Energy Efficient Ethernet Call-For-Interest

IEEE 802.3 Working Group

Dallas, TX

November 14, 2006

Objective for this Meeting

- **To *measure the interest* in starting a study group for Energy Efficient Ethernet**
  - ***Reduce power during low link-utilization***
  - ***Compatible with existing cabling infrastructure***

- **We don't need to**
  - **Fully explore the problem**
  - **Debate strengths and weaknesses of solutions**
  - **Choose any one solution**
  - **Create PAR or five criteria**
  - **Create a standard or specification**

- **Anyone in the room may speak / vote**
- **RESPECT… give it, get it**

Agenda

- **Presentations**
  - **“Network Energy Use,” Bruce Nordman**
  - **“Reducing Ethernet Energy Use,” Hugh Barrass**
  - **“The Feasibility of Energy Efficient Ethernet ,” Howard Frazier**
  - **“Why Energy Efficiency Now?,” Mike Bennett**
- **Discussion**
  - **Please hold questions till the end of the presentations**
- **Call for Interest**
- **Future Work**

Network Energy Use

Presented by
Bruce Nordman
Lawrence Berkeley National Lab
U.S. Department of Energy

IEEE 802.3 Working Group
Dallas, TX

November 14, 2006

The problem

Numbers represent U.S. only

- **"Big IT" – all electronics**
  - **PCs/etc., consumer electronics, telephony**
    - **Residential, commercial, industrial**
  - **200 TWh/year**
  - **$16 billion/year**
  - **Nearly 150 million tons of CO_2 per year**
    - **Roughly equivalent to 30 million cars!**

One central baseload power plant (about 7 TWh/yr)

PCs etc. are digitally networked now — *Consumer Electronics* (CE) will be soon

The problem

- **“ Little IT” — office equipment, network equipment, servers**
  - 97 TWh/year
    - **3% of national electricity**
    - **9% of commercial building electricity**
    - **Almost $8 billion/year**

	TWh/year
– **Network Equipment**	**13**
– **Servers**	**12**
– **PCs / Workstations**	**20**
– **Imaging (Printers, etc.)**	**15**
– Monitors / Displays	22
– UPS / Other	16

60% Networked Equipment

- **… However**
  - **Old data (energy use has risen)**
  - **Doesn’t include residential IT or networked CE products**

Note: Year 2000 data taken from Energy Consumption by Office and Telecommunications Equipment in Commercial Buildings--Volume I: Energy Consumption Baseline Roth et al., 2002 Available at: http://www.eren.doe.gov/buildings/documents

The problem

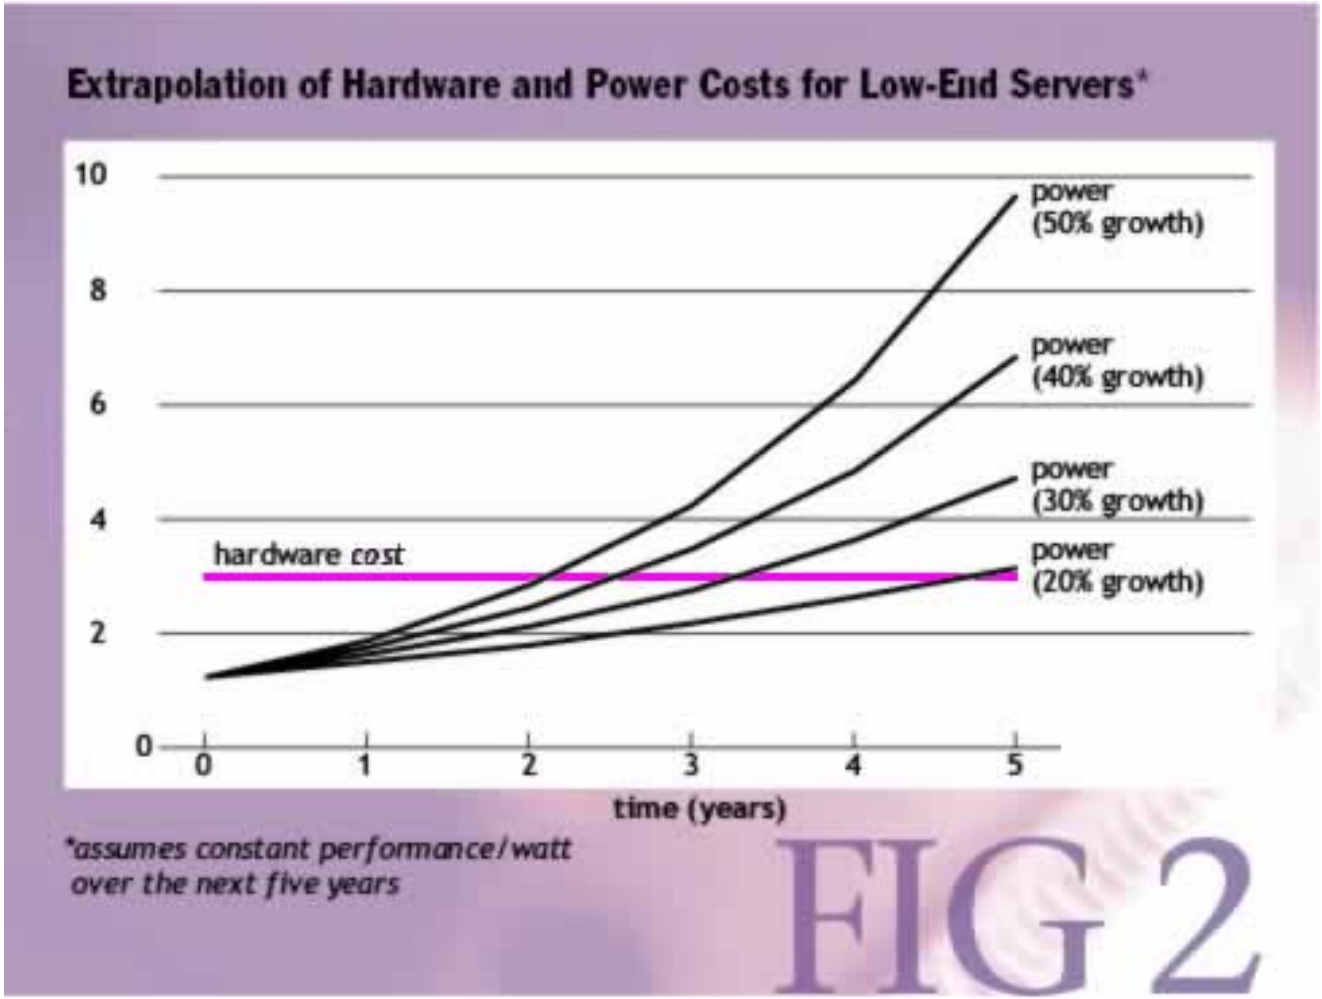

Unrestrained IT power consumption could eclipse hardware costs and put great pressure on affordability, data center infrastructure, and the environment.

