University of Victoria

ICSE 2015

SEAMS 2015 May 18-19 Firenze, Italy

Adaptive Management of Energy Consumption using Adaptive Runtime Models

Andreas Bergen, Nina Taherimakhsousi, Hausi A. Müller

The Short Version

- Automatic classification of energy consumption patterns connected to specific process type and resource usage
- Adaptive Reference Model to classify energy patterns across multiple devices
- Improved training method for energy consumption data and pattern classification.

Outline

- Motivation
- Background
- Experimental Setup
- Self-Adaptive Models
- Experimental Data
- Limitations
- Future Work
- Conclusions

Motivation (1)

- ~ 2% of global energy consumption is used for data centres
  [Kim2011, Koomey2011, Niles2008]

- Server power demands are not homogenous
  - Server A != Server B
  - Idle energy consumption is ~50% of consumption at full load
    [Meisner2009, Srikantaiah2008]

- Workload is not homogenous

- Sensors exist in data centers
  - PDU
  - Monitoring software
  - Energy monitoring (billing, infrastructure)

Motivation (2)

Research Question:

- Can we automatically determine the dominant hardware resource a software application is using, based only on the energy consumption profile?
- Can we adaptively schedule processes based on energy-optimal criteria?
- Can we do this dynamically across heterogeneous devices?

Background (1)

- Consolidation of VMs
- Various heuristics for VM management
- Dynamic Frequency and Voltage Scaling
- Efficient application algorithms
- Economics based models
- Mobile applications: app has control over some aspects of how much power it uses

Background (2)

- Power consumption measurements at different levels
  - CPU, memory,...individual hardware components

- Larger scale: Power grid models, global perspective

⇒ We measure “real” power at the individual server rack

Experimental Setup

- Real Data Centre (EDC2 at Uvic)
  - 3.000 standard servers
  - 1.26 mega watts
  - HPC, WestGrid, GENI
  - Three phase 208V PDS, each rack has 2 independent 30A breaker circuits

- We measure at the rack’s PDU
- We had full control over the servers in the rack
  - Homogenous hardware
- CPU, disk and memory intensive software

Self-Adaptive Models

- Ability to identify the resource utilization of a process based on energy consumption allows for dynamic scheduling decisions to be made.
- Knowing resource utilization, we need to identify the energy-optimal hardware to place this process.

⇒ Adaptive identification is needed, because not every device is the same. Static approach will not work

