Green Computing:
Hype or Hard Truth?

Yuvraj Agarwal
yuvraj@cs.ucsd.edu

http://mesl.ucsd.edu/yuvraj

Jan 25th, 2010

Outline:

- Where does the energy go in buildings?
  - Estimated > 70% total energy use in the US
  - Factors affecting building energy use
- Does IT power really matter?
  - Air-conditioning ? Lighting?
- Can we actually reduce our energy footprint?
  - Somniloquy, Sleep-Servers and beyond!

UCSD Campus as a Testbed

- 1200 acres, borders the Pacific Ocean
  - 450+ buildings
- 45,000 people = 29000 students + staff
  - 10,000 on campus residents
- Power Demand around 45MW!
  - USD (13MW), City of San Diego (~50MW)

UCSD is like a small town with its own utility

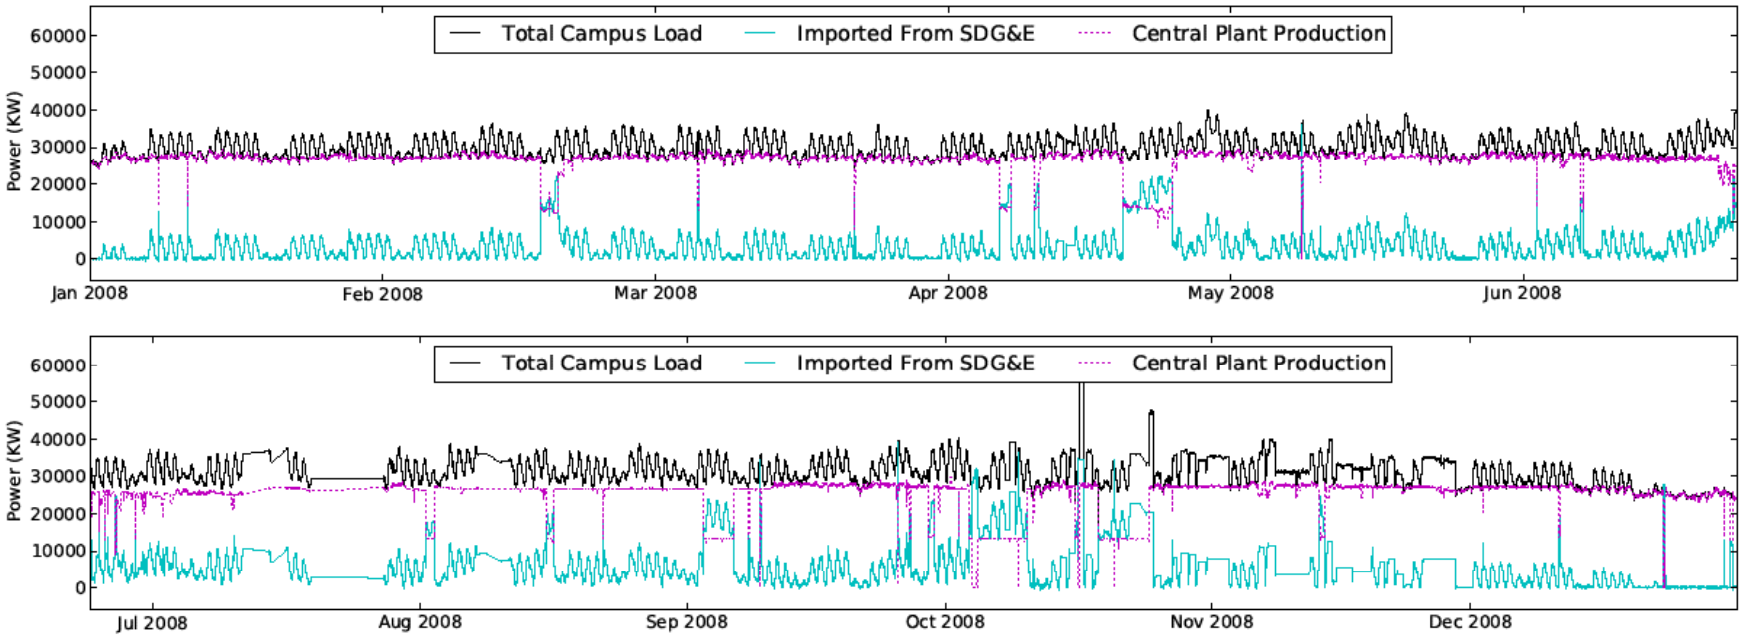

- We co-generates most (up to 80%) of our power
  - Import the rest from SDGE (5MW to 20MW)
  - Participate in demand response (DR)

