Statistical Modeling of Spot Instance Prices in Public Cloud Environments

Bahman Javadi , Ruppa K. Thulasiram , and Rajkumar Buyya

Cloud Computing and Distributed Systems (CLOUDS) Laboratory
Department of Computer Science and Software Engineering,
The University of Melbourne, Australia

Computational Financial Derivatives (CFD) Laboratory
Department of Computer Science,
University of Manitoba, Canada

4th IEEE/ACM International Conference on Utility and Cloud Computing (UCC 2011)

Agenda

- Introduction
- Modeling Approach
- Patterns of Spot Prices
- Global Statistics and Analysis
- Distribution Fitting
- Model Calibration
- Model Validation
- Conclusions

Introduction

- Cloud computing
  - Pay-as-you-go price model
  - Price-Performance trade-off
- New trade-offs
  - Fixed-price vs. Variable-price
  - Service Level Agreement (SLA)
- Amazon's EC2
  - On-demand instances
  - Reserved instances
  - Spot instances (SI)

SPOT INSTANCES

- Sell the Idle cycles of Amazon's data centers
- Price depends on VM demand within a data center
- Low price, but less reliability
- Competitive bidding option
  - Bid : max. price to be paid for an hour

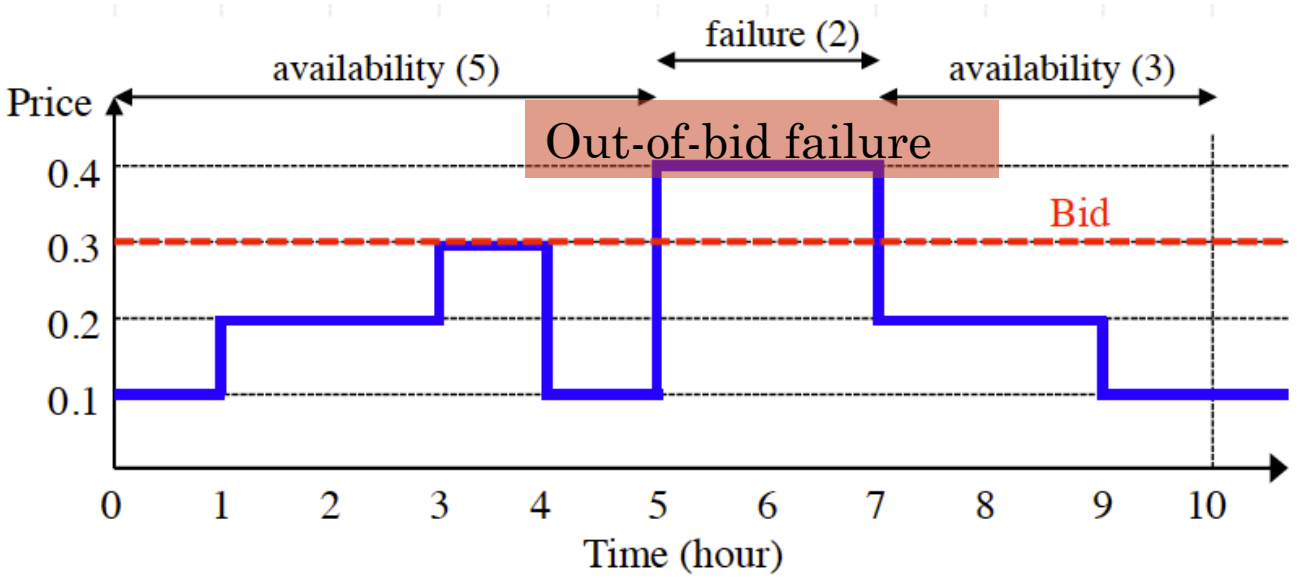

Research Goal

- Statistical Analysis of SIs
- Statistical Modeling of Sis
  - Generate synthetic prices
  - Design of stochastic scheduling algorithms
  - Fault tolerant mechanisms
    - Checkpointing
    - Replication
- IaaS Cloud providers that look forward to offer such a service in the near future

Modeling Approach

- Spot price (P_i)
- Inter-price Time ($T_i = t_{i+1} - t_i$)
- Price history: Feb. 2010-Feb.2011
  - Modeling Trace: Feb.2010-Nov.2010
  - Validation Trace: Dec.2010-Feb.2011
  - eu-west region, 8 Spot instances, Linux OS

