AUTOMATIC ENERGY EFFICIENCY MANAGEMENT OF DATA CENTER RESOURCES BY LOAD-DEPENDENT SERVER ACTIVATION AND SLEEP MODES

Paul J. Kuehn
University of Stuttgart, Germany

Maggie Mashaly
German University in Cairo, Egypt

All 4 Green 2nd International Workshop on
Energy-Efficient Data Centers, Berkeley, Ca., May 21, 2013
(co-located with ACM e-Energy 2013, May 22-24, 2013)

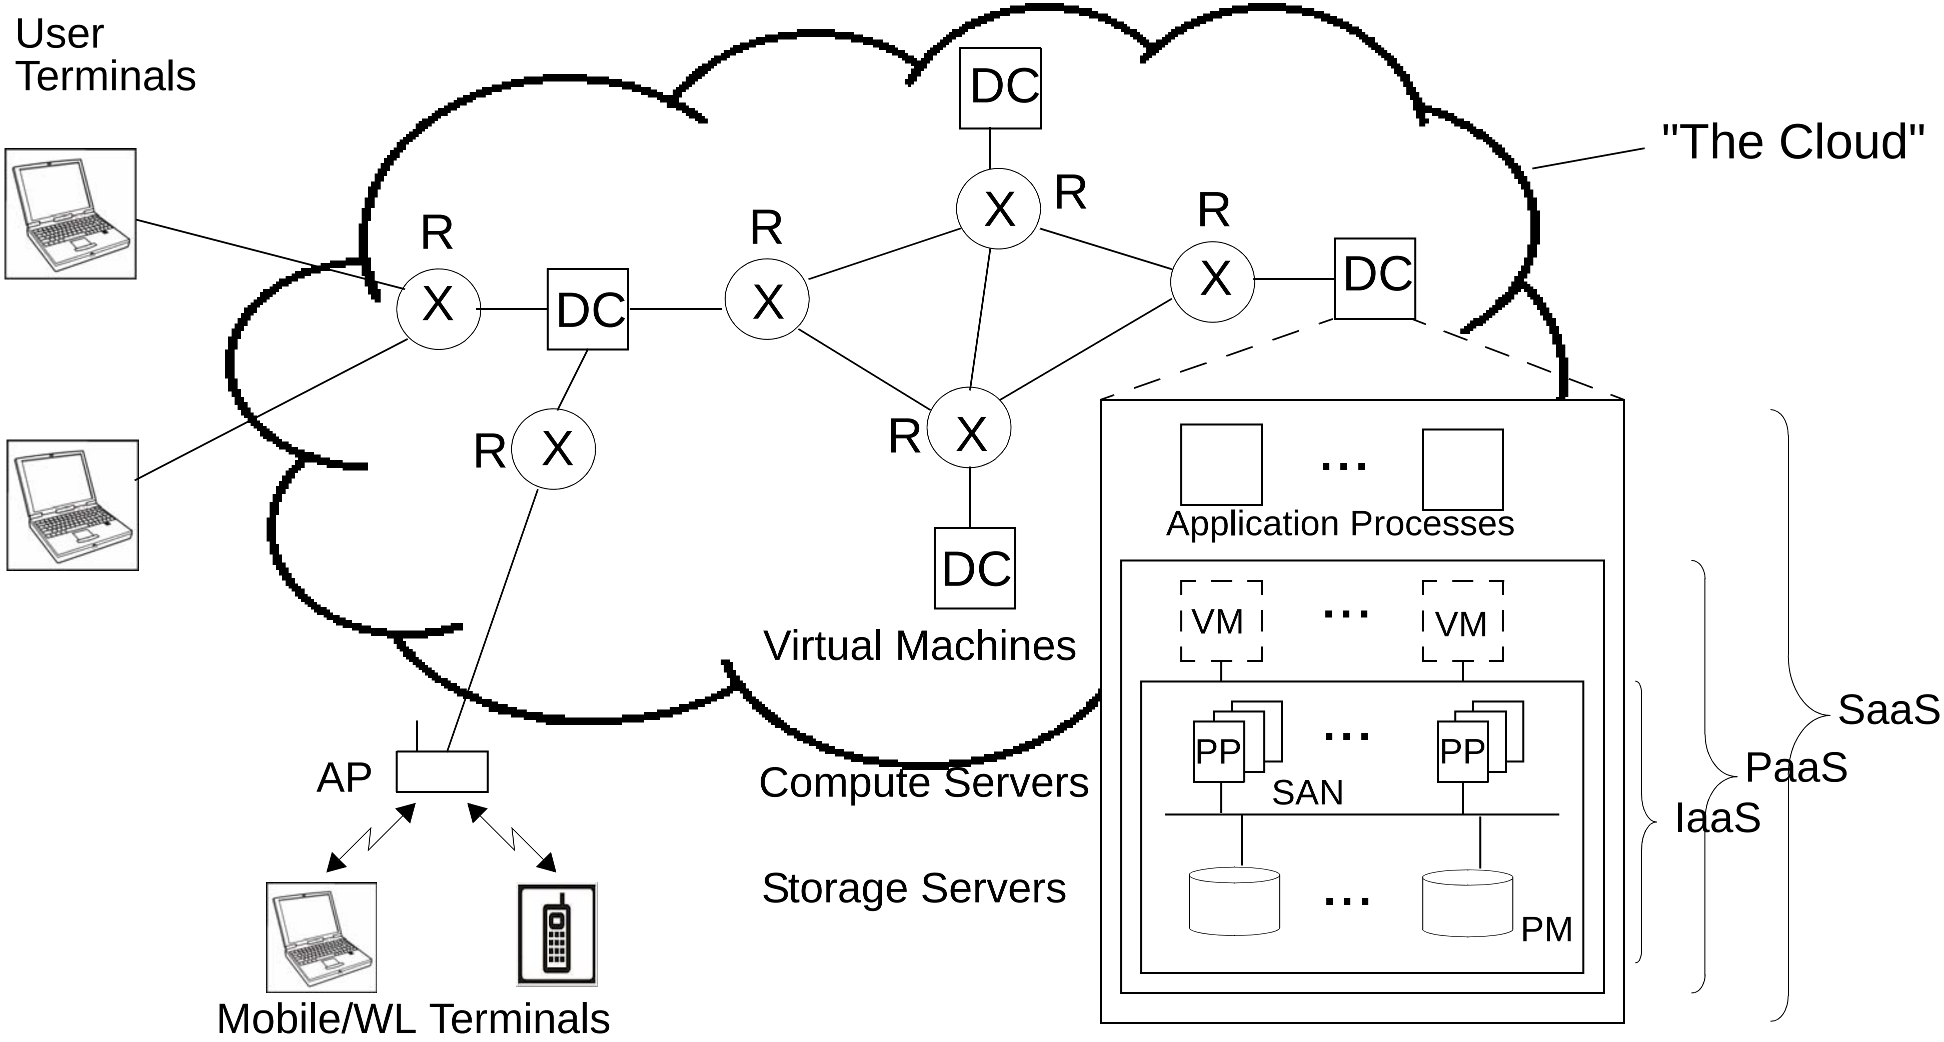
User
Terminals
"The Cloud"
DC
R
X
DC
X
DC
R
X
R
X
R
X
R
X
DC
Virtual Machines
Compute Servers
Storage Servers
AP
Mobile/WL Terminals
...
Application Processes
VM
...
VM
PP
...
PP
SAN
...
PM
SaaS
PaaS
IaaS

