IEEE 802.3az: the Road to Energy Efficient Ethernet

IMDEA Networks, Leganés, November 2011

P. Reviriego
Universidad Antonio de Nebrija
previrie@nebrija.es

UNIVERSIDAD NEBRIJA

Our Group

- Research Activities
  - Energy Efficient Networks
    - Energy Efficient Ethernet
    - Network Monitoring to detect switched on PCs
  - Soft Errors
    - Memory Reliability
    - Error Correction Codes
    - Fault tolerant signal processing circuits

More info on:
http://www.nebrija.es/~jmaestro/esa/researchPapers.htm

Outline

- **Energy Efficiency in wireline PHYs**
  - **Issues and techniques to improve the efficiency**
- **Energy Efficient Ethernet (EEE)**
  - **Background on Ethernet PHYs**
  - **Tutorial on the EEE standard**
  - **Performance details of the EEE (803.az) standard**
- **Conclusions**

Energy Efficiency in wireline PHYs

- They have been designed for
  - Performance (speed, reach, etc).
  - Peak power consumption (encapsulation, ventilation, port density).
- Energy consumption not a design objective
  - Over dimensioned in many cases (100m in Ethernet)
  - Continuous transmission to keep PHY ready

→ Current PHYs are very inefficient

Energy Efficiency in wireline PHYs

- Techniques to improve Energy Efficiency
  - Adapt the link rate to the actual data traffic
  - Put the link in low power mode when there is no data to send
  - Adapt the Tx and Rx to the channel
    - Short channel $\rightarrow$ reduce Tx amplitude
    - Excess SNR $\rightarrow$ disable Rx functionality
  - For parallel links adapt the number of active links to data traffic

Ethernet Overview

- Used for wired LANs.
- Supports cable and fiber connections.
- Most PCs are equipped with Ethernet.
- Each year well over 100million Ethernet ports are shipped.
- Energy consumption of several Tera Watts hour

Ethernet over UTP: Standards

- Four Speeds
  - 10 Mbps: 10BASE-T (IEEE 802.3i ,1990)
  - 100 Mbps: 100BASE-TX (IEEE 802.3u ,1995)
  - 1000 Mbps: 1000BASE-T (IEEE 802.3ab ,1999)
  - 10000 Mbps: 10GBASE-T (IEEE 802.3an ,2006)
- The issue: in all (speeds except 10Mbps) the devices are active and consuming energy even when no data is transmitted or received.

Consumption figures (PHYs)

- Consumption varies with manufacturer, product generation, etc.
- Some rough estimates
  - 10Mbps (<0.1 W)
  - 100Mbps (0.2-0.3 W)
  - 1Gps (0.4-1 W)
  - 10Gbps (3-6 W)

