



Advances in LAN transport over twisted pair media

Sanjay Kasturia & Jose Tellado

IEEE ComSoC SCV, September 8, 2010

Outline

- Twisted pair LAN evolution
- 10GBASE-T product evolution & roadmap
- Electromagnetic emissions & immunity
- Energy Efficient Ethernet

Twisted pair LAN evolution

- 10BASE-T
 - $\pm 2.5V$ PAM signaling
 - Manchester encoded
 - Category 3 and Category 5 cable
 - One pair used in each direction
- 100BASE-TX
 - 100BASE-T2 and 100BASE-T4 were commercial failures
 - +1V, 0V and -1V signaling
 - MLT3 encoding
 - Category 3 & Category 5 cable
 - One pair used in each direction
- Link margin went down
 - Dmin at transmitter reduced from 5V to 1V
 - Symbol rate increased from 10MHz to 125MHz
 - However it was still adequate for real world use
 - Widely used and robust, even after the reductions

Twisted pair LAN evolution –1000BASE-T

- 1000BASE-T
 - $\pm 1V$ PAM-5 signaling
 - Only Category 5e cable (Category 3 was dropped)
 - All four pairs used in each direction
 - Symbol rate remained at 125Msymb/sec but 4B5B style line code removed
 - 4x of the increase in rate comes from going from 1 lane to 4 lanes
- Margin reduction primarily due to higher modulation level
 - Dmin at transmitter (without FEC) reduced from 1V down to 0.5V
 - Reduction in Dmin is compensated by coding gain for Gaussian noise
 - 10V/m fields can create errors
 - Widely used & robust
- Was significantly more complex than 100BASE-TX
 - Echo/NEXT cancellation, FEC
 - Concerns about power, complexity, cost and performance were overcome by advances in technology

Twisted pair LAN evolution –10GBASE-T

- 10GBASE-T
 - $\pm 1V$ DSQ-128 (PAM ~12) signaling
 - Category 6, 6A and 7 cabling (Category 5 was dropped)
 - All four pairs used in each direction
 - Symbol rate increased from 125Msymb/s to 800Msymb/sec
 - Trellis code replaced by LDPC code
 - Equalization based on THP (gives gains of DFE while avoiding error prop)
- Margin reduction relative to 1000BASE-T
 - Dmin at transmitter (without FEC) reduced by ~7dB
 - Effective Insertion loss increased by ~15dB due to higher frequencies in TX spectrum
 - ~5dB more coding gain
 - Net is still ~17dB negative and has to be recovered by better equalization, better Echo, NEXT, FEXT cancellation, lower noise floor in receiver
- Is significantly more complex than 1000BASE-T
 - Concerns about power, complexity, cost and performance are being overcome by advances in technology

Teranetics 10GBASE-T Roadmap

GEN 1 PHY

130 nm



Single Port PHY
XAUI only

GEN 2 PHYs

65 nm



Single Port PHY
XAUI only



Dual Port PHY
XAUI and XFI



Single Port PHY
XAUI and XFI

GEN 3 PHYs

40 nm



Quad Port PHY
XAUI and SFI



Dual Port PHY
XAUI and SFI




Single Port PHY
XAUI and SFI

GEN 4 PHYs

28 nm



Quad Port PHY
XAUI and SFI
(Dual also available
in 27 mm)



Quad Port PHY
SFI only
(Dual also available
in 21 mm)



Single Port PHY
XAUI and SFI

10W/port @ 100m Reach
8W/port @ 30m Reach

5.7W/port (TYP) @ 100m
4W/port @ 30m Reach

3.7W/port (TYP) @ 100m
2.5W/port (MAX) @ 30m

Energy Efficient Ethernet
MACSEC
Enhanced Immunity

2.5W/port (TYP) @ 100m
1.5W/port (TYP) @ 30m

EEE
MACSEC

In Production

In Production

Sampling

Development



Electromagnetic Immunity (EMI)

