IEEE 802.3da – Carrier Sensing in Harsh Noise Environments

Piergiorgio Beruto

Outline

- In 802.3da we have had many discussions (some still ongoing) on the mixing-segment definition and how to achieve the target BER in various noise environments.
- In a half-duplex system, however, decoding bits from the line is not the only function that the PHY shall perform to guarantee the target BER and, in general, proper CSMA/CD operation.
- Carrier Sensing and Collision Detect are two essential *physical layer* functions that we must take care of.
- In 802.3cg this subject was discussed extensively, but the objectives of 802.3da create additional challenges (more reach, more nodes).
- This presentation explores the carrier sense function definition and requirements, proposing potential improvements for 10BASE-T1M and 10BASE-T1S.

Overview on Carrier Sense and Collision Detect

- In half-duplex CSMA/CD, a PHY shall provide an indication to the MAC layer of whether the line is "busy" (i.e., some node is transmitting)
- When the PHY reports a carrier, the MAC will defer any new transmission until the carrier event ends, and the inter-packet gap (>96 BT) elapses.
  - See the Deference Process in Clause 4.2.8 (normative) and 4.2.3.2.1 (informative)
  - This is a *best-effort* attempt at getting exclusive access to the medium
- Since carrier detection is not instantaneous (processing latency of the PHY, plus the line propagation delay), there is a time window, called the *collision window*, in which multiple nodes may attempt a transmission "simultaneously", resulting in a *collision*.
  - Collisions are not "errors", they are a *normal* part of the CSMA/CD mechanism.
- The PHY is required to detect when <u>its own transmission</u> results in a collision and report this information to the higher layers for the MAC to back-off and retry later.

Problem Statement

A special case: Receive Mode Collisions (RMC)

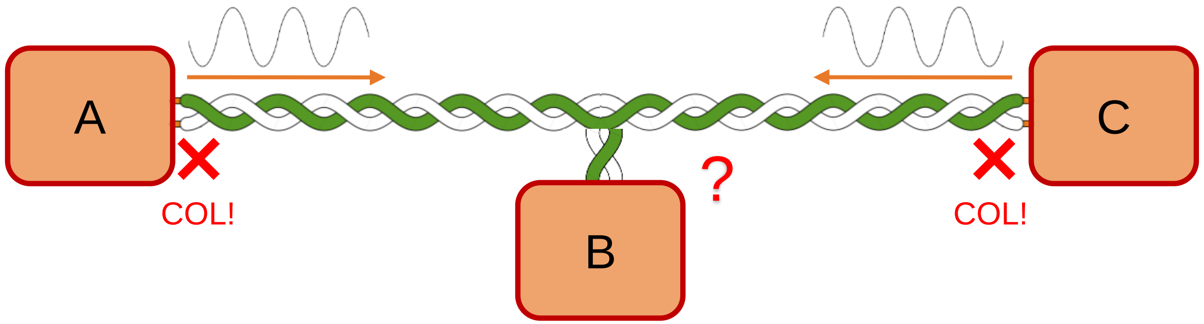

- Example: PHYs ‘A’ and ‘C’ start transmitting concurrently
  - both ‘A’ and ‘C’ detect a carrier and report a collision
- Node ‘B’ is just *listening* to ‘A’ and ‘B’: it is not transmitting
  - What should it report to the MAC?
- According to the normative ‘WatchForCollision’ and ‘TransmitLinkMgmt” procedures in Clause 4.2.8, the MAC ignores collisions when not transmitting (see also ‘collisionDetect’ in 4.2.3.2.4).
  - Therefore, B is –NOT– required to assert a collision
- What about Carrier Sense?

onsemi

A special case: Receive Mode Collisions (RMC) –contd–

- I could not find any requirement for CSMA/CD in Clause 4 regarding the generation of carrier sense during a collision, however…
- According to subclause 7.2.2.1.4 the PLS sublayer shall assert CARRIER_ON in the presence of the 'signal_quality_error (SQE)' message. This applies to MAUs supporting an AUI. Moreover, Tables 8-1 and 10-1 specify, for example, that (among others) 10BASE-2 and 10BASE-5 MAUs shall generate an SQE regardless of whether the MAU is transmitting.
- My understanding of the rationale: carrier sense *should* prevent listening nodes from “joining” the collision potentially too late within the collision window and ensure fragments stay below ‘slotTime’
  - Not joining the collision should yield somewhat better performance (fewer back-offs)
  - Additionally, it ensures that PHYs have enough “overlap” time to detect the collision
    - Anyway, this is not a strong guarantee and depends on the network propagation delays and PHY timings