Self-Adaptive Models

Pattern Classification

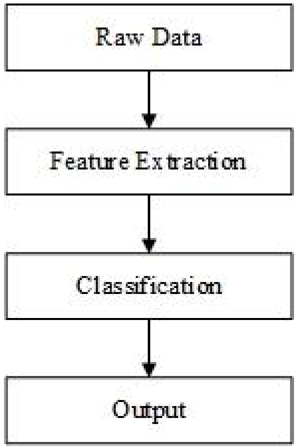

Self-Adaptive Models

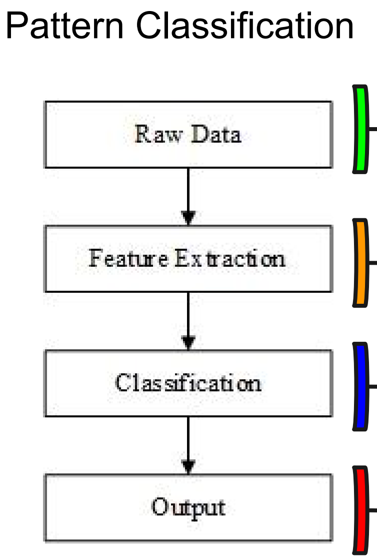

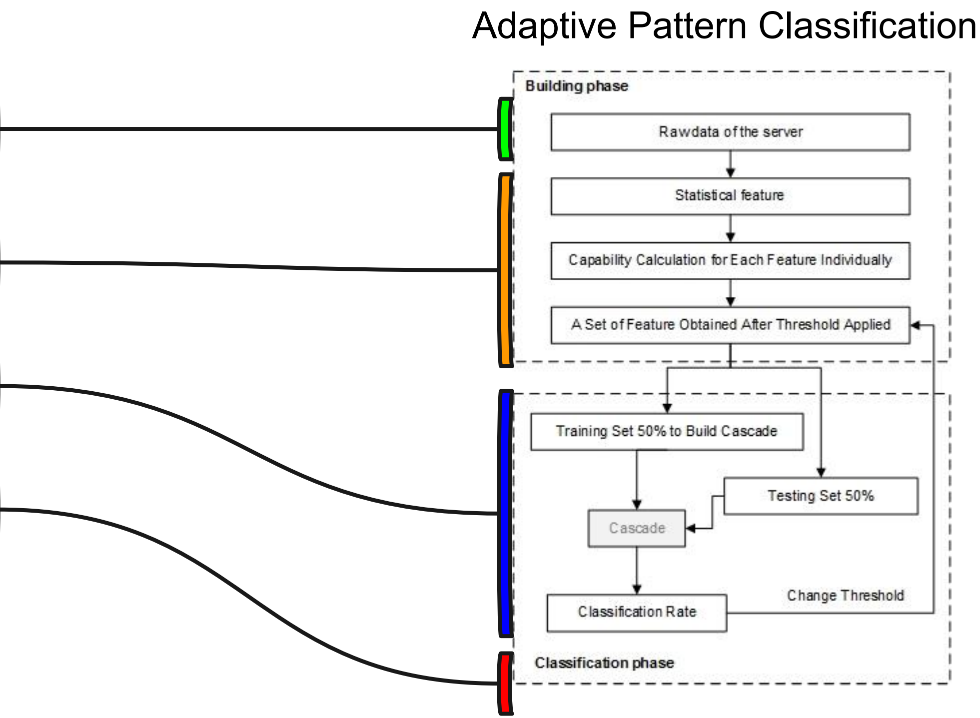

Self-Adaptive Models

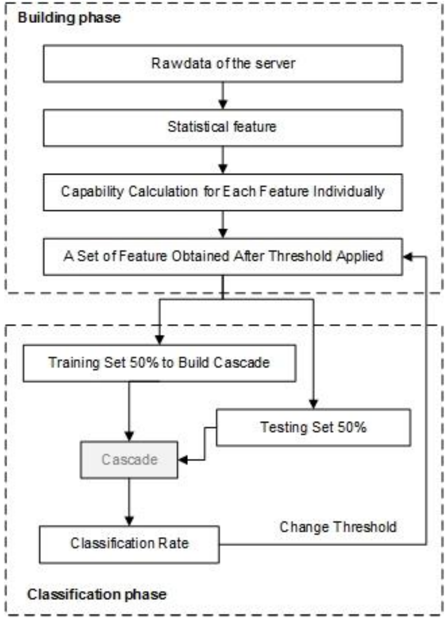

- Incorporate classification process into MIAC
  - dynamic classification on any device
- Redistribute processes based on overall energy consumption
- Dynamically update overall system

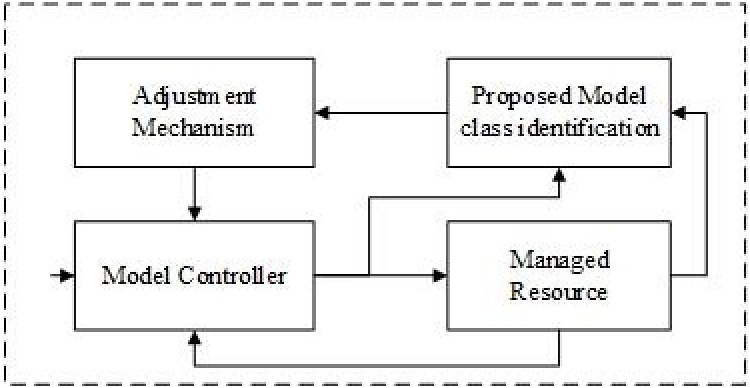

Self-Adaptive Models

- Machine learning:
  - SVM vs. Statistical Analysis vs. NN
    - theoretically both can achieve 100% correct identification
- Formula
- Using 50% of the data for training

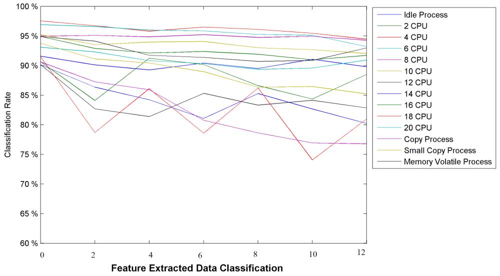

$$C_T = \frac{\mu_a - \mu_b}{\sqrt{\left(\frac{\sigma_a^2}{n_a}\right) + \left(\frac{\sigma_b^2}{n_b}\right)}}$$