IEEE ComSoC SCV, September 8, 2010

EM susceptibility

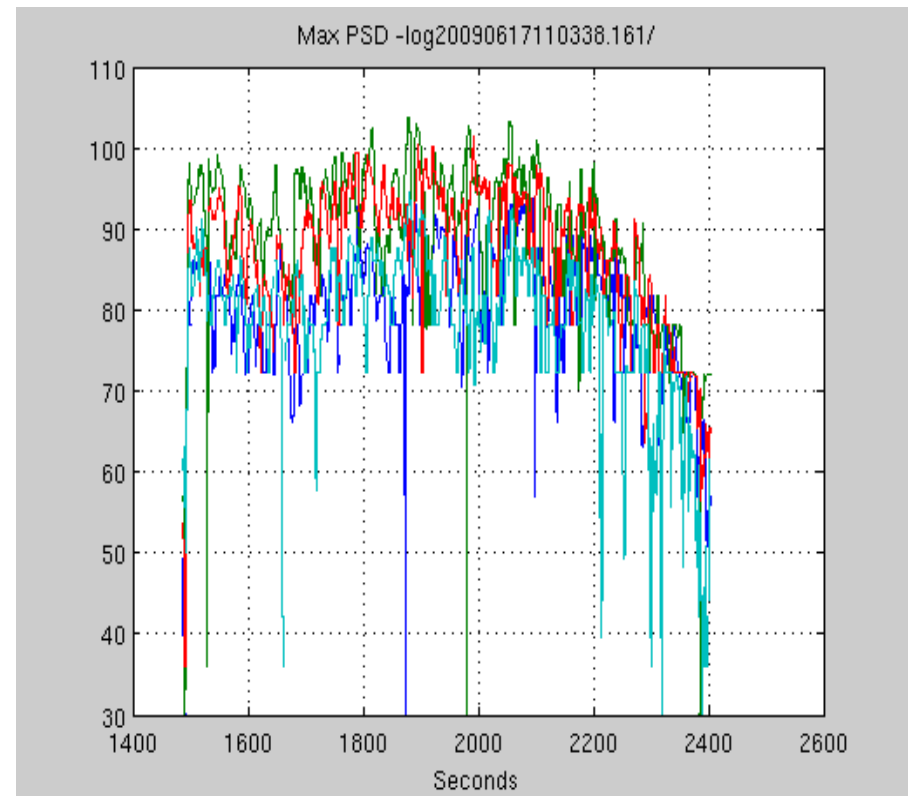
- External field couples to twisted pair
 - Coupling can be common and differential
 - Dominant coupling is to common mode
 - Differential coupling is more damaging but levels are lower
- Common mode coupling can cause errors
 - By overloading circuits so they are outside their normal mode of operation (easily handled)
 - By conversion to differential by imbalances in the channel or the receiver (this is the bigger challenge)
- Key protections against differential noise on the link are:
 - Common mode rejection from balance and/or shielding of the wire
 - Improve common mode rejection of the transformer
 - Additional attenuation of any residual common mode by common mode chokes in the magnetics
 - Improve common to differential performance of the connectors
- Industry specifications (covered in next slide) tend to be very conservative and passing those will remove risks

EMI susceptibility tests

- Conducted immunity
 - IEC 61000-4-6
 - 3V/m and 10V/m
 - Sweep from 150KHz to 80MHz in 1% increments
 - AM 80%, 1KHz
- Radiated immunity
 - CISPR24, Belcore GR1089, IEC 61000-4-3
 - 3V/m, 8.5V and 10V/m (aggregate)
 - Sweep from 80MHz-1000MHz in 1% increments (and discrete dwell freqs)
 - AM 80%, 1KHz
- EFT
 - IEC 61000-4-5
 - 500V, 1KV, 2KV
 - 5KHz and 100KHz repetition rate
- Teranetics' 2nd gen PHY (65nm) passed full EMC suite at all customers
 - Several different configurations and platforms and pass/fail criteria

Sample Differential level on 10GBase-T receiver

- Sample Differential level pickup
 - CISPR24 compliant EMC lab
 - 10V/m
 - Sweep 80-1000MHz, in 1% steps
 - cat6A
- 100dB ~ 10mVrms




10GBase-T EM susceptibility levels

- 10GBase-T transmit power decreases as cable loss improves (Power Backoff-PBO) to reduce alien xtalk
 - Transmit signal between 2000mVpkpk (PBO=0) to 400mVpkpk (PBO=14)
- Received signal for
 - short cable (~10m, PBO=10, meanIL~ 3dB) ~700mVpkpk (3s)
 - medium cable(~50m, PBO= 6, meanIL~ 9dB) ~ 900mVpkpk (5s)
 - long cable (~100m, PBO= 0, meanIL~15dB) ~1100mVpkpk(6s)
- Max received noise tolerance for
 - Long cable ~5mVpkpk (5s)
 - Medium cable ~14mVpkpk (5s)
 - Short cable ~ 35mVpkpk
- Worst case 10GBase-T noise sensitivity is at 100m (PBO=0)
 - Rx geometric signal is 30dB attenuated
 - Max noise tolerance is ~0.5mVrms (from IEEE spec of -141.9dBm/Hz)
 - Max EMI differential measurement at 10V/m is ~20mVrms
 - 30dB of EMI cancellation required

Handing of differential noise (out of band)

- Out of band can be filtered
 - 10GBaseT symbol rate is 800MHz
 - Majority of signal of interest is <450MHz
 - Significant filtering achievable >500MHz
 - Initial solutions used passive LC filters on PC board
 - Teranetics TN808X family integrates the filter
 - See Plot of VGA/Filter response



Handing of differential noise (in band)

