Latency and FEC options for 25G Ethernet

Adee Ran
Intel Corp.
August 2014

Goals

- Explore FEC encoding/decoding options
- Discuss FEC gain/latency effect on the 5C
- Suggest objectives

Common ground

- We would like to have channel types and reach similar to 802.3bj
- We can tolerate some latency, but we're not indifferent to latency
- We target MTTFPA > 1e10 years

Clause 91 RS-FEC

- Introduced in 802.3bj, adopted for P802.3bm too
- Enables operation over challenging 802.3bj channels, but has a latency impact – codeword time is ~51 ns with 4-lane striping
  - Quadruples if used on one lane
  - Error correction takes additional time, implementation dependent – assume ~50 ns more
- Some users don't like latency…
  - We added "bypass correction" mode, mainly to enable reduction of latency to just 1 codeword time; assuming BER<1e-12 without RS-FEC protection
  - Bypassing error *checking* altogether would eliminate the codeword time latency, but lead to severe MTTFPA degradation*; we prohibited that
- If the minimum latency is considered too high, users might be tempted to bypass error checking, regardless of what the standard allows
  - "non-standard feature by popular demand" ➔ MTTFPA trap
  - This will silently reduce MTTFPA of the network

* See cideciyan_01_0512

Compare: Clause 74 FEC (aka Fire code)

- Introduced in 802.3ap, adopted for P802.3ba too.
- Provides less protection than RS-FEC:
  - Mainly aimed at bursts due to DFE error propagation. Random error correction capability isn't strong (~ 1 per codeword).
  - Estimate: input BER < 1e-8 required for equivalent FLR. This can be used in COM to evaluate channels.
- Codeword time is only ~82 ns for a single-lane 25G PHY
  - Was ~205 ns for 10G/40G.
  - Error correction without marking uncorrectable errors requires 1 codeword time. Checking for uncorrectable errors and marking them creates additional delay. Assume 1 additional codeword time.
- What about those latency-sensitive users?
  - Assuming TX uses Clause 74 encoding, RX can count errors in parallel without marking them, eliminating the codeword time latency. The remaining latency is up to one 66-bit block.
  - In this case, the "transcoding" does not significantly impact MTTFPA (no header compression, and no "single bit error causes frame merging" scenario).

Some numbers

Property	Units	Clause 91 RS-FEC 4 lanes (100 Gb/s)	Clause 91 RS-FEC 1 lane (25 Gb/s)	Clause 74 FEC 1 lane (25 Gb/s)
Block size	Bits	5280	5280	2112
Block time	ns	51	205	82
Latency for error correction (marking bypassed)	# Blocks	~2	~1.25	1
	ns	~100	~250	82
	Equivalent m of cable	20	50	16
Latency for only error marking (correction bypassed)	# Blocks	1	1	1
	ns	51	205	82
	Equivalent m of cable	10	40	16
Input BER for FLR≈6e-10		1e-5	1e-5	1e-8
Supported cable length (26 AWG)	m	5	5	3

Other PMDs we may consider

- Optics?
  - The coding gain provided by Clause 74 FEC is small – ~2.1 dB, compared to RS-FEC 5.14 dB (in optical power, ~1 dB vs. 2.57 dB)
  - Fiber links are typically longer and PHY latency is less significant.
- KP?
  - Even higher coding gain needed to compensate for more dense PAM-4 constellation.
- Clause 74 FEC seems unsuitable for both optical and PAM-4 PMDs.

FEC gain/latency effect on the 5C

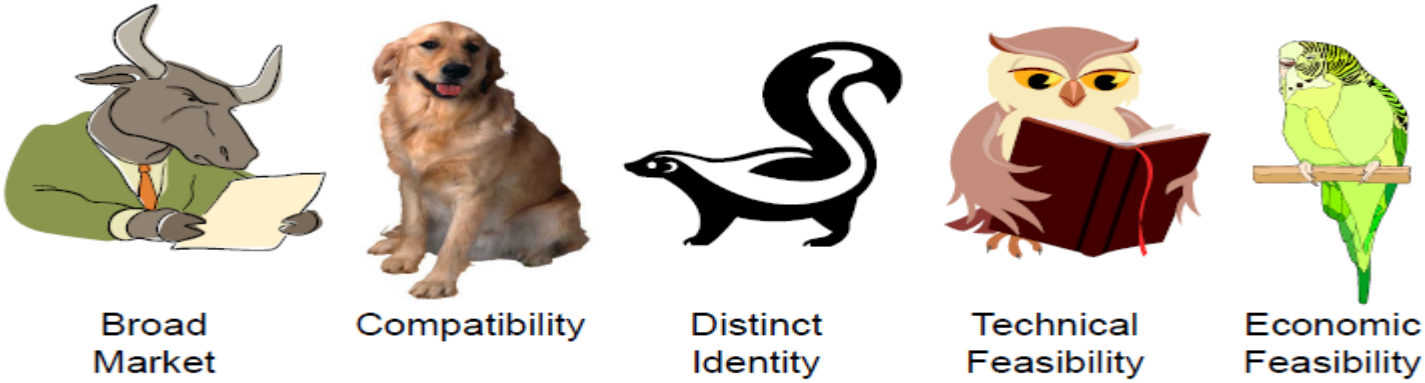

Source: frazier_01_0111

- Does a high-latency PHY have Broad Market Potential?
- Does the "MTTFPA trap" affect any of the critters?
- Can we have two FEC options and still meet Compatibility and Distinct Identity?

Possible way out

- RS-FEC encoding + decoding mandatory to implement
  - Guarantees interoperability
- Clause 74 as an option
  - Can be used for latency-sensitive applications with good signal integrity, replacing RS-FEC
- Negotiate using Clause 74 FEC via Clause 73 AN
  - Optical PMDs don't support AN, but Clause 74 isn't useful for this usage anyway
- Ultra-low latency can be achieved by transmitting with FEC encoding, and error checking/counting in the background, without marking errors
  - Unlike RS-FEC, this is MTTFPA safe. We don't have to prohibit, or even address, this type of usage.

Additional possible directions

- Stronger protection with even lower latency than Clause 74 FEC can be achieved using transcoding and a new, shorter RS-FEC code.
  - But this will require at least 1 codeword latency for error marking to be MTTFPA-safe.
- We can transmit 66b encoded data without any encoding.
  - This will simplify things for ultra-low latency applications
  - The downside is not having error counters – a useful diagnostic feature (unless we introduce AMs and BIP).

Suggested objectives

- All defined PHYs to meet:
  - **Frame loss ratio lower than X1**[1, 2]
  - **Mean Time To False Packet Acceptance higher than X2**[2]
  - Maximum delay lower than X3

Do we care?

- Define optional low-delay operation with maximum delay lower than X4 on links that permit it[2]

Do we care?

1. P802.3bs addresses FLR in the objectives (along with BER)
2. P802.3bn has an explicit FLR objective, an MTTFPA objective, and uses “channels that permit”

Thank you

- Discussion?