

The next frontier for communications networks:  
Power management

Kenneth J. Christensen  
Department of Computer Science and Engineering  
University of South Florida  
Tampa, Florida 33620  
*christen@csee.usf.edu*  
*http://www.csee.usf.edu/~christen*

KJCO01 (keynote.ppt - 09/09/03)



final

The ~~next~~ frontier for communications networks:  
Power management

Kenneth J. Christensen  
Department of Computer Science and Engineering  
University of South Florida  
Tampa, Florida 33620  
*christen@csee.usf.edu*  
*http://www.csee.usf.edu/~christen*

KJCO02 (keynote.ppt - 09/09/03)



Protecting the environment:

~~The next frontier for communications networks:~~

Power management

Kenneth J. Christensen  
Department of Computer Science and Engineering  
University of South Florida  
Tampa, Florida 33620  
*christen@csee.usf.edu*  
*http://www.csee.usf.edu/~christen*

KJC003 (keynote.ppt - 09/09/03)



Saving the power grid:

~~The next frontier for communications networks:~~

Power management

Kenneth J. Christensen  
Department of Computer Science and Engineering  
University of South Florida  
Tampa, Florida 33620  
*christen@csee.usf.edu*  
*http://www.csee.usf.edu/~christen*

KJC004 (keynote.ppt - 09/09/03)



"Always on" without being always (fully powered) on:

~~The next frontier for communications networks:~~

Power management

Kenneth J. Christensen  
Department of Computer Science and Engineering  
University of South Florida  
Tampa, Florida 33620  
*christen@csee.usf.edu*  
*http://www.csee.usf.edu/~christen*

KJC005 (keynote.ppt - 09/09/03)



The next big challenge for performance evaluation:

~~The next frontier for communications networks:~~

Power management

Kenneth J. Christensen  
Department of Computer Science and Engineering  
University of South Florida  
Tampa, Florida 33620  
*christen@csee.usf.edu*  
*http://www.csee.usf.edu/~christen*

KJC006 (keynote.ppt - 09/09/03)



## Topics

---

- Power management - what and why
- Power management at many levels
- A day in the life of a dormitory
- Power management for desktop computers
- A proxying Ethernet adapter
- Summary and future directions

KJC007



## What and why

---

- What is performance evaluation all about?
  - ➔ In short... optimizing scarce resources
- Traditionally these resources have been...
  - CPU
  - memory
  - storage
  - bandwidth
- Also...
  - logic gates on a chip
  - I/Os on a chip

KJC008



## *What and why* continued

---

- But...

- ➔ CPUs are fast (desktop  $\gg$  server)
- ➔ Memory is cheap (256M cheaper than 64M at Office Max)
- ➔ Storage is cheap (\$1 per gigabyte)
- ➔ Bandwidth is plentiful (for most applications...)

Moore's Law

KJC009



## *What and why* continued

---

- So, where is the challenge?

- ➔ Power consumption is increasing
- Mobile community is worried about battery life
- Heat is a limiting factor in CPU speed
- Web hosting community is worried about operating costs
- Entire community is worried about the large-scale effects
  - Global warming - Kyoto agreement
  - Stress on power grids

KJC010



## What and why continued

---

- The third wave of power plants may be upon us...

First wave - light bulb

Second wave - electric motor

Third wave - microprocessor

But, surely, ecommerce and telecommuting  
reduce overall energy usage...

KJC011



## Ecommerce saves energy?

What costs more... driving to the mall in your SUV to buy a book,  
or using ecommerce to express ship it from across the country?

Roughly speaking, the dollar cost of an item is proportional to  
the energy consumed to create and deliver it.

KJC012



### *What and why* continued

---

- How big is this "third wave"?

**"Some experts calculate that the demands of the Internet already consume some 8 to 13 percent of electricity. If demand grows at just the same pace as during the last decade, we'll need nearly 1,900 new plants by 2020 -- or more than 90 every year -- just to keep pace."**

- Spencer Abraham (U.S. Energy Secretary)

<http://usinfo.state.gov/topical/global/climate/01040201.htm>

KJC013



### *What and why* continued

---

- And...

**"...reasonable to project that half the power grid will be powering the digital-Internet economy within the next decade."**

- Mills and Huber (Forbes, 1999)

**"A year ago we estimated that some 13 percent of U.S. power output was being used to manufacture and run computers and the sprawling information technology infrastructure. It's more than that today."**

- Mills and Huber (Wall Street Journal, 2000)

<http://www.wired.com/news/technology/0,1282,40701-2,00.html>

KJC014



## *What and why* continued

---

- And more...

**“The current fuel economy rating: about 1 pound of coal to create, package, store, and move 2 megabytes of data.”**

- Mills and Huber (Forbes, 1999)

**“There is no empirical evidence to support those numbers. His estimates are absurd.”**

- Jonathan G. Koomey (LBNL, Energy End-Use Forecasting)

<http://www.wired.com/news/technology/0,1282,40701-2,00.html>

KJC015



## *What and why* continued

---

- Who is LBNL, Energy End-Use Forecasting?
  - LBNL = Lawrence Berkeley National Laboratory

The image is a screenshot of the 'ENERGY END-USE FORECASTING' website. The title is in large, bold, blue letters at the top. Below the title, there are several menu items: 'About End-Use Forecasting' (with sub-items: research focus, staff/contact info), 'Publications & Data' (with sub-items: publications, data), 'Project Areas' (with sub-items: ENERGY STAR program support, Technology data and modeling, Policy analysis), and 'Network for energy, environment, efficiency & the information economy'. On the right side, there is a 'Search' button. The background of the website features a collage of images including a city skyline, a person walking, and a network diagram. At the bottom right, there is a quote: "We transform market and technology data into information to provide a sound basis for policy decisions promoting the development and adoption of cost-effective energy-efficiency technologies." followed by a small portrait of Jon Koomey.

KJC016





## *What and why* continued

---

- The problem...

“We found that total direct power use by office and network equipment is about 74 TWh per year, which is about 2% of total electricity use in the U.S. When electricity used by telecommunications equipment and electronics manufacturing is included, that figure rises to 3% of all electricity use (Koomey 2000). More than 70% of the 74 TWh/year is dedicated to office equipment for commercial use. We also found that power management currently saves 23 TWh/year, and complete saturation and proper functioning of power management would achieve **additional savings of 17 TWh/year**. Furthermore, complete saturation of **night shut down for equipment not required to operate at night would reduce power use by an additional 7 TWh/year.**”

- Kawamoto et al. (LBNL, Energy End-Use Forecasting)

<http://enduse.lbl.gov/Info/LBNL-45917b.pdf>

KJC017



## *What and why* continued

---

- How much is 7 TWh/year???
  - At 8 cents per kWh... \$560 million per year
- Or...