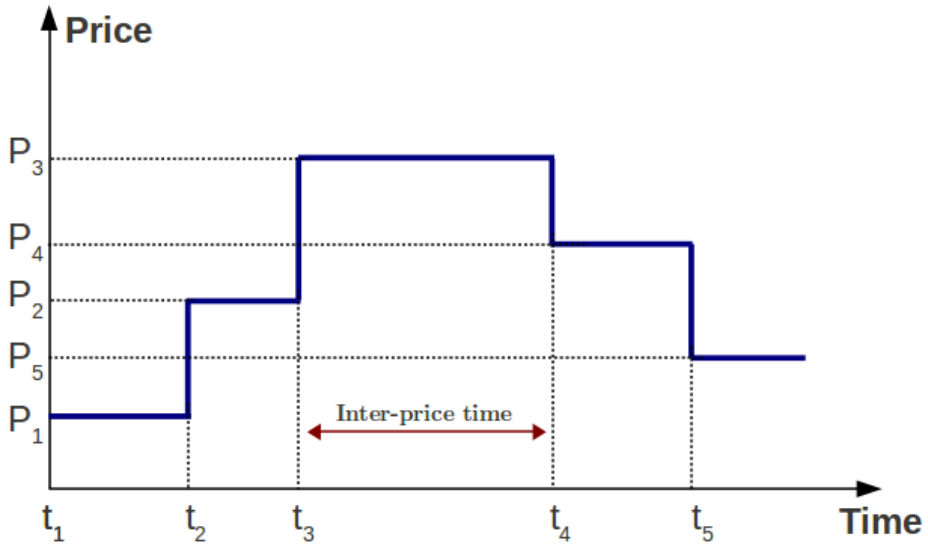

Patterns of Spot Prices

- Hour-in-day (eu-west)

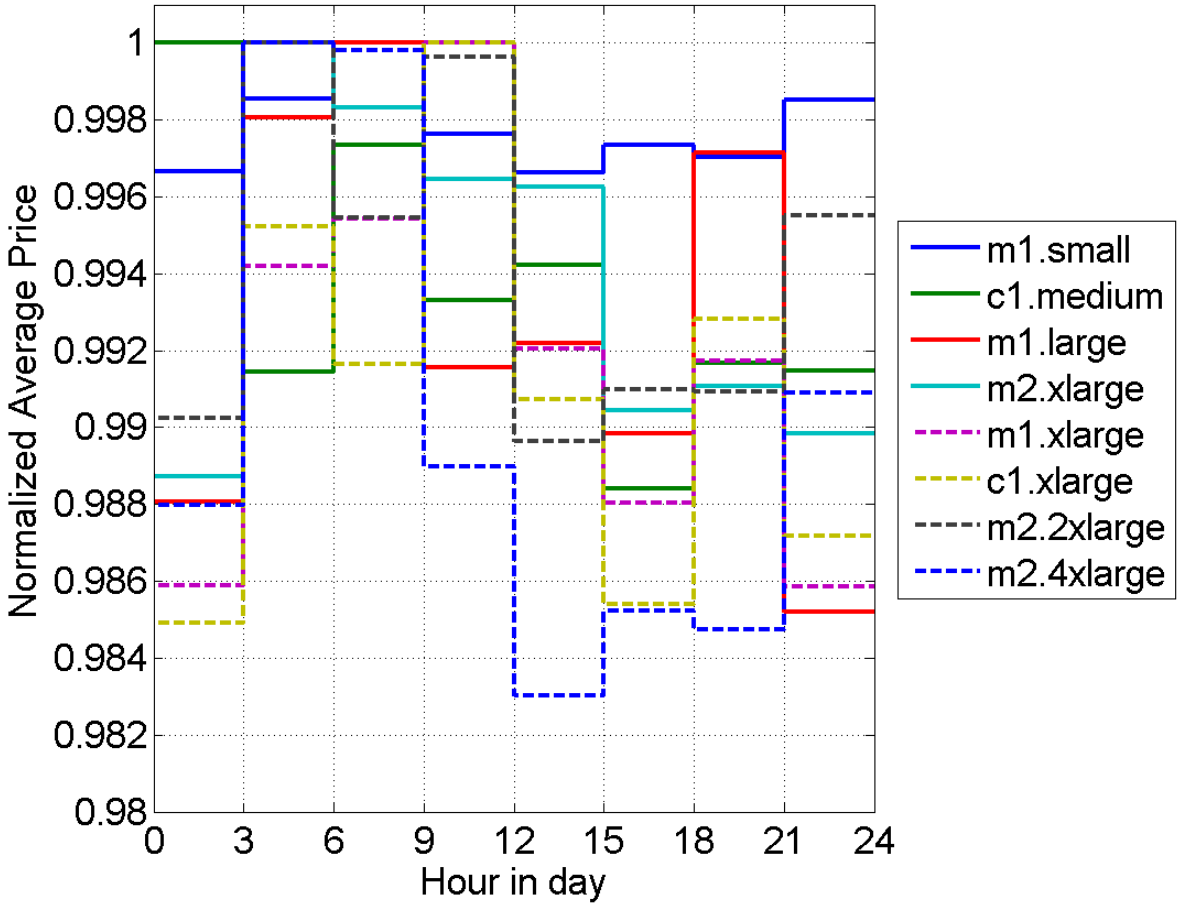

Patterns of Spot Prices

- Hour-in-day (us-east)

