Measuring and analyzing the energy use of enterprise computing systems

Paper written by:
Maria Kazandjievaa, Brandon Hellera, Omprakash Gnawalib, Philip Levisa, Christos Kozyrakisa, Stanford University and University of Houston

Presented by:
Simin Nadjm-Tehrani, Federica Uccello, Willem Meijer
Linköping University, Sweden

LINKÖPING UNIVERSITY

Motivation

- Researchers need to use per device power data in models or calculations for further conclusions
  - typically from earlier papers based on single measurement on some device

- Enterprises pay high power bills but no idea what is the main contributor to the energy use
  - education and office computing devices were 2% of US electricity consumption

Basic idea

“ … the more data and the better understanding of methodology we have, the better resulting insights and solutions will be.”

Contributions of this paper

- Powernet – a multi-year study at Stanford university CS building
  - Data over two years with 250 devices made available
- Methodology for creating a power measurement infrastructure and how to use it
  - How frequent, which devices, which metadata helps to give a big picture from a sample of measurements?

Three of the five floors

- Spread of devices monitored

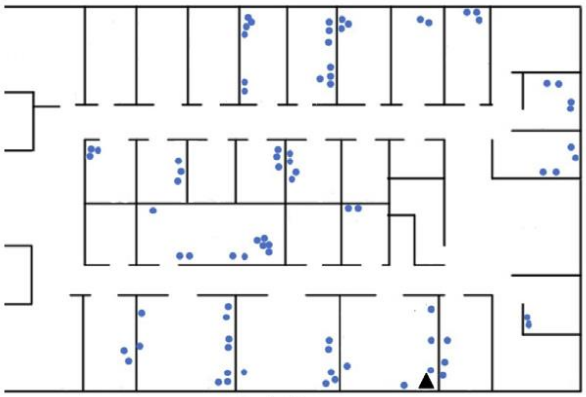
2nd floor

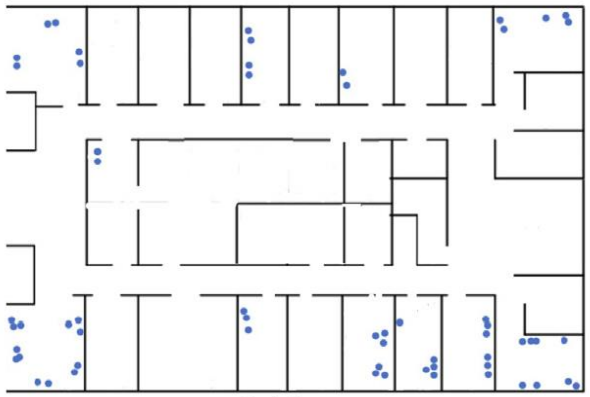
3rd floor

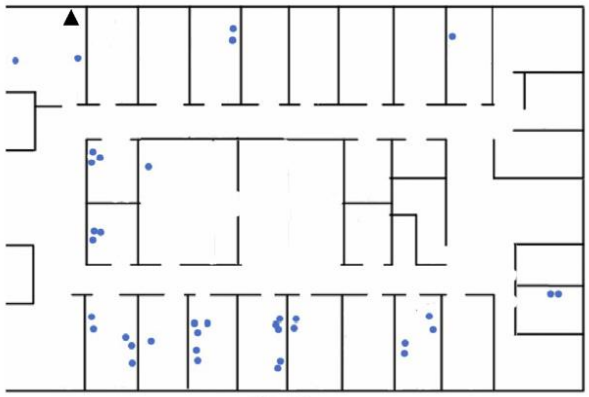
4th floor

Metering infrastructure

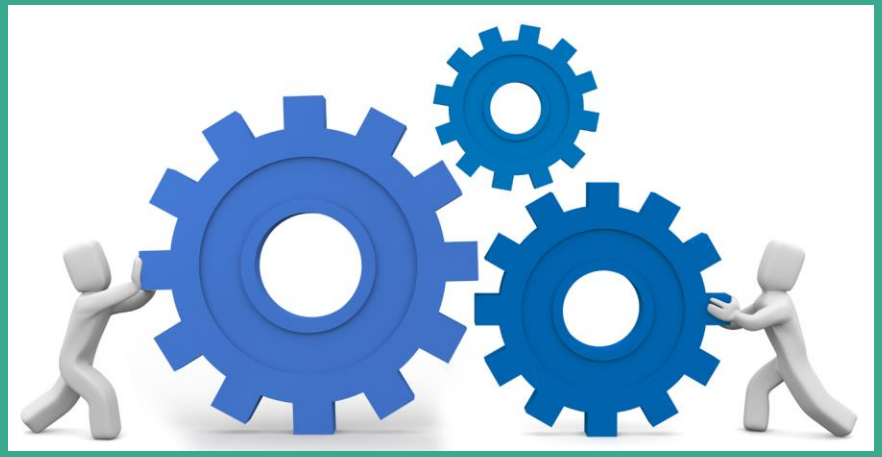

Initially: off-the-shelf meters

- Watts Up.Net power meters with built in Ethernet support to send the measurement per second
  - Difficulties in deployment
  - Bugs and maintenance
  - Proprietory firmware

Similar device type

Designed own meters based on motes

- Wireless, 802.15.4
- TinyOS
- Ad-hoc networking
- Configured to also send device usage data

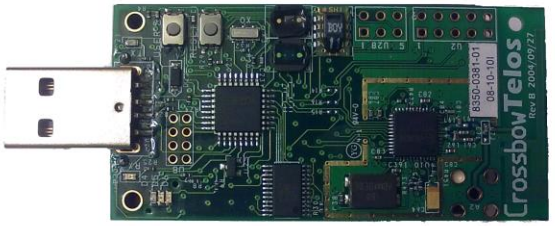

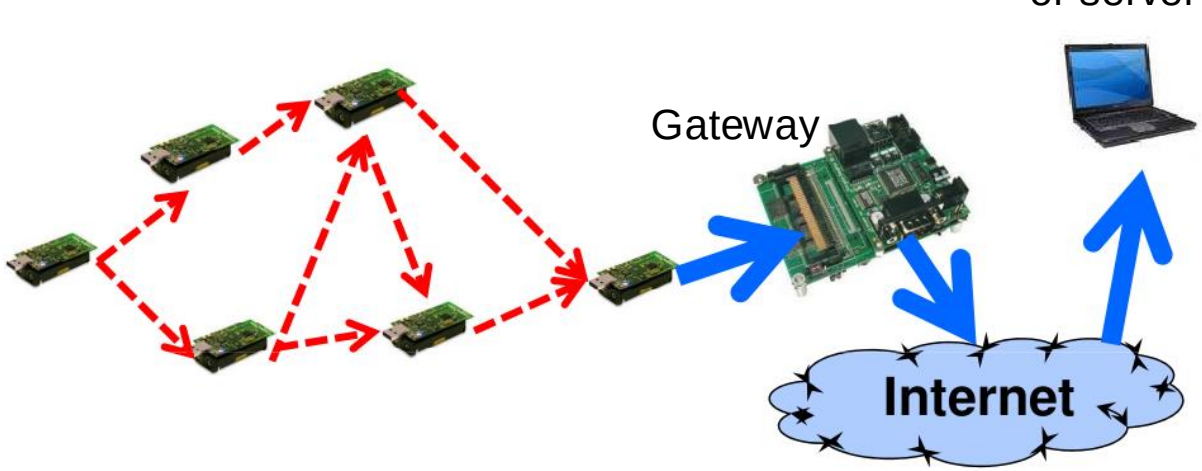

Image from: The Basics of Wireless Sensor Networking and its Applications
http://www.ida.liu.se/~rtslab/courses/wsn/Basics.pdf

Device types metered

Device type	Count
Desktop	75
Monitor	70
Laptop	28
Network switch	27
Printer	15
Server	36
Thin clients	12
Misc	3
Total	266

Usage values collected in the dataset

Sensing type	Num. datapoints
Power data	10 billion
CPU percent	400 million
User processes	2 billion
Network traffic	10 million

