OAM channel proposal for GEPOF

David Ortiz
Rubén Pérez-Aranda
(rubenpda@kdpof.com)

IEEE 802.3bv Task Force - Plenary, March 2015

Agenda

- Background and objectives
- Introduction
- Definitions
- TX operation
- RX operation
- PHY to PHY messages
- Step by step transmission example
- Transmission and handshaking block diagram
- Formal definition - state diagrams

Background & Objectives

- OAM was adopted as an objective for 802.3bp TF in Nov 2014 as a requirement of a big representation of car makers, but without any detailed description about the implementation (see [5])
- Automotive application was adopted as an objective for 802.3bv in Nov 2014, therefore we take the OAM channel requirement for 802.3bp from OEMs as a requirement also valid for 802.3bv TF
- The objective of this presentation is to define a mechanism to reliably exchange OAM messages between OAM peers sharing a link
- OAM message definition is not covered by this presentation and it is something up to the management entities implementer
- OAM message exchange is going to occur only in the PCS layer
  - Normally, when OAM is mentioned it occurs in the link layer or higher

OAM channel - introduction

- This presentation defines the OAM messages supported by the PHY. It defines the protocol to exchange the messages, and the mechanism needed to transport them. The format of the messages themselves is not covered in this presentation.
- The OAM messages are transmitted as part of the physical header data (PHD) block
  - The OAM message transmission channel is available as soon as both the local and remote PHY achieve the rcvr_hdr_lock = OK state; see [4]
  - It is active also during low power idle (LPI)
- It benefits from the robustness of the PHD transmission mechanism demonstrated in [4], allowing transmission of OAM messages in many situations in which the PHY devices may not reach the state of link_status = OK due to a bad channel condition.
- Two handshake mechanisms are defined to guarantee the integrity of the different OAM messages:
  - Transmission Handshake: there is a handshaking mechanism defined for transmitting a new message that guarantees the integrity of the message to be transmitted. It also implements flow control mechanism, locking the resources until they are free for transmission of a new message.
  - Reception Handshake: there is a handshaking mechanism defined to receive a new message that guarantees the integrity of the message, and also serves to lock the OAM Rx registers until the complete message has been completely processed by the Management Entity.

OAM channel - definitions

- Management Entity: each of the OAM agents that exchange information through the OAM message mechanism defined in this presentation.
- OAM Tx registers: registers used by the Management Entity to transmit a message to its peer at the other end of the link.
- OAM Rx registers: registers where the messages from the remote peer are reported.
- OAM message: set of 128 user data bits plus 12-bit user control bits plus the bits necessary to perform the message transmission and guarantee the required message integrity.

OAM channel - TX operation

- The Tx registers will be accessed through the management interface (i.e. Clause 45 MDIO), which provides atomic access to 16-bit width registers. A mechanism is implemented to guarantee message integrity, as it spans several MDIO registers.
- There are four control bits that are used to implement the OAM message handshake mechanism, and also give information on the status of previous message transmissions:
  - TXREQ: Transmission ReQuest. It is raised by the Management Entity to request the transmission of a new message. It is set to 0 by the local PHY to indicate that it has started the transmission of the message.
  - PHYT: PHY Toggle bit. Corresponds to the toggle bit of the last message received by the remote PHY.
  - MERT: Management Entity Read Toggle bit. Toggle bit of the last message read by the Management Entity attached to the remote PHY.
  - MSGT: MeSsaGe Toggle bit. Toggles with each successive transmitted message and is used for message identification.

OAM Message Tx Registers

TXREQ (RW)	PHYT (RO)	MERT (RO)	MSGT (RO)	OAM_CTRL (12 bits)	Register TxTBD0.15:0
OAM_DATA0 (16 bits)					Register TxTBD1.15:0
OAM_DATA1 (16 bits)					Register TxTBD2.15:0
OAM_DATA2 (16 bits)					Register TxTBD3.15:0
OAM_DATA3 (16 bits)					Register TxTBD4.15:0
OAM_DATA4 (16 bits)					Register TxTBD5.15:0
OAM_DATA5 (16 bits)					Register TxTBD6.15:0
OAM_DATA6 (16 bits)					Register TxTBD7.15:0
OAM_DATA7 (16 bits)					Register TxTBD8.15:0

OAM channel - TX operation

- When the Management Entity wants to transmit a new message:

1. Wait until bit TxTBD0.15 (TXREQ) is low, which indicates that the OAM Tx Registers are free to accept a new OAM Message.
2. Write message to transmit in registers TxTBD1 to TxTBD8.
3. Write OAM_CTRL field in TxTBD0.11:0, setting also TxTBD0.15 bit to 1 to indicate the request to transmit the new OAM Message.
4. When the PHY accepts the message for transmission it will lower TxTBD0.15 bit to 0. Simultaneously it will toggle bit TxTBD0.12 (MSGT) to notify the Management Entity of the MSGT to be associated with the message.
5. At this moment the OAM Tx Registers are free, and the local Management Entity can write a new message to them.
6. Bits PHYT and MERT indicate at every moment the last message being acknowledged by the remote PHY and remote Management Entity respectively.

OAM channel - TX operation

- One of the benefits of the proposed OAM message implementation is that the complete status of the OAM channel can be expressed by 4 bits, which are part of the OAM Tx control registers.
- There is no need to store previous values of the Tx registers to infer the situation of previously transmitted or outstanding OAM messages, since all the information is captured by the current value of the registers.
- Up to 3 OAM messages can live simultaneously in the system per direction, by using 3 different memory locations: Tx registers, PHD being transmitted and Rx registers.
- The table of the following slide exemplifies all the possible status combinations of the OAM channel with the corresponding values of the status bits of the OAM channel Tx control registers.
  - The value of variable ***a*** can take values 0 or 1 (**~** denotes logical not operator).
  - Each row is independent of each-other.
  - For each row:
    - Message K+1 corresponds to the message that has been written to the Tx Registers from the local Management Entity (if it exists). No MSGT has been associated with it yet.
    - Message K corresponds to the message being transmitted from the local PHY. It has MSGT value ***a***.
    - Message K-1 corresponds to the previous message transmitted from local PHY. It has MSGT value ***~a***.

OAM channel - TX operation

TXREQ	PHYT	MERT	MSGT	Message K+1 Status	Message K Status	Message K-1 Status
0	***a***	***a***	***a***	Not set.	Sent. ACK by remote PHY. ACK by remote ME.	Sent. ACK by remote PHY. ACK by remote ME.
0	***~a***	***~a***	***a***	Not set.	Sent. No ACK by remote PHY. No ACK by remote ME.	Sent. ACK by remote PHY. ACK by remote ME.
0	***a***	***~a***	***a***	Not set.	Sent. ACK by remote PHY. No ACK by remote ME.	Sent. ACK by remote PHY. ACK by remote ME.
0	***~a***	***a***	***a***	Not set.	Sent. No ACK by remote PHY. No ACK by remote ME.	Sent. ACK by remote PHY. No ACK by remote ME.
1	***a***	***a***	***a***	Written by local ME. Pending transmission by local PHY.	Sent. ACK by remote PHY. ACK by remote ME.	Sent. ACK by remote PHY. ACK by remote ME.
1	***~a***	***~a***	***a***	Written by local ME. Pending transmission by local PHY.	Sent. No ACK by remote PHY. No ACK by remote ME.	Sent. ACK by remote PHY. ACK by remote ME.
1	***a***	***~a***	***a***	Written by local ME. Pending transmission by local PHY.	Sent. ACK by remote PHY. No ACK by remote ME.	Sent. ACK by remote PHY. ACK by remote ME.
1	***~a***	***a***	***a***	Written by local ME. Pending transmission by local PHY.	Sent. No ACK by remote PHY. No ACK by remote ME.	Sent. ACK by remote PHY. No ACK by remote ME.

