PHYs, FEC modes and Cable assemblies: Can we do better?

Adee Ran

Intel Corp.

Abstract

- We seem to meet the CSD and the project's objectives with the current draft
- However, there are still claims of confusion and unnecessary burden on market that is not interested in low latency
- Consider possible improvements

What should we consider for success?

- IEEE 802.3 CSD
  - Managed Objects, Coexistence, Compatibility
  - Broad market potential
  - Economic feasibility
  - Distinct identity
- Project objectives, specifically
  - Define a single-lane 25 Gb/s PHY for operation over links consistent with copper twin axial cables, with lengths up to at least 3m
  - Define a single-lane 25 Gb/s PHY for operation over links consistent with copper twin axial cables, with lengths up to at least 5m

Do we fulfill the objectives?

- We defined the 25GBASE-CR PHY which can operate length up to 5 m with RS-FEC
  - 5 m meets both “at least 3 m” and “at least 5 m” – so we easily fulfill both objectives!
  - But we prefer to have better solutions, to address market needs.
- There is no objective to define a cable assembly…
  - Nevertheless, prior projects (802.3ba and 802.3bj) have defined a cable type as a complement for the PHY – so this seems to be in scope
  - This time we defined three modes of operation for the PHY, and defined a cable type for each mode
  - It seems to be in scope and doesn’t violate any objective.

Do we meet the CSD?

- Managed Objects, Coexistence, and Compatibility are trivial
- Technical feasibility seems to be in consensus
- We have addressed Broad Market Potential by including a low-latency mode that is important for cloud, HPC and other market segments
  - Despite not having an objective for that
  - Other market segments are also served
- Do two different PHYs per medium (e.g. CR and CR-S) create an issue with Distinct identity?
  - We can claim they don't, since each one addresses a separate objective
  - But what about KR and KR-S?
  - Wouldn't it be better to eliminate the question …?

Possible concerns

- "Technical/Economical" concern:
  - By requiring no-FEC performance, do we impose added cost on the whole market to address the needs of specific segments?
    - Some applications may never use no-FEC
    - As of D2.1, all PHYs must have no-FEC functionality and meet "enhanced" no-FEC performance even if they never operate in this mode!
- "Distinct Identity" concern:
  - Are the CA-N and CA-S different enough to justify two definitions?
- General argument:
  - 2 PHYs, 3 modes, 3 cable types – confusing!

So…

- No obvious problem in addressing the CSD and objectives
- But there is room for improvement…

FEC modes

- All PHYs must support BASE-R and no-FEC modes
- AN resolution rules make BASE-R the “default mode” for CR-S/KR-S
  - If one PHY requests it, the link will use BASE-R FEC
  - No way to “enforce” no-FEC
  - A deployed system may never use no-FEC!
- Likewise, BASE-R is the “default mode” for CR/KR assuming it interoperates with CR-S/KR-S (advertises both)
  - No way to “enforce” another mode
- This practically makes no-FEC and RS-FEC modes optional…
  - Not obvious, and confusing

Can we converge to one PHY?

- Having only one PHY with 3 modes can significantly reduce confusion
  - **BASE-R FEC mode** is in consensus for robustness and feasibility – and can work in many applications
  - **RS-FEC mode** addresses market segments that need higher reach and can tolerate the latency – broadens market potential
  - **No-FEC mode** addresses market segments that need low latency – broadens market potential
- One PHY with an optional RS-FEC mode can address both length objectives
- One PHY with an optional no-FEC mode eliminates imposing added cost on market segments that will never use no-FEC

RS-FEC and no-FEC optional?

- If both are optional modes, how does Auto-Negotiation resolve the link mode?
- A possible solution:
  - Use 2 bits to advertise support of the optional abilities (all 4 combinations are possible)
  - Use 2 other bits to request the mode from the LP (3 options)
  - Total of 4 bits – can be viewed as 2 bits (ability+request) per mode, as in 10G/lane PHYs (simplifies FEC resolution and reduces confusion)
  - A PHY shall request only a mode that it supports; request of both optional modes together is invalid
  - If the requested mode is not supported, the link defaults to BASE-R FEC mode – which is mandatory to support, and always “second best”
  - If both partners support both options and each one makes a different request, then priority is given to RS-FEC mode
- In engineered links, AN can be bypassed and then FEC “is controlled by implementation-dependent means”
- Detailed text in the backup section…

What about cable assemblies?