- 1 GB data every day

Bird's eye view

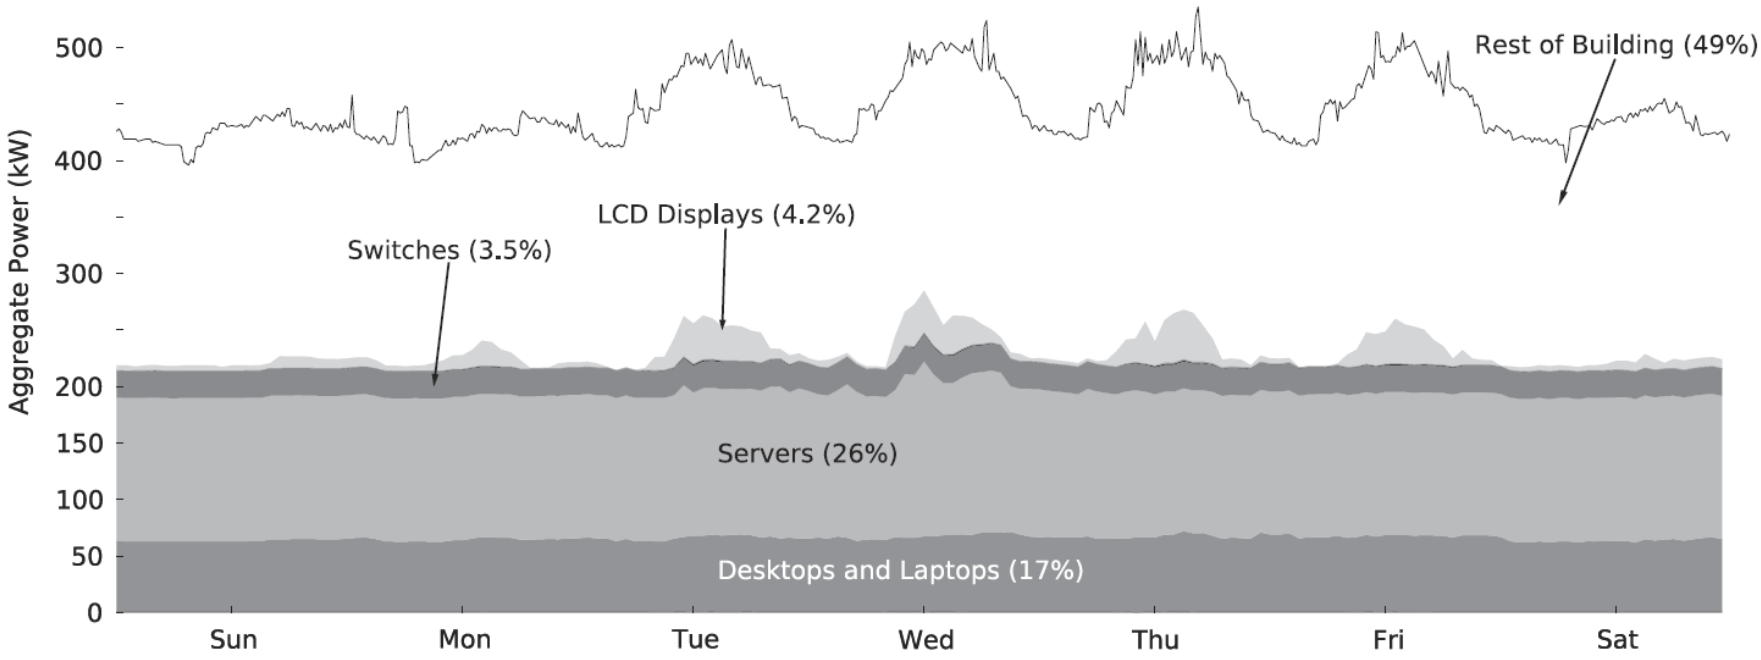

- Extrapolated power during a week: ~50% of total building power!

