Japan Industrial Imaging Association Guideline
日本インダストリアルイメージング協会 ガイドライン

JIIA CXPR-008-2023

CoaXPress over Fiber Bridge Protocol

Version 1.1
Enacted on 2023.11.17
2023 年 11 月 17 日 制定

The Standardization Committee—CoaXPress Working Group
Japan Industrial Imaging Association
標準化委員会 CoaXPress 分科会
一般社団法人 日本インダストリアルイメージング協会

This document is provided “as it is” with this current version.

None of JIIA, its members including contributions to the standard are also made by members of the G3 Group of Associations, the subsidiary companies of its members or the affiliates of its members are liable, whether express or implied, for any kinds of warranties of merchantability, fitness for a particular purpose, and non-infringement of third party property rights.

JIIA, its members including Contributions to the standard are also made by members of the G3 Group of Associations, the subsidiary companies of its members, or the affiliates of its members are exempted from taking responsibility and held harmless under the applicable law, whether express or implied, for any kinds of damages or losses arising from a course of usage of this document, including but not limited to passive damages, derivative damages, and incidental damages. This is applied even when any of JIIA, its members including Contributions to the standard are also made by members of the G3 Group of Associations, the subsidiary companies of its members or the affiliates of its members were informed of the possibilities of damages in advance.

JIIA, its members including Contributions to the standard are also made by members of the G3 Group of Associations, the subsidiary companies of its members or the affiliates of its members are exempted from taking responsibility and held harmless for any litigation, including but not limited to on defending, on cooperating, and on compensating, and infringement of intellectual property rights with the third party, incurred by the usage of this document.

CoaXPress is a registered trademark in Japan, the USA, the EU, and other countries and the property of the trademark belongs to Japan Industrial Imaging Association (JIIA). You need a trademark license to use the logo on any product or in commercial documents. If you want to use the trademark or have any questions about using trademarks, please contact a JIIA representative (sec@jiia.org).

この書面は“現状のまま”の状態で提供されます。
JIIA，またはG3加盟協会会員による規格化協力会員，会員の子会社もしくは会員の関係会社のいずれも，この書面の内容に関して，商品性，特定の目的への適合性，非侵害の保証を含め，いかなる保証も，明示たると黙示たるとを問わず一切行いません。
JIIA，またはG3加盟協会会員による規格化協力会員，会員の子会社もしくは会員の関連会社のいずれも，この書面の使用，または使用不能から生ずるいかなる損害（逸失利益，及びその他の派生的，または付随的な損害を含むが，これらに限定されない全ての損害を言います。）について，適用法で認められる限り，一切の責任を負わないものとします。
たとえJIIA，またはG3加盟協会会員による規格化協力会員，会員の子会社もしくは会員の関連会社がかかわる損害の可能性について知らされていた場合でも同様です。
JIIA，またはG3加盟協会会員による規格化協力会員，会員の子会社もしくは会員の関連会社のいずれも，この書面に起因して第三者との間に生じた，または生じうる知的財産権に関する紛争について，防御，協力，または補償する責任を負わないものとします。
CoaXPressは日本、アメリカ、EUその他の国に登録された商標であり、この商標権は一般社団法人日本インダストリアルイメージング協会に帰属します。製品または広報広告文書に使用する場合には使用許諾が必要となります。商標の使用を望まれる場合もしくは使用に関し疑義がある場合にはJIIA担当(sec@jiia.org)まで連絡をお願いします。

TABLE OF CONTENTS

1 Scope

The CoaXPress over Fiber Bridge Protocol was developed to allow employing optical fiber cables as transmission medium for the CoaXPress protocol. This document is intended to be used as a supplement to the existing CoaXPress standard.

2 References and Definitions of Terms

2.1 Normative References

The following referenced documents are essential for the application of this document. The latest edition of the referenced document (including any amendments) applies.

- Ref 1 IEEE 802.3: IEEE Standard for Ethernet
- Ref 2 JIIA CXP-001: CoaXPress Standard
- Ref 3 JIIA CXPR-007: Optical interface Guideline for CoaXPress

2.2 Definitions of Terms

8B/10B IDLE	Refers to the CoaXPress 8B/10B encoded idle word formed by the characters K28.5, K28.1, K28.1, and D21.5.
nGMII	Refers to both XGMII and 25GMII (see section 4.1).
nGMII IDLE	Refers to XGMII and 25GMII idle control characters.

2.3 Abbreviations

25GMII	25 Gigabit Media Independent Interface
CoF	CoaXPress over Fiber
CXP	CoaXPress
DW	Data Word, 4 bytes
EOP	End of Packet
FEC	Forward Error Correction
Gbps	Gigabit per second, 10^9 bits per second
HDP	High-Speed Data Payload
HKP	High-Speed K-Code Payload
IPG	Interpacket Gap
IT	Idle Transfer
LSP	Low-Speed Payload
MAC	Media Access Control
Mbps	Megabit per second, 10^6 bits per second
MDI	Medium Dependent Interface

PCIe	Peripheral Component Interconnect Express
PCS	Physical Coding Sublayer
PHY	Ethernet Physical Layer
PMA	Physical Medium Attachment
PMD	Physical Medium Dependent
RS-FEC	Reed-Solomon Forward Error Correction
RXC	Receive Control
RXD	Receive Data
SerDes	Serializer/Deserializer
SOP	Start of Packet
TXC	Transmit Control
TXD	Transmit Data
XGMII	10 Gigabit Media Independent Interface

2.4 Conventions

“Shall” means a mandatory requirement.

“Can” means an optional feature.

Comments written in italic text do not form part of the specification but are intended to help understand the requirements of the specification.

3 Introduction

CoaXPress is a successful interface protocol that is being largely used in the machine vision industry to connect high-end cameras to PCIe frame-grabbers. The main features that make CoaXPress appealing for this sector are its high bandwidth, highly deterministic triggers, long cable lengths, simplicity and robustness of cables and connectivity.

To meet the continuously increasing requirements of the machine vision industry, CoaXPress must cope with even higher bandwidths and longer cable lengths. Meeting these requirements using coaxial cables is unlikely to be feasible. Optical fiber cables are the most viable existing transmission medium that supports these requirements.

Optical fiber cables are massively used in Ethernet applications. Cables and peripheral components supporting very high bandwidths and extreme long distances have become accessible in terms of price and availability.

FPGAs already in use in CoaXPress applications include many features to support the Ethernet standard. This creates a natural path to employ optical fiber cables as transmission medium for the CoaXPress protocol.

The 10 Gigabit Media Independent Interface (XGMII), defined by IEEE Std 802.3 Clause 46, is the main access to the 10G Ethernet physical layer. The generic nature of this interface facilitates mapping the CoaXPress signaling into the PCS/PMA Ethernet sublayers. In addition, an XGMII interface has an effective bit rate equivalent to a CoaXPress CXP-12 connection without its 8B/10B encoding overhead. Slower CoaXPress bit rates can also be mapped into an XGMII interface by extending the spacing between two packets, known as the interpacket gap (IPG).

The use of the XGMII also allows CoaXPress to be ready to support higher bandwidths via the 25 Gigabit Media Independent Interface (25GMII), defined by IEEE Std 802.3 Clause 106. The 25GMII is a speeded-up version of the XGMII while being logically equivalent to it.

This document defines the CoaXPress over Fiber concept based on a physical layer topology and on a series of rules to map the CoaXPress protocol to the interface of the Ethernet physical layer. Based on these rules, a CoaXPress to Ethernet Physical Layer bridge (CXP-PHY bridge) is specified to allow the CoaXPress protocol to be transported over available Ethernet components without the need to establish new fiber components for the CoaXPress standard. The CXP-PHY bridge specification targets the following requirements:

- The CXP-PHY bridge must be transparent to the adjacent CoaXPress layer (Host or Device).
- The CXP-PHY bridge must (as much as possible) be independent from CoaXPress version variations (current and future ones).
- The CXP-PHY bridge must not perform any CoaXPress protocol decoding during mapping. Only K-code detection and high-priority packets identification are required.

For additional information regarding fiber optical connectors and cables, please refer to the Optical Interface Guideline for CoaXPress (see section 2.1 , Ref 3).

4 Ethernet – IEEE Std 802.3

The IEEE Standard for Ethernet (see section 2.1 , Ref 1) addresses Layer 1 (Physical) and Layer 2 (Data Link) of the Open Systems Interconnection model (OSI model) specified on ITU-T X.200.

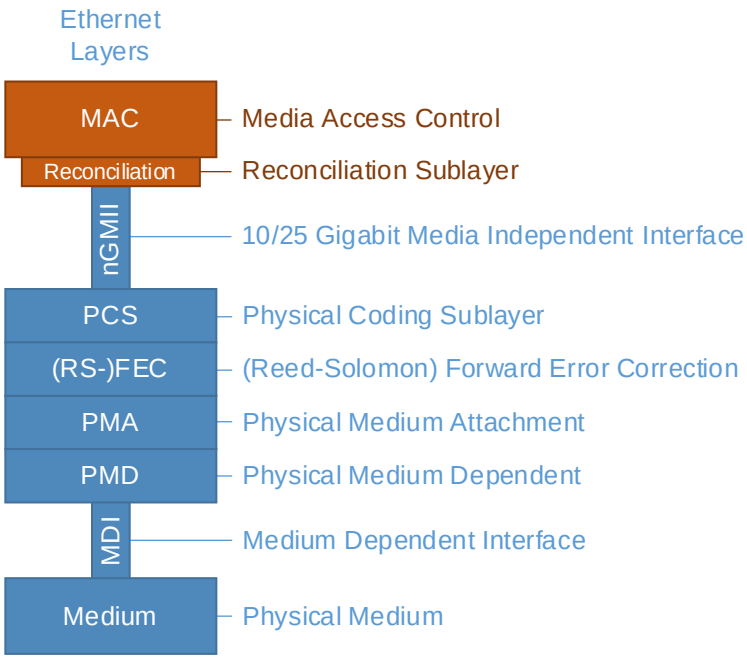

Figure 1 - Ethernet model for 10G and 25G BASE-R

To access the PCS/PMA sublayers, the CXP-PHY bridge replaces the Reconciliation Sublayer in the Ethernet model.

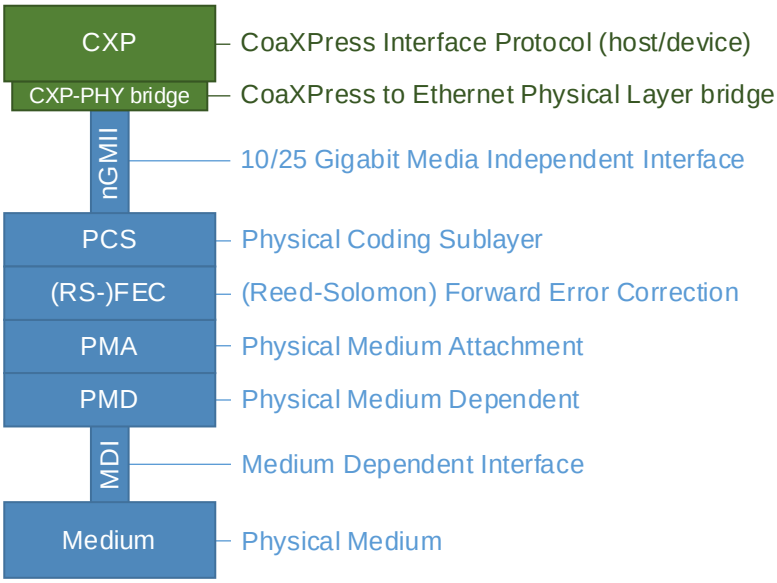

Figure 2 - Ethernet model in the CoaXPress over Fiber context

4.1 10G and 25G Media Independent Interfaces

The 10 Gigabit Media Independent Interface (XGMII) is defined by IEEE802.3 Clause 46 and provides an interconnection between the Media Access Control (MAC) and the 10G Ethernet Physical Layer (PHY).

The 25 Gigabit Media Independent Interface (25GMII) is a speeded-up version of the XGMII (2.5x faster) while being logically equivalent. It is defined by IEEE802.3 Clause 106 and provides an interconnection between the MAC and the 25G PHY.

XGMII TX/RX clock frequency is specified to be 156.25 MHz (+/- 0.01%) at DDR (Double Data Rate). In practice, the available FPGA PCS/PMA IPs work at SDR (Single Data Rate) at 312.5 MHz or with two aggregated interfaces at 156.25 MHz.

25GMII TX/RX clock frequency is specified to be 390.625 MHz (+/- 0.01%) at DDR. In practice, the available FPGA PCS/PMA IPs work at SDR with two aggregated interfaces at 390.625 MHz.

As both interfaces, XGMII and 25GMII, present the same encoding mechanism and identical structure of data and control paths, in this document they will be referenced as a single interface named *nGMII*.

The nGMII is composed of two independent 32-bit data paths (TXD and RXD) each one associated to a 4-bit control path (TXC and RXC). Each data path is organized into four 8-bit lanes.

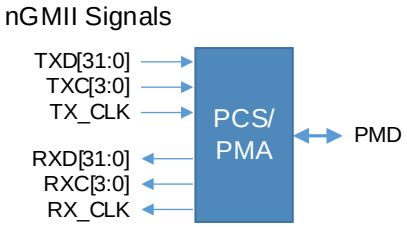

Figure 3 - nGMII signals

A nGMII transfer contains four characters where each character is formed by 8 bits of data and 1 bit of control. In a nGMII transfer, the first character is aligned to lane 0, the second character is aligned to lane 1, and so on.

