Energy Engineering for Protocols and Networks

IETF70 Vancouver Technical Plenary

Bruce Nordman and Elwyn Davies

Energy Efficiency in Protocols and Networks: *Why might the IETF care?*

- Keeping total consumption down
  - Lowering the cost to users
  - Reducing the impact on global energy use
- Limiting energy density in high end boxes
- Extending lifetime of battery power devices

Picture from Darth Vader Balloon

http://www.darthvaderballoon.be

Powered by the Internet?

Agenda

- Bruce Nordman from Environmental Energy Technologies Division of Lawrence Berkeley Labs
  - **"Networks, Energy and Energy Efficiency"**
  - http://efficientnetworks.lbl.gov
- Elwyn Davies
  - **"Long and Cool:**
    Engineering for Energy Efficiency"
  - Just a taster!

Networks, Energy, and Energy Efficiency

Bruce Nordman

Lawrence Berkeley National Laboratory

December 6, 2007

BNordman@LBL.gov — efficientnetworks.LBL.gov

Overview

- **How much energy does "The Internet" use**
- **How to think about networks / energy**
- **Current efficiency projects**
- **Energy saving opportunities**
- **IETF homework: Questions, Issues, Tasks**
- **Related topics**
- **Summary**

Key Collaborator: Ken Christensen, USF

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

ENERGY EFFICIENT DIGITAL NETWORKS

How much energy does the Internet use?

Forbes.com

Search: Forbes.com Quotes Video

U.S. EUROPE ASIA HOME PAGE FOR THE WORLD'S BUSINESS

HOME BUSINESS TECH MARKETS ENTREPRENEURS LEADERSHI

Video Blogs E-mail Newsletters Org Chart Wiki People Tracker Portfoli

E-mail | Print | Comments | Request Reprints | E-Mail Newsletters | My Yahoo! | RSS

Dig more coal -- the PCs are coming

Peter W. Huber and Mark P. Mills, 05.31.99

Southern California Edison, meet Amazon.com. Somewhere in America, a lump of coal is burned every time a book is ordered on-line.

The current fuel-economy rating: about a pound of coal to create, package, store and move 2 megabytes of data. The digital age, it turns out, is very energy-intensive. The Internet may someday save us bricks, mortar and catalog paper, but it is burning up an awful lot of fossil fuel in the process.

1999

"At least 100 million nodes on the Internet, … add up to … **8% of total U.S. demand**. … It's now reasonable to project that **half of the electric grid** will be powering the digital- Internet economy within the next decade."

emphasis added

Slashdot THE NUTS AND VOLTS OF NEWS FOR NERDS.

Log In | Create Account | Subscribe | Firehose

Sections

Main

Internet Uses 9.4% of Electricity In the US

Uclue got questions? (beta)

Home | Ask a Question | Categories | Browse Questions

★★★★★ ANSWERED on Fri 17 Aug 2007 - 6:17 pm MDT by davidsarokin

Question: Energy Use of Internet

2007

EFFICIENT DIGITAL NETWORKS

How much energy does the Internet use?

1999

"At least 100 million nodes on the Internet, … add up to … **8% of total U.S. demand**. … It's now reasonable to project that **half of the electric grid** will be powering the digital-Internet economy within the next decade."

emphasis added

Wrong Question

Wrong Answers

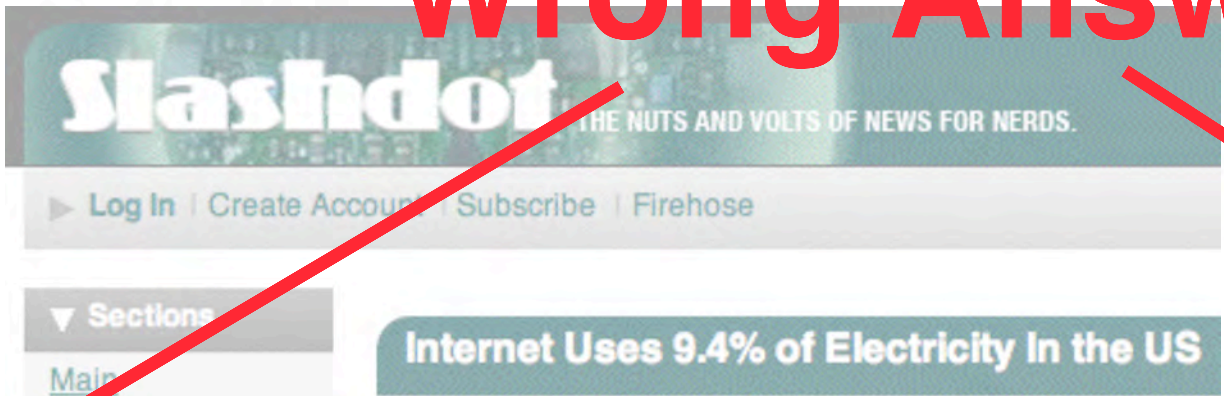

2007

EFFICIENT DIGITAL NETWORKS

How to think about energy quantities

Our needs only require approximations

• 1 year = 8,760 hours	~ 10,000 hours
• 1 kWh costs \$0.09	~ \$0.10
• 1 W for 1 year	~ \$1
• 1 TWh is 1 billion kWh	~ \$100 million
• U.S. annual consumption	~ 3,500 TWh

How much energy does network equipment consume?