Individual device types

Changing monitor brightness

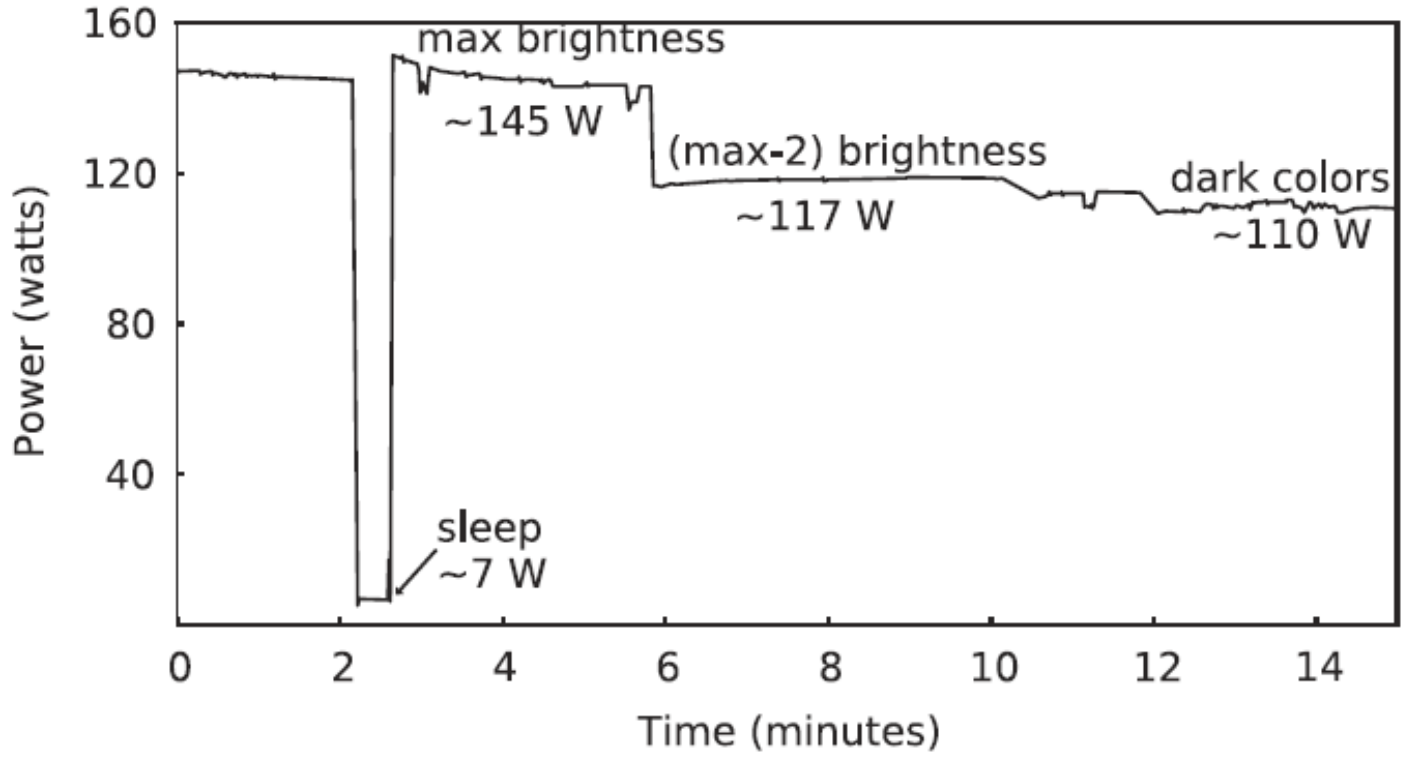

Network switches

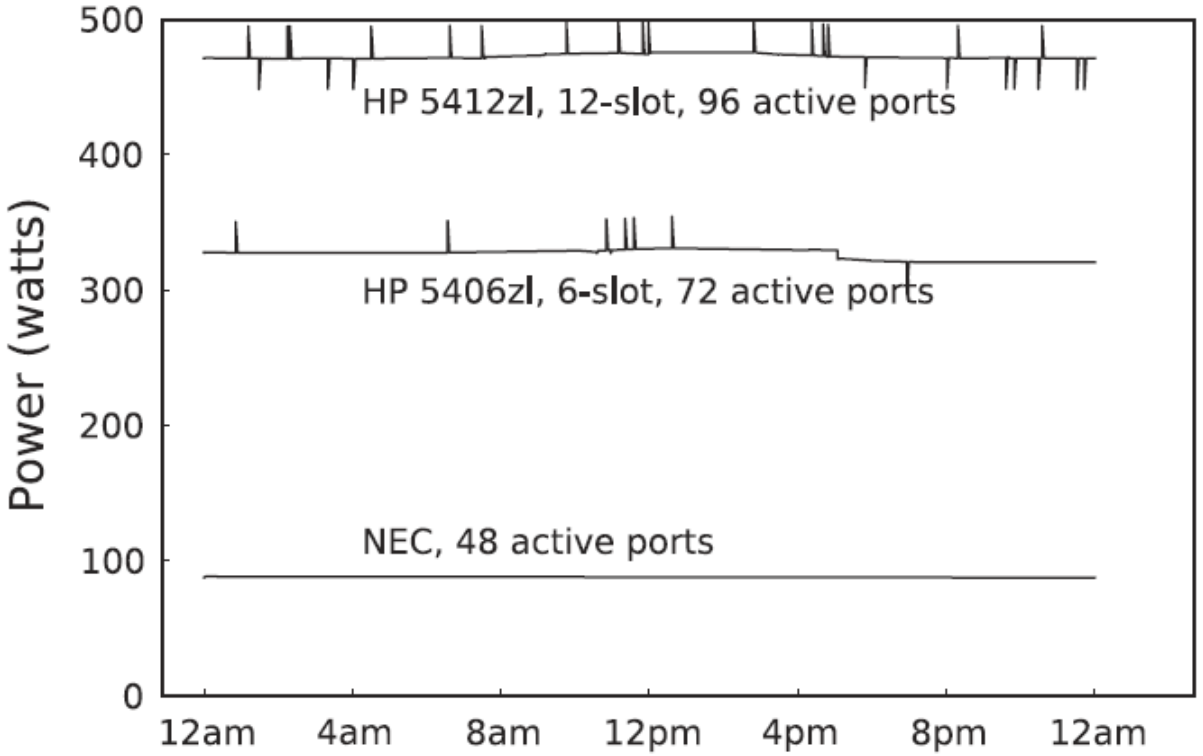

Idle power for three similar switches

Servers

- Less diversity among servers (compared to PCs that could be laptop, low-range, high-range)
- Still, a standard 1U rackmount could draw anything between 95W and 275W!
- Why it was important to have a high coverage/spreading of measurements among units
  - average power 233W

Power data alone is not enough!

Potential for power saving?

Idle power vs. active power

- In general, PCs were under-utilised
  - below 30% for 95% of the time