Table 1 - nGMII data and control paths

TXD/RXD	TXC/RXC	Lane
[7:0]	[0]	0
[15:8]	[1]	1
[23:16]	[2]	2
[31:24]	[3]	3

The permissible lane encodings are:

Table 2 - nGMII permissible lane encodings

TXC/RXC	TXD/RXD	Notation	Description
0	0x00 through 0xFF	/D/	Normal data transfer
1	0x06	/LI/	Only valid on all four lanes simultaneously to request/indicate Low Power Idle (LPI)
1	0x07	/I/	Idle (nGMII IDLE)
1	0x9C	/Q/	Sequence (only valid in lane 0)
1	0xFB	/S/	Start (only valid in lane 0)
1	0xFD	/T/	Terminate
1	0xFE	/E/	Error
1	Others	-	Reserved

nGMII packets shall respect the following list of rules to be considered valid. These rules are based on the IEEE802.3 Clauses 46 (XGMII) and 106 (25GMII).

- nGMII Rule #1: /S/ is only valid at lane 0.
- nGMII Rule #2: /T/ is valid in any lane while the following lanes in the transfer must present /I/ characters.
- nGMII Rule #3: The minimum IPG at the nGMII is five characters.
- nGMII Rule #4: Between a /S/ and a /T/, only /D/ characters are considered valid.
- nGMII Rule #5: /LI/, /Q/, and /I/ characters are only allowed within the IPG.

4.2 Forward Error Correction

Forward Error Correction is a method to correct errors in a serial link by inserting additional information (check bits) in the transmitter data stream while the receiver decodes this information to verify the correctness of the data and to correct possible errors. This method improves the Bit Error Rate (BER) performance of a link without requiring a retransmission of the corrupted data.

The IEEE802.3 defines different types of FEC depending on the type of PHY implemented. For 10G and 25G BASE-R PHYs, the IEEE802.3 defines the respective FEC types in Clauses 74 and 108.

4.2.1 IEEE802.3 Clause 74

The IEEE802.3 Clause 74 defines the FEC sublayer for 10GBASE-R PHYs. The FEC Clause 74, also known as *Firecode FEC* or *BASE-R FEC*, can correct burst errors of up to 11 bits per block of 32 64B/66B blocks.

Note: In the context of the CoaXPress over Fiber, the implementation of IEEE802.3 Clause 74 is mandatory in 10G applications.

4.2.2 IEEE802.3 Clause 108

The IEEE802.3 Clause 108 defines the Reed-Solomon FEC (RS-FEC) sublayer for 25GBASE-R PHYs. The RS-FEC can correct up to 7 symbols per FEC block.

Note: In the context of the CoaXPress over Fiber, the implementation of IEEE802.3 Clause 108 is mandatory in 25G applications.

5 CoaXPress - Link structure and protocol

This section summarizes the aspects of the CoaXPress protocol that are essential to define the mapping mechanism between CoaXPress and Ethernet PHY.

Note: It is not the intention of this section to reproduce all details of the CoaXPress standard. It is strongly recommended to refer to the latest review of the CoaXPress standard for details (see section 2.1 , Ref 2).

Note: All definitions, figures, and tables presented in this section are adapted excerpts from the CoaXPress 2.1 standard (JIIA CXP-001-2021) and may differ from other versions of the standard.

5.1 Link overview

A CoaXPress link is defined as a point-to-point interface composed of one or more connections. Each connection consists of a high-speed downconnection (Device to Host) and a low-speed upconnection (Host to Device). Optionally, a high-speed upconnection can be used to increase the control channel bandwidth and the maximum trigger rate supported.

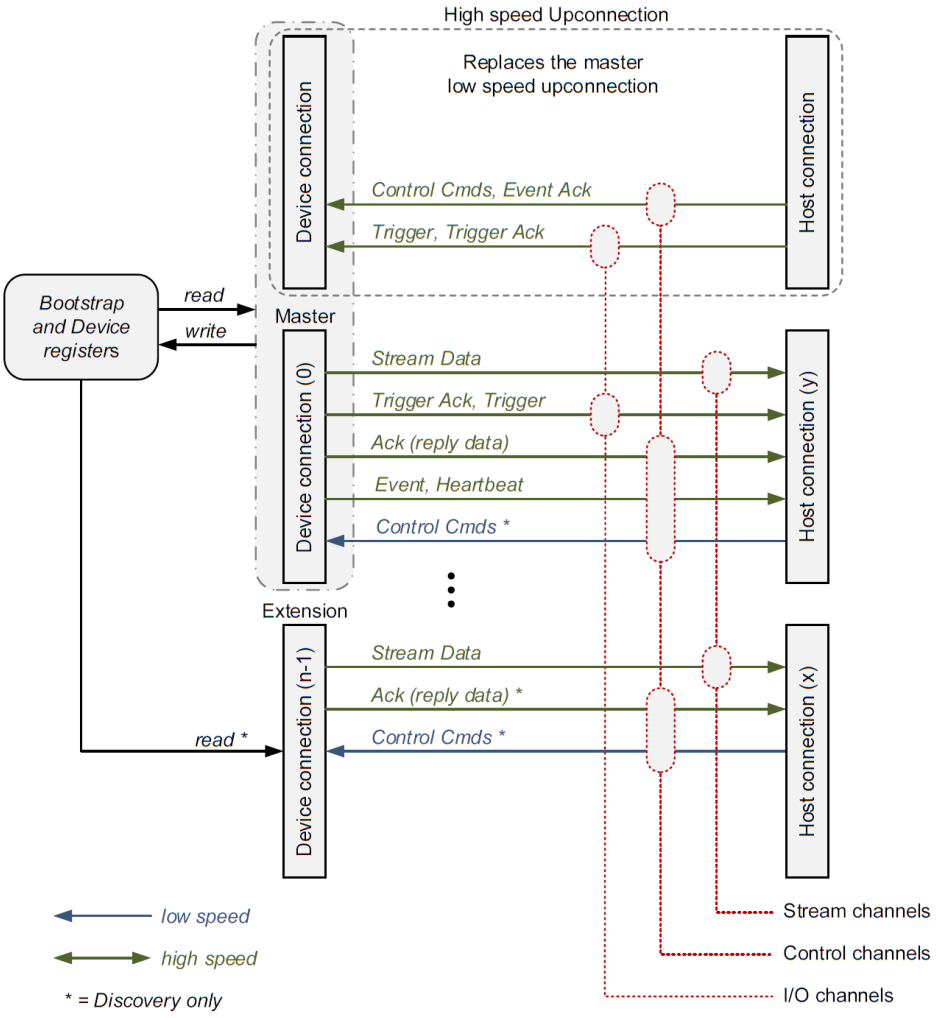

Figure 4 - CoaXPress link overview

Comment: The CoaXPress 2.0 standard (JIIA CXP-001-2019) does not support high-speed upconnection in contrast with other versions of the standard.

Comment: While for coaxial media the high-speed upconnection was costly to be implemented on both Hosts and Devices, for optical media it comes virtually with no cost since Ethernet-compliant components are typically symmetrical with up and downconnections working at the same bitrate.

5.2 Data organization

CoaXPress specifies two different data organizations: one for high-speed connections and one for low-speed connections. For high-speed connections, data transmission is organized in groups of four characters called a word. Each word consists of four consecutive 8B/10B characters labeled P0, P1, P2 and P3. For low-speed connections, data transmission is organized in single 8B/10B characters.

5.3 Packet transmission priority

CoaXPress defines three levels of packet transmission priorities specifying that a higher-priority packet can be inserted into a lower-priority one. The transmission of the lower-priority packet must be halted while the higher-priority packet is introduced in the data stream. The transmission of the lower-priority packet is resumed as soon as the higher-priority packet transmission is finished.

Table 3 - CoaXPress packet transmission priority

Priority Level	Packet type
0 (highest)	Trigger
1	I/O acknowledgment (for trigger packet)
2 (lowest)	All other packets

5.4 Idle transmission

CoaXPress specifies an 8B/10B encoded idle word (8B/10B IDLE word) that must be sent in both high-speed and low-speed connections when the respective connection is idle.

Table 4 - CoaXPress 8B/10B IDLE word format

Word	P0	P1	P2	P3
IDLE	K28.5	K28.1	K28.1	D21.5

8B/10B IDLE words are also regularly inserted in the data stream to guarantee word alignment and equalizer reference for the coaxial cable transceiver. In high-speed connections, 8B/10B IDLE words are sent at least once every 100 words and they can be embedded into any CoaXPress packet, except in trigger packets and I/O acknowledgement packets. In low-speed connections, they are sent at least once every 10,000 words and can also be embedded into any CoaXPress packet, except in trigger packets and I/O acknowledgement packets.

CoaXPress also permits stretching packets (other than trigger or I/O acknowledgment) in high-speed connections by inserting 8B/10B IDLE words. In low-speed connections, stretching packets is not allowed.

5.5 Packet format

The CoaXPress protocol employs four different packet formats to exchange data between Host and Device. Each one of these packet formats uses one or more 8B/10B K-code characters to indicate, for instance, the beginning or the end of a packet.

A summarized description of these packet formats is presented below in order to highlight their structure. For a detailed description of the CoaXPress packet formats, please refer to the CoaXPress standard.

5.5.1 Low-speed connection trigger

The packet format for low-speed connection triggers is only used in the low-speed upconnection.

Table 5 - CoaXPress low-speed connection trigger packet format

Char	Content	Description
0 – 2	*Either* K28.2 K28.4 K28.4	Trigger packet indication – LinkTrigger0
	or K28.4 K28.2 K28.2	Trigger packet indication – LinkTrigger1
3 – 5	3x Delay	Delay value between the trigger event and its transmission.

Comment: The CoaXPress 2.1 standard also defines an additional low-speed trigger packet format that repurposes the LSB of the delay value to enable two extra triggers, LinkTrigger2 and LinkTrigger3. As this additional packet format has a similar structure and character count as the one presented in Table 5, the CoaXPress over Fiber specification supports both packet formats.

5.5.2 High-speed connection trigger

The packet format for high-speed connection triggers consists of three words where each word contains one character replicated four times for redundancy. This packet format starts with a K28.2 character.

Table 6 - CoaXPress high-speed connection trigger packet format

Word	Content	Description
0	4x K28.2	Trigger packet indication
1	4x Delay	Delay value between the trigger event and its transmission.
2	4x LinkTriggerN	A number N between 0 and 15 which defines the trigger, corresponding to LinkTrigger0 to LinkTrigger15.

Comment: The current CoaXPress over Fiber specification does not support high-speed trigger packets from CoaXPress 1.x, which are formed by two words instead of three. The support of this packet format from CoaXPress 1.x was not a requirement for CoaXPress over Fiber since no products supporting this feature were registered with JIIA.

5.5.3 I/O acknowledgment

The packet format for I/O acknowledgment consists of two words where each word contains one character replicated four times for redundancy. This packet format starts with a K28.6 character.

Table 7 - CoaXPress I/O acknowledge packet format

Word	Content	Description
0	4x K28.6	I/O acknowledgment packet indication
1	4x Code	0x01 Trigger packet received OK

5.5.4 Data packet transmission

The data packet transmission format is used to exchange data between Device and Host. Typically, this packet format is predominant in terms of bandwidth utilization. It consists of three words where each word contains one character replicated four times for redundancy. Additional *D* words carry the data payload. The packet itself contains two K-code characters to mark start and end of packet (K27.7 and K29.7 respectively). Stream data packets (data packets with Type = 0x01) can also contain one or more K28.3 characters corresponding to stream markers.

Table 8 - CoaXPress data packet transmission format

Word	Content	Description
0	4x K27.7	Start of packet indication
1	4x Type	Data packet type
2… *D*+1	Data (4 bytes)	Data packet payload of *D* data words
D+2	4x K29.7	End of packet indication

5.6 Support of 25G fiber bit rate

The CoaXPress standard defines bit rates up to 12.500 Gbps (CXP-12). These bit rates correspond to the speeds supported by the serial data technology available for coaxial cables that guarantee a reliable link with a cable length going up to 200+ meters. With the introduction of CoaXPress over Fiber, Ethernet

compliant components can be used to go beyond the CoaXPress defined bit rates and cable lengths as in the case of the 25GMII (see section 4.1). To support the 25GMII, CoaXPress over Fiber defines the CXP Speeds *CXP-25 and CXP-31*. A comparison of the high-speed connection bit rates at the transmission media for each supported interface is indicated in the following table:

Table 9 – Supported high-speed connection bit rates

CXP Speed	Effective Bit Rate (Gbps)	Coaxial – 8B/10B Bit Rate (Gbps)	10G Fiber – 64B/66B Bit Rate (Gbps)	25G Fiber – 64B/66B Bit Rate (Gbps)
CXP-1	1.000	1.250	10.3125	25.78125
CXP-2	2.000	2.500	10.3125	25.78125
CXP-3	2.500	3.125	10.3125	25.78125
CXP-5	4.000	5.000	10.3125	25.78125
CXP-6	5.000	6.250	10.3125	25.78125
CXP-10	8.000	10.000	10.3125	25.78125
CXP-12	10.000	12.500	10.3125	25.78125
CXP-25	20.000	25.000	N/A	25.78125
CXP-31	25.000	31.250	N/A	25.78125

Comment: The bit rates for 10G and 25G Fiber are always the nominal for the transmission media and do not change in function of the selected CXP Speed (see section 9).