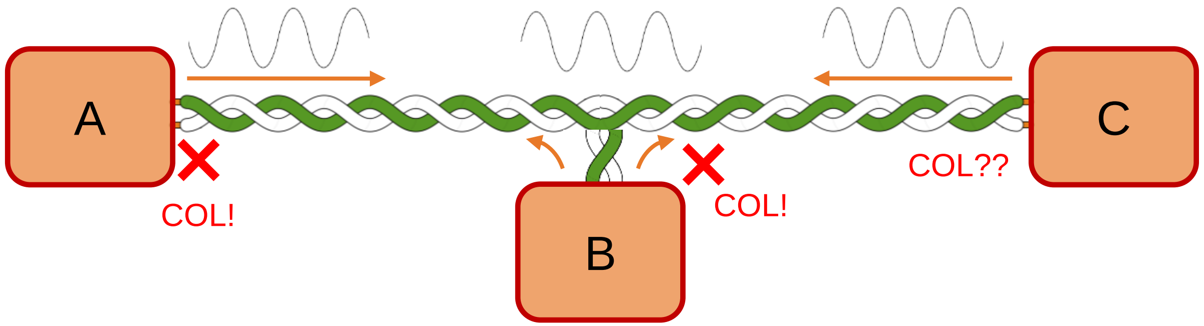

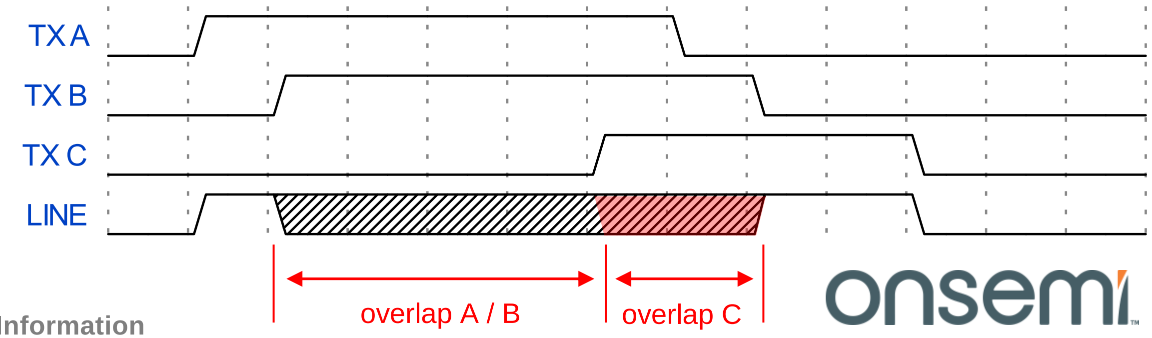

Why RMCs could be a problem in 802.3da?

- 10BASE-T1M is required to operate in harsh noise environments (industrial, transportation, building automation, etc.)
  - See https://www.ieee802.org/3/da/public/0722/beruto_3da_20220711_noise_env.pdf
- False carrier detections prevent the MACs from transmitting
  - The PHY could detect a carrier out of (high) differential noise
  - The measured BER would be 0… But just because the throughput is 0 (!!)
  - real-life experience with 802.3cg in automotive/industrial taught us that this is a real problem
    - which can be solved for 10BASE-T1S using PLCA along with additional implementation tricks, but it will not be enough for 802.3da and D-PLCA
- A well-known robust solution is to design a matched filter on DME and 4B/5B coding properties for generating the carrier sense indication
  - However, distinguishing “noise” from (multiple) “collisions” could be extremely challenging, if not impossible at all (e.g., in some cases the signal could cancel out completely)
  - RMC requirements effectively prevent PHYs from using more advanced DSP techniques for detecting carriers, forcing implementation to rely on energy detection only.
    - Matched filters could deliver around +10 dB better SNR compared to normal filtering (!)

