EEE and Multidrop

IEEE 802.3 Single Pair Multidrop Enhancements Study Group
George Zimmerman, CME Consulting, Inc
Piergiorgio Beruto, CanovaTech

Part 1: Energy Efficiency beyond the PHY

10BASE-T1S is inherently Energy Efficient

- What does this mean?
- The PHY transmits nothing when it has no packets to send
- The System at the far end knows nothing about when traffic may come at full rate
- PHY receivers are always active

IEEE Energy Efficient Ethernet - Overview

- Energy Efficient Ethernet (EEE) was introduced in BASE-T PHYs to help computers meet Energy Star power consumptions targets
  - Basic idea is to turn the link off as much as possible when there is no traffic, and turn it on as quickly as possible when there is traffic (50 us)
    - State is referred to as Low-power Idle (LPI)

EEE GOES BEYOND THE PHY!

A wider view

EEE study group has discussed saving power in the PHY

But whole system power measurements shown in CFI

- Power savings vs PHY speed > expected PHY power
- Even existing systems are saving more than PHY power

Examine current and potential system power savings

- "Reduction of power during low link-utilization"
- Where will this benefit from standards-based control?

July 16, 2007 20

IEEE 802.3 Tutorial, Energy Efficient Ethernet:
July 2007: H. Barrass, M. Bennett, W. Diab, D. Law, B. Nordman, G. Zimmerman

EEE Target was always beyond the PHY

- Estimates of > 50% power savings beyond the PHY, even with ‘fast-on’ EEE
- Multidrop, low-utilization sensors & actuators could be more
- ‘Sleep modes’ need this