Comment: CXP-12 has the same effective bit rate of 10 Gbps for Coaxial and 10G Fiber media when we exclude the encoding overhead (8B/10B and 64B/66B respectively).

Comment: CXP-31 has the same effective bit rate of 25 Gbps for Coaxial and 25G Fiber media when we exclude the encoding overhead (8B/10B and 64B/66B respectively).

To select the connection speeds CXP-25 and CXP-31 at the Device via the bootstrap registers *ConnectionConfig* and *ConnectionConfigDefault*, CoaXPress over Fiber defines an extension of the bit rate codes of the *Connection speed [15:0]* field as indicated in the following table:

Table 10 – Bit rate codes

CXP Speed	Bit Rate Code
CXP-1	0x28
CXP-2	0x30
CXP-3	0x38
CXP-5	0x40
CXP-6	0x48
CXP-10	0x50
CXP-12	0x58
CXP-25	0x70
CXP-31	0x78

6 CXP-PHY bridge specification

The CoaXPress-to-Ethernet Physical Layer bridge (CXP-PHY bridge) specification was designed to guide the development of a CXP-PHY bridge IP. The CXP-PHY bridge IP is intended to be placed between the CoaXPress cores (Device or Host) and the Ethernet PHY IP in order to map the CoaXPress protocol into the nGMII interface.

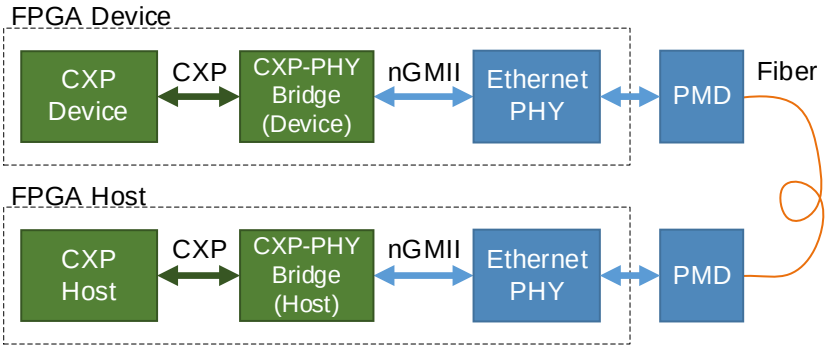

Figure 5 - CoaXPress over Fiber link overview

6.1 CoF physical layer topology

In terms of physical layer topology, the CXP-PHY bridge targets an asymmetrical configuration where the CoaXPress master connection requires two fibers (one upconnection and one downconnection) and any other extension connection only requires one fiber (downconnection).

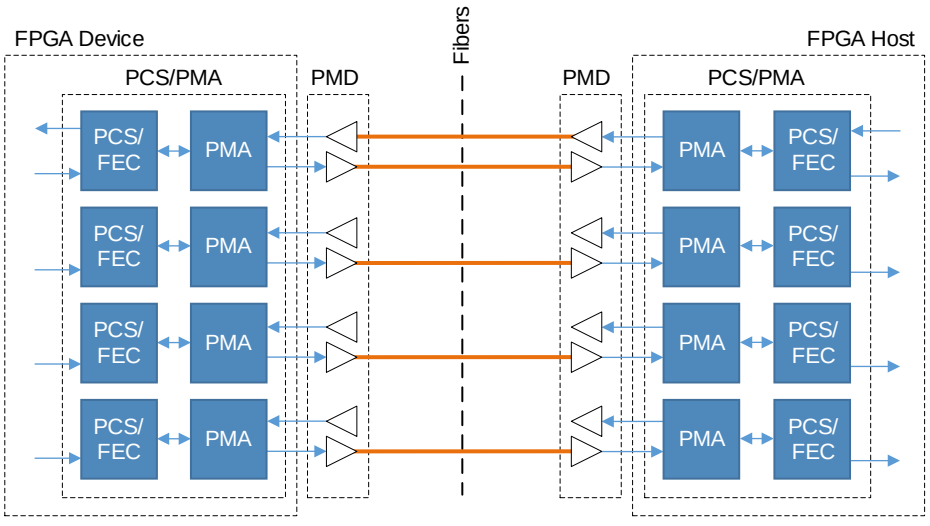

Figure 6 - CoF physical layer topology for a CoaXPress link with 4 connections

6.2 Bridge structure

To support the CoaXPress link specificities, two types of CXP-PHY bridge are specified, one for the Host and the other for the Device. In a CoaXPress multi-connection link, one CXP-PHY bridge is required per connection. Due to the asymmetrical nature of the targeted physical layer topology, the upconnection

mapping section in both bridge types (Host and Device) are only required at the CoaXPress master connections. Other connections only require the downconnection mapping section.

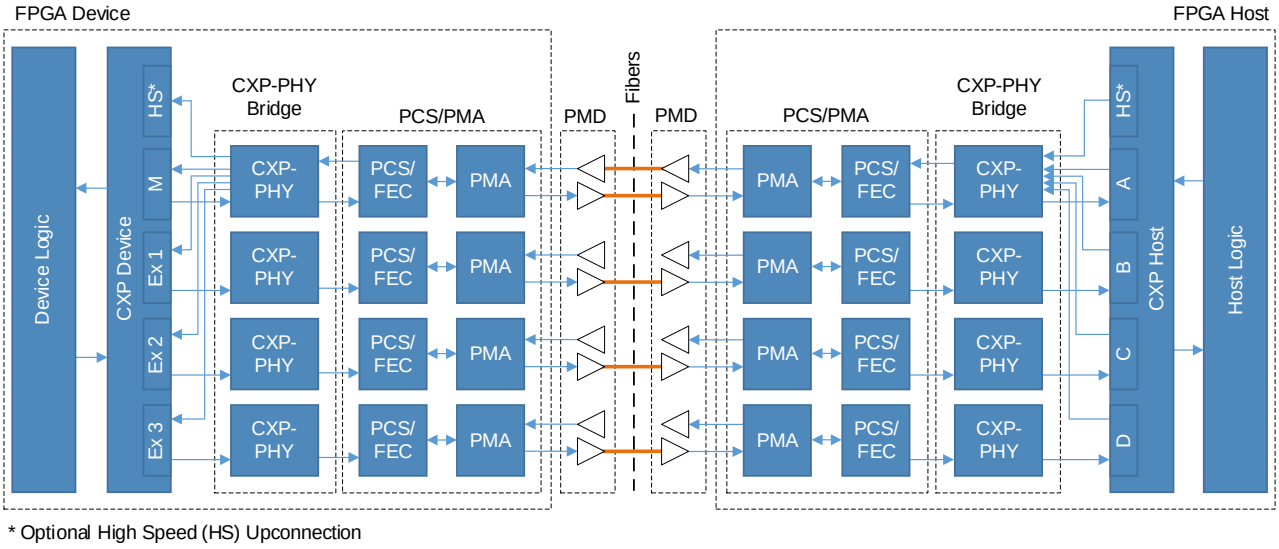

Figure 7 - CoaXPress over Fiber link via the CXP-PHY Bridge

Comment: It is important to highlight that both Host Logic and Device Logic must be dimensioned to support the maximum throughput of the Ethernet PHY to avoid the overflow of their internal coupler FIFOs.

6.2.1 CXP-PHY Bridge - Host

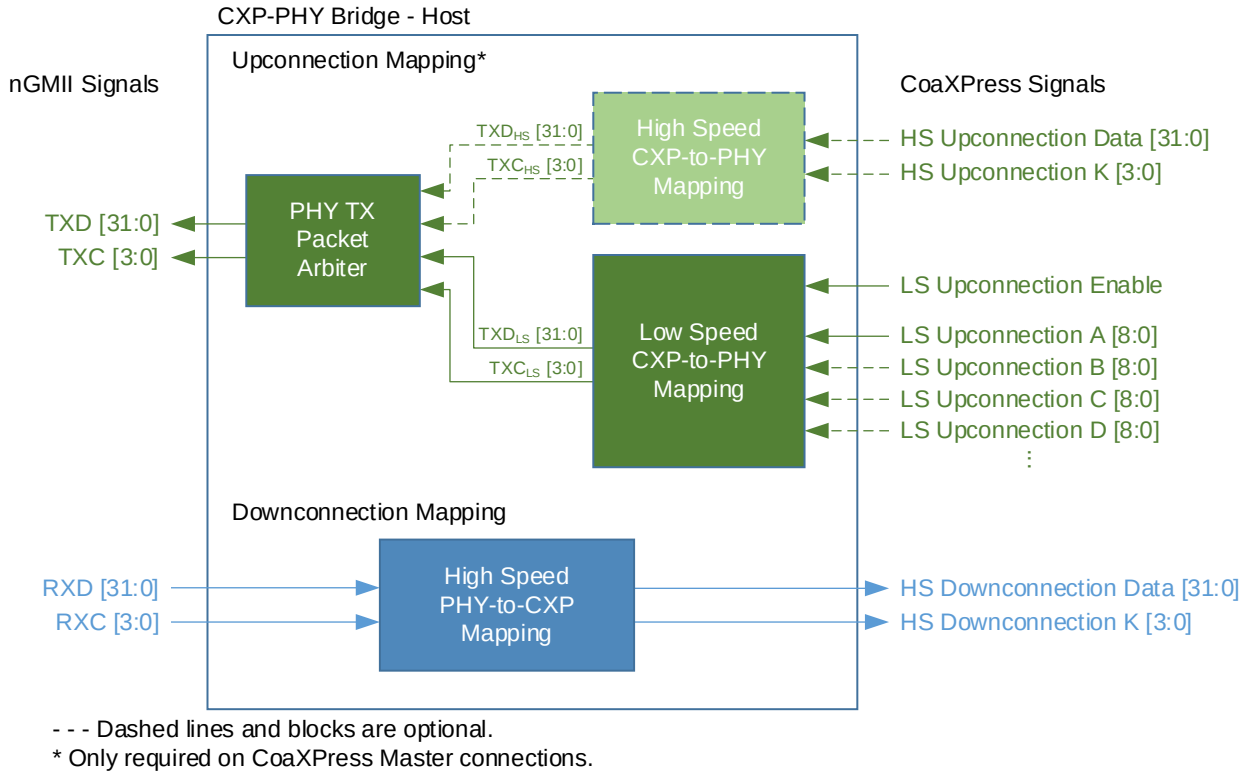

Figure 8 - CXP-PHY Bridge type Host

The CXP-PHY bridge type Host is divided into two parts, one for the CoaXPress downconnection mapping and the other for the CoaXPress upconnection mapping. The downconnection mapping part maps the nGMII RX signals into a CoaXPress high-speed downconnection compliant interface. The upconnection mapping part maps one or more CoaXPress low-speed upconnections and optionally a high-speed upconnection into the nGMII TX signals. The upconnection mapping part is only required on the CXP-PHY bridge connected to a Device master connection.

Comment: Hosts supporting multi-Device topologies must have multi-instances of the upconnection mapping mechanism, one for each CXP-PHY bridge connected to a Device master connection.

At the CoaXPress side, the high-speed upconnection and downconnection ports are composed of four 8-bit data signals and their corresponding K-code flags. The low-speed upconnection port is composed of a signal set (one 8-bit data signal and its K-code flag) for each low-speed upconnection. Due to the difference in bandwidth between the CoaXPress low-speed upconnection and the nGMII interface, an additional *Enable* signal is required to indicate when the data is available at this port.

The PHY TX Packet Arbiter maps one of its input signal sets to the nGMII TX signals. The arbitration mechanism is straightforward since CoaXPress does not allow a Host to use high-speed and low-speed upconnections at the same time, i.e., there are no conflicts between high-speed and low-speed packets.

Note: The CoaXPress signal set presented here is only an example to illustrate the CoaXPress interface. To be compatible with existing CoaXPress cores, additional control signals such as a SerDes *bit-clock lock*, may be necessary.

Note: The CoaXPress low-speed upconnection interface may require to deserialize and decode 8B/10B characters to be compatible with existing CoaXPress cores.

