Engineering Energy-Aware Web Services Toward Dynamically-Green Computing

Peter Bartalos
M. Brian Blake

Department of Computer Science and Engineering
University of Notre Dame

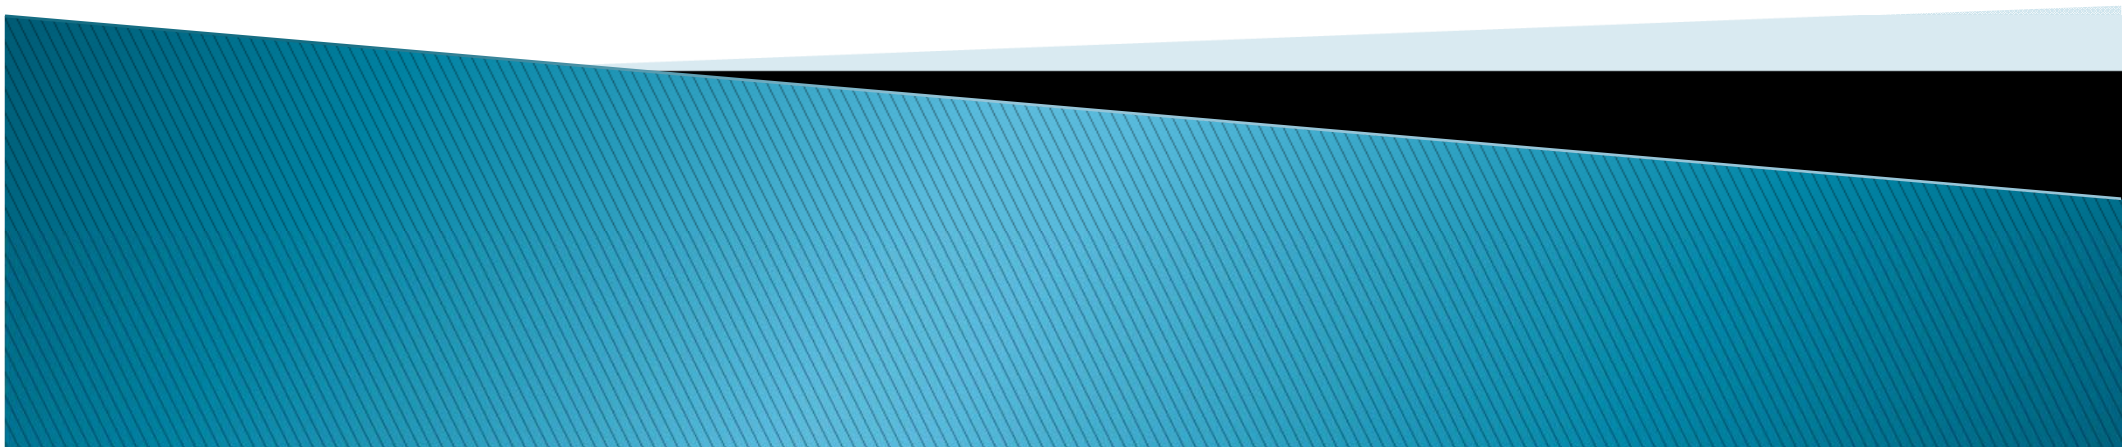

Goal

- Research the power consumption of web services

GREEN WEB SERVICES

Motivation

- The global CO2 emission in 2009 was estimated at 30.398 billion metric tonnes (http://co2now.org/)
- Gartner estimates that ICT is responsible for 2% (http://www.gartner.com/it/page.jsp?id=503867)
- This is equivalent to around $50 billion (considering $0.15/kWh).

Web services power consumption

- Multiple web service instances spread across regions
- Environmentally-aware web services

