Cloud Computing —
Challenges and Opportunities

Cheng-Zhong Xu
Cloud and Internet Computing Lab
Wayne State University
Detroit, USA
czxu@wayne.edu

http://www.cic.eng.wayne.edu

Mother Nature in Tear

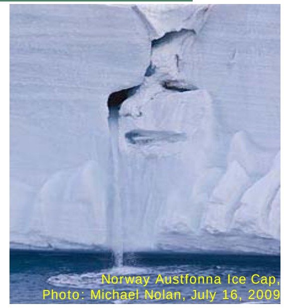

IT Liability

- ICT infra accounts for 2~3% of global electricity usage and greenhouse gas
  - About the emissions of airlines
  - Google search generated ~7g carbon emission; 200m search/day → 70,000 car's CO_2 emission
  - IEA updated a warning in 5/2009 that ICT energy use could double by 2022, and tripled by 2030

- More than half due to server and DC energy and emissions; 40% goes to PC/monitors.

Sustainability: we care, but …

- Server Usage
  - 10~15% utilization on average
- PC Usage
  - 70% in idle in office, 31% in idle at home; but only 4-6% in sleep state (50w in idle vs. 2~4w in sleep by Energy Star for PC, as of 2009)
  - Power management not widespread in IT (13% of business PCs)
- Very short life cycle (3~5 years)
  - 70% hazardous waste

Sustainability: Culture Issues

- General-purpose computer in open arch is hardly sustainable by its nature.
  - Reasons for replacing a desktop (consumers'08)
    - Too slow (50%), newer features (50%), upgrade/failure (30%), virus/spyware (14%)
  - Install new apps → Lib upgrade → OS upgrade→ current appl failure/too slow → Replace machine!
- Little chances for "good" enough computing
  - Be "good" enough for today's applications, but not for emerging appls. (vs TV, printer, camera, etc)

Sustainable Computing in Cloud

- "Boundless infrastructure", but you use what you need and you pay for what you use
- Virtualization on shared infra. leads to high utilization, and power conservation
- Server-based computing would create the culture of "good" enough computing

Cloud computing makes IT cool

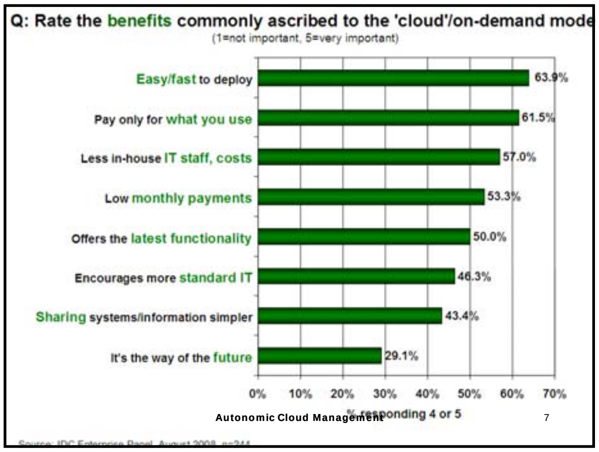

Outline

- Sustainable Computing in Cloud
- Cloud Computing and Challenges
- Automatic Cloud Management
  - Reinforcement Learning for Auto-configuration of virtual machines in cloud
- Cloud Computing in Retrospect

Cloud Computing

- A form of Internet computing, aiming to offer IT capabilities over the Internet as an on-demand, per-per-use service
  - Infrastructure-as-Service
    - E.g. Amazon's Web Services
- Platform as Service
  - Google AppEng,
  - Microsoft Azure
- Software as Service
  - Salesforce.com