Potential Solutions

Relax RMC requirements

- What is the real side-effect of not **requiring** a PHY to produce a carrier detection in the presence of a collision among two or more ***other*** stations?
  - Efficiency: we need to allow enough TX overlap time for collision detection
  - Robustness: we should ensure that a node transmitting at the end of the collision cannot generate (by accident) a ‘new’ valid frame as a continuation of the previous transmission.
    - Additionally, we must ensure that fragments resulting from collisions won’t exceed one ‘slotTime’ (512 BT). But this is true given the short propagation delays of these networks (< 6BT).
  - Performance: as explained earlier, more stations ***could*** join the ongoing collision generating additional back-offs
    - How bad is this really?
- NOTE: in 10BASE-T1S/M the line propagation delays are way shorter than in 10BASE-2/5 because we are not considering repeaters, and the target reach is also shorter

Efficiency: Time for detecting the collision

- If the PHY fails to detect a CRS during a collision among other nodes, and starts sending a new frame during the collision, two things could happen:
  1. The PHY detects the collision → no problem, normal back-off will follow
  2. The PHY does *not* detect the collision because the overlapping time is too short
     - We need to ensure that no frame is lost in this scenario

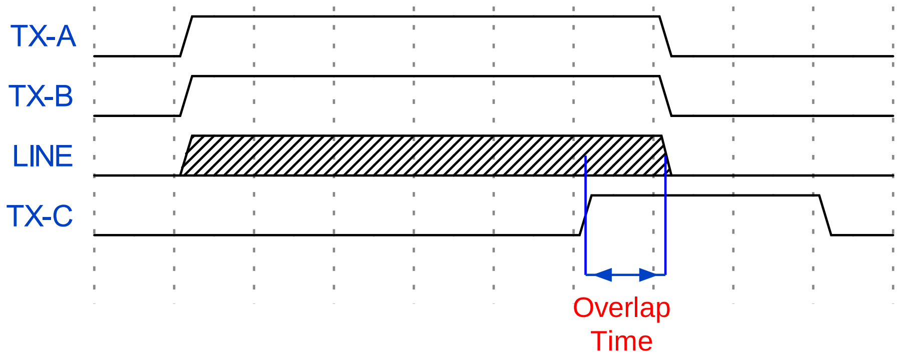

← collision between nodes A and B. C joins the collision at a “late” time because it does not see a CRS

- From 10BASE-T1S experience, a collision can be detected in roughly 2-3 BT
  – Let’s round this up to 5 to take some margin
- A frame preamble starts with the 5B symbol sequence JJHH (20 DME bits in total)

Efficiency: Time for detecting the collision -contd-

- The 10BASE-T1S PCS RX State Diagram (147.3.3.7) allows a frame reception to start if at least the ‘HH’ part of the preamble is received correctly. That is, the first two ‘J’ symbols could be “lost” in the PMA
  - Rationale: this tolerance was introduced to allow the PMA to lock the CDR on the preamble, and to allow any line transient to pass
- A collision detection time of 5 BT would “eat” exactly one ‘J’ symbol, leaving a ‘JHH’ to the PMA and PCS to initiate the frame reception
  - In other words, the collision detection time is shorter than the PMA/PCS margin for synchronizing on the incoming frame!

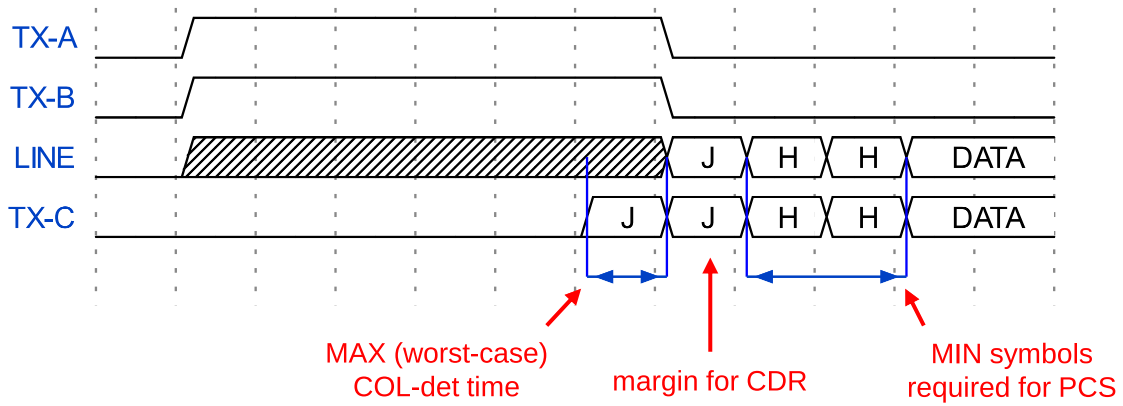

NOTE: in 10BASE-T1S the first ‘J’ is typically lost because of other reasons (line transients)

TAKE AWAY: if a node joins a collision at a late time, either it detects a collision in turn, or it sends out a decodable frame. In both cases, we’re good.