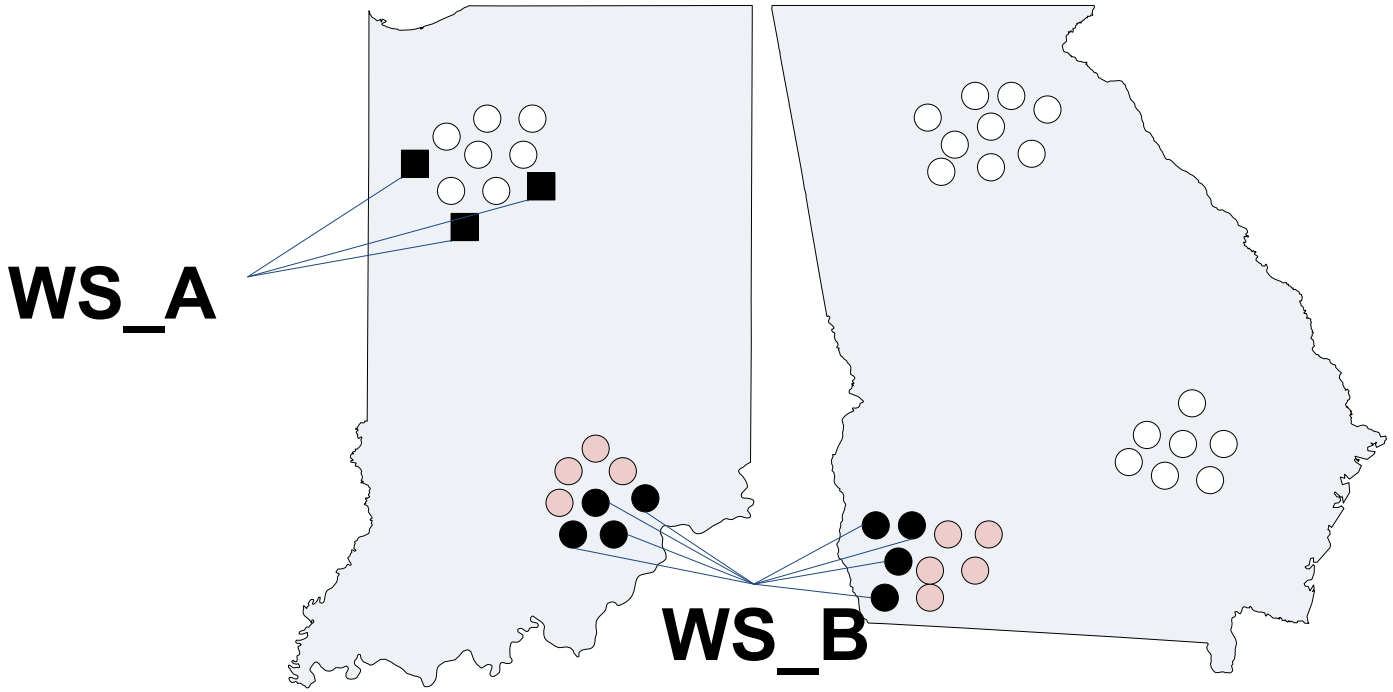

Preliminary experiments

- The nominal power consumption of web services varies according to the actual server load
- Dynamic power consumption determination is required