	$billion	TWh/year
Telecom	$0.80	8.0
Data center	$0.20	2.0
Residential	$0.73	7.3
Commercial (office)	$0.88	8.8
Subtotal	$1.80	18
IP Service providers (access, metro, core)	< ?	< ?

- **All of these figures rough estimates for 2006**
- *None of this includes cooling or UPS*
- *$0.10/kWh used for convenience*

- **U.S. only — Global figures probably 3-5 times larger**

How about all Electronics?

- **PCs/etc., consumer electronics, telephony**
  - **Residential, commercial, industrial**
- **250 TWh/year**
- **About 7% of U.S. total electricity**
- **Well over $20 billion/year**
- **Over 180 million tons of CO_2 per year**
  - **Roughly equivalent to 30 million cars!**

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

Numbers represent U.S. only

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

Networks and Energy

Network equipment ….

Modems, routers, switches, wireless APs, …

… vs networked equipment

PCs, printers, set-top boxes, …

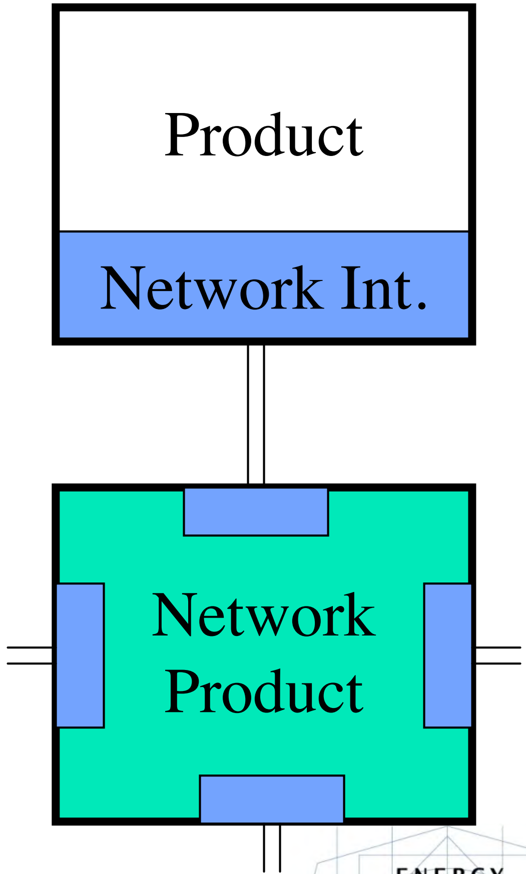

How networks drive energy use

- **Direct**
  - **Network interfaces (NICs)**
  - **Network products**
- **Induced in Networked products**
  - **Increased power levels**
  - **Increased time in higher power modes** (to maintain network presence)

