Proposed Text for 5GBASE-R PCS/PMA

William Lo
Marvell

To make editing instructions easier to understand for the purposes of assembling the initial Framemaker file the following conventions are used below:

Table 125–1 – Highlight yellow means it cross references something in the 802.3cb draft specification.

1.2.3 – Green text means it references a Clause in IEEE 802.3 that does not appear in the 802.3cb draft specification.

Insert these 2 rows in table 1-23 – Red text means editing instructions for person assembling the initial Framemaker file and should not appear in the document.

Underlined blue text – New text inserted in existing clauses. Note that text in new clauses are not underlined.

~~Cross out purple text~~ – Text that should be deleted in existing clauses.

Editor's Note: Insert following rows in Table 1-23 – Bold italic font look like editing instructions for the person assembling the Framemaker file, but they are instructions for the IEEE editor and are actually part of the of the 802.3cb draft specification as bold italic font. Leave these as is in the text and do not delete or take any action on these. Only act on the editing instructions in red.

Ignore this page with table of contents. It is included here to make sure all the headings in this document are formatted properly as headings. The final assembly of all the clauses in Framemaker should build the table of content in that document

Contents

201 Physical Coding Sublayer (PCS) and Physical Medium Attachment (PMA) sublayer for 64B/66B, type 5GBASE-R

201.1 Overview

201.1.1 Scope

This clause specifies the Physical Coding Sublayer (PCS) and the Physical Medium Attachment (PMA) sublayer that are common to a family of 5 Gb/s Physical Layer implementations known as 5GBASE-R. The 5GBASE-R PCS and 5GBASE-R PMA are sublayers of the 5 Gb/s BASE-R PHY listed in Table 125–1. The term 5GBASE-R is used when referring generally to Physical Layers using the PCS and PMA defined in this clause.

201.1.2 Relationship of 5GBASE-R to other standards

Figure 125–1 depicts the relationships among the 5GBASE-R sublayers, the Ethernet MAC and reconciliation layers, and the higher layers. The 5GBASE-R service interface is the XGMII, which is defined in Clause 46.

The 5GBASE-R PCS is identical to the 10GBASE-R PCS specified in Clause 49 with following exception:

> hi_ber is asserted if ber_cnt reaches 32 in a 125 microsecond period. This differs from the definition in 49.2.13.3 which defines hi_ber as occurring if ber_cnt reaches 16 in a 125 microsecond period.

The purpose of the serial PMA is to attach the PMD of choice to its client. The PMA service interface is defined in an abstract manner and does not imply any particular implementation. The nominal bit rate of the PMA is 5.15625 Gb/s.

201.1.3 Summary of 5GBASE-R sublayers

Figure 201-1 shows the relationship of the 5GBASE-R PCS sublayer with other sublayers to the ISO Open System Interconnection (OSI) reference model.

FM editor's note: Copy figure 125-1 from 802.3bz and draw diagram as sketched below:

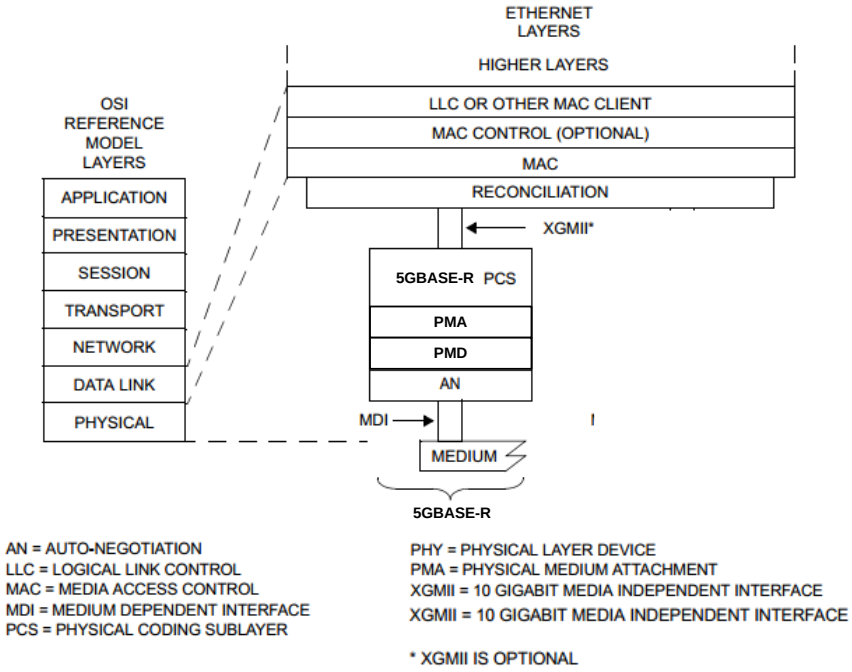

Figure 201-1 5GBASE-R PCS and PMA relationship to the ISO/IEC Open Systems Interconnection (OSI) reference model and IEEE 802.3 Ethernet model