KJC018



The speaker is in the wrong room,  
this is a *communications conference*.

- What does this have to do with communications?
- What does this have to do with traffic characterization?

It does! Wait and see...

KJC019



### *What and why* continued

- At SIGCOMM 2003...



#### Greening of the Internet

Maruti Gupta  
Department of Computer Science  
Portland State University  
Portland, OR 97207  
mgupta@cs.pdx.edu

Suresh Singh  
Department of Computer Science  
Portland State University  
Portland, OR 97207  
singh@cs.pdx.edu

#### ABSTRACT

In this paper we examine the somewhat controversial subject of energy consumption of networking devices in the Internet, motivated by data collected by the U.S. Department of Commerce. We discuss the impact on network protocols of saving energy by putting network interfaces and other router & switch components to sleep. Using sample packets

Device	Approximate Number Deployed	Total AECTW-h
Hubs	93.5 Million	1.6 TW-h
LAN Switch	95,000	3.2 TW-h
WAN Switch	50,000	0.15 TW-h
Router	3,257	1.1 TW-h
<b>Total</b>		<b>6.05 TW-h</b>

pp. 19-26

KJC020



### *What and why* continued

---

- Gupta and Singh describe annual US energy use...
  - From a study for the DOE

Device	Deployed	TWh/yr
Hubs	93,500,000	1.60
LAN switch	95,000	3.20
WAN switch	50,000	0.15
Router	3,257	1.10
Total		6.05

- 20K to 35K terabytes routed on the US Internet in December 2000
  - A. Odlyzko (University of Minnesota)

[http://www.eere.energy.gov/state\\_energy/technology\\_otherinfo.cfm?techid=17](http://www.eere.energy.gov/state_energy/technology_otherinfo.cfm?techid=17)

KJC021



### *What and why* continued

---

- Gupta and Singh discuss...
- Energy consumption of networking devices to increase
  - Increase of 1 TWh by 2005
- Packet traces show that 90% of time an interface can sleep
- High-level ideas for coordinating routing, QoS, and sleeping
  - Changes to OSPF to reduce messages sent
  - Aggregation to use fewer links
  - Activate links on an "as needed" basis

KJC022



### *What and why* continued

---

- Gupta and Singh argue that...
  - Internet is three times less efficient than 802.11
  - Significant because wireless links are *not* efficient

• Thus, there is room for improvement!

- With significant impact...

“The impact of saving energy is huge, particularly in the developing world where energy is a precious resource whose scarcity hinders widespread Internet deployment.”

- Gupta and Singh (2003)

KJC023



### *What and why* continued

---

- What really is the fuel rating???

➔ Between 5.5 to 9.7 Wh to send 2 megabytes

2.5 to 4.4 grams of coal  
(~ 1/10 of an ounce)

- Calculated using 1.25 TWh/yr for WAN switches and routers
- The weight of one penny is about 2.5 grams
  - Cost is about 0.01 cents for this much coal

Weight in coal =



KJC024



## Topics

---

- Power management - what and why
- Power management at many levels
  - Definitions
  - Methods and challenges
  - Cost to operate a PC
- A day in the life of a dormitory
- Power management for desktop computers
- A proxying Ethernet adapter
- Summary and future directions

KJC025



## At many levels

---

- Some quick definitions...
  - Power is  $W = V \times A$
  - Energy is  $Wh = \text{Power} \times \text{Time}$
- Consumed energy produces useful work... and heat
- Heat costs money in cooling
  - 25% of the cost of a web hosting facility is cooling
- For mobile devices, energy use consumes battery
  - Empty battery = *mobile user* not mobile anymore
  - Empty battery = *sensor network* node not sensing anymore

KJC026



### *At many levels* continued

---

- Three general methods for power management...

#### Method #1 - process and transmit less

- Transmitting is very expensive for wireless
- Sensor network community studying new routing protocols
  - » Source routing is back!

#### Method #2 - slow-down

- Process no faster than needed (be deadline driven)

#### Method #3 - turn-off "stuff" not being used

- Within a chip
- At a component level
- At a system level

KJC027



### *At many levels* continued

---

- Method #1 example - Eliminate routing updates
  - OSPF transmits updates whenever a link changes
  - Links change a lot in mobile ad hoc networks
  - Use source routing to discover the current best route
- Method #2 example - Voltage-frequency scaling at chip level
  - Reducing the voltage requires reducing the frequency
  - Process no faster than "fast enough"
- Method #3 example - Turning on and off servers in a cluster
  - Wish to maintain response time at a given level
  - No benefit in "too fast" a response time
  - Turn-off servers as a function of request load

KJC028



### *At many levels* continued

---

- Time scales of idle times
  - CPU and instruction level (nano to micro seconds)
  - Inter-packet (micro to milliseconds)
  - Inter-flow (seconds to hours)
- A flow is a TCP connection or other session
- Predicting, controlling, and making use of idle times is key
  - Inter-packet -- turn-off the processor
  - Inter-flow -- turn-off the system
- Power-down and power-up are *not* instantaneous
  - Function of technology used (getting faster...)
  - Power-up time can affect response time of a request

KJC029



### *At many levels* continued

---

- Key challenges to dynamic power management:

- 1) Predicting, controlling, and making the best use of idle times
- 2) Increasing the predictability of idle times
- 3) Creating additional idle time by bunching and/or eliminating traffic

KJC030



### *At many levels* continued

---

- What does it cost to operate a single PC?
  - About 100 W fully powered-on
  - System unit has increased while monitor has decreased
- Measured Dell 1.8-Ghz P4 with 256MB RAM and 19" LCD monitor
  - System unit powered-on = 60 to 85W
  - Monitor = 27W
  - Windows standby = 7 W
- A typical household is 10,219 kWh/yr (DOE)
  - A 100W PC always on 24/7/365 adds 876 kWh
  - Or, about 8.6% increase in household power consumption
- How many new PC's do you have in your house?
- Broadband is increasing the "always on" time

KJC031



### *At many levels* continued

---

- What does it cost to run a lot of PCs?
- Michigan State University asking students to shut-down PC's when going on winter break (two weeks)
  - “Shutting down the computers across campus over the winter break could save as much as \$20,000...”
  - Terry Link (Michigan State office of campus sustainability)

KJC032





## Topics

---

- Power management - what and why
- Power management at many levels
- A day in the life of a dormitory
  - Cisco NetFlow
  - Characterizing busy and idle times in USF dormitory PCs
- Power management for desktop computers
- A proxying Ethernet adapter
- Summary and future directions