Consumption beyond the PHY

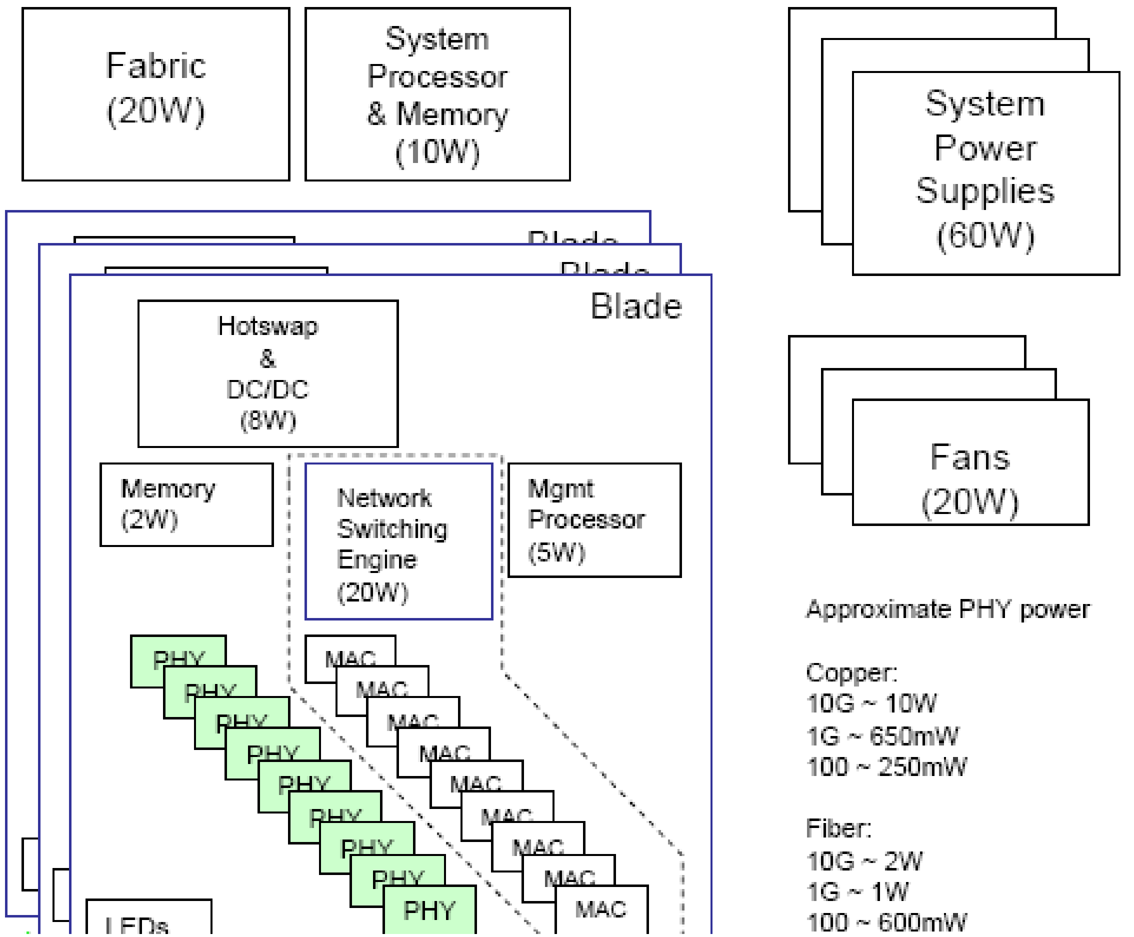

* From D. Dove "Energy Efficient Ethernet Switching Perspective" IEEE 802.3az January 2008 Meeting