Experimental Data

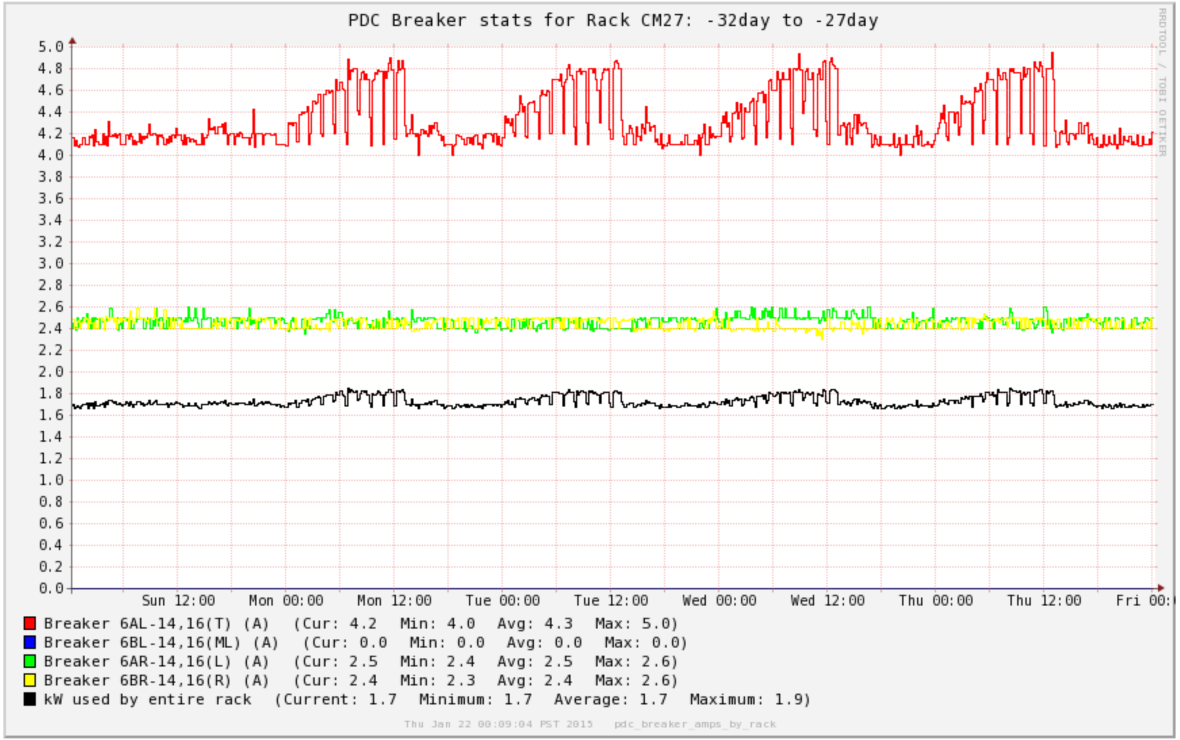

Limitations

- PDU → breakers connected to multiple servers
  - luckily we control all of them
  - can we deal with noise once we move to shared servers?
- Only “coarse” power measurements
  - data collection frequency
  - implications for types of measured processes
- Profiling of network I/O not yet completed

Future Work

- Complete profiling for network I/O
- Integrate profiling tools into a framework
- Migrate processes
  - measure actual energy profile of digital ecosystem
  - measure actual energy profile changes due to migration
- Profile multiple hardware

Conclusions

- Can we automatically determine the dominant hardware resource a software application is using, based only on the energy consumption profile? YES
- Can we adaptively schedule processes based on energy-optimal criteria? Theoretically yes
- Can we do this dynamically across heterogeneous devices? YES

[redacted]