Robustness

- Consider the situation of two colliding nodes (A, B), a third one (C) joining the collision “late” (as in the previous example), and a number of listening nodes. We need to ensure that the concatenation of the collision from A/B with the transmission from ‘C’ does not accidentally form a “valid” frame.
- In other words, the listening nodes **shall not** accept an invalid frame (corrupted by the collision) as valid.
- We need to explore two potentially “bad” corner cases
  1. The listening nodes do not ‘see’ a JJHH preamble before the collision
  2. The listening nodes ‘see’ a JJHH preamble before the collision
- Case #1 is a simple one: if the listening nodes don’t see a JJHH at the beginning of the collision, the PCS will **not** accept the incoming frame. See backup slide #21 for details.
  - Basically, we’re back to the example in slide #11.

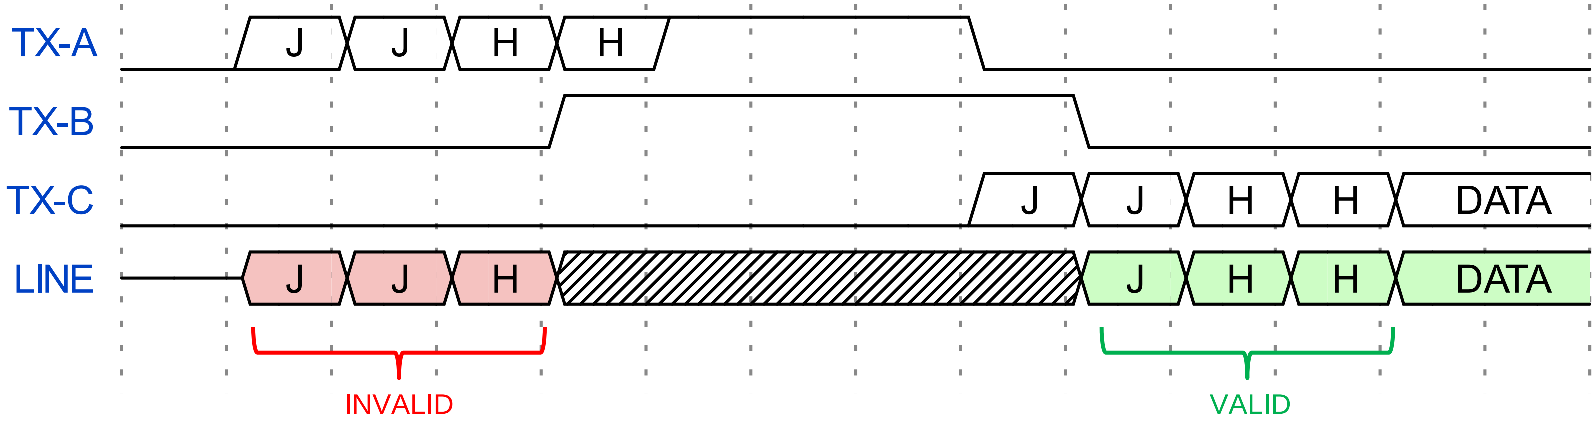

Robustness -contd-

- The below picture captures the second case (the nodes ‘see’ a JJHH)

