity reversal. Pin swapping is not allowed. Signal direction is defined as follows:

- Tx – is from the Host (CPU) perspective
- Rx – is from the Host (CPU) perspective
- Tx from Amphenol SlimSAS connectors to Rx of FPGA on OpenCAPI adapter
- Rx from Amphenol SlimSAS connectors to Tx of FPGA on OpenCAPI adapter
- AC Coupling must be done on both the 8 Rx and 8 Tx differential data signals on the carrier card

The following sideband signals are present:

- Ref Clk – Differential reference clock. Driven by the CPU. 156.25 MHz for OIF applications or 133.33 MHz for JEDEC applications. This differential clock is HSICL. External reference clock electrical specification can be found in the *OpenCAPI 25 Gbps Physical Signaling Specification* (Table 5-1).
- I2C (SCL/SDA) – I2C bus clock and data. I2C on the carrier card is only a slave implementation. The CPU drives these signals to a 3.3 V level from the main 3.3 V power supply. Recommended to have level translators on the carrier card to translate to 1.8 V. Pulled high (1.8 V) on the carrier card. These pins must be at a steady state voltage before the OPEN_CAPI device is taken out of reset.
- INT/RST – Reset to OpenCAPI FPGA controlled by the processor I2C bus. The CPU drives these signals to a 3.3 V level from the main 3.3 V power supply. This signal must have a pull-up resistor on the host side. Level translators are recommended on the carrier card to translate to 1.8 V. This signal

Approved

must have a pull-up resistor on the carrier side as well after the level translation.

- Cable Pre-Detect – Presence detect. Cable Pre-Detect must have a 49.9 Ω pull-down to GND resistor located on the adapter card.
- SPARE_1/2 – Spare pins. May be left floating.

4.5.3 Advanced accelerator cable loss budget

Table 4-1 lists the advanced accelerator cable loss budget.

Table 4-1. Advanced accelerator cable loss budget

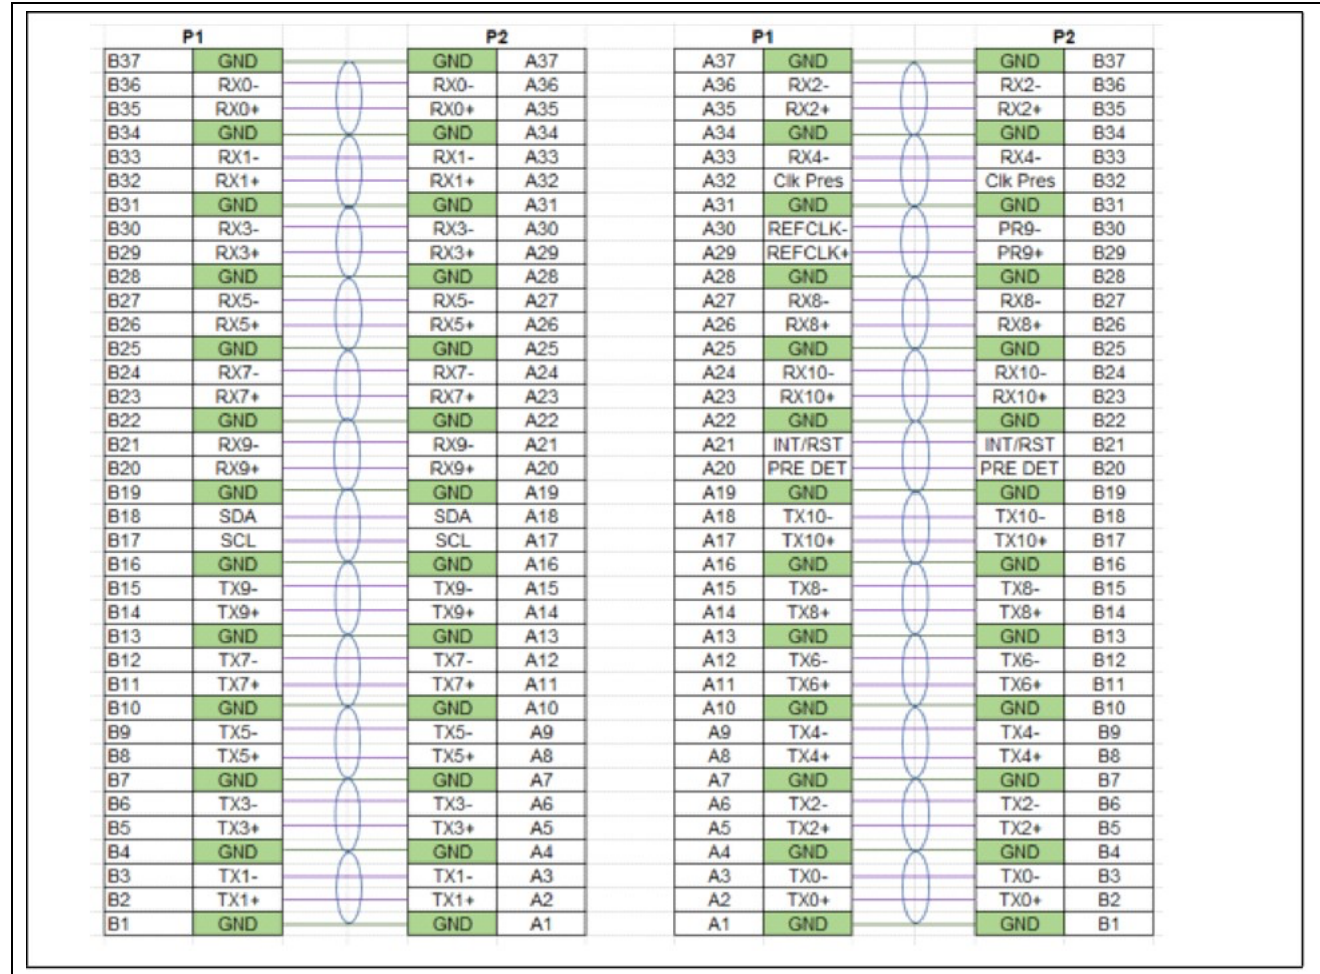
Component	Loss Budget (dB)	Notes
PCB connector	0.3	PCB connector
Cable connector	0.15	Connector on cable end
Cable	5.6	Cable
Cable connector	0.15	Connector on cable end
PCB connector	0.3	PCB connector
Total	6.5	Assembly