- In band solutions (0-500MHz)
 - Detection
 - Could be performed on one or more of the 4 differential receivers
 - Pros: Known signals could be removed (Echo, NEXT, FEXT, ISI)
 - Cons: The DSQ128 constellation is very dense, $d_{min} \sim \text{Order}(mV_{rms})$
 - Or additional common mode sensing receivers
 - Pros: Undesired signals are typically much stronger than known signals (10-30dB higher)
 - Cons: Undesired signal coupling functions are very frequency selective (>40dB)
 - Cancellation
 - Equalizer Notching
 - Additional common mode sensing receivers

Equalizer Notching

- A notch filter can be put into the FFF (equalizer feed forward filter) response
 - The notch will generate residual ISI (intersymbol interference)
 - Proportional to bandwidth of the notch
 - ISI can be compensated by a FBF (equalizer feed back filter) or by changing THP coefficients
- 802.3az specified Fast Retrain allows rapid recovery from strong electromagnetic interference via a quick exchange of THP coefficients
 - Normal retrain time (start-up) is 2 sec
 - Fast retrain max recommended time is 30ms
 - Could be executed in a few millisecs
- A 100m link, with properly designed hardware and algorithms, will run error free with 10V/m fields

Additional common mode sensing receivers

- Typically undesired signals have much stronger common mode components
 - A circuit that converts/samples the common mode components on the channel will have better interference to signal ratio
 - Could give very accurate representation for the interference that could be subtracted from the 4 differential receivers
 - Cons
 - Has large amounts of broadband interference
 - C2C xtalk is much higher than D2D or C2D xtalk
 - Known signals such as common mode residual echo, NEXT, FEXT, etc
 - Requires some level of echo/NEXT D2C rejection
 - Common mode Alien Xtalk
 - Requires additional HW on the components (transformer) and/or the PHY (vga, adc)
 - Because of frequency selectivity, Common mode components could be multipath cancelled at some frequencies



Energy Efficient Ethernet

IEEE ComSoC SCV, September 8, 2010

Energy Efficient Ethernet

- Energy Efficient options have been defined for all PHY rates and several different media (twisted pair, backplane and by extension, optical)
 - We only discuss twisted pair solutions
- After much discussion and debate, IEEE 802.3az task force settled on an approach based on the use of a Low Power Idle (LPI) mode
 - Exception being 10BASE-Te, which only reduces the voltage swing
 - Impact on link margin ameliorated by limiting media to Cat 5e
- The LPI approach requires
 - New signaling (Alert over the line, and to and from station management)
 - Fast transitions from LPI to active & vice versa

Link operation with LPI

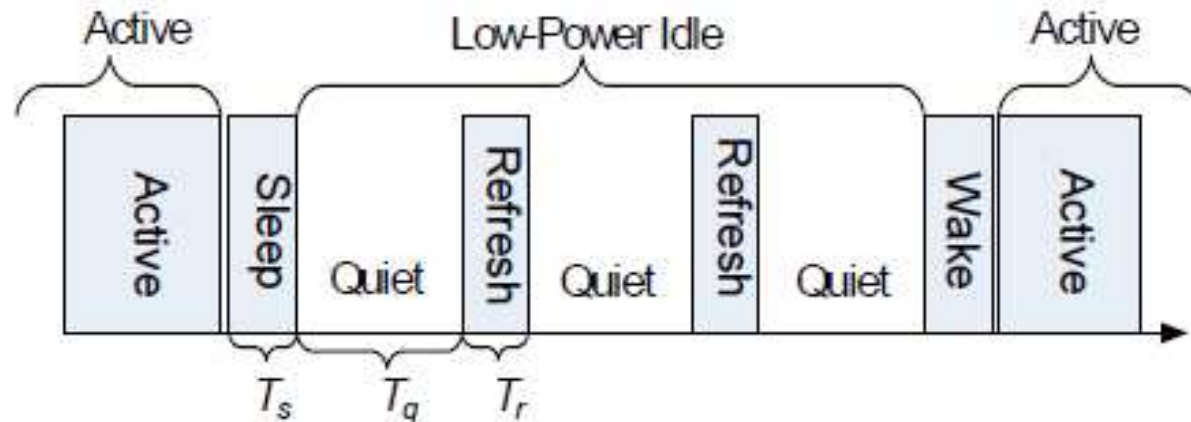


Figure 78-3—Overview of EEE LPI operation

- The LPI approach can scale easily from ~0% utilization to 100%
 - This was not possible with the sub-rate PHY approach
 - The draw back of the LPI approach, relative to the sub-rate PHY approach, is latency
- The Refresh is required to keep receiver parameters current (timing, equalizer coefficients, canceller coefficients etc)
 - This is critical to enable fast transitions from LPI to active
 - There is a tradeoff between idle power and latency