Source: Luiz André Barroso, (Google) "The Price of Performance," *ACM Queue*, Vol. 2, No. 7, pp. 48-53, September 2005.

(Modified with permission.)

Server industry response

- **Concept and real data showing how server power drops with computing load**
- **The SPECpower committee is currently defining a metric for this**

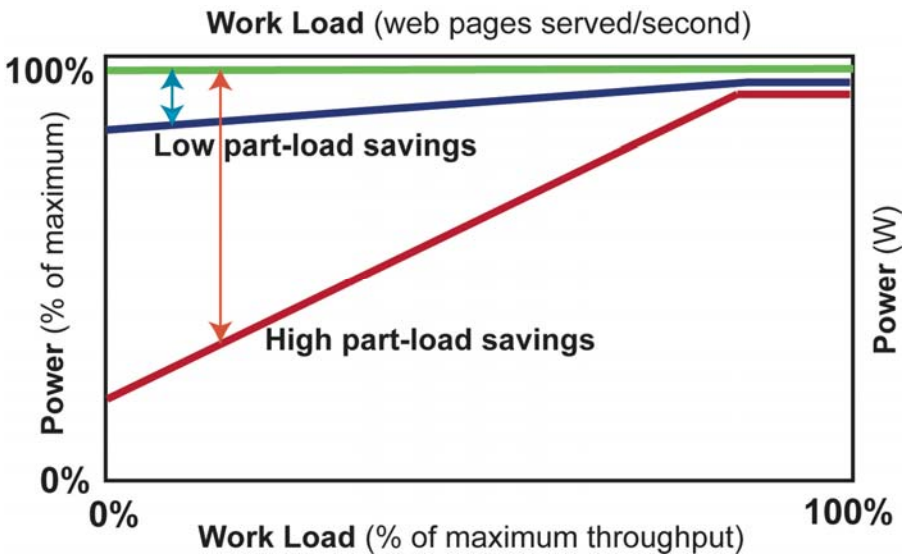

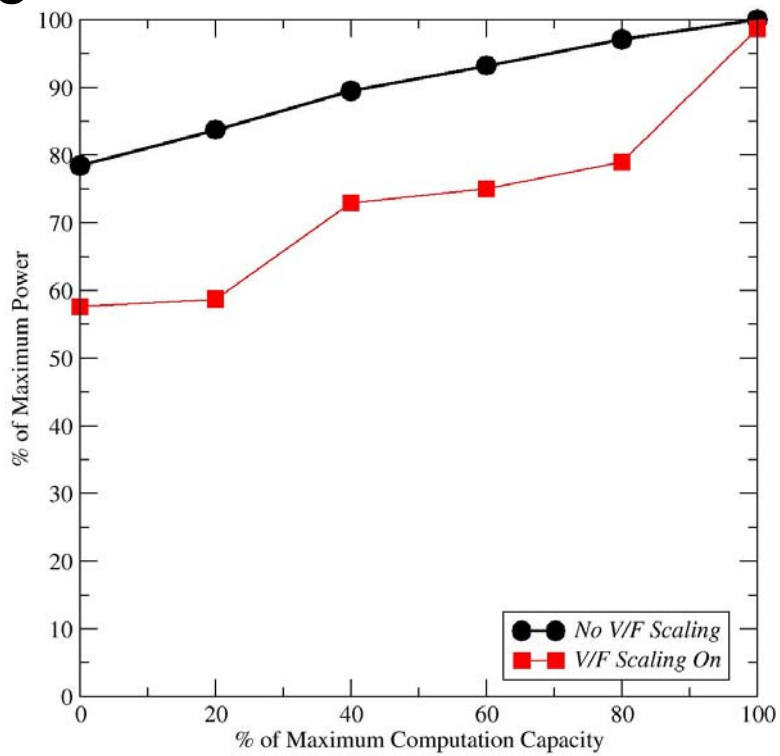

Energy industry responds

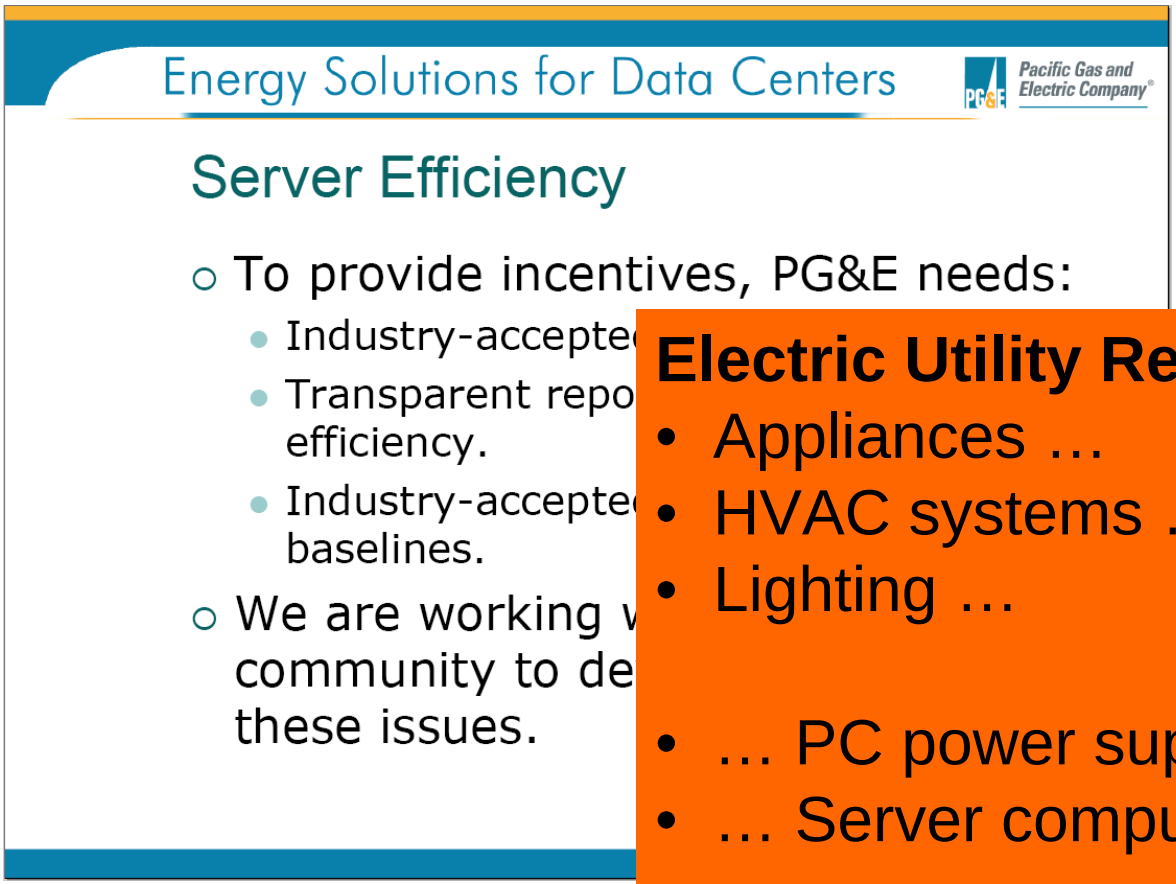

Electric Utility Rebates

- Appliances …
- HVAC systems …
- Lighting …

- … PC power supplies (2005)
- … Server computers (2006)
- …

Reference: http://www.pge.com/docs/pdfs/biz/rebates/hightech/DataCenters_slides.pdf

PG&E (California) provides rebates for more energy-efficient servers

Link power

Results from (rough) measurements

— all incremental AC power

— measuring 1st order

- **Typical switch with 24 ports 10/100/1000 Mb/s**

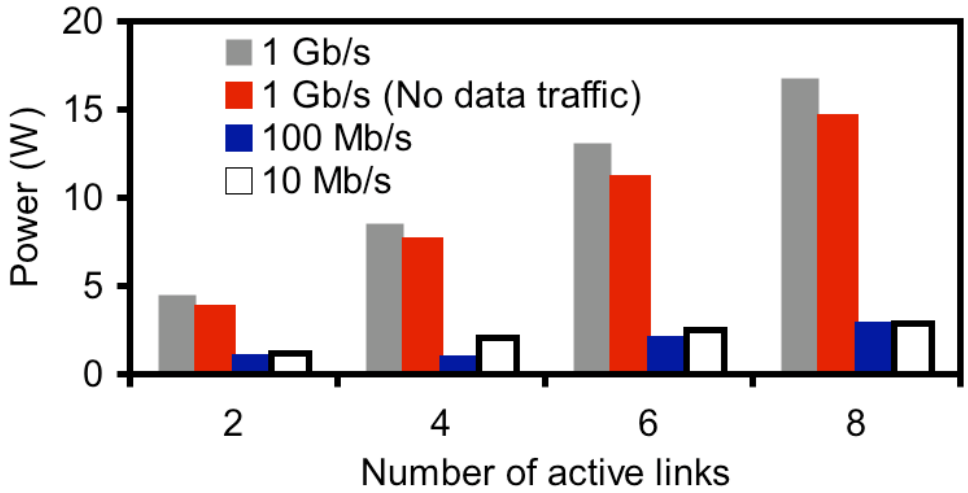

- **Various computer NICs averaged**

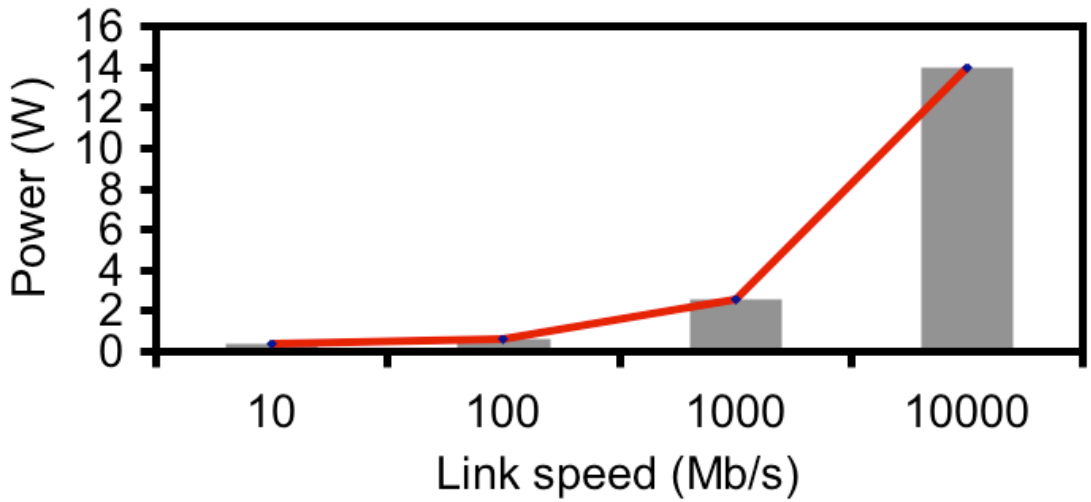

Data center context

- Networking power consumption significant portion of energy use in data centers
- Other components increasingly can modulate power to needed performance
- More and more data centers running into power and cooling infrastructure limits
- Server manufacturers all have jumped onto energy efficiency bandwagon

- Network industry can ‘go green’ with Energy Efficient Ethernet

Network industry design criteria

- **Past/current practice**
  - **Design for maximum performance and ensure maximum power condition can be powered / cooled.**

- **Future practice needs**
  - **Design for normal usage, ensuring maximum energy efficiency at that operating point**
  - **Lower energy use at lower utilization**
  - **Design for minimum energy usage over operational lifetime**

Network Energy Use

Presented by
Hugh Barrass
Cisco
(for Ken Christensen, University of South Florida)

IEEE 802.3 Working Group
Dallas, TX

November 14, 2006

This material is based upon work funded by the National Science Foundation under grant CNS-0520081 for Christensen.