Use Modality of Buildings: Macro Scale

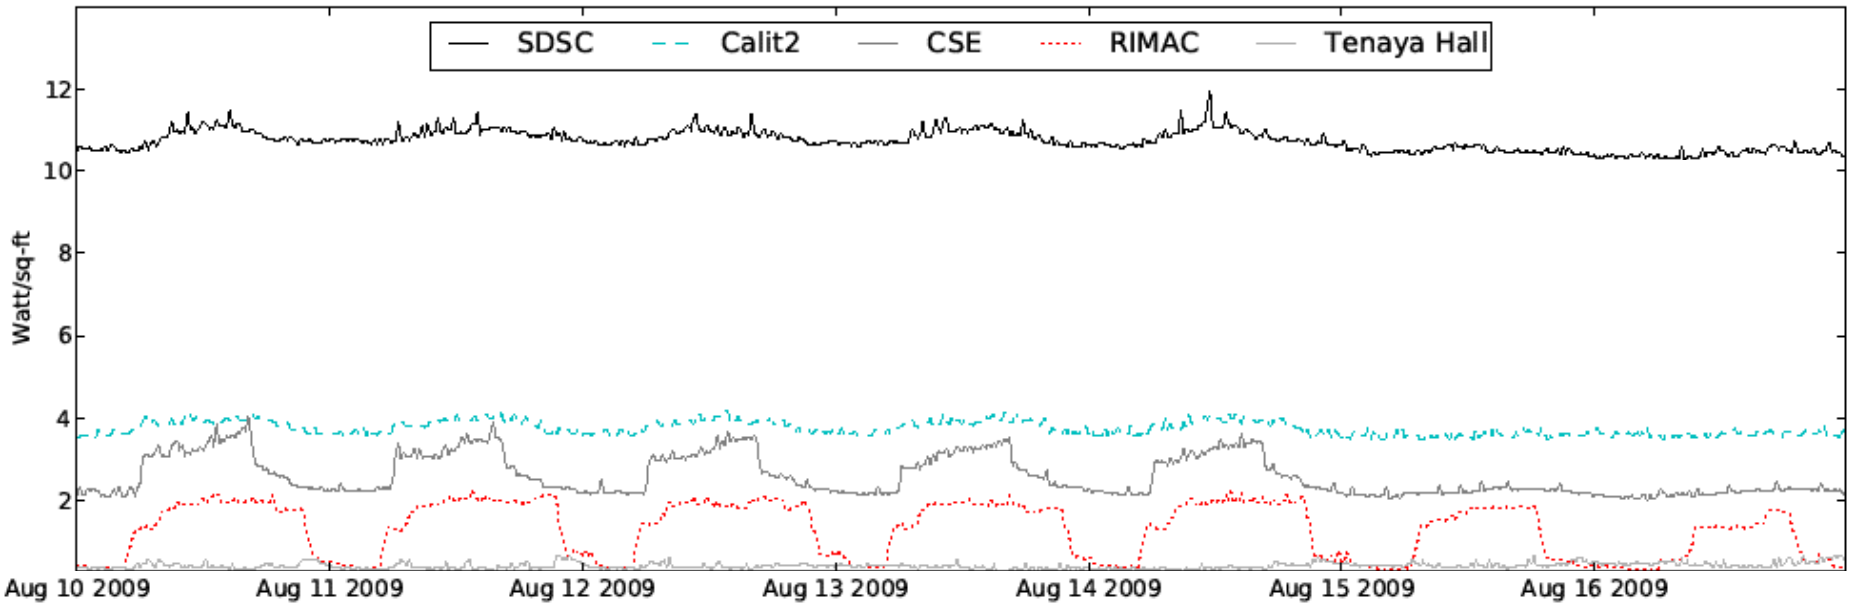

- Compare different builds across campus
  - Residential, “mixed-use”, highly IT centric
- As IT load increases, so does watts/sq-ft!

Mixed-Use Buildings: CSE Department!

- Large mixed use building:
  - Nominally 600 people, increases by 1.7x
  - 750+ desktops (+200 undergrad labs) + servers
  - Closed loop and zonal climate control
  - Served by campus chilled and hot water loop

- Look at trends: yearly, weekly, daily?

CSE Building Energy Use Trends:

