Green Networks: The Next Steps

Ken Christensen
Department of Computer Science and Engineering
University of South Florida
Tampa, Florida USA 32620
christen@cse.usf.edu
http://www.csee.usf.edu/~christen

Funding for this work from NSF (CNS-0520081 and CNS-0721858) and Cisco

USF UNIVERSITY OF SOUTH FLORIDA

Invited talk at LatinCom 2010 – Bogotá, Colombia

Thank you

Thank you for inviting me to give a talk. I am very honored and excited to be here.

Acknowledging my students

Some of the work presented here was done by past and present students including,

- Chamara Gunaratne (PhD in 2008)
  - Early Proxying and Ethernet work
- Miguel Jimeno (PhD in 2010)
  - Proxying (especially for applications)
- Mehrgan Mostowfi (Current PhD student)
  - Recent Ethernet work

From Colombia!

USF UNIVERSITY OF SOUTH FLORIDA

Why green networks?

One of the most urgent challenges of the 21st century is to investigate new technologies that can enable a transition towards a more sustainable society with a reduced CO_2 footprint.

→ We need to reduce energy consumption

Here is one reason why

Sea level in 2100 under high emissions scenario

I live here

From U.N. Intergovernmental Panel on Climate Change

The challenge to ICT

What role will ICT play in this grand challenge?

- Reduce energy consumed by ICT
- Enable energy savings in non-ICT

Subject of other talks today

ICT = Information Communications Technology

Conservation versus efficiency

Two ways to consume less energy…

1) Have and do less = conservation
2) Improve performance = efficiency

What we seek

"I mean using less energy for identical performance, measured in whatever way the consumer wishes."
- Richard Muller (*Physics for Future Presidents*, 2008)

➡ We seek to have the same QoS/QoE for less energy

Roadmap of this talk

This talk has four major topics

- Quantifying the energy use of ICT
- Reducing direct energy consumption
- Reducing induced energy consumption
- Future challenges

Key definitions

Direct energy use

- Energy used by network links and equipment, but not hosts

Induced energy use

- Incremental additional energy used for a higher power state of hosts needed to maintain network connectivity

Quantifying the energy use of ICT

How much energy does ICT consume?

... the Internet is part of this

A view from the Climate Group

The SMART 2020 report

- Focus is on ICT’s role in reducing greenhouse gases
  - Both of and by ICT
- A view of the world in 2020
  - Taking into account “likely” technology developments
- Supporting organizations
  - Include Cisco, Intel, HP, Sun, national telecoms, and telecom operators