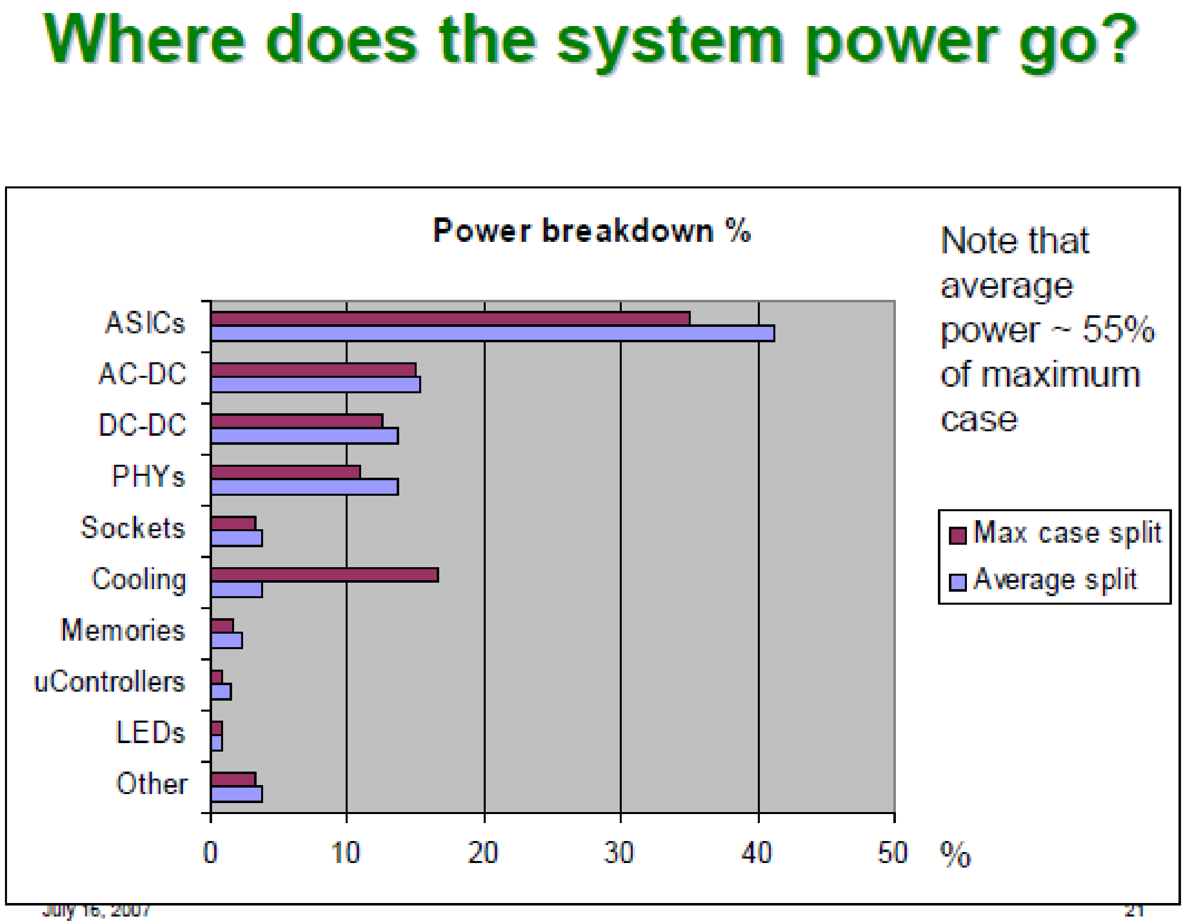

Source: “IEEE 802 Tutorial – Energy Efficient Ethernet”, Hugh Barrass, et al. (IEEE 802 tutorial, July 2007)

How EEE Works

- EEE is controlled by the Low Power Idle (LPI) client
  - Allows the MAC & higher layers to go to sleep
  - Physical layer only knows it is in LPI state and is quiet most of the time
- LPI client signals and controls the transition in/out of Low Power Idle state
- Signalling to LPI client via primitives is the KEY to system power savings

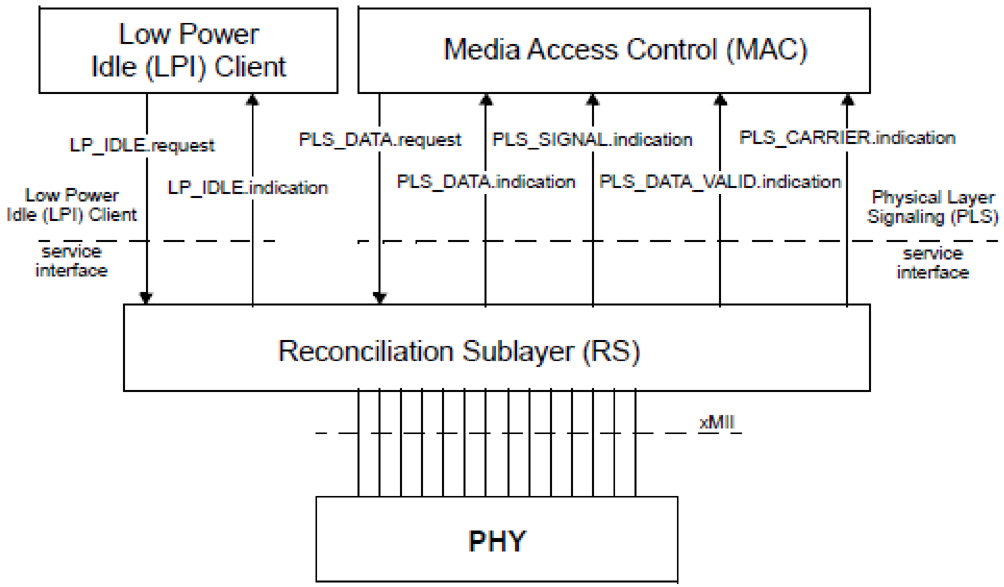

Figure 78–1—LPI Client and RS interlayer service interfaces

IEEE 802.3az Designed for bursty data – with rapid wake up

- Designed for bursty data and requiring minimal buffering
- Fast, application-transparent recovery
  - “The link status shall not change as a result of the transition”