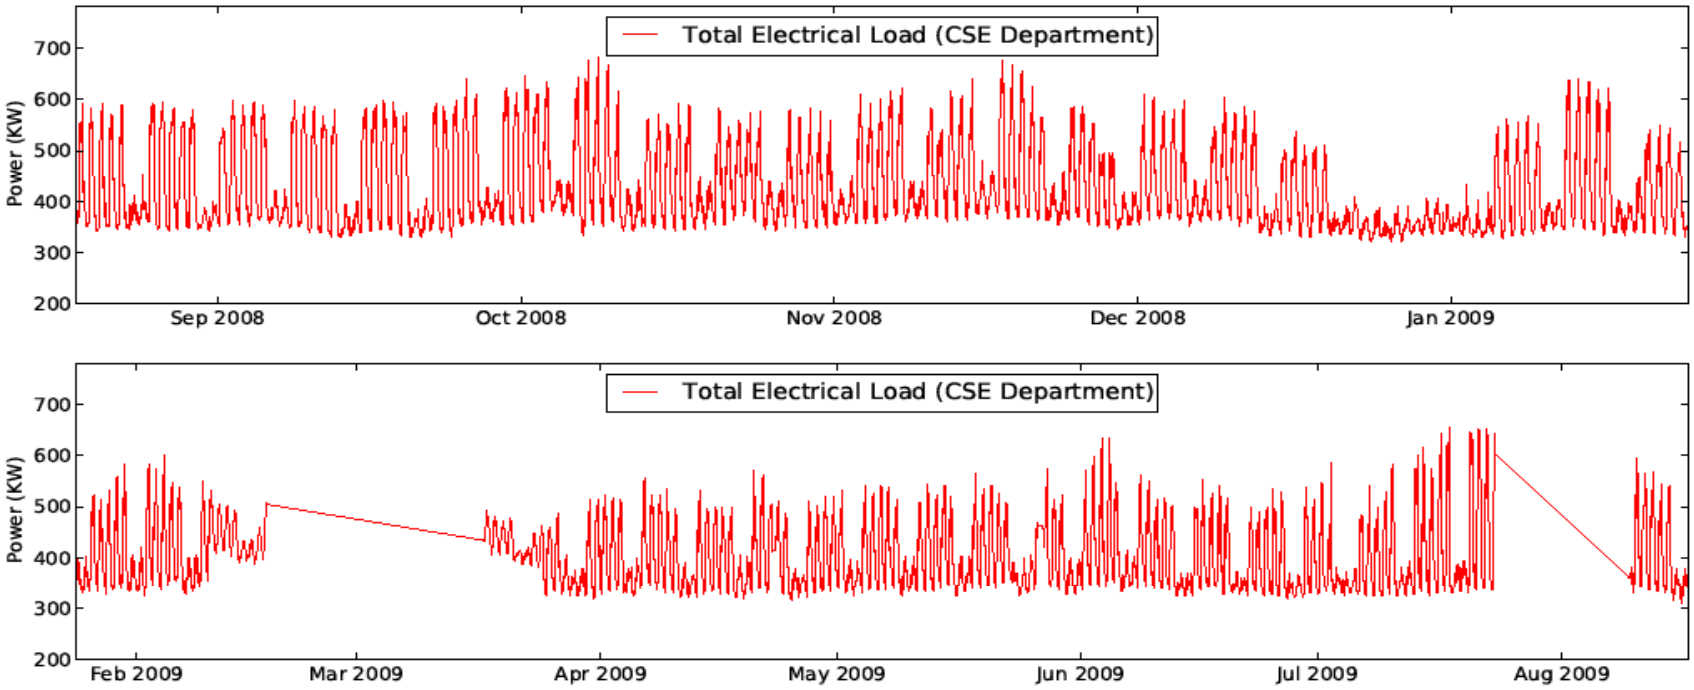

- Yearly variation not much (across seasons)
  - "Base load" of CSE is over 325KW all year!
  - Seasons don't have much effect, neither do breaks

CSE Building: Weekly variation

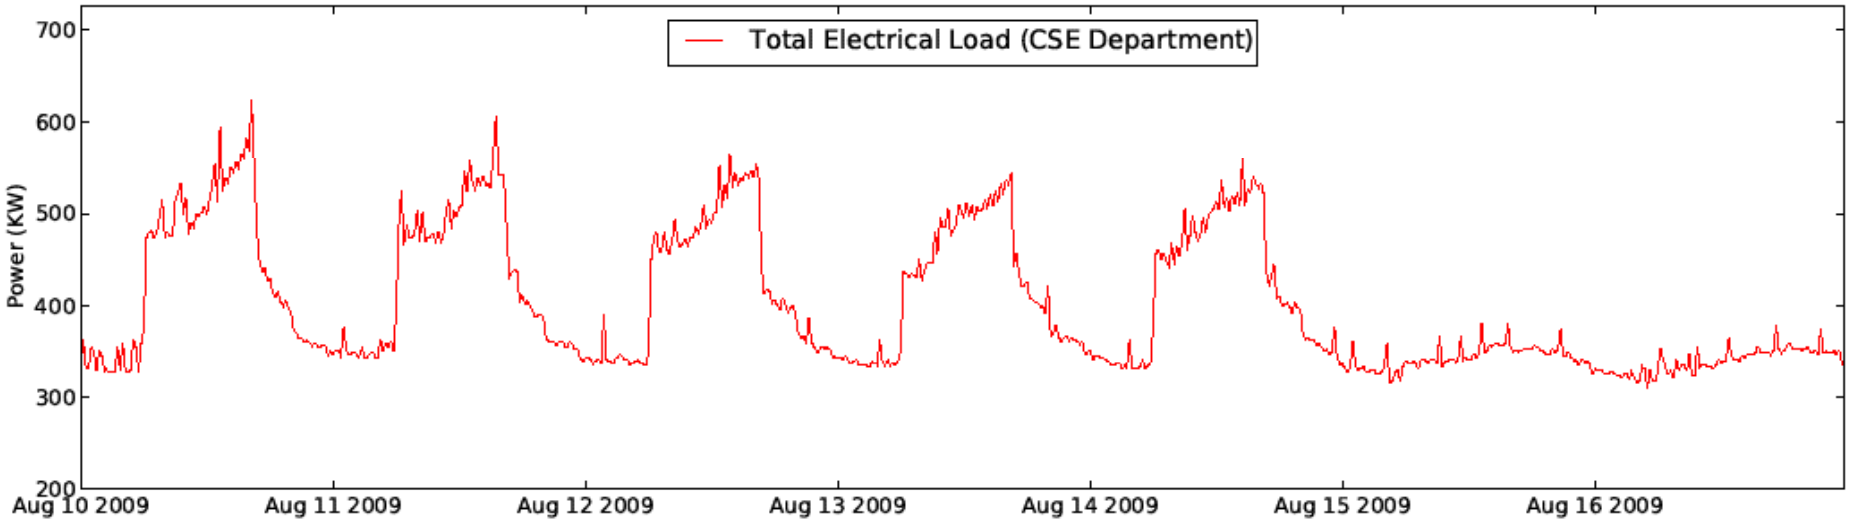

- Weekly Variations:
  - Effect of Weekends is clearly observable
- Dynamic variation: 325KW to 600KW at peaks
  - "Base load" still remains at 325KW!

Further Breakdown of CSE base loads

- Instrument the CSE department
  - Sensors to measure power
  - Breakdown based on :
    - Subsystems: Air Conditioning, Lighting, …
    - Individual floors, circuits (Macro Scale)
- Measuring individual PCs
  - Power/Energy Usage (Micro Scale)
- Objectives:
  - Build better energy models
  - *Measure* effectiveness of energy management schemes