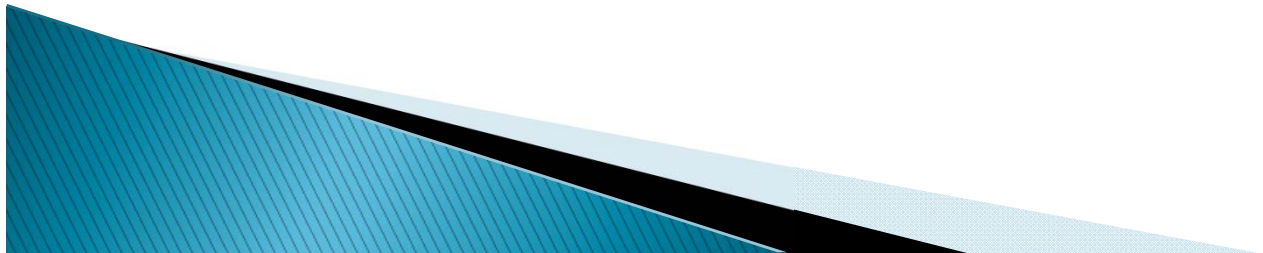

Measured power consumption

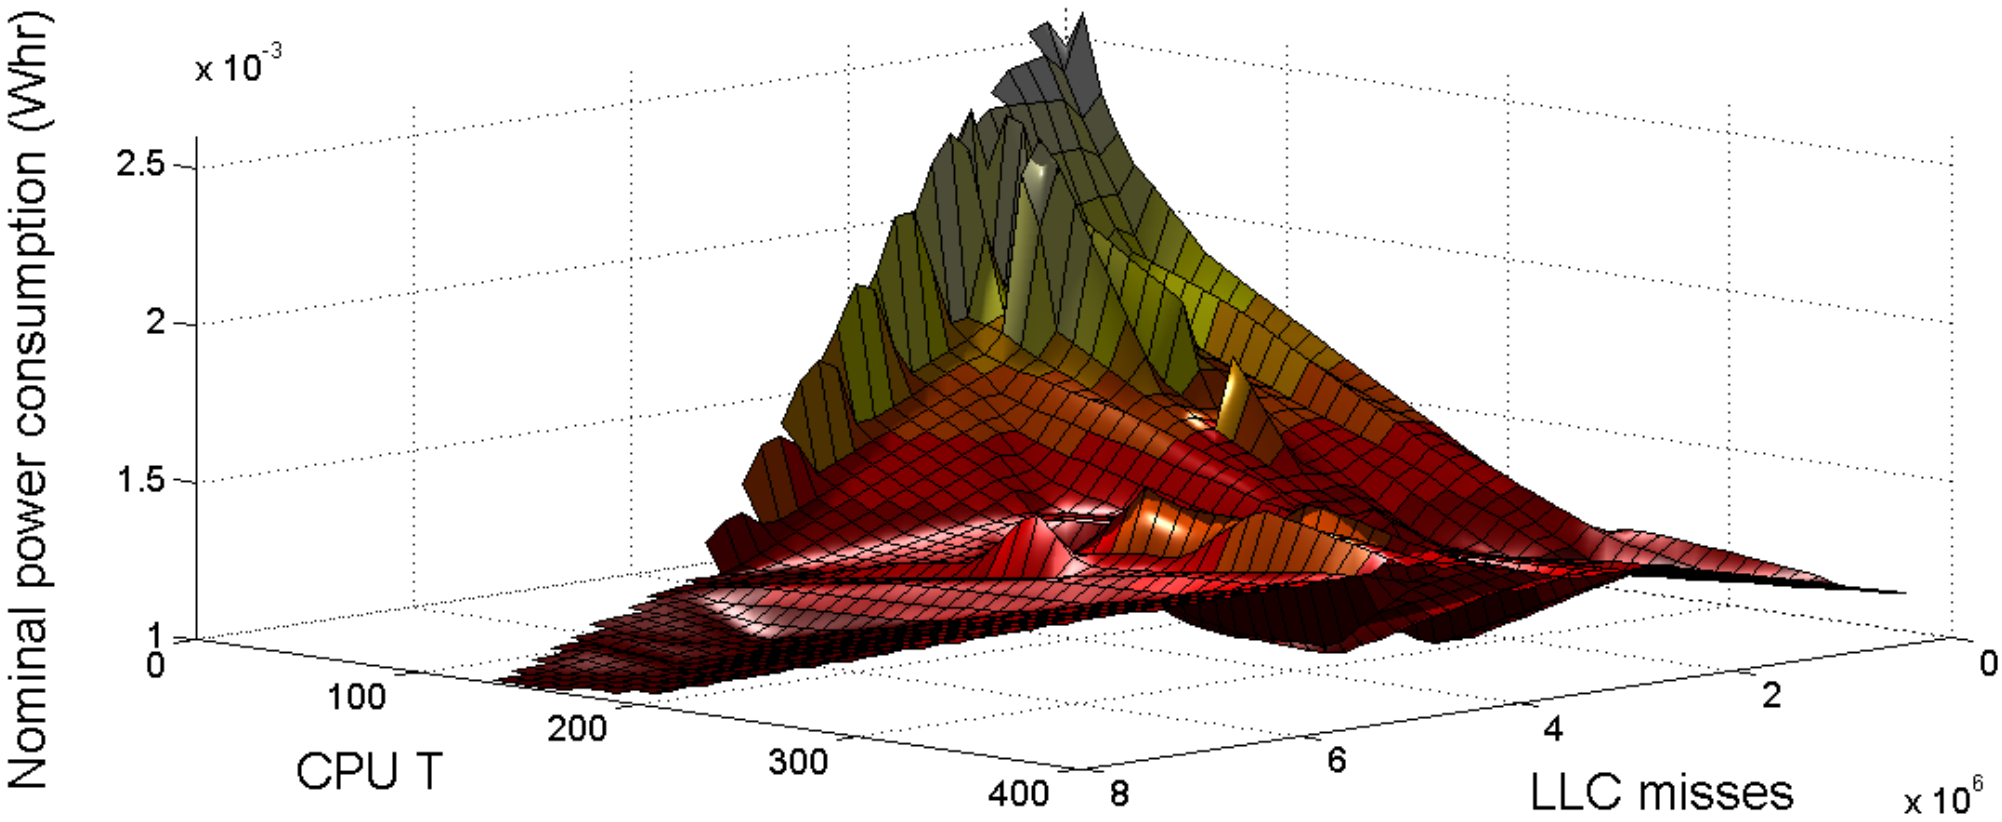

WS power model

$$\textbf{power} : Platform \times State \times WS \times I \longmapsto\longrightarrow R$$

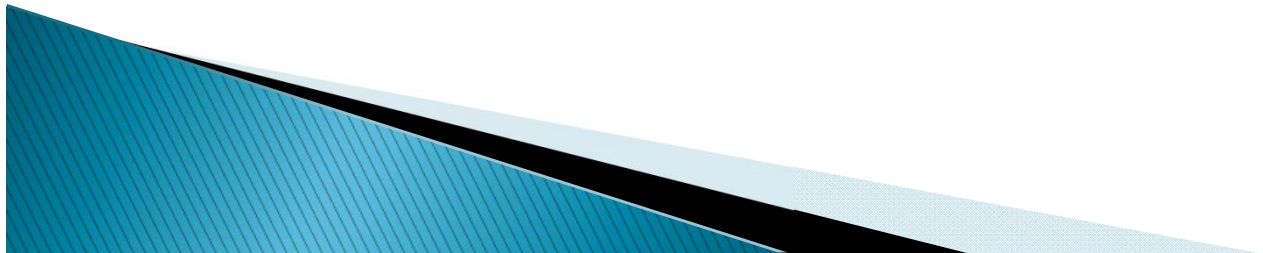

WS power model

$$\widehat{\mathbf{power}} : \widehat{State} \longmapsto R$$

- The utilization state is captured by a vector of hardware performance counters

Performance counters

- CPU utilization
- Last level cache misses
- HDD read/write bytes
- Network traffic – I/O bytes