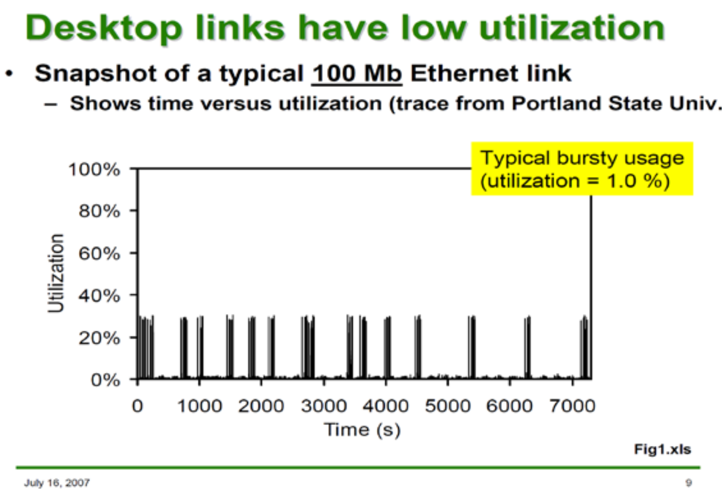

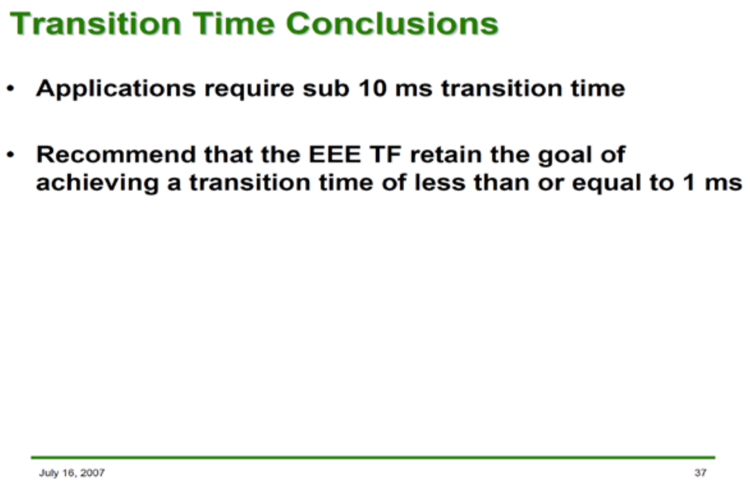

Source: “IEEE 802 Tutorial – Energy Efficient Ethernet”, Hugh Barrass, et al. (IEEE 802 tutorial, July 2007)

DO MULTIDROP APPLICATIONS REQUIRE THIS?

Multidrop and EEE are Synergistic

- Multidrop is the right solution when nodes do NOT regularly consume the bandwidth
- Managing energy at the system level is the right solution when activity is not always required
- *Both are consistent with long periods of inactivity at each node*
- Multidrop nodes can be ‘put to sleep’ or ‘woken up’
- Right now we don’t have the mechanism in Ethernet

Adjusting EEE parameters for low-duty-cycle or asymmetric traffic (802.3ch ‘slow wake’)

- PHY gets to specify system wake up timing parameters
- While Tw_phy is a minimum, Tw_sys_tx can be set longer, to accommodate longer PHY wake up.

Figure 78–5 illustrates the relationship between the LPI mode timing parameters and the minimum system wake time.

