



PIERGIORGIO BERUTO
ANTONIO ORZELLI

IEEE802.3cg TF

PLCA overview

November 7th, 2017

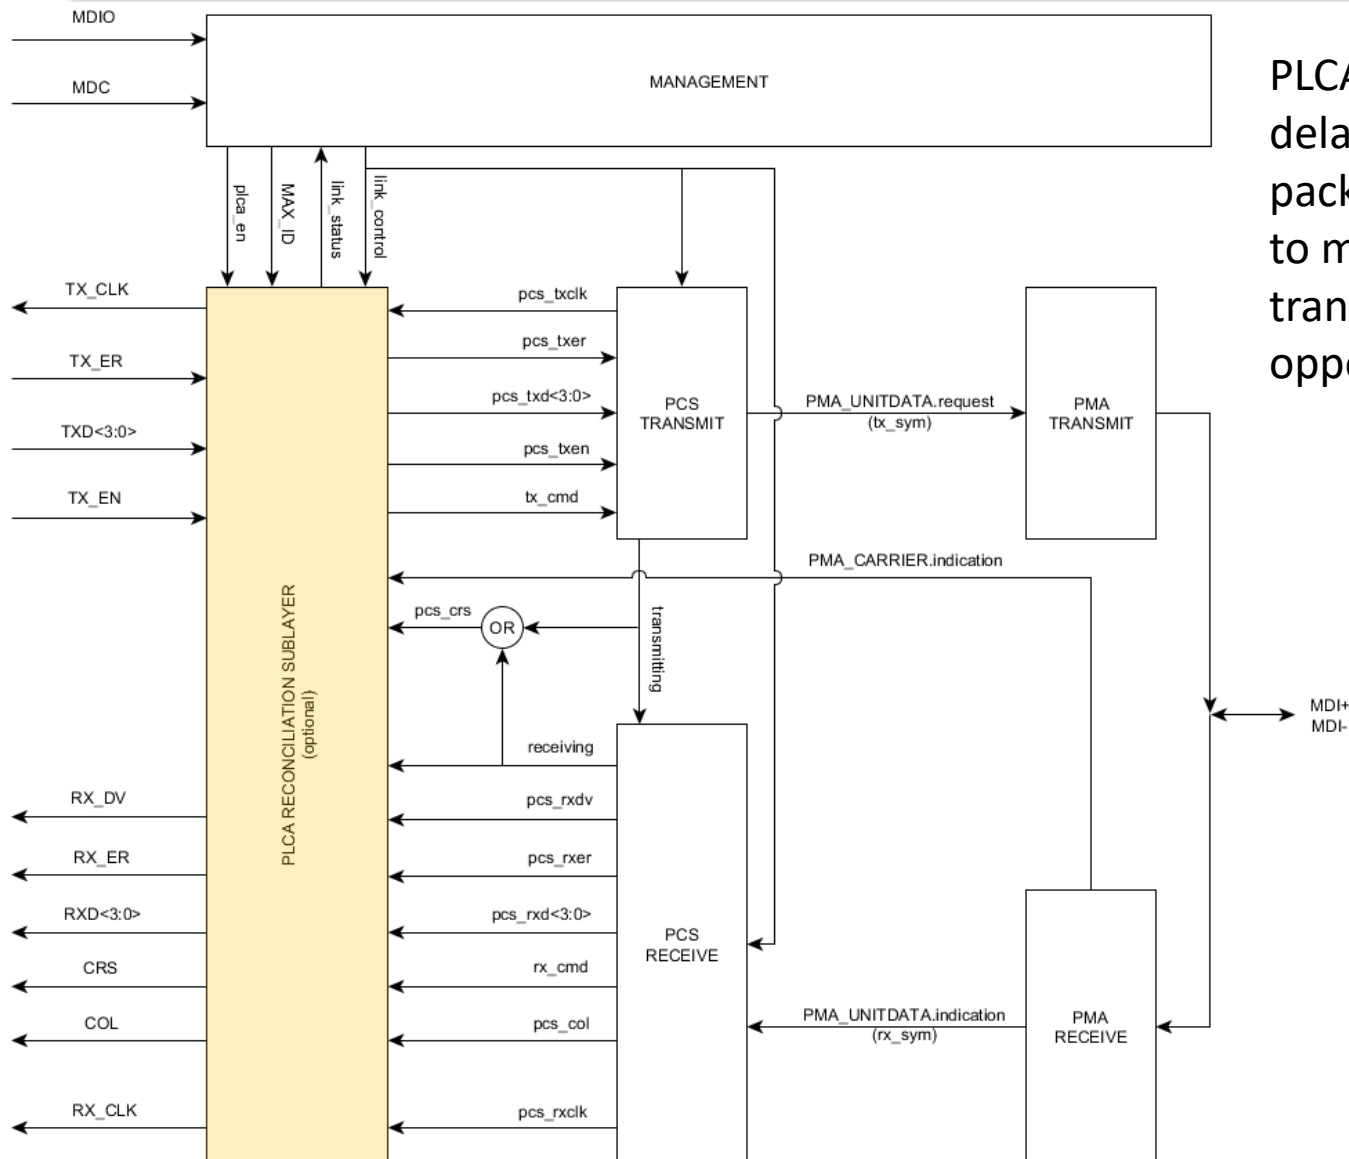


- What is PLCA?
 - PHY-Level Collision Avoidance is meant to provide improved performance (throughput, max latency, fairness) over standard CSMA/CD method for multidrop Ethernet networks featuring low number of nodes (< 16) and low propagation delays (short cables).
 - CSMA/CD functions are provided by the MAC
 - PLCA functions are provided by the PHY
 - Distributed protocol based on defining transmit opportunities to avoid physical collisions on the line.
 - Defined as an optional Reconciliation Sublayer
 - Proposed for adoption in 802.3cg group for the 10BASE-T1S PHY
 - Specified for 4B/5B coding but can be specified with different schemes as well
- What PLCA is not
 - Not a replacement of CSMA/CD → PLCA relies on it
 - Not a replacement of TSN → TSN expected to work on top of PLCA