MANAGING CDNs - VIRTUALIZATION AND VM MIGRATION

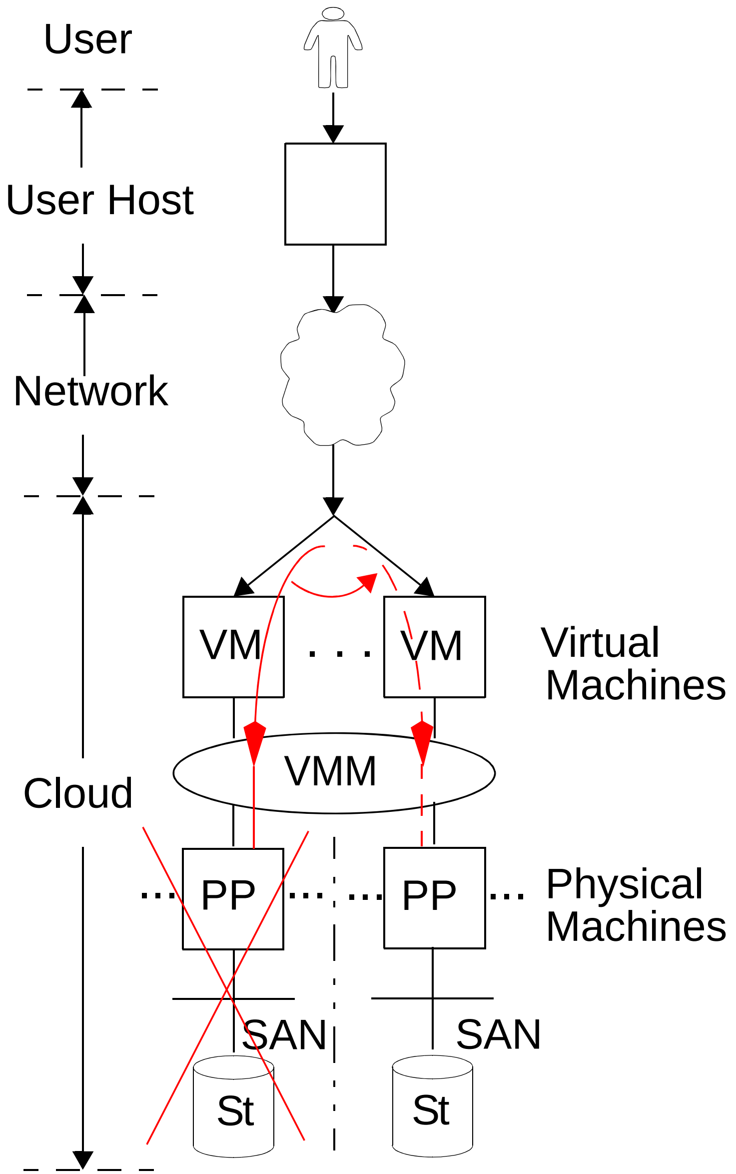

Cloud: Pool of Physical Resources
Interconnected by Network

VM: Virtual Machine
Virtualized View on the Resource Pool

VMM: VM Monitor ("Hypervisor")
Mapping of VM to PM

VM Migration:

- Change of Assignment VM --- PM
- Different Migration Strategies

  "Suspend-and-Copy"

  "Pre-Copy"

  "Post-Copy"

MANAGING CDNs - DYNAMIC PROVISIONING OF PHYSICAL RESOURCES

- Incentives

Hot Spot Mitigation	---->	Overload Avoidance
Load Balancing	---->	Economic Capacity Utilization, Energy Saving
Server Consolidation	---->	Avoiding "Sprawling" of Resources
Performance/SLA	---->	Meeting RT Requirements
Economics	---->	Trade-off between Storage Cost -- Communication Cost in Case of Content Storage Replication

- Content Location: Centralized or Decentralized

MODELING CLOUD DATA CENTERS (1)

Modeling Assumptions:

- Cloud with Distributed Data Centers
- NNC Address Resolution by Publish/Subscribe Service
- Multi-Server Model for DC Content Delivery
- Sleep Mode + Activation Delays for Multi-Core Nodes
- Self-Adapting Activation/Deactivation of Core Nodes within each DC (state-dependent; can be extended to Measurement- or Forecast-Based Operation)

MODELING CLOUD DATA CENTERS (2)

BASIC IDEAS:

- Self-Adapting Operation of Data Center Resources
  - Local Monitoring of Load Development
  - Local Control of Resource Activation/Deactivation by FSM

BASIC MODEL:

- Uniform Services, N Data Centers
- Focus on Processing Resources only
- (n_i, ρ_i) Resource/Utilization Vector of DC_i, $i \in [1, N]$