Link utilization

- **Desktop-to-switch links**
  - **Are mostly idle**
    - **Lots of very low bandwidth "chatter"**
  - **High bandwidth needed for bursts**
    - **Bursts are often seconds to hours apart**
- **Server links are also often not fully utilized**
  - **Higher speed links offer more opportunity to save energy**
  - **This is an area where more data is needed**
- **Evidence of low utilization (desktop users)**
  - **LAN link utilization is generally in range 1 to 5% [1, 2]**
  - **Utilization for "busiest" user in USF was 4% of 100 Mb/s**

[1] A. Odlyzko, "Data Networks are Lightly Utilized, and Will Stay That Way", *Review of Network Economics*, Vol. 2, No. 3, pp. 210-237, September 2003.

[2] R. Pang, M. Allman, M. Bennett, J. Lee, V. Paxson, and B. Tierney, "A First Look at Modern Enterprise Traffic," *Proceedings of IMC 2005*, October 2005

Desktop links have low utilization

- **Snapshot of a typical 100 Mb Ethernet link**
  - **Shows time versus utilization (trace from Portland State Univ.)**

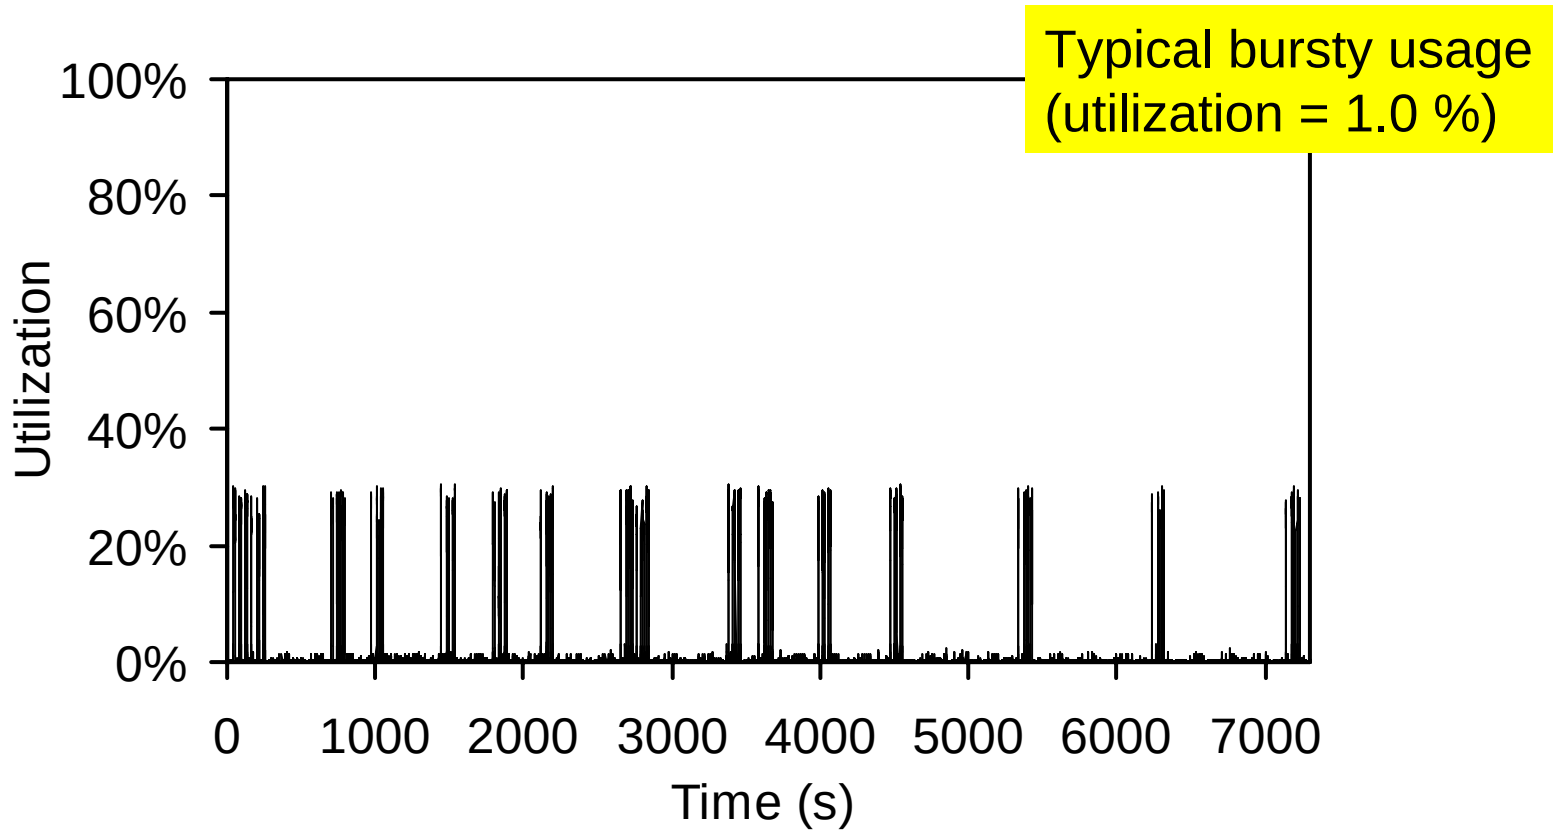

Fig1.xls

Reducing the link rate

- Can (and does) save energy
- Some NICs drop link rate when a laptop is battery powered
  - Or, when a PC goes into sleep state
  - Turns-off PHY if no signal on link
- Match the link rate to utilization
  - High utilization = high link rate
  - Low utilization = low link rate
- Currently implemented using auto-negotiation
  - Set the Technology ability bits/message codes and then reset the link
  - Takes about 1000 milliseconds (a looooooooong time)

Need for fast transitions

- **Can extend the benefits of link data rate reduction**
  - **By making the data rate transition faster**

- **Need a faster way than auto-negotiation**
  - **When I need high data rate, I need it now**
  - **Can't advertise the desire to change to a higher speed**
- **Need a mechanism that is transparent to upper layers**
- **Need a *standard mechanism* to rapidly transition:**
  - **From low to high data rate**
  - **From high to low data rate**
  - **Within the capabilities established by Auto Negotiation**