KJC033



## Traffic characterization

---

- College students are at the cutting-edge of network applications
  - Lots of peer-to-peer file sharing!
- Collected one day (about 24 hours) of flow data from top 100 PCs
  - Top 100 users by volume (tracked by MRTG)
  - From USF dormitories (about 3000 PCs)
- Goal was to characterize the flow-level traces
  - Idle and busy periods for each host

### Definitions:

*Idle period* = no flows active

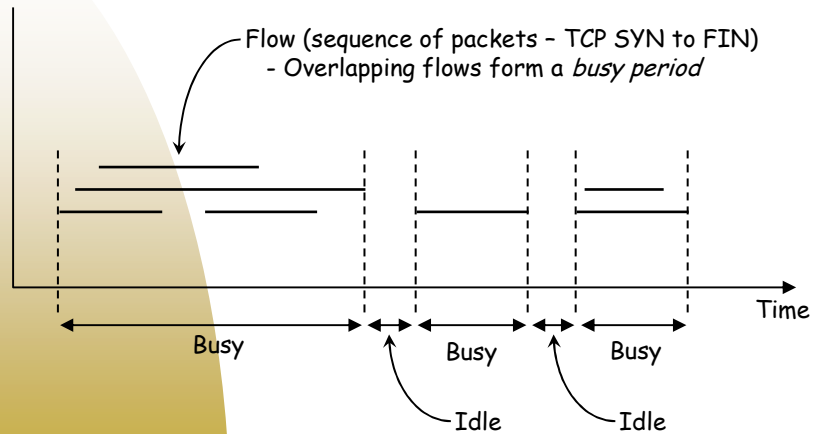
*Busy period* = one or more flows active

KJC034



## Traffic characterization continued

- Idle and busy periods



KJC035



## Traffic characterization continued

- Flows are Cisco NetFlow records
- Natively collected by Cisco routers
- A flow is a unidirectional sequence of packets
  - Delimited by SYN and FIN for TCP

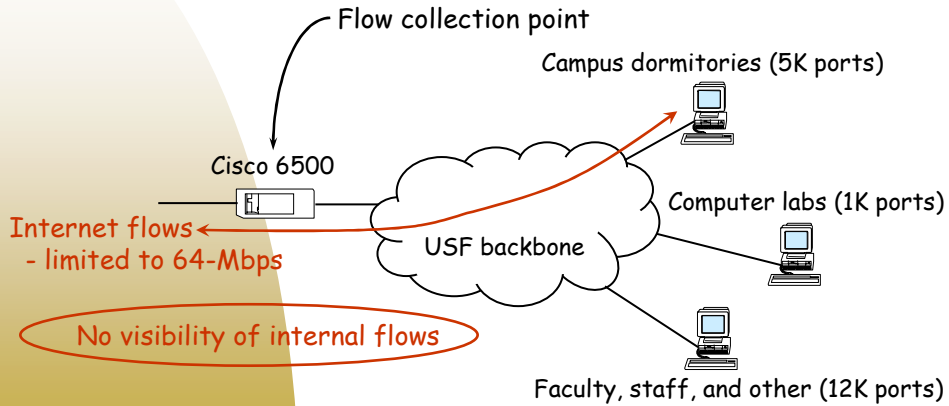
Start time/date: Time and date  
End time/date: Time and date  
SrcIPAddress: Source IP address  
SrcP: Source port number  
DstIPAddress: Destination IP address  
DesP: Destination port number  
P: Protocol  
Pkts: Number of packets  
Octets: Number of octets

KJC036



*Traffic characterization* continued

- Network configuration for top 100 PCs flow collection
  - Flows collected on March 27, 2003

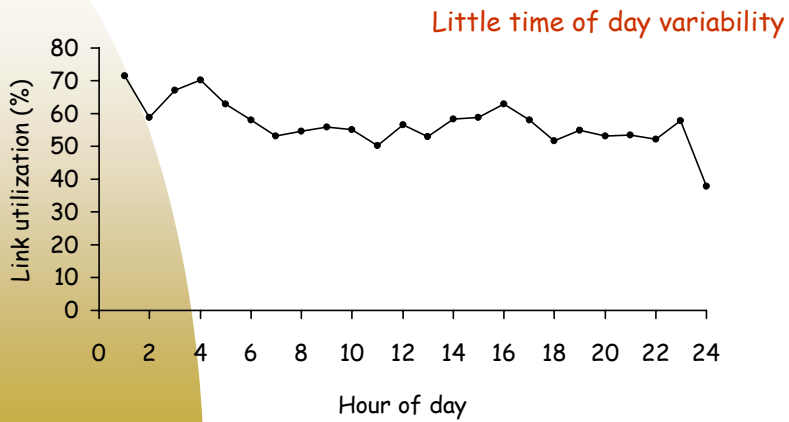


KJC037



*Traffic characterization* continued

- Link utilization (64-Mbps = 100%)
  - Outgoing = 12.4%, incoming = 44.0%



KJC038



### Traffic characterization continued


- Volume of data sent and received in 24 hours
  - Received = 469 Gbytes
  - Sent = 132 Gbytes
- By application
  - eDonkey = 2%
  - Kazaa = 19%
  - Web = 1%
  - Unknown = 78% (port hopping of above applications?)
- By protocol
  - TCP = 99%
  - UDP = 1%
  - Negligible amounts of ICMP

KJC039



### Traffic characterization continued

- Volume of data sent and received by PC (ranked)
  - Max = 72 Gbytes, mean = 5.8 Gbytes

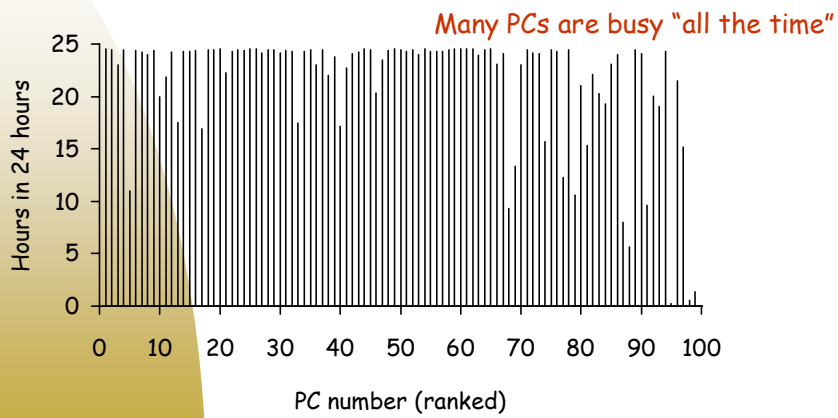


KJC040



*Traffic characterization* continued

- Total busy time (for 100 PCs)
  - Max = 24.5 hrs, mean = 21.2 hrs,  $\sigma = 6.0$  hrs