6.2.2 CXP-PHY Bridge – Device

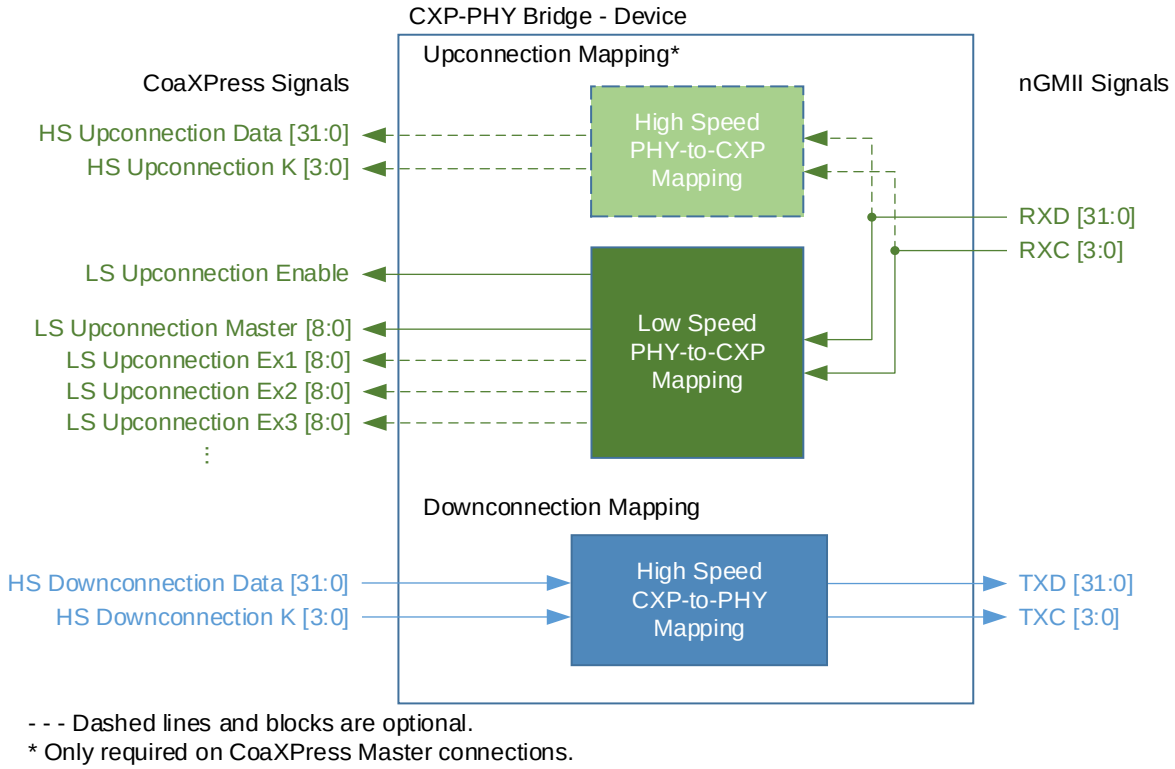

Figure 9 - CXP-PHY Bridge type Device

The CXP-PHY bridge type Device is divided into two parts, one for the CoaXPress downconnection mapping and the other for the CoaXPress upconnection mapping. The downconnection mapping part maps the CoaXPress high-speed downconnection into a nGMII TX compliant interface. The upconnection mapping part maps the nGMII RX signals into one or more CoaXPress and low-speed upconnections and optionally to a high-speed upconnection. The upconnection mapping part is only required on the CXP-PHY bridge connected to the Device master connection. No arbitration on the nGMII RX signals is required since the low-speed and high-speed PHY-to-CXP mapping modules only process their corresponding packet types.

At the CoaXPress side, the high-speed upconnection and downconnection ports are composed of four 8-bit data signals and their corresponding K-code flags. The low-speed upconnection port is composed of a signal set (one 8-bit data signal and its K-code flag) for the master upconnection and one signal set for each existing extension connection. Due to the difference in bandwidth between the CoaXPress low-speed upconnection and the nGMII interface, an additional *Enable* signal is required to indicate when the data is available at this port.

Note: The CoaXPress signal set presented here is only an example to illustrate the CoaXPress interface. To be compatible with existing CoaXPress cores, additional control signals such as a SerDes *bit-clock lock*, may be necessary.

Note: The CoaXPress low-speed upconnection interface may be required to encode the output data to 8B/10B characters and then serialize them to be compatible with existing CoaXPress cores.

6.3 Packet structure

The words that constitute the CXP-PHY packet stream are briefly described in the following table:

Table 11 - Word types of the CXP-PHY packet

Word Type	Notation	Size (characters)	Description
Start of Packet	SOP	4	Marks the start of a CXP-PHY packet
End of Packet	EOP	8	Marks the end of a CXP-PHY packet
Idle Transfer	IT	4	Interpacket gap (IPG)
High-speed Data Payload	HDP	4	High-speed payload word containing 4 bytes of data
High-speed K-code Payload	HKP	4	High-speed payload word containing 4 K-code characters
Low-speed Payload	LSP	4	Low-speed payload word containing 2 bytes of control and 2 bytes of data and/or K-code characters

All CXP-PHY packets are formed as SOP + *Payload* + EOP.

All CXP-PHY packets shall contain at least one *Payload* word (HDP, HKP, or LSP).

The CXP-PHY bridge uses two types of packets in nGMII stream: high-speed and low-speed packets.

High-speed packets are used to transmit data from a *High-Speed CXP-to-PHY Mapping* module to a *High-Speed PHY-to-CXP Mapping* module. The payload of CXP-PHY high-speed packets is formed by one HKP and/or one or more HDP words. An HKP word is always preceded by a SOP and followed by an HDP or an EOP.

Low-speed packets are used to transmit data from a *Low-Speed CXP-to-PHY Mapping* module to a *Low-Speed PHY-to-CXP Mapping* module. The payload of CXP-PHY low-speed packets is formed by one or more LSP words.

Transfer	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**
Word	IT	IT	SOP	HKP	HDP	HDP	EOP		IT	IT
	<== IPG ==>			<===== Payload ====>					<== IPG ==>	
			<=============== Packet ================>							

Figure 10 - Example of a CXP-PHY high-speed packet

Transfer	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**
Word	IT	IT	SOP	LSP	LSP	EOP		IT	IT	IT
	<== IPG ==>			<= Payload =>				<=== IPG ===>		
			<=========== Packet ============>							

Figure 11 - Example of a CXP-PHY low-speed packet

A description of all valid words in a CXP-PHY packet is listed below.

Note: In a CXP-PHY packet, all reserved or not in use fields in a word shall be set to “0”.

6.3.1 SOP – Start of Packet

The SOP word marks the start of a CXP-PHY packet at the nGMII interface for both up and down connections.

Table 12 – Start of Packet (SOP) word format

Word	Lane	TXC/RXC	TXD/RXD
0	0	1	[7:0] Start
	1	0	[7:0] SopCtrl
	2	0	[7:0] SopData0
	3	0	[7:0] SopData1

The SOP fields are described below.

6.3.1.1 Start

Start[7:0] - Start control character respecting nGMII Rule #1.

0xFB => nGMII control character /S/

6.3.1.2 SopCtrl

SopCtrl[7] - Packet type

When “0” => Low-speed packet (see section 8 Low-speed mapping for details)

When “1” => High-speed packet (see section 7 High-speed mapping for details)

SopCtrl[6:4] - Reserved

SopCtrl[3:0] - Control field dependent on the packet type (SopCtrl[7])

For low-speed packets (SopCtrl[7] = “0”), the Control field informs the upconnection status:

SopCtrl[3] - Update flag

When “0” => No update

When “1” => Upconnection status (SopCtrl[1:0]) updated

SopCtrl[2] - Reserved

SopCtrl[1] - Low-speed rate

When “0” => 20.83˙ Mbps[1]

When “1” => 41.6˙ Mbps[2]

SopCtrl[0] - High-speed upconnection state

When “0” => Deactivated

When “1” => Activated

This control field is used in the upconnection mapping mechanism to inform the CXP-PHY Bridge in the Device about the upconnection status in the Host side (see section 9). The bits SopCtrl[0] and SopCtrl[1]

[1] The dot after the number 3 means “3 recurring”, i.e. the bitrate is 20.833333… Mbps.

[2] The dot after the number 6 means “6 recurring”, i.e. the bitrate is 41.666666… Mbps.

must report the upconnection status in every low-speed packet. The bit SopCtrl[3] is set to “1” only in packets where the Upconnection status is different from the previous one, otherwise this bit shall be set to “0”.

For high-speed packets (SopCtrl[7] = “1”), the Control field informs the next word type in the stream:

SopCtrl[3:1] - Reserved

SopCtrl[0] - Next word type

When “0” => HDP

When “1” => HKP

6.3.1.3 SopData0

SopData0[7:0] - Embedded K-code (replaces a CoaXPress K-code replicated four times)

When “0x00” => No embedded K-code

When “0x5C” => K28.2 - High-speed trigger

When “0xDC” => K28.6 - I/O acknowledgment

When “0xFB” => K27.7 - Start of data packet indication

When others => Reserved

When a K28.2 (high-speed trigger) is embedded into the SopData0, the SopData1 replaces the *Delay* word of the CoaXPress high-speed trigger packet and the *LinkTriggerN* word is mapped in the HDP word that follows the SOP word.

When a K28.6 (I/O acknowledgment) is embedded into the SopData0, the SopData1 is reserved and the *Code* word of the CoaXPress I/O acknowledgment packet is mapped in the HDP word that follows the SOP word.

When a K27.7 (start of data packet indication) is embedded into the SopData0, the SopData1 replaces the *Type* word of the CoaXPress data packet.

Note: The SopData0 field is only valid on high-speed packets (SopCtrl[7] = “1”). On low-speed packets, the SopData0 field is reserved.

6.3.1.4 SopData1

SopData1[7:0] - Embedded Data (replaces a CoaXPress byte replicated four times)

The SopData1 field is used when SopData0 is “0x5C” (K28.2 – High-speed trigger) to replace the *Delay* word of the CoaXPress high-speed trigger packet. When the SopData0 is “0xFB” (K27.7 – Start of data packet indication), the SopData1 field replaces the *Type* word of the CoaXPress data packet. For all other cases, the SopData1 field is reserved.

Note: The SopData1 field is only valid on high-speed packets (SopCtrl[7] = “1”). On low-speed packets, the SopData1 field is reserved.

6.3.2 EOP – End of Packet

The EOP word marks the end of a CXP-PHY packet at the nGMII interface for both up and down connections.

Table 13 - End of Packet (EOP) word format

Word	Lane	TXC/RXC	TXD/RXD
0	0	0	[7:0] EopData0
	1	0	[7:0] EopData1
	2	1	[7:0] Terminate
	3	1	[7:0] nGMII IDLE
1	0	1	[7:0] nGMII IDLE
	1	1	[7:0] nGMII IDLE
	2	1	[7:0] nGMII IDLE
	3	1	[7:0] nGMII IDLE

The EOP fields are described below.

6.3.2.1 EopData0

EopData0[7:0] - Embedded K-code (replaces a CoaXPress K-code replicated four times)

When “0x00” => No embedded K-code

When “0x7C” => K28.3 - Stream marker

When “0xFD” => K29.7 - End of data packet indication

When others => Reserved

When the EopData0 has no embedded K-code, it indicates that the EOP was originated by 8B/10B IDLEs or high-priority packets embedded into the CoaXPress stream (see sections 7.1.2 and 7.1.3).

When a K28.3 is embedded into the EopData0, it indicates that the EOP was originated in order to introduce a CoaXPress stream marker into the nGMII stream (see section 7.1.2).

When a K29.7 is embedded into the EopData0, it indicates that the EOP was originated by a CoaXPress end of data packet (see section 7.1.2).

Note: The EopData0 field is only valid on high-speed packets (SopCtrl[7] = “1”). On low-speed packets, the EopData0 field is reserved.

6.3.2.2 EopData1

EopData1[7:0] - Reserved

6.3.2.3 Terminate

Terminate[7:0] - Terminate control character

0xFD => nGMII control character /T/

6.3.2.4 nGMII IDLE

nGMII IDLE[7:0] - Idle control character

0x07 => nGMII IDLE control character /I/

Note: The EOP contains 5 nGMII IDLE characters respecting nGMII Rules #2 and #3.

6.3.3 IT – Idle Transfer

IT words form the IPG and are used when there are no packets to be transferred in the nGMII data stream. IT words are formed by four nGMII IDLE control characters /I/ (0x07).

Table 14 - Idle Transfer (IT) word format

Word	Lane	TXC/RXC	TXD/RXD
0	0	1	[7:0] nGMII IDLE
	1	1	[7:0] nGMII IDLE
	2	1	[7:0] nGMII IDLE
	3	1	[7:0] nGMII IDLE

6.3.4 HDP – High-Speed Data Payload

Word containing 4 bytes of data payload.

Table 15 - High-Speed Data Payload (HDP) word format

Word	Lane	TXC/RXC	TXD/RXD
0	0	0	[7:0] Data stream Byte 0
	1	0	[7:0] Data stream Byte 1
	2	0	[7:0] Data stream Byte 2
	3	0	[7:0] Data stream Byte 3

6.3.5 HKP – High-Speed K-Code Payload

Word containing four K-code characters of payload.

Table 16 - High-Speed K-Code Payload word format

Word	Lane	TXC/RXC	TXD/RXD
0	0	0	[7:0] Data stream K-code 0
	1	0	[7:0] Data stream K-code 1
	2	0	[7:0] Data stream K-code 2
	3	0	[7:0] Data stream K-code 3

Note: The HKP word includes only the K-code corresponding hexadecimal value without the K-code bit flag.

6.3.6 LSP – Low-Speed Payload

LSP words form the payload of nGMII low-speed packets.

Table 17 - Low-Speed Payload (LSP) word format

Word	Lane	TXC/RXC	TXD/RXD
0	0	0	[7:0] LspCtrlSlot0
	1	0	[7:0] LspCharSlot0
	2	0	[7:0] LspCtrlSlot1
	3	0	[7:0] LspCharSlot1

The LSP fields are described below.