Network Structure

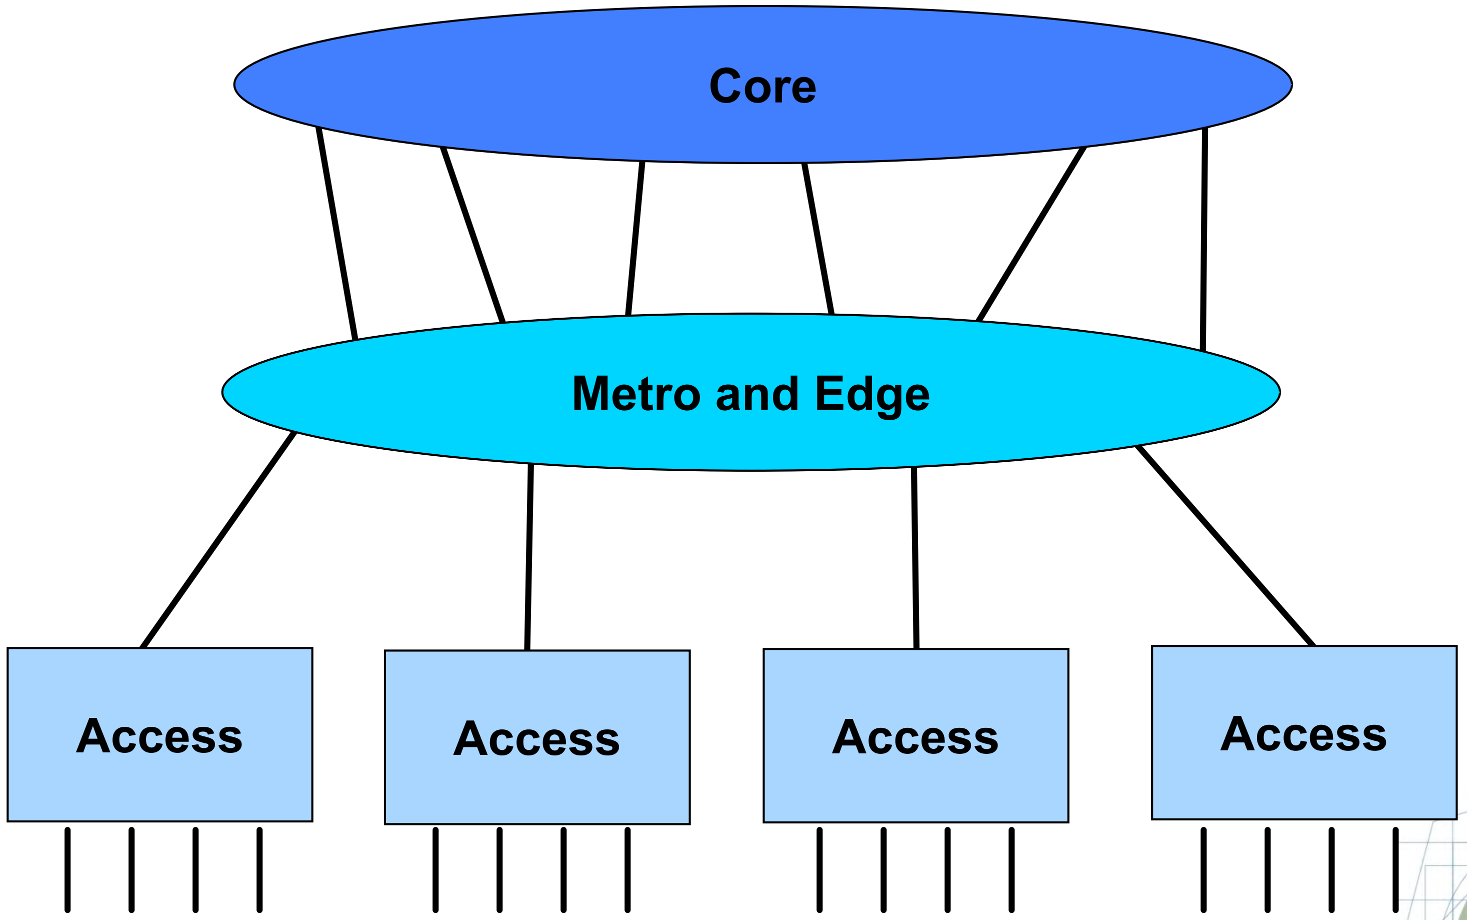

Source: Tucker et al., 2007

How should we think about networks and energy?

Approaches / Focus

- **Device**
  - **AC*-powered products**
- **Link**
  - **Capacity, usage, distance, technology**
- **Throughput**
  - **Traffic totals, patterns, distribution**
- **Application / Protocol**
  - **Drivers of infrastructure, nodes**
- **Context**
  - **In-use / not, time-sensitive / not, etc.**

Essential to use all approaches simultaneously

See Suresh Singh "Greening of the Internet" for good introduction

Efficiency Approaches

Product Focus | **Network Product Focus** | **Interface Focus** | **Protocol / Application Focus**