Global ICT CO_2 footprint

Today ICT is 2% of global CO_2

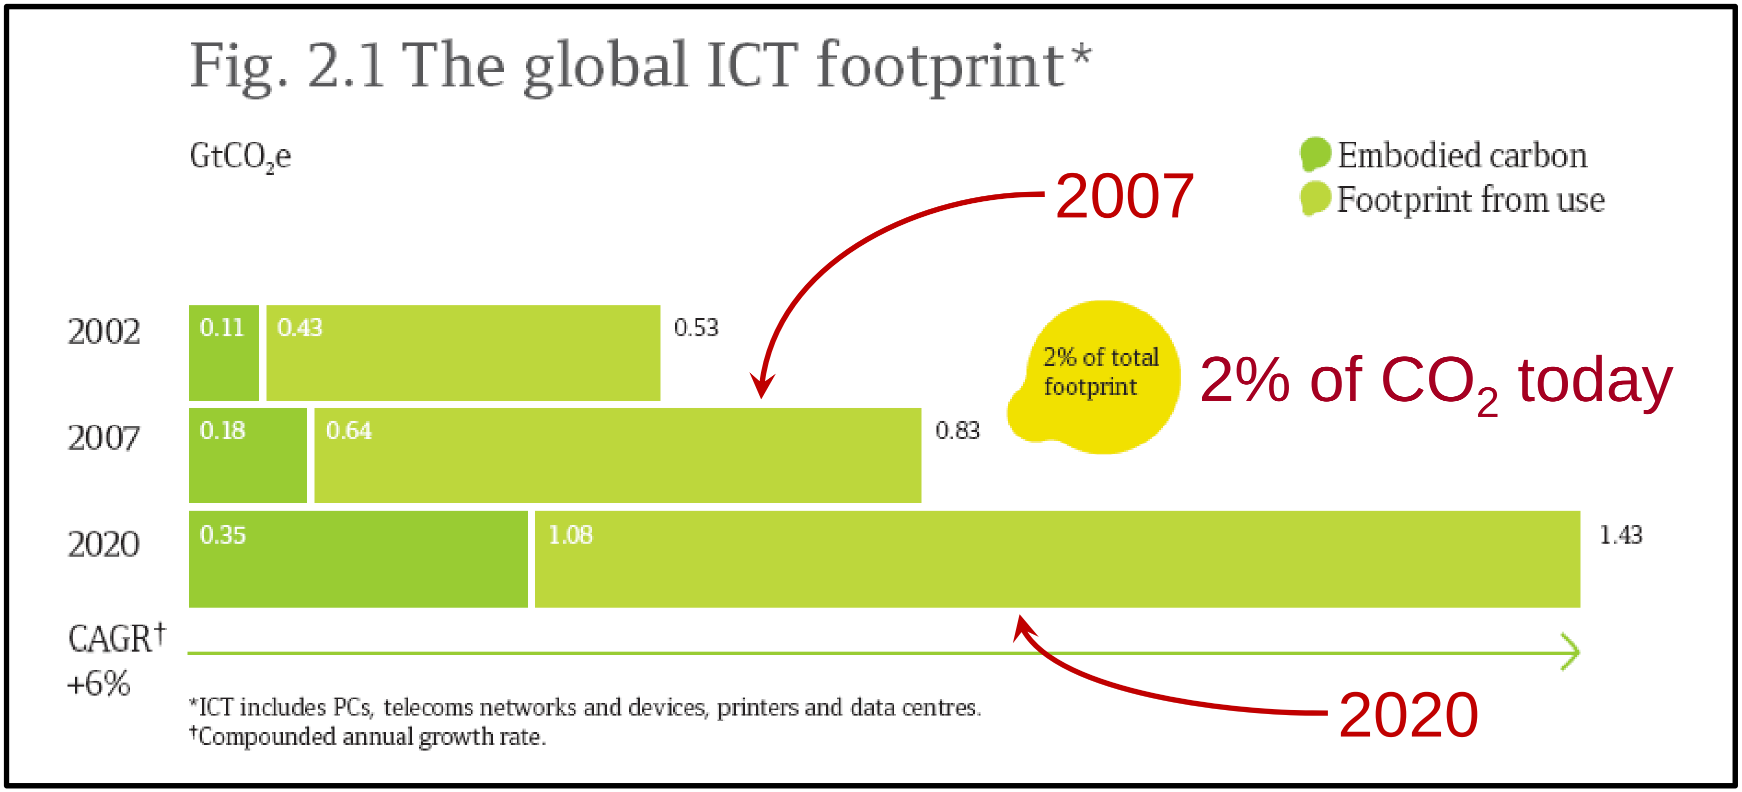

From SMART 2020 report

Global ICT CO_2 footprint <u>continued</u>

Telecom and PCs are the major contributors

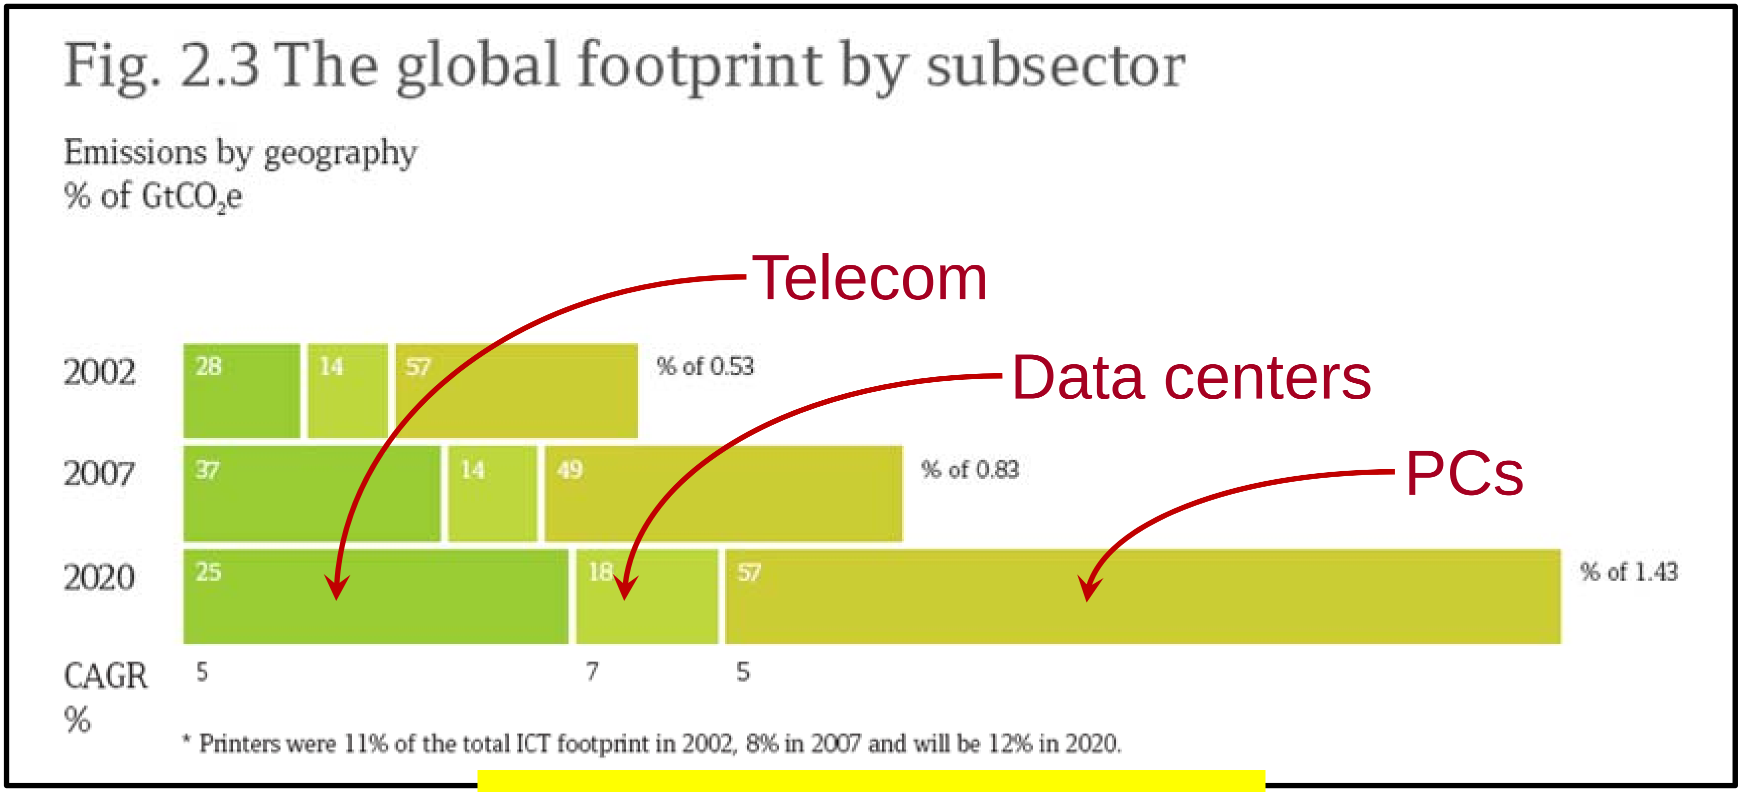

Data centers are less than 15%

From SMART 2020 report

ICT CO_2 > Aviation CO_2

A very significant statistic…

"The global information and communications technology (ICT) industry accounts for approximately 2 percent of global carbon dioxide (CO_2) emissions, a figure equivalent to aviation."

- Gartner Group, Inc. (2007)

Most energy use is from the end user

More significant statistics…

"Desktop computing accounts for 45 percent of global carbon emissions from information technology."
- govtech.com

"Most PC energy use in the US occurs when no one is there, and this is greater than the total energy use of all network equipment."
- Bruce Nordman (LBNL)

Statistics from Italy – broadband

17.5 million broadband users, overall population is 60 million

(A) 2015-2020 NETWORK FORECAST: DEVICE DENSITY AND ENERGY REQUIREMENTS IN THE BUSINESS-AS-USUAL CASE (BAU). EXAMPLE BASED ON THE ITALIAN NETWORK.

	power consumption [W]	*number of devices* [#]	*overall consumption* [GWh/year]
Home	10	17,500,000	1,533
Access	1,280	27,344	307
Metro/Transport	6,000	1,750	92
Core	10,000	175	15
		Overall network consumption	1,947

From: R. Bolla, R Bruschi, K. Christensen, F. Cucchietti, F. Davoli, and S. Singh, "The Potential Impact of Green Technologies in Next Generation Wireline Networks – Is There Room for Energy Savings Optimization?," to appear in *IEEE Communications*.

Statistics from the USA – big picture

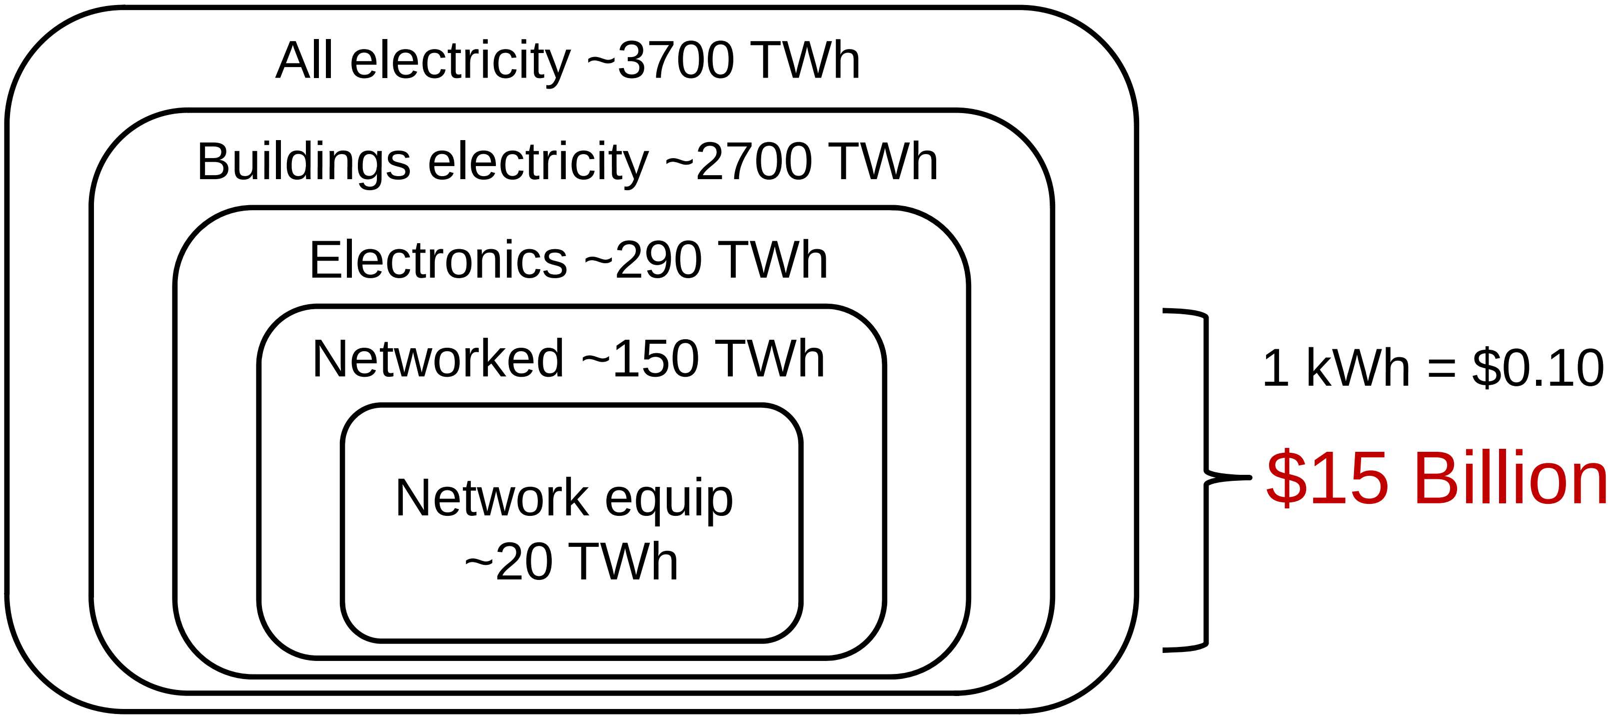

Where (exactly) is this electricity used?

How much of it is wasted?

How much can be saved?

From Bruce Nordman, LBNL, 2010.

How much is wasted?

Most energy consumed by networks is wasted

- Fact #1: Networks are generally lightly utilized
  - Over provisioned for peak and redundancy
  - 1% to 5% utilization typical at edges

- Fact #2: Network elements have high base power
  - Base power is power draw when idle
  - 80% base power is typical for PCs, routers, links, etc.