Cloud Service: Current and Future

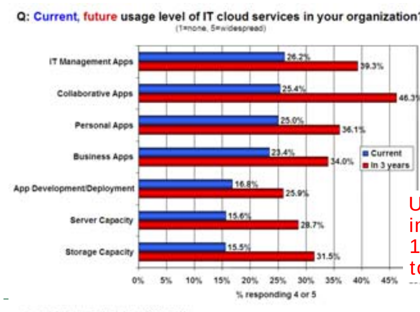

Usage increase from 15~25% to 30~40%

Cloud Service: Cross the Chasm

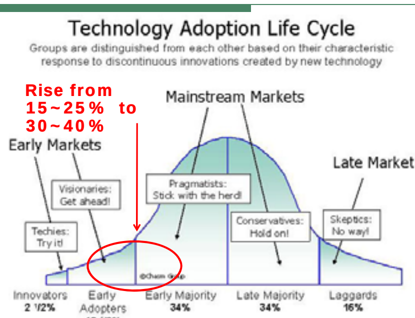

Key Enabling Technologies

- Server/Datacenter
  - Exponential growth of processing/storage capacity
  - Industries race to build next-gen DCs
  - DCs become the next computing platform

- High bw and pervasive connectivity
  - 100 gigbit, 1 terabit networks
  - [Ulta] wideband wireless mobile access

Key Enabling Technologies (cont')

- Server virtualization
  - Consolidation of servers in DC
  - Realize economics of scale
- Pay-as-you-go business model
  - Service-oriented computing
  - Low-touch, low commitment self-service
- Autonomic management
  - Scaling up and down with load on auto-pilot, self-configuration, self-*
  - Quality of service (QoS) provisioning

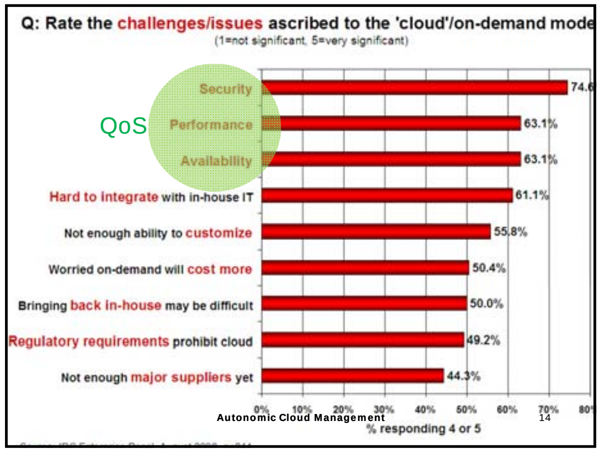

Quality of Service

- Performance
  - Perf. isolation and differentiation
- Availability
  - Server behavior in stress conditions?
  - Resilient to failure
- Security
  - Prevention like ID/AC is not enough for flash-crowd like attacks
  - Admit good, block bad, and contain suspicious ones. How to contain?

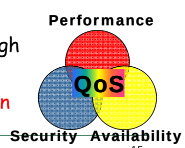

Challenges in Management

- Systems are too large, too complex
  - Multiple levels of abstractions and interactions between the components
- Too much data, less info
  - Little actionable info
- Heuristic knowledge based diagnosis,
  - not enough to deal with complex systems
- Need for online decision makes the problem even harder
  - Some problems could take days/weeks to resolve

Examples of the Challenge

- SLA Example: Amazon EC2
  - 99.95% availability for the EC2 service on a yearly basis (4 hours and 23 min outage per year)
  - Unavailability in a 5-minute period
  - 10% credit, if SLA violation is proved
- Revenue Loss Per Down Hour
  - Amazon outage on June 6, 2008 for 2 hrs; cnet.com estimated a loss of $16,000/minute and $2M in total
  - eBay search engine down for 1.5 hrs on Aug 16, 08
  - Google Gmail down for 2+ hrs on Aug 11, 2008

Worldwide Server/DC Spending

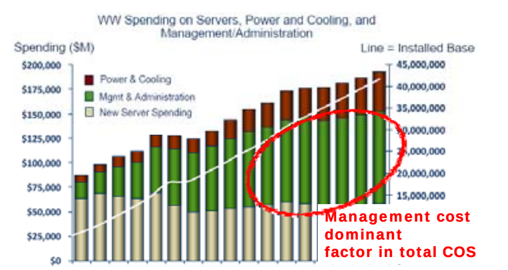

Source: IDC, June 2009

ACM @ Wayne State

- Autonomic Cloud Management for rapid deployment and management of cloud
  - Adaptive to workload change, client requirements, resource supplies, system failures, power cap/energy budget, etc
- Machine Learning, Optimization and Control
  - Auto-configuration of virtual machines
  - Service quality assurance and adaptation
  - Proactive failure management
  - Coordinated power/perf management

