1000BASE-RH OAM CLARIFICATION

David Ortiz
Rubén Pérez-Aranda

Background and Objectives

- In 802.3bv TF Interim meeting of January 2015 and Plenary meeting of March 2015 several parts of the baseline were adopted:
  - OAM Channel proposed in [1].
  - Transmission scheme defined in [2].
  - Content of Physical Header Data defined in [3].
- During the presentation of the OAM Channel it was questioned why the OAM channel proposed for 1000BASE-T1 were not also adopted for 1000BASE-RH.
- In this presentation the proposal adopted for 1000BASE-RH is compared to the proposal adopted for 1000BASE-T1 and the reasons why the proposal explained in [1] and [3] are better suited for 1000BASE-RH.

Main Differences

- Both 1000BASE-RH and 1000BASE-T1 have adopted a mechanism to allow the MAC entities at the ends of a single link exchange information, not during the link establishment (as is possible through page exchanges during auto-negotiation in different Ethernet clauses), but during normal operation.
- Both proposals try to satisfy the same end-user requirement. Because of that both proposals are similar.
- However there are important differences between them, as the top level requirement is implemented over completely different side-channels. The main differences can be summarized as:
  - OAM Message Size.
  - Ping Functionality.
  - SNR/Link status reporting.
  - Parity bit + CRC16.
  - OAM Message acknowledgment Implementation.
- This presentation try to explain the reasons for those differences and the reasoning behind the OAM channel proposal adopted for 1000BASE-RH.

OAM Message Transmission Side-Channel

- To understand the differences between the two OAM implementations it is essential to know the differences in the physical layer used to carry them in each standard.
  - In 1000BASE-RH the OAM messages are transmitted as payload of the PHD.
    - The length of the PHD is fixed in 1000BASE-RH. It has to carry some essential information needed for synchronization, channel equalization, etc.
    - The available bandwidth in the PHD side-channel free for accessory functionality (like OAM messages and remote status reporting) is approximately 1Mbps.
    - As shown in [4] the PHD transmission side-channel is more robust than the channel used to transmit the GMII data (operates with extra margin larger than 10dB). The data contained in a single PHD is protected by a CRC16.
    - As the length of the PHD is fixed in 1000BASE-RH, and tied to the structure of the transmit block, the latency of this side-channel it is also fixed to the duration of a transmitter block, and equal to approximately 695 usec. This latency is fixed independently of the length of the information sent (up to the limit of the data that fits in a single PHD).
    - The transmit block provides the synchronization mechanism needed to properly decode and interpret the information contained in a PHD.
  - In 1000BASE-T1 the OAM messages are transmitted using the OAM9 bits available in each RS code-word that are not used to transmit user data. In EEE operation when the link is in LPI state, the information is sent as part of the refresh cycle.
    - The total bandwidth available in that channel is about 2.5Mbps when link is operating in normal mode. Less if the link is operating in LPI mode.
    - The side-channel does not provide synchronization. To decode information sent through this message some sort of message synchronization is needed.
    - The robustness of the side-channel is equal to the robustness of the channel used to transport the GMII information in normal mode, robustness is lower during the refresh cycles.

Differences I

- OAM Message size:
  - In the OAM proposal adopted as baseline for 1000BASE-RH, the number of user data bits per OAM Message is 140 bits, whereas in the OAM proposal adopted for 1000BASE-T1 the number of user-data bits is lower.
  - The reason of going for a bigger message size in 1000BASE-RH is to optimize the available bit rate to transmit the OAM message as selecting a smaller message size would have resulted in a reduction of the bitrate of this function. The bitrate available for OAM Messages in 1000BASE-RH is approximately 200 Kbps.
  - The selection of a bigger message size does not impact the latency of the transmission. At the same time and thanks to the robustness of the channel, it also does not impact the error probability, as shown in [4] this probability is negligible.
- Ping Functionality:
  - In 1000BASE-RH the PHD.RX.HDRSTATUS already indicates if the PHD side-channel is available. This information is available to the MAC entities, and also to the OAM partners. Thanks to the robustness of the PHD side-channel this renders unnecessary any ping functionality.
- SNR/Link Status reporting.
  - In 1000BASE-RH the status reporting is integrated with the rest of the remote status reporting that is indicated in the PHD fields and not part of the OAM functionality.
- Parity bit + CRC16.
  - In 1000BASE-T1 the parity and CRC16 are needed to provide a synchronization mechanism and to protect the messages against the possible corruption caused during transmission. In 1000BASE-RH the PHD side-channel already provides the needed synchronization, and already provides a PHD level CRC16. The OAM Message is carried just like any other field in the PHD and is protected by the same mechanisms.