→ Significant potential for large energy savings

Notion of energy proportional computing

Relationship of power use and utilization

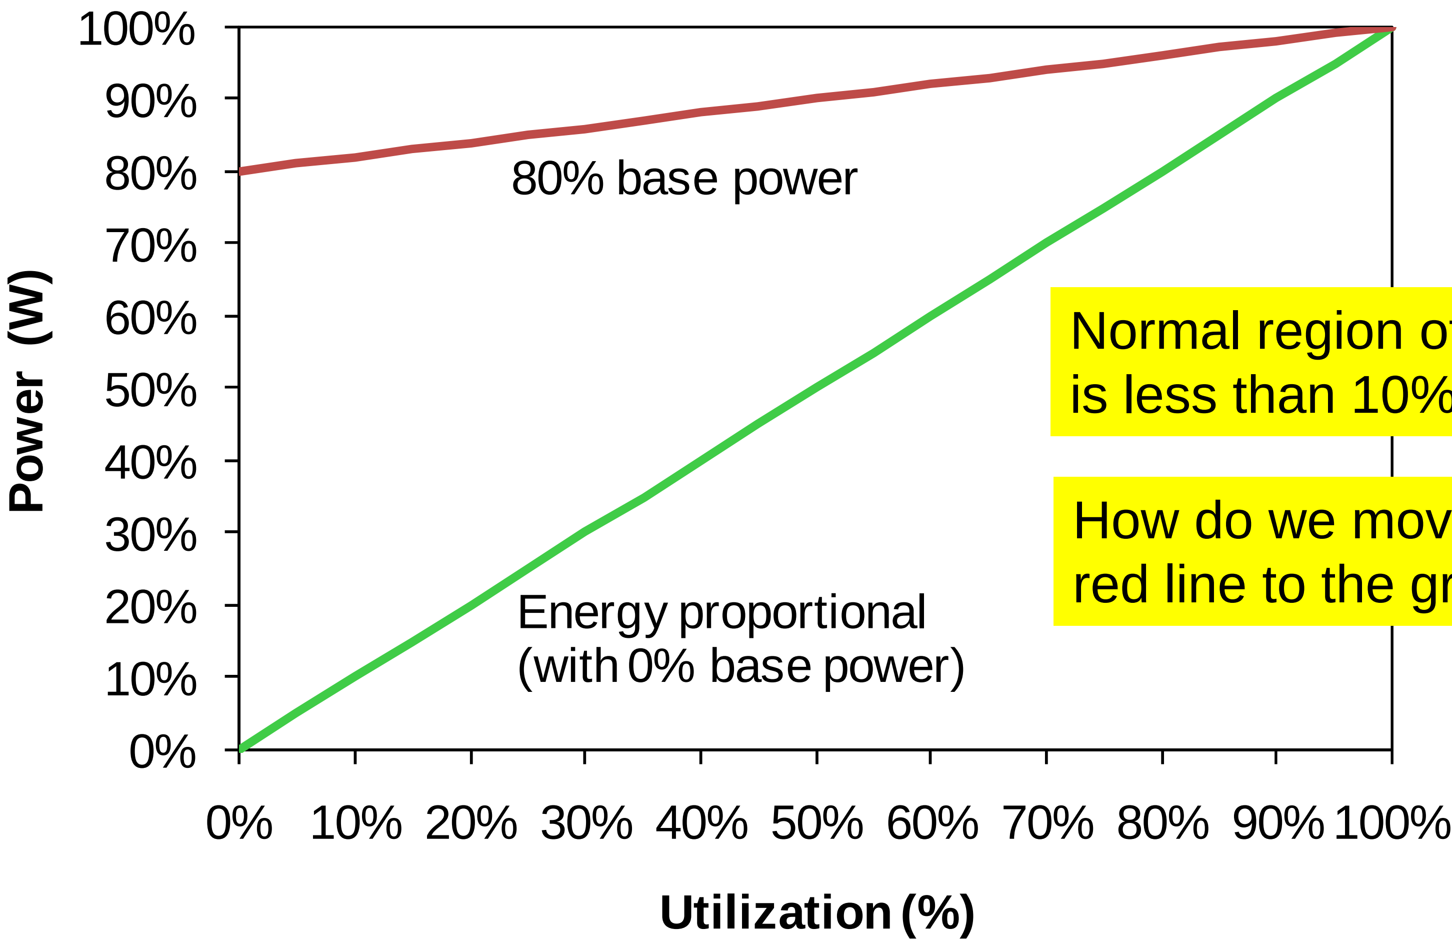

Normal region of operation is less than 10% utilization

How do we move from the red line to the green line?

Energy proportional applies everywhere

EL HERALDO.COM.CO | Locales

Barranq

Inicio Secciones Opinión Blogs Especiales Revistas Entretenimiento Semanarios Multimedia Chat Pelaos Edició

En consumo de energía, el estrato es lo de menos

Challenge is in "doing nothing well" (David Culler)

Es usual que la Universidad del Atlántico luzca así de iluminada, como ocurrió el viernes a las 8 y 30 de la noche. Foto Jairo Buitrago

Por Nadia Nájera Ricardo

Los usuarios de la energía en Barranquilla consumen alrededor de 220 millones de kilovatios mensuales, por lo que pagan cerca de 57 mil millones de pesos, teniendo en cuenta que el precio del kilovatio está en unos 260 pesos.

Summary of ICT energy use

ICT consumes and wastes a lot of energy

- ICT contributes about 2% of human emitted CO_2
  - About equal to aviation industry
  - Rapidly growing
- Most of this energy consumption comes from the edge
  - From edge networks, edge network equipment, and PCs
  - Not from data centers
- Most of the energy consumed is wasted
  - Provisioning for peak resulting in low average utilization
  - High base power

Next Steps

The next steps are…

- Our goal is energy proportional behavior for all of our systems
  - Both ICT and non-ICT

Reducing direct energy consumption

Can we reduce energy used by Ethernet?

... can Ethernet be made energy proportional?

Reducing energy use of Ethernet

Key observations:

- Most Ethernet links are lightly utilized (1% to 5%)
- Ethernet power consumption is independent of utilization
  - Not energy proportional

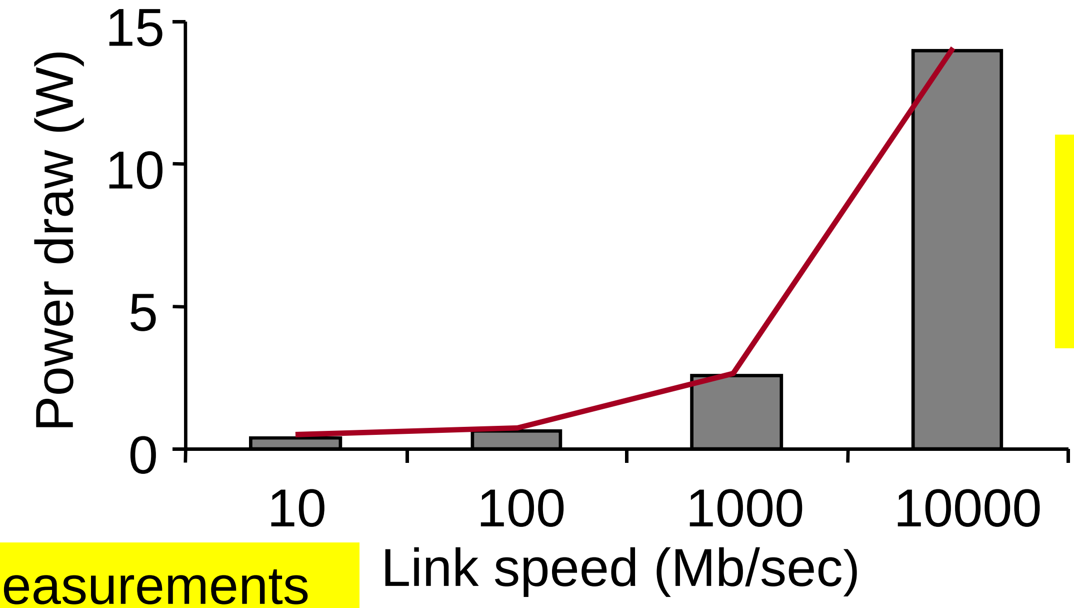

Power draw increases with link rate

Actual measurements

Two ideas to reduce energy use

Can we adapt power use to utilization?

- Idea #1: Adaptive Link Rate (ALR)

- Idea #2: Low Power Idle (LPI)

Idea #1: Adaptive Link Rate (ALR)

Proposed in 2005 by Nordman and Christensen

Adaptive link rate (ALR) continued

Independent of PC power management

Goal: Save energy by matching link data rate to utilization

- **Change (or adapt) data rate in response to utilization**
  - Use 10 or 100 Mb/sec during low utilization periods
  - Use 1 or 10 Gb/sec during high utilization periods
- **Need new *mechanism***
  - Current auto-negotiation is not suitable (too slow)
    - Designed for set-up (e.g., boot-up time), not routine use
- **Need *policies* for use of mechanism**
  - *Reactive policy* possible if can switch link rates “quickly”
  - *Predictive policy* is needed otherwise

72

The link switches between rates as a function of link utilization

From Bruce Nordman and Ken Christensen, “Reducing the Energy Consumption of Networked Devices,” IEEE 802.3 tutorial, July 19, 2005 (San Francisco).