Savings beyond the PHY

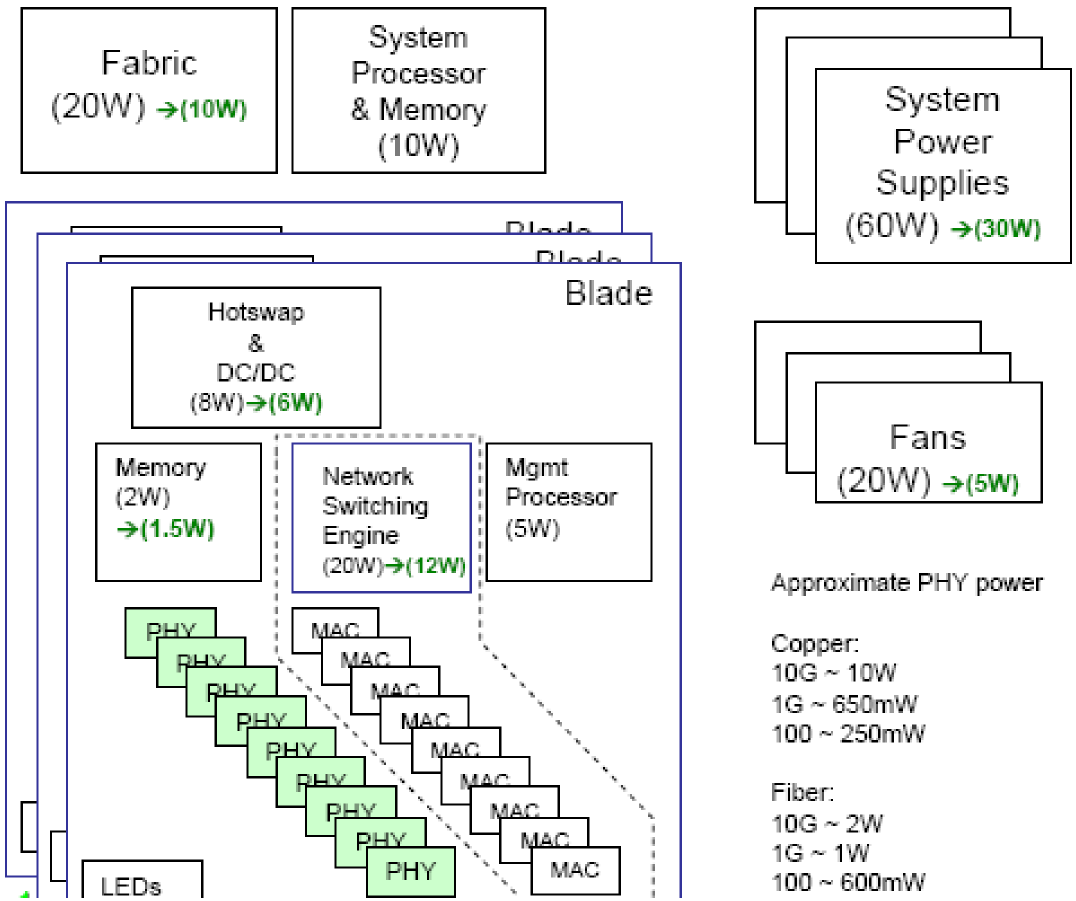

* From D. Dove “Energy Efficient Ethernet Switching Perspective” IEEE 802.3az January 2008 Meeting

Example of Energy Wasted

- User PC with Ethernet configured at 1Gbps.
- Traffic is less than 10Mbps for over 99% of the time.

