

Some slides are from a presentation given at Cisco by Ken Christensen.

An Energy Efficient Internet: Some Ongoing Work

Francisco Blanquicet

Department of Computer Science and Engineering

University of South Florida

Tampa, FL 33620

fblanqui@cse.usf.edu

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

Thank you

- **Thank you for inviting me to this brown bag lunch**
- **I am a graduate student at the University of South Florida**
- **My advisor is Ken Christensen**
 - <http://www.csee.usf.edu/~christen/christen.html>
- **We work extensively with Bruce Nordman at LBNL**

Our project


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

The Energy Efficient Internet Project - Mozilla Firefox

File Edit View History Bookmarks Tools Help

The Energy Efficient Inter...

USF UNIVERSITY OF SOUTH FLORIDA

 UNIVERSITY OF FLORIDA

The Energy Efficient Internet Project


- [Project description](#)
- [People](#)
- [Publications and talks](#)
- [Energy use](#)
- [Maintain](#)
- [has](#)
- [ity for](#)


This project addresses the increasingly critical need to improve the energy efficiency of the Internet by focusing on the primary and often neglected energy consumer, edge devices. Unfortunately, due to limits of existing protocols and architectures, networked desktop computers typically remain powered-up during frequent and often lengthy periods of idleness. As network devices, they are prevented from operating in an energy-efficient manner due to their need to respond to network transactions of various types without warning. In

Focus is on edge devices, not on Internet core or data centers. The edge (desktop PCs, etc.) consumes more energy than data centers.

sleeping PCs.

- The [UPnP Forum](#) released standardized service descriptions for [Low Power V1.0](#) in August 2007. This includes a UPnP proxy to which we made contributions (see [here](#)). Jakob Klama (student) is specifically listed as a contributor on the UPnP Forum standard.
- The EPA [ENERGY STAR Program Requirements for Computers: Version 4.0, Draft](#) now states that "Computers shall reduce the speed of any active 1 Gb/s Ethernet network links when transitioning to Sleep or Standby." (page 11). This follows directly from our Ethernet Adaptive Link Rate (ALR) work described below.
- An IEEE 802.3 [Energy Efficient Ethernet Study Group](#) was established in November 2006 and is now the [IEEE 802.3az task](#)





Where this short talk is going

- **Energy usage by IT equipment**
- **Reducing induced energy use**
- **Reducing direct energy use**
- **Additional efforts at USF**

Electricity production and costs

- **Let's set the stage...**


- 1 kWh \approx \$0.10 (average US residential cost for electricity)
- 1 Wyr = \$0.88 \approx \$1.00
- 1 TWh = \$100 million
- 1 TWh = 0.75 million tons of CO₂

Crystal River nuclear plant (about 7 TWh/yr)



Electricity use in the USA

- **2006 US electricity usage*** (not to scale)



* B. Nordman, "Networks, Energy, and Energy Efficiency," presentation at *Cisco Green Research Symposium*, March 2008.

Electricity use by IT equipment today

- **How much electricity do PCs consume?**
 - EPA estimates about 2% of all electricity consumed*
- **How much electricity do data centers consume?**
 - About 1.5% of all electricity consumed in US**
 - Doubled since 2000
- **How much electricity does the Internet consume?**
 - “The Internet accounts for 5% of all the power we consume – in a couple of years, that figure will be 10%.” ***

This might be too high?

* “EPA Announces New Computer Efficiency Requirements,” Release date: 10/23/2006, Contact: Enesta Jones.

** Jon Koomey quoted in <http://gigaom.com/2008/06/25/structure-08-data-center-power-guru-jonathan-koomey/>

*** Institute for Energy Efficiency, UC Santa Barbara, 2008.

PC energy impact


- **The typical US home consumes 10,654 kWh/yr***
- **One 100W PC on 24/7 for one year is...**
 - 876 kWh/yr
 - This is 8.2% of entire home consumption



Each PC adds about 8% to your power bill!

* Energy Information Administration, "U.S Household Electricity Report," July 2005.