OAM channel - RX operation

- The Rx registers will be accessed through the management interface (i.e. Clause 45 MDIO), which provides atomic access to 16-bit registers. A mechanism is implemented to guarantee message integrity, as it spans several MDIO registers, and the Management Entity needs several access cycles to read the message. All the Rx registers are Read-Only.
- There are two control bits that are used to implement the OAM message handshake mechanism:
  - RXVAL: Reception VALid. It is raised by the PHY to indicate to the Management Entity that a new message has arrived and is ready for processing. The message will be locked in the Rx Registers until the Management Entity access register RxTBD8, at which point the PHY will clear bit RXVAL and will also acknowledge the message processing to the link partner (using bit MERT).
  - MSGT: MeSsaGe Toggle bit. Indicates the toggle bit associated with the received message, and serves as message identification. It will toggle with every new received message.

OAM Message RX Registers

RXVAL			MSGT	OAM_CTRL (12 bits)	Register RxTBD0.15:0
OAM_DATA0 (16 bits)					Register RxTBD1.15:0
OAM_DATA1 (16 bits)					Register RxTBD2.15:0
OAM_DATA2 (16 bits)					Register RxTBD3.15:0
OAM_DATA3 (16 bits)					Register RxTBD4.15:0
OAM_DATA4 (16 bits)					Register RxTBD5.15:0
OAM_DATA5 (16 bits)					Register RxTBD6.15:0
OAM_DATA6 (16 bits)					Register RxTBD7.15:0
OAM_DATA7 (16 bits)					Register RxTBD8.15:0

OAM channel - RX operation

- When the Management Entity wants to read a new message:
  1. Read register RxTBD0
  2. If bit RxTBD0.15 (RXVAL) is high it indicates that a new OAM message has arrived and is ready to be processed by the Management Entity, then
     - The Management Entity reads the registers RxTBD1 to RxTBD8, in this order
     - Read operation of RxTBD8 produces the PHY clears the bit RXVAL and acknowledges to the remote peer that the OAM message has been processed
  3. If bit RxTBD0.15 (RXVAL) is low, go to step 1

OAM channel - PHY to PHY message

- The OAM message is carried as part of the PHD. See [3] for full PHD content defintion.
- When no new message information has to be transmitted, the PHD keeps retransmitting the same information of the previous PHD block, ensuring that the information reaches the destination even in the event of a corruption in the PHD transmission. See [4] for detailed PHD reliability analysis.
- The OAM message information transmitted in the PHD can be divided in two sub-groups:
  - Message identification bit (MSGT) plus OAM Message contents. They carry the contents of the last message sent by the local Management Entity through the Tx Registers. When no message has been transmitted since the last reset of the physical layer they are set to 0.
  - OAM Status reporting bits (PHYT plus MERT). They contain respectively the MSGT bit corresponding to the last message received by this PHY from the link partner, and the last message being processed by the Local Management Entity through the OAM Rx Registers.

PHD OAM fields

OAM. PHYT	OAM. MERT	OAM. MSGT	OAM.CTRL (12 bits)	15 bits
OAM.DATA0 (16 bits)				16 bits
OAM.DATA1 (16 bits)				16 bits
OAM.DATA2 (16 bits)				16 bits
OAM.DATA3 (16 bits)				16 bits
OAM.DATA4 (16 bits)				16 bits
OAM.DATA5 (16 bits)				16 bits
OAM.DATA6 (16 bits)				16 bits
OAM.DATA7 (16 bits)				16 bits

OAM channel - example of transmission A to B

- Step 1: Management Entity A writes OAM message Tx registers and raises TXREQ bit