  - The most power hungry PC for an admin staff drawing 150W was used 3.1% for 95% of the time
- They were kept on all the time!
- Similar story for NW switches…

Why use a 96 1-Gbs active HP switch when the 200Mbps observed traffic could be carried by an edge switch?

Research implications

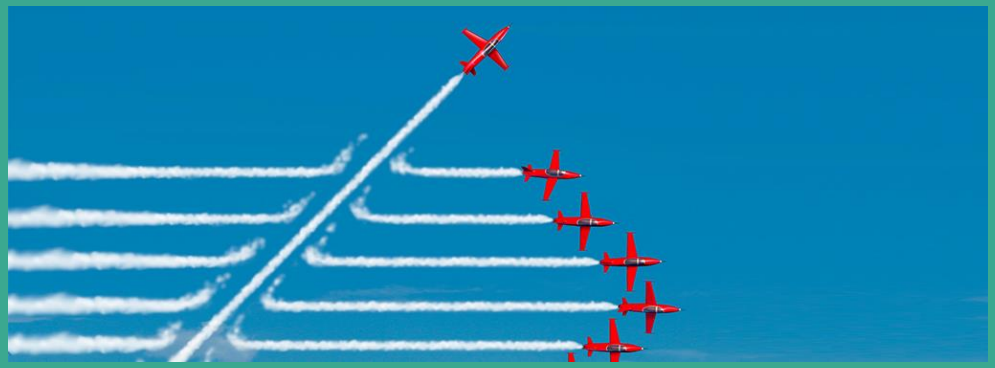

Lessons learnt

- High resolution for measurements strains the measurement infrastructure
  - But may reveal anomalies that do not show up in averages
- Comparing power draw even for identical models may show diversity
  - Need utilisation values as additional data
- More data and frequent samples gives closer values to ground truth