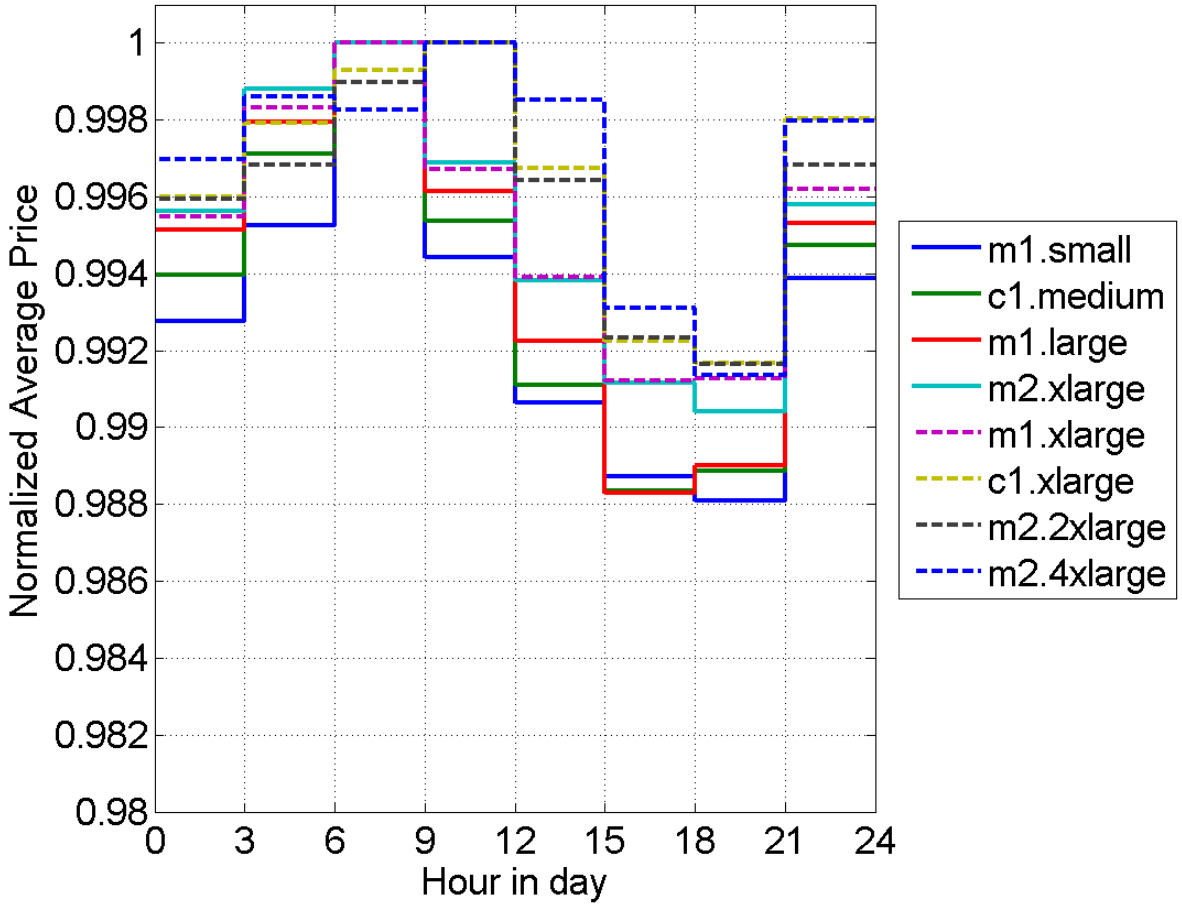

Patterns of Spot Prices

- Day-of-week (eu-west)