201.1.3.1 Physical Coding Sublayer (PCS)

The PCS service interface is the XGMII, which is defined in Clause 46 running at 5Gb/s. The 5GBASE-R PCS provides all services required by the XGMII including Encoding (decoding) of XGMII data octets to (from) 64B/66B blocks for communication with the underlying PMA.

201.1.3.2 Physical Medium Attachment (PMA) sublayer

The PMA provides a medium-independent means for the PCS to support the use of serial-bit oriented physical media. The 5GBASE-R PMA performs the following functions:

In the transmit direction (i.e., transmitting data from the PMA client to the PMD), the PMA performs the following functions:

a) Provide transmit source clock to PMA client.
b) Serialization of 16-bit data to serial bit stream.
c) Transmission of serial data to PMD.

In the receive direction (i.e., serial data from PMD to the PMA client), the PMA performs the following functions:

a) Bit clock recovery of serial data from PMD.
b) Provide receive clock to PMA client.
c) Deserialization of serial data to16-bit parallel data.
d) Transmission of parallel data to PMA client.
e) Provide link status information.

201.1.4 Inter-sublayer interfaces

There are a number of interfaces employed by 10GBASE-R. Some (such as the PMA service interface) use an abstract service model to define the operation of the interface. The PCS service interface is the XGMII running at 5Gb/s that is defined in Clause 46. The XGMII has an optional physical instantiation.

The upper interface of the PCS may connect to the Reconciliation Sublayer through the XGMII. The lower interface of the PCS connects to the PMA sublayer to support a 5GBASE-R PMD. The nominal rate of the PMA service interface is 322.27 Mtransfers/s, which provides capacity for the MAC data rate of 5 Gb/s.

The nominal rate of the PMD service interface is 5.15625 Gb/s.

If the optional Energy-Efficient Ethernet (EEE) capability is supported (see Clause 78, 78.3) then the interface with the PMA sublayer includes rx_mode and tx_mode to control power states in lower sublayers and energy_detect that indicates whether the PMD sublayer has detected a signal at the receiver.

It is important to note that, while this specification defines interfaces in terms of bits, octets, and frames, implementations may choose other data-path widths for implementation convenience.

201.2 Physical Coding Sublayer (PCS)

201.2.1 Functions within the PCS

The 5GBASE-R PCS shall have all the functionality of the 10GBASE-R PCS specified in Clause 49.

The BER monitor state diagram shown in Figure 49–15 still applies but it shall test ber_cnt for a value of 32 rather than 16 in the exit conditions from state BER_BAD_SH. So:

1) The definition of “ber_cnt” in 49.2.13.2.4 is replaced with “Count up to a maximum of 32 of the number of invalid sync headers within the current 125 us period”.

2) The definition of “hi_ber” in 49.2.13.2.2 is replaced with “Boolean variable which is asserted true when the ber_cnt exceeds 32 indicating a bit error ratio >10-4”.

The PCS encodes data and control information into 66-bit blocks. The relationship of block bit positions to the XGMII, PMA, and other PCS constructs is illustrated in Figure 201-2 for transmit and Figure 201-3 for receive.

FM editor’s note: Copy figure 49-5 and modify text as sketched below:

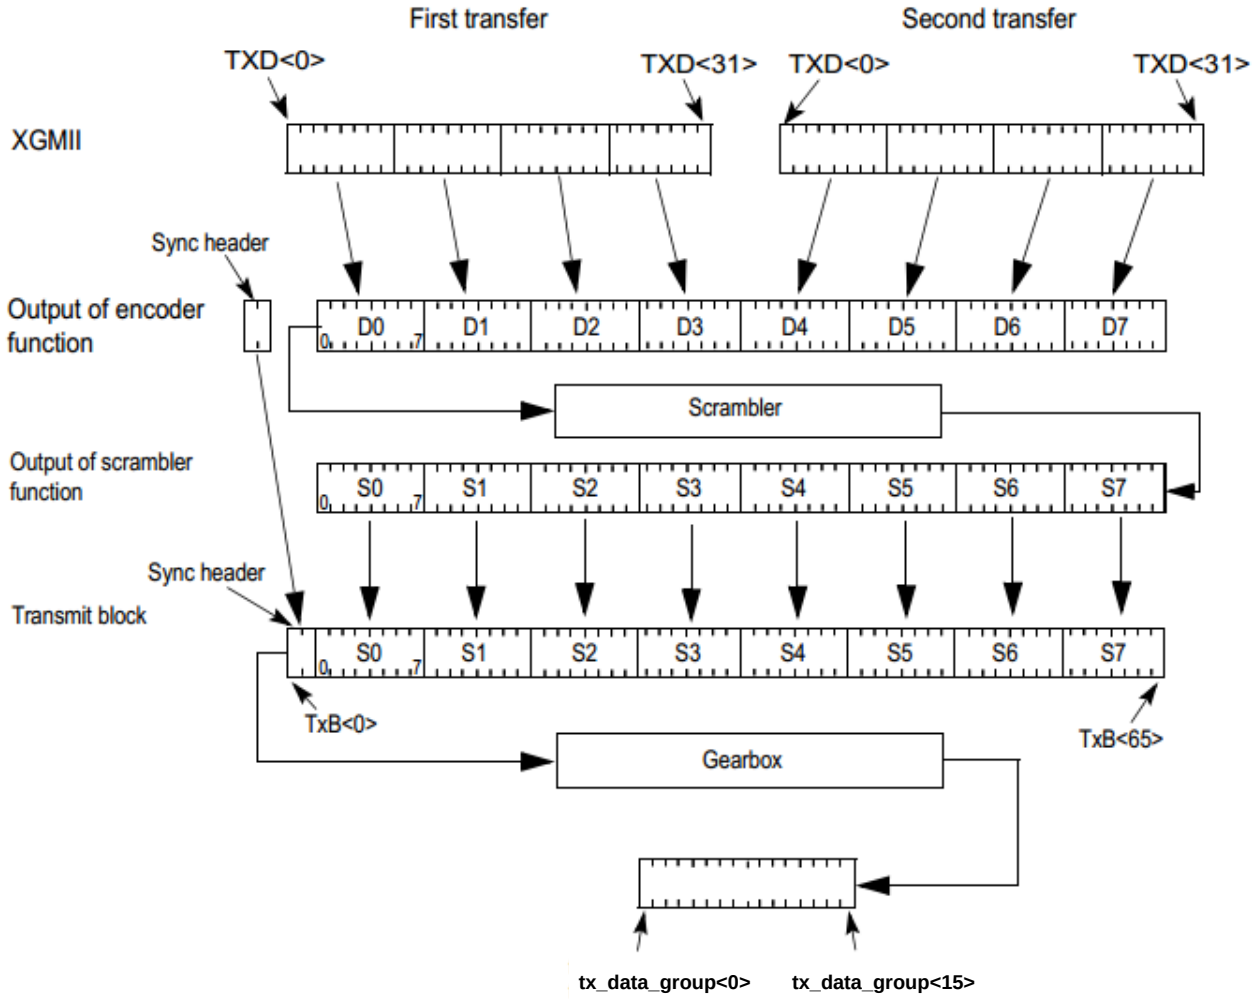

Figure 201-2 PCS Transmit bit ordering

FM editor's note: Copy figure 49-6 and modify text as sketched below:

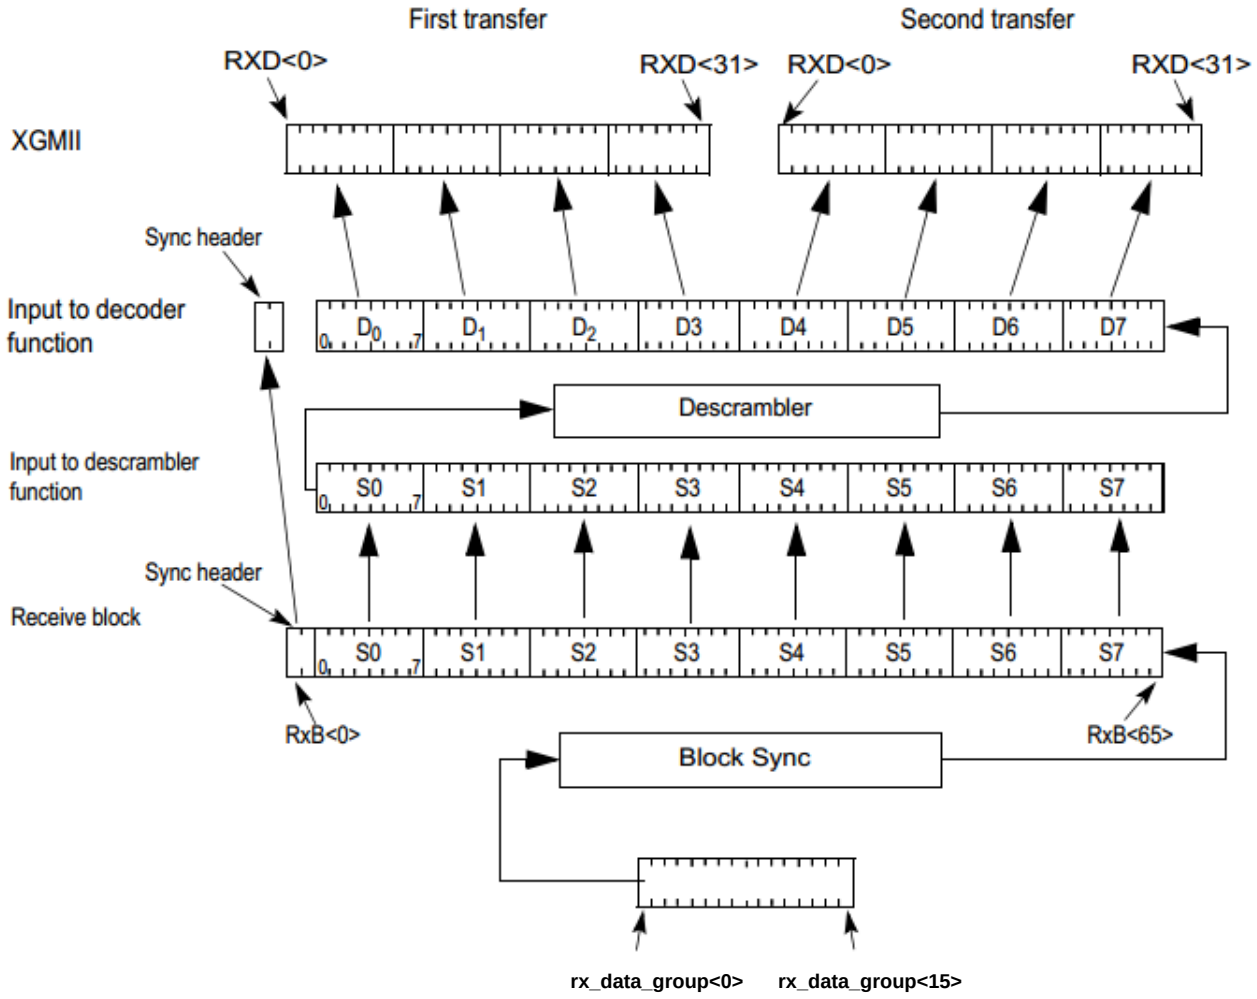

Figure 201-3 PCS Receive bit ordering

201.2.2 Notation conventions

Values represented in binary are shown with the first transmitted bit (the LSB) on the left.

Two consecutive XGMII transfers provide eight characters that are encoded into one 66-bit transmission block.

201.2.3 Transmission order

Block bit transmission order is illustrated in Figure 201-2 and Figure 201-3. Note that these figures show the mapping from XGMII to 64B/66B block for a block containing eight data characters.

201.2.4 Low Power Idle

If the 25GBASE-R PCS is part of a PHY configured for EEE operation, the PCS shall follow the state diagrams specified in Figure 49-12 and Figure 49-13.

The timer values for these state diagrams are shown in Table 201-1 for transmit and Table 201-2 for receive.

Table 201-1 Transmitter LPI timing parameters

FM editor's note – copy table 49-2 as is.