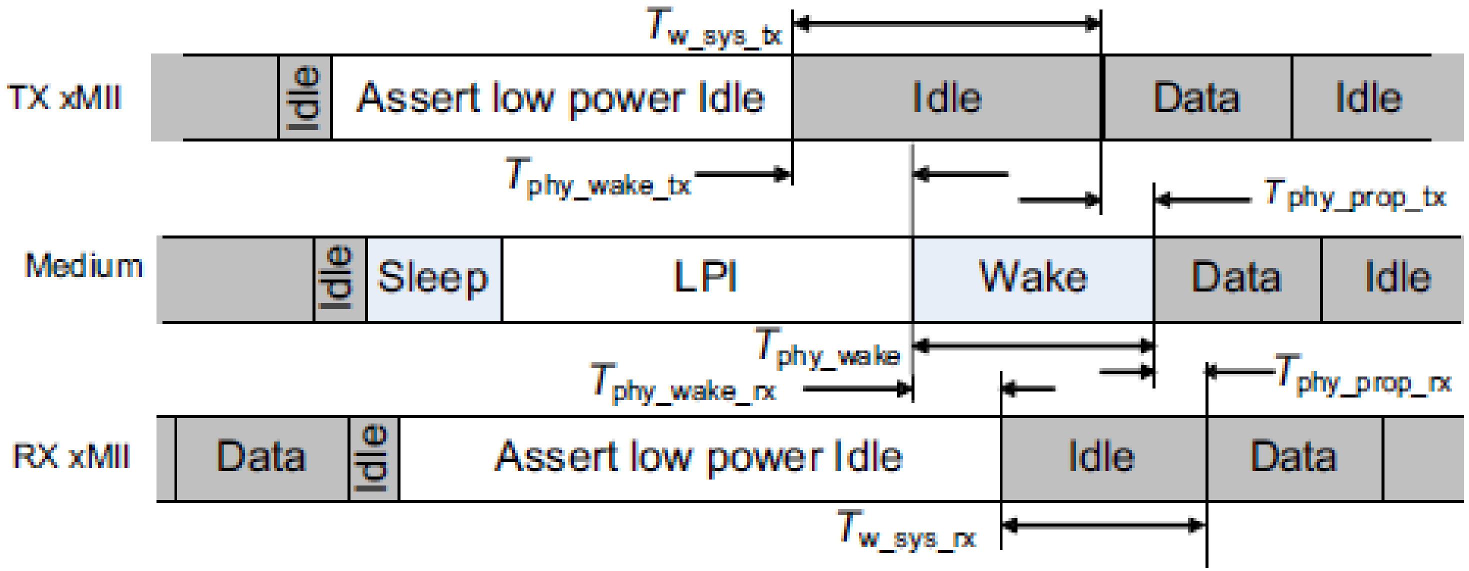

What could we add?

- Right now, 10BASE-T1S treats ‘Assert LPI’ and IDLE identically
- A remote PHY will never see a link partner indicating it is about to ‘Assert LPI’
- LPI client primitives cannot be exercised

Table 22–1 summarizes the permissible encodings of TXD<3:0>, TX_EN, and TX_ER.

Table 22–1—Permissible encodings of TXD<3:0>, TX_EN, and TX_ER

TX_EN	TX_ER	TXD<3:0>	Indication
0	0	0000 through 1111	Normal inter-frame
0	1	0000	Reserved
0	1	0001	Assert LPI
0	1	0010 through 1111	Reserved
1	0	0000 through 1111	Normal data transmission
1	1	0000 through 1111	Transmit error propagation

Table 22–2—Permissible encoding of RXD<3:0>, RX_ER, and RX_DV

RX_DV	RX_ER	RXD<3:0>	Indication
0	0	0000 through 1111	Normal inter-frame
0	1	0000	Normal inter-frame
0	1	0001	Assert LPI
0	1	0010 through 1101	Reserved
0	1	1110	False Carrier indication
0	1	1111	Reserved
1	0	0000 through 1111	Normal data reception
1	1	0000 through 1111	Data reception with errors

EEE

- EEE could be implemented to allow deep sleep mode of the PHY and system
- How?
  - Simple: LPI could be mapped to any unused 4B/5B code
  - Assert PI
- That's it!

Part 2: PLCA & EEE

PIERGIORGIO BERUTO

SPMD

PLCA & EEE

September 11th, 2019

Outline

- 10BASE-T1S is intrinsically power efficient
  - it's silent (no TX activity) when idle (MAC not sending data)
- PLCA sends a BEACON periodically
  - one each cycle of TOs, that is once each 20 * (N+1) bit times, with N = number of nodes
  - The BEACON is 20 bit times long
- Power consumption of a 10BASE-T1S implementation could be ~30 mA when transmitting and ~1 mA when receiving.
  - Avg power when idle (sending only the BEACON) for an 8 nodes network is ~4 mA
  - Avg power when idle (sending only the BEACON) for a 20 nodes network is ~2 mA
  - Avg power when silent is ~1 mA
- Stop sending the BEACON could save some more power and also reduce emissions
  - Is it worth it?