Typical commercial PC energy use



$$P_{on} \gg P_{sleep}$$
$$P_{sleep} \approx P_{off}$$

Consumption is driven by on time, not by usage

From http://www.energystar.gov/ia/partners/prod_development/revisions/downloads/computer/TierII_Network_Issue_Slides.pdf

PC energy use is “induced”

- **Why is the typical PC fully powered on when idle?**
 - For usability reasons?
 - For network connectivity reasons?

 This is induced energy use

Where this talk is going


- Energy consumption by IT equipment
- Reducing induced energy use
- Reducing direct energy use
- Additional efforts at USF

Reducing energy use of edge devices

- **Observation**: Devices are often left fully powered-on to maintain network connectivity or presence
 - Usually devices are not in active use
- **Idea**: A low-power network connectivity proxy
 - Maintains *full network connectivity* for a sleeping device
 - Enable devices to sleep more often
 - Not related to existing Wake-on-LAN
- **Key issues**: Lots of issues...
 - Definition of “connectivity”
 - Wake-ups (not too many, not too few)
 - Packet loss (during wake-up)

Network connectivity proxy

- **Proxy covers for sleeping device**
 - Proxy could be in local NIC or in the network



Steps:

- 1) PC determines it is time to sleep
- 2) PC state transferred to proxy
- 3) PC sleeps, proxy maintains presence
- 4) Proxy determines need to wake-up PC
- 5) PC awakes and proxy state transferred

Proxying for reducing energy use

- Started to think about proxying 10 years ago
 - Now hoping to define a direction towards a standard

1998

Enabling Power Management for Network-attached Computers

Power management is an emerging area of interest for network management. This article reviews current

Proxying of ARPs to maintain reachability.

Energy consumption in the USA, representing between 5% and 20% of the electrical load in office buildings.¹² 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

Keywords: 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 architecture. He has over 50 conference and journal publications, over 15 patents, and is a member IEEE member. Homepage: <http://www.com.usf.edu/~christen>. E-mail: "B" Gullidge is a graduate student at the University of South Florida pursuing an M.S. in Computer Science. His research is in the area of power management of network-attached computers. Homepage: <http://www.com.usf.edu/~gullidge>.

¹²Correspondence to: Kenneth J. Christensen, Department of Computer Science and Engineering, University of South Florida, 4302 East Fowler Avenue, ENP 118, Tampa, FLA 33620, USA.

© 1998 John Wiley & Sons, Ltd

CCC 1055-7148/98/020120-13\$17.50

2007



ethernet alliance

Improving the Energy Efficiency of Ethernet-Connected Devices: A Proposal for Proxying

Version 1.0 September 2007

Authors:

Bruce Nordman, Lawrence Berkeley National Laboratory

Ken Christensen, University of South Florida

ethernet alliance | p.o. box 200757 | austin, tx | 78720-0757 | usa
www.ethernetalliance.org

Key protocols that proxying might cover

- **Layers 1 and 2**
 - Already covered by PHY/MAC
- **Layer 3**
 - ARP, ND, ICMP, IGMP, IPSec, etc.
- **Layer 4**
 - TCP connection request (SYN)
 - TCP connections (keep-alives)
- **Higher layers including applications**
 - NetBIOS, SMB, DHCP, SNMP, SSDP, VPN, SSH, etc.
 - Application semantics

Network connectivity from the EPA

- **Future EPA Energy Star Program Requirements**
 - Version 5.0, Draft 1 Comment Summary* (for computers)

“Full Network Connectivity: The ability of the computer to maintain network presence while in sleep and intelligently wake when further processing is required. Maintaining network presence may include obtaining and/or defending an assigned interface or network address, responding to requests from other nodes on the network, or sending periodic network presence messages to the network all while in the sleep state. In this fashion, presence of the computer, its network services and applications is maintained even though the computer is in sleep. (Note: More information on this can be found at: <http://efficientnetworks.lbl.gov/enet-proxying.html>)

* From http://www.energystar.gov/ia/partners/prod_development/revisions/downloads/computer/Defs_Summary_Computer_Comment_Resp.pdf

Proxying next steps...