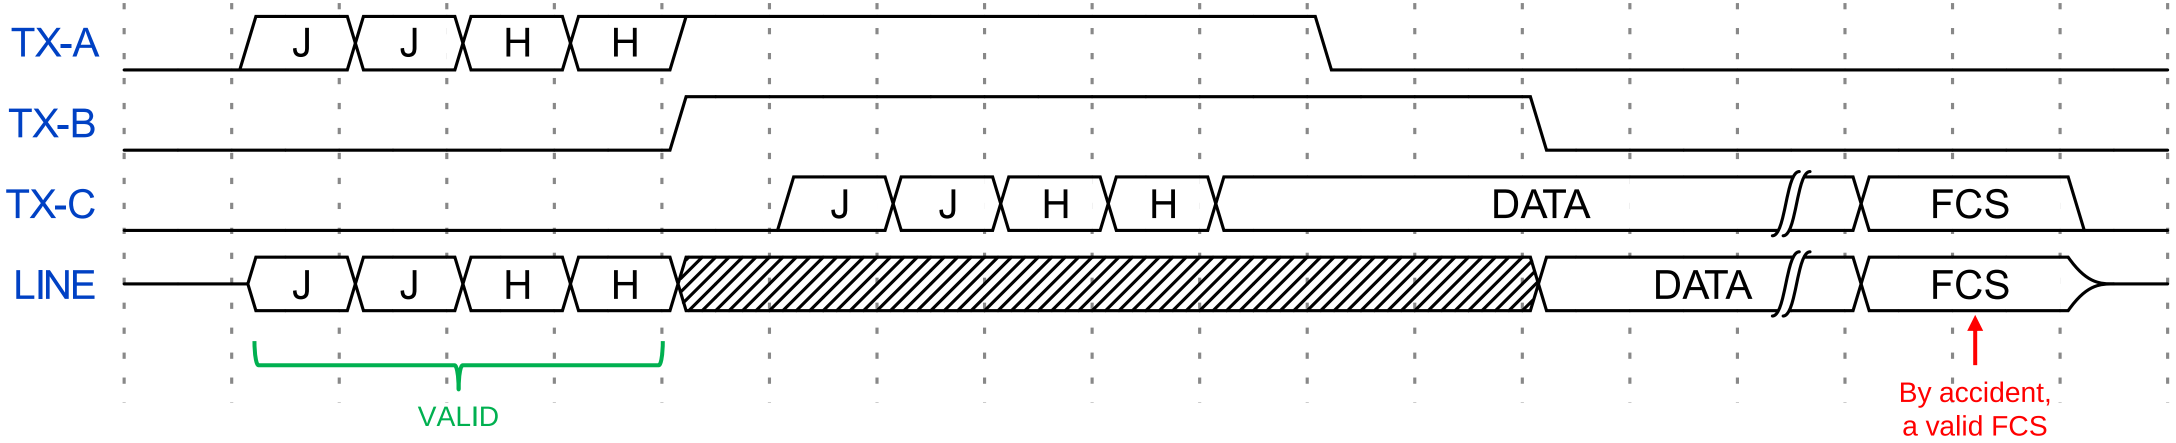

- In principle, a listening node could interpret the first JJHH as the beginning of a frame. The colliding data, concatenated with the transmission from C, could form a valid frame. It’s extremely unlikely, but it may not be less than one chance over the universe’s life, as the chance of acceptance of invalid frames requires.
- However, <u>this situation cannot really happen</u>! Here’s why…

Robustness -contd-

- If node “C” sees a **clean** JJHH, it will assert CRS after max 6 BT. Therefore, the MAC would not start a transmission after max 8 BT from the CRS assertion (see backup slide #20).
- On the other hand, nodes A and B would not stop transmitting until the 8-byte MAC preamble has been transmitted, plus the JAM time (in total, 96 BT).