Idea

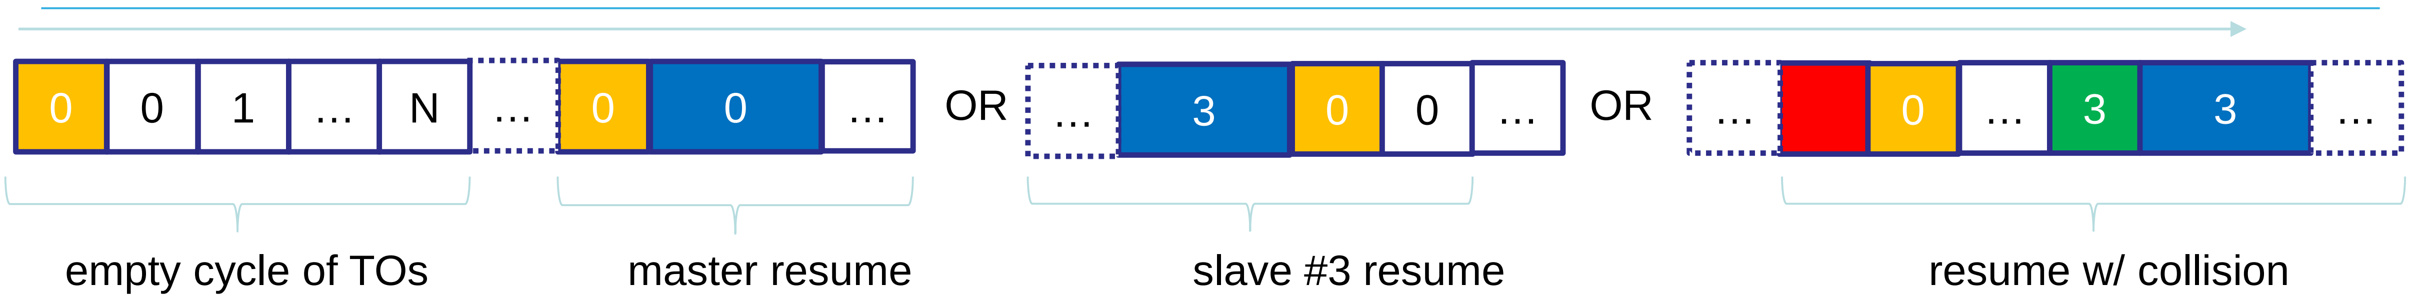

- The master node stops sending the BEACON after a certain amount of “empty” cycles of TOs
- Slave nodes revert to plain CSMA/CD mode at the same time
- If the master has a packet to send, it sends out a BEACON first.
- If a slave node has a packet to send, it does so and the master resumes sending the BEACON
  - This may cause an initial collision when multiple nodes have a packet to send at the very same time
  - After this potential collision, a BEACON is sent during IPG (already in 802.3cg) so no more collisions are possible

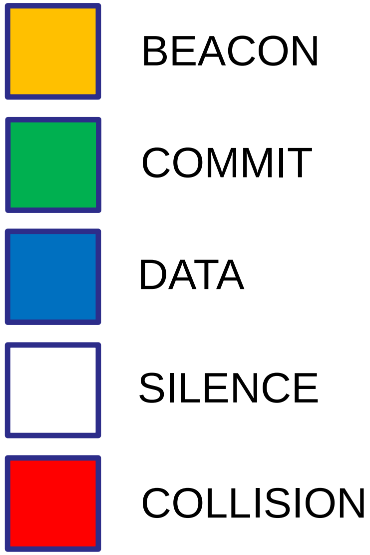

IEEE STANDARDS ASSOCIATION IEEE

THANK YOU!

Consensus
WE BUILD IT.

Connect with us on:

- **Facebook:** https://www.facebook.com/ieeesa
- **Twitter:** @ieeesa
- **LinkedIn:** http://www.linkedin.com/groups/IEEESA-Official-IEEE-Standards-Association-1791118
- **IEEE-SA Standards Insight blog:** http://standardsinsight.com
- **YouTube:** IEEE-SA Channel

IEEE
standards.ieee.org
Phone: +1 732 981 0060 Fax: +1 732 562 1571

Sources for EEE Overviews

- "An Overview of Energy Efficient Ethernet" Mike Bennett, youtube talk (TeamNANOG, June 2014)
  - Includes discussion of fast wake/deep sleep, optics, recent projects
- "IEEE 802.3az: The road to energy efficient Ethernet", Ken Christensen, et. Al, IEEE Communications Magazine, December 2010
- "Energy Efficiency and Regulation" Bruce Nordman, et al. (IEEE 802 tutorial, July 2008)
  - (see slides 27-50) EEE and WiFi energy management
- "Baseline Summary", IEEE 802.3az Task Force Home Page, March 2009
  - A nice summary of the 802.3az PHY modes
- "IEEE 802.3az Energy Efficient Ethernet Task Force Update", 802.1/802.3 joint meeting Mike Bennett, July 2008
  - Update on 802.3az approaches, scope related to systems issues
- "IEEE 802 Tutorial – Energy Efficient Ethernet", Hugh Barrass, et al. (IEEE 802 tutorial, July 2007)
  - Objectives, explanation of 802.3az motivations and approaches
- A Brief Tutorial on Power Management in Computer Systems, Dave Chalupsky & Emily Qi (IEEE 802.3 EEE Study Group, March 2007)
  - Covers both link and Systems level power management