- **Progress is being made with**
 - Definition
 - Trace analysis
 - Prototyping
 - Use case
- **Involvement with several companies to do above**
 - Work being done is mostly not by us
- **Plan to approach to Ecma to host/develop standard**
 - Bruce Nordman (LBNL) driving this

Where this talk is going


- Energy consumption by IT equipment
- Reducing induced energy use
- Reducing direct energy use
- Additional efforts at USF

Reducing energy use of links

- **Observation #1**: Most edge links are lightly utilized
 - 1% to 5% on average
- **Observation #2**: Higher rates consume more power
 - About 2 to 4 W for 1 Gb/s versus 100 Mb/s
 - Much more for 10 Gb/s versus 1 Gb/s
- **Idea**: Match link data rate with utilization
- **Key issue**: Time to change between data rates
 - Can buffer overflow occur during transition?
 - What impact might this packet loss have?

Ethernet Adaptive Link Rate (ALR)

- **Two parts to the problem**
 - 1) *Mechanism* for how to switch link rate
 - 2) *Policy* for when to change link rate



Save energy in MAC/PHY at each end
(potential for deeper savings)

Ethernet ALR publications

- We published the idea and some results
 - From a 2005 paper and a 2006 whitepaper

2005

INTERNATIONAL JOURNAL OF NETWORK MANAGEMENT
 Int. J. Network Mgmt 2005, 18: 297-310
 Published online in Wiley InterScience (www.interscience.wiley.com). DOI: 10.1002/nem.565

Managing energy consumption costs in desktop PCs and LAN switches with proxying, split TCP connections, and scaling of link speed

By Chamara Gunaratne, Ken Christensen** and Bruce Nordman

The IT equipment comprising the Internet in the USA uses about \$6 billion of electricity every year. Much of this electricity use is wasted on idle, but fully powered-up, desktop PCs and network links. We show here to recover a large portion of the wasted electricity with improved power management methods that are focused on network issues. Copyright © 2005 John Wiley & Sons, Ltd.

1. Introduction

A growing expense and impact of the Internet is its energy use. Current estimates are that 2% of electricity consumption in the USA goes to powering the Internet.¹ In Germany it is estimated that energy consumption by IT equipment will be between 2% and 5% in 2010.² The 2% estimate for the USA totals more than 74 TWh/year or \$6 billion per year. It is predicted that energy use of IT equipment is growing faster than energy use of any other type within buildings.³ Much of this energy use is wasted. Energy use by IT equipment is not proportional to utilization of the equipment. A recent study by Lawrence Berkeley National Laboratory (LBNL) showed that 60% of all desktop PCs in commercial buildings remain fully powered-on during nights and weekends⁴ with existing power management almost always disabled. Beyond the PC are the Ethernet link and workgroup switch. At present, these energy consumers have almost no

means of power management. Existing Internet protocols including discovery and routing are also 'energy unaware'; future protocols need to be made energy aware. For existing protocols that cannot be changed, methods of accommodating current operation must be developed. In previous work we have shown that there exists the potential for savings of billions of dollars per year in the USA alone.⁵⁻⁷ These savings are summarized in Section 6 of this paper. Energy costs are a part of the total cost of ownership of an IT operation. Savings in these costs are of interest to IT managers and companies are beginning to respond with network management products (such as Vordium with its centralized power management controller⁸) to address this need.


An efficient device consumes energy proportional to its output or utility. Thus, an idle or lightly utilized PC or Ethernet link should not consume the same energy as one that is highly utilized. In this paper, we develop several new methods to reduce energy consumption of PCs,

Chamara Gunaratne is a graduate student in the Department of Computer Science and Engineering at the University of South Florida.
 Ken Christensen is an Associate Professor in the Department of Computer Science and Engineering at the University of South Florida.
 Bruce Nordman is a Principal Research Associate in the Energy Analysis Department, Environmental Energy Technologies Division, of Lawrence Berkeley National Laboratory Berkeley, California.