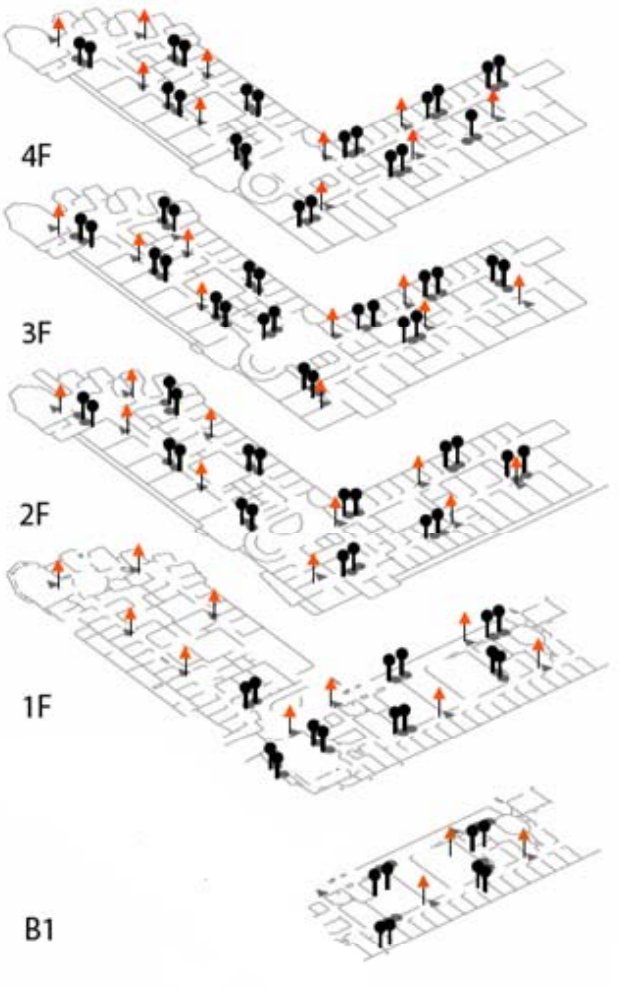

CSE Building Floorplan

Contributors of the CSE base load:

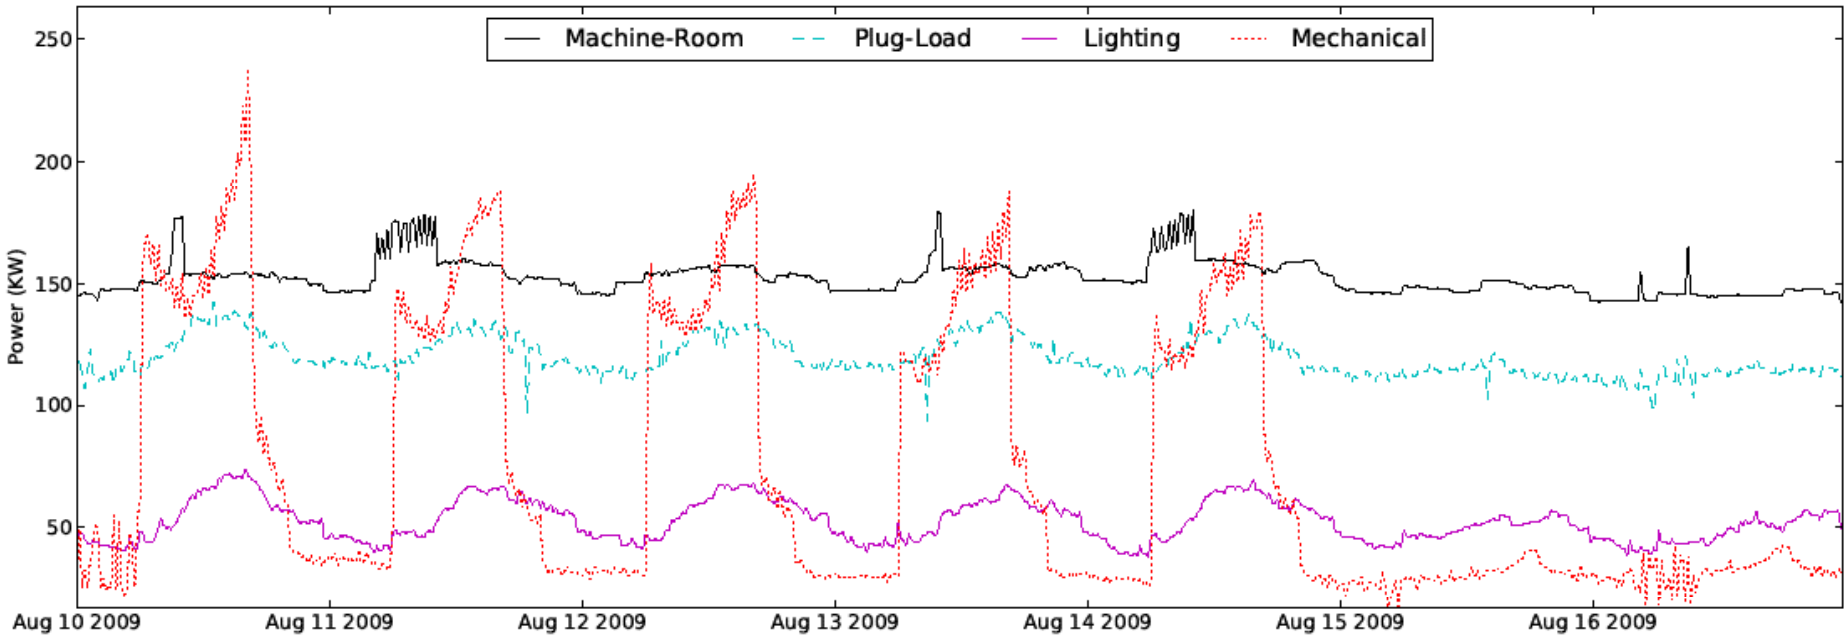

- 16+ individual meters, grouped into 4 categories
  - Machine Room (Servers + HVAC)
  - Plug-Loads: Mostly PCs (we are a CSE department!)
  - Lighting: Includes everything, including emergency
  - Mechanical: Air-handlers, chilled water pumps, etc

Contributors of the CSE base load:

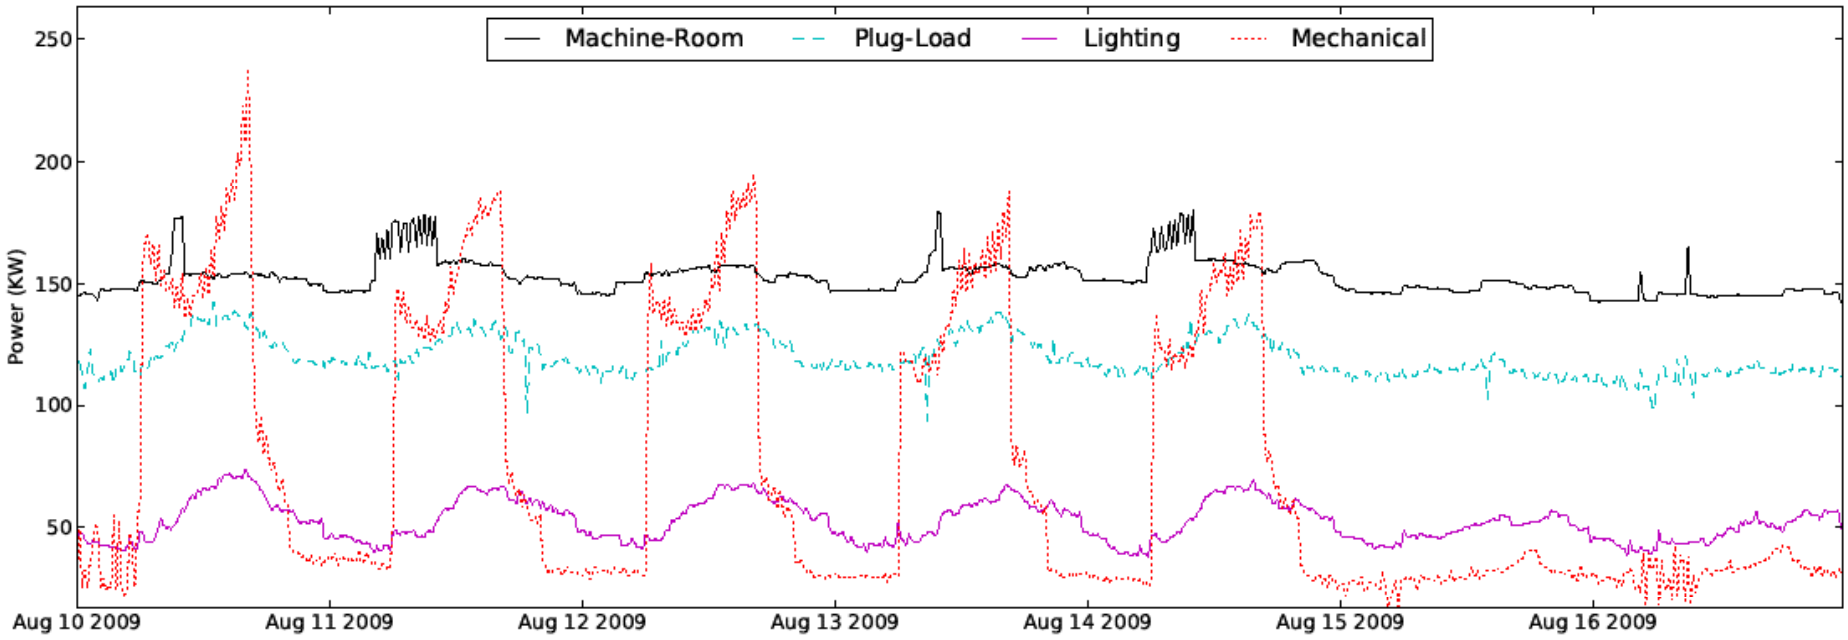

- IT loads account for **50%** (peak) to **80%** (off-peak)!
  - Includes machine room + Plug Loads
- Most IT equipment not powered off or put to S3
- Largest dynamic variation in "mechanical loads"

Outline:

- Where does the energy go in buildings?
  - Estimated > 70% total energy use in the US
  - Factors affecting building energy use
- Does IT power really matter?
  - Air-conditioning ? Lighting?
- Can we actually reduce our energy footprint?
  - Somniloquy, Sleep-Servers and beyond!

IT Equipment Consumes Significant Power

- Yet, shutdown opportunities are rarely used
- Studies and our own data shows:
  - 67% of office PCs are left on after work hours
    - "Sleep" modes used in less than 4% of these PCs! [1]
  - Home PCs are left on for 34% of the time
    - 50% of the time they are not being used
- Confirmed by our measurements: CSE@UCSD
  - 600+ desktops left always on (total=700+)
    - @150W each ➔ 100kW (25% of total energy bill)
    - Used network measurements (ARPs + ICMPs)

[1] J. Roberson et al. "After-hours Power Status of Office Equipment and Energy use of Miscellaneous Plug-load Equipment. *Lawrence Berkeley National Laboratory, Berkeley, California. Report# LBNL-53729-Revised, 2004*

Saving Power Runs into Usability

- Reasons why users do not switch off their PCs
  - – Maintain state: desktop and applications preferences
  - ✖ – Occasional access
    - Remote desktop/SSH, accessing files
    - Administrative: updates, patches, backups
  - ✖ – Active applications running
    - Maintaining presence: e.g. incoming Skype call, IM
    - Long running applications: Web downloads, BitTorrent

Cannot be handled by low-power modes (e.g. Sleep, Hibernate)

Power Management vs. Use Models

- Current design trends in power management:
  - Hosts (PCs): either *Awake* (Active) or *Sleep* (Inactive)
    - Power consumed when Awake = 100X power in Sleep!
  - Network: Assumes hosts are always "Connected" (Awake)

- What users really want:
  - Provide functionality of an *Awake* (active) host…

    ….While consuming power as if in *Sleep* mode
  - Resume host to Awake mode only if needed

Change the fundamental distinction between Sleep and Active states…

Augment the Network Interface of PCs

- Objective: Make PCs responsive even when asleep
  - Maintain availability across the entire protocol stack
    - E.g. ARP(layer 2), ICMP(layer 3), SSH (Application layer)
  - Without making changes to the infrastructure or user behavior