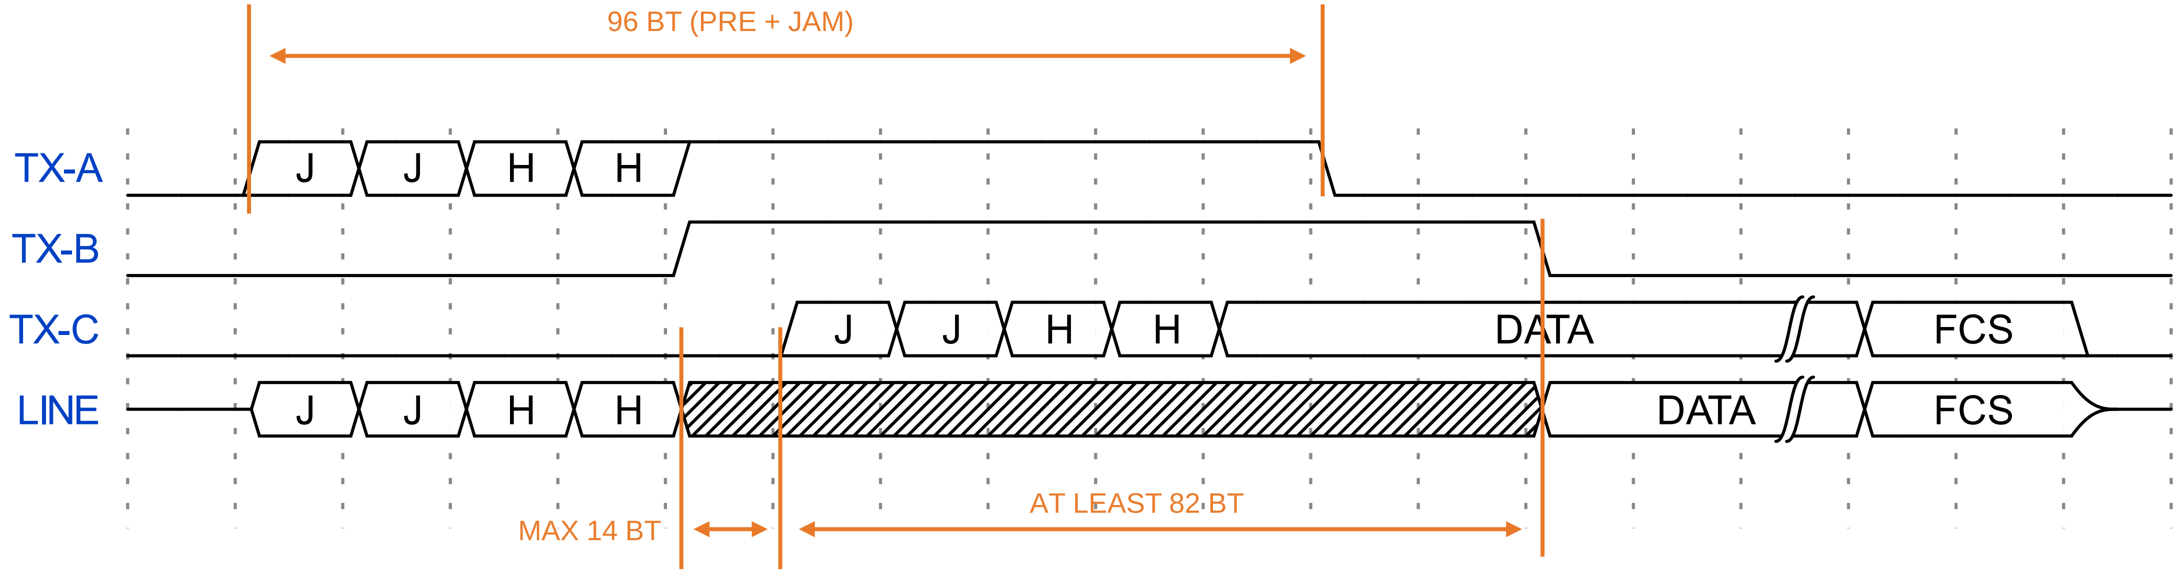

- However, since node “C” would join a collision lasting for at least 82 BT, it *will* report a collision to its MAC, resulting in a fragment. The listening nodes’ respective MACs will discard it.

Performance

- As explained, not asserting a carrier during a collision could lead, in principle, to more nodes joining the collision
- This effect has been assessed by means of simulations and measurements on real systems (source code of simulator available on request)
  - Tried with 16 and 32 nodes and average bus loads between 10% and 90%
  - The simulation explored all corner cases by randomizing the allowed CRS assertion and de-assertion times, as well as the line TPD.
- Despite what one could expect, the actual decrease in performance (increase of collisions) is barely measurable
  - Rationale: the most likely case that triggers a collision is when a PHY has just transmitted a frame, and all other PHYs ‘see’ the CRS fall “at the same time” (in a time windows determined by the difference between CRS assert/de-assert time and the line propagation delay).
  - In this case, the MACs will basically transmit concurrently anyway, generating a collision before the PHYs have a chance of detecting a carrier.
    - NOTE: this is exactly why CSMA/**CA** randomizes the transmit time after the IPG

Proposed text changes

- Relaxing the RMC requirements takes a single-word change…

147.3.5 Collision detection

When operating in half-duplex mode, the 10BASE-T1S PHY shall detect when a transmission initiated locally results in a corrupted signal at the MDI as a collision. When collisions are detected, the PHY shall assert the signal COL on the MII for the duration of the collision or until TX_EN signal is FALSE.