Approved

4.5.4 Accelerator-cable pin definition

Figure 4-3 illustrates the accelerator cable pin definition.

Figure 4-3. Accelerator cable pin definition

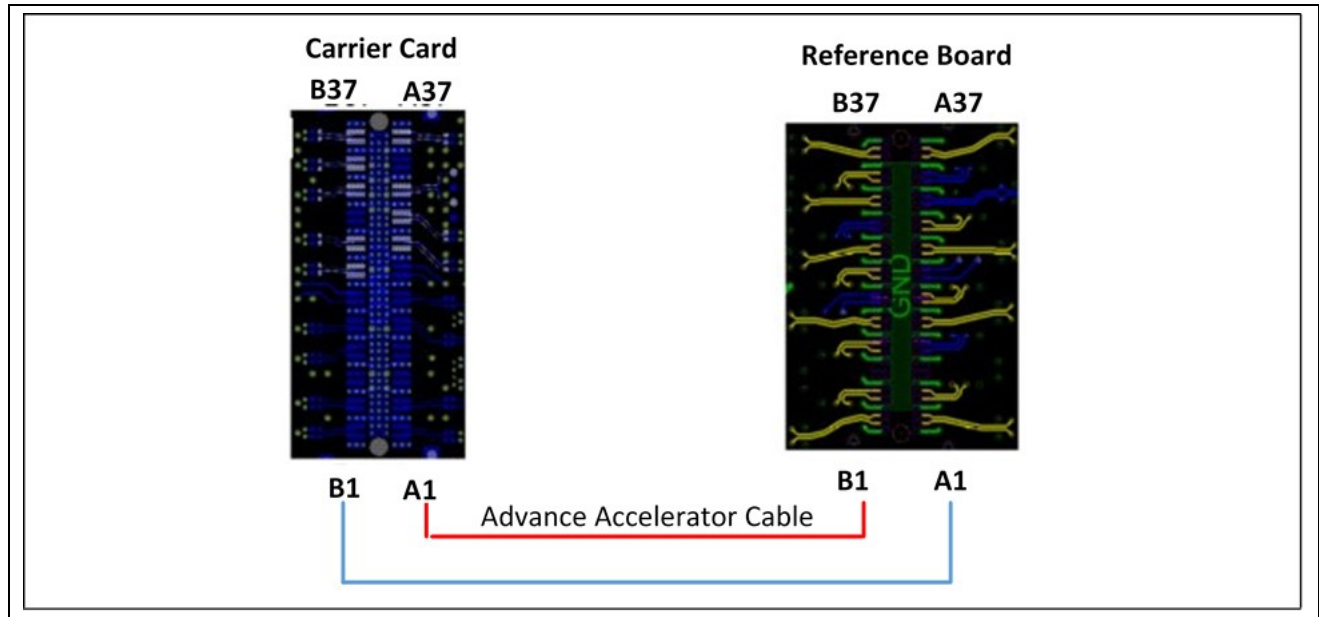


Approved

4.5.5 Internal cable pin ordering and mapping

The interconnect allows lane and polarity reversal. Pin swapping is not allowed. The GRD board pin ordering and list are shown in *Figure 4-4*. Also shown is a generic pin list for the cable connection.