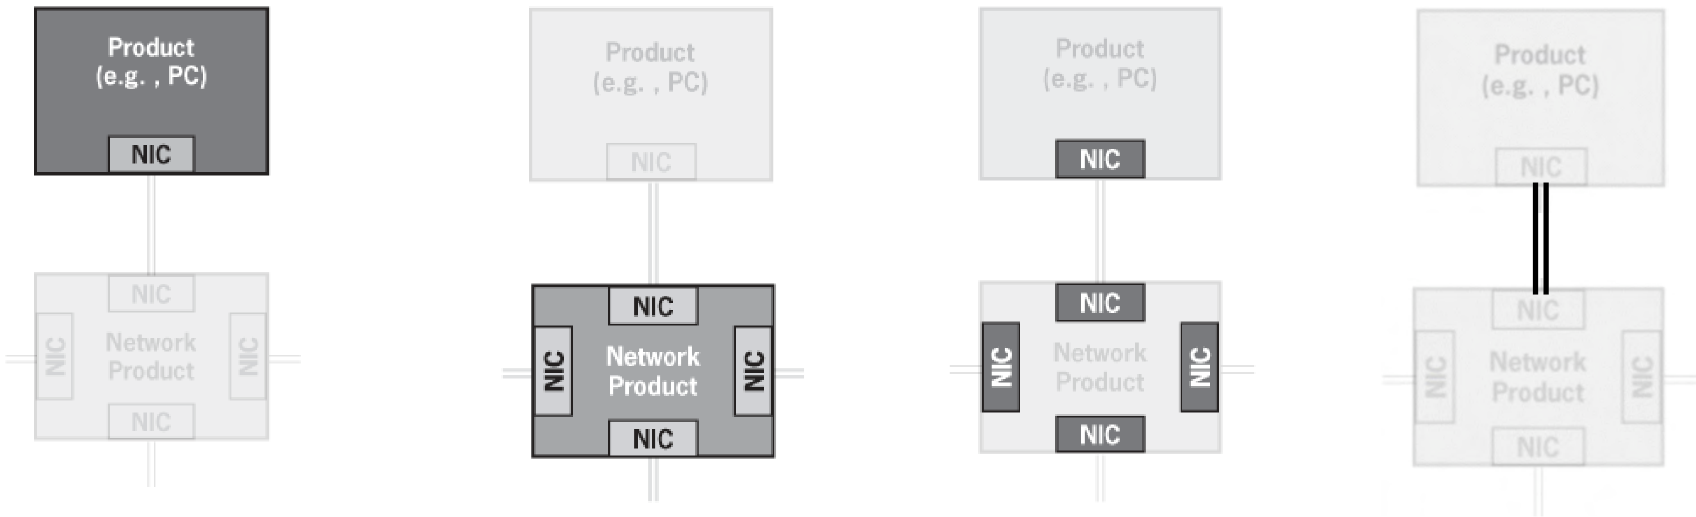

Need all approaches

Power Consumption Patterns

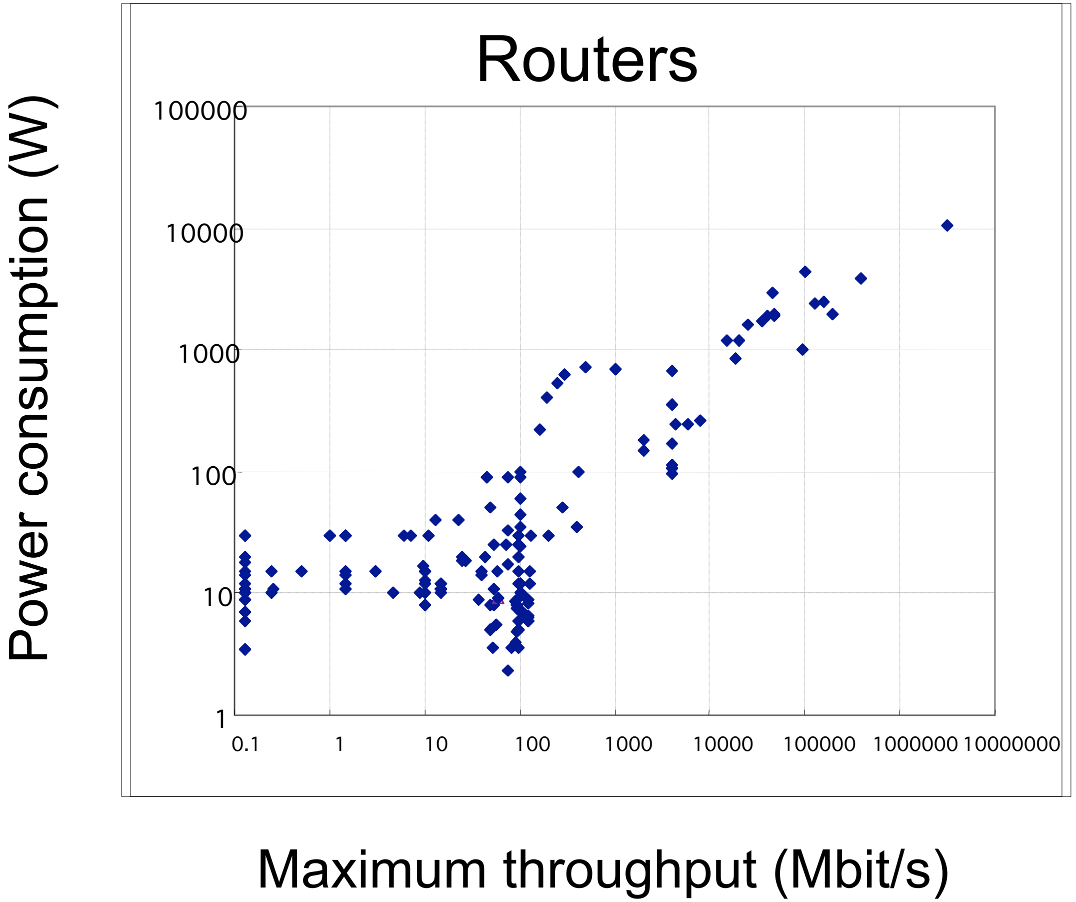

Source: METI, 2006

Sample utilization graphs

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

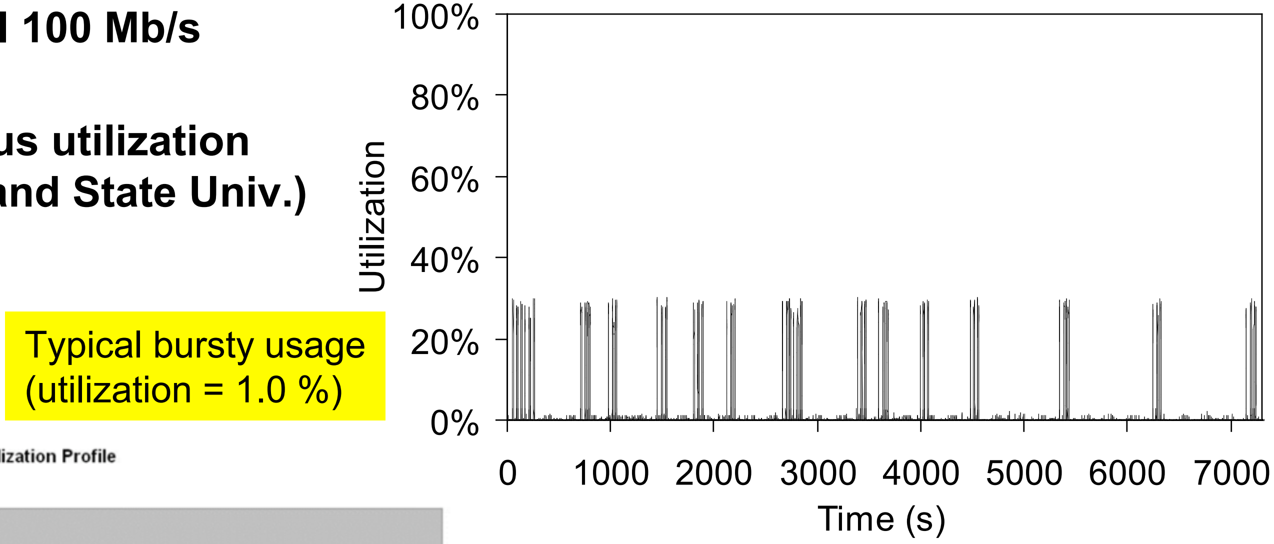

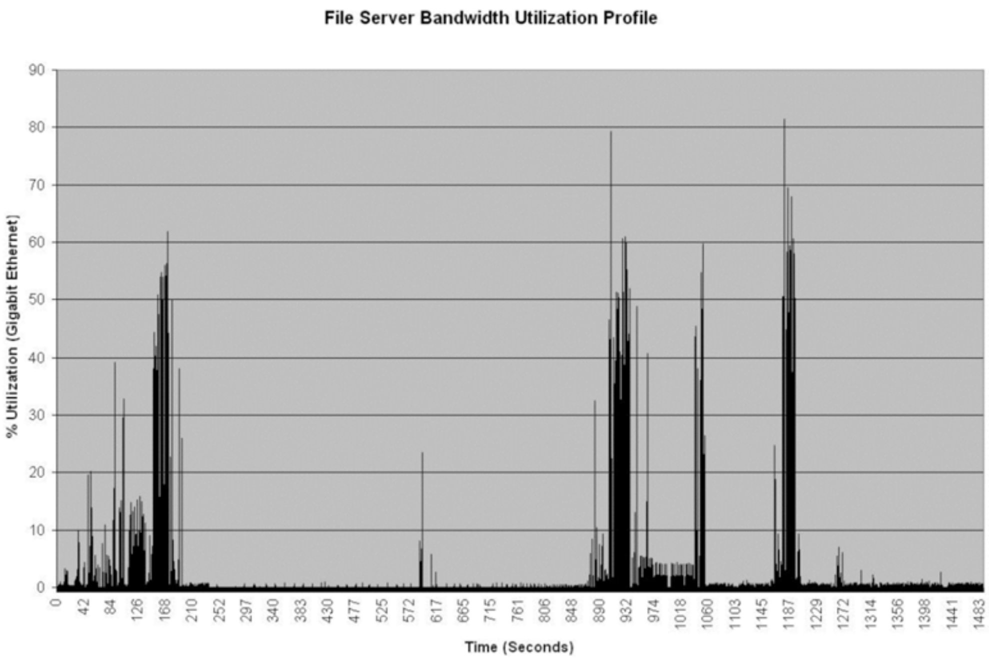

- **File server link utilization (daytime)** *(Bennett, 2006)*

Conclusions for edge links only
- **Burstiness**
- **Very low average utilization**

Utilization

- *Data networks are lightly utilized, and will stay that way*, A. M. Odlyzko, *Review of Network Economics*, 2003

Network	Utilization
AT&T switched voice	33%
Internet backbones	15%
Private line networks	3~5%
LANs	1%