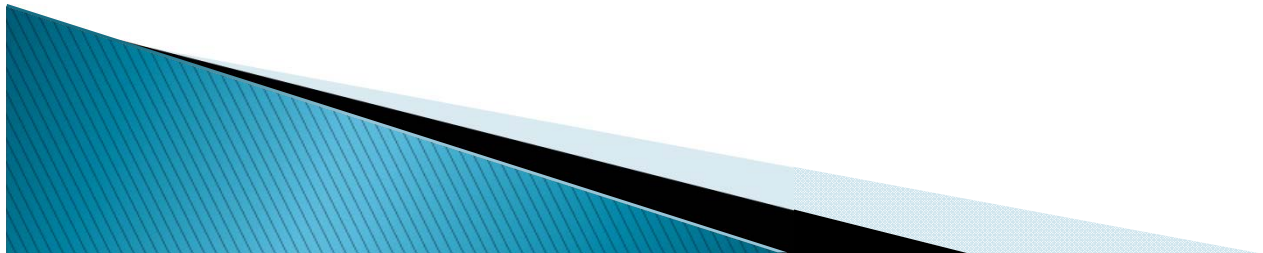

Calculating nominal power consumption

$$C = \frac{P_E}{P_E + P_B} \frac{(P_I + P_{EB})\Delta t}{|WSE|_{EB}}$$

WS power model

$$\widehat{\mathbf{power}} : \widehat{State} \longmapsto R$$

- Regression models, neural networks
  - Training phase
  - Operation phase

$$P = P_{idle} + \beta_1 X_1 + \beta_2 X_2 + \ldots + \beta_n X_n$$

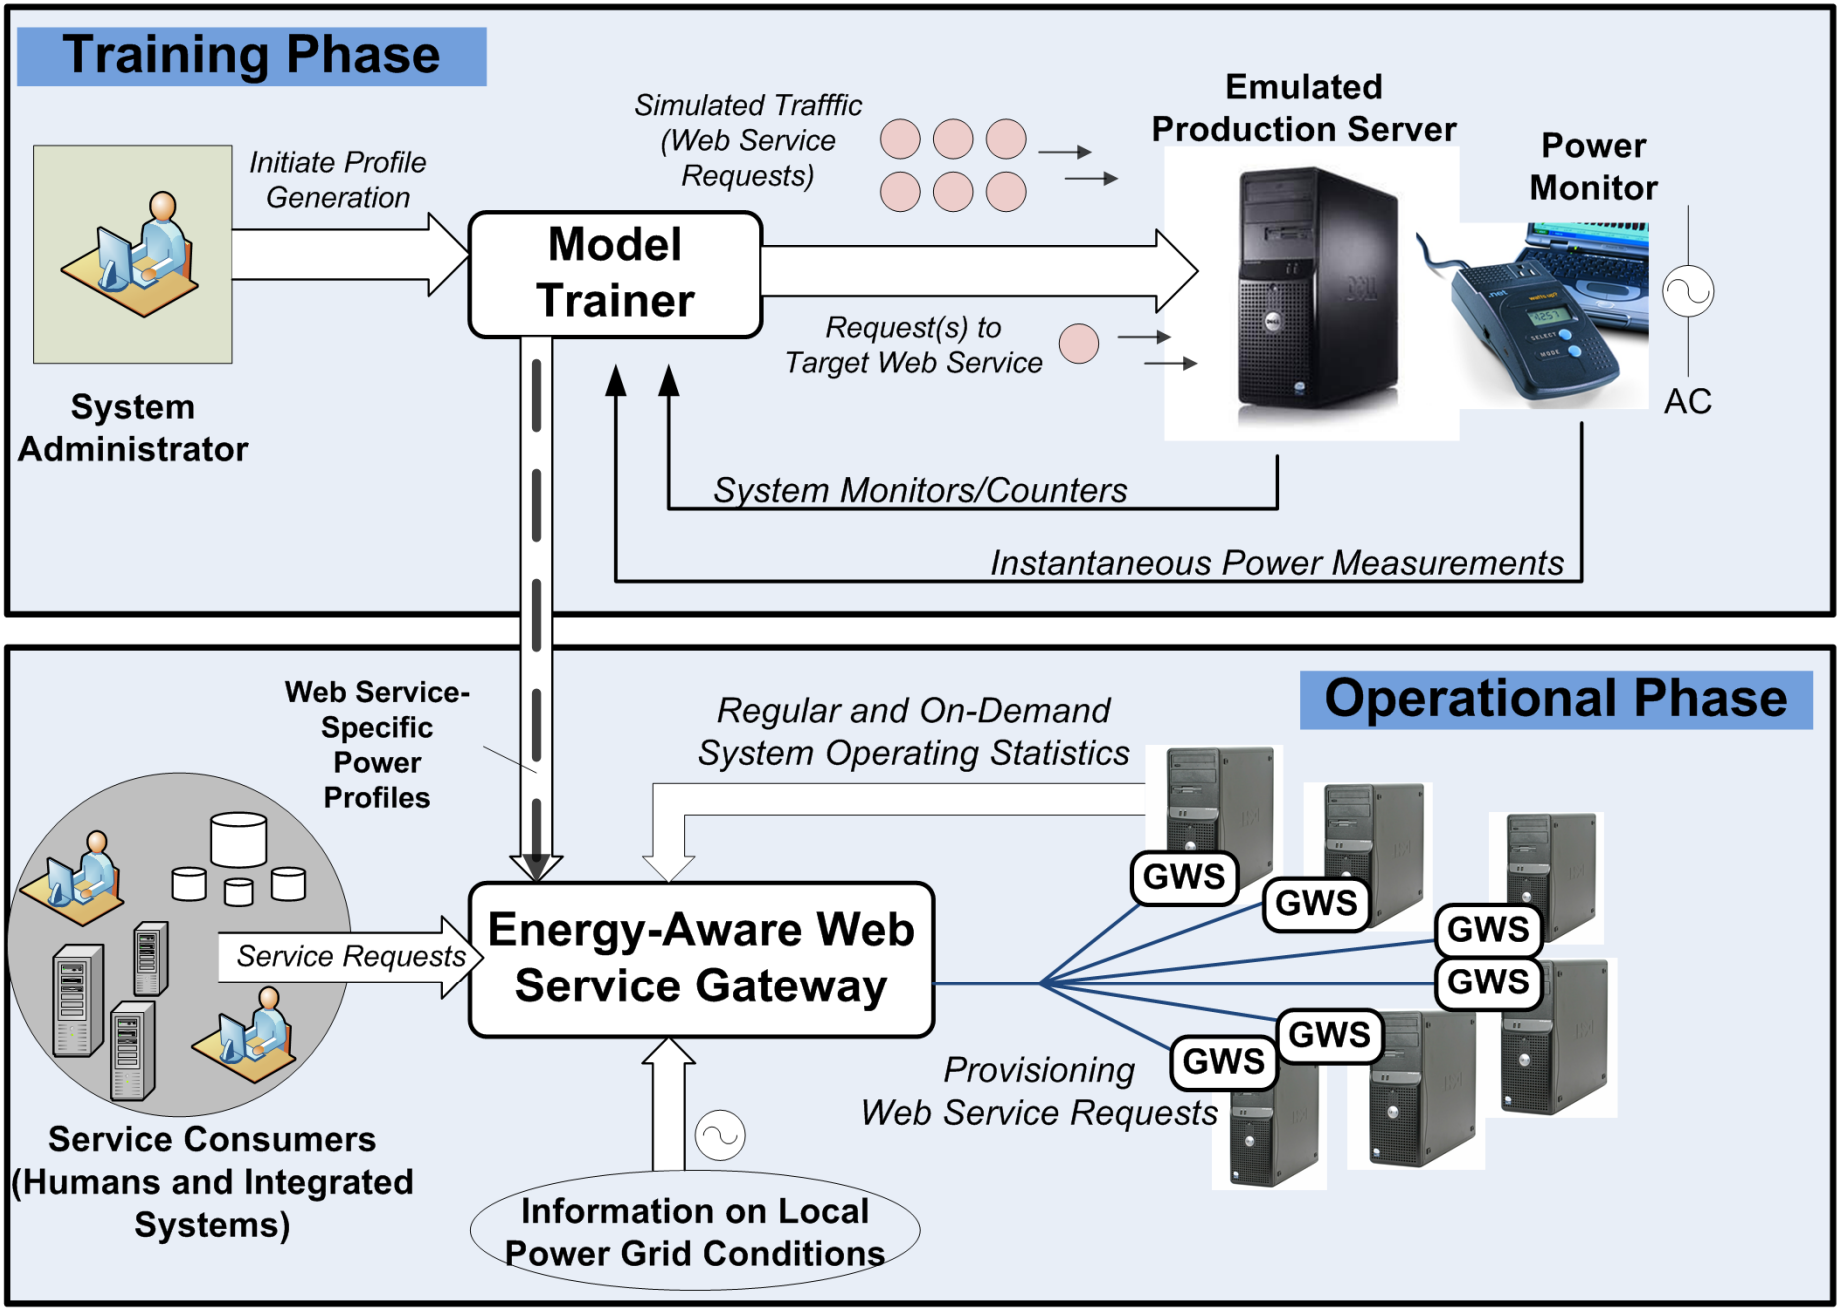
Training Phase
Initiate Profile Generation
System Administrator
Model Trainer
Simulated Trafffic (Web Service Requests)
Request(s) to Target Web Service
Emulated Production Server
Power Monitor
AC
System Monitors/Counters
Instantaneous Power Measurements
Operational Phase
Web Service-Specific Power Profiles
Regular and On-Demand System Operating Statistics
Service Requests
Energy-Aware Web Service Gateway
GWS
GWS
GWS
GWS
GWS
GWS
Provisioning Web Service Requests
Service Consumers (Humans and Integrated Systems)
Information on Local Power Grid Conditions