Summing up the method

- Identify which devices show large diversity before deploying metering infrastructure
  - This decides the placement of sensors
- Decide the frequency and duration of measurements
  - When need to choose, more devices is better than longer interval
- Make sure accurate metadata is available
  - IT dept should know how many devices of any sort are being used in buildings

Classification

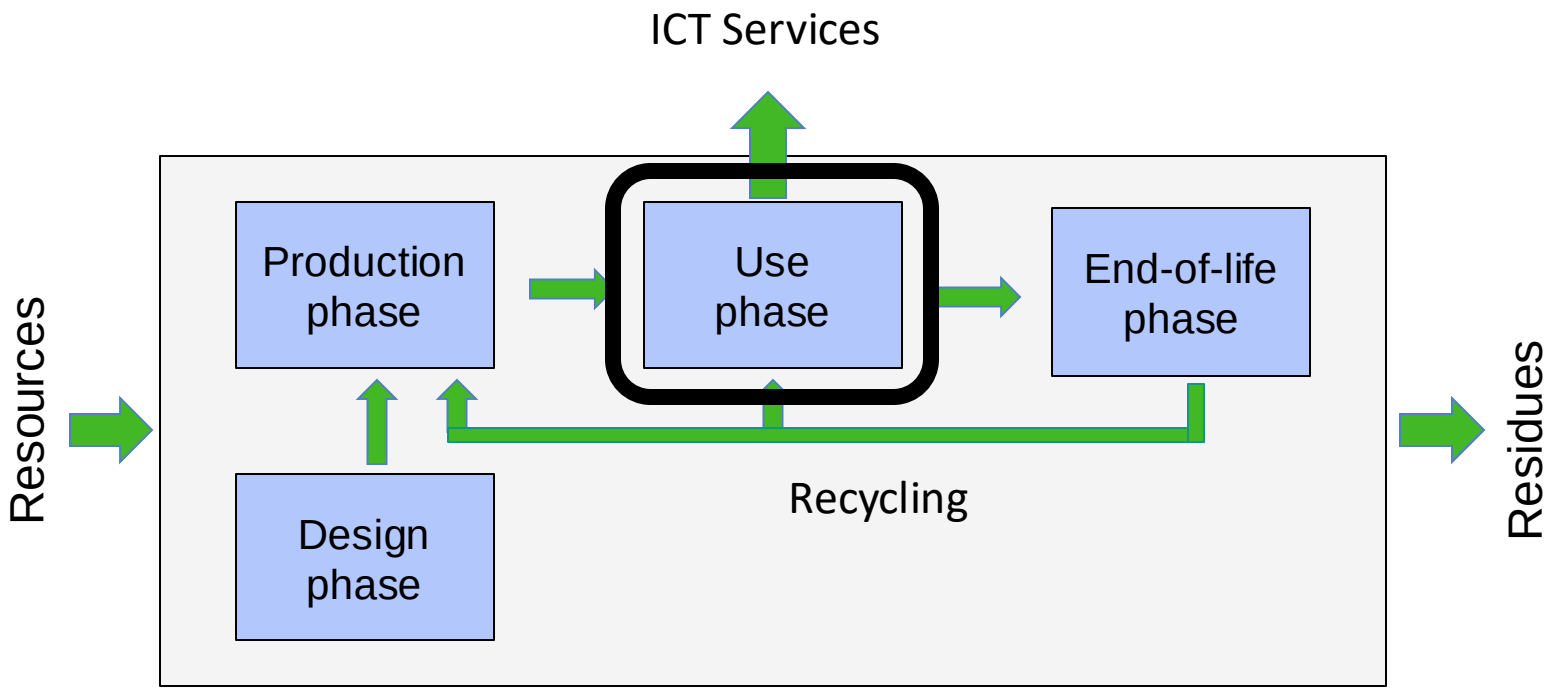

Questions?

LINKÖPING UNIVERSITY