- The CA types have a practical meaning: they signify the “minimal” FEC mode that can be used safely with the cable assembly.
  - CA-N means “you can use this cable with no-FEC” (and also BASE-R and RS-FEC)
  - CA-S means “you can use this cable with BASE-R FEC” (and RS-FEC, but not no-FEC)
  - CA-L means “you can use this cable with RS-FEC” (but not no-FEC nor BASE-FEC)
- The length label assigned to each CA is only descriptive; cable compliance is specified in terms of COM parameters.
  - A 3-meter cable assembly may be wide-gauge and comply with CA-N, or narrow-gauge and comply only with CA-S (or even CA-L).
  - Other parameters, such as RL, may also affect COM and compliance with the different specifications.
- Based on several presentations, we assume users know what FEC mode they want to use
  - Having a CA specification to match each FEC mode would be useful
  - Merging CA-S with CA-N might create even more confusion!

Summary

- Although formally this project is OK, there is room for improvement.
- Consider the following proposal:
  - Remove the 25GBASE-CR-S and 25GBASE-KR-S PHY definitions – to have one PHY per medium.
    - Significant changes in clauses 110 and 111, but mostly deletions and simplifications.
  - Define RS-FEC mode and no-FEC mode as optional abilities.
    - Make PHY TX and RX functionality and receiver interference tolerance requirements based on supported abilities.
    - Use AN to advertise/resolve FEC mode.
- Based on level of consensus, to be either implemented in D2.2 or discussed further for possible implementation in D3.x.

BACKUP

Changes in clause 73

- In subclause 73.6.5 (FEC capability):
  - Change the list to define
    - F2 as 25G RS-FEC ability
    - F3 as 25G RS-FEC request
    - F4 as 25G no-FEC ability
    - F5 as 25G no-FEC request
  - And change descriptive text accordingly.
- In subclause 73.6.5.1 (FEC resolution for 25G PHYs):
  - Delete existing text, and insert instead:

    For 25G PHYs, two optional FEC modes can be negotiated: RS-FEC and no-FEC. Each mode has an ability bit and a request bit.
    - When the RS-FEC ability bit F2 is set to logical one, it indicates that the PHY has RS-FEC ability. When the RS-FEC requested F3 bit is set to logical one, it indicates a request to enable RS-FEC on the link.
    - When the no-FEC ability bit F4 is set to logical one, it indicates that the PHY has no-FEC ability. When the no-FEC requested F5 bit is set to logical one, it indicates a request to enable no-FEC on the link.

    Since the local device and the link partner may have set the 25G FEC capability bits differently, the priority resolution function shall be used to select FEC mode in both PHYs:
    - The RS-FEC mode is enabled if 25GBASE-KR or 25GBASE-CR is the HCD technology (see 73.7.6), both devices advertise RS-FEC ability on the F2 bit, and at least one device requests RS-FEC on the F3 bit.
    - Otherwise, the no-FEC mode is enabled if 25GBASE-KR or 25GBASE-CR is the HCD technology (see 73.7.6), both devices advertise no-FEC ability on the F4 bit, and at least one device requests no-FEC on the F5 bit.
    - Otherwise, the BASE-R FEC mode is enabled.
  - (Consider swapping 73.6.5.1 and 73.6.5.2 so that 25G comes after 10G)

Draft Proposal

- Define RS-FEC mode and no-FEC mode as optional abilities.
  - Add variables for “ability” and “enable” for each one, in clauses 110 and 111.
  - Define two AN bits (ability+request) for each one. Change advertisement and resolution rules in clause 73, as described in previous slide.
  - In Receiver interference tolerance test subclauses (110.8.4.2 and 111.8.3.1), make RS-FEC and no-FEC tests conditional with having the respective abilities.
- Remove the 25GBASE-CR-S and 25GBASE-KR-S PHY definitions in clauses 110 and 111, and all references to these PHYs in other clauses and annexes.
- Proposed detailed text to be provided (?)