Overview

MII

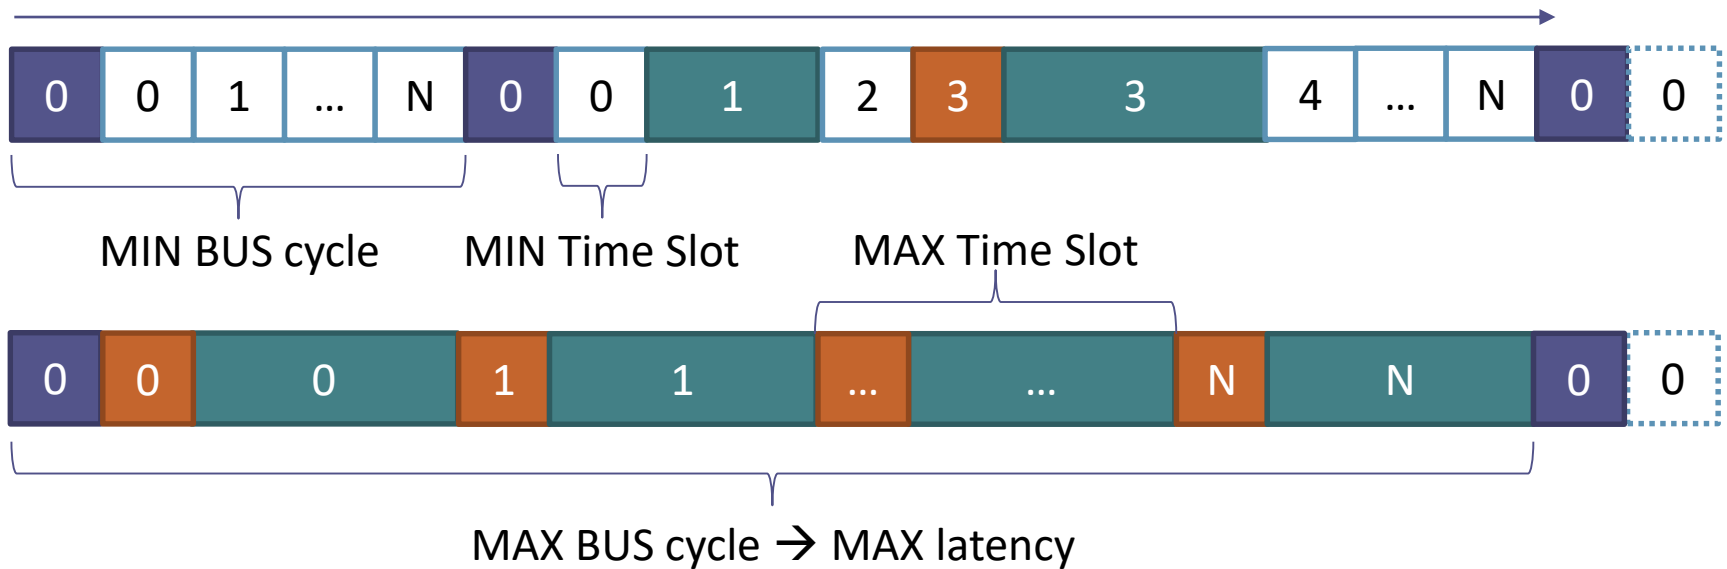


PLCA adds variable delays to incoming packets at the MII to meet specific transmit opportunities



In a nutshell

LINE



- PHYs are statically assigned unique IDs [0..N]
- PHY with ID = 0 is the master
 - Sends BEACON to signal the start of a BUS cycle and let slaves synchronize their timers
- A BUS cycle consists of N+1 **variable size** time slots plus the initial BEACON
 - PHYs are allowed to transmit only during the time slot which number matches their own ID
 - Time Slots end if nothing is transmitted within “MIN Time Slot” period or at the end of any transmission
 - PHYs are allowed to pad their own time slots with IDLE to compensate for any MAC latency (e.g. IFG)
- In numbers
 - BEACON time == MIN Time Slot == ~20 bT (bit-times) → ~1.6μs assuming 4b/5b + DME encoding
 - MAX latency == BEACON time + PHYs LATENCY + MAX BUS Cycle (all PHYs transmit one packet of MAX size, e.g. 1542 Bytes including IFG) → ~12.500 bT * (N+1) → ~1ms per PHY assuming 4b/5b + DME
- Round-robin scheduling guarantees fairness