Thank you

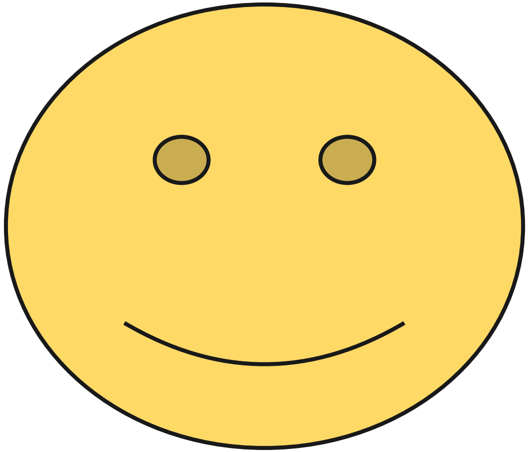

Backup Slides

Pattern Identification

Cross-validation	Idle Process	2 CPU	4 CPU	6 CPU	8 CPU	10 CPU	12 CPU	14 CPU	16 CPU	18 CPU	20 CPU	Copy Process	Copy Small	Memory Violate
50% no theshold	94.68	96.29	98.86	98.22	96.43	97.30	92.32	92.07	93.53	93.85	93.96	94.68	96.69	95.49
50% threshold	91.59	94.91	97.55	96.91	94.94	95.21	90.16	90.13	90.65	91.51	93.09	90.62	93.76	94.86
70% no theshold	91.07	96.25	98.00	97.40	96.16	96.27	92.5	92.44	92.3	92.99	94.21	93.77	96.2	95.39
70% threshold	90.55	95.21	97.22	96.7	95.22	95.13	91.43	92.29	90.71	91.52	93.15	91.19	93.8	94.71
80% no theshold	94.08	96.19	98.01	96.37	96.16	95.77	91.84	91.39	92.67	92.31	93.39	93.78	95.76	95.46
80% threshold	90.98	94.77	97.20	95.70	95.14	95.16	89.52	90.08	91.63	90.15	93.58	91.34	93.9	94.94
90% no theshold	91.37	96.96	97.35	96.29	94.8	94.93	92.35	93.22	92.31	92.84	93.79	93.72	95.68	95.47
90% threshold	89.16	95.52	96.92	96.04	94.25	94.84	91.3	93.25	91.43	90.78	93.02	91.49	93.77	95.32

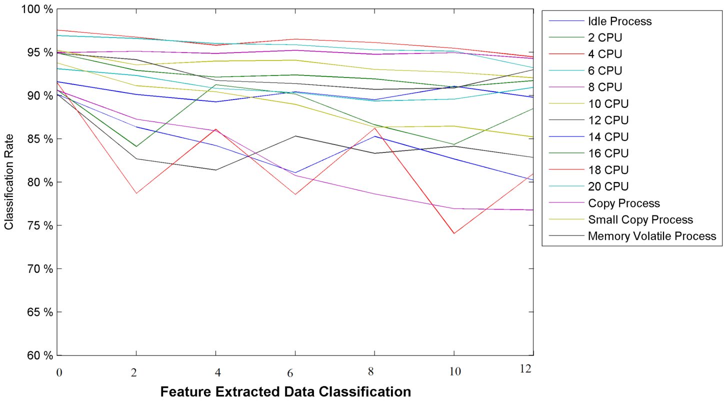

feature: amp up and down

CPU - Energy Consumption

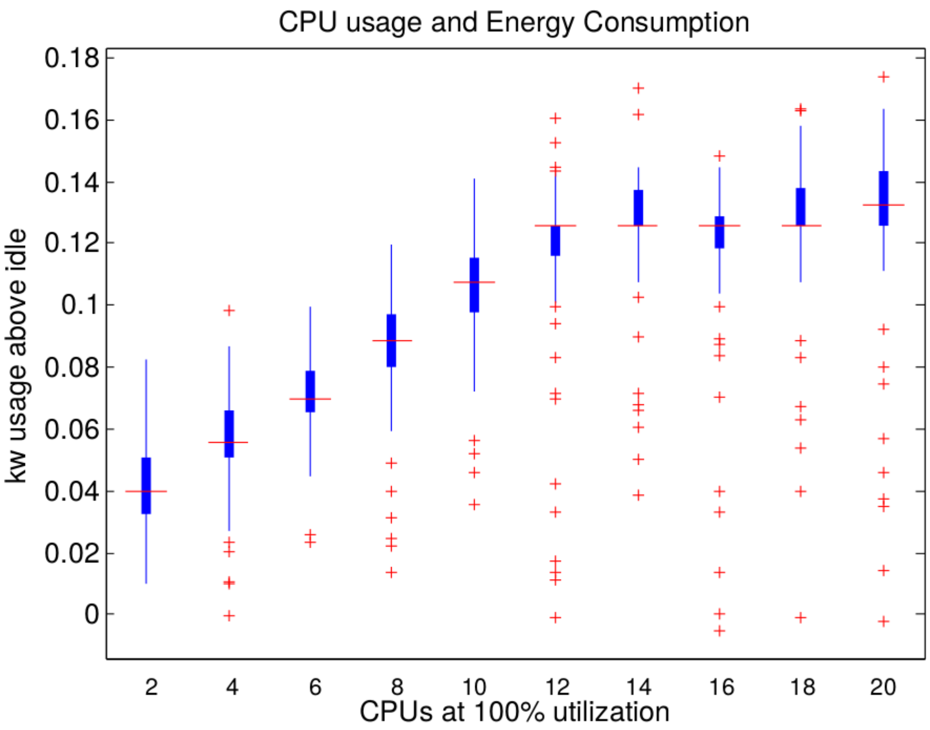

Conclusions

- Power profiles are distinguishable
- Power consumption can be treated as “utility” in resource allocation
- Even at our size data center, we can have significant energy savings
- CPU energy consumption ← step function
- Adding resources does not always improve performance
- Heavy memory usage: changing memory is what drives energy consumption
- Disk I/O is currently NOT optimized at all

Smart Adaptive Green Data Centre

- Adaptive job scheduling and resource provisioning
- We need a power consumption framework
- Formalize/automate informed/smart decision making → adhere to SLAs
- Need more information, sensors and monitoring access points
- Dynamic frameworks and models