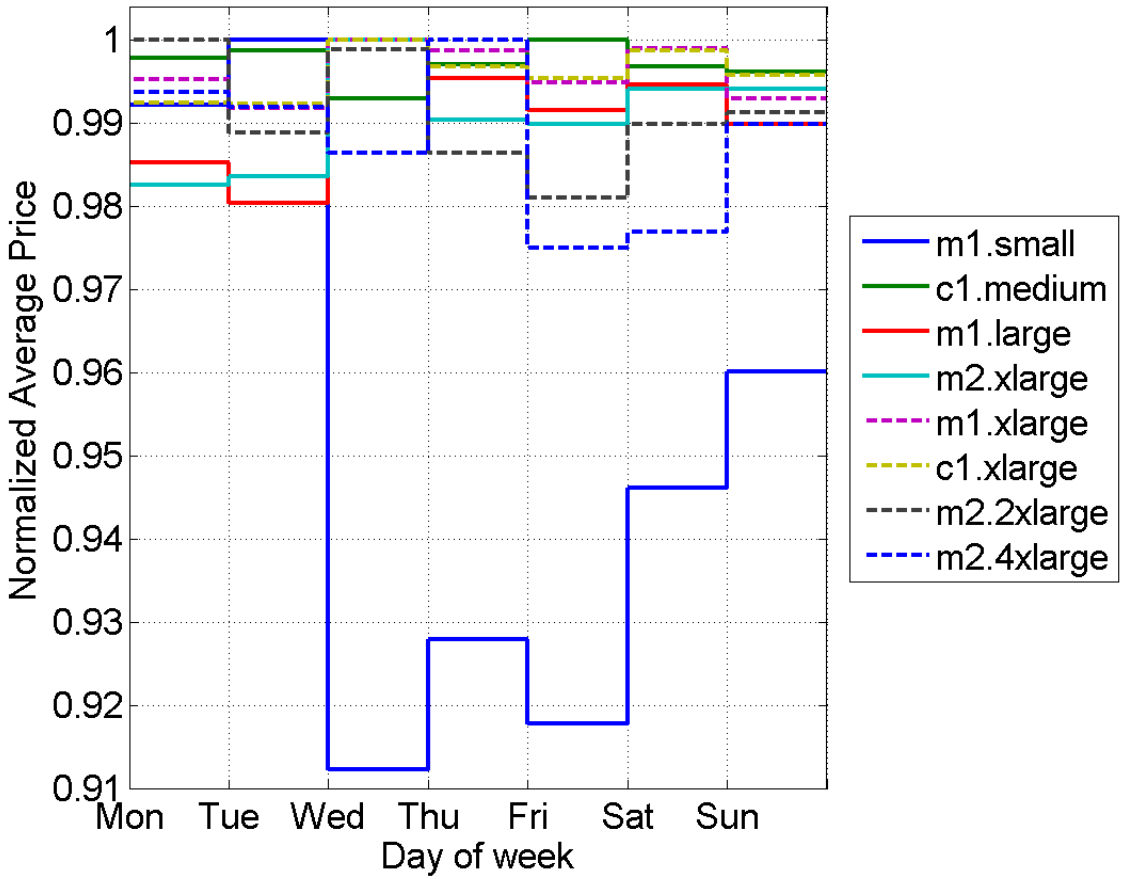

Patterns of Spot Prices

- Day-of-week (us-east)

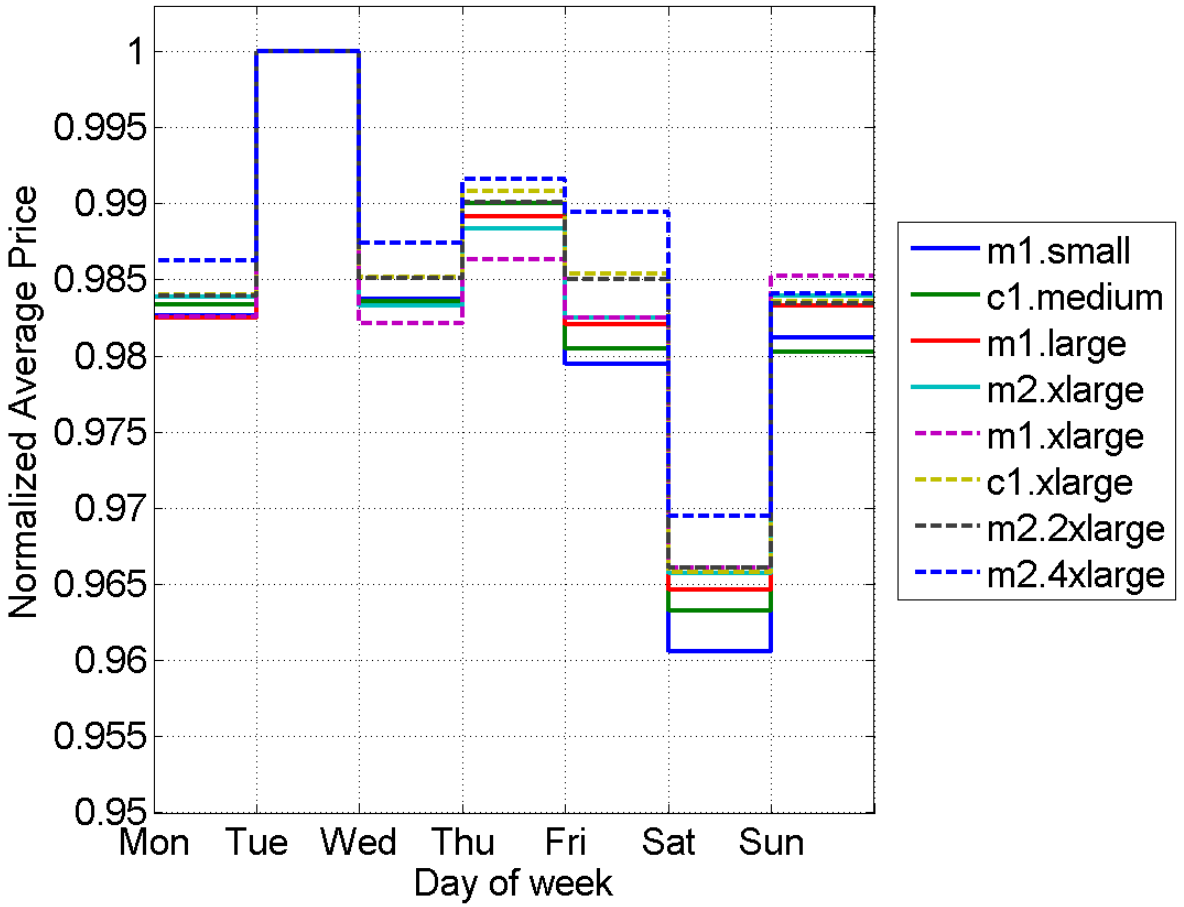

Global Statistics and Analysis

- Statistics for Spot prices (cents)