6.3.6.1 LspCtrlSlot*n*

LspCtrlSlot*n*[7:2] - Reserved

LspCtrlSlot*n*[1:0] - Indicates the character type at the Character Slot *n* field

When “00” => Not in use

When “01” => Data

When “10” => K-code

When “11” => Reserved

6.3.6.2 LspCharSlot*n*

Character from a corresponding CoaXPress low-speed connection.

7 High-speed CXP-PHY mapping

The high-speed CXP-PHY mapping mechanism covers CoaXPress high-speed connections (up and down) mapping to/from a nGMII interface. This section describes the mapping procedure for each element present in the CoaXPress data stream. This section also presents an analysis of the bandwidth overhead introduced by the CXP-PHY mapping mechanism.

7.1 CXP-to-PHY mapping

Note: All examples in this section do not consider the relative latencies between the *CXP Stream* and the *nGMII Stream*. For detailed examples, please refer to Annex A.

7.1.1 CoaXPress 8B/10B IDLE words

All 8B/10B IDLE words are removed from the CoaXPress stream at the CXP-to-PHY mapping. This saves at least 1% of the bandwidth in the nGMII since a minimum of one 8B/10B IDLE word must be inserted into the CoaXPress data stream every 100 words. Removing 8B/10B IDLE words also simplifies the mapping since they differ from regular CoaXPress K-code words, which include one K-code replicated four times.

When there is no data to be transferred, IT words are introduced into the IPG to fill the nGMII stream.

7.1.2 CoaXPress data packets

CoaXPress data packets are formed as follows:

Transfer	**0**	**1**	**2**	**3**	**4**	…	***n*-1**	***n***
CXP Stream	4xK27.7	4xType	DW 0	DW 1	DW 2	…	DW *m*	4xK29.7

Figure 12 - Example of a CoaXPress data packet

The procedure to map data packets into the nGMII stream is to replace the 4xK27.7 and 4xType words by an SOP word and to replace the 4xK29.7 word by an EOP word. The CoaXPress Data Words (DWs) are simply mapped into HDP words.

As a result, we obtain:

Transfer	**0**	**1**	**2**	**3**	…	***n*-2**	***n*-1**	***n***
nGMII Stream	SOP	HDP	HDP	HDP	…	HDP	EOP	
Lane 0	/S/	DW 0	DW 1	DW 2		DW *m*	0xFD	4x /I/
Lane 1	0x80						0x00	
Lane 2	0xFB						/T/	
Lane 3	Type						/I/	

Figure 13 - Example of a CoaXPress data packet mapped into the nGMII stream

CoaXPress data packets can also include one or more stream marker words that contain a K28.3 character replicated four times and 8B/10B IDLE words for packet stretching.

A data buffer is required in the CXP-to-PHY mapping module to guarantee that nGMII packets are continuous. If this buffer becomes empty, the packet must be terminated since nGMII IDLE characters are not allowed within a nGMII packet. In case 8B/10B IDLE words are embedded in a CoaXPress data packet, the packet may have to be split into two nGMII packets to introduce IT words until more data is available for transmission. The following example illustrates this case considering that the internal buffer at the CXP-to-PHY mapping module becomes empty:

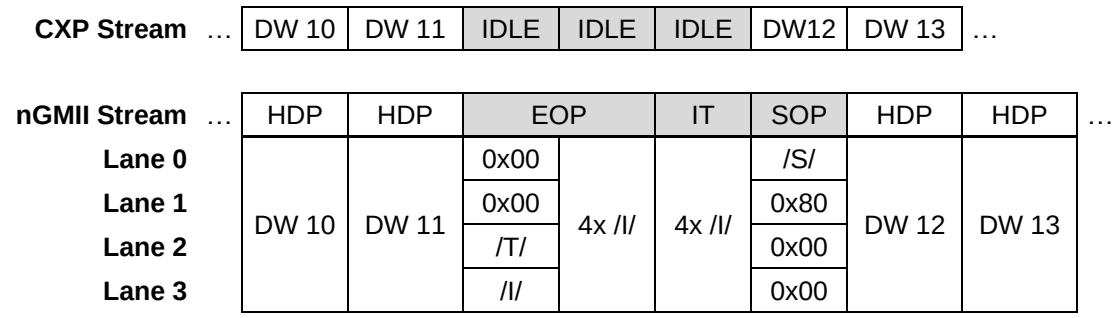

Figure 14 - Mapping a CoaXPress data packet with embedded 8B/10B IDLE words

When one or more stream marker words (4xK28.3) are embedded in a CoaXPress data packet, the marker is replaced by an EOP followed by a SOP to resume packet transmission.

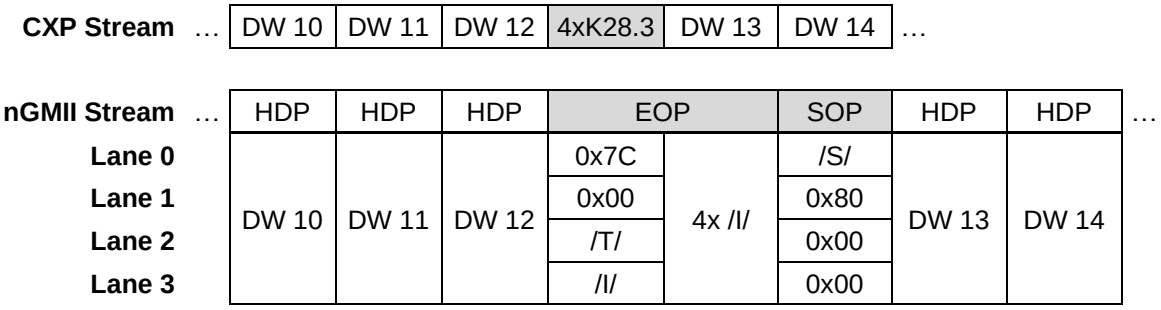

Figure 15 - Mapping a CoaXPress data packet with embedded stream marker words

When splitting a CoaXPress packet into two or more nGMII packets, it is mandatory that the resulting packets have at least one word of payload.

7.1.3 High-priority packets

High-priority CoaXPress packets are specific packets that can be formed by two words (I/O acknowledgement) or three words (high-speed trigger). They do not present regular start and end of packet markers (K27.7 and K29.7 respectively). The high-speed CXP-PHY mapping mechanism shall map high-priority packets into the nGMII stream as soon as they are detected in the CoaXPress stream, thus bypassing any data buffer in the CXP-to-nGMII mapping module.

Three cases must be taken into consideration during mapping:

- Isolated high-priority packets
- High-priority packets embedded in data packets
- High-speed trigger packet embedded in a I/O acknowledgement packet

Isolated high-priority packets are present in the CoaXPress data stream outside of a data packet. In this case, there is no packet in the nGMII stream to be terminated. The following example illustrates this case:

CXP Stream	…	IDLE	IDLE	4xK28.2	4xDelay	4xTrgN	IDLE	IDLE	…

nGMII Stream	…	IT	IT	SOP	HDP	EOP		IT	…
Lane 0		4x /I/	4x /I/	/S/	4xTrgN	0x00	4x /I/	4x /I/	
Lane 1				0x80		0x00			
Lane 2				0x5C		/T/			
Lane 3				Delay		/I/			

Figure 16 - Mapping an isolated high-priority packet

When a high-priority packet is embedded in a data packet, the corresponding nGMII packet must be terminated before inserting the high-priority packet. Once the high-priority packet data is inserted into the new nGMII packet, the latter needs to be terminated only if the following word in the CoaXPress stream is a K-code or no other word is available for transmission. The following examples illustrate the high-priority packets mapping procedure:

CXP Stream	…	DW 10	DW 11	4xK28.6	4xCode	DW 12	DW 13	…

nGMII Stream	…	HDP	HDP	EOP		SOP	HDP	HDP	HDP	…
Lane 0		DW 10	DW 11	0x00	4x /I/	/S/	4xCode	DW 12	DW 13	
Lane 1				0x00		0x80				
Lane 2				/T/		0xDC				
Lane 3				/I/		0x00				

Figure 17 - Mapping a CoaXPress data packet with embedded I/O acknowledge packet

Note: The 4xCode word is not embedded in the SOP because, as stated before, a nGMII packet must be formed by at least one word of payload.

CXP Stream	…	DW 10	DW 11	4xK28.2	4xDelay	4xTrgN	DW 12	DW 13	…

nGMII Stream	…	HDP	HDP	EOP		SOP	HDP	HDP	HDP	…
Lane 0		DW 10	DW 11	0x00	4x /I/	/S/	4xTrgN	DW 12	DW 13	
Lane 1				0x00		0x80				
Lane 2				/T/		0x5C				
Lane 3				/I/		Delay				

Figure 18 - Mapping a CoaXPress data packet with embedded high-speed trigger packet

In the case where a high-speed trigger packet is embedded into an I/O acknowledge packet, the high-speed trigger has priority over the I/O acknowledge and must be mapped into the nGMII stream as soon as possible. The following example illustrates this case:

CXP Stream	…	DW10	4xK28.6	4xK28.2	4xDelay	4xTrgN	4xCode	DW11	DW12	DW13	…

nGMII Stream	…	DW10	EOP		SOP	HDP	EOP		SOP	HDP	…
Lane 0		DW10	0x00	4x /I/	/S/	4xTrgN	0x00	4x /I/	/S/	4xCode	
Lane 1			0x00		0x80		0x00		0x80		
Lane 2			/T/		0x5C		/T/		0xDC		
Lane 3			/I/		Delay		/I/		0x00		

Figure 19 - Mapping a high-speed trigger embedded in an I/O acknowledge packet

To avoid introducing jitter on high-speed trigger packets, it is mandatory to have a fixed latency between the start of a high-speed trigger at the CXP stream (4x K28.2) and its mapped version at the nGMII stream (SOP). Once the start of a high-speed trigger packet is detected in the CXP stream, the mapping mechanism shall wait three transfer cycles before introducing its corresponding SOP. These three cycles are used by the mapping mechanism to properly end a packet that is being transferred in the nGMII stream while respecting the packet structure described in section 6.3 . In the case where there is no other packet being transferred while the start of a high-speed trigger packet is detected in the CXP stream, these three cycles shall be filled with IT words.

For more detailed examples of CXP-to-nGMII mapping, please refer to Annex A.

7.1.4 Use of HKP words

HKP words are used to map K-code words into the nGMII stream when these words cannot be embedded into an SOP or an EOP. One example of this situation is when there is no CXP-PHY packet being transferred and the only data available in the data buffer is a K-code word. In this case, an HKP word is used in order to respect the minimum payload size of one word.

CXP Stream	…	DW10	IDLE	IDLE	IDLE	4xK29.7	IDLE	…

nGMII Stream	…	HDP	EOP		IT	SOP	HKP	EOP		IT	…
Lane 0		DW10	0x00	4x /I/	4x /I/	/S/	0xFD	0x00	4x /I/	4x /I/	
Lane 1			0x00			0x81	0xFD	0x00			
Lane 2			/T/			0x00	0xFD	/T/			
Lane 3			/I/			0x00	0xFD	/I/			

Figure 20 - Example of the use of HKP words

Note: An HKP word is always preceded by a SOP and followed by an HDP or an EOP (see section 6.3).

7.2 PHY-to-CXP mapping

The PHY-to-CXP mapping mechanism is responsible for the reconstruction of the CoaXPress data stream based on nGMII packets generated by the CXP-to-PHY mapping module. A data buffer is required in this module to be able to halt the input stream when decoding a SOP word, inserting 8B/10B IDLE words or a high-priority message into the CoaXPress data stream.

8B/10B IDLE words are reinserted into the data stream at a minimum rate of one 8B/10B IDLE word every 100 words and whenever no data is available at the receiver. As the 8B/10B IDLE words are not present in the nGMII stream, the PHY-to-CXP mapping module freely determines the position where they are reinserted in the CoaXPress data stream. The only restriction is that 8B/10B IDLE words cannot be inserted into a high-priority packet as stated in the CoaXPress standard.

7.3 Overhead analysis

The CoaXPress bandwidth usage is dominated by data packet transfers. The bandwidth necessary to transfer high-priority packets and mandatory 8B/10B IDLE words is virtually negligible (~1%). The CXP-to-PHY mapping mechanism basically does not introduce overhead when mapping data packets into nGMII packets. However, for any K-code embedded in a CoaXPress data packet (i.e., stream markers and high-priority packets), the mapping mechanism must terminate the current nGMII packet to announce the new K-code in a SOP word. This process introduces two words of overhead.

The ratio of embedded K-code characters in comparison with regular data words is strongly dependent on the size of lines being transferred from Device to Host. The following examples illustrate that.

7.3.1 Example 1 – Area Scan

The CoaXPress data stream parameters for this example are:

Pixel format = Mono8
Number of embedded I/O ack. = 2 per frame
Number of CXP connections = 1

Table 18 - nGMII mapping overhead for area scan applications

Xsize	Ysize	TotalSize (words)	MappingOverhead (words)	MappingOverhead (%)
64	64	1,183	122.2	10.3
128	128	4,383	218.2	5.0
256	256	16,927	348.7	2.1
512	512	66,591	364.1	0.5
1,024	1,024	264,223	-588.2	-0.2
2,048	2,048	1,052,703	-6,425.0	-0.6
4,096	4,096	4,202,527	-33,827.1	-0.8
8,192	8,192	16,793,631	-151,546.3	-0.9

In this table, the *TotalSize* is the total number of words for a given frame size and is obtained as follows:

TotalSize = ImageHeader + StreamMarker + Ysize(Xsize/4 + StreamMarker) + 2*Acknowledge*

where *ImageHeader* is 25 words, *StreamMarker* is 2 words, and *Acknowledge* is 2 words.