Energy cost is a function of capacity, not throughput

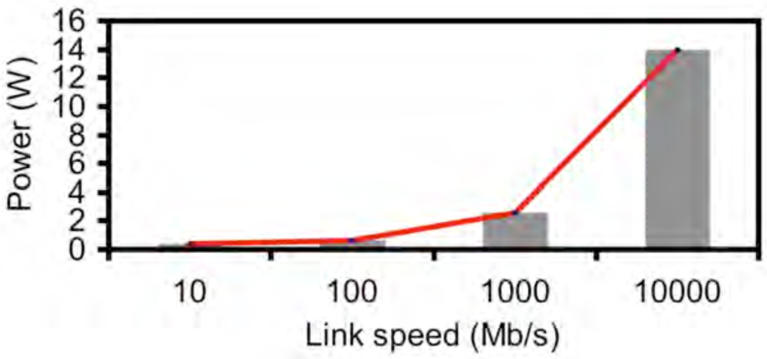

Measured power of various computer NICs (averaged)

Source: Christensen, 2005

Adaptive Link Rate —
Energy Efficient Ethernet

- **Concept**
  - **Add power management to Ethernet**
    - **Scale capacity to need**

Energy
Efficient
Ethernet

- **Method**
  - **Reduce link rate at times of low traffic levels**
    - **Most time on most links is low traffic levels**
  - **Quick transitions and seamless operation essential**
  - **This has implications on protocols that use/report link rate (link metrics, congestion control, MIBs etc.)**

- **Energy Savings**
  - **In network interface hardware and rest of system**
  - **In homes, commercial buildings, and data centers**
  - **U.S. direct savings — $ several hundred million/year**

- **Status**
  - **In midst of IEEE 802.3 standards process**
    http://ieee802.org/3/az/
  - **Hardware should be available in several years**

ENERGY
EFFICIENT
DIGITAL
NETWORKS

Network Presence Proxying

Core Fact: Most PC energy use occurs when no one present

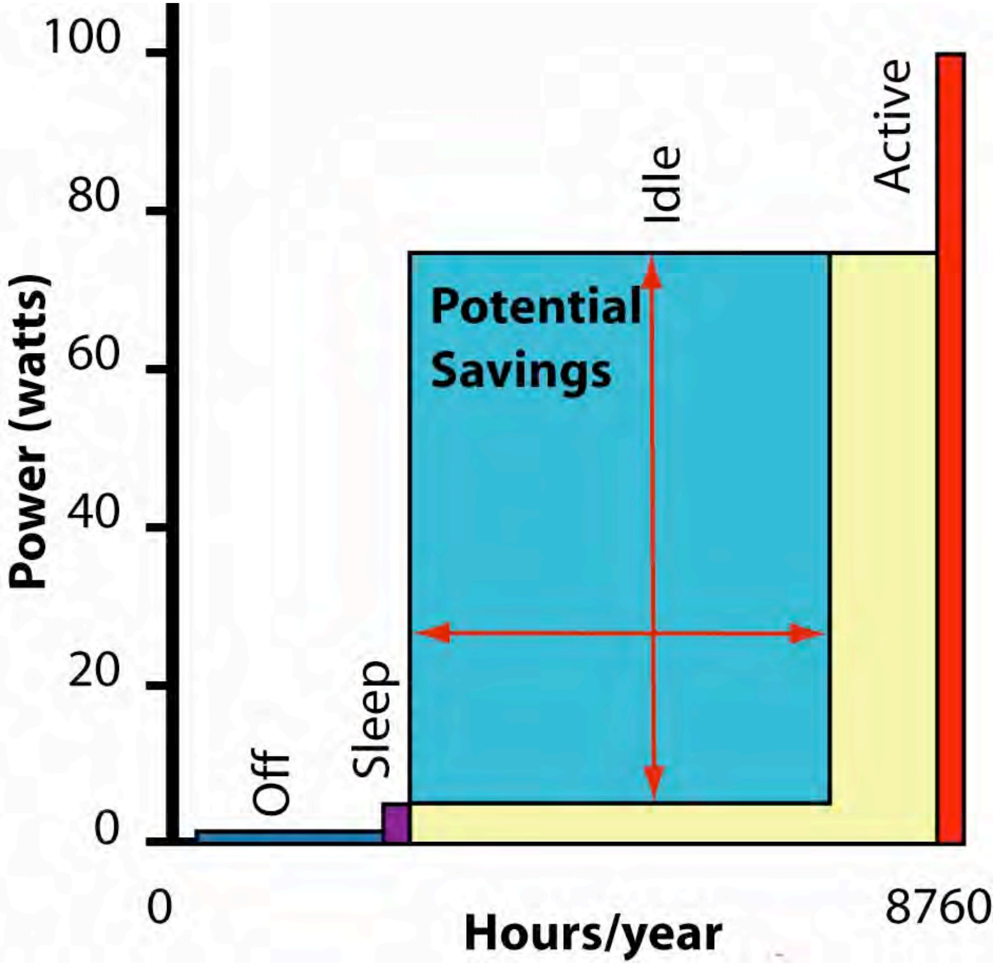

All time for year sorted by power level

Most of time when idle, could be asleep

PC savings potential is **most** of current consumption

Similar patterns apply to set-top boxes, for TVs, printer, …

Network Presence Proxying, cont.