Instances	Mean	TrMean	Median	Std	CV	IQR	Max	Min	Skewness	Kurtosis	No.
m1.small	4.00	4.00	4.00	0.19	0.05	0.20	9.50	3.80	9.44	242.97	3702
c1.medium	8.00	8.00	8.00	0.27	0.03	0.40	10.10	7.60	0.28	3.91	3812
m1.large	16.04	16.02	16.10	0.85	0.05	1.00	50.00	15.20	21.55	792.41	3875
m2.xlarge	24.04	24.03	24.10	1.03	0.04	1.40	57.10	22.80	12.91	387.69	3763
m1.xlarge	32.05	32.01	32.10	1.60	0.05	2.00	76.00	30.40	15.34	415.47	3917
c1.xlarge	32.04	32.03	32.10	1.07	0.03	2.00	45.00	30.40	0.54	8.27	3658
m2.2xlarge	56.04	56.04	56.20	1.83	0.03	3.42	76.00	53.20	0.25	4.99	4001
m2.4xlarge	112.08	112.08	112.50	3.62	0.03	6.80	150.00	106.40	0.21	4.55	3912

- Statistics for Inter-price time (hours)

Instances	Mean	TrMean	Median	Std	CV	IQR	Max	Min	Skewness	Kurtosis	No.
m1.small	1.96	1.61	1.35	2.66	1.35	0.30	109.08	0.02	19.94	727.54	3701
c1.medium	1.91	1.59	1.34	1.86	0.97	0.32	22.81	0.02	4.53	30.63	3811
m1.large	1.88	1.57	1.33	1.79	0.95	0.31	30.94	0.02	5.02	42.02	3874
m2.xlarge	1.79	1.53	1.34	1.56	0.87	0.30	22.83	0.02	4.93	38.54	3762
m1.xlarge	1.86	1.58	1.34	1.78	0.96	0.31	38.20	0.02	7.34	101.43	3916
c1.xlarge	1.99	1.56	1.34	7.22	3.63	0.30	378.19	0.02	44.38	2169.40	3657
m2.2xlarge	1.82	1.55	1.33	1.60	0.88	0.31	29.02	0.02	5.11	45.75	4000
m2.4xlarge	1.86	1.58	1.34	1.71	0.92	0.31	26.51	0.02	5.20	44.28	3911

Distribution Fitting

- Probability Density Function (PDF)
  - Example: c1.medium

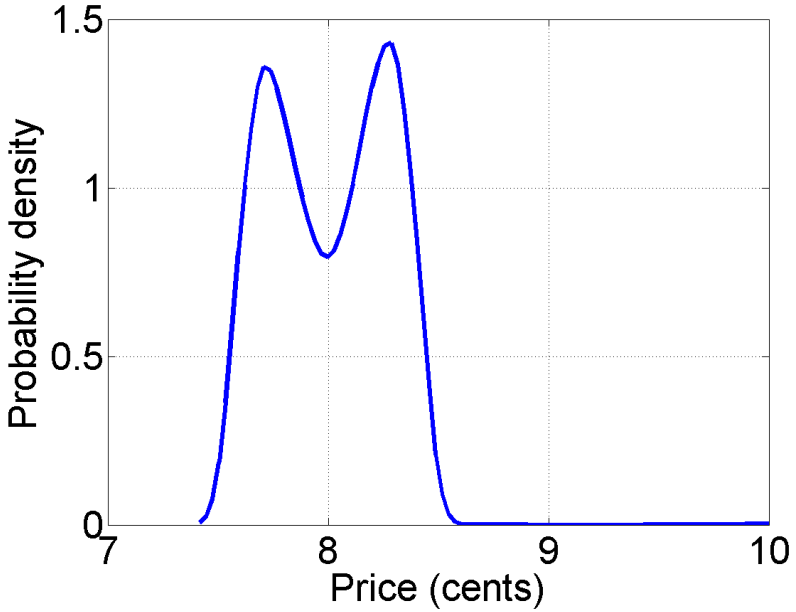

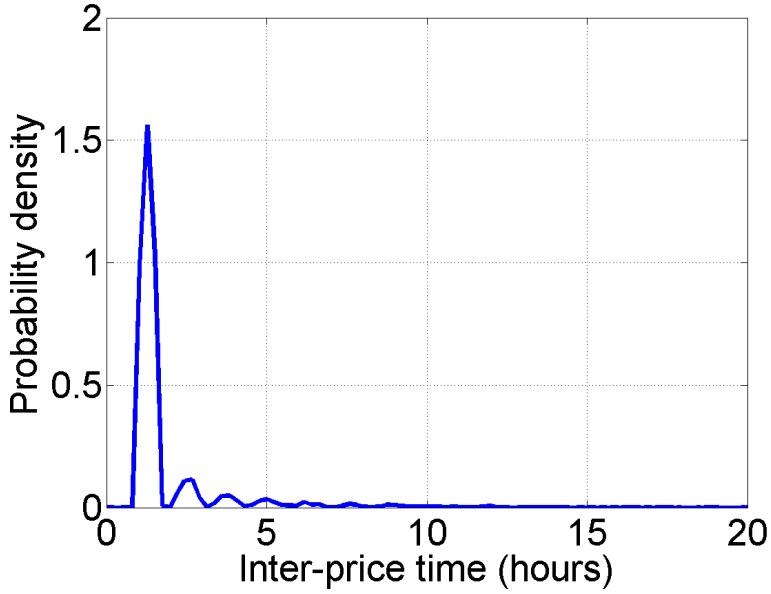

- Bi-modality (Multi-modality)

