Low-cost Disaggregation of Sub-system Power

Yuwen Sun, Lucas Wanner, Puneet Gupta, and Mani Srivastava
UC Los Angeles

Motivation: sub-system power

- Fine-grained power info enables optimization & adaptation for OS and Apps
- Power sensors for each component in the system: simple, but expensive
- ***Objective: sub-system power with only ONE system-wide power meter***

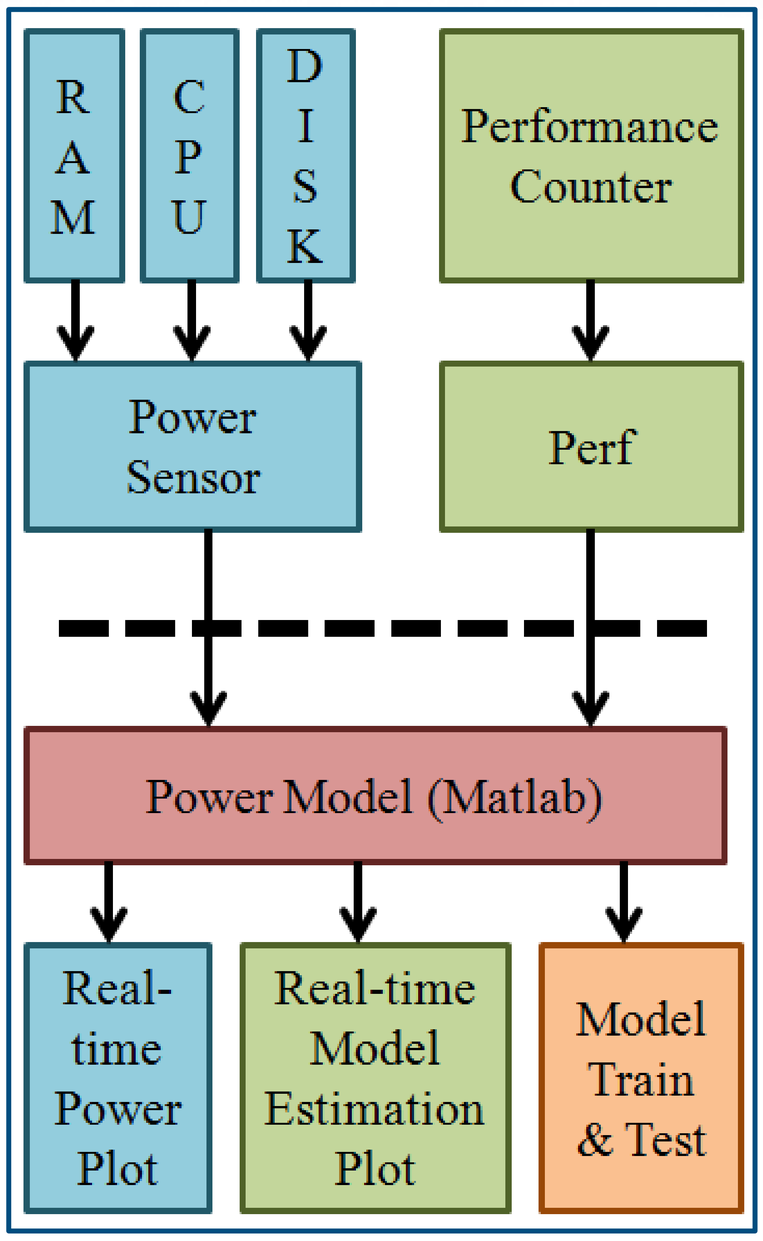

Approach: single power meter plus event counters

- Event counters record information about hardware usage:
  Context switches, cache misses, branches, etc.
- Hypothesis: ***A linear model can predict sub-system power***
  using event counters and system-wide power info
- ***Model: Low cost, transplantable, robust , and adaptable to variation***
  Self-training cannot be achieved by traditional static model

Linear Model

$$\text{Power}_{\text{Total}} = \left(\alpha \times \begin{bmatrix} C_{RAM}1 \\ C_{RAM}2 \\ . \\ C_{RAM}n \end{bmatrix}^T + \beta \times \begin{bmatrix} C_{CPU}1 \\ C_{CPU}2 \\ . \\ C_{CPU}n \end{bmatrix}^T + \gamma \times \begin{bmatrix} C_{Disk}1 \\ C_{Disk}2 \\ . \\ C_{disk}n \end{bmatrix}^T\right) \times \begin{bmatrix} \text{Per-Counter 1} \\ \text{Per-Counter 2} \\ . \\ \text{Per-Counter n} \end{bmatrix}$$

System Design

Preliminary Results

- Testbed: Sub-system power sensors as ground truth and system-wide power sensor as input to linear model
- Training: Ground truth power and event counters are recorded for CPU and Memory-intensive benchmarks
- ***Testing: Real-time disaggregation of power with accuracy of 95%***
- ***Live Demo***

Future Works

- Sophisticated model
- Different hardware/software information
- Various training algorithms & tools

Data Processing

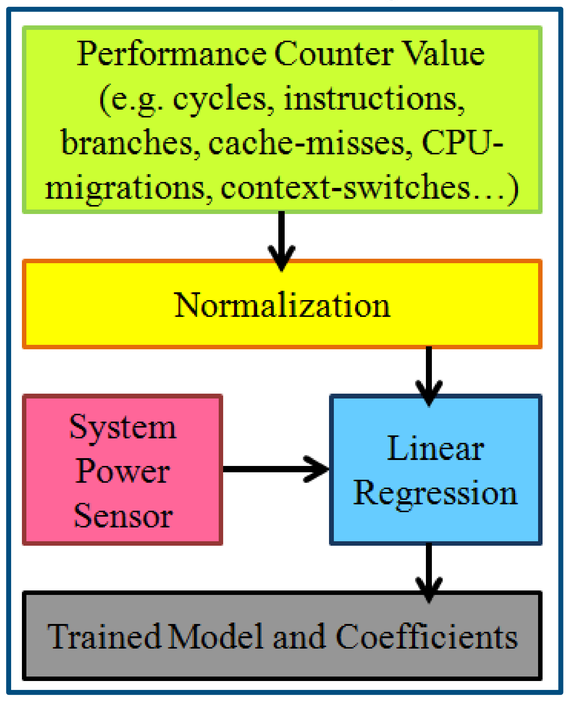

CPU Power DAQ Plot (8 hours, Ground-Truth)