Matching Link Rate to Utilization

- **Rapid PHY Selection (RPS) includes**
  - **A PHY selection mechanism**
  - **A control protocol**

- **Open PHY-level challenges**
  - **How fast to handshake?**
  - **How to re-synchronize for 1 Gb/s**
  - **How to re-synchronize for 10 Gb/s?**

- **Speed control policy issues need to be studied**

Rapid PHY Selection

- **One possibility… MAC frame handshake**

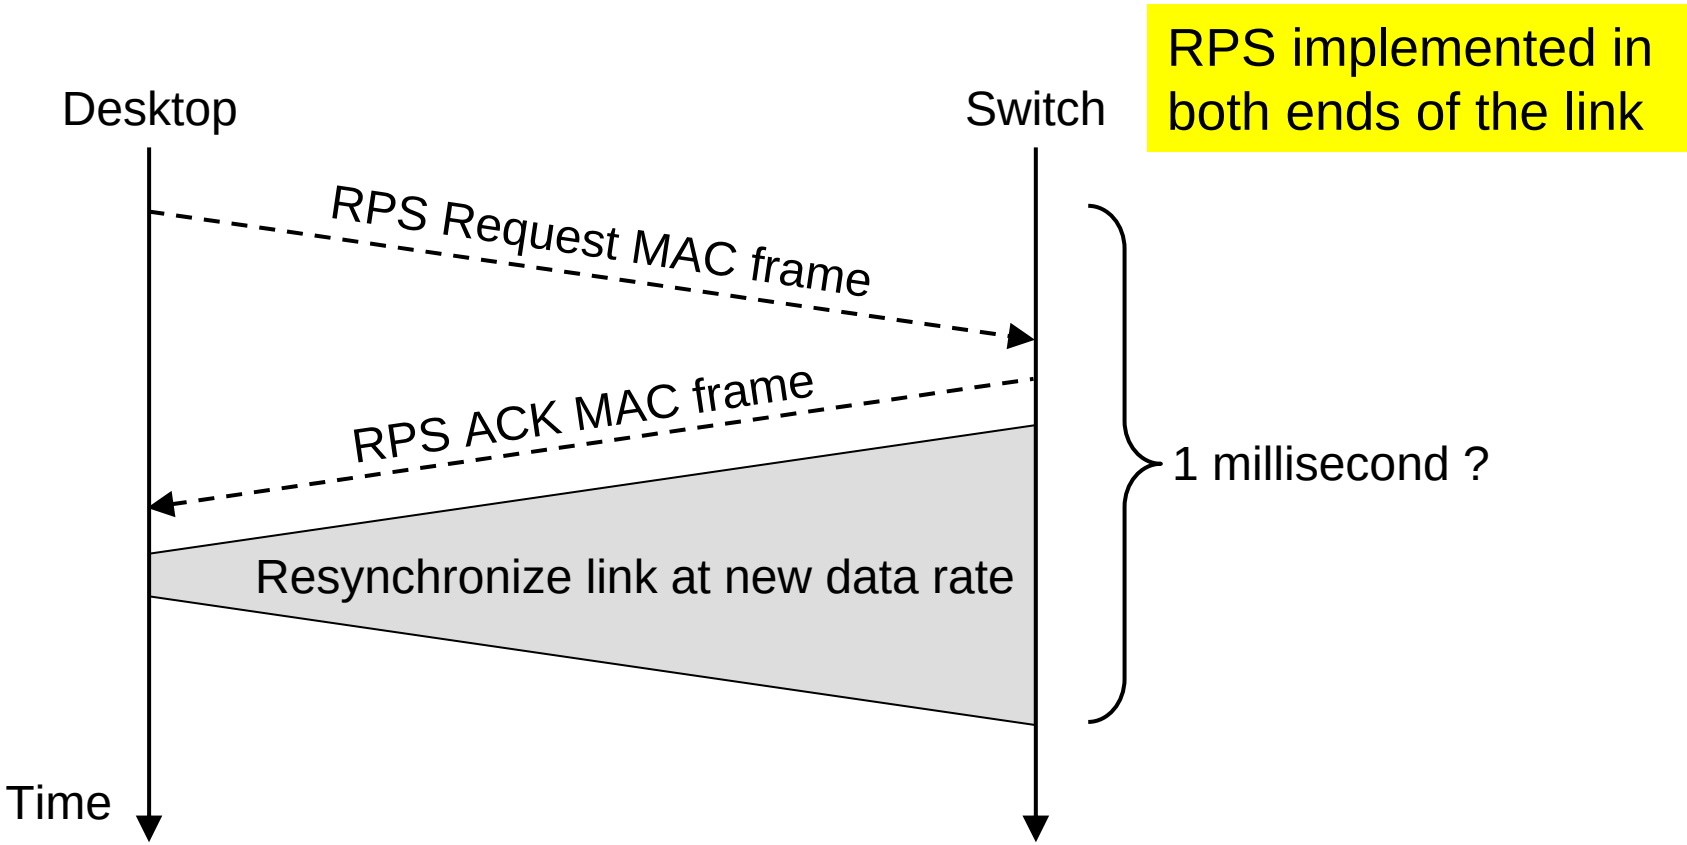

Control policy use of RPS

- **Control policy determines when to transition data rate**
  - **RPS can support many different control policies**
  - **Need to consider but not define**

- **Trade-off of energy saved versus packet delay**
  - **Energy savings achieved by operating at low data rate**
  - **Delay occurs during transition from low to high data rate**