Acknowledgment

- Ph.D. Graduates/Research Associates
  - H. Shen (PhD'06, U. of Clemson), J. Wei (PhD'06, Yahoo! Technical), X. Zhong (PhD'07, Microsoft), S. Fu (PhD'08, New Mexico Tech.), B. Yu (Postdoc'07-09, GM R&D), X. Zhou (Postdoc'02-04, U. of Colorado)
- Current Students/Post-Docs
  - M. Dawoodi, E. Harash, J. Rao, J. Gong, X.Bu, M. Xu, K Wang, Q. Lan, C. Li, Z. Kong (post-doc)
- Funded by NSF and Sun:
  - CNS-0914330 , CNS-0851856, CRI-0708232, CNS-0702488, DMS-0624849, CCF-0611750
  - Sun Microsystems' Center of Excellent in Open Source Computing and Applications (OSCA)

Auto-Configuration of Cloud

(ICDCS'09 and ICAC'09, in collaboration with J. Rao and X. Bu)

Needs for VM Dynamic Config.

- Server virtualization
  - an abstraction layer that hides physical characteristics of a "real" machine.
  - Server consolidation to reduce TCO
- VM creation on demand in a real-time fashion
  - Static capacity planning is insufficient to deal with traffic dynamics

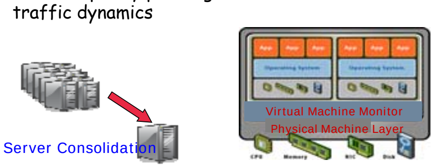

Challenges of VM Config

- Performance Interference between VMs

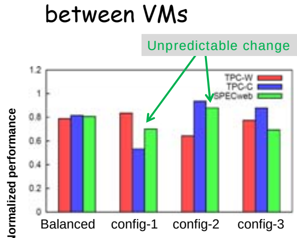

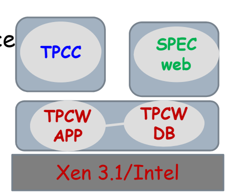

Config-1: move 1G mem From TPCC to TPCW

Challenges of VM Config (con't)

- Delayed effect: Reconfiguration takes effect after certain delays
- Performance inertia due to memory reconfig

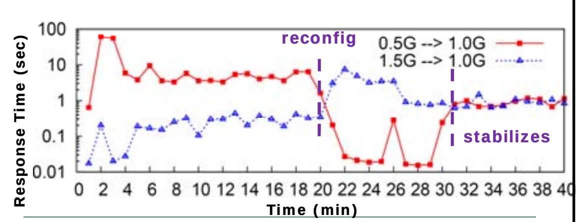

Dynamic Application Config

- Network apps contains many perf-critical params
  - Both Apache and Tomcast have 100+ configurable params
- Default settings (one-size-fit-all) may not work for different types of input traffic/VM configs

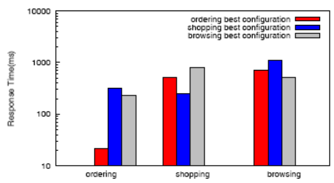

Reinforcement Learning Approach

- Learning by interaction with env
  - State: configuration of VMs (cpu, mem, time, etc)
  - Action: reconfiguration (increase/decrease/nop of resrc)
  - Immediate reward: summarized perf. of VMs w.r.t. response time or throughput
- Objective
  - For a given state, find an maximize long-run return

$$Q(s,a) = E\{\sum_{k=0}^{\infty} \gamma^k r_{t+k+1} | s_t = s, a_t = a\}$$

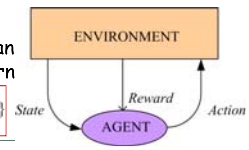

Parallel Q-Learning

- Q-learning to approximate the long-run return iteratively:

$$Q(s_t,a_t) = Q(s_t,a_t) + \alpha * [r_{t+1} + \gamma * Q(s_{t+1},a_{t+1}) - Q(s_t,a_t)]$$

- Neural-networks for Q-table initialization
- NN for approximation of unvisited states
- See [JOTA'08] for properties of convergence and convergence rate of asynchronous parallel Q-learning.