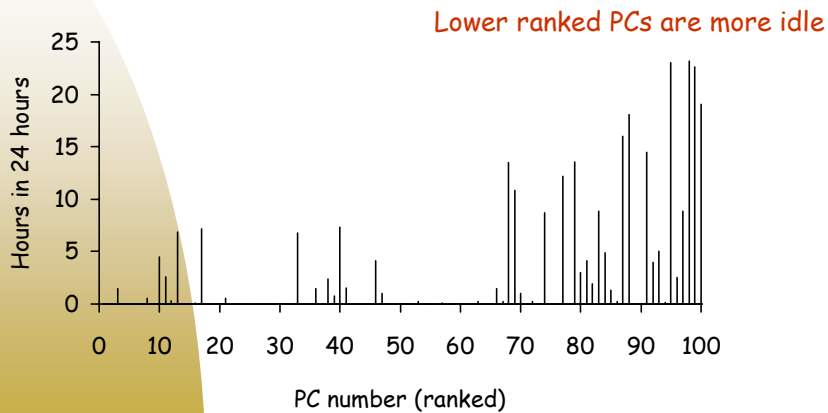


KJC041



*Traffic characterization* continued

- Total idle time (for 100 PCs)
  - Max = 23.2 hrs, mean = 2.9 hrs,  $\sigma = 5.6$  hrs




KJC042



*Traffic characterization* continued

- Distribution of flow size (for 100 PCs - 11,048,003 flows)
  - Max = 0.97 Gbytes, mean = 57 Kbytes,  $\sigma = 1.72$  Mbytes

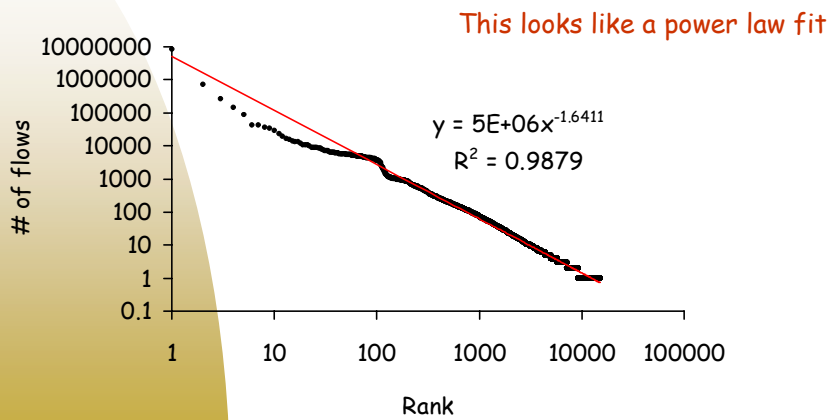


KJC043



*Traffic characterization* continued

- Power law fit for flow size (for 100 PCs)

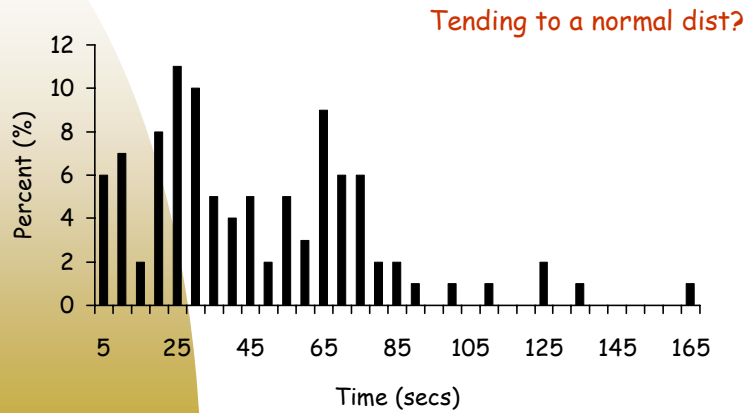


KJC044



*Traffic characterization* continued

- Distribution of flow length (for 100 PCs)
  - Max = 2.7 min, mean = 44 sec,  $\sigma$  = 31 sec

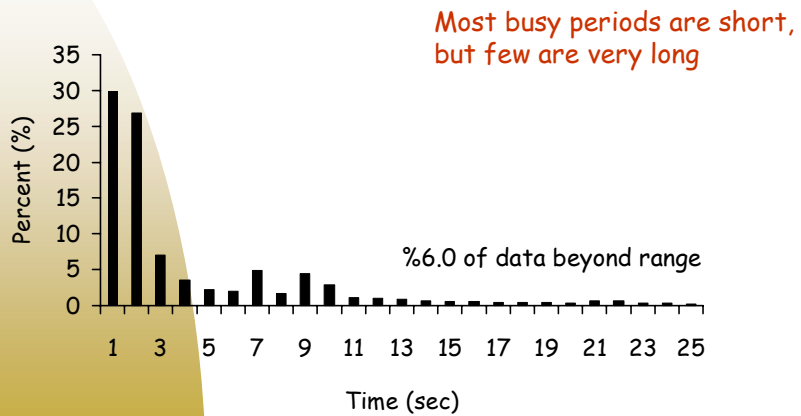


KJC045



*Traffic characterization* continued

- Distribution of busy periods (for 100 PCs - 70,533 busy periods)
  - Max = 24.5 hrs, mean = 1.8 min,  $\sigma$  = 41.6 min