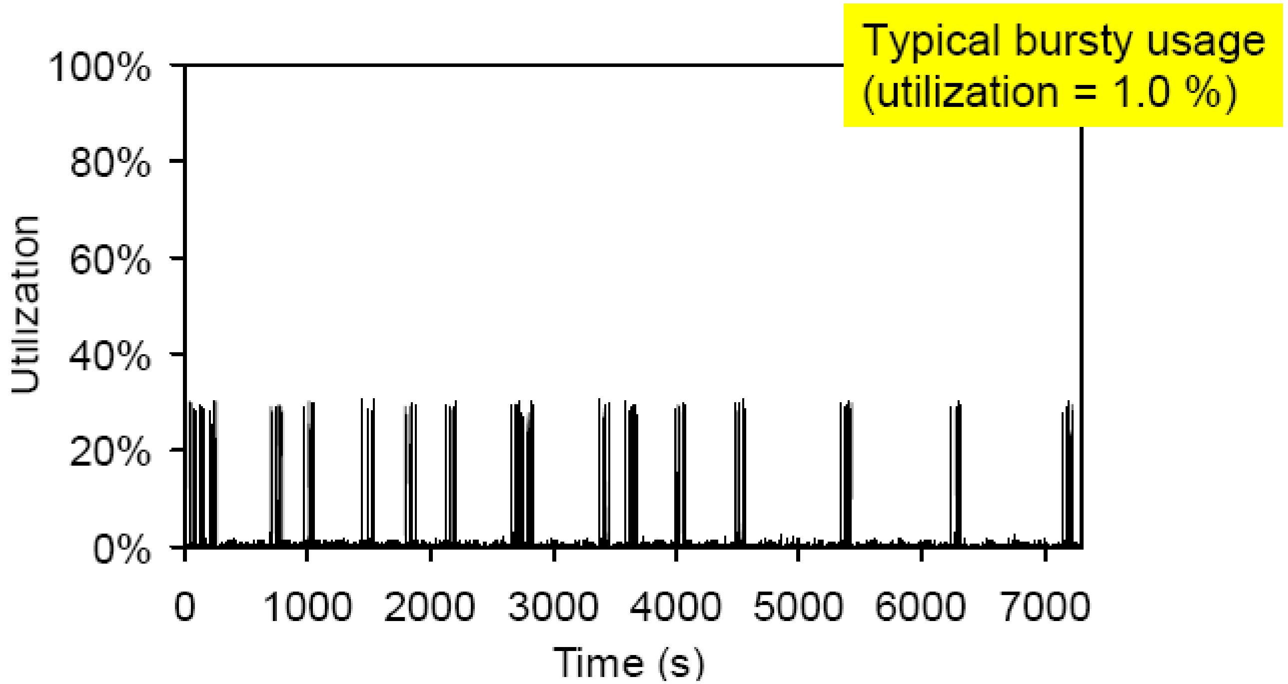

* From M. Bennett "Server Bandwidth Utilization plots" IEEE 802.3 meeting March, 2007

Example of Energy Wasted II

- The traffic load is low most of the time even for many servers.

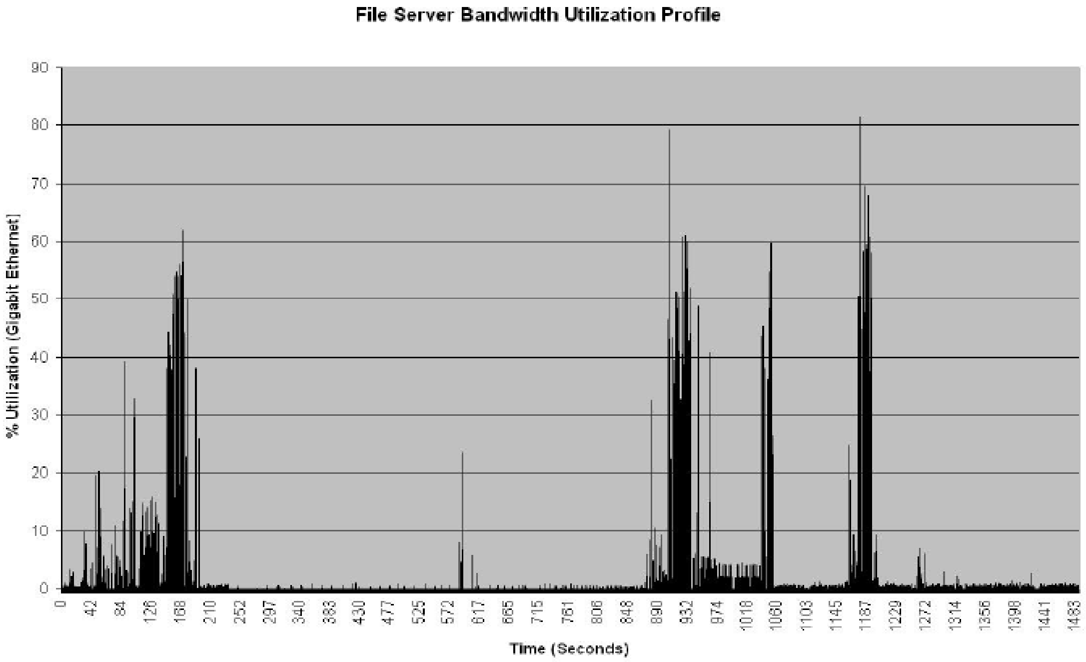

* From M. Bennett "Server Bandwidth Utilization plots" IEEE 802.3 meeting March, 2007

Energy Efficient Ethernet (EEE)

- New standard developed by IEEE
  - If there is no data to transmit or receive the device is put in low power mode.
  - Large energy savings can be achieved.
- Only effective when new devices that implement the standard are used.
  - Standard approved in 2010.
  - Supported now by most vendors in new products.
  - Only works if both ends of the link implement the standard.

* More info in K. Christensen, et al "IEEE 802.3az: The Road to Energy Efficient Ethernet", IEEE Communications Magazine, November 2010.

Energy Efficient Ethernet (EEE)

- Uses a low power mode when there is no data to transmit/receive:
  - Transition times between modes in the order of microseconds.
  - Savings close to 90% in low power mode.
  - Significant energy consumption during transitions between modes.

Energy Efficient Ethernet (EEE)

- Transitions between modes
  - Standard provides mechanisms, implementation decides when to change.
  - To maximize savings: sleep when there are no frames.
  - Standard also defines refresh periods to keep receiver aligned with transmitter.