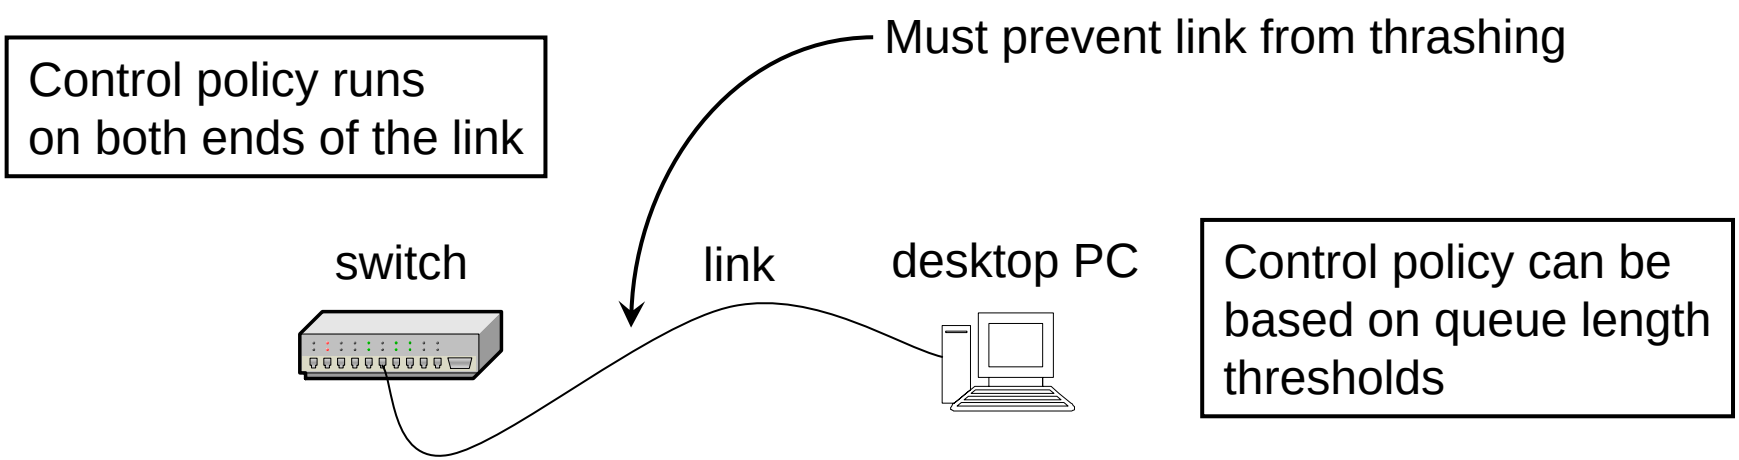

RPS – a picture tells the story

- **Snapshot of a typical Ethernet link with simulated RPS**

Mean packet delay with RPS is 0.67 ms
Mean packet delay without RPS is 0.12 ms
Utilization-threshold policy is used

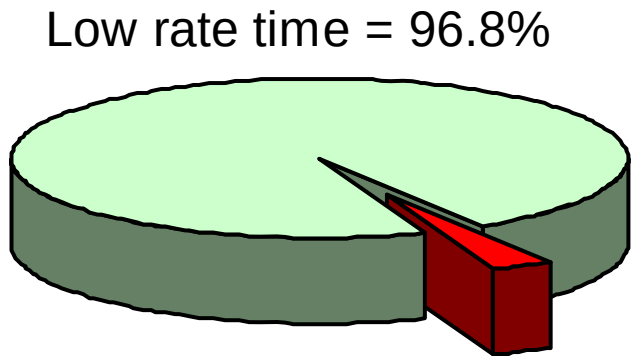

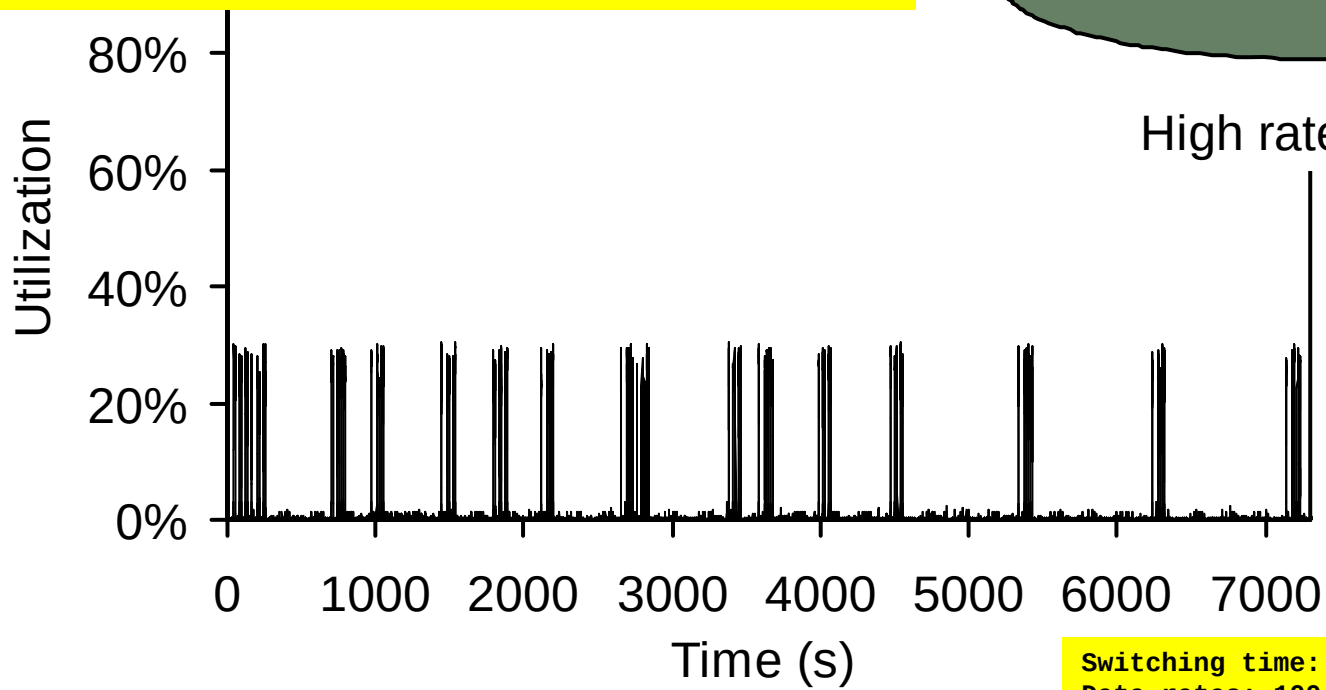

Switching time: 1ms
Data rates: 100 and 10 Mb/s
Low/High thresh: 0KB and 32KB

Benefits of Energy Efficient Ethernet

- **1 Gb/s**
  - **Most NICs and most energy to be saved**
  - **Substantial benefits for homes and offices**
  - **Battery life benefit for notebooks**
- **10 Gb/s (copper)**
  - **Reduces power burden in data centers**
  - **Reduces cooling burden in data centers**
  - **May increase switch/router port capacity**
- **Generally…**
  - **Provides real economic benefit through energy savings**