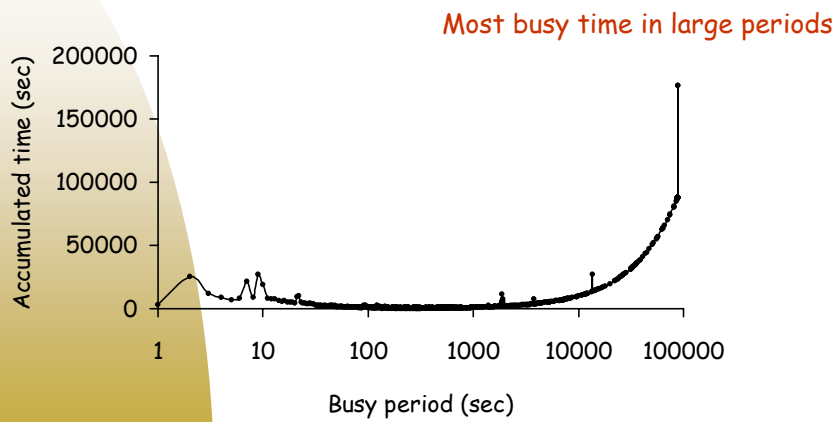


KJC046



*Traffic characterization* continued

- Accumulated time plot for busy periods

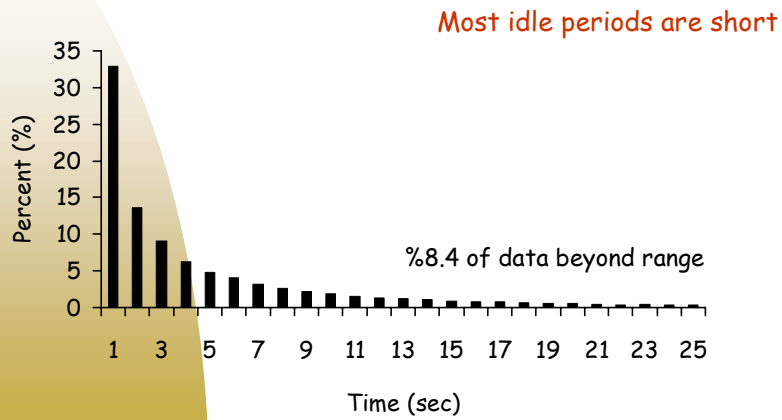


KJC047



*Traffic characterization* continued

- Distribution of idle periods (for 100 PCs - 70,434 idle periods)
  - Max = 4.7 hrs, mean = 15.0 sec,  $\sigma = 2.1$  min



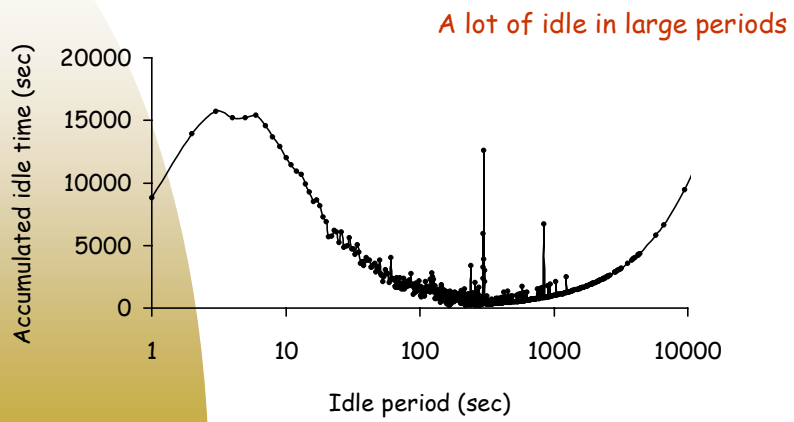
KJC048





*Traffic characterization* continued

- Accumulated time plot for idle periods

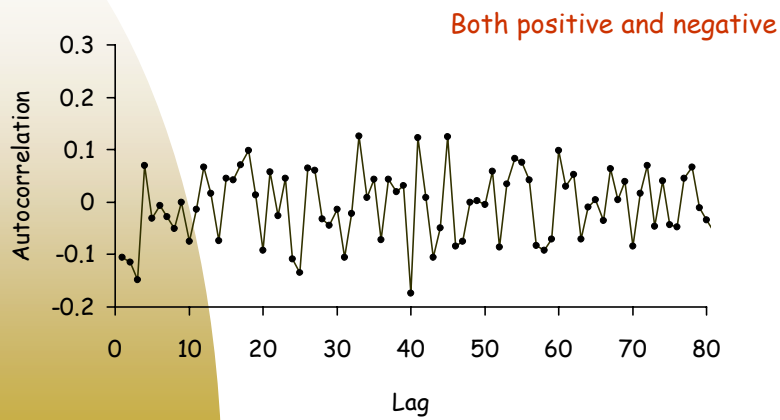


KJC049



*Traffic characterization* continued

- Autocorrelation of busy period (for 100 PCs)

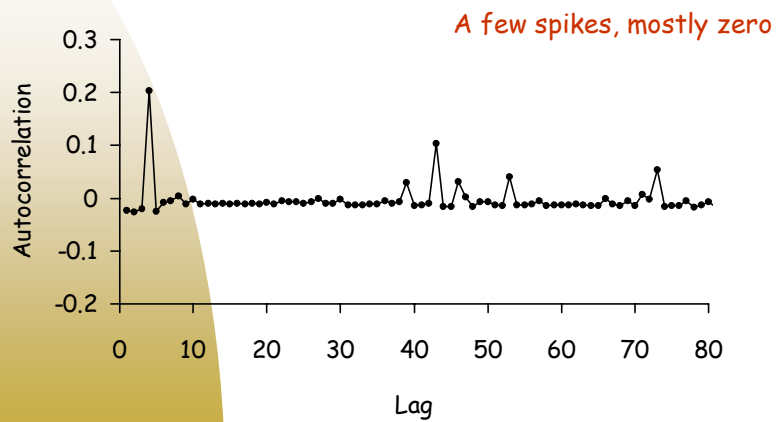


KJC050



### Traffic characterization continued

- Autocorrelation of idle periods (for 100 PCs)



KJC051



### Traffic characterization continued

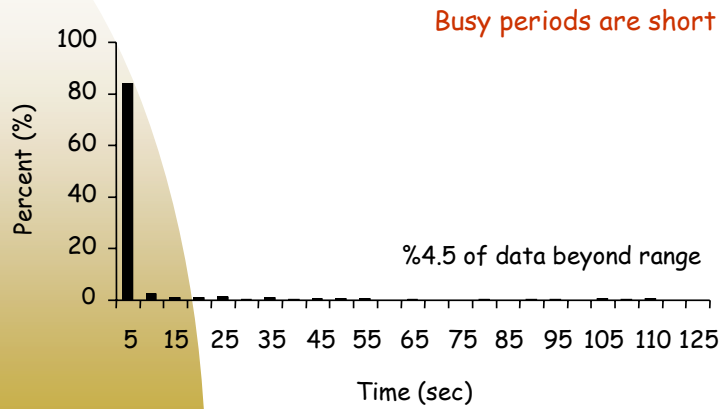
- Correlation statistics (for 100 PCs)
  - Busy to next idle = -0.12
  - Busy number of flows to next idle = -0.33
  - Busy number of flows to current busy = 0.24

KJC052



*Traffic characterization* continued

- Distribution of busy period (for "typical" PC)
  - Max = 6.7 hrs, mean = 47.2 sec,  $\sigma = 12.6$  min