Figure 4-4. GRD board pin ordering




Note: The raw wire within the cable is 12 inches, 30 AWG (6 dB/meter). The I/O assignments are referenced from the CPU board side.

4.5.6 AAC connector escape routing example

Figure 4-5 shows some AAC connector escape routing examples. The connector antipad goes through the top two GND plane layers. The signal must not cross the antipad void.

Connector break-out P and N: The connector break out P and N should be matched immediately to the first via.

Figure 4-5. AAC connector escape routing example



4.6 Cabled CPU to carrier card interconnect and loss budget

4.6.1 CPU-to-QSFP carrier card configuration

The end-to-end AAC interconnection consists of three channels: the reference design board, the AAC cable, and the carrier card as shown in Figure 4-6. The path includes the internal cable.

The board topology consists, from right-to-left: the processor module socket; the CPU board trace; vertical connector; 12 in., 30 AWG twin-ax cable; vertical and right-angle connector; carrier card trace; and the connector receptacle. Also shown is the loss-budget break down for each electrical component in the link.

The worst-case loss budget requirement is 20 dB. Table 4-2 and Table 4-3 show examples of end-to-end loss budget for the CPU-to-QSFP carrier card and the CPU-to-accelerator topologies.

Diagram illustrating the Carrier Card architecture:

- 25Gb Cable** connects to the **25Gb Connector** on the **Carrier Card**.
- The **Carrier Card** is connected to the **PCIe Slot**.
- The **Carrier Card** connects to the **CPU** via **74 Pin Right Angle Connectors** and **74 Pin Vertical Connectors**.
- An **Advanced Accelerator Cable** (12"; 30 AWG Twin-ax) connects the Carrier Card to the CPU.
- A dimension line indicates the **CPU through Carrier Card** path.

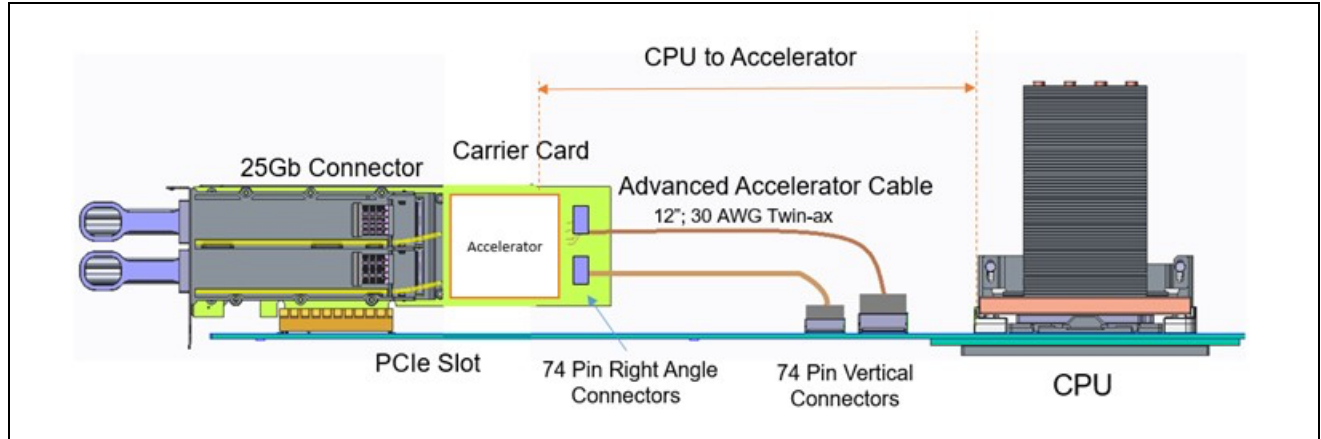
Component	Loss Budget (dB)	Notes
CPU socket	4.1	34 mm trace, LGA 7-2-7 laminate package
CPU board trace	5.1	3.2"; trace and via
Cable assembly	6.5	12.5" 30 AWG twin-axial cable (with connectors)
Carrier card (PCIe)	1.6	1" trace; blind via PCB
Cable plug	1.6	1.2" trace; blind via PCB
Worst-case PCB	0.9	Account for ± 10 PCB tolerance
Total	19.8	Receiver specification
Note: The worst-case loss budget requirement is 20 dB.		