RL Arch. for Autoconfiguration

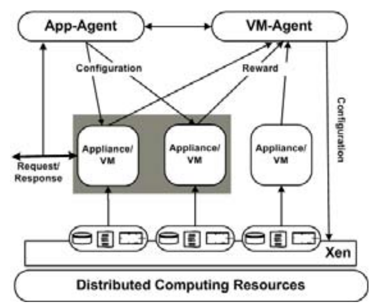

- VM-Agent
  - VM config on same PM
- App-Agent
  - App config for app in the same or diff VMs

Experimental Settings

- VMs ran on a cluster of Xen servers (2-sockets, quad-core)
- Benchmarks
  - TPC-C: online trans.
    - Disk/net I/O-intensive
  - TPC-W: e-commerce
    - Browsing input mix
    - Shopping input mix
    - Ordering input mix
  - SPEC-Web
    - Static vs dynamically generated web pages

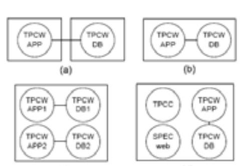

Adaptability to Workload Change

- Two TPC-W apps run concurrently

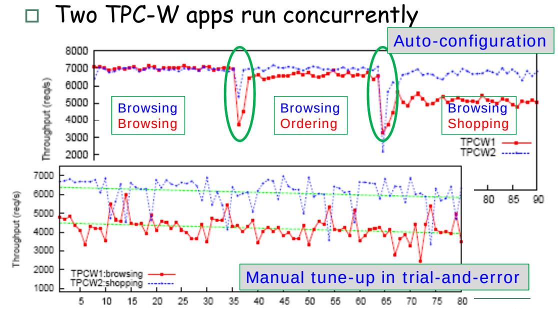

Coordinated Reconfig

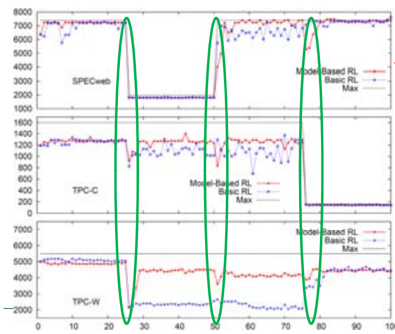

Throughput

Adaptability to Traffic and VM

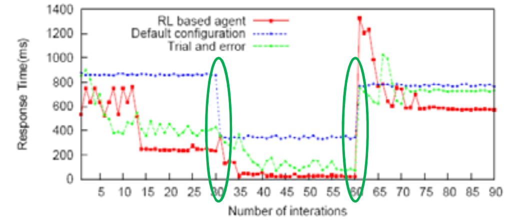

Shopping
High VM Config

Ordering
High VM Config

Ordering
Low VM Config

See ICAC'09 and ICDCS'09 for more details

Cloud Management: In Summary

- Reinforcement learning for VM Autoconfiguration [ICAC'09] and web appliances [ICDCS'09]
- Machine learning for proactive server failure management [SC'07][SRDS'07]
- Statistical modeling for server anomaly detection (overload) using OS and hardware-level parameters [IPDPS'08][ICDCS'08]
- Service quality assurance and adaptation [TC'05, TPDS'06, Computer'08, ICDCS'08]
- Power-aware resource management [ICDCS'08, INFOCOM'08, LCTES'08, TC'06, ACM TECS'08]

Outline

- Sustainable Computing in Cloud
- Cloud Computing and Challenges
- Automatic Cloud Management
  - Reinforcement Learning for Auto-configuration of virtual machines in cloud
- Cloud Computing in Retrospect

Cloud Computing: In Retrospect

Centralized
→ Decentralized Resource
→ Centralized Management ?