Experimental setup

WS power model – training

- Required data:
  - Performance counters
  - WS statistics
  - Power measurements

- Q1: Which performance counters?
- Q2: How to access data?
- Q3: How to synchronize the samples?
- Q4: How to train the model?

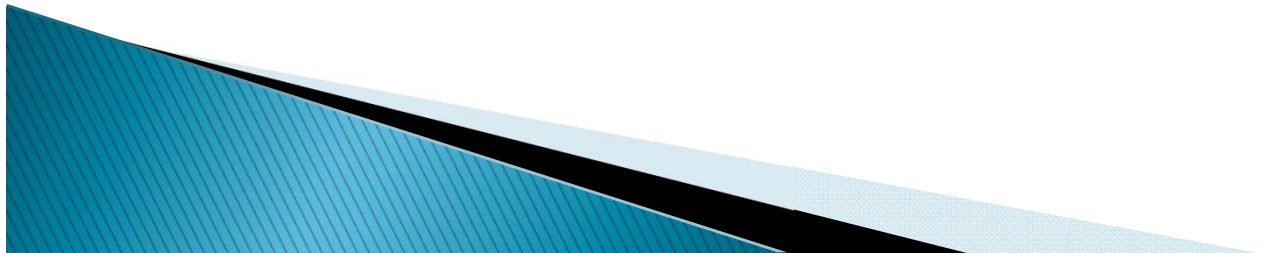

Q1: Which performance counters?