Idea #2: Low Power Idle (LPI)

Proposed in 2007 by Intel

7

Active/Idle Toggling with 0BASE-x Concept

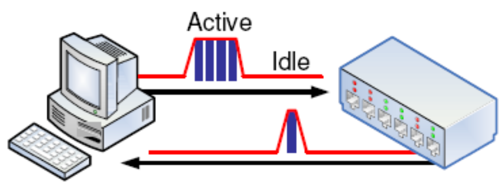

The link sleeps between packets

- Principle: Transmit data at fastest rate then return to idle
  - Energy savings come from power cycling between active/idle states
- Active/Idle toggling could be used *instead* of PHY rate shifting
  - Offers the best energy efficiency on links with lower utilization
  - Integrates well with existing PC power management schemes (e.g. ACPI)
  - Clock & power gating (on/off) is easier than rate shifting
- Asymmetrical operation would provide even better energy efficiency
  - Each direction could enter active & idle states independently
  - Most end-node traffic is heavily weighted toward either send or receive
  - Tx & Rx data paths already operate independently above the PHY

Energy Efficient Ethernet

intel

From Robert Hays, "Active/Idle Toggling with 0BASE-x for Energy Efficient Ethernet," presentation to IEEE 802.3az Task Force, November 2007.

ALR, LPI, and IEEE 802.3az

- Opportunity for energy savings to IEEE 802.3 in 2005
  - Presented need for energy savings and idea of ALR
  - A Study Group was formed
  - Mike Bennett from LBNL became the chair
- Became “Energy Efficient Ethernet”
  - Became IEEE 802.3az task force
- ALR became RPS, which then became LPI
- Standard based on LPI to be ratified in late 2010
- Vendors are now sampling products (based on LPI)
  - Broadcom and Realtek

How LPI works

PHY goes to sleep between packets

- Sleep is idle = about 10% of full power
  - Periodic refreshes to keep synchronized
  - Has wake-up and sleep transitions
    - » First packet after an idle incurs a wake-up transition (T_w)
    - » After last packet in a burst a go to sleep transition (T_s)

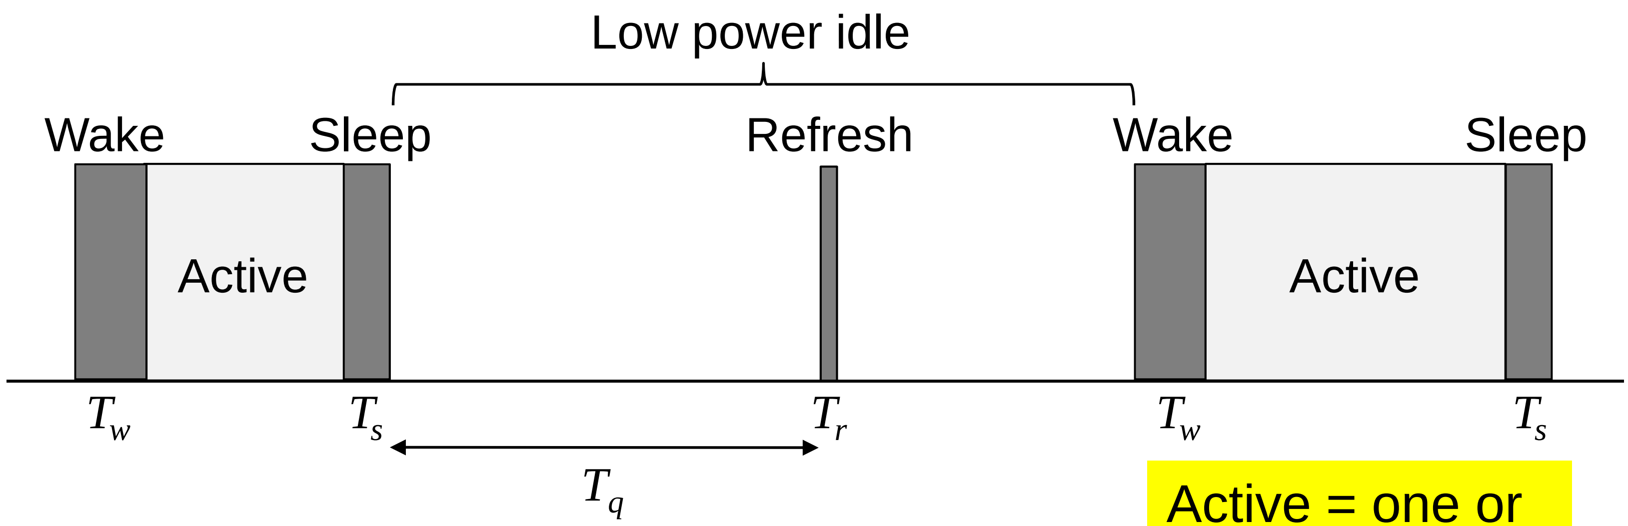

Active = one or more packets

Effect of LPI overhead

Efficiency for single packet case

$$Efficiency = \frac{T_{Frame}}{T_{Frame} + T_w + T_s}$$

Protocol	Min T_w (μs)	Min T_s (μs)	T_{Frame} (1500B) (μs)	Frame eff.	T_{Frame} (150B) (μs)	Frame eff.
100Base-Tx	30	100	120	48%	12	8.5%
1000Base-T	16	182	12	5.7%	1.2	0.6%
10GBase-T	4.16	2.88	1.2	14.6%	0.12	1.7%

From P. Reviriego, J. Hernandez, D. Larrabeiti, and J. Maestro, "Performance Evaluation of Energy Efficient Ethernet," *IEEE Communications Letters*, Vol. 13, No. 9, pp. 1-3, September 2009.

Effect of LPI overhead continued

Efficiency for 10 Gb/s (Poisson arrivals)

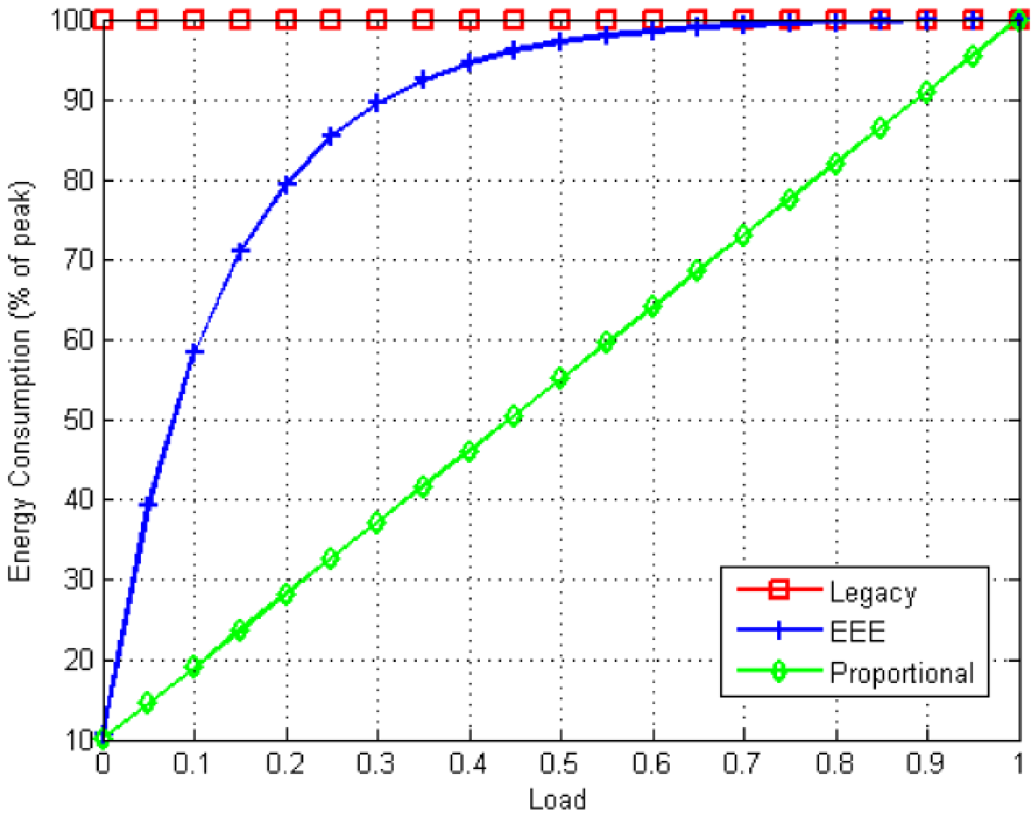

Simulation. Fixed length 1250 byte packets.

There is potential for improvement!

From P. Reviriego, J. Hernandez, D. Larrabeiti, and J. Maestro, "Performance Evaluation of Energy Efficient Ethernet," *IEEE Communications Letters*, Vol. 13, No. 9, pp. 1-3, September 2009.