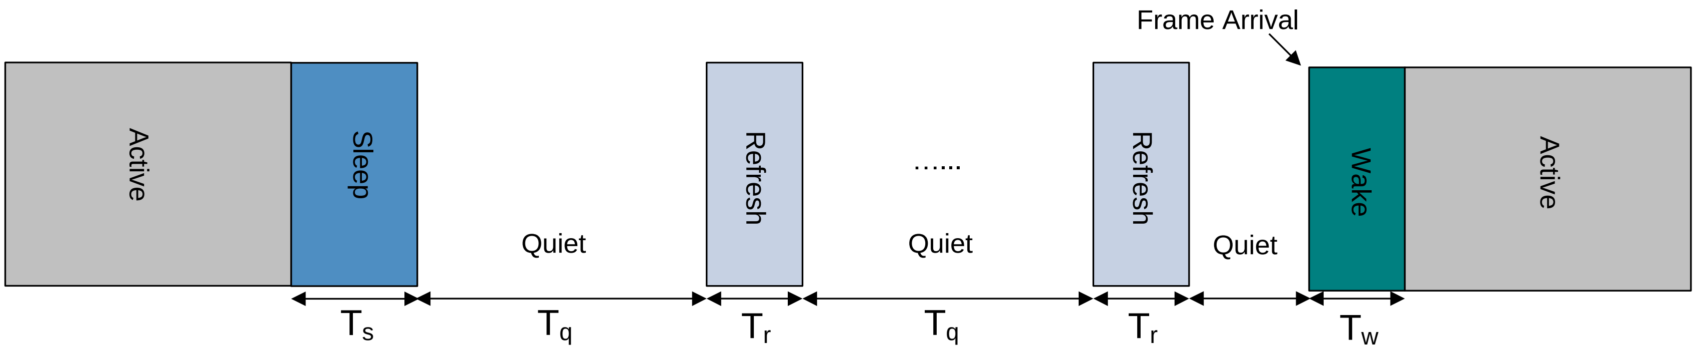

Energy Efficient Ethernet (EEE)

- Savings will be proportional to the percentage of time that the link is in low power mode.
- This will depend on the number of frames and the inter-arrival pattern as transition times are significant.

Technology	100Base-TX	1000Base-T	10GBase-T
T_w (min)	$30.5\mu s$	$16.5\mu s$	$4.48\mu s$
T_s (min)	$200\mu s$	$182\mu s$	$2.88\mu s$
T_{frame} (1500 bytes)	$120\mu s$	$12\mu s$	$1.2\mu s$
Tx efficiency (1500 bytes)	34.2%	5.7%	14%
T_{frame} (150 bytes)	$12\mu s$	$1.2\mu s$	$0.12\mu s$
Tx efficiency (150 bytes)	4.9%	0.6%	1.6%

Energy Efficient Ethernet (EEE)

- Let us assume
  - Sleep if there are no frames wake on frame arrival.
  - Energy consumption during transitions is equal to that of active mode
  - Savings in Low Power mode are 90%
  - Poisson Arrivals
  - 1250 bytes frames.

Energy Efficient Ethernet (EEE)