INDIVIDUAL DC MODEL

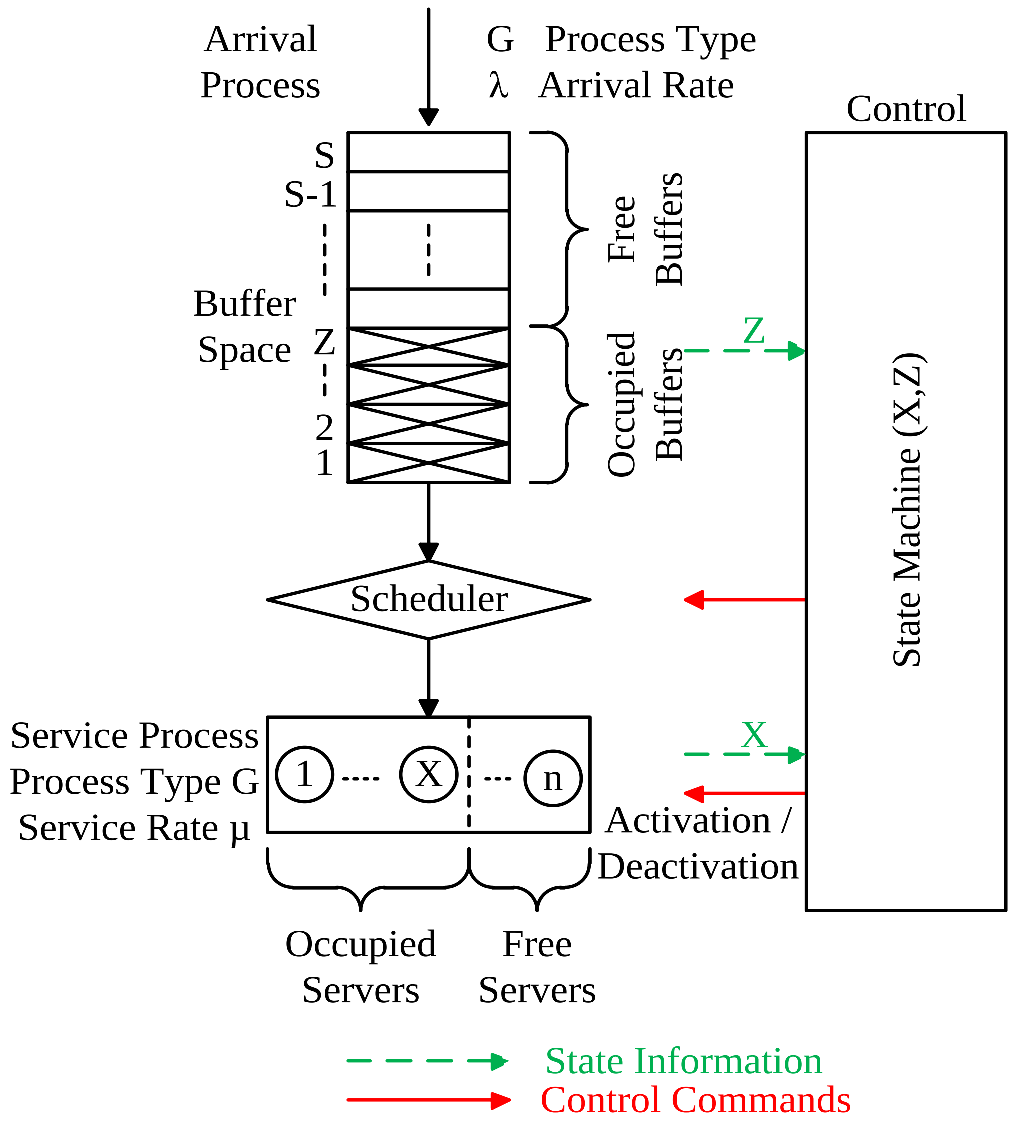

NON-ADAPTING MODEL BY FSM

(1) SINGLE HYSTERESIS MODEL

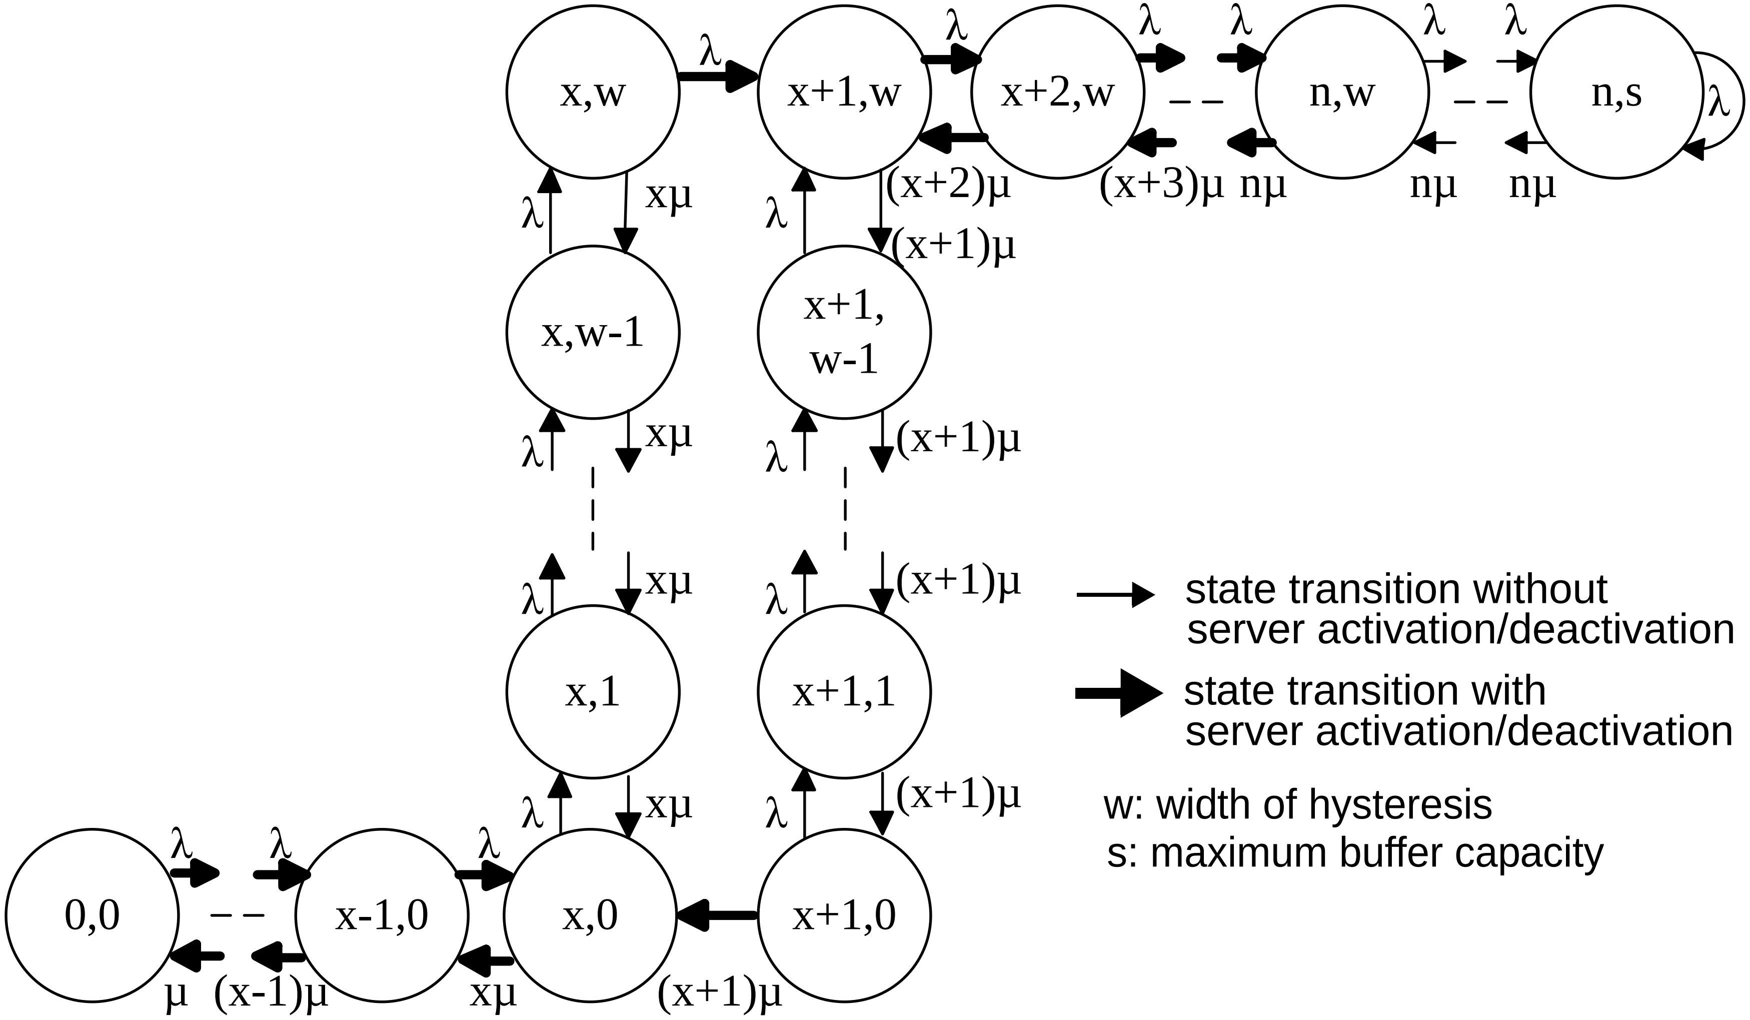

SELF-ADAPTING MODEL BY FSM

(2) MULTIPLE SERIAL HYSTERESIS MODEL

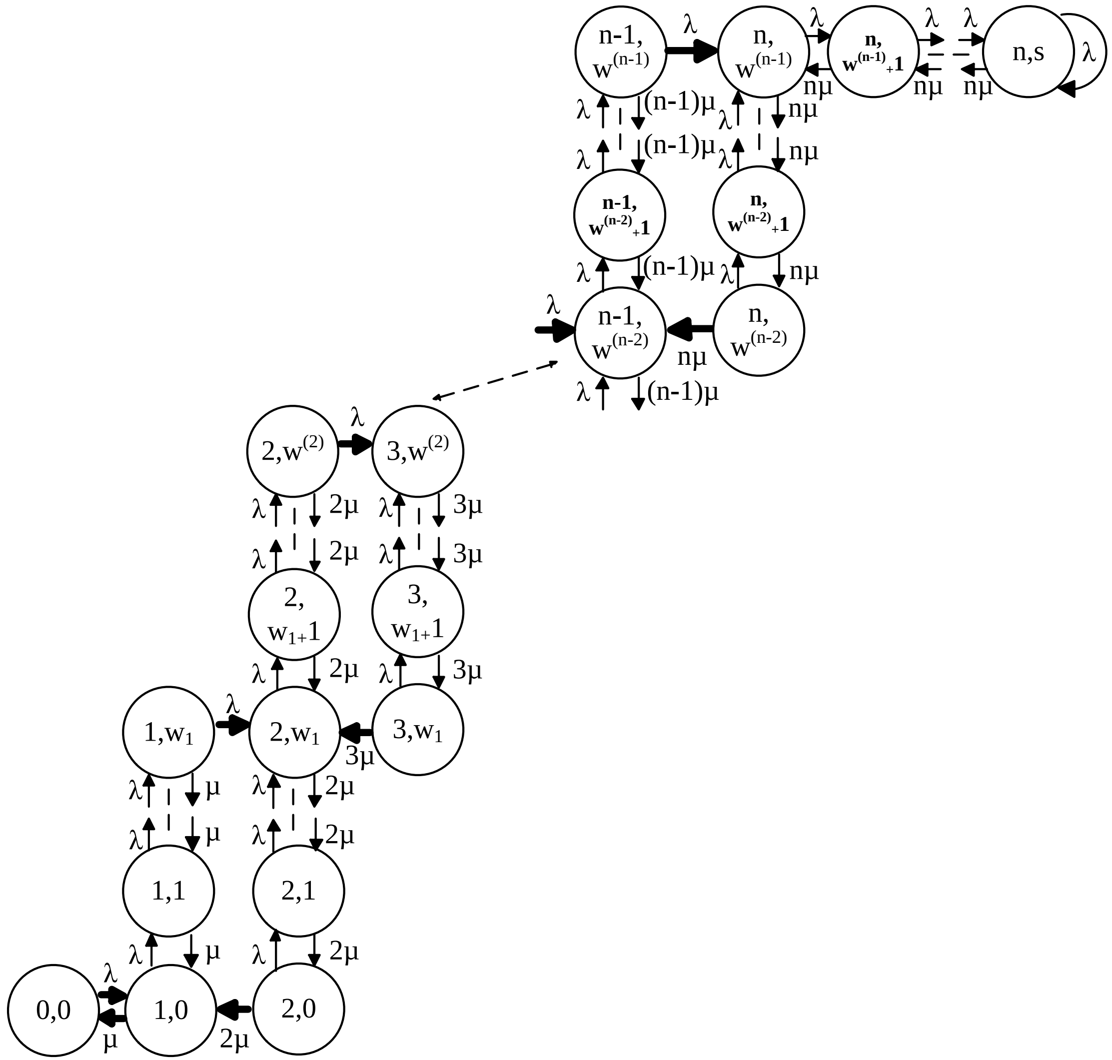

SELF-ADAPTING MODEL BY FSM / (3) MULTIPLE PARALLEL HYSTERESIS MODEL

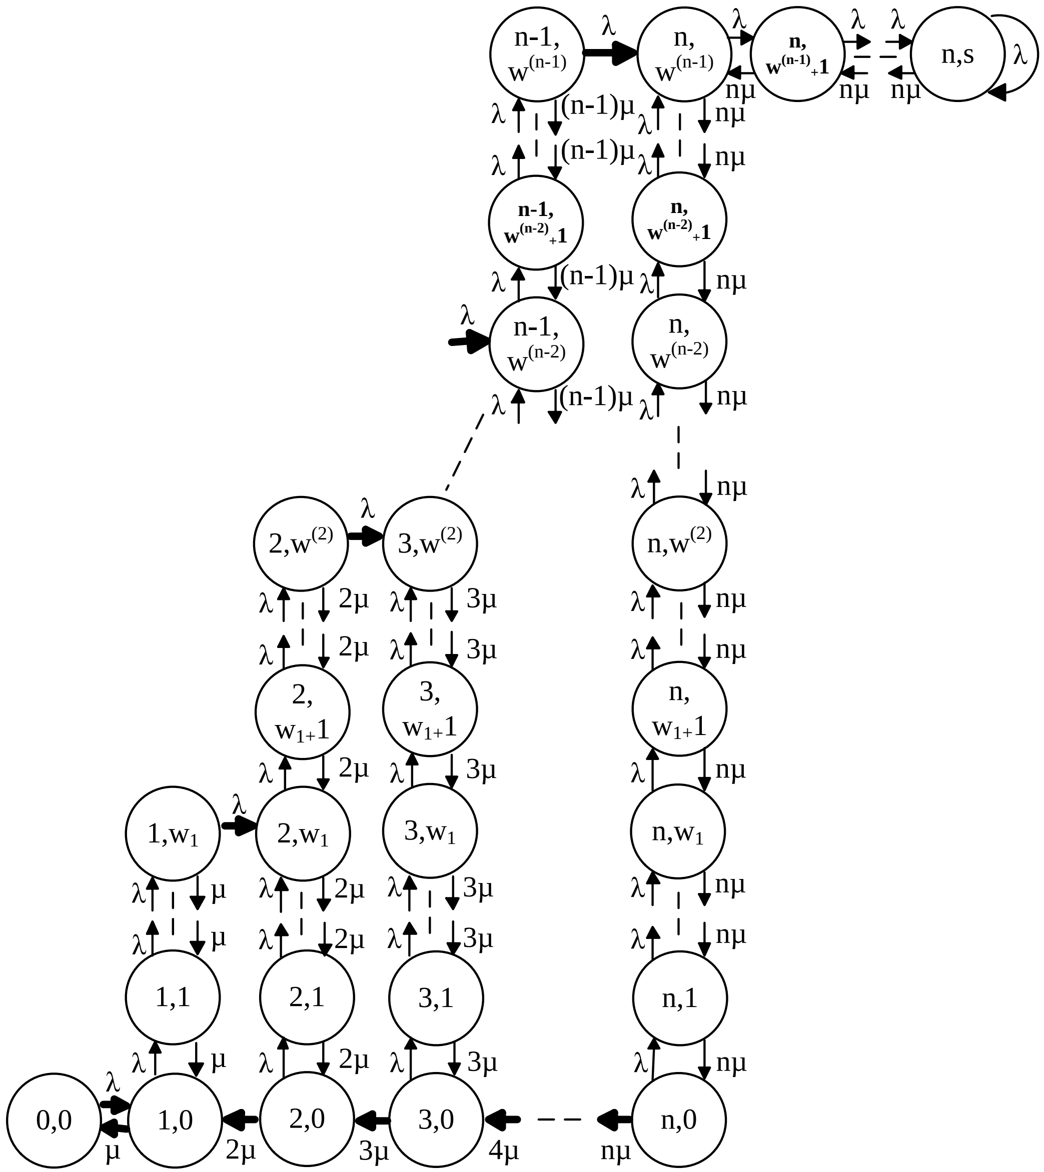

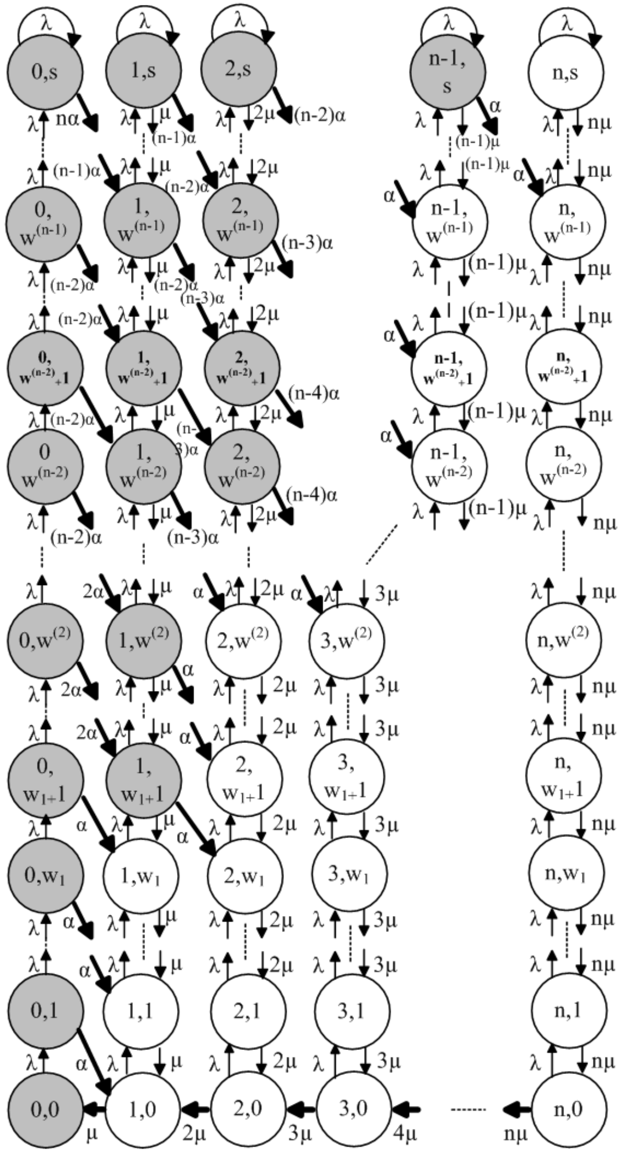

Multiple Parallel Hystereses
Multi-Server Queuing System
with/without Activation Overhead

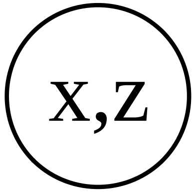

without Activation Overhead

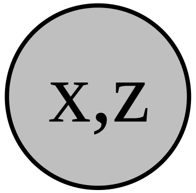

with Activation Overhead

MODELING CLOUD DATA CENTERS (8)

Conditions for the FSM Control:

- Multiple hysteresis thresholds for automatic adaptation to variable load
- Buffering of requests to throttle down frequent server activations
- Serving of tasks with maximum possible service rates by activated servers
- Throttling of server deactivations by Dynamic Frequency Scaling (DFS)