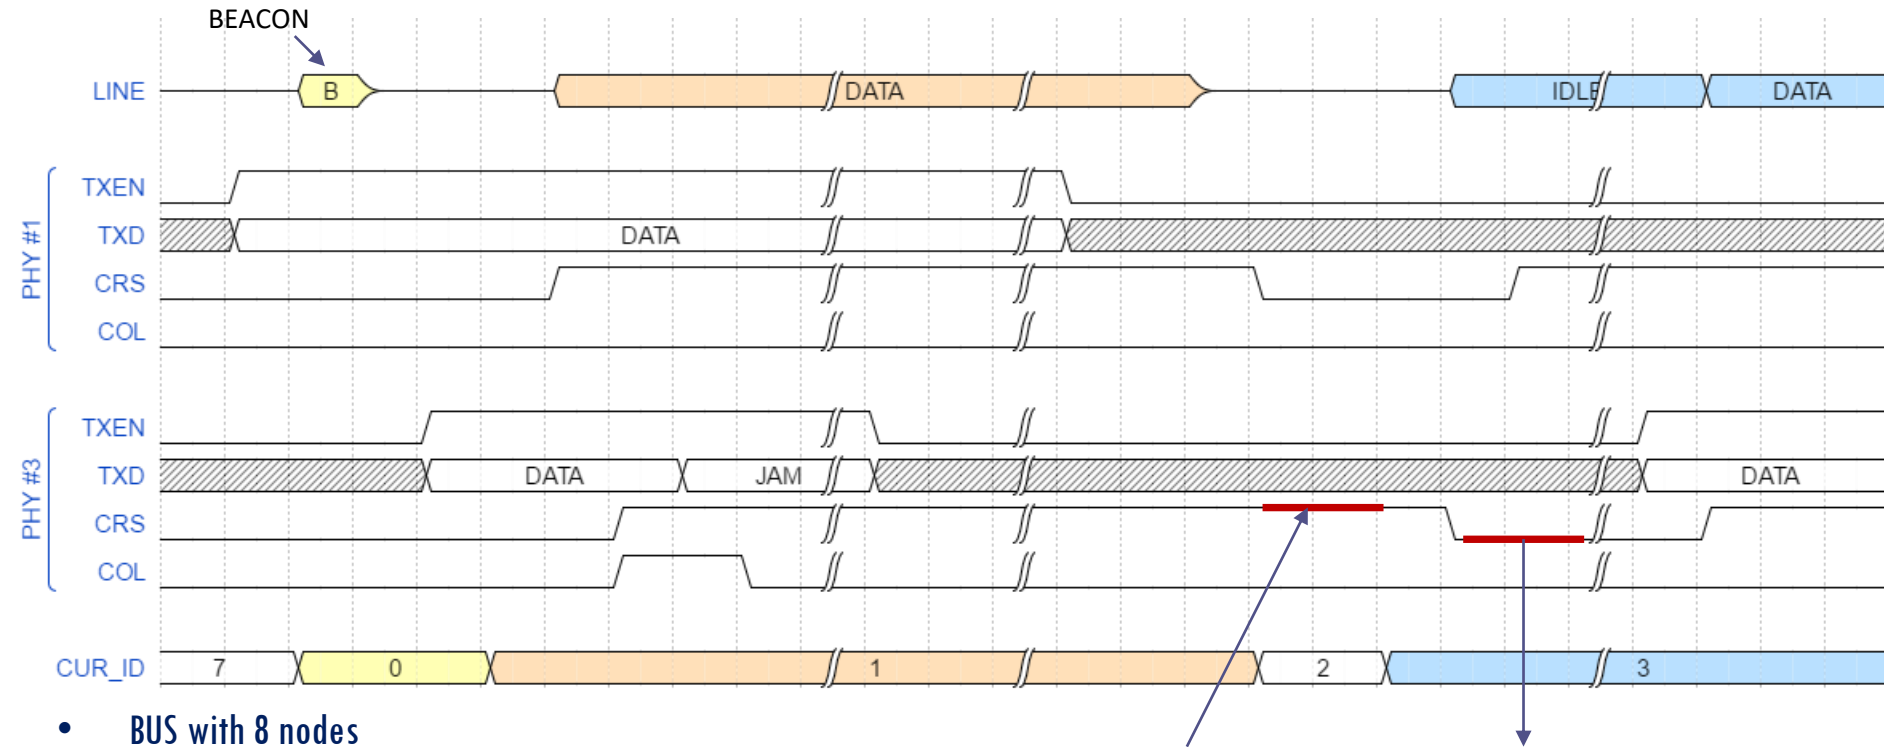


Relationship to CSMA/CD

- CSMA/CD MAC transmit process (from IEEE 802.3, clause 4)
 - If line is busy ($\text{CarrierSense} = 1$) \rightarrow wait (defer transmission)
 - Wait IPG (at least 96 bits)
 - Start transmitting, despite line becoming busy again
 - If a collision is detected ($\text{COL} = 1$) \rightarrow backoff:
 - Send jam for 32 bit times, stop transmission
 - Retry after $\text{random}(0, \text{ATTEMPTS}) * 512$ bit times
 - If $\text{ATTEMPTS} > \text{attemptLimit}$ \rightarrow give up (discard packet)
- CRS / COL can be used to have the MAC defer transmission until next handshaking
 - Use CRS to have the MAC defer transmission
 - Use COL at most once and only at beginning of a packet
 - MAC is ready to re-send in at most $32 + 512 = 544$ bit times
 - Less than minimum packet size (576 bits)



Example Waveforms




- BUS with 8 nodes
- Node #1 and #3 want to transmit data, others are silent
 - PHY #1 just defers TX until its own time slot is available
 - PHY #3 signals a collision because PHY #1 is transmitting, however:
 - No physical collisions on the line
 - Actual TX occurs immediately after PHY #1 transmission with no additional delay ($\text{MAX backoff} + \text{latency} < \text{MIN packet size}$)

CRS forced HIGH to prevent the MAC from transmitting until CUR_ID = 3

CRS forced LOW to have the MAC deliver the packet



Example Waveforms



MAC #1, 3 start transmitting.
PHY #1, 3 framePending <= TRUE

PHY #1 time slot begins, data is put
on the line since framePending = TRUE

PHY #3 signals a collision to its MAC
since PHY #1 carrier is sensed

MAC #3 initiates backoff and
sends JAM in response

MAC #3 backoff time always ends before
PHY #1 transmission is over (attempt = 1)
MAC #3 does not perform a new attempt
(yet) because CRS is asserted