- P. Bartalos and M. B. Blake. Green web services: Models for energy-aware web services and applications. In *KASTLES'11.*

- Performance counters best capturing the server power consumption:
  - Percentage of elapsed time the CPU spends to execute a non-Idle thread.
  - LLC misses.
  - Number of unhalted core cycles of the CPU.
  - Number of instructions executed by the CPU.

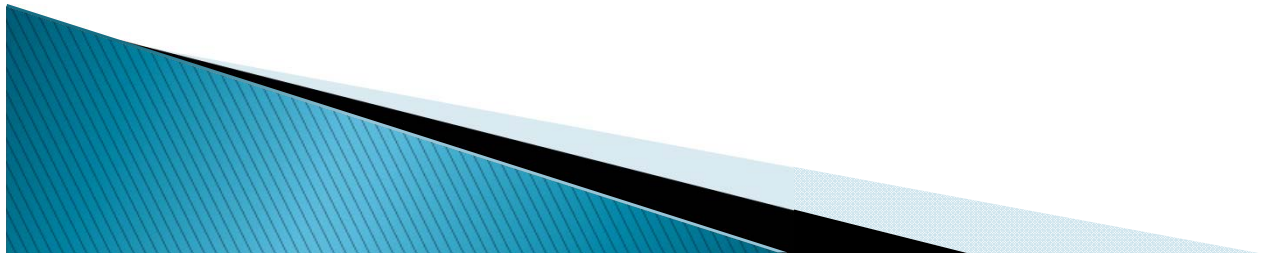

Q2: How to access data?

- CPU
  - Intel Performance Counter Monitor
- WS statistics
  - Servlet filters
- Power measurements
  - Watts UP .net power meter + custom app

Q3: How to synchronize the samples?

- The communication with the power meter is slow and unreliable
- Under high CPU loads the inter-app communication is slow and unreliable

- Central Java application manages the apps
- Synchronization according to a signal sent by the power meter app
- Inter-app communication only at the beginning and the end of data capturing

Q4: How to train the model?

- The training procedure significantly affects the accuracy of the power model
- We need samples from a variety of server states

Q4: How to train the model?

- Workloads generated by benchmarks
  - The mean error was around 4.84%
- Synthetic web service workloads
  - Web service selection strategy
    - If fewer samples are used to train the model, S2 performs better
    - If the training set is large enough, S1 is better
  - Local vs. remote invocation
    - Remote clients achieve better results
  - The mean error is about 1.7%

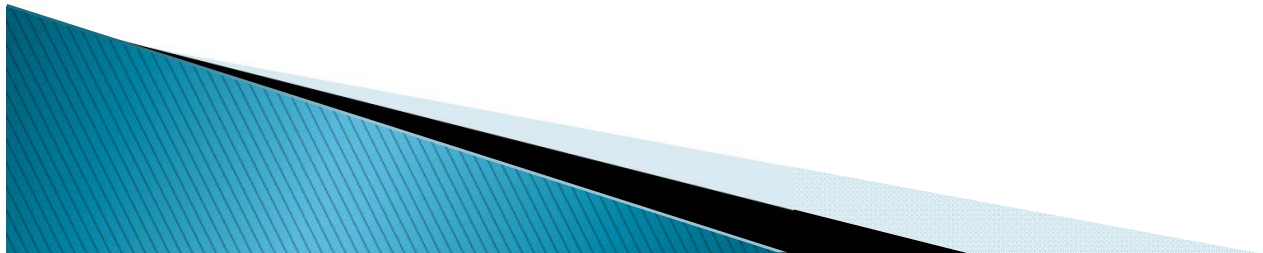

Summary

Peter Bartalos, Brian Blake

- Web services **power model**
- **Dynamic** power consumption determination
- **Environmentally-aware** web service discovery, and composition

GREEN WEB SERVICES

peter.bartalos@gmail.com