On the Science of Power Management: Encouraging Sustainability R&D

Erez Zadok

Dept. of Computer Science

Stony Brook University

http://www.fsl.cs.sunysb.edu/

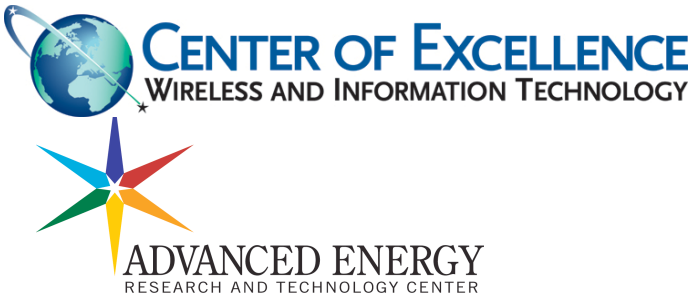

STONY BROOK
STATE UNIVERSITY OF NEW YORK

NSF SciPM Workshop 2009

- Science of Power Management
- **http://scipm.cs.vt.edu/**
- Bring multi-disciplinary people:
  - Theory, practice, industry, academia, government.
  - Identify, prioritize, and recommend promising research directions
  - Over 80 participants
- 7 key findings

1: Observe Systems

- Simply measure and analyze what systems are doing
- At all levels from chip, to system, to data center, and beyond
- Disseminate results widely
- Encourage prototyping

- Required for modeling and optimization

2: Develop Metrics

- How can you demonstrate benefits?
- Need for useful, clear metrics
  - ops/sec, total watts/joules, ops/watt
  - ops/watt-second?
  - dollars?
- How to account for long term effects?
  - e-waste, carbon footprints
  - longer hardware lifetimes, IT manpower costs

3: Models

- Systems too complex today
- Models help simplify and understand
  - Make simulations useful
- Challenge: model the most significant factors
  - After you observe and develop metrics
- Need for models at all levels:
  - Hardware and software
  - Chip, system, data center, Internet wide

4: Optimization

- Too many “point” solutions
  - Short term incremental benefits
  - How useful to others?
- Systems are complex
  - Multi dimensional: power, performance, reliability, security, usability, ...
  - Multi-variate: lots of h/w and s/w knobs to tweak
  - Non-linear: e.g., power/perf. can go together or opposite

4: Optimization (cont.)

- Need rigorous analytical techniques
  - Algorithms
  - Control theory
- Global view optimization
  - Across all layers of s/w and h/w

5: Education

- Few IT classes
- Little education on power management
  - Special grad topics
- Need undergrad curriculum
  - Brought down to core topics
- For now: integrate into existing classes
- Example: security education in 1995 vs. 2010?
  - Cannot wait 15 years...

6: Develop a Scientific Community

- Cross all sub-disciplines of computer science
- Multi-disciplinary interactions
- Need more cross-disciplinary workshops and conferences

- E.g., NSF sponsorship of student travel for SustainIT'10 (thanks!)

7: Beyond IT

- Help beyond just computing and data centers
- Need lots of software, techniques, and tools for example:
  - Smart buildings
  - Smart power grids
  - Automated transportation systems
  - Tele-presence
  - Climate and weather modeling

Every Great Journey Starts with...

- ... peeling onion (layers)
- Develop optimal software
  - Applications, middleware, OSs, clusters
- but first: understand interactions of hardware, software, and workloads of complex distributed systems
- but first: understand simple clusters
- but first: understand client-server systems
- but first: understand standalone systems
- but first: understand individual components

Survey 1: Can Compression Help?

- Idea: if you compress all data, less to write and trasmit, but costs in CPU
- Studied diff. hardware, compression tools/algorithms, and data types
- Conclusions [ACM SYSTOR 2009]
  - Improve energy/perf. by 10–40% at best
  - Worst case hurt energy/perf by 10–100x!
  - Heavily depends on hardware, software
  - Depends on workloads:
    - Data type, read to write ratios

Survey 2: Workload Effects on Servers

- Studied different server machines
  - Try different file system configurations
  - Workloads: Web, mail, database, etc.
- Found large perf/energy variations:
  - From 6–8% to 9 times better!
  - Small one-time reconfigurations needed
- Depends on exact hardware, software, configuration, and workloads
- Plug: FAST'10 paper, Friday 2/26 11am

Survey 3: Workload Effects on Client/ Server Network File Systems

- NFSv4 standard and interoperable, but
  - Different implementations
- Studying mix of NFS clients and servers
  - BSD, Linux, Solaris
  - Workloads: Web, email, database, etc.
- Found 2–3x performance variations
  - Depends on hardware, software, configuration, and workloads
- Plug: NFSv4 study, FAST'10 Poster session

Conclusions

- Very complex systems
- Hard to understand and optimize
- Lots of waste in software
- Great opportunities to improve
  - Research opportunities
  - Commercial tools and services

Let’s get to work...