- Two server deactivation modes:
  - Server Cold Standby (CSB) ---> Booting required for activation
  - Server Hot Standby (HSB) (Sleeping Mode) ---> Warmup required for activation (Realized by Dyn. Voltage Scaling, DVS)
- All requirements can be met by a pseudo-2-dimensional FSM
- Exact analysis by fast recursive algorithm under Markovian traffic Assumptions
- Parameters:
  - λ task (job) arrival rate ($1/\lambda$ mean interarrival time)
  - μ task service rate ($1/\mu$ mean service time)
  - α server activation rate ($1/\alpha$ mean activation time for booting/warmup)
  - μ^* reduced service rate by DFS

MODEL ASSUMPTIONS

- Load-Dependent Activation / Deactivation of Resources -
- Multiple Parallel Hysteresis Model with Server Activation Overhead
- Server Activation: after Server Booting, Queue Threshold Crossing, Process Migration
- Server Deactivation: only when a Server Becomes Idle or the System Becomes Empty (Server Consolidation)
- Notations:

λ	Arrival Rate (Requests, Data Units, ...)
μ	Service Rate of a Server
α	Activation Rate of a Triggered Server Activation
ρ	Utilization Factor ($\rho = \alpha/\mu$)
$E[T_W \mid T_W>0]$	Mean Waiting Times of Delayed Requests
R_A	Server Activation/Deactivation Rate
$W(>t)/W$	Compl. DF of Buffered Requests

PERFORMANCE ANALYSIS AND RESULTS (2)

NUMERICAL EVALUATION

- 1st Choice: Based on the fundamental solutions of Ibe/Keilson by Green’s Function

  - **Result:** Numerically too complex

- 2nd Choice: Based on the fundamental solutions of Lui/Golubchik by Stochastic Complement Analysis

  - **Result:** Numerically too complex

- 3rd Choice: New solution by iterative recursions

  - **Result:** Extremely fast and numerically stable
    Extension to DF of delays
    Optimization wrt performance constraints

- Extensions

  In all solution methods certain generalizations are possible as

  – bulk arrivals
  – inclusion of activation overhead
  – inclusion of look-ahead activations

PERFORMANCE ANALYSIS AND RESULTS (3)

NUMERICAL RESULTS (without Activation Overhead)