- 100Base-TX and 1000Base-T

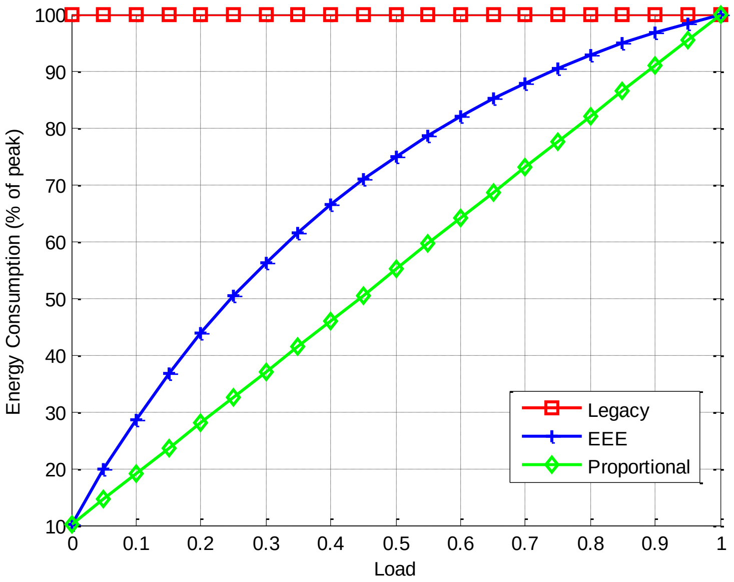

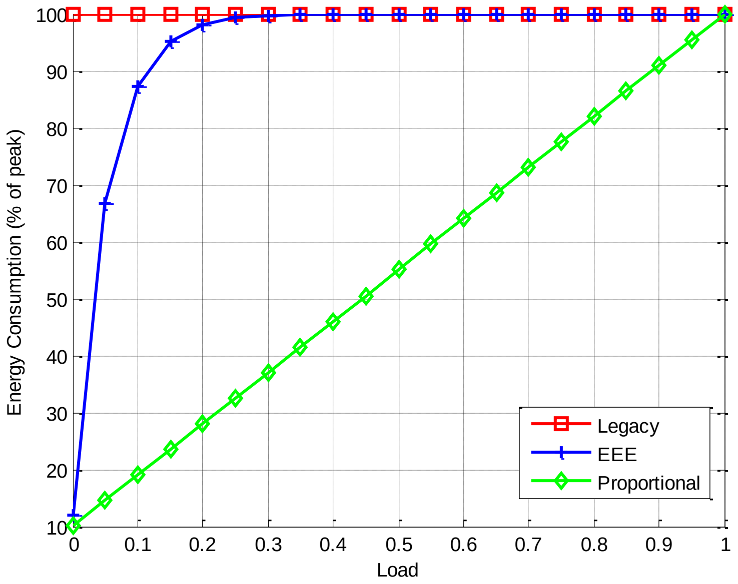

* From P. Reviriego, et al, "Performance Evaluation of Energy Efficient Ethernet", IEEE Communications Letters, September 2009,

Energy Efficient Ethernet (EEE)

- 10GBase-T

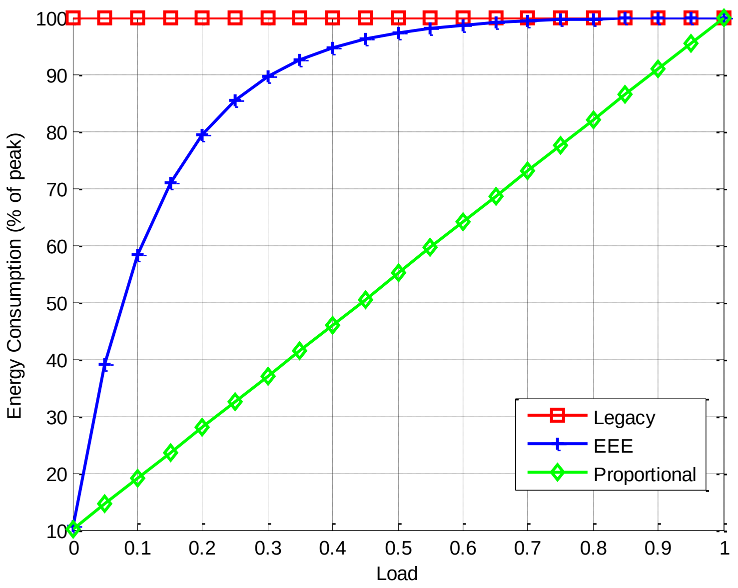

* From P. Reviriego, et al, "Performance Evaluation of Energy Efficient Ethernet", IEEE Communications Letters, September 2009,

Energy Efficient Ethernet (EEE)

- Simulations show
  - Energy consumed during transitions is significant.
  - Therefore inter arrival patterns are important.
- But
  - Poisson assumes independent arrivals, this is not the case for LANs with low aggregation.
  - Only one frame size was used (1250 bytes).