*Correspondence to: Ken Christensen, Department of Computer Science and Engineering, University of South Florida, Tampa, FL 33620, USA
 E-mail: chrisk@usf.edu

Copyright © 2005 John Wiley & Sons, Ltd.

2006



ethernet alliance

Improving the Energy Efficiency of Ethernet: Adaptive Link Rate Proposal

Version 1.0, July 15, 2006

Authors: Mike Bennett
 Lawrence Berkeley National Laboratory
 Ken Christensen
 University of South Florida
 Bruce Nordman
 Lawrence Berkeley National Laboratory

ethernet alliance | p.o. box 200757 | austin, tx | 78720-0757 | usa
 www.ethernetalliance.org

Where is EEE now...

- **IEEE 802.3az Task Force**
 - Mike Bennett from LBNL is chair
- **Moving forward to becoming a standard**
 - 2009 timeline for completion
 - Is in EPA Energy Star Version 5.0 Draft 1 for computers
- **Current proposal is for “Active-Idle” approach**
- **Active-Idle idea is from Intel**
 - Use a low-power idle between packets
 - Switch to high data rate when a packet is queued
 - About *10 microseconds* to transition out of low-power idle

Some press and a logo...



Energy*
Efficient
Ethernet

The collage features several news articles related to energy-efficient Ethernet. Key headlines include: 'The Net's Going Green: Multipronged Approach Might Save Costs, Energy — and the Climate', 'Researchers seek energy efficient nets, devices', 'Inefficient Ethernet wastes over \$1bn a year', 'IEEE Works on Energy-Efficient Ethernet', and 'IEEE Seeks For Ethernet To Go Green'. The articles discuss the IEEE's efforts to reduce energy consumption in network infrastructure, highlighting the potential for significant cost savings through more efficient Ethernet technology.


* Logo by Glen Kramer of Teknovus, Inc. (full IEEE permission for use granted via email dated January 27, 2007)

What is next?

- **EEE requires both ends to participate**
 - Full EEE deployment is still many years off
- **Can we do something simpler?**
 - Something backwards compatible?
- **Idea: PAUSE Power Cycle (PPC)**
 - Use PAUSE to proactively cycle links on and off
 - During link off time power-down LAN switch


PAUSE Power Cycle (PPC)

- **Basic idea is to periodically send PAUSE frames**
 - Power “stuff” down during PAUSE interval



PPC publications

- **Winner of Ethernet Alliance white paper contest**
 - Shown at EA booth in Interop 2008 at Las Vegas



Where this talk is going

- Energy consumption by IT equipment
- Reducing induced energy use
- Reducing direct energy use
- Additional efforts at USF

SNMP Power MIB

- **We want to expose and control power state**
- **Can we use SNMP to do this?**
- **Developing a Power MIB for desktop PC**
 - Power management capabilities
 - Power management settings
 - Total time for idle, busy, and sleep
 - Current elapsed time for idle, busy, or sleep
 - Statistics on wake-up events (network, user, etc.)
 - Statistics on sleep events
 - Actual power and energy use if metering is implemented