MULTIPLE SERIAL HYSTERESIS MODEL **Probabilities of State**

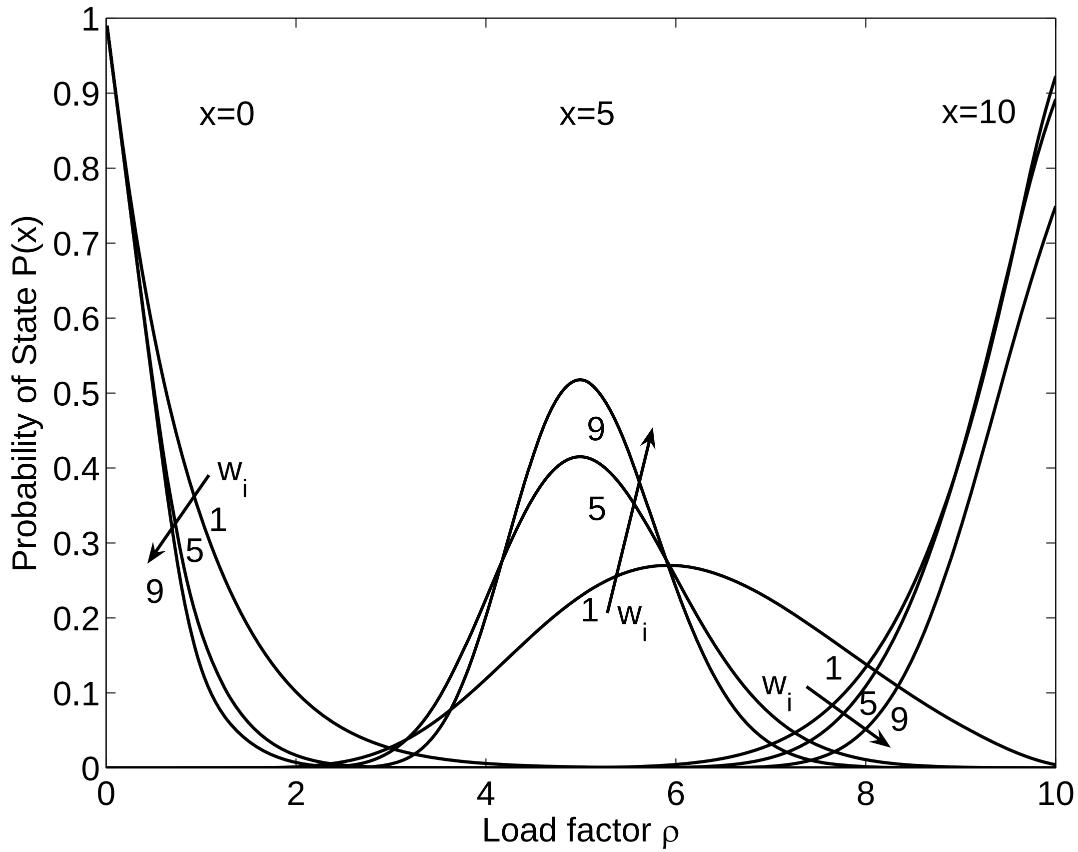

NUMERICAL RESULTS (without Activation Overhead)

MULTIPLE SERIAL HYSTERESIS MODEL **Server Activation / Deactivation Rate**

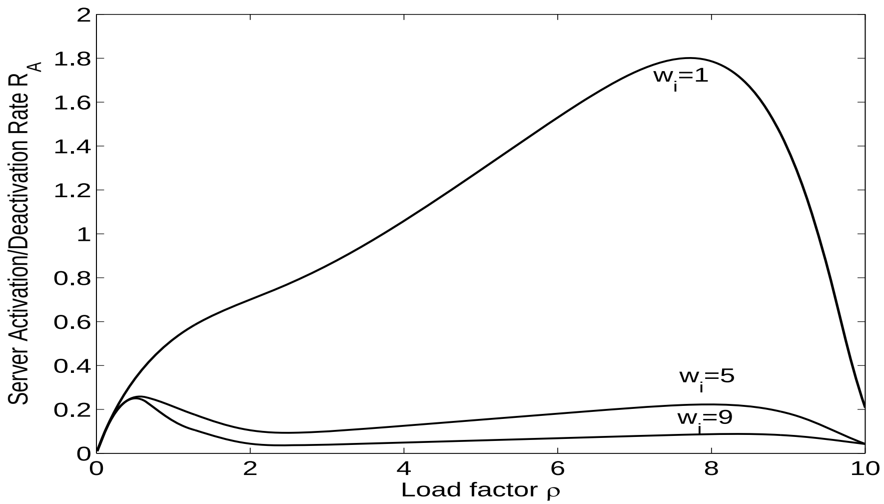

PERFORMANCE ANALYSIS AND RESULTS (5)

NUMERICAL RESULTS (without Activation Overhead)

MULTIPLE SERIAL HYSTERESIS MODEL **Mean Waiting Time of Delayed Requests**

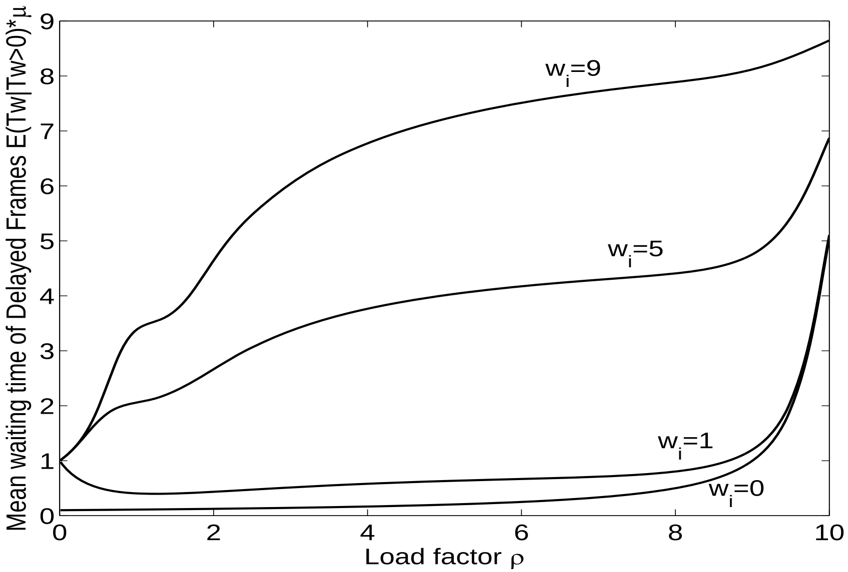

NUMERICAL RESULTS (without Activation Overhead)

MULTIPLE PARALLEL HYSTERESIS MODEL **Mean Waiting Time of Delayed Requests**

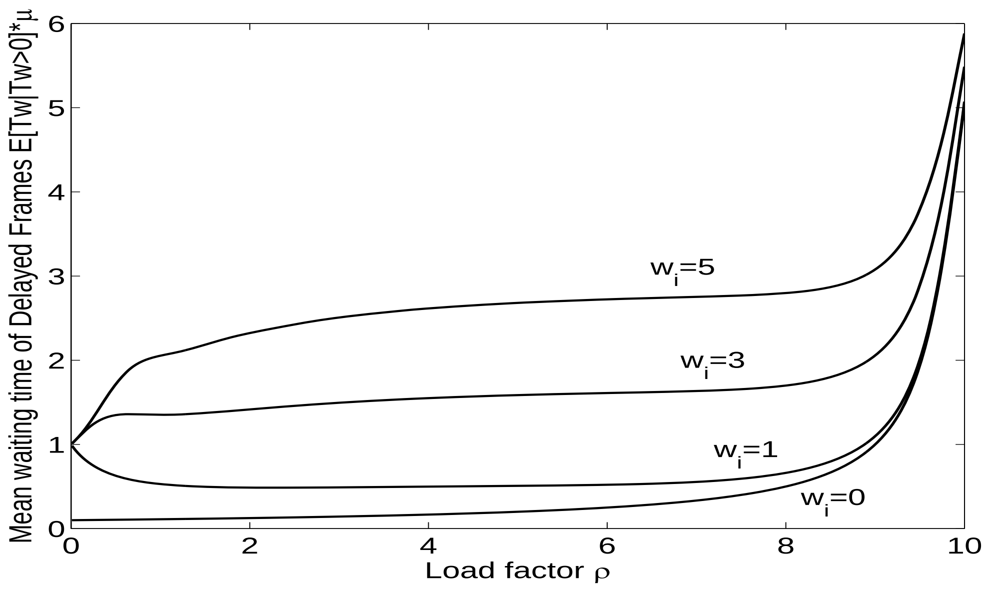

NUMERICAL RESULTS (without Activation Overhead)