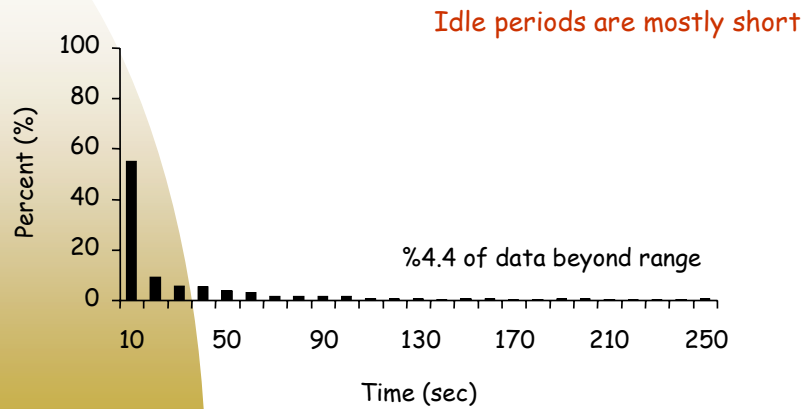


KJC053



*Traffic characterization* continued

- Distribution of idle period (for "typical" PC)
  - Max = 5.0 min, mean = 38.2 sec,  $\sigma = 72$  sec




KJC054



### Traffic characterization continued

- Autocorrelation of busy and idle periods (for "typical" PC)



KJC055



### Traffic characterization continued

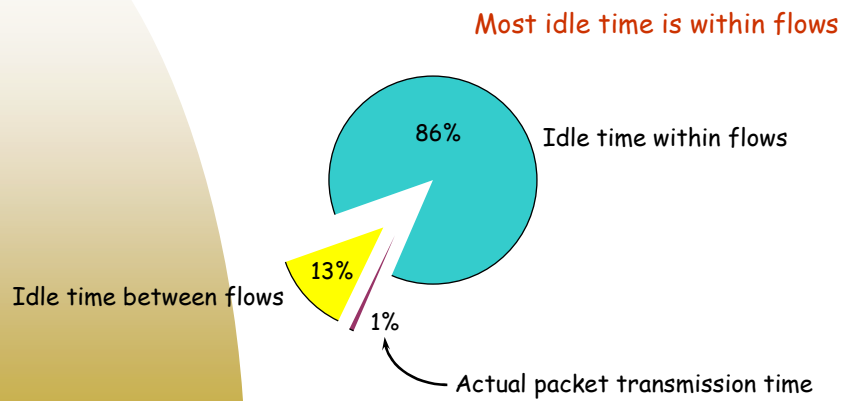
- Correlation statistics (for "typical" PC)
  - Busy to idle = 0.04
  - Busy number of flows to idle = -0.13
  - Busy number of flows to current busy = 0.89

KJC056



*Traffic characterization* continued

- For the total idle time for the 100 PCs...
  - How much idle time is *within* and *between* flows?

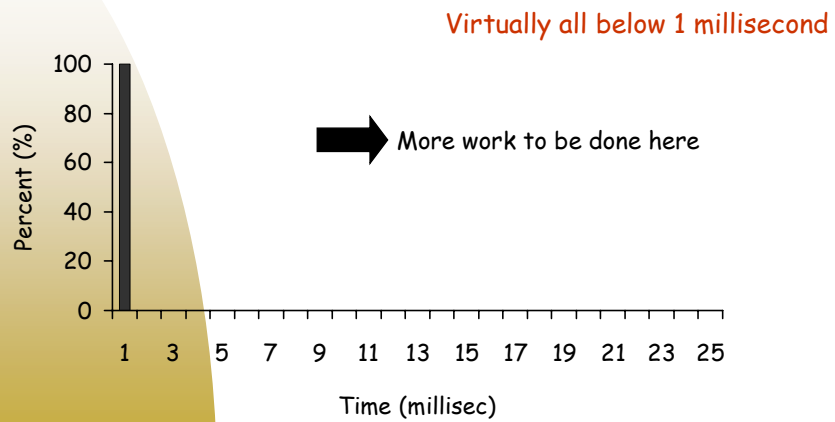


KJC057



*Traffic characterization* continued

- Distribution of idle within a flow
  - FTP on a 100-Mbps link in the lab



KJC058



## Topics

---

- Power management - what and why
- Power management at many levels
- A day in the life of a dormitory
- Power management for desktop computers
  - Energy Star and an executive order
  - The problem - disabling of power management
  - Industry directions (including Wake on LAN)
  - Power management with time-out

A proxying Ethernet adapter

- Summary and future directions

KJC059



## Power management

---

- EPA Energy Star program...

**“ENERGY STAR is a government-backed program helping businesses and individuals protect the environment through superior energy efficiency.”**

- EPA (2003)



**ENERGY STAR®**  
Money Isn't All You're Saving

KJC060



### *Power management* continued

---

- EPA Energy Star for office equipment (started in 1991)
- EPA MOU ("spec") to manufacturers for logo
  - Government purchases must be Energy Star logo'ed
- Key criteria (from EPA):

- Automatically enter a low-power "sleep" mode after a period of inactivity
- Energy-efficiency specifications based on power supply
- Include mechanisms through which the low-power modes of qualified monitors can be activated

KJC061



### *Power management* continued

---

- Executive Order 13221

**“By the authority vested in me as President by the Constitution and the laws of the United States of America ... it is hereby ordered as follows:**

**Section 1. Energy-Efficient Standby Power Devices.**

**Each agency, when it purchases commercially available, off-the-shelf products that use external standby power devices, or that contain an internal standby power function, shall purchase products that use **no more than one watt in their standby** power consuming mode. ...”**

- President Bush (signed by)

KJC062



## *Power management* continued

---

- The problem is...

**“PC and Monitor Night Status: Power Management Enabling and Manual Turn-off from 2000 estimates that the enabling rate is at most 25% for PCs, and 60% for displays.”**

- Bruce Nordman et al. (LBNL, 2000)

➔ In other words, users *disable* Power Management

<http://www.aceee.org/conf/00ss/00sstoc7.pdf>

KJC063



## *Power management* continued

---

- Power management is disabled because...