Distribution Fitting

- Mixture of Gaussians (MoG)

$$cdf(x; k, \vec{p}, \vec{\mu}, \vec{\sigma^2}) = \sum_{i=1}^{k} \frac{p_i}{2} \left(1 + erf\left(\frac{x - \mu_i}{\sigma_i \sqrt{2}}\right)\right)$$

- Parameter estimation
  - Model Based Clustering (MBC)
  - $2 \leq k \leq 4$
  - Goodness of Fit (GOF) tests
    - Kolmogorov-Smirnov (KS)
    - Anderson-Darling (AD)
    - Graphical tests (Probability-Probability Plot)

Distribution Fitting

- Graphical test
  - Probability-Probability Plot (PP Plot)
  - Example: Spot price for m2.xlarge

- k=2

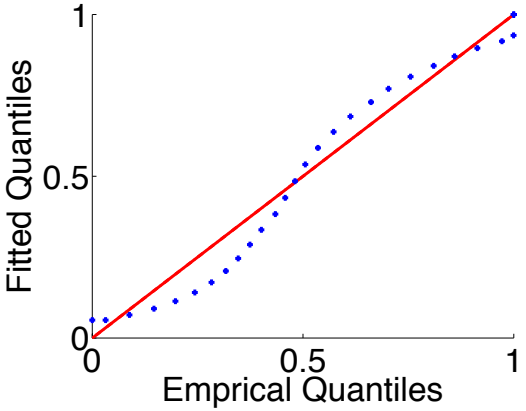

k=3

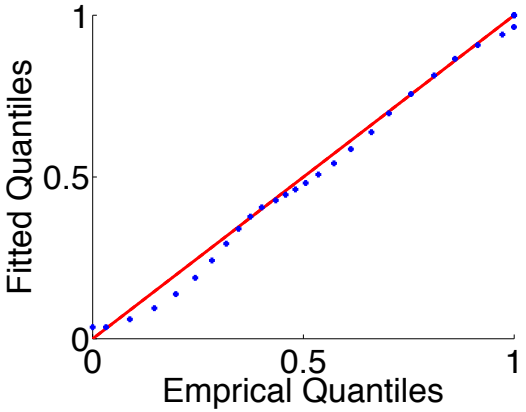

k=4

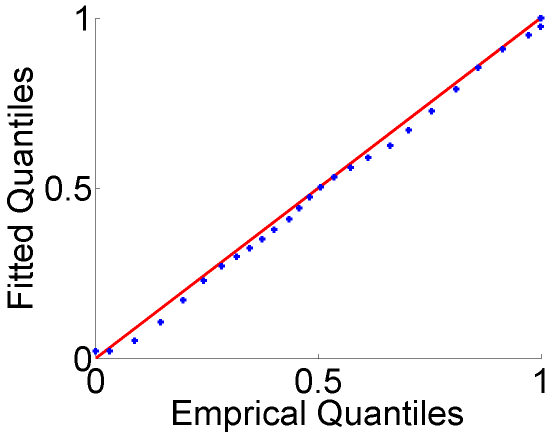

Distribution Fitting

- Graphical test
  - Probability-Probability Plot (PP Plot)
  - Example: Inter-price time for c1.xlarge
  - k=2