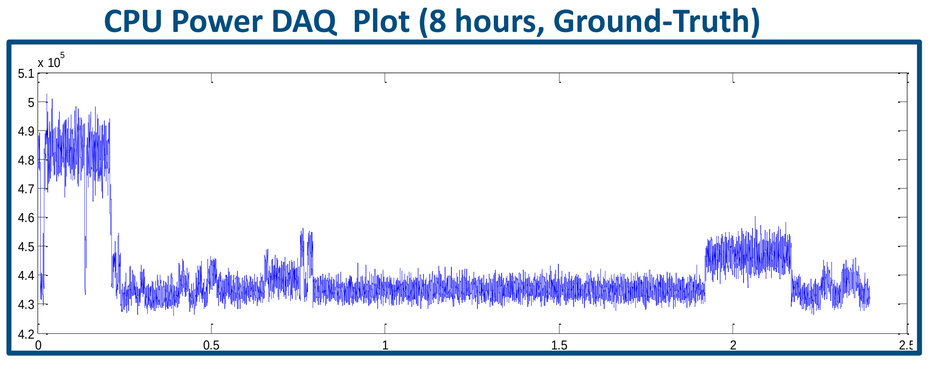

RAM Power DAQ Plot (8 hours, Ground-Truth)

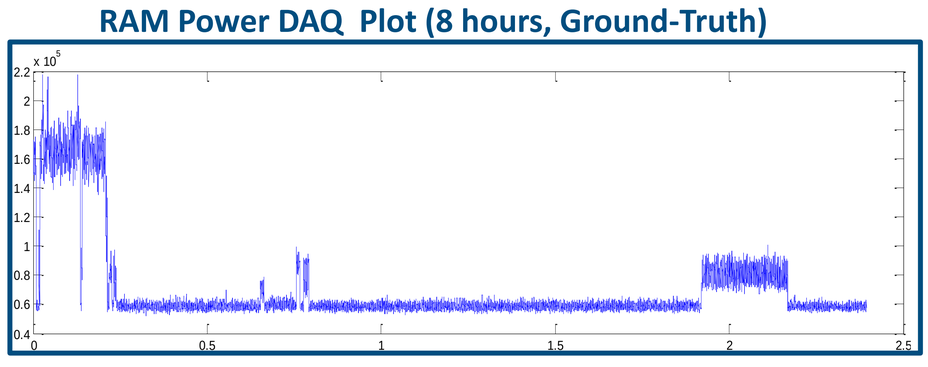

CPU Power Model Estimation Plot

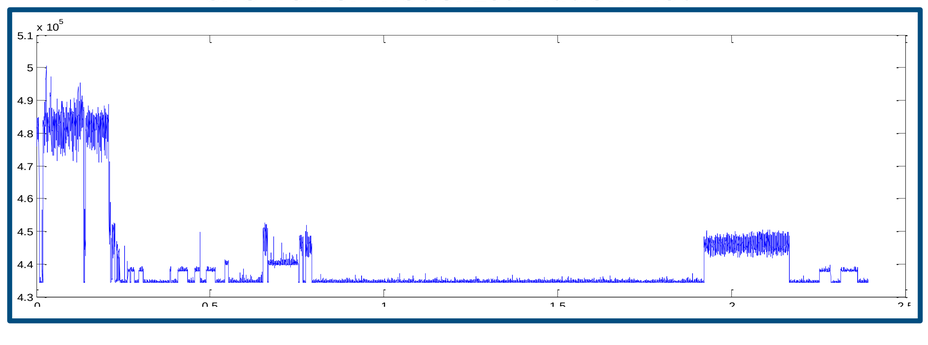

RAM Power Model Estimation Plot

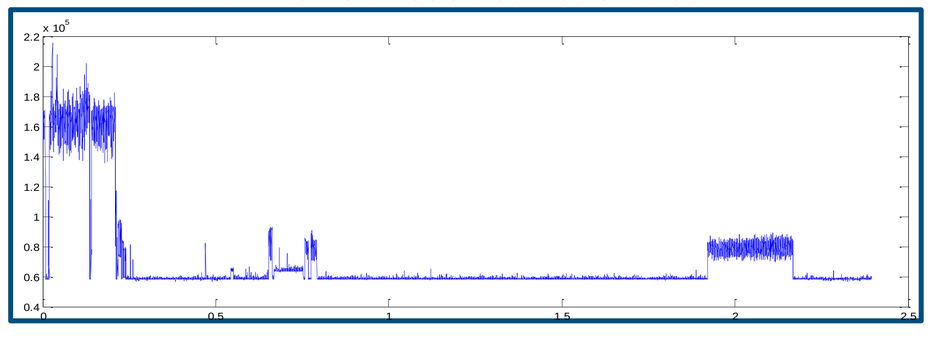

NSF Expedition in Computing, Variability-Aware Software for Efficient Computing with Nanoscale Devices http://variability.org

UC San Diego · UCLA · Michigan · Stanford University · UC Irvine · University of Illinois at Urbana-Champaign