- Cannot remotely manage or access the PC
- Cannot share files
- Lost work if sharing to other computers

KJC064





## *Power management* continued

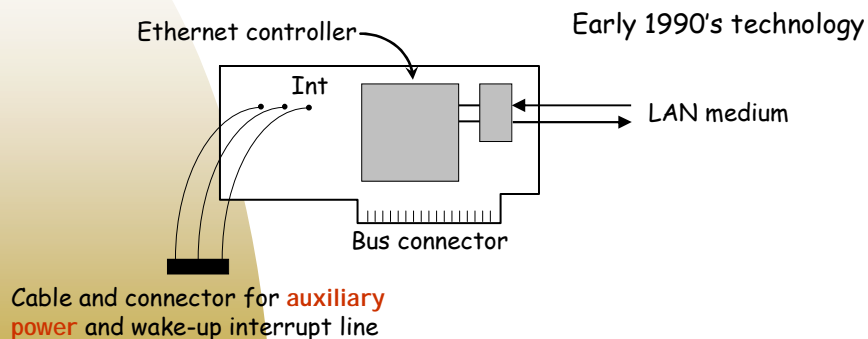
- Many industry initiatives...
- Advanced Configuration and Power Interface (ACPI)
  - Compaq/Intel/Microsoft/Phoenix/Toshiba initiative
- Develop industry common interfaces
  - Device and system power management
  - Operating System Power Management (OSPM)
- Define global power states
  - Working / Sleeping / Soft off / Mechanical off
  - Ongoing effort to standardize nomenclature and symbols

KJC065



## *Power management* continued

- Wake-on-LAN (WOL) for Ethernet
  - A specially defined packet to trigger a wake-up interrupt
  - WOL-cable is built-in to most PC motherboards



KJC066



### *Power management* continued

---

- Wake-on-LAN (WOL) packet is a "MAC frame"
  - Cannot be routed
  - Need to know MAC addresses
  - Not part of TCP/IP protocol implementations
- Desirable to wake-up on existing protocol events
  - Wake-up on a valid TCP SYN packet?
- Should PC wake-up on every incoming packet?
  - Clearly, not

KJC067



### *Power management* continued

---

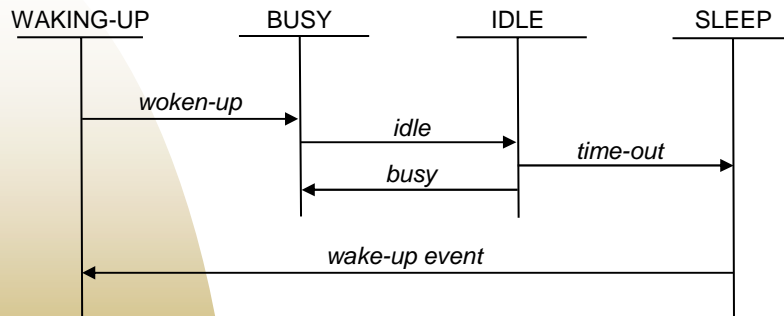
- System power-down using an *inactivity timer*
  - Time-out value is fixed or determined adaptively
- Currently used for monitor and disk power-down in PCs
- Assume that a mechanism to wake-up system exists
- We apply inactivity timer to our dormitory PCs
  - Simulated using packet traces

KJC068



## Power management continued

- FSM for power management with network inactivity time-out



### Notes:

- In the IDLE state an idle (inactivity) timer is started
- A *wake-up event* can be a connection request or other activity
- The time needed to wake-up is a performance penalty
- In the SLEEP state achieving low power does not occur instantly

KJC069



## Power management continued

- Waking-up takes time
  - Time continues to decrease as technology improves
- Full system power-up in seconds
  - MRAM may reduce to 10's of milliseconds
- CPU power-up in 100's of milliseconds

➔ Performance impact (on response time)

### Definitions:

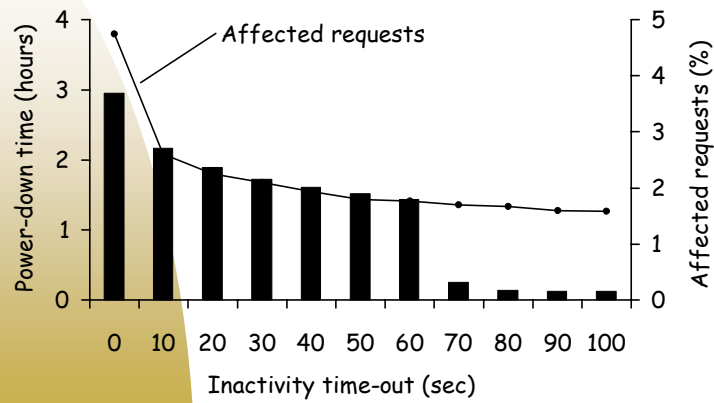
*Affected request* = The flow immediately following an idle period that triggers a wake-up

KJC070



*Power management* continued

- Results for fixed inactivity time-out (for 100 PCs)  
- Powered-down time and affected requests




KJC071



*Power management* continued

- Results for fixed inactivity time-out (for "typical" PC)  
- Powered-down time and affected requests



KJC072



### *Power management* continued

---

- Can we improve on a fixed time-out?

➡ It appears not! (But, more work is needed)

- We implemented existing methods used for disk drive spin-down
  - Results were worse than fixed time-out
- Possibly not enough idle and busy correlation?
  - Very low autocorrelation and correlation

KJC073



### *Power management* continued

---

- How to determine time-out value?
- Want to preset the percentage of affected requests

➡ This is an online percentile estimation problem

- Existing estimation methods are complex and/or require memory
  - First work by Jain and Chlamtac (CACM 1985)
  - Seminal work by Greenwald and Khanna (SIGMOD 2001)
- We are investigating simple methods that need no memory

KJC074



## *Power management* continued

---

- Online percentile estimation algorithm idea
- Increment and decrement around an *estimated percentile*
  - For each new value
    - » If value is larger than estimated, increase estimate
    - » If value is smaller than estimated, decrease estimate
- What should the increment size be?
  - *Ideally*, the difference of real percentile and its adjacent value
  - We estimate this...
    - » The estimate is "self correcting"

KJC075



## *Power management* continued

---

- Algorithm (shown here for median estimation)

```
while (values to read) do
  read value // Get a value

  // Adjust the increment size (delta)
  low = LOW * est_med
  high = HIGH * est_med
  if ((value >= low) && (value <= high)) count++
  if (count > 0) delta = (high - low) / count

  // Estimate the median
  if (value > est_med) est_med = est_med + delta
  else if (value < est_med) est_med = est_med - delta


  output est_med // Output estimated median
```

KJC076



*Power management* continued

- Results for idle periods for all 100 PCs concatenated
  - Estimating the median idle period time

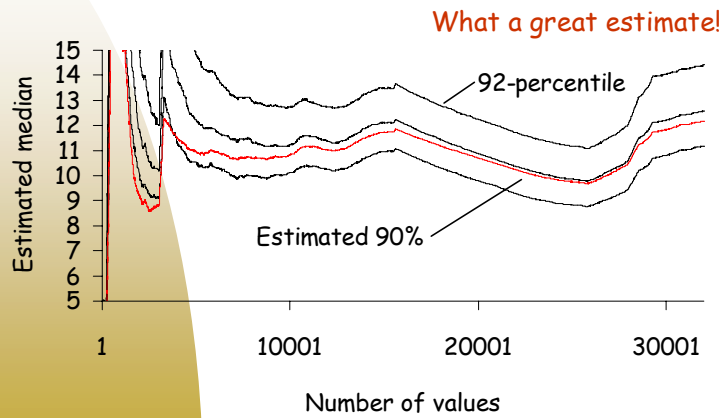


KJC077



*Power management* continued

- Results for idle periods for all 100 PCs concatenated
  - Estimating the 90th percentile idle period time




KJC078



### *Power management* continued

- Comparison to Jain et al. method for median estimation
  - Used implementation by Hoermann and Leydold (2000)



KJC079



### *Power management* continued

- More work to be done in percentile estimation
- This is a "classic" problem with much existing literature
- Compare against existing methods
- Can we bound the estimation error?