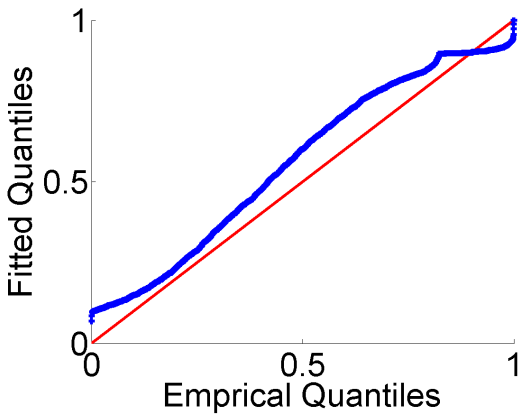

k=3

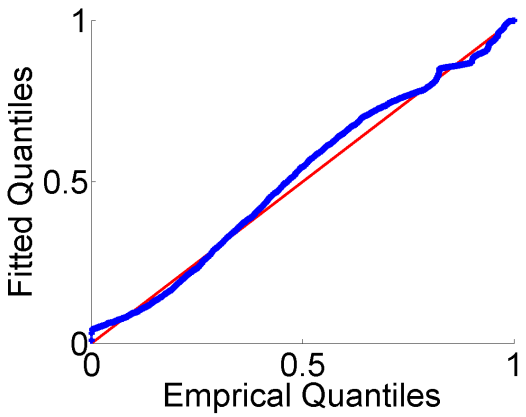

k=4

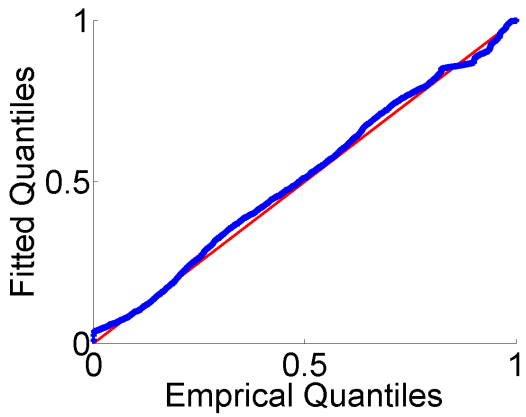

Distribution Fitting

- GoF tests for Spot price

Instances	MoG ($k = 2$)	MoG ($k = 3$)	MoG ($k = 4$)
m1.small	0.016 0.791	0.017 0.789	0.053 0.803
c1.medium	0.211 0.779	0.217 0.791	0.224 0.790
m1.large	0.113 0.678	0.319 0.752	0.354 0.754
m2.xlarge	0.139 0.616	0.356 0.721	0.415 0.734
m1.xlarge	0.134 0.570	0.369 0.708	0.431 0.706
c1.xlarge	0.394 0.681	0.444 0.705	0.421 0.707
m2.2xlarge	0.420 0.648	0.469 0.682	0.450 0.672
m2.4xlarge	0.429 0.617	0.463 0.637	0.476 0.653

- GoF tests for Inter-Price time

Instances	MoG ($k = 2$)	MoG ($k = 3$)	MoG ($k = 4$)
m1.small	0.347 0.476	0.415 0.592	0.489 0.627
c1.medium	0.382 0.546	0.390 0.566	0.380 0.566
m1.large	0.390 0.552	0.387 0.573	0.400 0.574
m2.xlarge	0.389 0.556	0.393 0.566	0.405 0.585
m1.xlarge	0.369 0.526	0.391 0.564	0.406 0.581
c1.xlarge	0.221 0.319	0.399 0.561	0.467 0.602
m2.2xlarge	0.376 0.532	0.426 0.570	0.463 0.610
m2.4xlarge	0.368 0.529	0.383 0.569	0.395 0.573