Time slot #1 ends when PHY#3
falls silent again

PHY #3 keeps CRS asserted as
framePending == TRUE

Time slot #2 is yielded

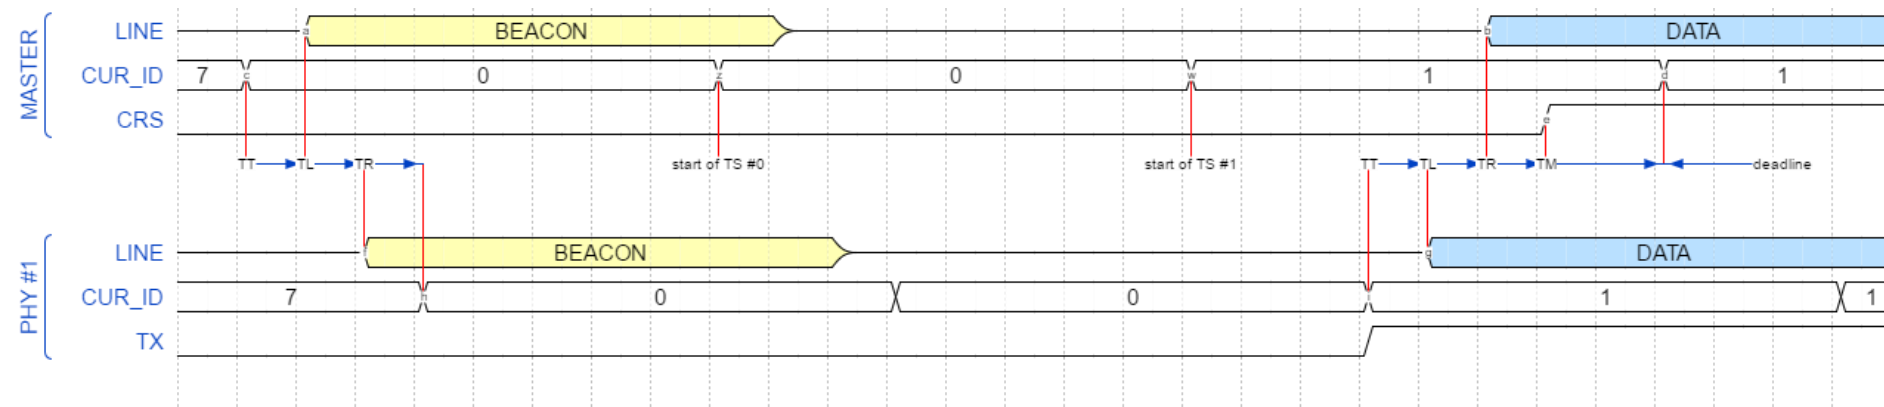
PHY #3 de-asserts CRS (allow MAC to
perform new attempt)

Phy #3 puts IDLE on the line to extend
the time slot until MAC is transmitting

MAC #3 waits for IFG then attempts
transmission again. DATA is
eventually put on the line