Energy Efficient Ethernet (EEE)

- Analysis of real frame traces: University Access Link at 1Gbps

	Direction	Link Load %	Average Frame size (Bytes)	Energy ideal % of peak	Energy % of peak
University Link	Download	10.94	679	19.84	92.80
University Link	Upload	17.66	919	25.89	96.20

* From P. Reviriego, et al, "Performance Evaluation of Energy Efficient Ethernet", IEEE Communications Letters, September 2009,

Energy Efficient Ethernet (EEE)

- Analysis of real frame traces: Datacenter servers at 1Gbps

Server Type	Direction	Link Load %	Average Frame size (Bytes)	Energy ideal % of peak	Energy % of peak
File & Search	Input	1.22	87	11.10	65.90
File & Search	Output	52.21	1497	56.99	72.92
Search	Input	8.51	945	17.66	45.28
Search	Output	7.23	934	16.51	42.30
File and App	Input	0.65	130	10.58	61.37
File and App	Output	4.02	749	13.62	57.10

* From P. Reviriego, et al, "Performance Evaluation of Energy Efficient Ethernet", IEEE Communications Letters, September 2009,

Energy Efficient Ethernet (EEE)

- Real Measurements
  - Using Realtek NICs

Figure. 1 Photograph of the NIC used in the experiments

Energy Efficient Ethernet (EEE)

- Real Measurements
  - Using Realtek NICs
  - Focusing on 1Gbps
  - Power measured at PCIe supply
  - Measurements include PHY+MAC+PCIe interface
  - Traffic shaped with the Linux TC tool

Energy Efficient Ethernet (EEE)

- Real Measurements

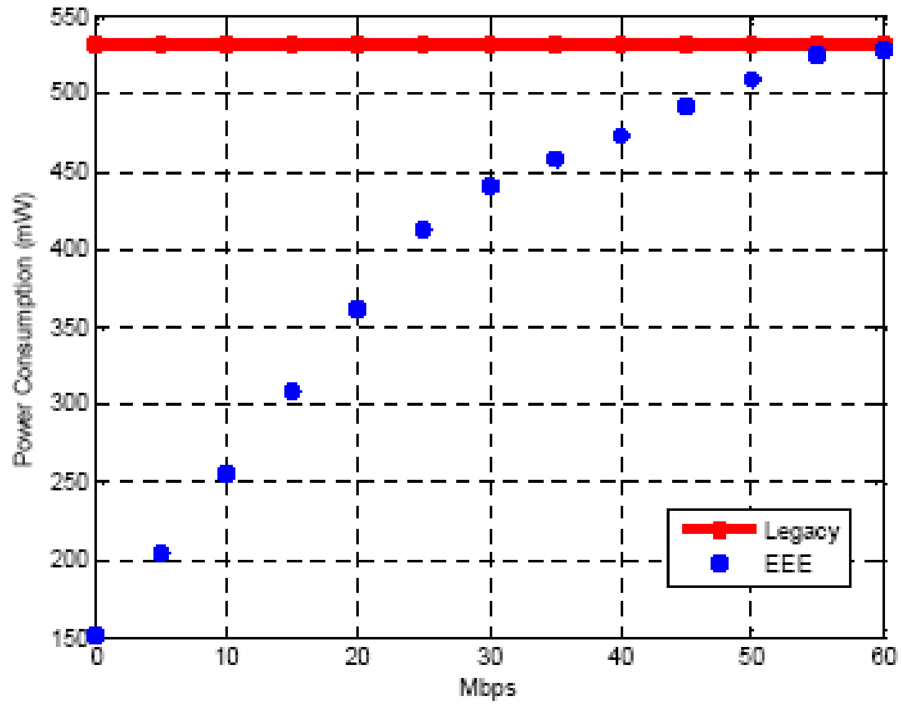

Figure. 2 NIC Power Consumption versus load for 1000Base-T (1 Gbps)

* From P. Reviriego et al "An Initial Evaluation of Energy Efficient Ethernet", IEEE Communications Letters, May 2011.