Note: The *TotalSize* is an approximation to simplify the understanding of the mapping overhead. For a precise value, other parameters (e.g., packet size) must be taken into consideration.

The *MappingOverhead* is the total number of additional words that must be inserted into the nGMII stream to perform the CXP-to-PHY mapping. *MappingOverhead* is obtained as follows:

MappingOverhead = StreamMarkerOvh(Ysize + 1) + 2*AcknowledgeOvh – NumberOfIDLEs*

where *StreamMarkerOvh* is the overhead due to the stream marker mapping and it corresponds to 2 words; *AcknowledgeOvh* is the overhead due to the I/O acknowledge packet mapping and it corresponds to 2 words as well; *NumberOfIDLEs* is number of mandatory 8B/10B IDLE words embedded in the CoaXPress stream and it corresponds to *TotalSize/100* words.

Note: The *MappingOverhead* assumes negative values when the number of mandatory 8B/10B IDLE words in the CoaXPress stream is larger than the overhead introduced by the CXP-to-PHY mapping mechanism. This happens because all the 8B/10B IDLE words are removed from the CoaXPress stream by the CXP-to-PHY mapping mechanism (see section 7.1.1).

7.3.2 Example 2 – Line Scan

The CoaXPress data stream parameters for this example are:

Pixel format = Mono8
Number of embedded I/O ack. = 1 per line
Number of CXP connections = 1

Table 19 - nGMII mapping overhead for line scan applications

Xsize	TotalSize (words)	MappingOverhead (words)	MappingOverhead (%)
64	20	3.8	19.0
128	36	3.6	10.1
256	68	3.3	4.9
512	132	2.7	2.0
1,024	260	1.4	0.5
2,048	516	-1.2	-0.2
4,096	1,028	-6.3	-0.6
8,192	2,052	-16.5	-0.8

In this table, the *TotalSize* is the total number of words for a given line size and is obtained as follows:

$$TotalSize = Xsize/4 + StreamMarker + 1*Acknowledge$$

where *StreamMarker* is 2 words and *Acknowledge* is 2 words.

Note: The *TotalSize* is an approximation to simplify the understanding of the mapping overhead. For a precise value, other parameters (e.g., packet size) must be taken into consideration.

The *MappingOverhead* is the total number of additional words that must be inserted into the nGMII stream to perform the CXP-to-PHY mapping. *MappingOverhead* is obtained as follows:

$$MappingOverhead = 1*StreamMarkerOvh + 1*AcknowledgeOvh – NumberOfIDLEs$$

where *StreamMarkerOvh* is the overhead due to the stream marker mapping and it corresponds to 2 words; *AcknowledgeOvh* is the overhead due to the I/O acknowledge packet mapping and it corresponds to 2 words as well; *NumberOfIDLEs* is number of mandatory 8B/10B IDLE words embedded in the CoaXPress stream and it corresponds to *TotalSize/100* words.

8 Low-speed CXP-PHY mapping

The bandwidth difference between the CoaXPress low-speed upconnections and the nGMII permits multiplexing several low-speed connections into a single nGMII channel. For instance, a CoaXPress low-speed upconnection with a bandwidth of 41.6˙ Mbps (@ CXP-12) is 300 times slower than the XGMII.

The procedure to map one or more low-speed upconnection streams into a single nGMII channel is to create a new nGMII packet when the periodic LS Upconnection Enable signal (see section 6.2.1) indicates that a new set of upstream characters is valid and stable at the input of the low-speed CXP-to-PHY mapping module. This nGMII packet is formed as in the following example for four connections:

CXP LS Stream (A)	…	DataA0	DataA1						…
CXP LS Stream (B)	…	DataB0	DataB1						…
CXP LS Stream (C)	…	DataC0	DataC1						…
CXP LS Stream (D)	…	DataD0	DataD1						…

nGMII Stream	…	IT	IT	SOP	LSP	LSP	EOP		IT	…
Lane 0		4x /I/	4x /I/	/S/	0x01	0x01	0x00	4x /I/	4x /I/	
Lane 1				0x00	DataA1	DataC1	0x00			
Lane 2				0x00	0x01	0x01	/T/			
Lane 3				0x00	DataB1	DataD1	/I/			

Figure 21 - CoaXPress low-speed upconnection mapping example

The data available in each CoaXPress low-speed upconnection is mapped into one Character Slot of an LSP word. The order of slots corresponds to the order of connectors in the Device when mapping back to the CoaXPress low-speed stream. There must be enough LSP words in the nGMII packet to accommodate the number of connections available on the Host side.

In the low-speed CXP-to-PHY mapping mechanism, the data in the CoaXPress stream may not be synchronized as in the previous example. To cope with that case, the low-speed CXP-to-PHY mapping mechanism shall latch the data available in its inputs when the LS Upconnection Enable signal (see section 6.2.1) is set.

Note: The first character slot is mapped to the Master upconnection in the Device side, i.e., triggers and control packets are mapped in this slot.

Comment: The order that the low-speed upconnections are mapped to the Character Slots depends on the connection topologies supported by the Host.

Comment: For Hosts supporting multi-Devices and/or multi connection topologies, it may be necessary to have a more advanced upconnection mapping mechanism for each CXP-PHY bridge.

The mapping procedure maps all data available in the CoaXPress stream. The K-code characters (the corresponding hexadecimal value without the K-code bit flag) are mapped into the LspCharSlot fields while the character type (K-code or data) is announced in the LspCtrlSlot fields (see section 6.3.6). With this mechanism, all packets and 8B/10B IDLE words are mapped character by character without any modification.

9 Discovery considerations

To be transparent to the adjacent CoaXPress layers (Host or Device), the CXP-PHY bridge supports the different tasks related to the Device discovery. This section describes how the CXP-PHY bridge behaves in each step of the Device discovery.

9.1 Connection state

The CoaXPress standard states that each transmitter connection has an associated connection state: Activated or Deactivated. A transmitter connection is considered Activated when 8B/10B IDLE words are sent on the connection and Deactivated when no 8B/10B IDLE words are sent on the connection.

In the CXP-PHY bridge, the high-speed CXP-to-PHY mapping mechanism replaces 8B/10B IDLE words by IT words. When the minimal number of 8B/10B IDLE words are not detected in the CoaXPress stream (see section 5.4), the CXP-to-PHY mapping mechanism shall not generate IT words. Instead of generating IT words, a word containing four times the data 0x00 shall be used. The low-speed CXP-to-PHY mapping mechanism is always active and shall insert IT words into the nGMII stream independently of the activity in the CoaXPress interface.

Comment: When transferring four times the data 0x00 instead of IT words, the PCS Ethernet sublayer at the reception side will decode the incoming data even if they are not part of a well-formed packet. Depending on the PCS implementation, it may also report a framing error.

The CoaXPress standard also defines that each receiver connection has an associated connection state: Detected or Undetected. A receiver connection is considered Detected when it achieves low-level lock, and Undetected when the low-level lock is not achieved. The detection of the low-level lock depends on the CoaXPress implementation, but it is closely related to the receiver bit clock lock and word and character alignment. The verification of the presence of 8B/10B IDLE words is also recommended to be used as a condition to achieve low-level lock.

In a nGMII connection, word and character alignments are intrinsic to the PCS encoding and 8B/10B IDLE words are not present in the stream. The solution to detect the low-level connection lock in a nGMII connection is to monitor the presence of IT words and valid CXP-PHY packets in the nGMII stream. The CXP-PHY bridge shall achieve low-level connection lock when IT words and/or valid CXP-PHY packets are present in the nGMII stream. The CXP-PHY bridge shall lose low-level connection lock when no IT words nor valid CXP-PHY packets are present in the nGMII stream within a timeout period.

Comment: The value of the timeout period is beyond the scope of this standard. In the CXP-PHY bridge, the high-speed PHY-to-CXP mapping mechanism shall only start regenerating 8B/10B IDLE words in the CoaXPress stream when it achieves low-level connection lock. It shall stop regenerating 8B/10B IDLE words in the CoaXPress stream when it loses low-level connection lock.

When the high-speed upconnection is implemented in the Device, its corresponding high-speed PHY-to-CXP mapping mechanism shall also use the *High-speed upconnection state* (SopCtrl[0], see section 6.3.1.2) as an additional condition to start regenerating 8B/10B IDLE words.

9.2 Match Device transmitter and Host receiver bit rates

The procedure to match Device transmitter and Host receiver bit rates starts with the Host writing a connection reset command into the Device via the low-speed upconnection at 20.83˙ Mbps. The low-speed CXP-to-PHY mapping mechanism at the Host side should not change its behavior either if its input is at 20.83˙ Mbps or at 41.6˙ Mbps, i.e., it keeps mapping into a nGMII packet any new data available at its input. The low-speed PHY-to-CXP mapping mechanism at the Device side also does not change its behavior.

Regarding the CoaXPress high-speed downconnections, during discovery they work at one of the discovery bit rates (1.25 Gbps or 3.125 Gbps). The method to adapt the bit rate between CoaXPress and nGMII could be based on coupling FIFOs. However, according to the CoaXPress implementation in use, the method can vary and falls outside the scope of this specification.

Once the Host and Device have established a connection, the following tasks of the Device discovery procedure can be performed normally:

- Discover devices and connection topology
- Negotiate the CoaXPress version used
- Negotiate maximum packet sizes
- Setting the operating bit rate

9.3 Detection of high-speed upconnection

When the procedure to detect a high-speed upconnection starts, the Host and Device have already established a connection (low-speed upconnection and high-speed downconnection). Since the connection is already established, the first steps of the detection can be performed normally. The first steps consist in reading the *Capability* register bit *HsUpconnectionSupported* via the low-speed connection. If the value of this bit is “1”, the Host shall set the *HsUpconnectionEnable* bit in the *FeatureControl* register to “1” and wait for the acknowledgement. Once the Device acknowledges the FeatureControl access, the Host shall activate its high-speed upconnection. In the CoaXPress standard, this activation corresponds to start sending 8B/10B IDLE words via the high-speed upconnection while keep sending 8B/10B IDLE words via the low-speed upconnections.

Note: The CoaXPress over Fiber Bridge Protocol document assumes that the CoF Device and the CoF Host implement the procedure to detect the high-speed upconnection as described in Annex B.

When the Host CXP-PHY bridge detects that the high-speed upconnection is activated, the PHY TX Packet Arbiter module shall stop sending new low-speed packets, with the exception of low-speed packets with *Update flag* (SopCtrl[3]) set to “1”. These packets are used to control the behavior of the Low-Speed PHY-to-CXP Mapping module at the Device CXP-PHY bridge. The PHY TX Packet Arbiter module shall resume sending low-speed packets when the high-speed upconnection is deactivated.

The Low-Speed CXP-to-PHY Mapping module at the Host CXP-PHY bridge must be able to detect the state of the high-speed upconnection. When the high-speed upconnection state is updated, the Low-Speed

CXP-to-PHY Mapping module shall set the *Update flag* to “1” in the next packet to be transmitted and update the bit *High-speed upconnection state* (SopCtrl[0]).

The Low-Speed PHY-to-CXP Mapping module at the Device CXP-PHY bridge shall start generating 8B/10B IDLE words automatically as soon as it receives a packet with the bit *High-speed upconnection state* set to “1”. This is necessary because no other low-speed packets will be available in the nGMII stream. When it receives a packet with the bit *High-speed upconnection state* set to “0”, it shall stop the automatic generation of 8B/10B IDLE words and resume the decoding of the nGMII input stream.

The mechanism to detect a high-speed upconnection shall proceed the following:

- All packets generated by the High Speed CXP-to-PHY Mapping module shall be transferred to the nGMII interface.
- All packets generated by the Low Speed CXP-to-PHY Mapping module with SopCtrl[0] = “0” shall be transferred to the nGMII interface.
- All packets generated by the Low Speed CXP-to-PHY Mapping module with SopCtrl[3] = “1” shall be transferred to the nGMII interface.

Note: Packets generated by the Low Speed CXP-to-PHY Mapping module with SopCtrl[3] = “1” have the purpose of activating the generation of 8B/10B IDLEs at the Low-Speed PHY-to-CXP Mapping module in the Device bridge, as well as, changing its bitrate to 20.83˙Mbps or 41.6˙Mbps.

- The Low Speed CXP-to-PHY Mapping module shall set the High-speed upconnection state (SopCtrl[0]) to “1” once the High Speed CXP-to-PHY Mapping module starts generating IT words, i.e., when the Host activates its high-speed upconnection.

After the Host activates the high-speed upconnection, the next steps of the procedure can be performed normally since the connection is already established.

Note: The Low Speed CXP-to-PHY Mapping module shall set the High-speed upconnection state (SopCtrl[0]) back to “0” (default value) when the High-speed upconnection is Deactivated, i.e., no IT words nor high-speed packets are generated by the High Speed CXP-to-PHY Mapping module for a timeout period.