MULTIPLE PARALLEL HYSTERESIS MODEL **Compl. DF of Buffered Requests**

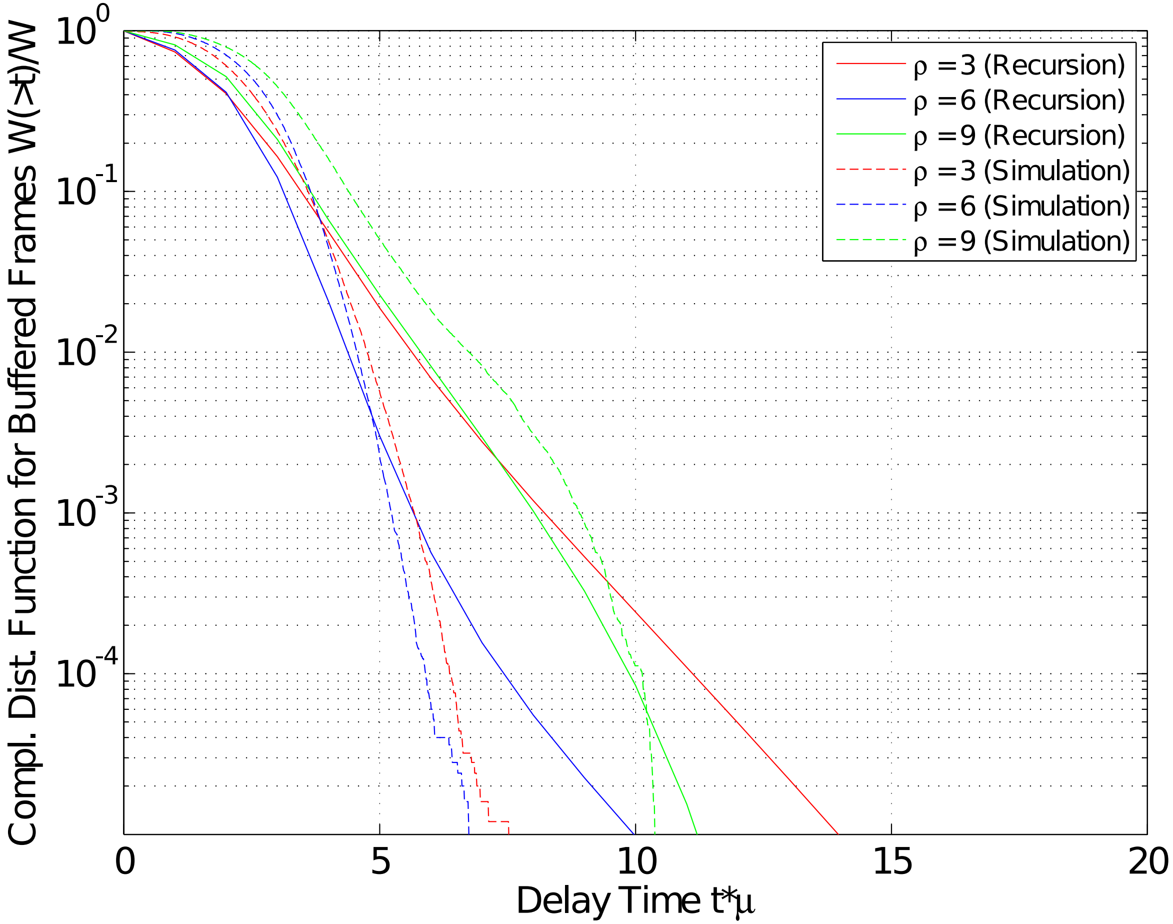

PERFORMANCE ANALYSIS AND RESULTS (8)

NUMERICAL RESULTS (with Activation Overhead): *Probability State Distributions*

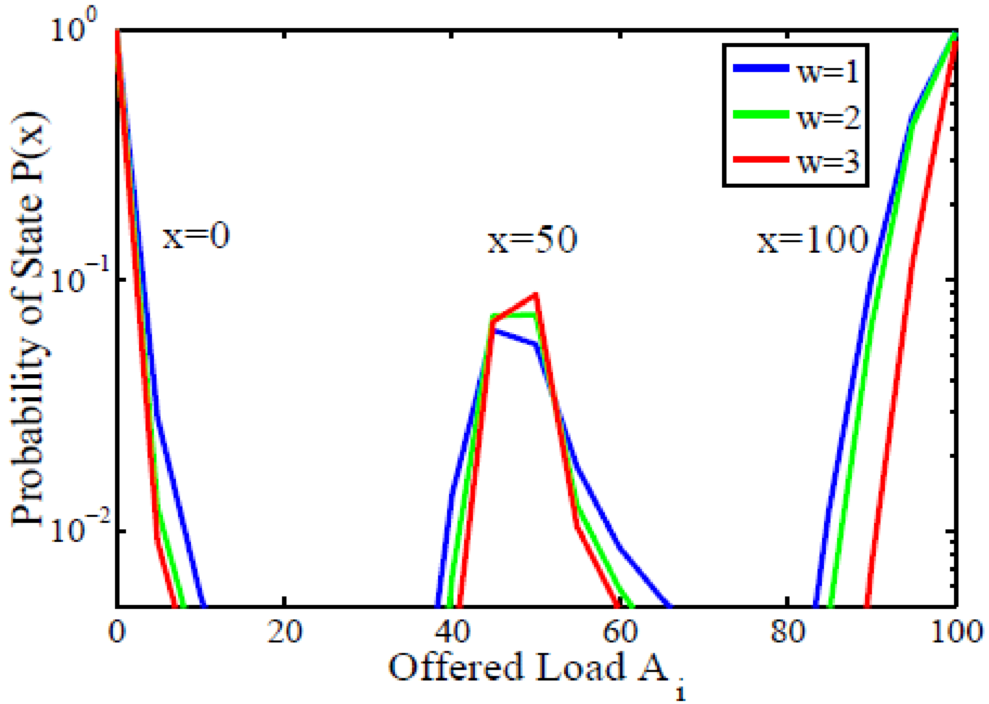

Figure: Probability of 'x' active servers vs. offered load
n = 100, s = 300, α = 1, variable w