Energy Efficient Ethernet (EEE)

- One Alternative is to use Burst Transmission

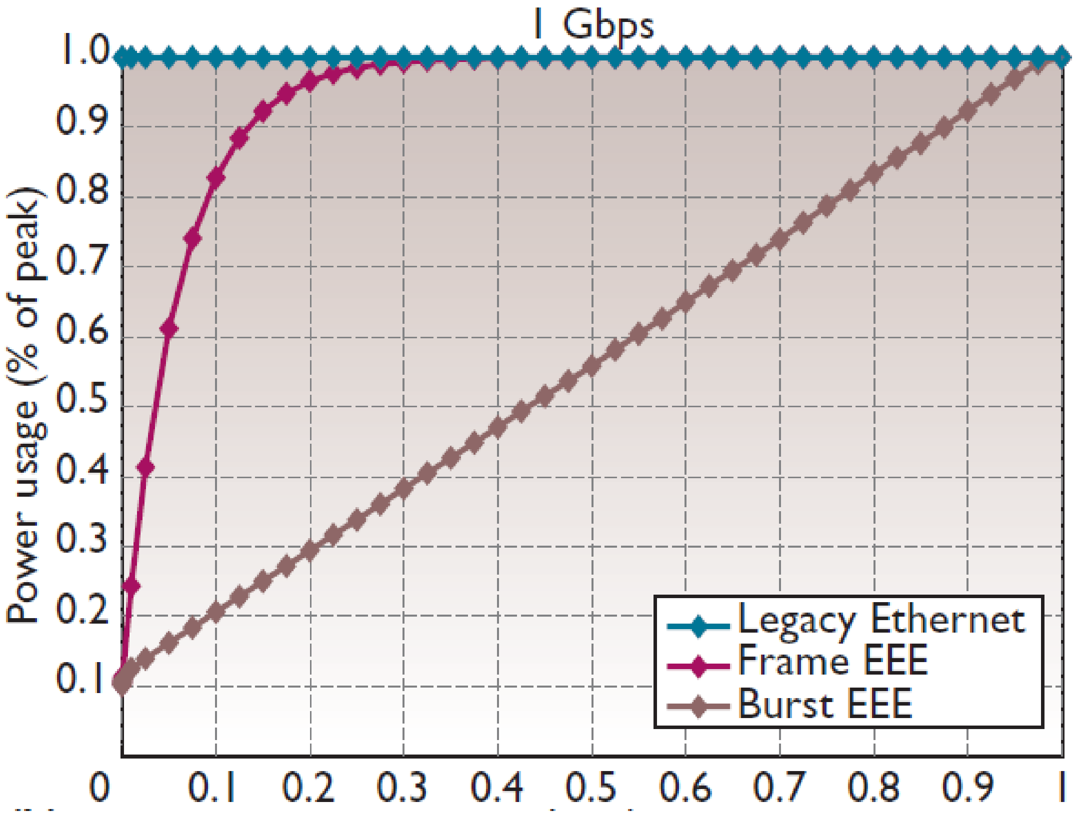

* From P. Reviriego et al “Burst Transmission for Energy Efficient Ethernet", IEEE Internet Computing, August 2010.

Energy Efficient Ethernet (EEE)

- But impacts
  - Packet Delay
  - Queue dynamics
  - Probability of frames being discarded

Conclusions

- Energy Efficient Ethernet will be deployed massively in a few years
- Energy Efficient Ethernet enables savings of the order of some Tera Watts hour per year

but

- There are performance issues at medium loads

Future Work

- Currently the IEEE 802.3 is working on a 100G standard for Datacenter Connectivity (802.3bj)
- The new standard
  - Connects servers in the data center
  - Uses backplane and short reach copper
  - Will consider EEE as a design goal
  - Implementing LPI is challenging (Load and Speed)

To learn more

- Our publications in this area

  http://www.nebrija.es/~jmaestro/esa/researchPapers.htm

- Energy Efficient Ethernet standard

  http://www.ieee802.org/3/az/index.html

- Energy Efficiency in the Internet

  http://www.csee.usf.edu/~christen/energy/pubs.html