9.3.1 Setting the bit rate via the high-speed upconnection

The bit rate at low-speed upconnection varies between 20.83˙ Mbps and 41.6˙ Mbps according to the bit rate at high-speed connections:

CXP-1, CXP-2, CXP-3, CXP-5, and CXP-6 => 20.83˙ Mbps
CXP-10 and CXP-12 => 41.6˙ Mbps

When the high-speed upconnection is activated, the Low-Speed PHY-to-CXP Mapping module at the Device CXP-PHY bridge keeps generating 8B/10B IDLE words automatically at the same rate as before activating the high-speed upconnection.

The Low-Speed CXP-to-PHY Mapping module at the Host CXP-PHY bridge must be able to detect rate changes in its input (this will occur when, for instance, the high-speed downconnections change from CXP-

6 to CXP-12 or vice-versa). When an input rate change is detected, the Low-Speed CXP-to-PHY Mapping module must set the *Update flag* to “1” in the next packet to be transmitted and update the bit *Low-speed rate* (SopCtrl[1]).

Comment: At the Host side, the input rate can be detected based on the character rate at the CoaXPress low-speed upconnection interface.

The Low-Speed PHY-to-CXP Mapping module at the Device CXP-PHY bridge must be able to change the generation rate of 8B/10B IDLE words as soon as it receives a packet with the bit *Low-speed rate* set at a rate different from the one in use.

Annex A: Examples of CXP-to-PHY mapping

A.1 Mapping isolated CXP packets

Isolated Data Packet

	0	1	2	3	4	5	6	7	8	9	10	11
CXP	IDLE	4xK27.7	4xType	DW 0	DW 1	DW 2	DW 3	4xK29.7	IDLE	IDLE	IDLE	IDLE
nGMII	IT	IT	IT	IT	SOP	HDP	HDP	HDP	HDP	EOP		IT
	4x /I/	4x /I/	4x /I/	4x /I/	/S/	DW 0	DW 1	DW 2	DW 3	0xFD	4x /I/	4x /I/
					0x80					0x00		
					0xFB					/T/		
					Type					/I/		

Note: The nGMII data packet can only start when there is at least one DW available for mapping.

Isolated I/O Acknowledge Packet

	0	1	2	3	4	5	6	7	8	9	10	11
CXP	IDLE	4xK28.6	4xCode	IDLE	IDLE	IDLE	IDLE	IDLE	IDLE	IDLE	IDLE	IDLE
nGMII	IT	IT	IT	SOP	HDP	EOP		IT	IT	IT	IT	IT
	4x /I/	4x /I/	4x /I/	/S/	4xCode	0x00	4x /I/	4x /I/	4x /I/	4x /I/	4x /I/	4x /I/
				0x80		0x00						
				0xDC		/T/						
				0x00		/I/						

Note: The nGMII I/O acknowledge packet can only start when the 4xCode word is available for mapping.

Isolated High-Speed Trigger Packet

	0	1	2	3	4	5	6	7	8	9	10	11
CXP	IDLE	4xK28.2	4xDelay	4xTrgN	IDLE	IDLE	IDLE	IDLE	IDLE	IDLE	IDLE	IDLE
nGMII	IT	IT	IT	IT	IT	SOP	HDP	EOP		IT	IT	IT
	4x /I/	4x /I/	4x /I/	4x /I/	4x /I/	/S/	4xTrgN	0x00	4x /I/	4x /I/	4x /I/	4x /I/
						0x80		0x00				
						0x5C		/T/				
						Delay		/I/				

Note: The nGMII high-speed trigger packet can only start after a fixed latency of 3 transfer cycles is met.

A.2 Mapping I/O acknowledge packets

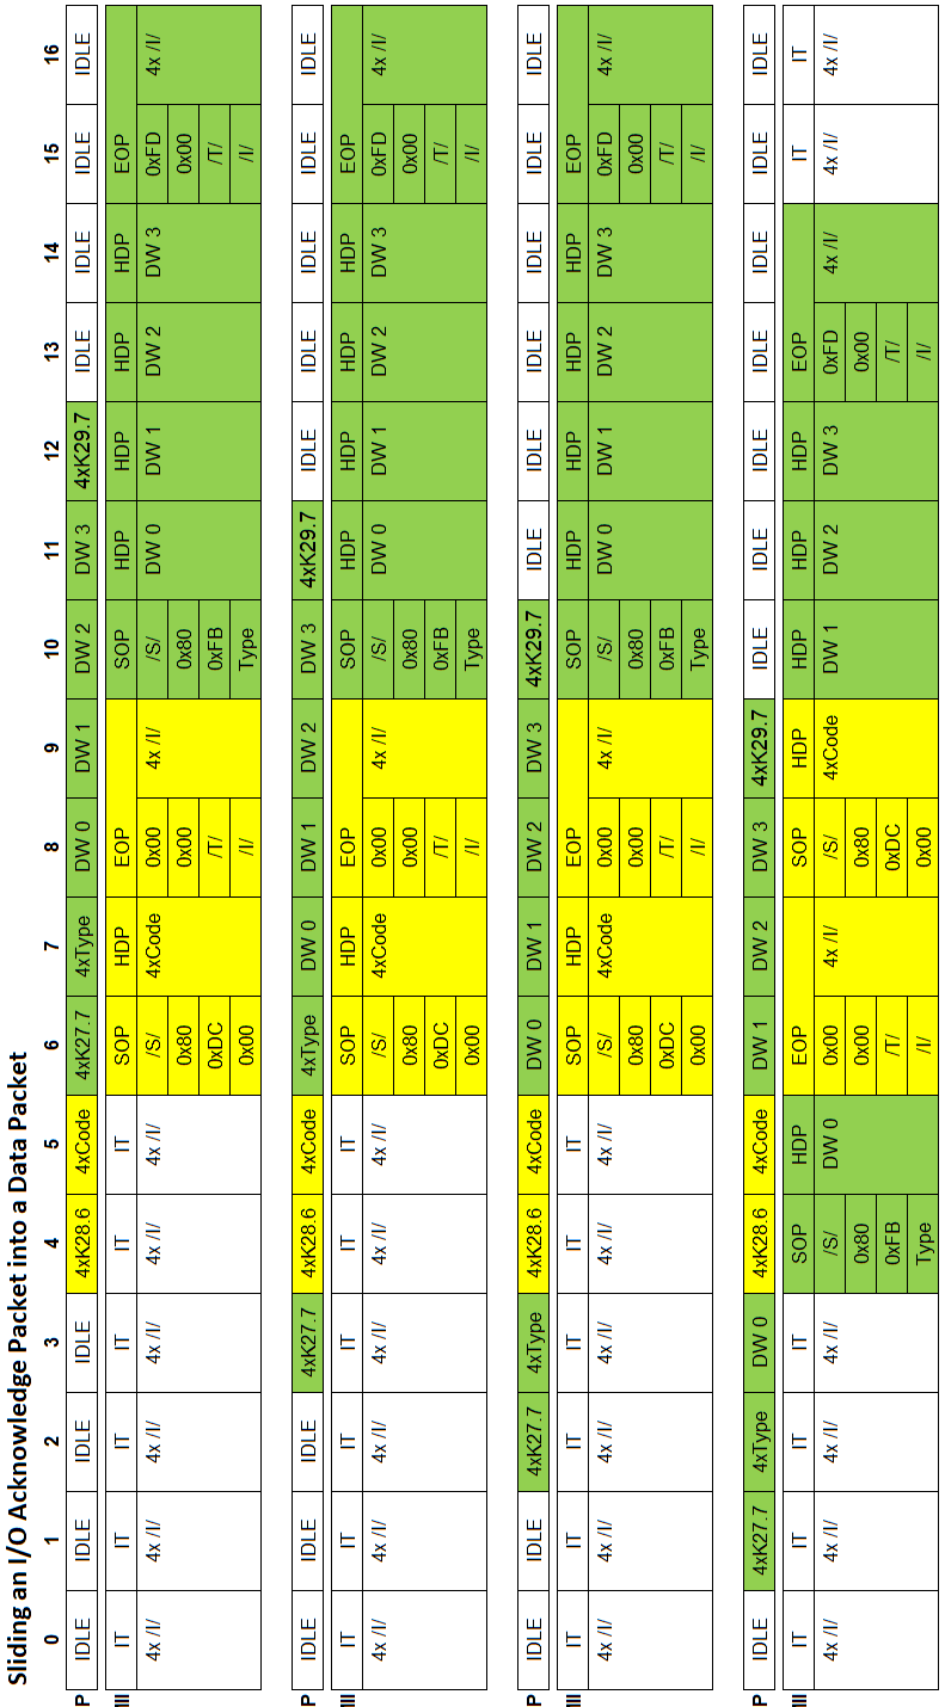

Comment: *Mapping I/O acknowledge packets is straightforward since CoaXPress does not require a delay compensation for these packets.*

A.3 Mapping high-speed trigger packets

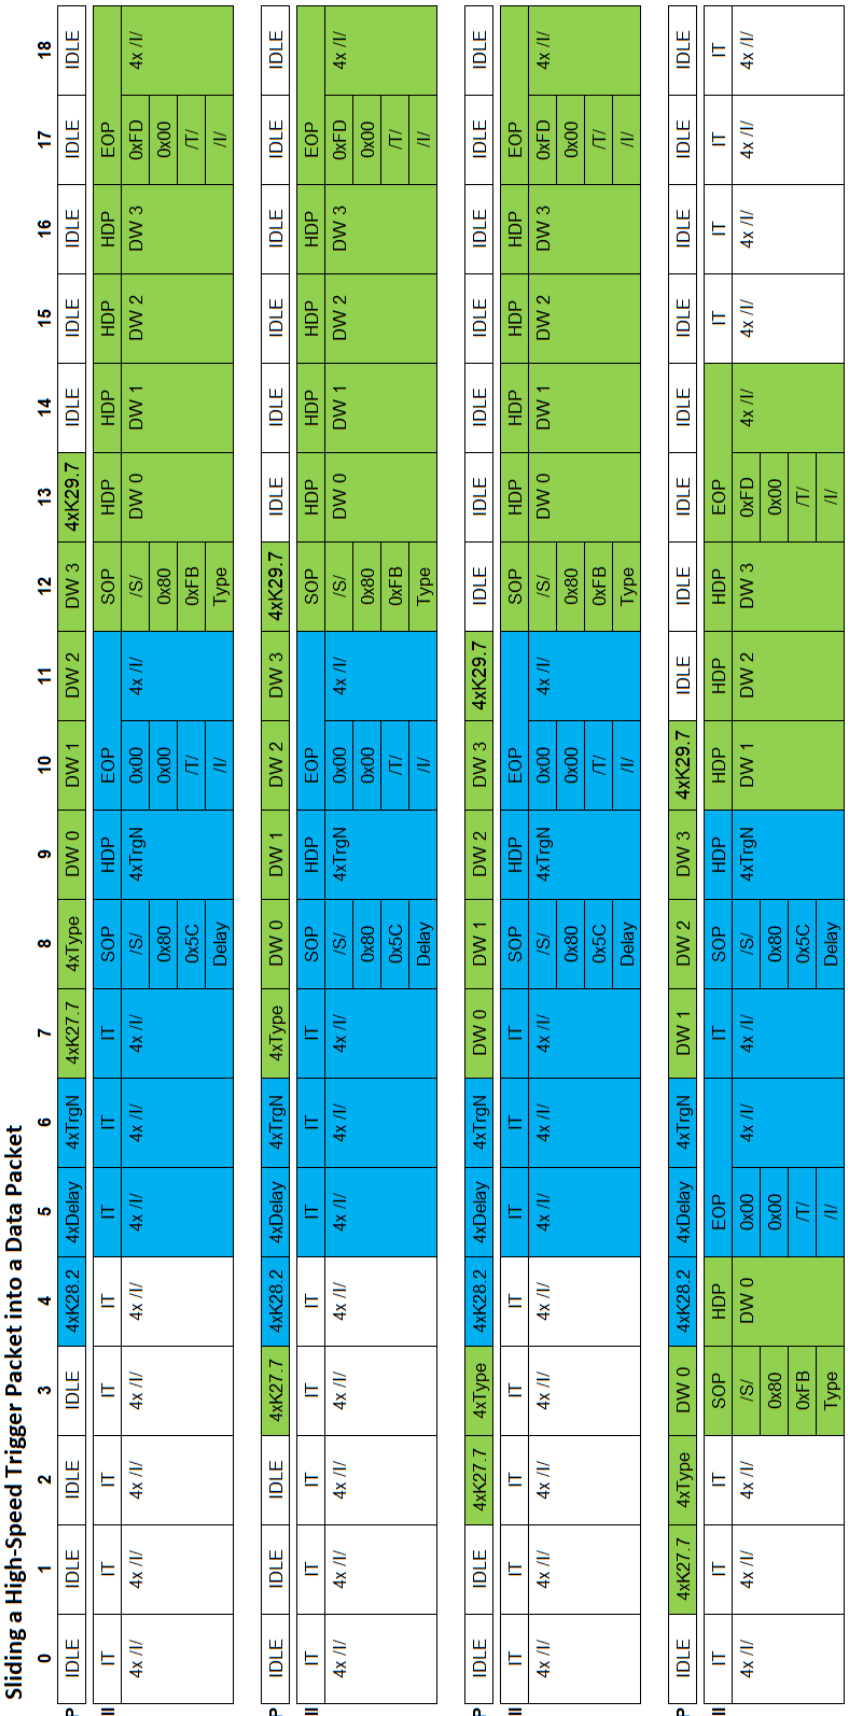

Comment: *Mapping high-speed trigger packets requires a fixed latency between the start of packet at the CXP stream and the corresponding SOP word at the nGMII stream. To guarantee this fixed latency, it is necessary to wait three transfer cycles before starting the packet at the nGMII stream. These three cycles of latency cover all mapping scenarios.*