Active State : >140 W
Idle State : 100 W
Sleep State : 1.2 W

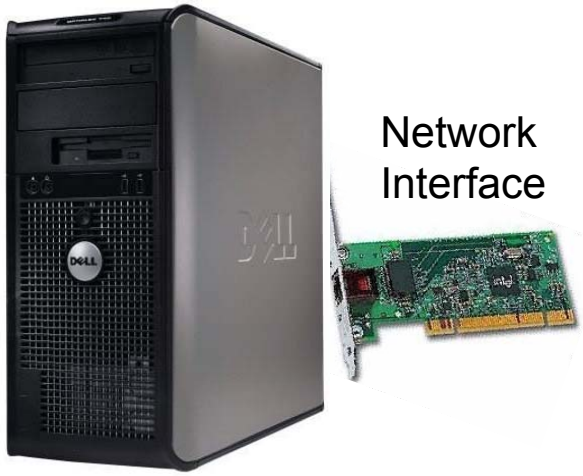

Requirements:
- Functionally similar - masquerade as the host
- Much lower power

Secondary Processor
(Power in active state ~1W)

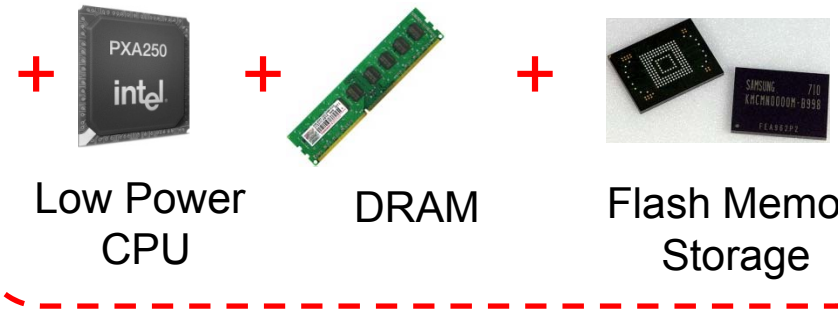

*Somniloquy**: PCs that Talk in their Sleep

- Augment network interfaces:
  - Add a separate power domain
    - Powered on when host is asleep
    - Processor + Memory + Flash Storage + Network stack
  - Same MAC/IP Address
- Wake up Host when needed
  - E.g. incoming connection
- Handle some applications while PC remains asleep
  - Using “application stubs”