Key EEE time intervals

- T_s is the minimum time sleep must be signaled
- T_q is the quiet time
- T_r is the refresh time

Protocol	$T_s \mu\text{s}$		$T_q \mu\text{s}$		$T_r \mu\text{s}$	
	min	max	min	max	min	max
100BASE-TX	200	220	20,000	22,000	200	220
1000BASE-T	182	202	20,000	24,000	198	218.2
10GBASE-T	2.88	3.2	39.68	39.68	1.28	1.28

EEE operation with LPI

- Requires a new service interface to station management
 - Two new primitives : LP_IDLE.request and LP_IDLE.indication
 - These are carried through the RS to the line side of the link
- When these are used is not within the scope of the standard
 - Application specific

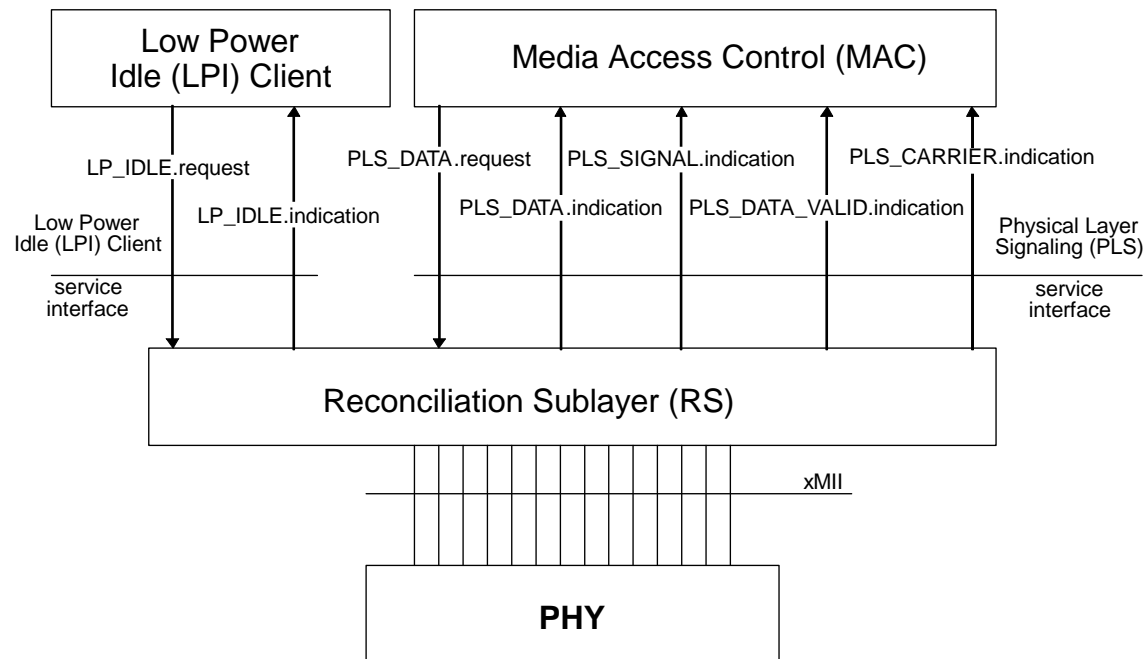


Figure 78-1—LPI Client and RS interlayer service interfaces

Wake times

- Shorter PHY wake times are better, but could limit the depth of the power reduction in LPI mode
- System wake times could vary widely
 - Clause 78 of 802.3AZ provides some default values
 - Clause 78 & 79 specify how the Link Layer Discovery Protocol (LLDP) can be used to negotiate these

PHY Type	Case	$T_{w_sys_tx}$ (min), in μs	T_{w_phy} (min), in μs	$T_{phy_shrink_tx}$ (max), in μs	$T_{phy_shrink_rx}$ (max), in μs	$T_{w_sys_rx}$ (min), in μs
100BASE-TX		30	20.5	5.0	15	10
1000BASE-T	Case-1	16.5	16.5	5.0	2.5	1.76
	Case-2	16.5	16.5	12.24	9.74	1.76
10GBASE-T	Case-1	7.36	7.36	4.48	0	2.88
	Case-2	4.48	4.48	1.6	0	2.88

EEE standard status

- It has been in the works for a LONG TIME!
 - The last two ballot cycles have resulted in no changes and the project has been submitted to RevCom which will make a recommendation to the IEEE Standards board end of this month, at which point it becomes an approved standard
- Call for interest Nov 2006
- Project authorized July 2007
- Baseline agreed upon July 2008
- Draft 1.0 Sept 2008
- Balloting by 802.3 July 2009
- Last technical change July 2010
- Publication is expected October 2010

Summary

- Its Ethernet, never perfect but will be everywhere!
- Volume will overcome complexity on the cost front
- Technology, ingenuity and investment will bring the power down