Packet coalescing to fix inefficiency

Addressed EEE improvements in a recent work

IEEE 802.3az: The Road to Energy Efficient Ethernet

Authors: K. Christensen, P. Reviriego, B. Nordman, M. Bennett, M. Mostowfi, and J.A. Maestro

- Explored coalescing of packets at transmitter
  - Reduce overall wake and sleep overhead
  - Trade-off of energy savings and packet delay

- Paper to appear in special issue on Green Communications in *IEEE Communications* magazine

How packet coalescing works

Coalescing is the intentional bunching of packets

- First packet arrival to an empty transmit queue starts a *timer* and a *counter*
- When the timer expires or the counter reaches maximum, send all the packets
  - Timer set to $t_{coalese}$ (times down to 0)
  - Counter set to 0 (counts up to *max*)
- Already often implemented on receive side to reduce interrupt overhead

Evaluation of EEE with coalescing

Reproduced previous 10 Gb/s experiment

- EEE parameter values
  - T_WAKE = 4.16 μs
  - T_SLEEP = 2.88 μs
  - For 1250 byte packet service time = 1.0 μs
- Coalescing parameter values
  - max = 10 packets or $t_{coalese}$ = 12 μs
  - max = 100 packets or $t_{coalese}$ = 120 μs
- Assume that idle power use is 10% of full power use
- Vary offered load from 0% to 95%
  - Poisson arrivals, fixed length packet