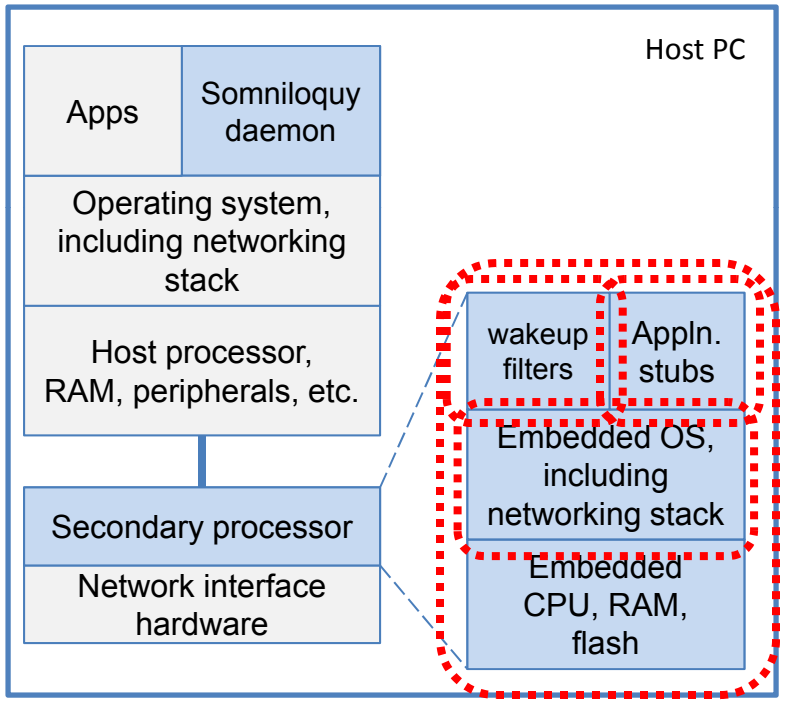

*som-nilo-quy (noun): The act or habit of talking while asleep

Prototype

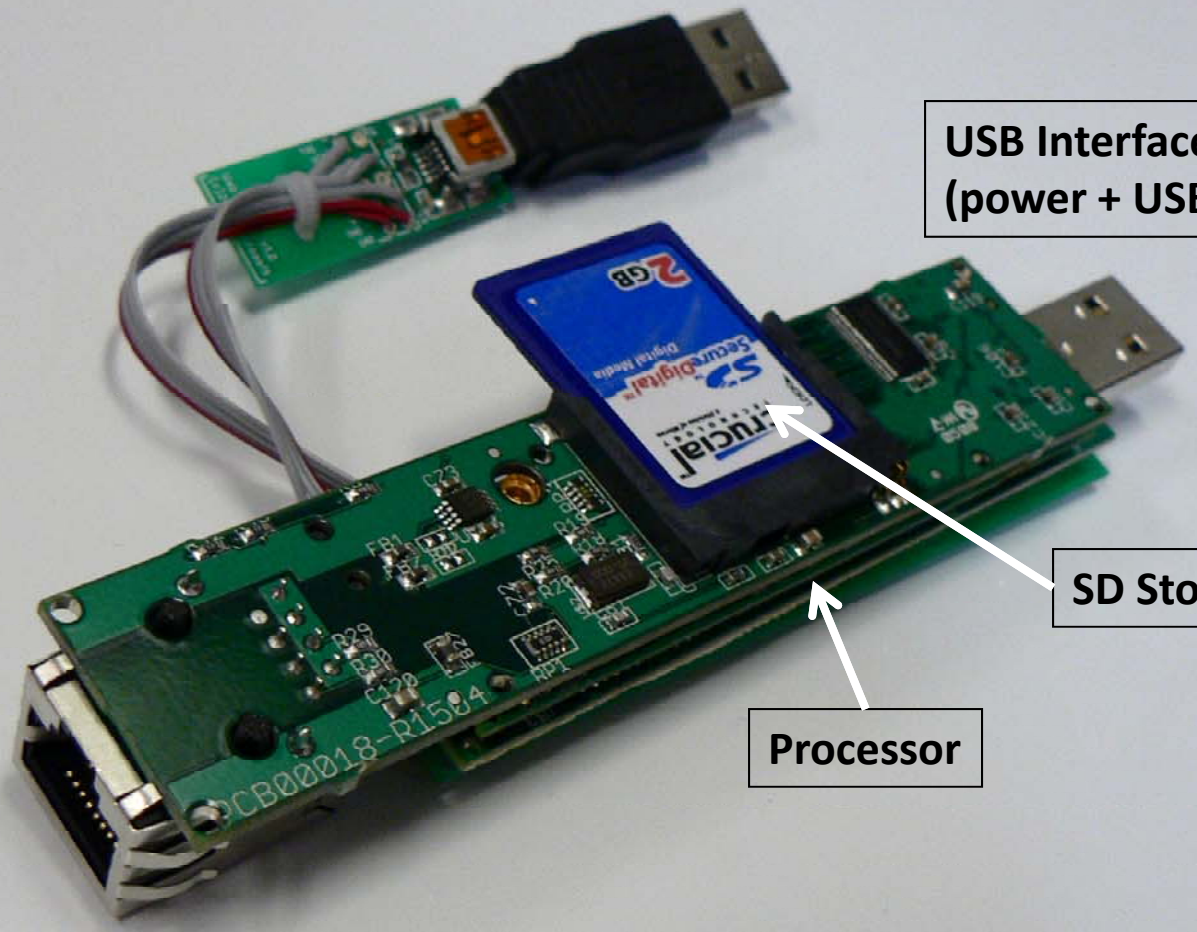
USB Interface
(Wake up Host + Status + Debug)
USB Interface
(power + USBNet)
SD Storage
Processor
100Mbps Ethernet Interface

Desktops: Power Savings with Somniloquy

State	Power
Normal Idle State	102.1W
Lowest CPU frequency	97.4W
Disable Multiple cores	93.1W
"Base Power"	93.1W
Suspend state (S3)	**1.2W**

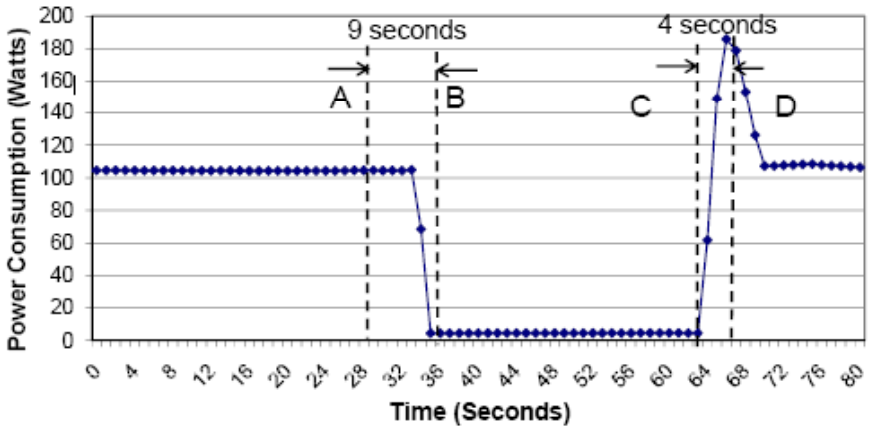

Dell Optiplex 745 Power Consumption and transitions between states

Using Somniloquy:

- Power drops from >102W to <3W
- Assuming a 45 hour work week
  - 620kWh saved per year
  - US $56 savings (at 9cents/KWHr)

Laptops: Extends Battery Lifetime

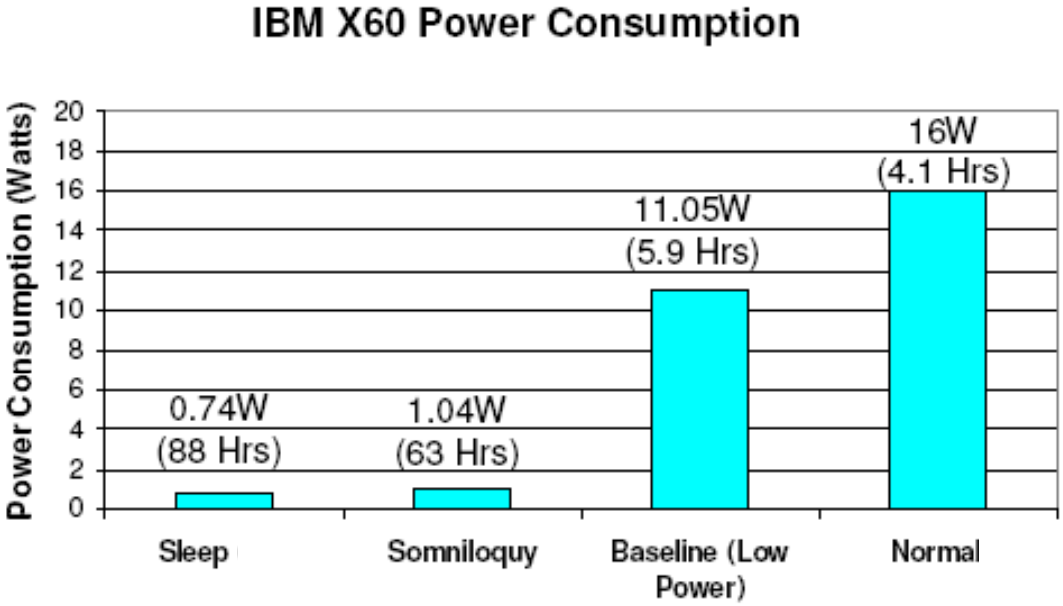

Using Somniloquy:

- Power drops from >11W to 1W,
  - Battery life increases from <6 hours to >60 hours
- Provides functionality of the "Baseline" state
  - Power consumption similar to "Sleep" state

Stateful Application: Energy Savings

- Web download "stub" on the gumstix
  - 200MB flash, download when Desktop PC is asleep
  - Wake up PC to upload data whenever needed

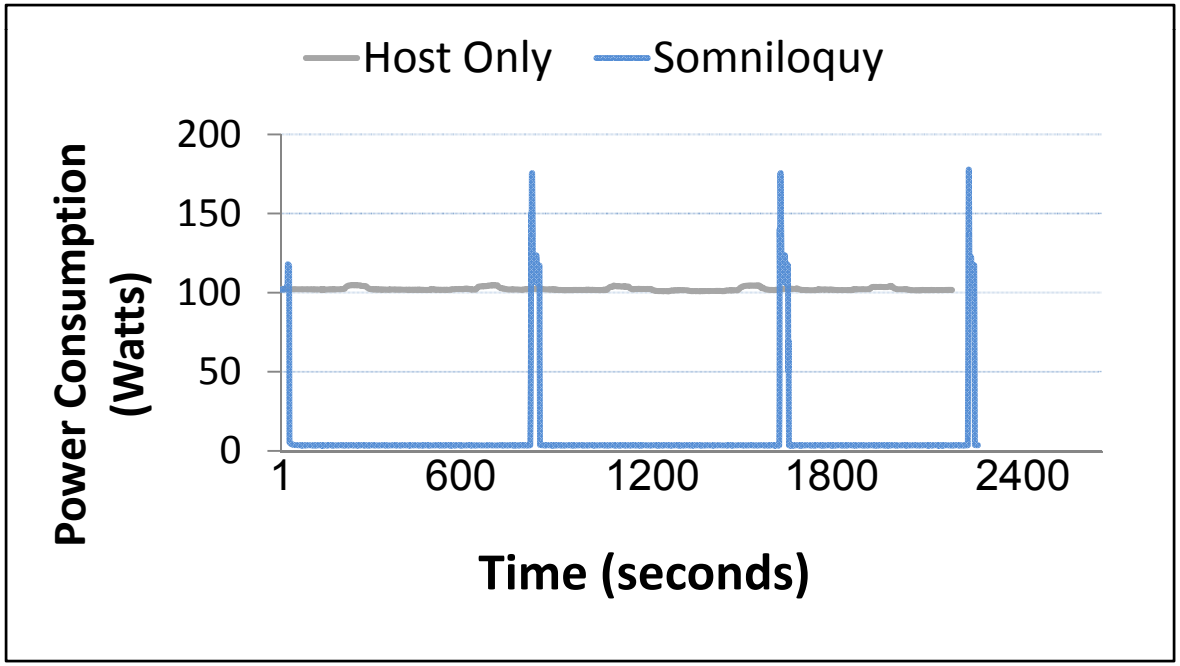

  - 92% less energy than using the host PC for download

Deploying Somniloquy at UCSD

- ***Enterprise*** environments:
  - Adding Somniloquy hardware to desktop PCs costly
  - Management, maintenance costs

- "Somniloquy" without adding hardware?
  - Yes – It is possible….
  - However, need some help from the network
  - Stubs: Can we generate them automatically?

- Individual desktops/Laptops: *Somniloquy* needed

Enter Sleep-Servers:

- For each “asleep” host PC
  - Instantiate a minimal *image* on a Sleep-Server
- When host transitions to Sleep
  - Its *image* on the Sleep Server can:
    - Maintain network presence
    - Watch for triggers to wake up
    - Execute application stubs

Sleep Servers: Energy Savings

- Server class machine:
  - Can support hundreds of desktops (200+)
  - One server (~300Watts), *each* desktop (~150W)
  - Single point of management
  - No hardware additions to host PCs!
  - > 100x Energy Savings, easily deployable in enterprises
- UCSD / CSE : 650 desktops PCs
  - Potential savings: $37,000/year (8 hours/day + weekends)

Sleep-Servers Status

- Currently we have:
  - An implementation for Windows, Linux host PCs
  - UCSD patent pending for our architecture
- Further Steps:
  - Scale up to hundreds of host PCs on single server
  - Deploy widely within CSE building (Need you!)
  - Implement for other OSes (MAC, …)
  - What about servers in the machine room?

Energy Dashboard: Detailed Accounting

- Measuring individual plug loads (30 meters)
  - Commercial *WattsUP* meters (~150$)
  - Building our own wireless power meter (< 20$)
    - Measurement and Management
    - Objective: **Educate** users about their power usage
  - Easy custom data logging interface to our servers

- Measuring aggregate power loads (15 meters)
  - Commercial Schneider Meters (~3000$)
  - Interfacing a challenge: proprietary interfaces

Energy Dashboard: Demo!

http://power.ucsd.edu

Thank you!

Questions?

http://mesl.ucsd.edu/yuvraj

* Come find me (CSE 2116) if you want to be a tester for Sleep-Servers!
*** Also come find me if you want to work on cool projects.**

Wake-on-LAN – Limitations

- Runs into usability
  - Need MAC address, modifications at remote end
- Requires BIOS/Operating System support
- Requires router infrastructure support
  - Cannot go across subnets, (subnet broadcast)
- Granularity is very coarse
  - Too few or too many wakeups
- Very limited in functionality
  - No support for application stubs