OAM Tx registers

TXREQ	PHYT	MERT	MSGT
1	*a*	*a*	*a*

OAM Rx registers

RXVAL	MSGT
0	*a*

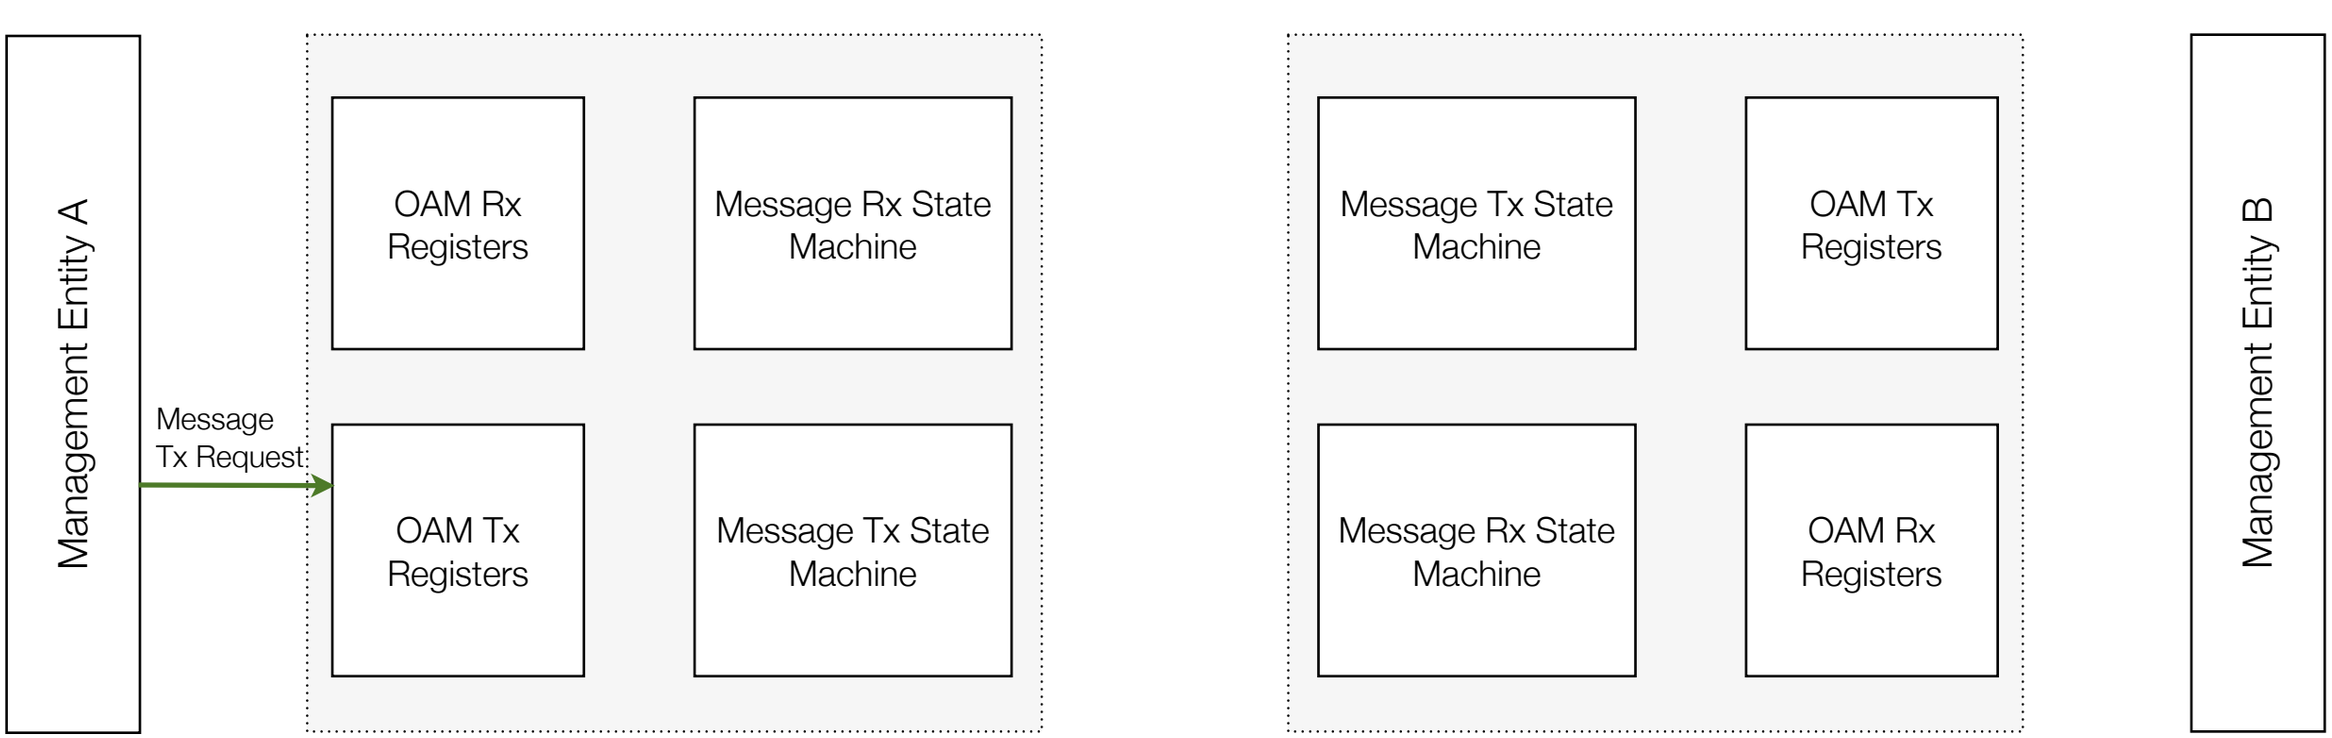

OAM channel - example of transmission A to B

- Step 2: local PHY (A) acknowledges the message and transmit it to the remote PHY (B)

OAM Tx registers

TXREQ	PHYT	MERT	MSGT
0	*a*	*a*	*~a*

OAM Rx registers

RXVAL	MSGT
0	*a*

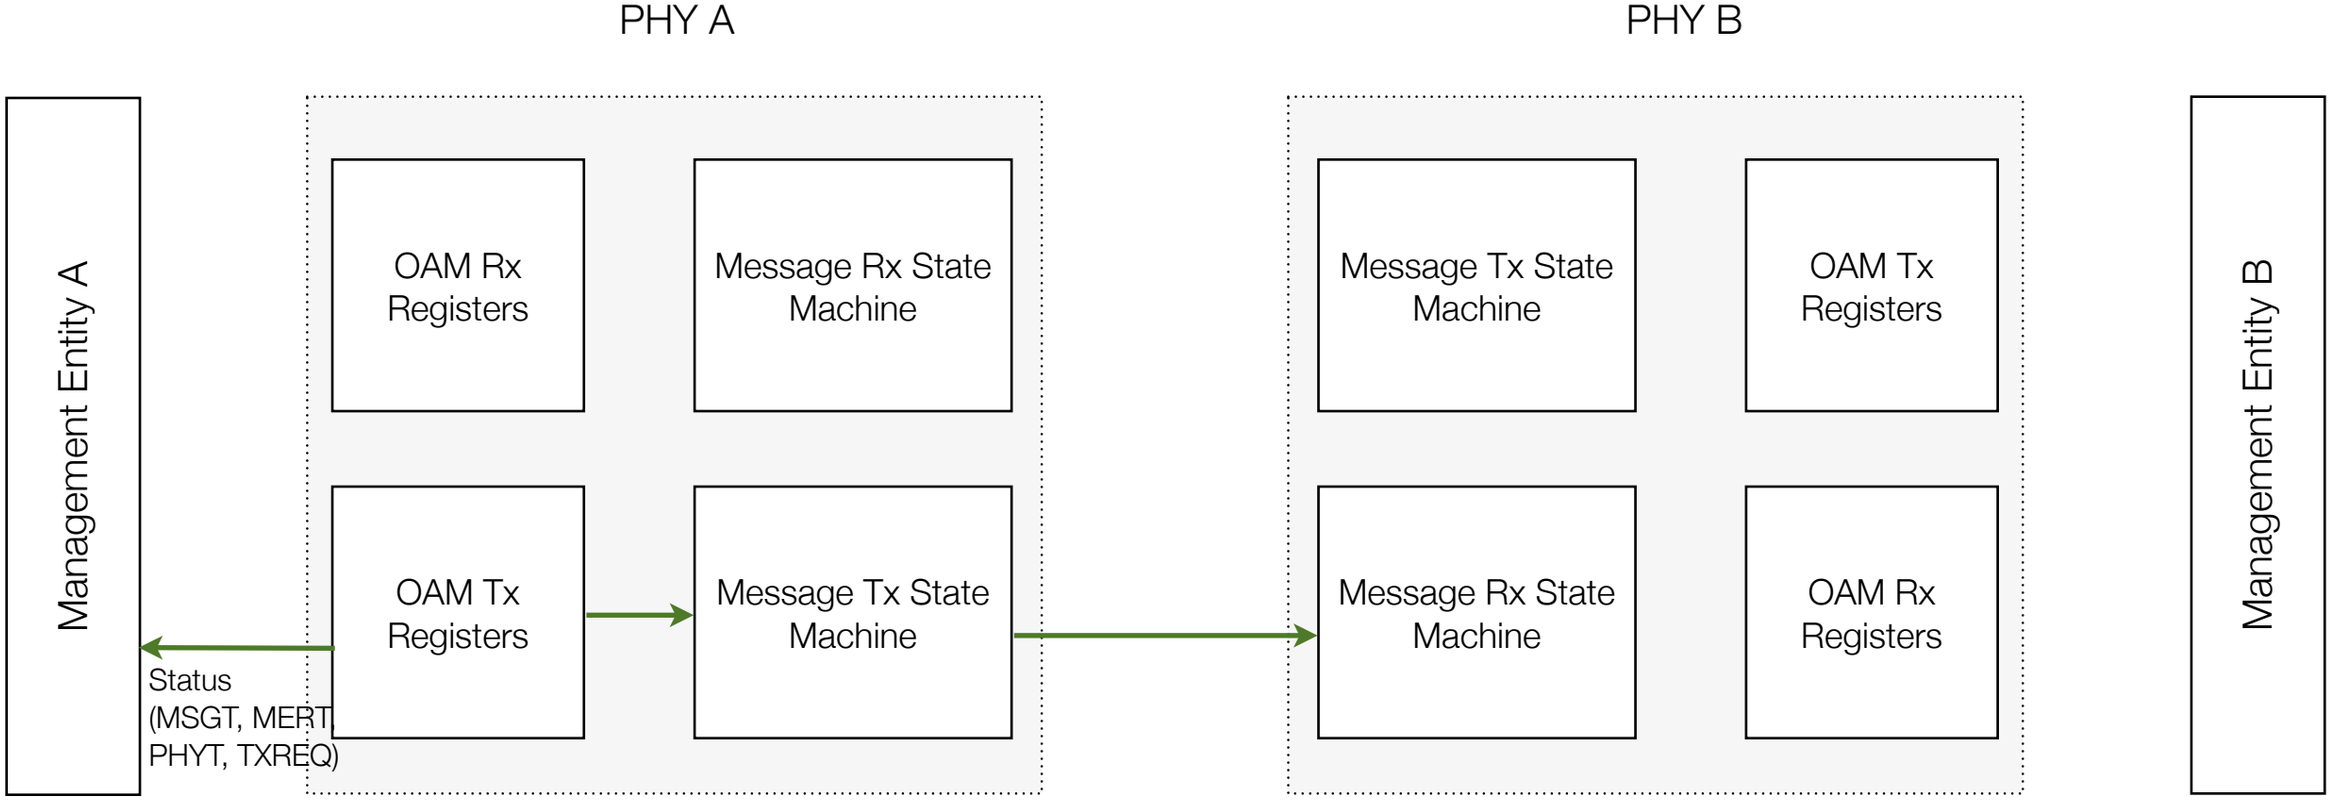

OAM channel - example of transmission A to B

- Step 3: remote PHY (B) receives the message. It signals the reception of the message to the remote Management Entity (B) and sends PHYT bit to the PHY A