"The computation and the data and so forth are in the servers. ... We call it cloud computing." (Erick Schmidt, 2006)

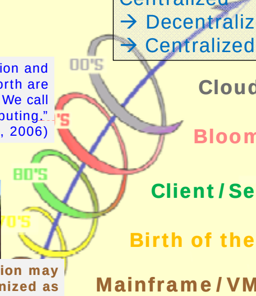

Cloud Computing / VM

Bloom of the Internet

Client / Server Computing

Birth of the Internet

"computation may someday be organized as a public utility" (John McCarthy, 1960)

Mainframe / VM
IBM 370 VM / CMS

Big Switch in Early 20th Century

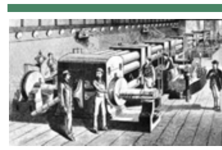

Key Enabling technologies
- •Power generating capacity
- •Alternative-current (1888)
- •Demand metering (1894)

Edison's Pearl Street Station (NYC): the first electric lighting station opened in 1882

Transition from private plants to utilities
- •1900: 50,000 private electric plants in U.S. 3,600 central stations
- •1907: 40% in utilities
- •1920: 70%
- •1930: 80%~90%

Big Switch in Early 21st Century

- Industries race to build datacenters (capacity)
- High BW netork and pervasive connectivity
  - Terabit network in sight
  - Wideband wireless mobile access
- Virtualization help realize the economics of scale
- Measurement and management

Thanks.

Cloud and Internet Computing Laboratory
Center for Networked Computing Systems
Sun Center of Excellence in Open Source Computing and Applications (OSCA)
Wayne State University, Detroit, Michigan
HTTP://www.cic.eng.wayne.edu
Cheng-Zhong Xu, czxu@wayne.edu

Backup Slides

Clouds vs Grids– A Personal View

- **Grids**
  - Born to share resources among different admin domains for problems beyond the individual capacity
  - Primary usages:
    - End-users applications are run a batch mode
    - Best effort job scheduling
    - Resource reservation
  - Key enabling tech.
    - Interoperability of services
    - Resource management
    - etc
- **Clouds**
  - Born to enable flexible renting (sharing) of an IT infrastructure
  - Primary usages:
    - On-demand real-time allocation
    - Pay as-you-go
  - Key enabling Tech.
    - Virtualization
    - **Auto-management:** Scaling up and down with load on auto-pilot; Auto-configuration; Failure repair

Challenge due to scale and real-time.

QoS Assurance and Adaptation

MLOC for QoS

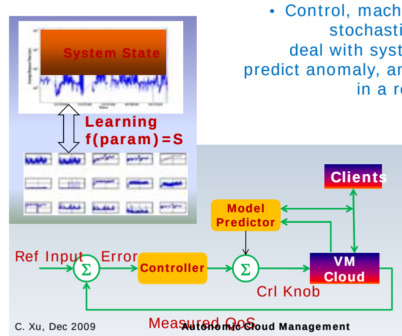

- Control, machine learning, and stochastic optimization to deal with systems uncertainty, predict anomaly, and scale resource in a real-time manner.
- Multi-agent coordination for tradeoff

Model-Free, Self-Tuning Fuzzy Control

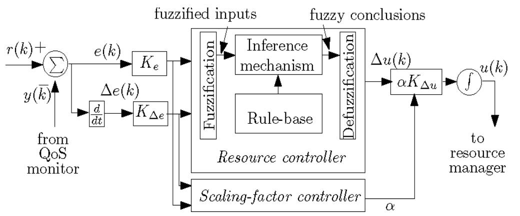

See [TC'05, Computer'08] for stability analysis and other details

Transient Behavior on PlanetLab

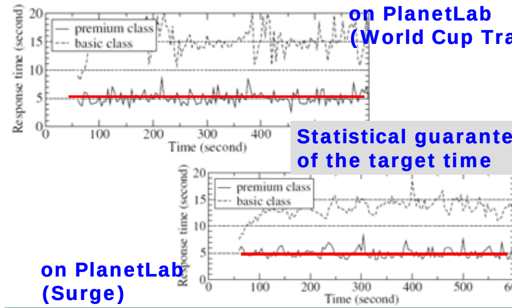

on PlanetLab (World Cup Trace)