A.4 Mapping high-priority packets

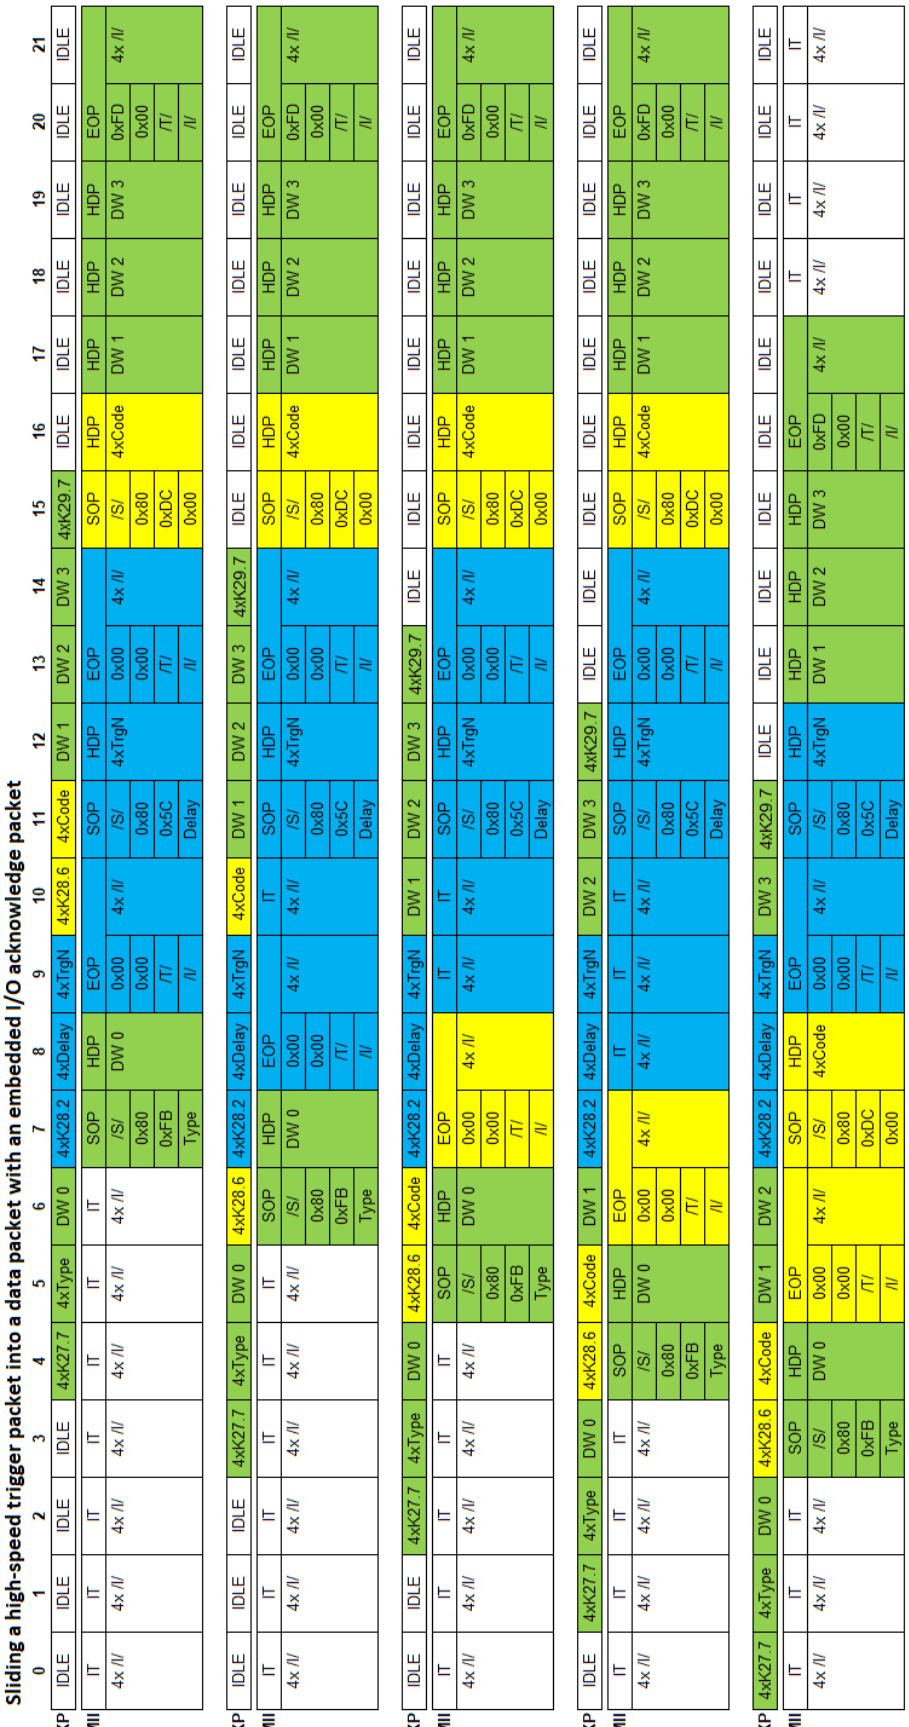

Comment: *This example shows how to map a high-speed trigger while keeping its fixed mapping latency.*

Annex B: CoaXPress Standard Update

B.1 Detection of High Speed Upconnection

The CoaXPress over Fiber Bridge Protocol document assumes that the CoF Device and the CoF Host implement the following procedure to detect the high-speed upconnection. This text modification of section 12.1.5 (JIIA CXP-001-2021) was submitted to the CoaXPress Technical Committee of the JIIA CoaXPress Working Group and is planned to be applied in the next revision of the CoaXPress Standard.

B.1.1 JIIA CXP-001 - Section 12.1.5

Note: Text modification in **bold**.

12.1.5 Detection of High Speed Upconnection

If the Host supports a high speed upconnection, it needs to discover if the Device also supports one, and if so, if there is a cable present to form the connection. A Host supporting a high speed upconnection shall therefore perform the following process.

- The Host shall read the *CapabilityRegister* bit *HsUpconnectionSupported* via the low speed connection.
- **If the value of bit 0 of the register is "1", then the Host shall set the *HsUpconnectionEnable* bit in the *FeatureControl* register to 1 and wait for the acknowledgement.**
- **After receiving the acknowledgement, the Host shall enable its high speed upconnection at the current downconnection speed and send IDLE words on it.**
- When the Device achieves connection state "Detected" (i.e bit and word lock) on its high speed upconnection, it shall switch all upconnection communications from the low speed connection to the high speed connection. Until the Device achieves connection state "Detected" it shall continue to use the low speed connection.
- The Host shall wait 200ms for the Device to achieve lock on the high speed connection.
- The Host shall then read the *CapabilityRegister* via the high speed connection. If the read succeeds, then the Host shall switch all upconnection communications to the high speed connection from the low speed connection. If the read fails then the Host shall disable its high speed upconnection (so not sending IDLE words), and continue to use the low speed upconnection. It shall also set the *HsUpconnectionEnable* bit in the *FeatureControl* register back to 0.

…

Annex C: Descriptive Clause

C.1 Background of CoaXPress over Fiber

The CoaXPress over Fiber initiative was driven by the increasing demand from the market to extend the use of the CoaXPress protocol in terms of bandwidth and cable reach. In 2019, during the IVSM in Stresa, Euresys s.a. and Sensor to Image GmbH presented the original proposal to address these demands by creating a CXP-PHY bridge. This proposal was discussed and accepted by the CoaXPress Technical Committee and the document CoaXPress over Fiber - Bridge Protocol is the result of this work.

C.2 Version History

C.2.1 v1.0

Initial release.

C.2.2 V1.0 to v1.1

Section numbers listed refer to the v1.1 CoaXPress over Fiber - Bridge Protocol document.

v1.1 defines CXP-25 as an extension of the existing CXP Speeds and its corresponding bit rate code to allow selecting a connection speed the supports the 25GMII nominal bit rate (see section 5.6).

Key changes:

- Added section 5.6 describing the CXP-25 and CXP-31 extensions for CoaXPress.

Detailed changes:

- “CoaXPress Specification” changed to “CoaXPress Standard” in the document, to match the official JIIA title.
- Added CoF abbreviation for CoaXPress over Fiber. See section 2.3.
- Removed the number of symbol errors that the RS-FEC can detect. See section 4.2.2.
- Updated Figure 4 as in CXP-001-2021. See section 5.1.
- Title of section 6.1 Multi-lane configuration modified to 6.1 CoF physical layer topology.
- Clarified text regarding the physical layer topology. See sections 6.1 and 6.2.
- Caption of Figure 6 - Physical medium topology in a 4-connection configuration modified to Figure 6 - CoF physical layer topology for a CoaXPress link with 4 connections. See section 6.1.
- Added a note to remind that the CXP-to-nGMII examples in section 7.1 do not consider the relative latencies of the mapping mechanism. See section 7.1.
- Clarified text regarding how to map high-priority CXP packets into the nGMII stream. See section 7.1.3.
- Updated Figure 8 to highlight the optional parts. See section 6.2.1.
- Clarified text regarding the optional features of the CXP-PHY Bridge – Host. See section 6.2.1.
- Updated Figure 9 to highlight the optional parts. See section 6.2.2.
- Clarified text regarding the optional features of the CXP-PHY Bridge – Device. See section 6.2.2.
- Clarified the use of the bitfield SopCtrl[3:0] in low-speed packets. See section 6.3.1.2.

- Replaced the fixed period of 10 ms to lose low-level connection by a timeout. See section 9.1.
- Clarified the text on the procedure to detect the high-speed upconnection. See section 9.3.
- Added “Annex A” with detailed examples of CXP-to-nGMII mapping.
- Added “Annex B” describing the proposed modification in the “Detection of High Speed Upconnection” section of the CoaXPress Standard.
- Other minor wording improvements throughout the document.

C.3 Members Involved

C.3.1 The main authors of the original CoaXPress over Fiber – Bridge Protocol

The CoaXPress over Fiber – Bridge Protocol v1.0 was edited by Paulo Possa of Euresys s.a.

The main authors of the original CoaXPress over Fiber – Bridge Protocol were:

- Euresys s.a. Paulo Possa
- Sensor to Image GmbH Jiri Halak, Werner Feith

In addition, the following people contributed to the CoaXPress over Fiber initiative:

- Active Silicon Ltd. Chris Beynon
- Adimec Advanced Image Systems B.V. Frederik Voncken, Masahito Watanabe
- AVAL DATA Corporation Masahide Matsubara
- Basler AG Frank Krehl
- Euresys s.a. Serge Peterges, Xavier Creuen
- Matrox Electronic Systems Ltd. Eric Boisvert, Jonathan Lavergne
- Sensor to Image GmbH Jan Pech
- Silicon Software GmbH Björn Röger
- TM-Research Inc. Masaru Naganuma
- Vision Research Inc. Lee Denaro

C.3.2 Contributors for v1.1

The CoaXPress over Fiber – Bridge Protocol v1.1 was edited by Paulo Possa of Euresys s.a.

The following people have contributed to the CoaXPress over Fiber – Bridge Protocol v1.1:

- Active Silicon Ltd. Chris Beynon
- Adimec Advanced Image Systems B.V. Frederik Voncken, Masahito Watanabe
- AVAL DATA Corporation Masahide Matsubara
- Basler AG Björn Röger
- Bitflow, Inc. Michael Fortuna
- Euresys s.a. Paulo Possa
- EASii IC SAS Benjamin Pussacq
- FLOVEL Co., Ltd. Masaru Naganuma
- Sensor to Image GmbH Jiri Halak
- Zebra Technologies Corporation Eric Boisvert, Jonathan Lavergne

The standards which the Japan Industrial Imaging Association publishes are enacted whether it infringes or not, on the industrial property such as patents and new utility designs and so on.
一般社団法人日本インダストリアルイメージング協会が発行している規格類は，工業所有権（特許，実用新案など）に関する抵触の有無に関係なく制定されています。

The Japan Industrial Imaging Association is not responsible for any industrial property rights related to the contents of this standard.
一般社団法人日本インダストリアルイメージング協会は，この規格類の内容に関する工業所有権に対して，一切の責任を負いません。

JIIA CXPR-008-2023

Published on November 17, 2023
2023 年 11 月 17 日発行

Published by
Japan Industrial Imaging Association
Room 512, Kikai Shinko Kaikan 3-5-8, Shibakoen,
Minato-ku, Tokyo 105-0011 Japan
TEL. 81-3-6403-1440

発行
一般社団法人 日本インダストリアルイメージング協会
〒105-0011 東京都港区芝公園 3 丁目 5-8 機械振興会館 512 室
TEL:03-6403-1440