The method for detecting a collision is implementation dependent but the following requirements have to be fulfilled:

a) The PHY shall assert COL when it is transmitting, and one or more other stations are also transmitting at the same time.
b) The PHY ~~shall~~ assert CRS in the presence of a signal resulting from a collision between two or more other stations. should

- Sometimes a small text change means a whole world of things!
  - This is just the nature of the IEEE specifications language …

onsemi

Conclusions

Conclusions

- The 10BASE-T1M needs to operate in harsh noise environments
  - High differential noise could lead to false carrier detections, preventing nodes from transmitting for potentially very long times (seconds, or even more)
- Receive-Mode Collision requirements (RMC) prevent PHY implementations from using advanced DSP techniques for detecting real carriers
- Relaxing RMC requirements could be a solution
  - This presentation explored potential issues in terms of robustness and performance
  - No real draw-backs could be identified for 10BASE-T1M/S
- Mandating the use of (D)PLCA is another potential solution, which would also simplify the PHY design and reduce the complexity of D-PLCA.

onsemi™

Intelligent Technology. Better Future.

Follow Us @onsemi

www.onsemi.com

Collision window

- Normally, a collision can only happen at the beginning of a transmission, before the carrier information is conveyed and elaborated by the MAC layer

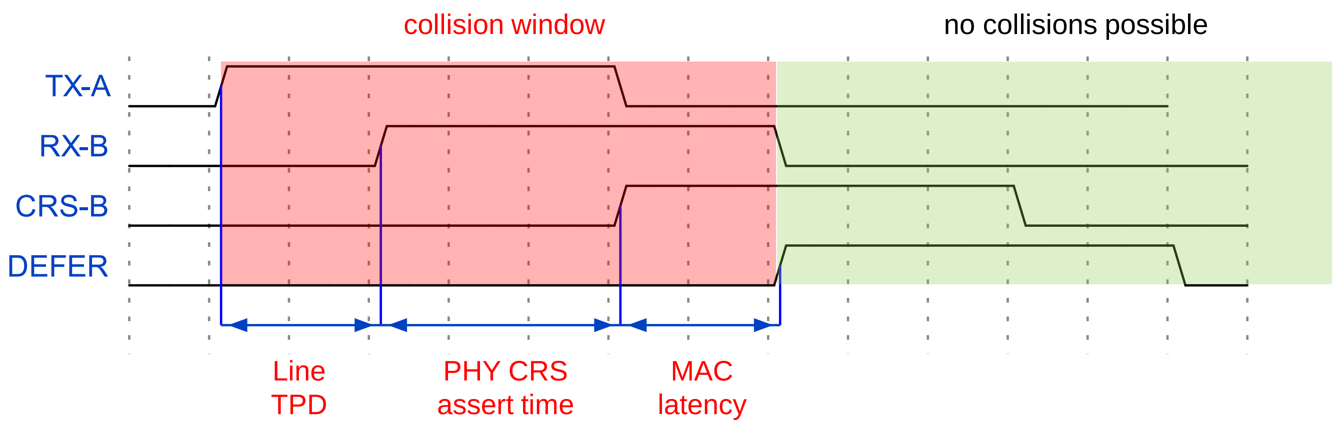

- A 100m link would yield 600 ns of propagation delay (considering a worst-case of 6 ns/m)
  - This is 6 BT at 10 Mb/s.
- The PHY CRS assertion time is specified as max 1040 ns in table 147-6 (T1S)
  - This is ~10 BT at 10 Mb/s
- The MAC latency is assumed to be max 8 BT in table 21-2
- This gives a collision window of max **24 BT**