Potential Savings from EEE

Assume 100% adoption (U.S. Only)

- **Residential**
  - **PCs, network equipment, other**
  - **1.73 to 2.60 TWh/year**
  - **$139 to $208 million/year**
- **Commercial (Office)**
  - **PCs, switches, printers, etc.**
  - **1.47 to 2.21 TWh/year**
  - **$118 to $177 million/year**
- **Data Centers**
  - **Servers, storage, switches, routers, etc.**
  - **0.53 to 1.05 TWh/year**
  - **$42 to $84 million/year**

These figures do **not** include savings from cooling/power infrastructure

Total: $298 to $469 million/year

Summary

- **IT Energy use accounts for a substantial portion of overall consumption**
  - **Networks are an increasing part of this**
- **Customers want to lower operational expenses**
  - **Reducing energy consumption saves $$**
  - **Secondary effects: lower cooling requirements – saves $$**
- **Energy Efficient Ethernet targets low-hanging fruit (cost, time, market introduction)**

The Feasibility of Energy Efficient Ethernet

Presented by
Howard Frazier – Broadcom

IEEE 802.3 Working Group
Dallas, TX
November 14, 2006

Premise

- **There are many ways to save power when the link utilization is low**
  - **Goal is to do so while minimizing impact on standard and industry**
- **This presentation assumes the use of a simple, "brute force" approach**
  - **Use a high data rate PHY when the utilization is high**
  - **Use a low data rate PHY when the utilization is low**
  - **Assume that a low data rate PHY consumes less power than a high data rate PHY**

Transition time

- A key topic for further study …
- Factors to be considered
  - switch/bridge buffer size
    - Shorter transition time minimizes disruption
- Use 1 ms transition time as a starting point for discussion
- Auto-negotiation is at least two orders of magnitude too slow

Sequence

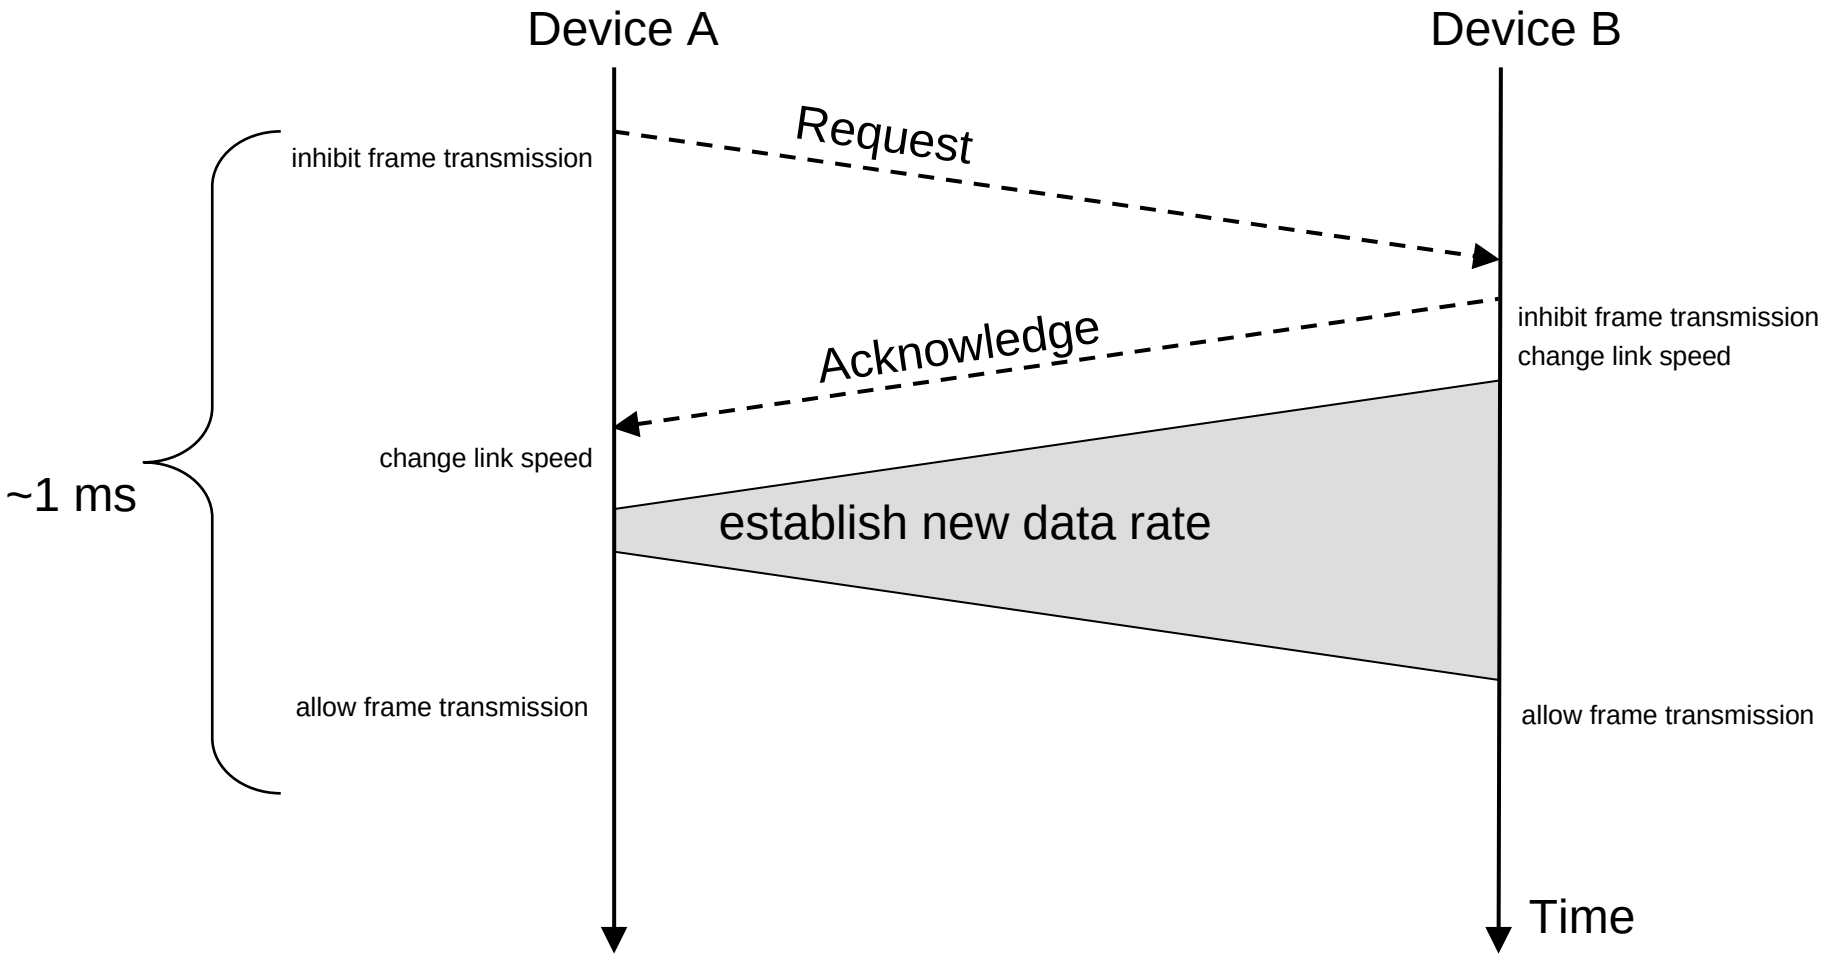
Device A
Device B
Request
inhibit frame transmission
Acknowledge
inhibit frame transmission
change link speed
change link speed
~1 ms
establish new data rate
allow frame transmission
allow frame transmission
Time