Clock Skew

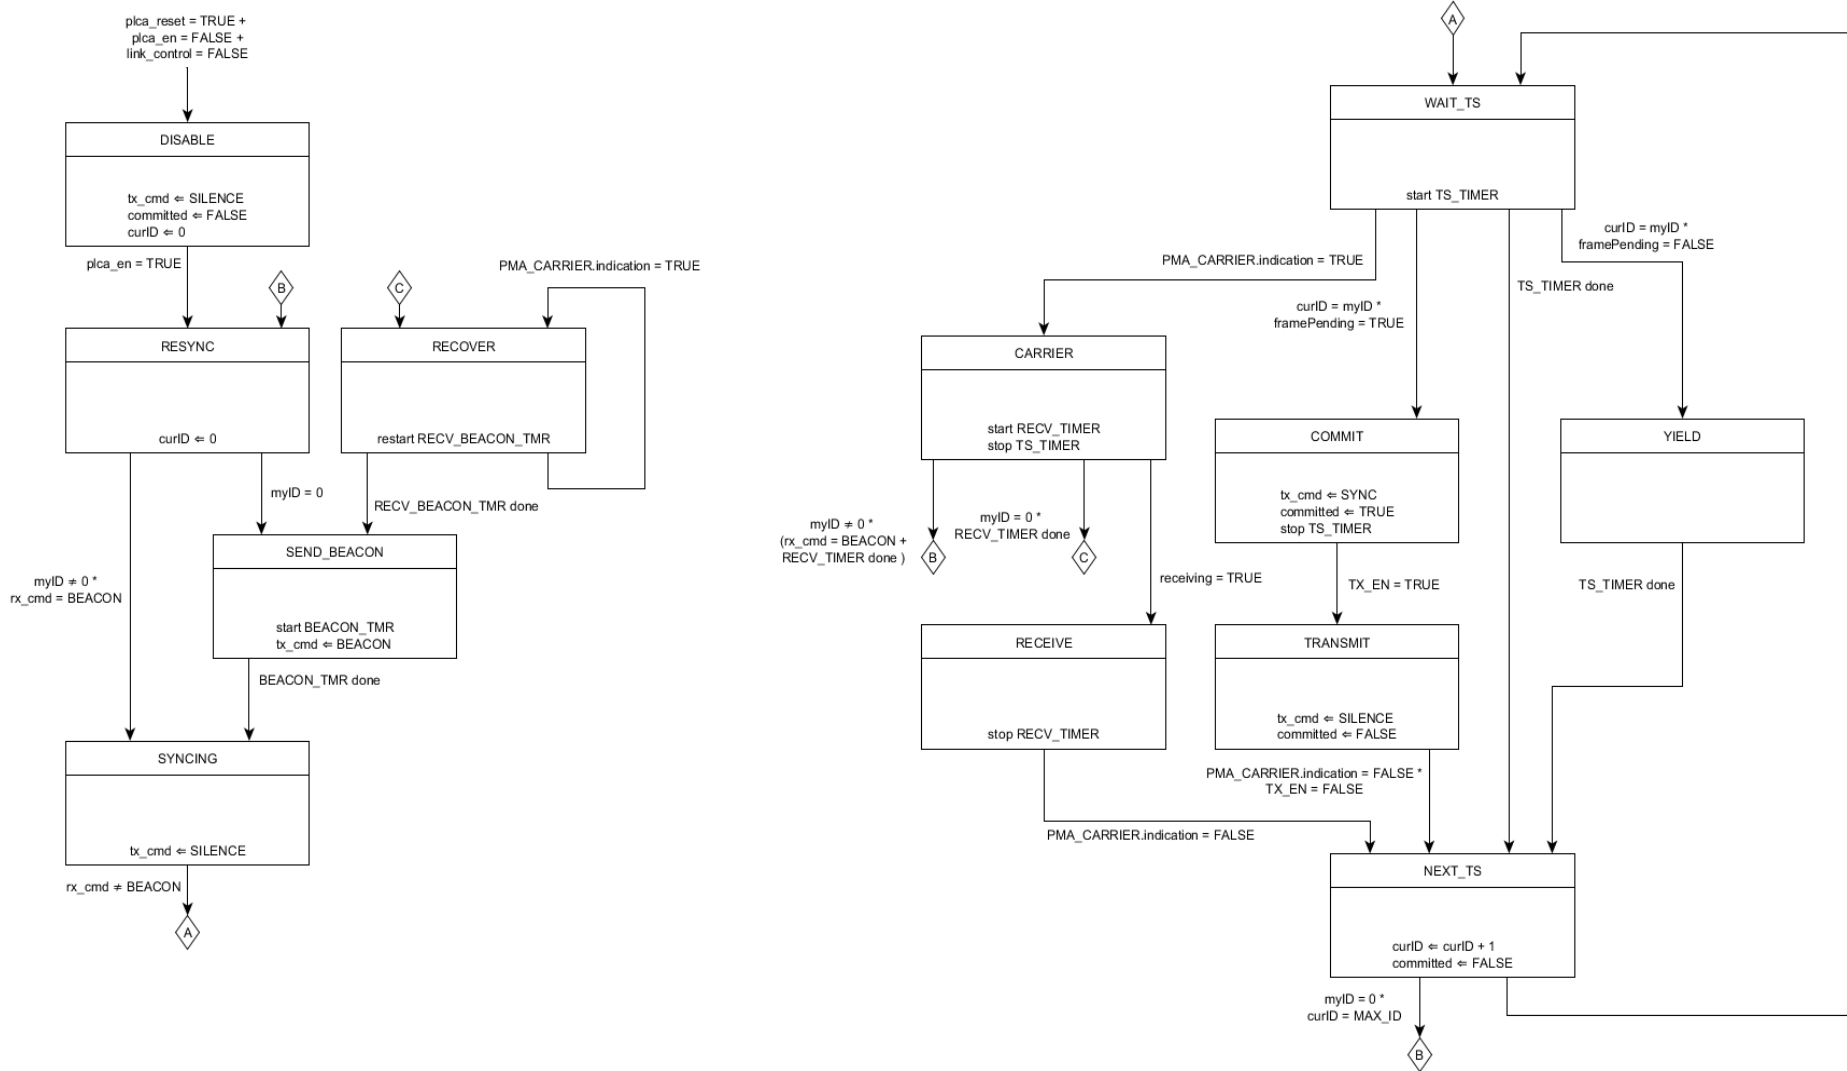


PARAM	DESCRIPTION
TT	PHY TX latency
TL	Line propagation delay
TR	PHY RX latency
TM	margin

- Total clock skew = $2 * (TT + TL + TR)$
- Margin > 0 for the system to work
- $TS_TIMER > TOTAL\ SKEW$