PCS Receive State Diagram

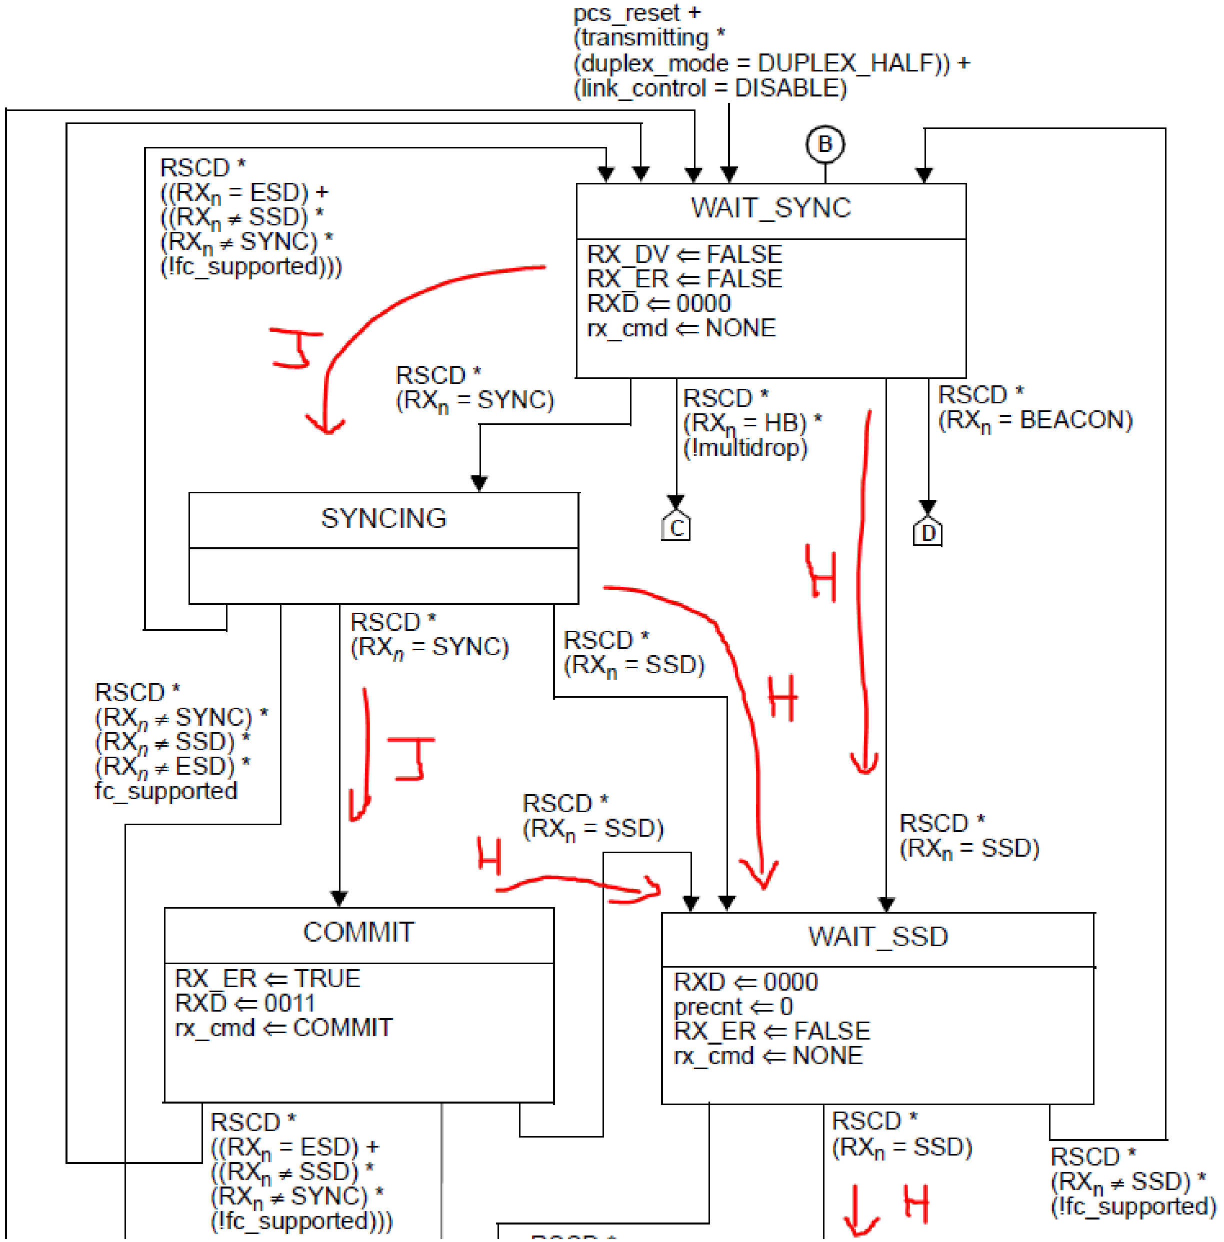

- The PCS accepts a frame when detecting
  - JJJ…HH
  - JJHH
  - JHH
  - HH
- any other combination brings the state diagram back to WAIT_SYNC