KJC080





## Topics

---

- Power management - what and why
- Power management at many levels
- A day in the life of a dormitory
- Power management for desktop computers
- A proxying Ethernet adapter
  - Development of a solution
- Summary and future directions

KJC081



## A proxying Ethernet adapter

---

- The problem:
  - Sleeping computers lose their network connectivity
    - Cannot respond to routine protocol messages
    - Cannot be connected to
- The solution:
  - Smarter Ethernet adapters or NICs
    - Proxy for routine protocols messages
    - Wake-up computer when needed

KJC082



## *A proxying Ethernet adapter* continued

---

- The website at LBNL for this work...



<http://eetd.lbl.gov/Controls/network>

KJC083



## *A proxying Ethernet adapter* continued

---

- Need compatibility with humans
  - Symbols
  - Meaning of *on*, *off*, *sleep*, *standby*, and *hibernate*
- LBNL is proposing...
  - *Off*, *Sleep*, and *On*
- Symbol for Sleep is a crescent moon



<http://www.lbl.gov/Science-Articles/Archive/EETD-simple-symbols.html>

KJC084



### An offensive symbol for sleep?

Is the crescent moon a religious symbol? This was an issue that was brought-up and featured in *Technology Review* magazine. Nine professors of Islamic studies were polled, eight said this would not be offensive.

<http://eetd.lbl.gov/Controls/publications/moon.pdf>

KJC085



### *A proxying Ethernet adapter* continued

---

- Proxying Ethernet adapter (NIC)...
- Move some protocol functions to the NIC
  - Including ARP, ICMP ping, DHCP
- "Intelligently" wake-up on existing protocol messages
  - For example, on a TCP SYN to an open port
- Two enabling technologies...

➡ Very low cost processors (we estimate \$10)

➡ Very fast wake-up of PCs (MRAM)

KJC086



### *A proxying Ethernet adapter* continued

---

- What size processor on Proxying Ethernet NIC?
- We are characterizing broadcast and other traffic
  - As seen by an idle host
- For a PC on a university network we saw...
  - About 4 packets per second
  - 1/3 of all packets are broadcast ARPs
  - 1/6 of all packets are routing related

➡ More work to be done here

KJC087



### *A proxying Ethernet adapter* continued

---

- Four phases to this project
- Phase #1 - Emulate a proxying NIC with a second PC
  - Phase #2 - Use an existing processor-full NIC in a single PC
    - Ethernet controller
  - Phase #3 - FPGA development of a proxying NIC
  - Phase #4 - Address power use of a NIC and of TCP/IP
    - A "Green TCP/IP" is the goal

➡ Partial phase #1 will soon be available via the web

➡ Phase #3 and #4 at the University of Florida

KJC088



## *A proxying Ethernet adapter* continued

---

- Mock-up of phase #1...



KJC089

**USF** UNIVERSITY OF  
SOUTH FLORIDA

## *A proxying Ethernet adapter* continued

---

- Why phase #4?
- Assume 100 million PCs in the USA
  - Estimate that PCs are on half the time
  - NIC in each PC as 2W

➔ Almost 1 TWh/yr just for NICs!

KJC090

**USF** UNIVERSITY OF  
SOUTH FLORIDA

## Topics

---

- Power management - what and why
- Power management at many levels
- A day in the life of a dormitory
- Power management for desktop computers
- A proxying Ethernet adapter
- Summary and future directions
  - Wrap-up of our tour
  - Future directions
  - Acknowledgements

KJC091



## Summary and future directions

---

- A grand tour of an important problem
  - We went down many side roads
- PC's are left on at night for no good reason
  - Wasted electricity = output of one nuke plant
  - It is a network problem!
- Many interesting analysis and evaluation problems
  - Characterization of idle times
  - Investigation of time-out schemes
  - Estimation of percentiles for fixed time-out values
  - Evaluation of effect of wake-up time on response time
- Interesting technology problems
  - How to proxy host functions in an adapter
  - How to recognize wake-up events
  - How to build "instant on" PCs

KJC092



### *Summary and future directions* continued

---

- Our short-term future directions...

- Move proxying NIC to implementation
  - Added cost of proxying recovered in few months
- Proposal submitted to NSF STI program in April 2003
  - With Alan George at University of Florida
- Work with LBNL, Microsoft, and Intel

KJC093



### *Summary and future directions* continued

---

- The long-term future...
- Need "Green" protocols
  - "Green TCP/IP"
  - "Green OSPF"
- Need "Green" network controllers
  - Focus on networked devices
    - » The Internet-connected doorbell is not far off
- Need to explore proxying on a larger scale
  - One router proxies for many sleeping hosts?

KJC094



*Summary and future directions* continued

---

- Beyond the desktop...

**“By 2010, 95% of Internet-connected devices will NOT be computers.  
How will they be connected?”**

- Ipsil Incorporated (2003)

- The problem is only becoming worse, not better!

KJC095



*Summary and future directions* continued

---

- What we want in the future...

Always on without being always (fully powered) on

KJC096






## Summary and future directions continued

- Or, do we want them on at all?

**“Energy wasted by computers and monitors costs public school districts in this country more than \$145 million dollars each year.”**

- EPA (2003)



*"Our NCLB plan calls for LITERACY and nothing but LITERACY. So we've been conserving energy by turning off the screensavers."*

KJC097

## Acknowledgements

- Colleagues and students:
  - Bruce Nordman (LBNL)
  - Alan George (UF)
  - Mamatha Kumar (MS student, USF)
  - Chamara Gunaratne (PhD student, USF)

**"If a king or emperor says "we did it", it means "I did it." When a professor says "we did it" it means "my graduate student did it."**

- V. E. Bondybey

KJC098

**USF** UNIVERSITY OF  
SOUTH FLORIDA