SNMP Power MIB continued

NAME	SYNTAX	ACCESS	DESCRIPTOR	UNITS
totalOffTime:1	Integer	Read-Only	Total time in off state	seconds
totalSleepTime:2	Integer	Read-Only	Total time in sleep state	seconds
totalOnTime:3	Integer	Read-Only	Total time in on state	seconds
totalInactiveTime:4	Integer	Read-Only	Total time in inactive state	seconds
totalActiveTime:5	Integer	Read-Only	Total time in active state	seconds
lastOfftime:6	Integer	Read-Only	Duration of last off time period	seconds
lastSleepTime:7	Integer	Read-Only	Duration of last sleep time period	seconds
lastOnTime:8	Integer	Read-Only	Duration of last on time period	seconds
lastInactiveTime:9	Integer	Read-Only	Duration of last inactive time period	seconds
lastActiveTime:10	Integer	Read-Only	Duration of last active time period	seconds
goToSleepCount:11	Counter	Read-Only	Counts number of transitions to sleep state	number
lastGoToSleepType:12	String	Read-Only	Describes what caused last transition to sleep	System/User
goToOffCount:13	Counter	Read-Only	Counts number of transitions to off state	number
lastGoToOffType:14	String	Read-Only	Describes what caused last transition to off	System/User
currentPowerState:15	String	Read-Write	Current state the system is in	Off/Sleep/On
powerSource:16	String	Read-Only	Current power source	AC/Battery
totalTimeACSource:17	Integer	Read-Only	Total time on AC power	seconds
totalTimeBattSrc:18	Integer	Read-Only	Total time on battery power	seconds
lastACTime:19	Integer	Read-Only	Duration of last time period on AC power	seconds
lastBattTime:20	Integer	Read-Only	Duration of last time period on battery power	seconds
powerLevel:21	Gauge32	Read-Only	Current power consumption level	Watts
energyConsumption:22	Gauge32	Read-Only	Total energy used	Watt-hours
energyLeft:23	Gauge32	Read-Only	Energy left (if on battery)	Hours
utilization:24	Guage32	Read-Only	Current utilization	percentage
inactivityOffTimer:25	Integer	Read-Write	Inactivity time-out to go to Off state	seconds
inactivitySleepTimer:26	Integer	Read-Write	Inactivity time-out to go to Sleep state	seconds
currentInactivityVal:27	Integer	Read-Only	Current inactivity timer value	seconds
acpiSupportStates:28	String	Read-Only	ACPI supported states	ACPI defined states
acpiEnabledStates:29	String	Read-Only	ACPI enabled states	ACPI defined states
acpiCurrentState:30	String	Read-Only	ACPI current state	ACPI defined states
acpiLastState:31	String	Read-Only	ACPI last state	ACPI defined states



Adaptive power management for PCs

- **PC operating systems use inactivity time-out**
 - For example, to put system to sleep
 - Use a fixed value for time-out
- **Can we do better with an adaptive time-out?**
 - Greater energy savings and less user annoyance
- **Have experimented with using past activity history**
 - Adaptively set the time-out value based on prediction
 - No conclusive results yet
 - Need to do more characterization of PC users

Lots of potential here!

Summary

- **Energy savings can be achieved**

- **Induced energy savings**

- Letting devices sleep without losing network presence

 How can we enable this for all end devices?

- **Direct energy savings**

- Matching link data rate to link utilization

 How can we exploit this for deeper savings?

Microsoft involvement

- **Need more Microsoft involvement in proxying standard**
 - Contact Bruce Nordman (bnordman@lbl.gov)
 - He will be in Redmond area on July 7th through 9th
- **Need funding to continue work in adaptive time-out**
 - Contact Ken Christensen (christen@csee.usf.edu)
 - Working with Ann-Gordon Ross at UF on this

My advisor said to always ask for money 😊.

Acknowledgments

- **Folks who have contributed...**

- Bruce Nordman
- Francisco Blanquicet
- Miguel Jimeno
- Jakob Klamra
- Jeremy Blackburn
- Many others

} Students in the group

Thank you!

Questions?

Francisco Blanquicet

Department of Computer Science and Engineering

University of South Florida

Tampa, FL 33620

fblanqui@cse.usf.edu

Key links

- **USF project website**
 - <http://www.csee.usf.edu/~christen/energy/main.html>
- **LBNL project website**
 - <http://efficientnetworks.lbl.gov/>
- **Project publications and invited talks**
 - <http://www.csee.usf.edu/~christen/energy/pubs.html>
- **ALR whitepaper at Ethernet Alliance**
 - http://www.ethernetalliance.org/technology/white_papers/alr_v10.pdf
- **Proxying whitepaper at Ethernet Alliance**
 - http://www.ethernetalliance.org/technology/white_papers/Proposal_for_Proxying_edit.pdf
- **IEEE 802.3az taskforce**
 - <http://www.ieee802.org/3/az/index.html>
- **UPnP Forum Low Power V 1.0**
 - <http://www.upnp.org/specs/lp.asp>