OAM Tx registers

TXREQ	PHYT	MERT	MSGT
0	*~a*	*a*	*~a*

OAM Rx registers

RXVAL	MSGT
1	*~a*

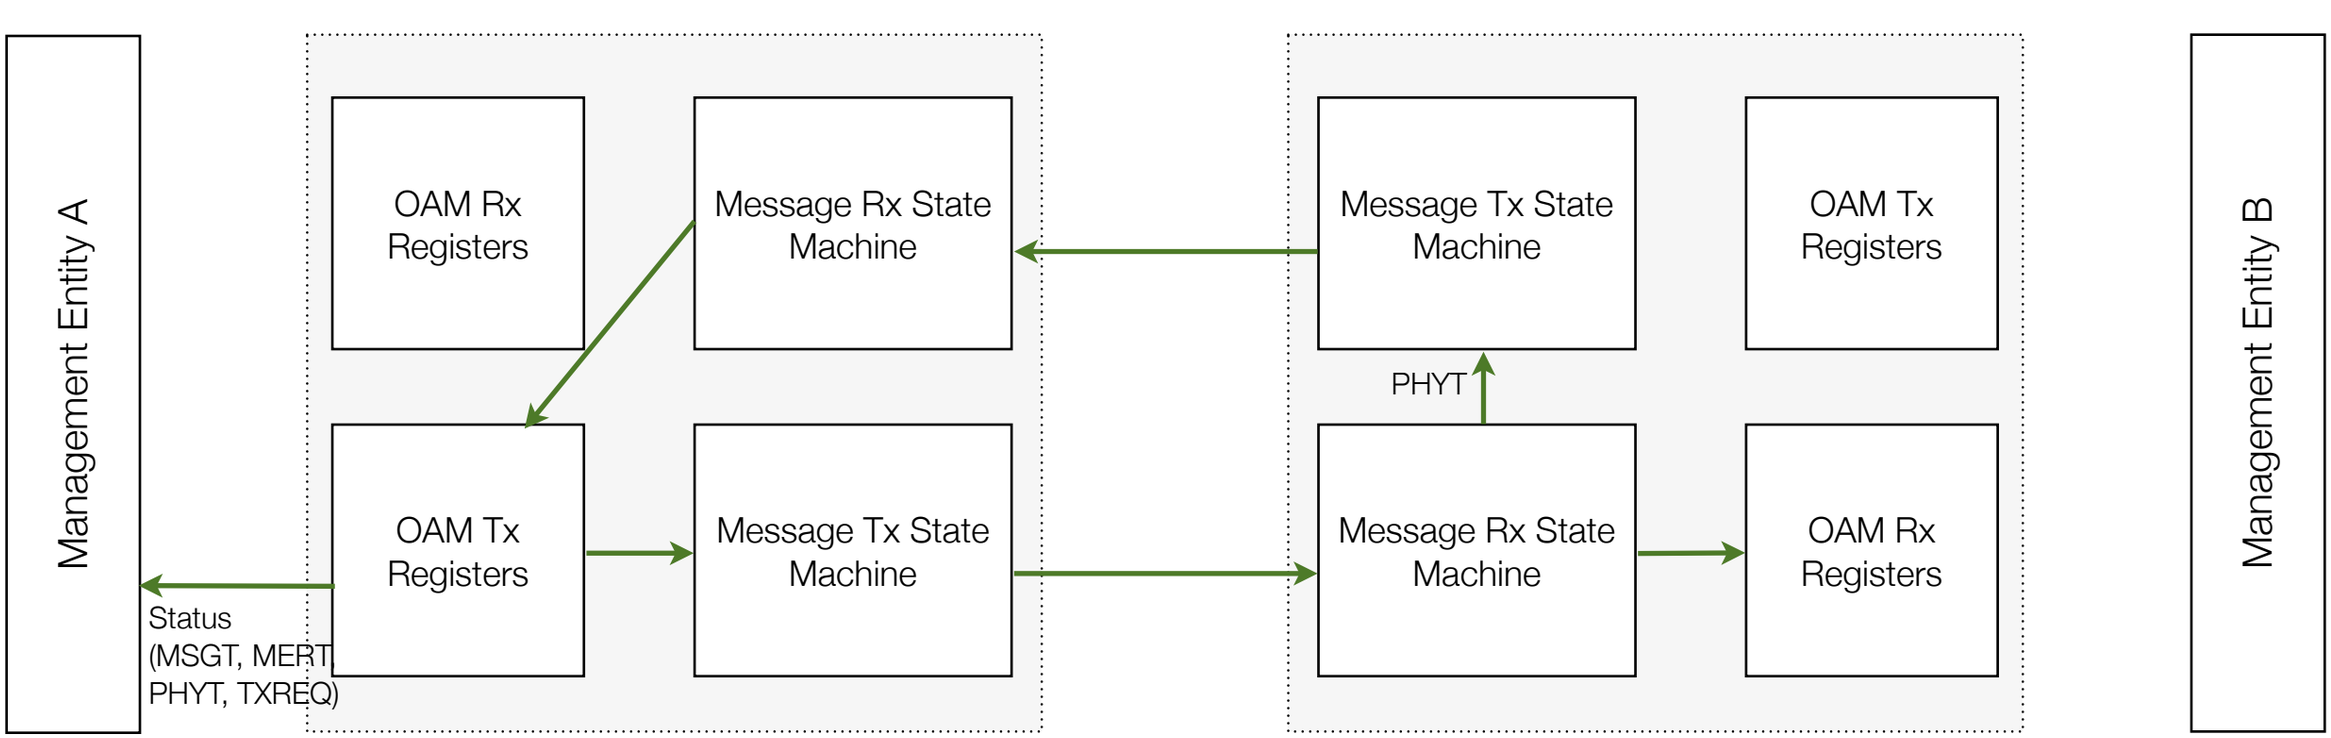

OAM channel - example of transmission A to B

- Step 4: remote Management Entity (B) reads the message, and the message read toggle bit reaches the local PHY (A) and is signaled in the local OAM Tx registers