Parameter	Description	Min	Max	Units
T_{SL}	Local Sleep Time from entering the TX_SLEEP state to when tx_mode is set to QUIET	4.9	5.1	µs
T_{QL}	Local Quiet Time from when tx_mode is set to QUIET to entry into the TX_ALERT state	1.7	1.8	ms
T_{WL}	Time spent in the TX_WAKE state	10.9	11.1	µs
T_{1U}	Time spent in the TX_ALERT and TX_SCR_BYPASS states	1.1	1.3	µs

Table 201-2 Receiver LPI timing parameters

FM editor's note – copy table 49-3 as is.

Parameter	Description	Min	Max	Units
T_{QR}	The time the receiver waits for energy_detect to be set to TRUE while in the RX_SLEEP and RX_QUIET states before asserting receive fault	2	3	ms
T_{WR}	Time the receiver waits in the RX_WAKE state before indicating a wake time fault (when scr_bypass_enable = FALSE)	—	11.5	µs
T_{WR}	Time the receiver waits in the RX_WAKE state before indicating a wake time fault (when scr_bypass_enable = TRUE)	—	13.7	µs
T_{WTF}	Wake time fault recovery time	—	10	ms

201.3 Physical Medium Attachment (PMA) sublayer

201.3.1 Service Interface

The Serial PMA provides a Service Interface to the 5G BASE-R PCS. These services are described in an abstract manner and do not imply any particular implementation. The PMA Service Interface shall support the exchange of data-groups between the PMA and the PMA client. The PMA converts data-groups into bits and passes these to the PMD, and vice versa. It also generates an additional status indication for use by its client.

The following primitives are defined:

PMA_UNITDATA.request(tx_data-group<15:0>)
PMA_UNITDATA.indication (rx_data-group<15:0>)
PMA_SIGNAL.indication(SIGNAL_OK)
PMA_RXMODE.request(rx_mode)
PMA_TXMODE.request(tx_mode)
PMA_ENERGY.indication(energy_detect)

The definitions of these primitives are found in Clause 51.2.

201.3.2 Functions within the PMA

The PMA comprises the PMA Transmit and PMA Receive processes for 5GBASE-R. The PMA Transmit process serializes the tx_data-groups and passes them to the PMD for transmission on the underlying medium. Similarly, the PMA Receive process deserializes received data from the PMD and presents the data as rx_data-groups to the PMA client. The PMA receiver continuously conveys sixteen-bit data-groups to the PMA client, independent of data-group alignment.

201.3.2.1 PMA transmit function

The PMA Transmit function passes data unaltered (except for serializing) from the PMA client directly to the PMD. Upon receipt of a PMA_UNITDATA.request primitive, the PMA Transmit function shall serialize the sixteen bits of the tx_data-group<15:0> parameter and transmit them to the PMD in the form of sixteen successive PMD_UNITDATA.request primitives.

201.3.2.2 PMA receive function

The PMA Receive function passes data unaltered from the PMD directly to the PMA client. Upon receipt of sixteen successive PMD_UNITDATA.indication primitives, the PMA shall assemble the sixteen received bits into a single sixteen-bit value and pass that value to the PMA client as the rx_data-group<15:0> parameter of the primitive PMA_UNITDATA.indication. The PMA receive function does not align rx_datagroup<15:0> to the original tx_data-group<15:0> from the remote end of the link.

201.3.3 PMA loopback mode (optional)

PMA loopback is optional. If PMA loopback is implemented it shall conform to the requirements of clause 51.8.

201.4 Compatibility considerations

There is no requirement for a compliant device to implement or expose any of the interfaces specified for the PCS or PMA. Implementations of a XGMII shall comply with the requirements as specified in Clause 46.

201.5 Delay constraints

Predictable operation of the MAC Control PAUSE operation (Clause 31, Annex 31B) demands that there be an upper bound on the propagation delays through the network. This implies that MAC, MAC Control sublayer, and PHY implementers must conform to certain delay maxima, and that network planners and administrators conform to constraints regarding the cable topology and concatenation of devices.

The sum of transmit and receive delay contributed by the 5GBASE-R PCS and PMA shall be no more than 3584 BT.

The reference point for all MDI measurements is the 50% point of the mid-cell transition corresponding to the reference bit, as measured at the MDI.

201.6 Environmental specifications

Editor's Note: Need input on whether this is adequate

All equipment subject to this clause shall conform to the requirements of 51.9.

201.7 Protocol implementation conformance statement (PICS) proforma for Clause 201, Physical Coding Sublayer (PCS) and Physical Medium Attachment (PMA) sublayer, type 5GBASE-R

Editor's Note: To be inserted later by PICS editor after text stabilizes