Statistical guarantee of the target time

on PlanetLab (Surge)

Robustness

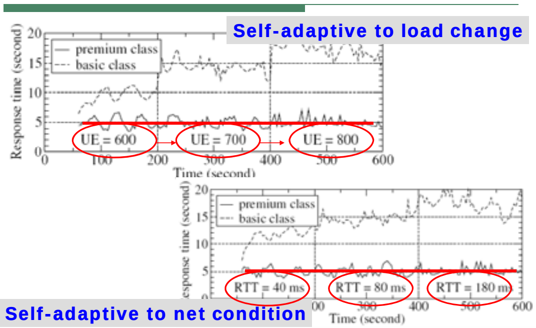

Self-adaptive to load change

Self-adaptive to net condition

Proactive Failure Management

Dealing with Failures

- Preventive measures
  - Keep enhancing system components reliability
  - Simplify systems design
  - Disable components which are prone to failures
    - BlueGene/L in LLNL disabled L1 cache in each node when jobs larger than a few hours were running
- Checkpoint-based recovery measure
  - Frequent checkpointing is costly and conservative
- Proactive management to handle failures before they occur

Proactive Failure Management

- Key is failure prediction/detection
  - Predict failure occurrences in the near future based on statistics of observed failures and dependence on performance states
- Opportunities
  - Failure occurrences display uneven inter-arrival time
  - They are correlated in time and space domains
  - But, no simple and general model for failure dynamics on production systems

Multi-time Scale Predictor

- Multiscale spherical covar for time loc.
- Aggregate model for spatial locality

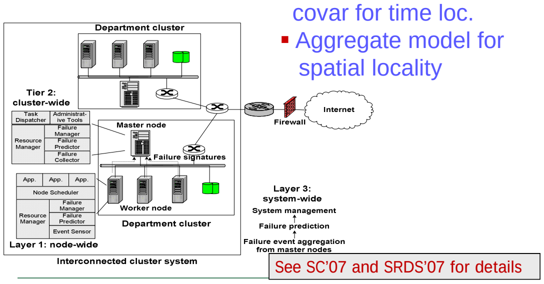

See SC'07 and SRDS'07 for details

Failure Prediction on LANL Trace

- 20 failures of Node 1

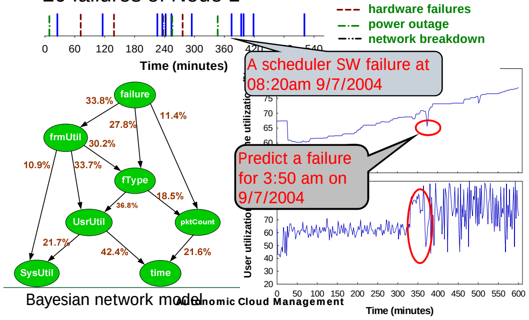

Bayesian network model

Online Failure Prediction

- On-line prediction from 5/12/2006 to 4/2/2007 on the Wayne State Grid.

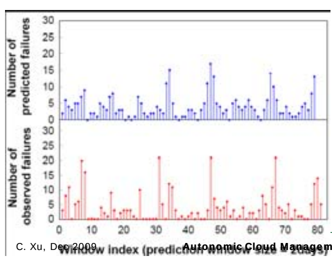

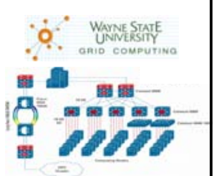

40 high-end compute servers in 3 clusters: ISC, CIT, CHM (16, 16, 8)

QoS toward Effective Cloud Comp

- Performance
  - Perf. isolation & differentiation
- Availability
  - Graceful perf degradation in stress conditions
  - Resilient to failure
- Security
  - ABC admission policy: Admit good, Block bad, and Contain suspicious ones.