Protocol Candidates

- **LLDP (802.1AB)**
  - **with anticipated enhancements**
- **OAM (802.3ah)**
- **other slow protocol (802.3ad)**
- **MAC Control frame**
- **Physical layer signaling**
  - **e.g. special idle sequences or ordered sets**

Stored link parameters

- **Another subject for further study …**
- **E.g. equalizer coefficients**
- **May permit faster link speed changes.**
- **The power consumption associated with parameter storage is much less than the power consumption associated with operating the link**

Summary

- **Use 1 ms transition time as a starting point for discussion**
- **Several candidate protocols exist**
- **Link parameter storage requirements are reasonable**
- **Energy Efficient Ethernet is feasible**

Energy Efficient Ethernet
Why Now?

Presented by
Mike Bennett
Lawrence Berkeley National Lab
U.S. Department of Energy

IEEE 802.3 Working Group
Dallas, TX
November 14, 2006

Why Energy Efficiency Now?

- **Networking industry has an opportunity to catch up with the server industry in this area**
  - **April 19, 2006 "Green Grid" formed**
    - **"**A group of technology industry leaders form The Green Grid to help reduce growing power and cooling demands in enterprise datacenters."
  - **Energy industry incentives for efficient products**
- **Energy Star**
  - **Requirements coming in 2009**
    - **"All computers shall reduce their network link speeds during times of low data traffic levels in accordance with any industry standards that provide for quick transitions among link rates"**
- **Customers like saving energy because it reduces operating costs**

Reference: ENERGY STAR® Program Requirements for Computers (final draft, tier 2 requirements) available at http://www.energystar.gov/ia/partners/prod_development/revisions/downloads/computer/ComputerSpec_Final_Draft.pdf

Why Energy Efficiency Now?

- **Energy Efficiency gets U.S. congressional recognition**
  - **July 12, 2006 House Resolution 5646 passes**
    - **"To study and promote the use of energy efficient computer servers in the United States"**
  - **What's next?**
- **The market for energy efficient Ethernet**
  - **Driven by customer's desire to save energy costs**
  - **Ethernet is used in markets where saving energy is crucial**
  - **Accelerate deployment for new applications**
  - **Enables use of incentives by energy industry**
  - **Ultimately these translate to increased demand**

Things to be studied …

- **Reducing energy consumption during periods of low link utilization is low hanging fruit**
  - **This will be the focus of a study group, if formed**
- **Issues that need study**
  - **How to minimize transition time?**
  - **How to avoid thrashing between speeds?**
  - **Interaction with higher-layer protocols**
  - **Link utilization on servers**
  - **Interaction with control policy**

Why Energy Efficiency Now?

- **Reducing energy consumption during periods of low link utilization is technically feasible**
  - **Can leverage existing work, e.g. LLDP, OAM, etc.**
- **802.3 can best define Energy Efficient Ethernet**
- **End users want lower power-consuming products**
  - **Power saving modes have been a desirable feature in electronic devices for quite a long time**
- **We have a chance to do something good for the planet**
  - **Not a bad thing to do**

Supporters

David Law	3Com
Brad Booth	AMCC
Bill Woodruff	Aquantia
Scott Powell	Broadcom
Wael Diab	Broadcom
Li Tienan	China Standard Certification Center
Claudio DeSanti	Cisco
Fred Schindler	Cisco
Andrew Fanara	EPA
Paolo Bertoldi	European Commission DG JRC
Joel Goergen	Force10 Networks
John D'Ambrosia	Force10 Networks
Steve Carlson	HSP Design
Petar Pepeljugoski	IBM Research
Ilango Ganga	Intel
David Chalupsky	Intel
Mike McConnell	KeyEye Communications
Ted Sopher	Lawrence Berkeley National Lab
Gopi Sirineni	Marvell
Alan Flatman	LAN Technologies
Geoff Thompson	Nortel Networks
Joseph Babanezhad	Plato Networks
Geoff Garner	Samsung
Eric Ryu	Samsung
George Zimmerman	Solarflare
Shimon Muller	Sun
Mark Bowman	Tennessee Valley Authority
Rahul Chopra	Teranetics
Sanjay Kasturia	Teranetics
Bob Noseworthy	UNHIOL
Adam Bechtel	Yahoo!

Questions?

Straw Polls

Call-For-Interest

- **Should a Study Group be formed for “Energy Efficient Ethernet”?**

  Y: **N:** **A:**

Participation

- **I would participate in the "Energy Efficient Ethernet" Study Group in IEEE 802.3.**

  Tally:

- **My company would support participation in the "Energy Efficient Ethernet" Study Group in IEEE 802.3**

  Tally:

Future Work

- **Ask 802.3 to form EEE SG on Thursday**
- **If approved**
  - **802 EC authorization of EEE SG on Friday**
  - **First EEE SG meeting, week of January 2007 IEEE 802.3 Interim.**

Thank You!