Approved

4.6.3 CPU-to-accelerator configuration

Figure 4-7 illustrates a CPU-to-accelerator configuration.

Figure 4-7. Example of a CPU-to-accelerator configuration



4.6.4 CPU-to-accelerator loss budget

CPU ↔ Host board ↔ Cable ↔ Carrier card ↔ OpenCAPI module (for example, an:FPGA).

Table 4-3 an example of the CPU-to-accelerator loss budget.

Table 4-3. Example of CPU-to-accelerator loss budget (OpenCAPI module)

Components	Loss Budget (dB)	Notes
CPU socket	4.1	LGA plus 34 mm trace
CPU board trace	5.1	3.2"; trace and via
74-pin Vertical Connector	0.45	
Cable assembly	6.5	12.5" 30 AWG twin-ax cable
74-pin Vertical Connector	0.45	
Carrier card	2.5	2" trace, 0201 AC capacitor, blind via PCB
Advance accelerator socket/package	1.8	11 mm trace, laminate package
Total	20.9	Receiver specification

Note: The worst-case loss budget requirement is 20 dB.

4.7 Accelerator card design supplement


This section is intended to provide additional information that is necessary for accelerator card design.

Systems based on the host processor are designed so that each x8 OpenCAPI bus has a total card wiring skew budget (processor module die C4 to accelerator module die C4) of 10 UI. This skew budget is the same for both the Tx and Rx directions. When designing an accelerator card for an existing system, skew budget information is provided below for each of three topologies. Note that hardware manufacturing tolerances and PCB glass-weave skew effects should be considered as part of the allotted skew budget.

4.7.1 Topology 1

Figure 4-8 illustrates the direct-attach accelerator card. The MEG-array connector resides on the same planar with the host processor. Skew budget allowance for the accelerator card is 3.3 UI.

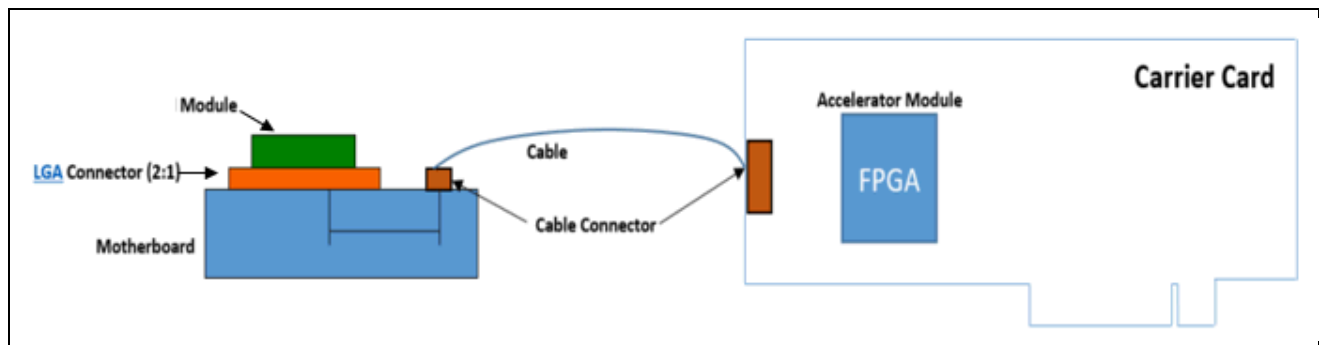
Figure 4-8. Example of a direct-attach accelerator card



4.7.2 Topology 2

Figure 4-9 illustrates a cable-attached accelerator carrier card. The CPU module and accelerator FPGA reside on different cards with a cable connection between them. The skew budget allowance for the accelerator carrier card is 3.3 UI and the skew budget allowance for the cable is 1.5 UI.

Figure 4-9. Example of a cable-attached accelerator carrier card



4.8 Carrier card stack-up

Figure 4-10 is an example stack-up for the PCIe carrier card that supports the advance accelerator cable interface. Skip or blind via technology is assumed. The signal should be on the top layers where the connector resides.

Figure 4-10. Example of a carrier card stack up