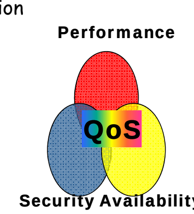

Anomaly Detection

System Anomaly Detection

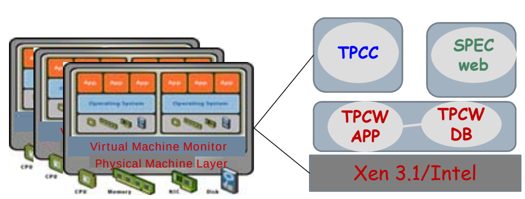

How to detect/predict system anomaly: overload, failure, SLO miss?

Dynamics of a Multi-tier Website

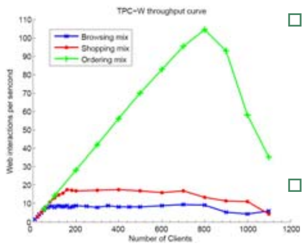

- Throughput is input traffic dependent
  - Browsing has burden on DB server
  - Ordering put extra work on App server
- Perf. bottleneck may shift between tiers

How to measure the system capacity?

Limitations of Response Time

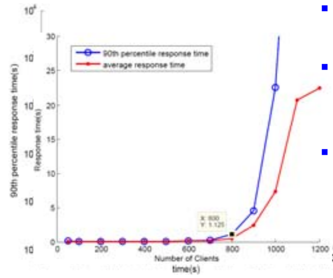

- Threshold setting for admission control?
- Long dead time: RT won't be available until request is completed
- Response time provides little insight into constrained resources for online trouble shooting.

Machine Learning for Anomaly Detection

- Use a set of OS or HW-level system perf. metrics as indicator
- Using machine learning to correlating the metrics to high-level system state in a model
- Given a snapshot of the low-level parameters, predict the system high-level state based on the model

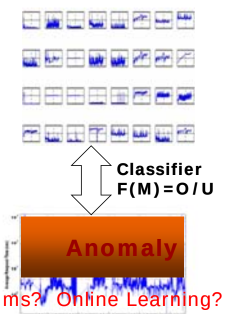

What parameters? What algorithms? Online Learning?

Coordinated ML for OL Detection

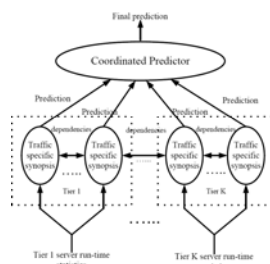

OS Level Parameters:

Appl Server
Dentunusd, kbbuffers kbcached,%util,%idle

Database Server
ldavg-1,renq-sz,%idle %user

HW Level Counters:

X87 FP RETIRED, X87 FP UOP, L2 REF, L2 MISS, INS RETIRED, CPU STALL, ITLB REF, ITLB MISS, RETIRED MISPRED BRANCH

Prediction Accuracy

- Overload prediction under various inputs
- Bottleneck detection under various inputs

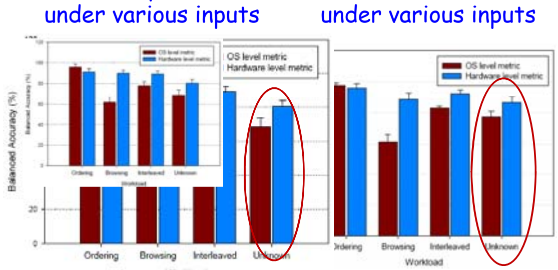

Prediction Agility

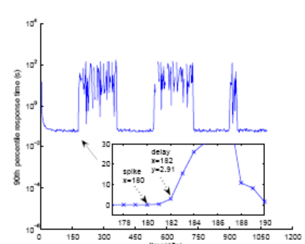

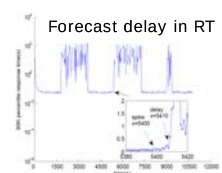

Forecast delay reduced from 10 sec to 2 sec

See IPDPS'08 and ICDCS'08 for details