- Concept
  - Allow sleeping hosts to remain fully network connected
  - Network does not know that host is asleep
  - Initially PCs - then printers, game consoles, set-top boxes, …
- Method
  - Define standard for how network interface can maintain "full network presence"
- Energy Savings
  - Likely < 1 W extra for proxy hardware
  - Avoids > 50 W for PC being on
  - U.S. direct savings — Easily > $1 billion/year
- Status
  - Working with industry to draft content of proxying standard
    http://www.ethernetalliance.org/technology/white_papers/
  - If you are more interested, contact me

ENERGY EFFICIENT DIGITAL NETWORKS

Network Presence Proxying, cont.

Proxy operation

1. PC awake; becomes idle
2. PC transfers network presence to proxy on going to sleep
3. Proxy responds to routine network traffic for sleeping PC
4. Proxy wakes up PC as needed

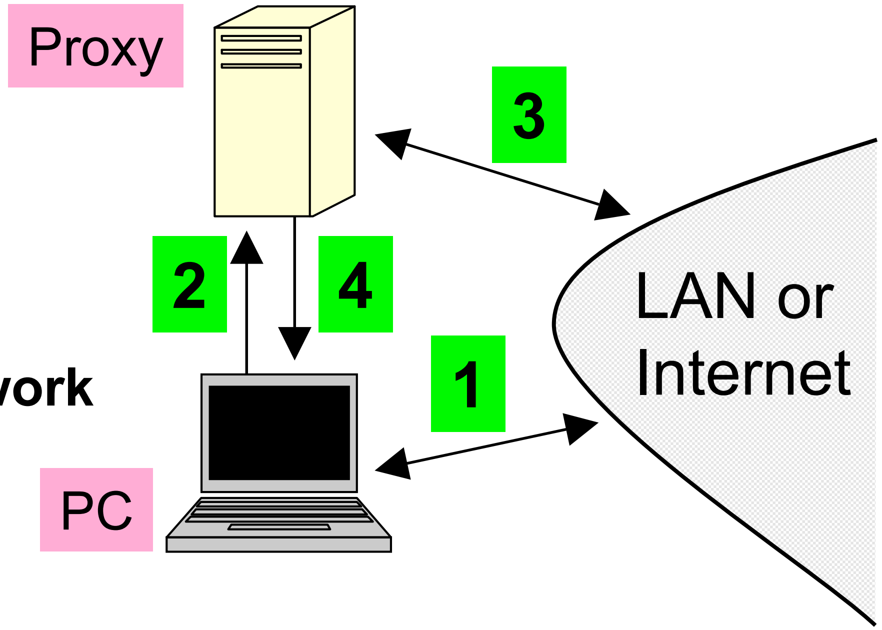

Proxy can be internal (NIC), immediately adjacent switch, or "third-party" device elsewhere on network

Proxy does: ARP, DHCP, TCP, ICMP, SNMP, SIP, ….

Exposing Power State

NSF/FIND Project

Enabling an Energy-Efficient Future Internet Through Selectively Connected End Systems

Allman, Christensen, Nordman, Paxson

- **Explore exposing power state to network and embodying in protocols**
- **Key issue: Distinct sleep state with reduced network connectivity**

http://www.icir.org/mallman/research/proj-energy-arch.html

Finding Energy Savings Opportunities

- Relax assumptions commonly made about networks
  - when feasible (rarely in core)
- These assumptions drive systems to peak performance
  - average conditions require less energy
  - many assumptions tied to latency
- Design for average condition, not just peak

Finding Energy Savings Opportunities

Common Internet design assumptions

- **Edge devices always fully present**
  - Facilitate multiple forms of reduced presence
- **Network links always at maximum speed**
  - Enable optional reduced speeds
- **Latencies always to be minimized**
  - Expose knowledge of acceptable latencies
- **Rate of increase infinite**
  - Determine when/how to facilitate slower acceleration
- **Maximize interconnections**
  - Facilitate powering-down links when capacity not needed
  - Tolerate non-trivial wake-up time for links
- **Always avoid intentionally losing packets**
  - Determine when some routine packet loss is OK

ENERGY EFFICIENT DIGITAL NETWORKS

What can IETF do ?

- What existing and developing protocols have features that:
  - –inadvertently work against energy saving ?
  - –facilitate energy saving ?
- What “guiding principles” might ensure that protocols maximize energy efficiency ?
- What revisions to existing protocols are warranted?
- What extensions to IETF scope are merited ?