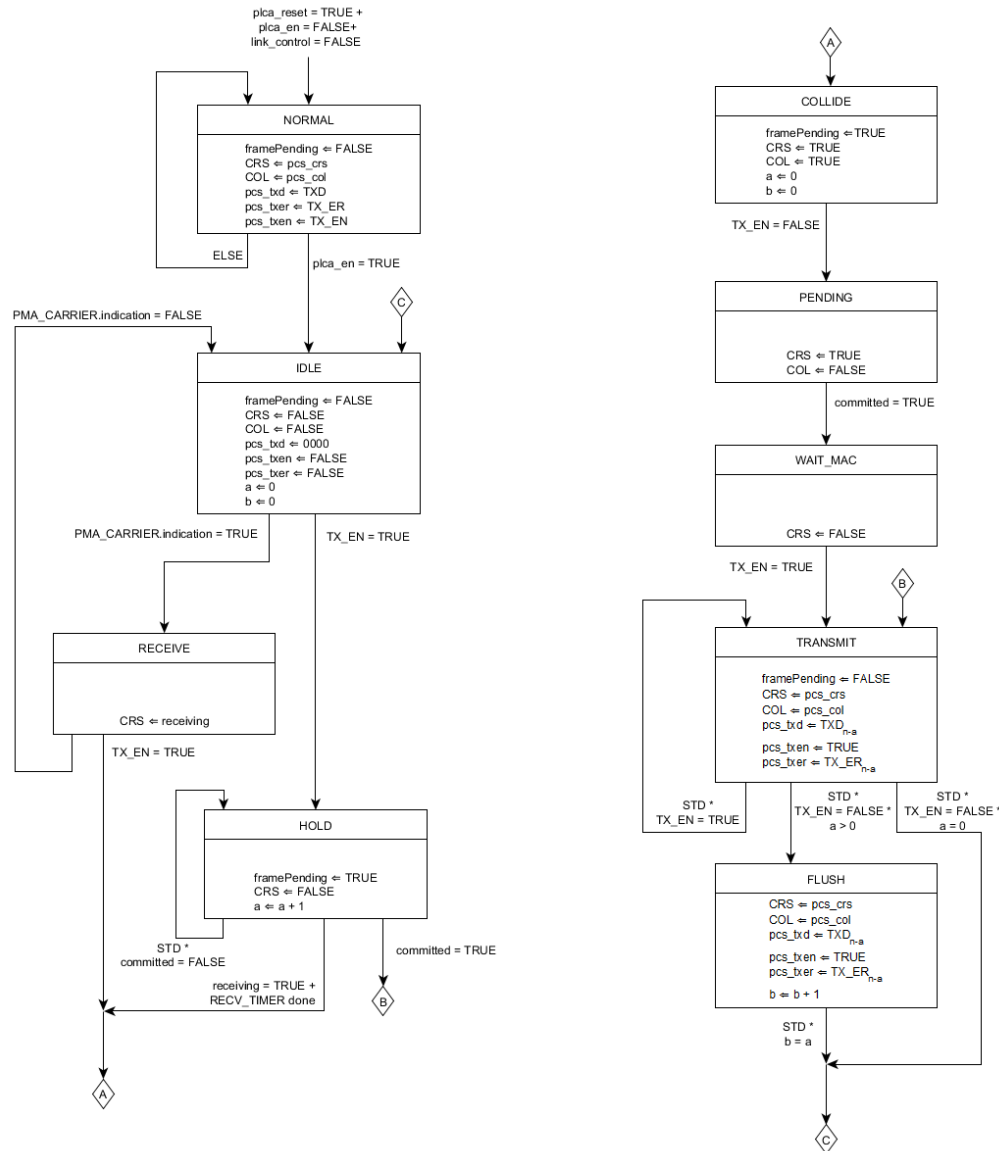


PLCA functions



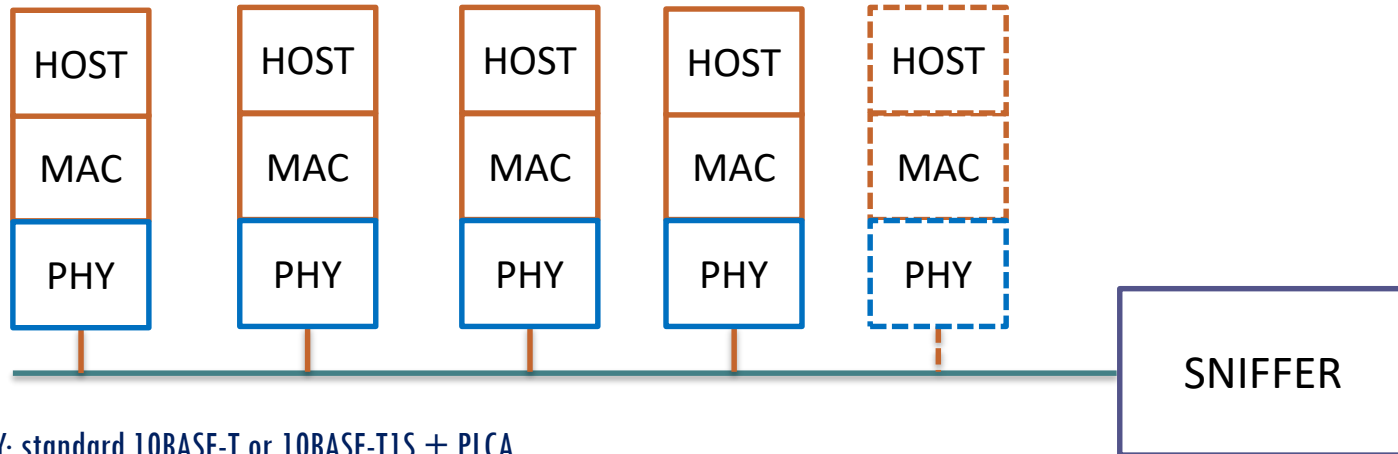


PLCA functions






- Full digital simulation (verilog)
 - 4b/5b encoding + DME



- PHY: standard 10BASE-T or 10BASE-T1S + PLCA
- MAC: standard CSMA/CD capable MAC (802.3 clause 4)
 - host interface: DPRAM (one frame) + busy indication + size + trigger
 - PHY interface: MII (txd, txclk, txen, txer, rxd, rxclk, rxdv, rxer, **col**, **crs**)
- HOST: simple transmitter
 - wait for MAC BUSY = 0
 - wait random time between 0 and MTP (sim parameter, 0 = MAX speed)
 - write random data in DPRAM of size PKTSZ (sim. parameter $60 < PKTSZ < 1500$) or random size
- SNIFFER: measuring throughput, latency
 - throughput: number of received bytes (excluding FCS, PREAMBLE) / total simulation time
 - latency: time between MAC BUSY = 1 and MAC BUSY = 0 for each node




Bitrate, N = 1






Bitrate, N = 2





Simulations


Bitrate, N = 3






Simulations


Bitrate, N = 4






Bitrate, N = 5









Bitrate, MTP = 0








Bitrate, MTP = 2000





Bitrate, MTP = 5000





- 500 pkts, size = 60B, variable MTP, 6 nodes. Latencies in μs .
- Comparison between simple half-duplex 10base-T and PLCA

MTP	MAX_LAT	AVG_LAT	STDEV
0	57595.6	553.3	4826.0
500	59692.8	1034.2	4637.4
2000	29387.5	618.9	2298.2
5000	19645.4	264.0	1035.7

10base-T

MTP	MAX_LAT	AVG_LAT	STDEV
0	443.4	441.1	26.2
500	546.4	186.4	90.7
2000	269.2	74.8	31.6
5000	223.7	64.0	17.8

PLCA

Thank You !