OAM Tx registers

TXREQ	PHYT	MERT	MSGT
0	*~a*	*~a*	*~a*

OAM Rx registers

RXVAL	MSGT
0	*~a*

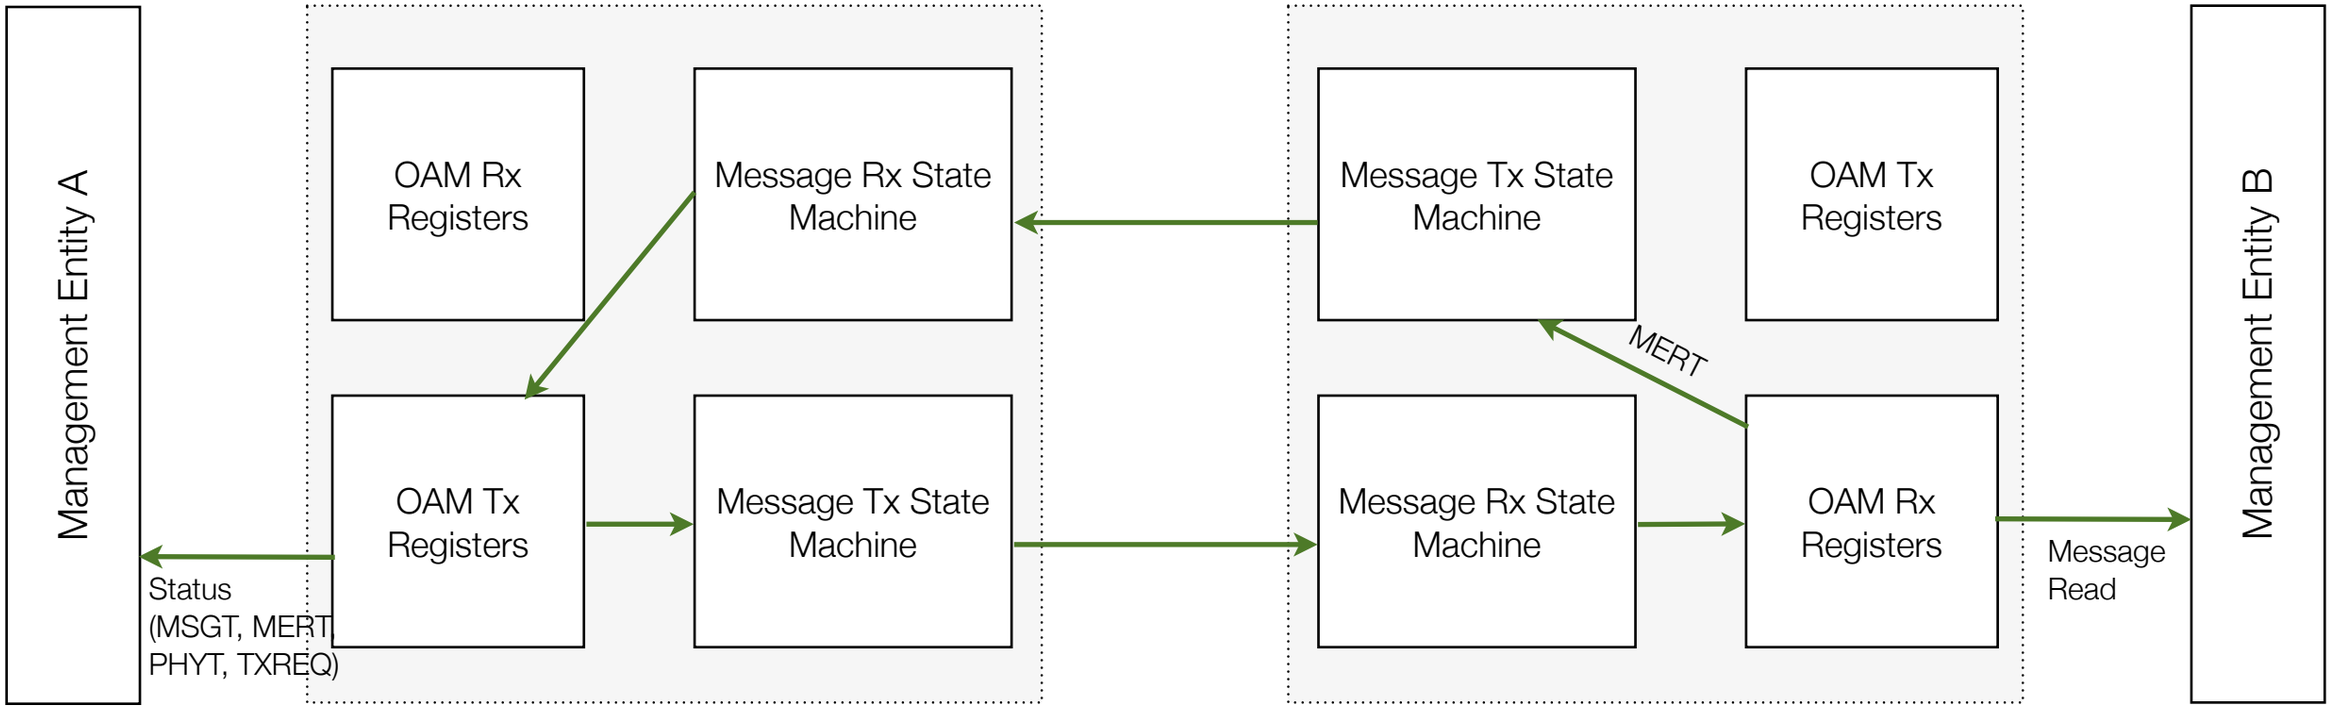

OAM channel - transmission and handshaking

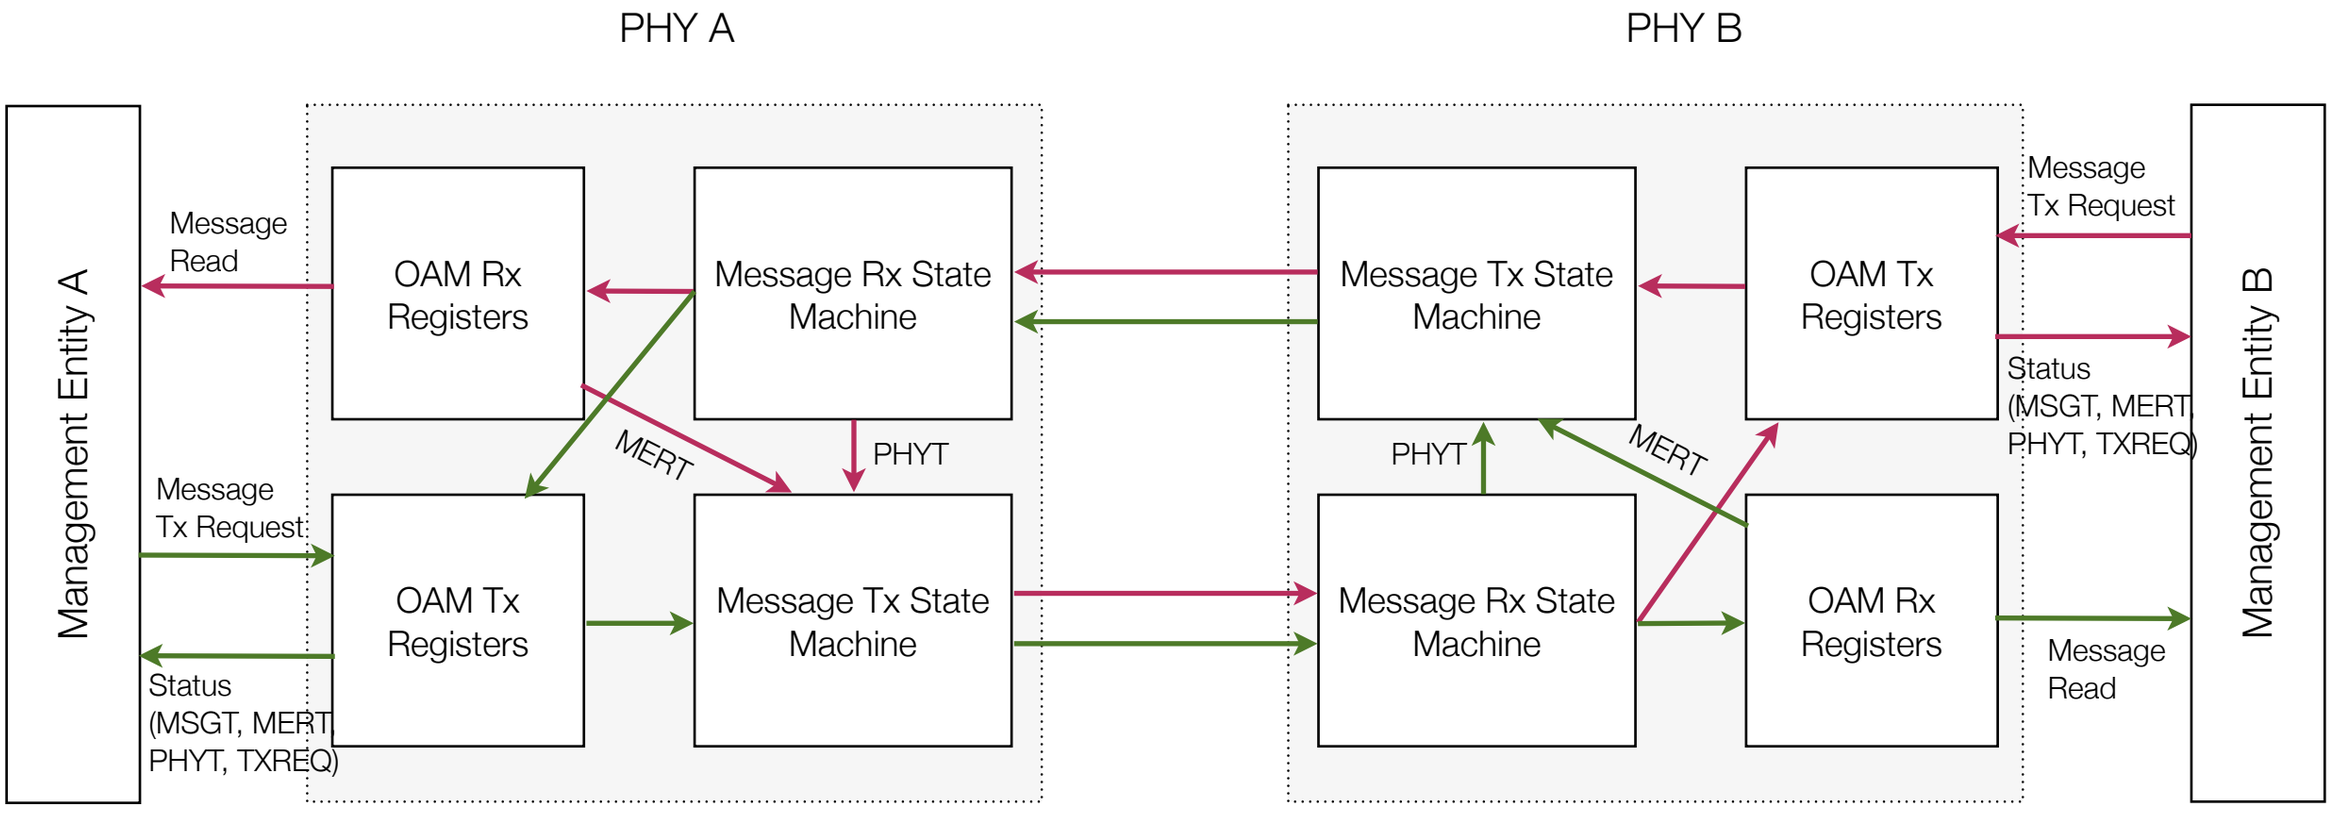

OAM channel - TX state diagram

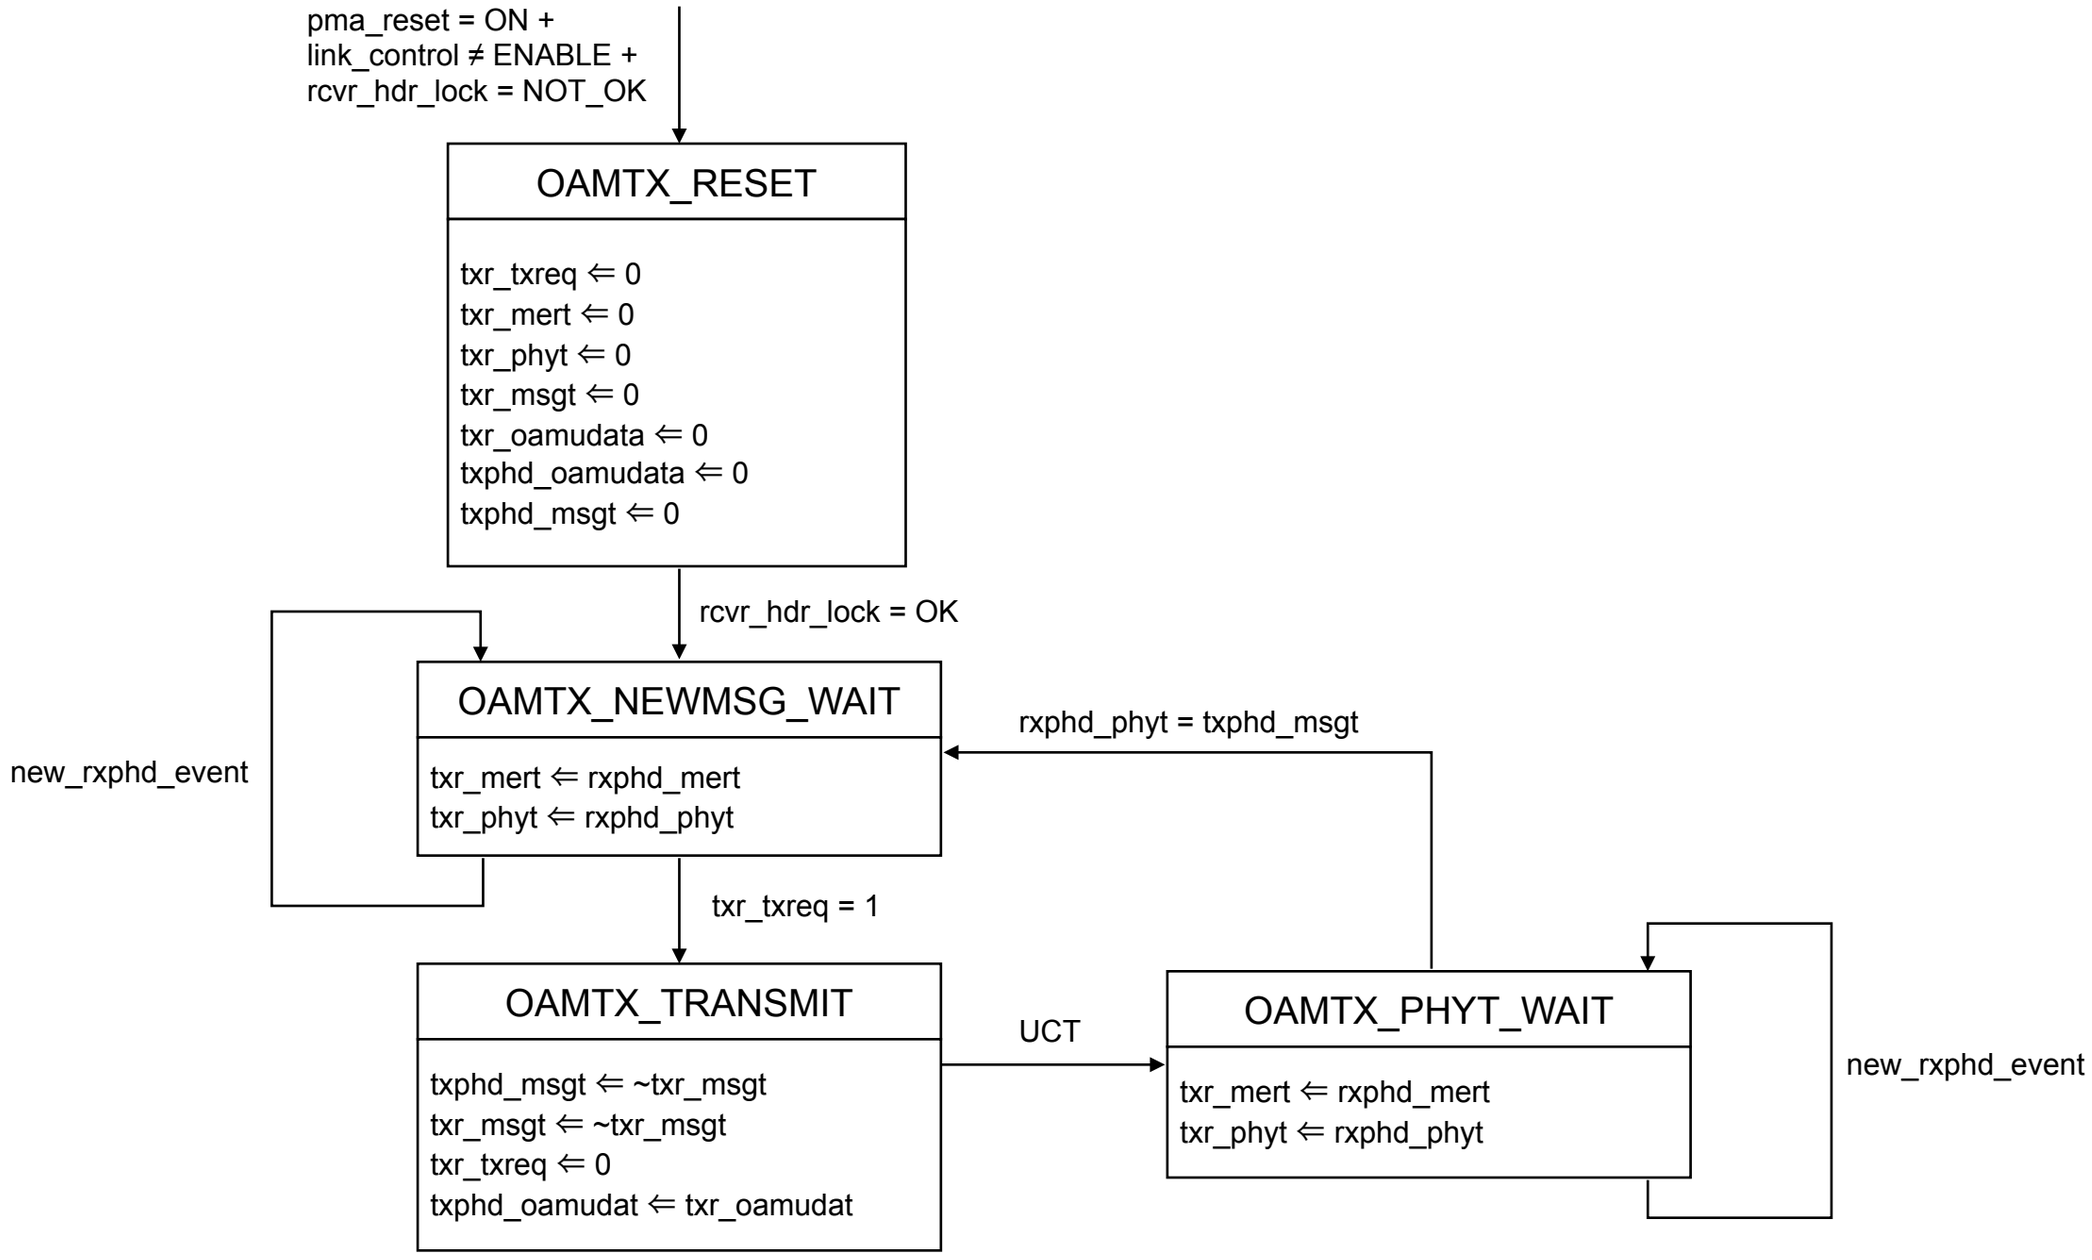

OAM channel - TX state variables

- pma_reset: allows reset of all the PMA functions. It is set by the PMA reset. PMA reset function is intended to be executed whenever one of power on or reset from management entity conditions occur. All state diagrams take the open-ended pma_reset branch upon execution of PMA Reset.
  - ON: reset is asserted
  - OFF: reset de-asserted
- link_control: controls the connection of PMA to the PMD. This control variable is foreseen for an eventual coexistence of the PHY with an autonegotiation sub-system.
  - DISABLE: isolates the PMA from the PMD
  - ENABLE: connects the PMA to the PMD (both transmitter and receiver)
- txr_txreq: value of bit TxTBD0.15. This bit indicates if the local Management Entity is requesting the transmission of a new OAM message