EEE with coalescing results

Efficiency of 10 Gb/s with coalescing

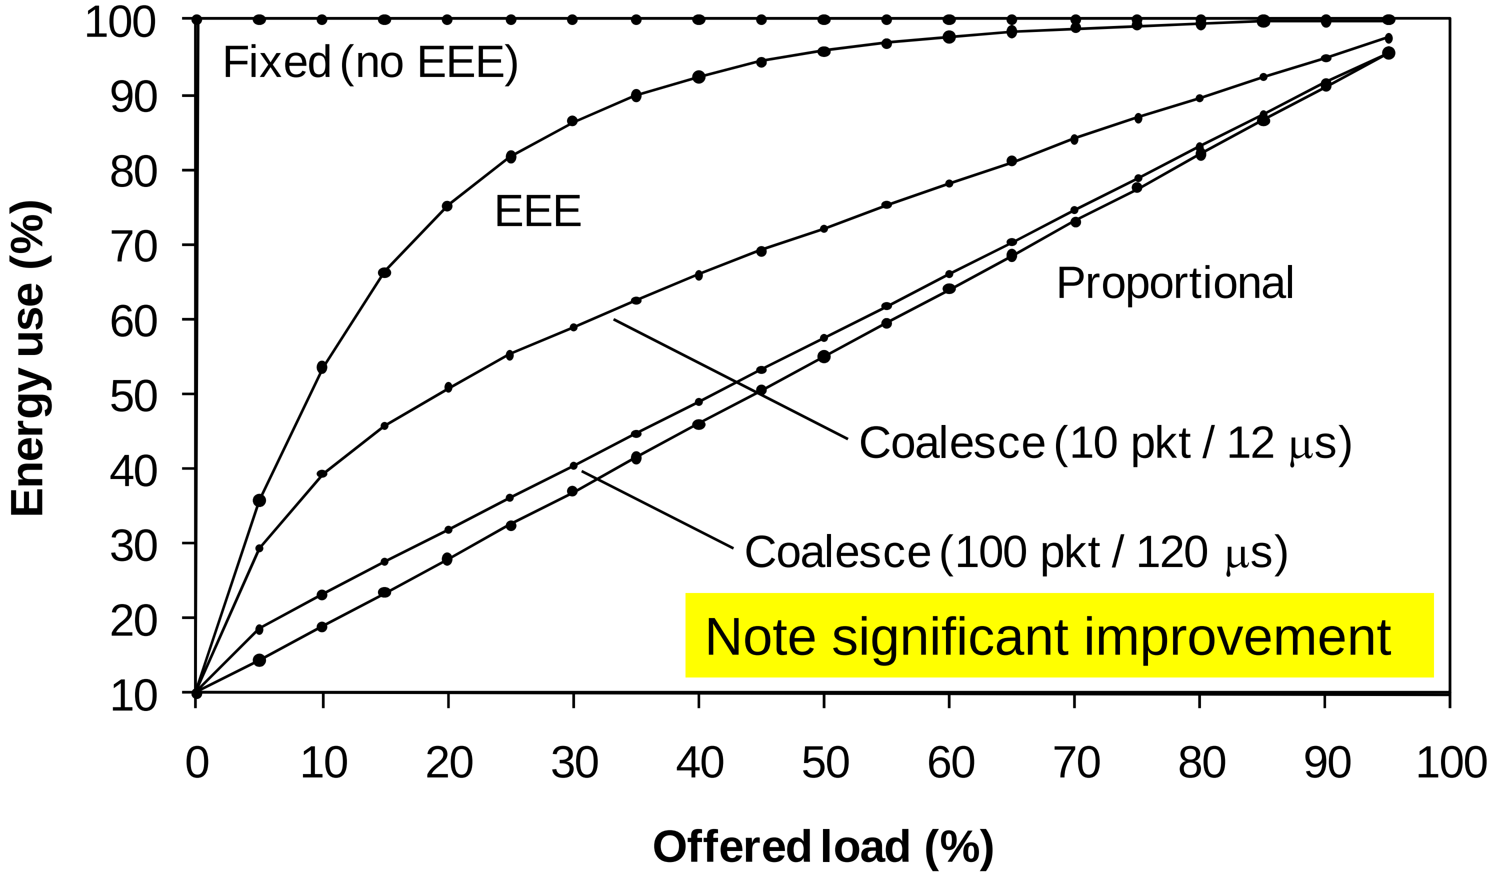

EEE with coalescing results continued

Packet delay for 10 Gb/s with coalescing

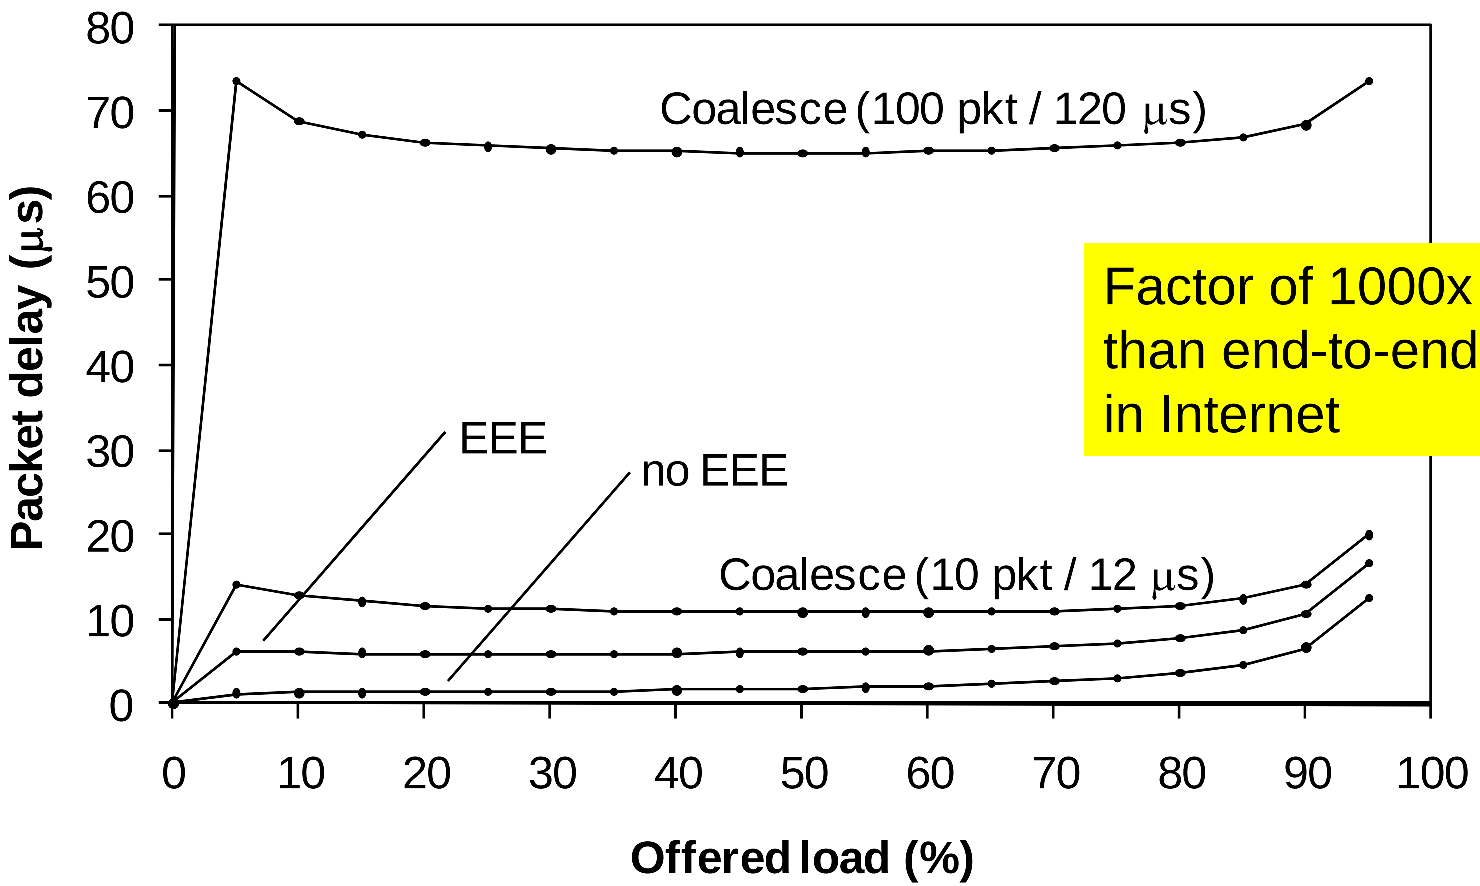

Factor of 1000x less than end-to-end delay in Internet

Expected savings

Energy savings have been estimated for USA

- Assume 2008 stock of Ethernet links as the "future"
  - Assume all interfaces support EEE
  - 250 million 1 Gb/s and 65 million 10 Gb/s
  - Per link savings of 1 W for 1 Gb/s and 5 W for 10 Gb/s
  - Get efficiency values from simulation graphs

➡ EEE savings per year in the USA = $410 million

➡ Additional savings from coalescing = $80 million

Summary of EEE

IEEE 802.3az improves energy efficiency of Ethernet

- Ethernet links typically have low utilization
- EEE = Energy Efficient Ethernet (based on LPI)
  - Will soon be a ratified standard
  - Vendors are starting to ship parts
- Packet coalescing can achieve energy proportionality
  - Added delay is in 10s of microseconds – probably not an issue for end-to-end delay in an Internet connection

Next Steps

The next steps are…

- Explore how other link technologies be made energy proportional
  - With sleeping or rate adaptation
- Explore possible side effects of coalescing
  - TCP ACK compression is one possible issue

Reducing induced energy consumption

Can we reduce energy used by hosts?

... a problem of network presence

Reducing energy use of network hosts

Key observation

- “Today, billions of dollars’ worth of electricity are used to keep Ethernet (and other) connected devices fully powered on at all times only for the purpose of maintaining this connectivity.” (Bruce Nordman, 2007)
- The need for network presence is driving PCs to be left fully powered-on at all times

Defining “network presence” is a key challenge

Network Connectivity Proxy

How can we maintain network presence?

- Two possible approaches
  1) Redesigning protocols and applications
  2) Encapsulating intelligence for maintaining network presence in an entity other than the core of the network devices

- Approach (2) best in the near-term

- A proxy is "an entity that maintains full network presence for a sleeping device"
  - Host appears to other devices as fully operational

High level view of a proxy

Operation of a proxy

1) Host awake; becomes idle
2) Host transfers state to proxy on going to sleep
3) Proxy responds to routine traffic for sleeping host
4) Proxy wakes up host as needed

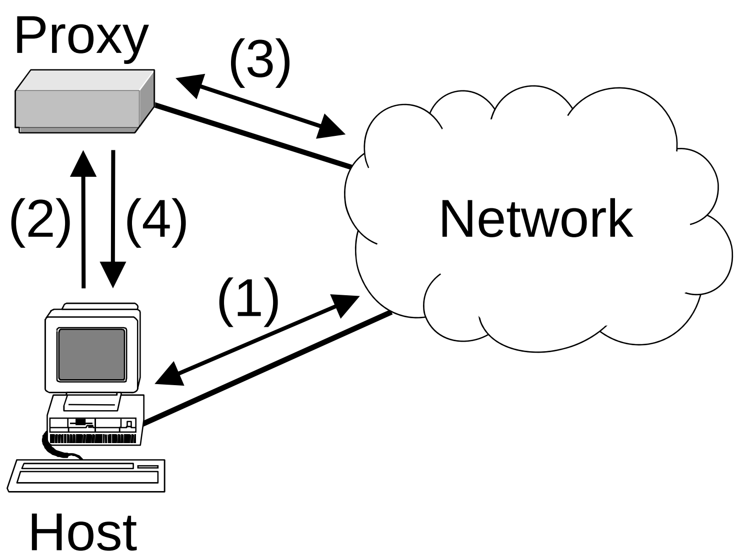

Proxy can be in separate entity, or within host NIC

The first work on proxying

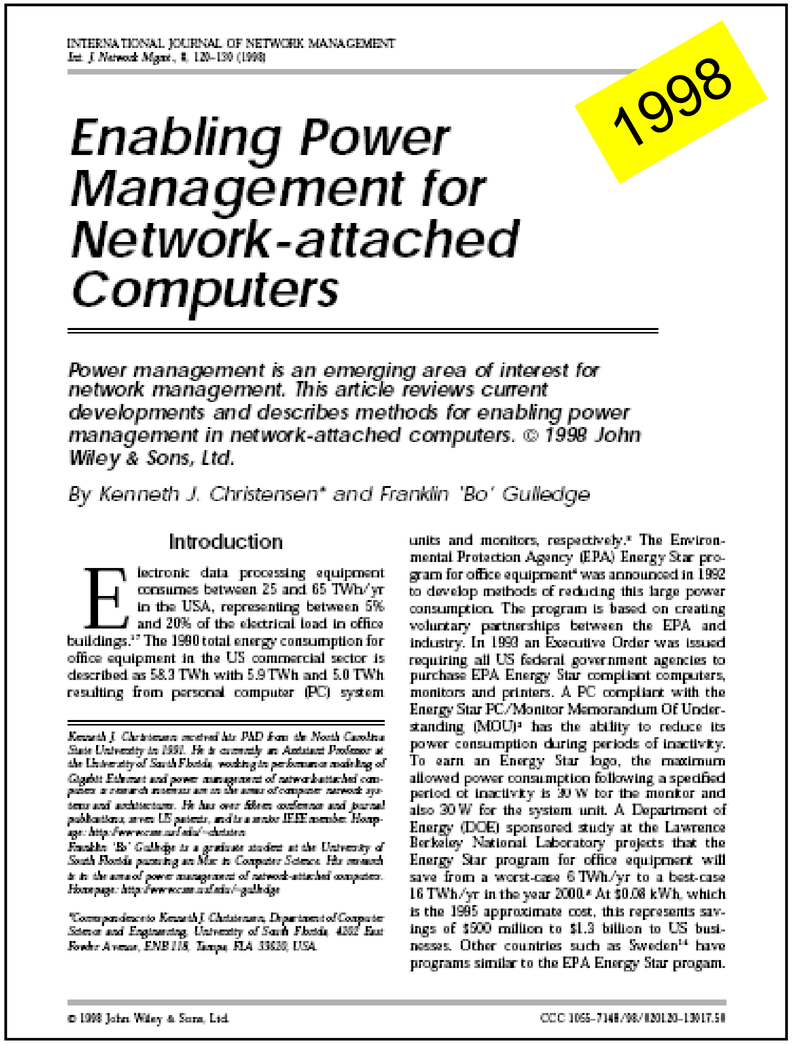
INTERNATIONAL JOURNAL OF NETWORK MANAGEMENT
Int. J. Network Mgmt., 8, 120–130 (1998)

1998

Enabling Power Management for Network-attached Computers

Power management is an emerging area of interest for network management. This article reviews current developments and describes methods for enabling power management in network-attached computers. © 1998 John Wiley & Sons, Ltd.

By Kenneth J. Christensen and Franklin 'Bo' Gulledge*

Introduction

Electronic data processing equipment consumes between 25 and 65 TWh/yr in the USA, representing between 5% and 20% of the electrical load in office buildings.[1-7] The 1990 total energy consumption for office equipment in the US commercial sector is described as 58.3 TWh with 5.9 TWh and 5.0 TWh resulting from personal computer (PC) system units and monitors, respectively.[8] The Environmental Protection Agency (EPA) Energy Star program for office equipment[4] was announced in 1992 to develop methods of reducing this large power consumption. The program is based on creating voluntary partnerships between the EPA and industry. In 1993 an Executive Order was issued requiring all US federal government agencies to purchase EPA Energy Star compliant computers, monitors and printers. A PC compliant with the Energy Star PC/Monitor Memorandum Of Understanding (MOU)[3] has the ability to reduce its power consumption during periods of inactivity. To earn an Energy Star logo, the maximum allowed power consumption following a specified period of inactivity is 30 W for the monitor and also 30 W for the system unit. A Department of Energy (DOE) sponsored study at the Lawrence Berkeley National Laboratory projects that the Energy Star program for office equipment will save from a worst-case 6 TWh/yr to a best-case 16 TWh/yr in the year 2000.[8] At $0.08 kWh, which is the 1995 approximate cost, this represents savings of $500 million to $1.3 billion to US businesses. Other countries such as Sweden[14] have programs similar to the EPA Energy Star progam.