Difficult Issues

- Security
  - In tussle with energy efficiency, security likely to win
- Management
  - Inattention to management will increase costs and reduce energy savings
- Bandwidth increases
  - IPTV
  - Widespread telepresence
  - ???

Related Topics

- **Consumer Electronics**
  - **–currently a mess from network protocol perspective**
- **Networking of non-electronic products (buildings)**
  - **lighting, climate control, …**
  - **–poised to become a similar mess**

- Both suffer from non-interoperability
- Both would benefit from an IETF-like structure

Real-world layering

User Interface

Applications

Concepts

Network Infrastructure

- Common standards for UI and concepts
- Concepts: presence, lights, windows, spaces, schedules, prices, …

Summary

- **Network energy use neither huge nor small**
  - **induced larger than direct**
- **Most energy use is at the edge**
- **Energy use is affected by applications and protocols**
  - **not just hardware**
- **Most opportunity is at non-peak conditions**
- **Protocols** (enhanced and new) **can greatly reduce energy use**
- **Should power state be widely exposed to the network?**
- **Need to extend protocols to the “real world”**
- **Make power/energy a key concern for all protocol design** (like security)

Thank you!

efficientnetworks.LBL.gov

Bruce Nordman
Lawrence Berkeley National Laboratory
BNordman@LBL.gov
510-486-7089

(m: 510-717-2916)

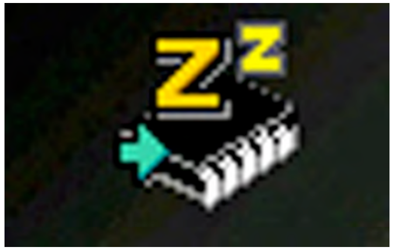

ENERGY EFFICIENT DIGITAL NETWORKS

Agenda

- Bruce Nordman from Environmental Energy Technologies Division of Lawrence Berkeley Labs
  - **"Networks, Energy and Energy Efficiency"**
  - http://efficientnetworks.lbl.gov
- Elwyn Davies
  - **"Long and Cool:
Engineering for Energy Efficiency"**
  - Just a taster!

Keeping the Core Cool

- Energy Density vs Total Consumption
  - Total usage in core routers not a big problem
    - Data centers are another matter
- Energy density in core routers is very high: High Capacity, Many Interfaces
  - Rack level
    - How to get the power in and the heat out?
  - Chip level
    - Core routers live on the electronic bleeding edge
      - Lots of the fastest ASIC technology
      - Large amounts of fast memory

Costs of Energy Density

- Energy 'Burden Factor'
  - 1.8 to 2.5 times power consumption of the electronics that does useful work
  - Complex cooling design needed
  - Power supply inefficiencies and loss
- Incompatibility with existing buildings
  - Too much power/cooling needed per rack
- Hot chips
  - High temperatures reduce reliability/life time

Red Hot Silicon

- High capacity routers rely on highest available speed and capacity ASICS

  $\Rightarrow$ high temperature, high power needs

- Latest ASIC technology 'wastes' ~50%
  - Fast $\Rightarrow$ Small (60nm/45nm) $\Rightarrow$ High Leakage
  - New 'High K' technology (Hafnium) helps
    - but not yet.. won't reach ASICs for a while

Every extra useful watt pulls in ~4 useless watts

Complexity and Featureitis

- From RFC 1958 - Architectural Principles of the Internet

  Section 3.2: ***If there are several ways of doing the same thing, choose one.*** *... Duplication of the same protocol functionality should be avoided as far as possible....*

- Extra complexity, more/duplicated features
  - in software...

    cost development time & reliability
    (and a bit of memory and power)
  - in hardware...

    add more silicon real estate &
    lots more power consumption

Helping the Core

- IP Layer and Routing have most effect
- Energy Density:
  An added concern but not a new story!
- Good Old Internet Principles still apply
  - KISS
  - One solution is better than two
  - Do it at the edge

Making it Last

(the opposite end of the scale)

- Many IP devices rely on battery power
  - Wireless Sensors
  - Mobile phones
  - etc.
- Need long periods of autonomous ops
  - Days or months
- Only possible if 'asleep' almost always!
  - and minimal transmission when awake

Recipes for Long Life

- Avoid
  - chatty protocols
    - no keepalives or periodic refreshes
    - minimize number of packets per transaction
  - 'waking up' whenever a packet comes by
  - naive routing protocols, e.g.
    - # of routing proto packets □ # of data packets
    - routing through power-challenged nodes
- Proxies (see Bruce's work) could help

Finally...

- 'Energy effectiveness' for a protocol design has lots of dimensions
  - Only skimmed the surface tonight
- Thanks to Bruce Nordman

Thanks to Dave Thaler and David Ward for their assistance

Thank You!

Questions?