NUMERICAL RESULTS (with Activation Overhead): *Server Activation Rate*

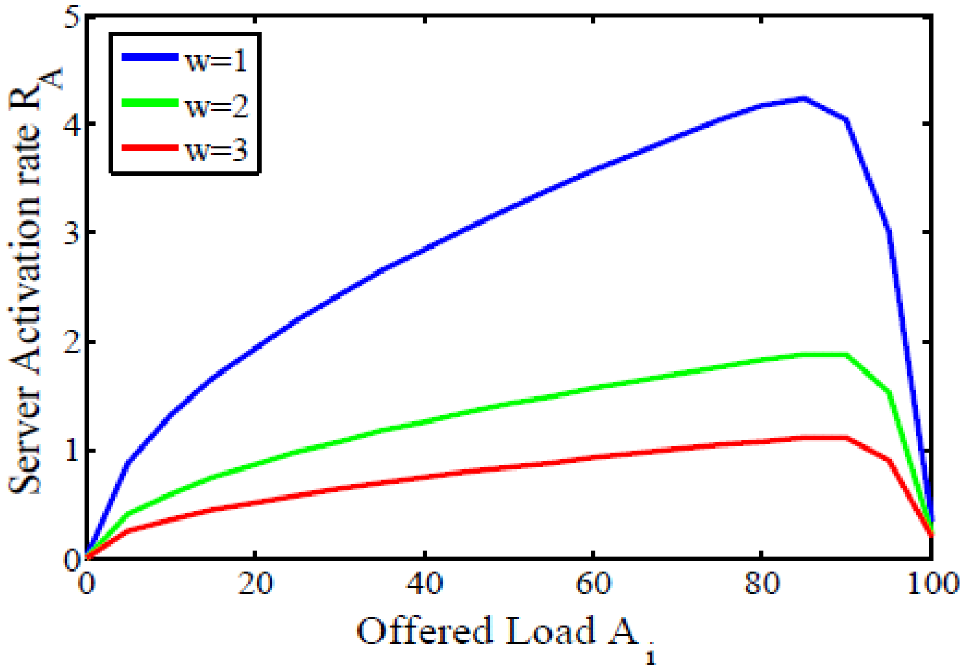

Figure: Server activation rate vs. offered load

n = 100, s = 300, α = 1, variable w

NUMERICAL RESULTS (with Activation Overhead): *Mean Delay*

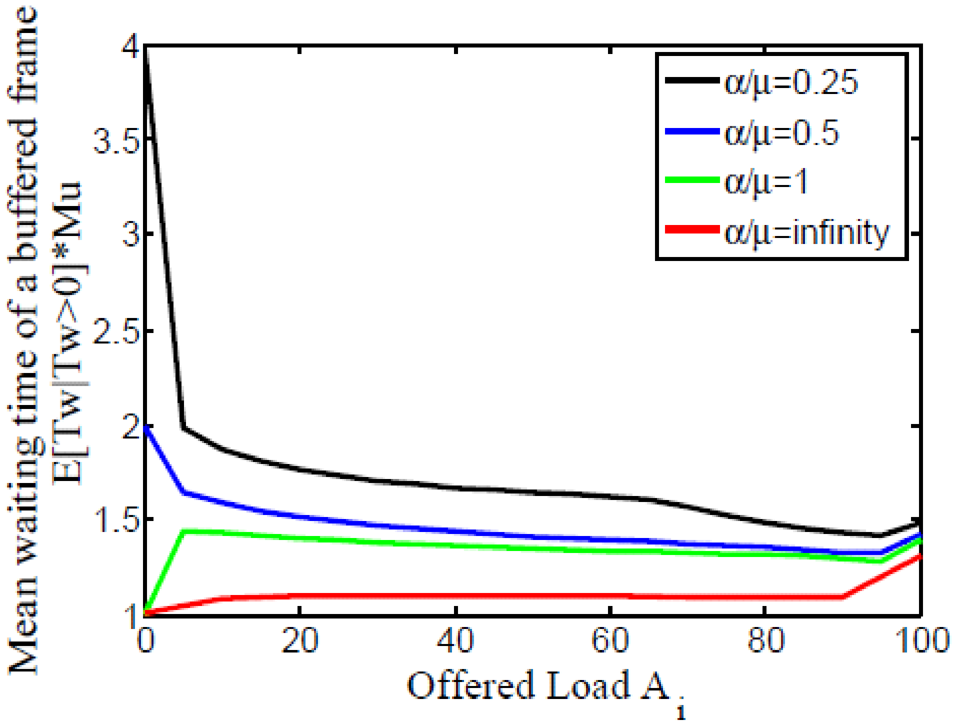

Figure: Mean delay of delayed frames vs. offered load

n = 100, s = 200, w = 2, variable α/μ

PERFORMANCE ANALYSIS AND RESULTS (11)

NUMERICAL RESULTS (Cold/Hot Standby, with Activation Overhead): *Power-Saving Efficiency*

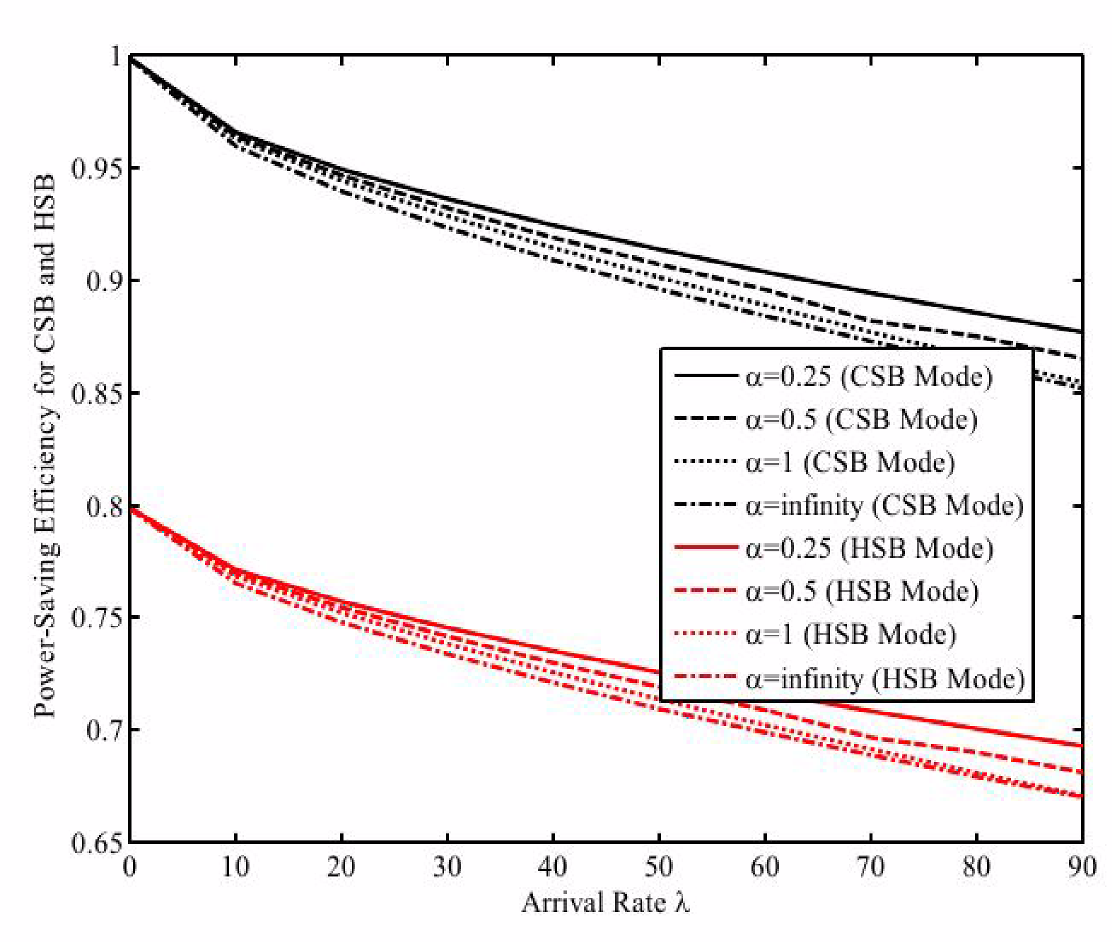

SUMMARY AND OUTLOOK

- Cloud Server Virtualization Allows for Flexible Content Distribution and Access
- Models for Self-Adapting DC Server Activation/Deactivation (Server Consolidation)
- Trade-off between Power Reduction and Performance
- Extension to Include Activation Overhead, Hot/Cold Standby, and DVFS
- Recursive Algorithms for Markovian Assumption by new Recursive Algorithms

Outlook

- Realistic Cloud Application Classes
- Refined Models for DC Architectures and Operations
- Cost Optimization