  - 0: there is no message transmission request from the local Management Entity.
  - 1: local Management Entity has written a new message to the Tx registers and is requesting the transmission to the local PHY.
- txr_msgt: value of bit TxTBD0.12. This bit is written by the OAM TX state machine to reflect the identifier of the message being sent by the local PHY. It alternates between values 0 and 1.
- txr_phyt: value of bit TxTBD0.14. This bit is written by the OAM TX state machine to reflect the identifier of the last message being acknowledged by the remote PHY. It alternates between values 0 and 1.
- txr_mert: value of bit TxTBD0.13. This bit is written by the OAM TX state machine to reflect the identifier of the last message that has been processed by the remote Management Entity. It alternates between values 0 and 1.

OAM channel - TX state variables

- txr_oamudat: payload of the OAM message to be transmitted. It corresponds to the content of registers TxTBD0.11:0 and TxTBD1.15:0 to TxTBD8.15:0. It can take any value.
- txphd_msgt: identifier of the message being transmitted by the local PHY. It corresponds to the content of header field PHD.OAM.MSGT. It alternates between values 0 and 1.
- txphd_oamudat: content of header fields PHD.OAM.CTRL and PHD.OAM.DATA0 to PHD.OAM.DATA7 being transmitted by the local PHY. It can take any value.
- rxphd_phyt: content of field PHD.OAM.PHYT of the last PHD that has been correctly received by the local PHY. It alternates between values 0 and 1.
- rxphd_mert: content of field PHD.OAM.MERT of the last PHD that has been correctly received by the local PHY. It alternates between values 0 and 1.
- new_rxphd_event: variable set by the PCS receive function to indicate the arriving of a new correct (i.e. CRC16 is OK) PHD block received from the link partner
  - 1: indicates the event of a new PHD received from link partner. The value 1 extends one receive symbol period
  - 0: indicates no new PHD was received
- rcvr_hdr_lock: variable set by the PMA Receive function to indicate whether reliable transmission and reception of PHD are detected
  - OK: PHD transmission and reception are reliable.
  - NOT_OK: PHD transmission or reception are unreliable.