Kenneth J. Christensen received his PhD from the North Carolina State University in 1991. He is currently an Assistant Professor at the University of South Florida, working in performance modeling of Gigabit Ethernet and power management of network-attached computers. His research interests are in the areas of computer network systems and architectures. He has over fifteen conference and journal publications, seven US patents, and is a senior IEEE member. Homepage: http://www.csee.usf.edu/~christen

Franklin 'Bo' Gulledge is a graduate student at the University of South Florida pursuing an MSc in Computer Science. His research is in the area of power management of network-attached computers. Homepage: http://www.csee.usf.edu/~gulledge

**Correspondence to Kenneth J. Christensen, Department of Computer Science and Engineering, University of South Florida, 4202 East Fowler Avenue, ENB 118, Tampa, FLA 33620, USA.*

CCC 1055–7148/98/020120–13\$17.50

- Described proxying for ARP and TCP keep-alives
- Described a centralized proxy covering for many hosts on a shared Ethernet LAN

Early work: A prototype ARP/SYN proxy

Emulated proxy to allow a web server to sleep

From K. Christensen, P. Gunaratne, B. Nordman, and A. George, "The Next Frontier for Communications Networks: Power Management," *Computer Communications*, Vol. 27, No. 18, pp. 1758-1770, December 2004.

Recent work: A proxy for SIP phones

IP phones are a new energy consumer

- IP phones need to maintain network presence
  - In order to receive a “ring” signal on incoming call
- IP phone draws about 10 to 20 W (so, $10 to $20 per year)
- Can also use a PC to make a “soft phone”
  - PC then needs to remain powered-up at all times

Example: The Magic Jack product

A new product to replace landline telephone service

- USB device to plug an analog phone into a PC
  - Then use a SIP-based IP telephony service
  - Uses your Broadband service “for free”

Millions sold

Requires PC to be fully powered-on to be able to participate in SIP protocol

Power costs can exceed savings from canceling landline service

USF UNIVERSITY OF SOUTH FLORIDA

The "SIP catcher" – system view

Developed a proxy within a Linksys router

- Knows sleep/wake state of a soft phone PC or IP phone
- Handles SIP protocol and wakes IP phone as needed

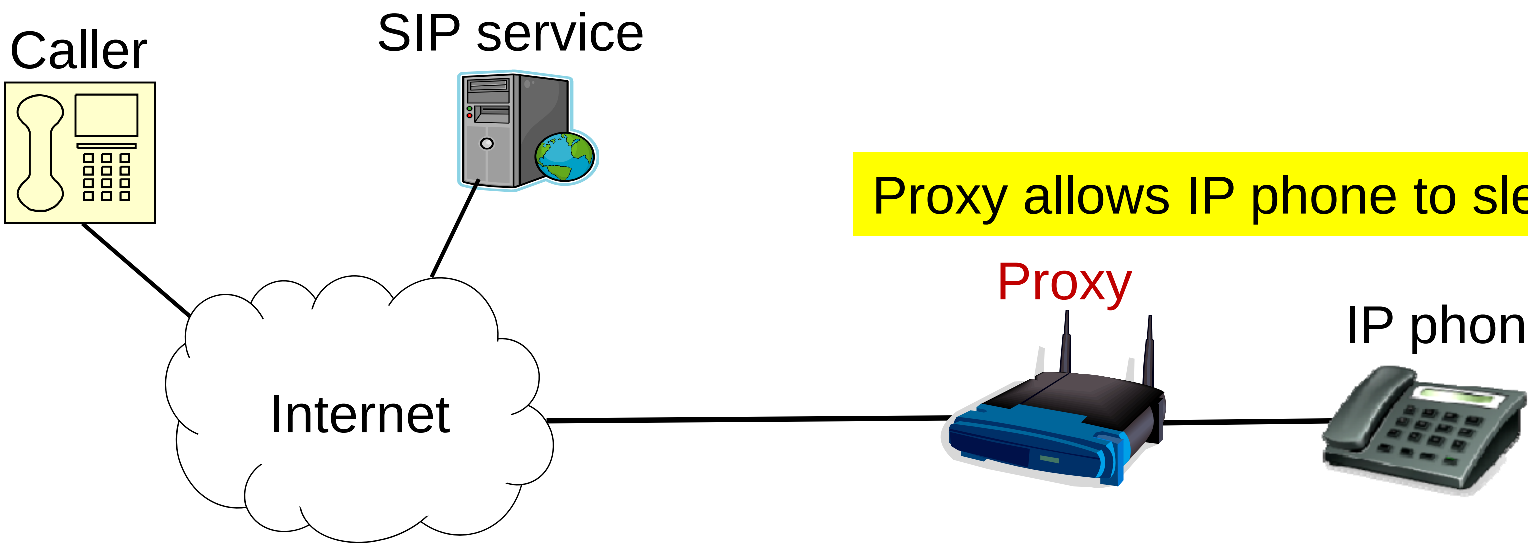

The “SIP catcher” – packet flow view

Key steps:
1) Wakes up phone when call detected (incoming INVITE)
2) Responds on behalf of phone (TRYING)
3) Forwards INVITE to phone when it is awake

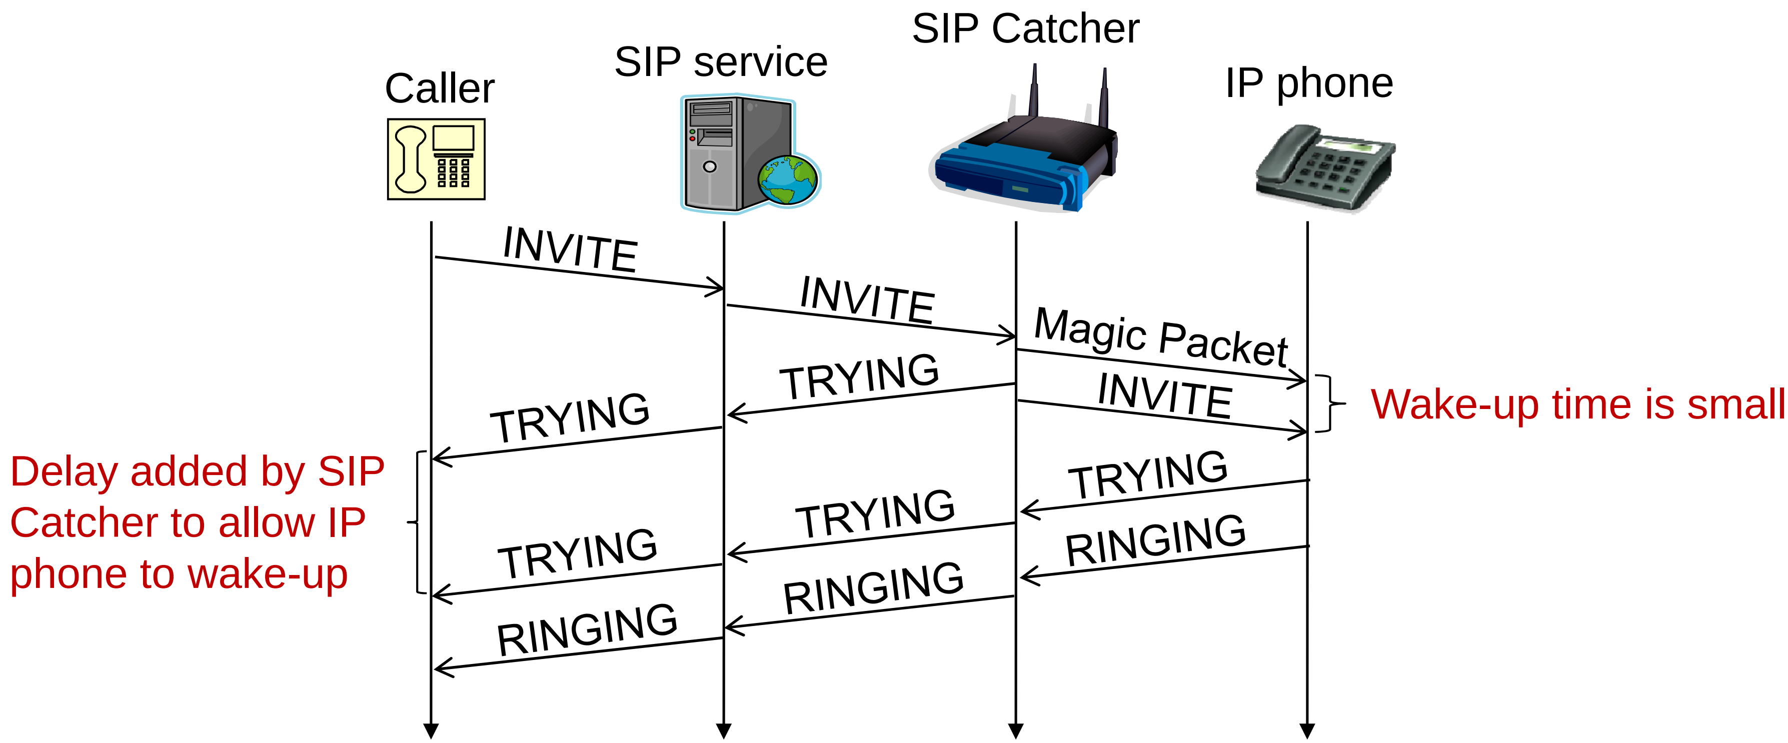

Full 8 minute version on YouTube (featuring Miguel Jimeno)