Differences II

- OAM Message Acknowledgment Implementation:
  - In 1000BASE-T1 4 bits are sent through the OAM side-channel which are used for OAM Message Management. Those 4 bits provide the following functionality:
    - Message identification.
    - Indication of when the remote PHY has loaded the message in the remote MDIO registers.
    - Flow Control.
  - In 1000BASE-RH three bits are sent through the side-channel for OAM Message management. Those three bits provide the following functionality:
    - Message identification.
    - Indication of when the remote PHY has loaded the message in the remote MDIO registers.
    - Direct indication of when the remote Management Entity has read the message in the remote MDIO registers.
    - Flow Control.
- OAM Message MDIO register handshaking:
  - In both proposals the OAM message occupies several MDIO registers, so access from the Management Entity to the OAM message registers is not atomic.
  - It is essential to ensure integrity of the message handling between Management Entity and the PHY, both for message write and for message read. To maintain the message integrity both proposals implement a handshaking mechanism.

Differences III

- OAM Message status:
  - In both solutions the Management Entities need to know the status of previous OAM messages. The implementation is different in each proposal:
    - In 1000BASE-T1 the management entity needs to perform the following steps to check that a message is available for read by the remote Management Entity:
      - Read and store the toggle bit of the transmit message to be checked in register bit 3.TBD0.14 (97.7.3.1.2 of http://www.ieee802.org/3/bp/public/feb15/Lo_3bp_01a_0215.pdf).
      - Poll bit 3.TBD0.13 until it is 1, that indicate that the message with toggle bit indicated at 3.TBD0.12 has reached the remote Phy. The bit is self-cleared after read. There is no indication of when/if the remote Management Entity has read the message.
    - In 1000BASE-RH the Management Entity can not only know if the Transmit registers are available to accept a new message, but also know the status of all the potentially outstanding previous messages. To do this the Management Entity has to read register TxTBD0 register. Bits [15:12] indicate the status as indicated in slide 9 of [1]. The Management Entity does not need to hold any internal state (previous values or message toggle bit values). The status reporting does not have any self-clear bit, which eases debugging of Management Entity implementations.

Conclusions

- This presentation has compared the two proposals adopted in 1000BASE-RH and 1000BASE-T1 for implementing OAM.
- It has also analyzed the differences, and the reasons why the OAM proposal of 1000BASE-T1 was not well suited for 1000BASE-RH and why a more optimized implementation were chosen:
  - A direct reuse of the OAM proposal adopted in 1000BASE-T1 is not possible because of the different characteristics of the side-channel used to transport the OAM messages.
  - The side-channel used to transmit the OAM messages in 1000BASE-RH is very robust (much more robust than the channel used to transmit the Ethernet packet themselves), and also exists in situations where the partners do not reach a link_status = OK situation.
  - Because of this it is possible to simplify the message handling mechanism. At the same time the handshaking and remote status management has been modified to provide two benefits:
    - The status of the three possible outstanding messages is explicitly provided in the TxTBD0 register. The OAM proposal adopted in 1000BASE-T1 does not provide explicit indication of when a message is read by the Remote Entity.
    - OAM channel status (of all the possible 3 outstanding messages) is always available through a single read of the TxTBD0 register, the Management entity does not need to store any previous state information to decode the information.
  - This results in a simple implementation of the OAM Message mechanism as can be shown in the state machines needed to implement it as shown in [1].

References

- [1] *David Ortiz, et al., “OAM Proposal for GEPOF”, IEEE 802.3bv TF, Plenary Meeting, March 2015.*
- [2] *Rubén Pérez-Aranda, “Transmission scheme for GEPOF”, IEEE 802.3bv TF, Interim Meeting, January 2014*
- *[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,”PMA control state machines for GEPOF”, IEEE 802.3bv TF, Plenary Meeting March 2015*