OAM channel - RX state diagram

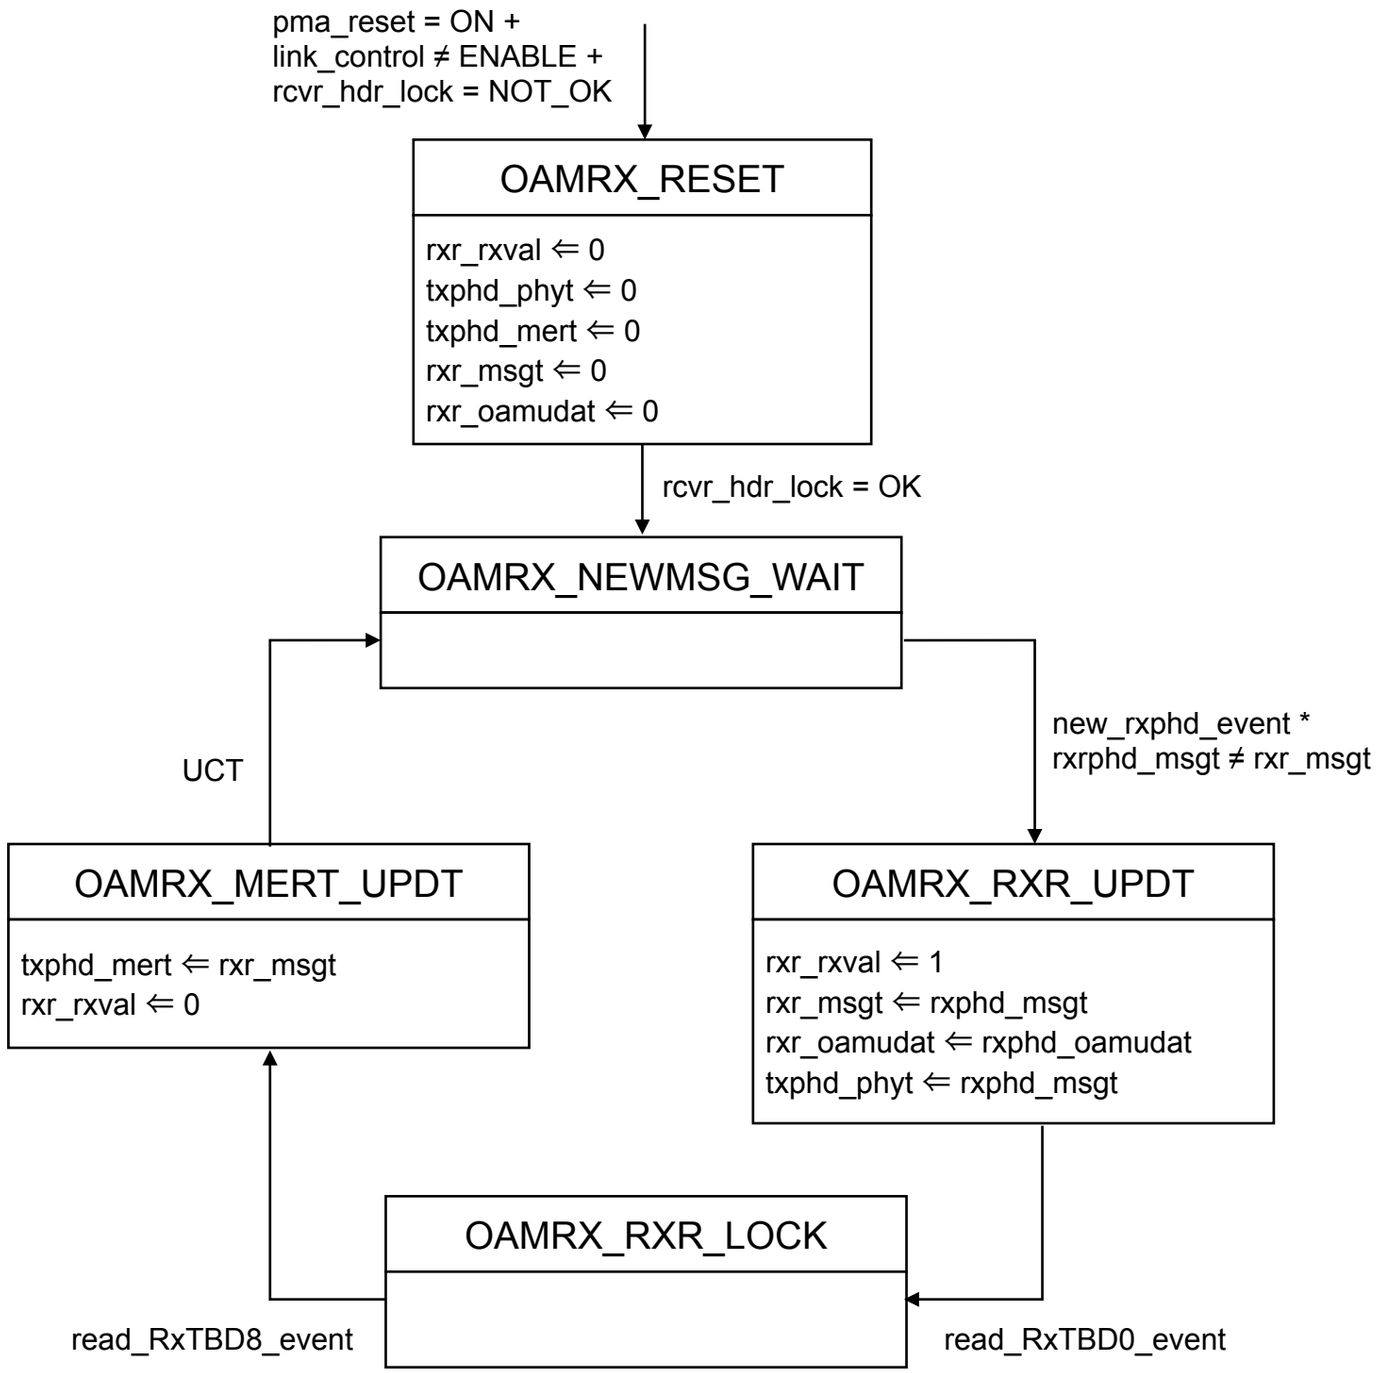

OAM channel - RX state variables

- rxr_rxval: value of bit RxTBD0.15. This bit indicates the presence of a valid message in the Rx registers
  - 1: there is a valid message on the Rx registers which is pending for processing from the local Management Entity.
  - 0: there is no valid message on the Rx registers.
- rxr_msgt: value of bit RxTBD0.12. This bit is the identifier of the message that is contained in the Rx registers. It alternates between 0 and 1.
- rxr_oamudat: content of registers RxTBD0.11:0 and RxTBD1.15:0 to TxTBD8.15:0. It can take any value.
- txphd_phyt: content of field PHD.OAM.PHYT being transmitted by the local PHY. It indicates to the remote PHY which is the last message that has been written to the Rx registers to be processed by the local Management Entity. It alternates between 0 and 1.
- txphd_mert: content of field PHD.OAM.MERT being transmitted by the local PHY. It informs to the remote PHY which is the last message that has been processed by the local Management Entity. It alternates between 0 and 1.
- rxphd_msgt: content of field PHD.OAM.MSGT of the last PHD that has been correctly received by the local PHY. It contains the message identifier of the OAM message carried in that PHD. It can take any value.
- rxphd_oamudat: content of fields PHD.OAM.CTRL and PHD.OAM.DATA0 to PHD.OAM.DATA7 of the last PHD that has been correctly received by the local PHY. It carries the payload of the OAM message. It can take any value.
- read_RxTBD0_event: event to indicate the RxTBD0.15:0 has been read.
- read_RxTBD8_event: event to indicate the RxTBD8.15:0 has been read.

References

- [1] *Rubén Pérez-Aranda, et al., "High spectrally efficient coded 16-PAM scheme for GEPOF based on MLCC and BCH", IEEE 802.3bv TF, Interim Meeting, January 2015*
- [2] *Rubén Pérez-Aranda, "Transmission scheme for GEPOF", IEEE 802.3bv TF, Interim Meeting, January 2015*
- [3] *Rubén Pérez-Aranda, et al., "Physical Header Data content for PCS encoding, PHY control and OAM channel in GEPOF", IEEE 802.3bv TF, Plenary Meeting, March 2015*
- [4] *Rubén Pérez-Aranda, et al., "PMA control state machines for GEPOF", IEEE 802.3bv TF, Plenary Meeting, March 2015*
- [5] *Kirsten Matheus, et al., "Proposal for an OAM channel", IEEE 802.3bp TF, Plenary Meeting, November 2014*

Questions?