The "SIP catcher" – a demonstration

Ecma proxying standard

Ecma-393 ProxZzzy for sleeping hosts

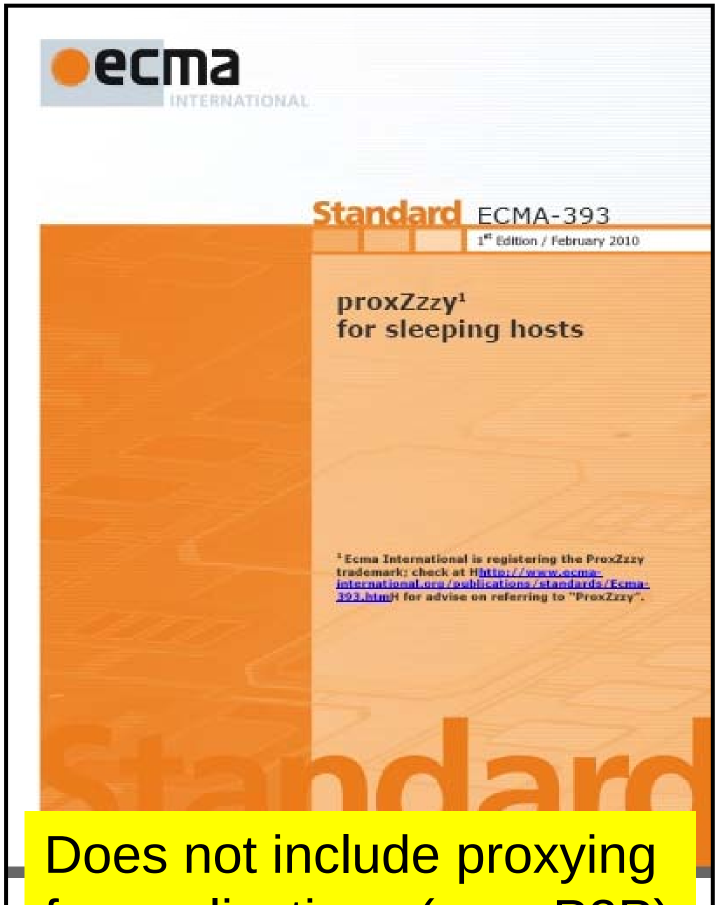

Does not include proxying for applications (e.g., P2P)

- "… maintenance of network connectivity and presence by proxies to extend the sleep duration of hosts"
- Addresses low layers
- Satisfies EPA Energy Star "platform-independent industry standard"
- Approved in February 2010

Proxying in EPA Energy Star

EPA Energy Star for Computers, Version 5.0

- "Proxying refers to a computer that maintains Full Network Connectivity as defined in Section 1 of this specification. For a system to qualify under the proxying weightings above, it must meet a non-proprietary proxying standard that has been approved by the EPA and the European Union as meeting the goals of ENERGY STAR."*

The Ecma standard is key to this

From ENERGY STAR® Program Requirements for Computers, Version 5.0, EPA, 2009.

Proxying in products

Apple Snow Leopard

- "Wake on Demand. This is Apple's name for a new networking feature that lets a Snow Leopard Mac go to sleep while a networked base station continues to broadcast Bonjour messages about the services the sleeping computer offers."*

Bonjour Sleep Proxy, supports ARP, file and print serving, and SSH login initiation.

From "Wake on Demand lets Snow Leopard Sleep with One Eye Open," MacWorld, August 28, 2009

Expected savings

Energy savings have been estimated for USA

- For desktop PCs most time is spent as on and idle
- Proxying could save more than half of energy used by these PCs and PC-like products

Savings potential for desktop PCs = $0.8 to $2.7 billion

Summary of Proxying

Proxying reduces induced energy use of hosts

- Hosts usually idle but fully powered on to maintain “presence”
- Idea of a network connectivity proxy
  - Based on low-power hardware covering for high-power hardware
  - Estimated savings are on the order of $1 billion per year in the US
- Proxying for lower layers is now real
  - Ecma standard, EPA ES statement, and Apple products
- Proxying can reduce energy costs of deploying IP phones
  - With the “SIP catcher”

Next Steps

The next steps are…

- Generalize proxying to the notion of selective connectivity
  - Explore architectural implications to Internet

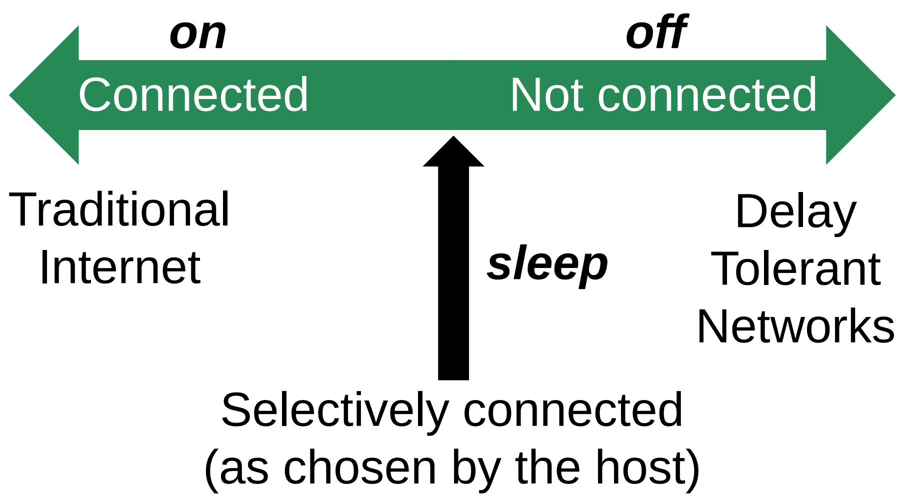

From M. Allman , K. Christensen, B. Nordman, and V. Paxson, "Enabling an Energy-Efficient Future Internet through Selectively Connected End Systems,“ *Sixth Workshop on Hot Topics in Networks (HotNets-VI)*, November 2007.

Future challenges

Where do we go from here?

… energy savings of and by ICT

Future challenges in green networks

Future challenges in four areas

1) General

2) Network core and edge

3) Network hosts

4) Distributed applications

Future challenges <u>continued</u>

General

- Metrics
  - How do we measure energy-performance trade-offs?
- Models
  - How do we model energy-performance trade-offs?
- Exposing power and usage state of equipment
  - Need to be able to remotely determine power/use state
  - How to know when something is idle?
- Architectures for selective connectivity
  - Need mechanisms/protocols for selective connectivity
  - Includes notions of proxying

Future challenges <u>continued</u>

Network core and edge

- Energy efficient routers and switches
  - Support sleep states and rate adaptation
- Energy efficient links
  - Adapt link rates to load
- Traffic shaping
  - Shape traffic to enable short-term sleeping during idle periods
- Traffic engineering
  - Consolidate routes for long-term sleeping of idle routes
- Data caching
  - Cache popular data to reduce load on network and servers

Future challenges <u>continued</u>

Network hosts

- Discovery of devices, capabilities, content, and services
  - Need to be able to discover low-power substitutes

Distributed applications

- Move computing work to where power is cheapest
  - “Follow the moon” for data center activity
- P2P, multiplayer games, and virtual worlds
  - When idle should sleep
- Webcams and sensors everywhere (Internet of things)
  - When idle should sleep

Where are the “best” challenges?

My thoughts…

- The biggest challenges are at the edge
  - Most energy use and most opportunity for making changes
- Be careful to not work on problems already solved
  - Much has now been solved (the “low hanging fruit”)
  - Always be able to quantify expected savings and argue that they are sufficient to be of interest
- The really biggest challenges may be in the “other 98%”
  - Many open networks problems for Smart Buildings

Conclusions

- ICT has large and growing energy use
  - Estimated to be 2% of human generated CO_2
- EEE will reduce direct energy use
  - Hundreds of millions of dollars per year in US expected
- Packet coalescing will improve efficiency of EEE
  - Tens of millions of dollars per year saving in US possible
- Proxying will reduce induced energy use by hosts
  - Potential for billions of dollars per year savings in the US
- There are future challenges to be addressed

Any questions?

Ken Christensen

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