Model Calibration

- Time evolution of Spot prices
  - Example: m2.xlarge

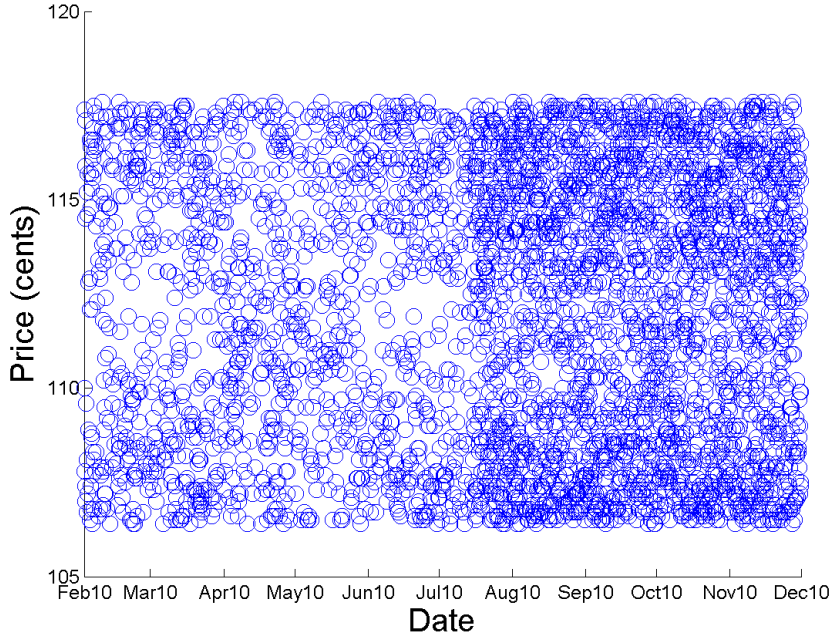

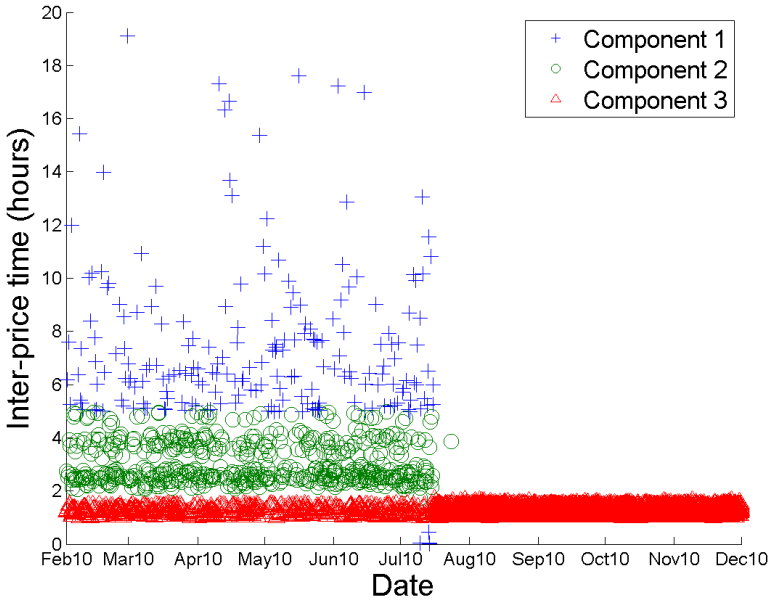

Model Calibration Algorithm

Algorithm 1: Model Calibration Algorithm

Input: $Trace_{inst}, k$
Output: $CalDate, \overrightarrow{RCmps}$

$T_s \leftarrow Trace_{inst}.start.time$;
$T_e \leftarrow Trace_{inst}.end.time$;
$n \leftarrow Sizeof(Trace_{inst})$;
$\overrightarrow{index} \leftarrow (c_1, c_2, \ldots, c_n) \quad c_i \in \{1, \ldots, k\}$;
$\overrightarrow{date} \leftarrow (d_1, d_2, \ldots, d_n) \quad d_i \in \{T_s \ldots T_e\}$;
$q_{a,b} \leftarrow$ probability of component a in month b;
$\overrightarrow{Q} \leftarrow \{q_{a,b} | a \in \{1, \ldots, k\}, b \in \{T_s \ldots T_e\}\}$;
$\overrightarrow{Q_m} \leftarrow \{q_{f,e} | q_{f,e} < q_0, q_{f,e} \in \overrightarrow{Q}\}$;
$\overrightarrow{Cmps} \leftarrow \{g | q_{g,h} \in \overrightarrow{Q_m}\}$;
$\overrightarrow{RCmps} \leftarrow \{1, \ldots, k\} - \overrightarrow{Cmps}$;
$m \leftarrow min\{h | q_{g,h} \in \overrightarrow{Q_m}\}$;
//$Trace_{inst}(m)$ is the trace for month m;
$T_{m_s} \leftarrow Trace_{inst}(m).start.time$;
$T_{m_e} \leftarrow Trace_{inst}(m).end.time$;
$z \leftarrow Sizeof(Trace_{inst}(m))$;
$\overrightarrow{Sindex} \leftarrow (c'_1, c'_2, \ldots, c'_z) \quad c'_i \in \{1, \ldots, k\}$;
$\overrightarrow{Sdate} \leftarrow (d'_1, d'_2, \ldots, d'_z) \quad d'_i \in \{T_{m_s} \ldots T_{m_e}\}$;
$t \leftarrow max\{r_l | \overrightarrow{Sindex}(r_l) == g, l \in \{1, \ldots, z\}\}$;
$CalDate \leftarrow \overrightarrow{Sdate}(t)$;

- Results for k=3

Instances	Calibration Dates	Remaining Components
m1.small	24-July	3
c1.medium	15-July	1
m1.large	15-July	3
m2.xlarge	13-July	1
m1.xlarge	23-July	1
c1.xlarge	23-July	1
m2.2xlarge	23-July	1,2
m2.4xlarge	15-July	3

$$p_j = \frac{p_j}{\sum_{\forall i} p_i} \quad i, j \in \overrightarrow{RCmps}$$

Model Validation

- CloudSim simulator
- LCG workload trace (first 1000 jobs)
  - 400 hours on a m1.small
- EC2 compute unit ≈ 1000 MIPS
- Bid = a large value (e.g. on-demand price)
  - No out-of-bid failure
- MoG (k=3) for Price and Inter-Price time
- Results are collected for 50 runs
- Confidence level of 95%
- Metric : Monetary cost to execute the workload
  - Price History
  - Calibrated Model
  - Non-Calibrated Model

Model Validation (Feb.10-Nov.10)

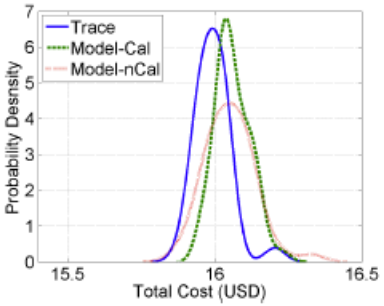

(a) m1.small

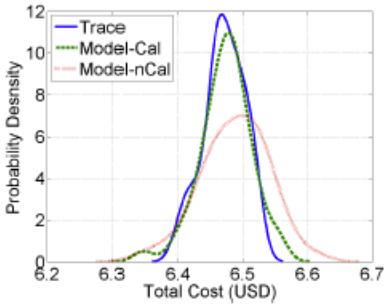

(b) c1.medium

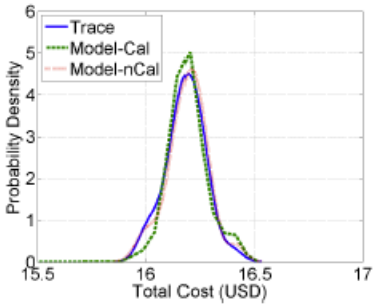

(c) m1.large

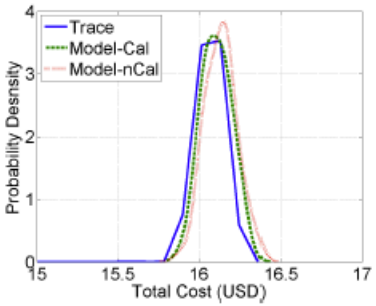

(d) m2.xlarge

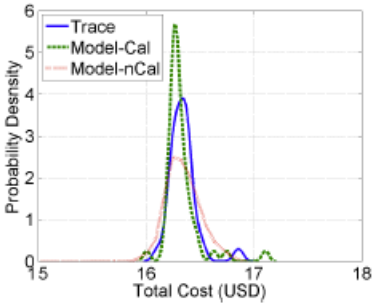

(e) m1.xlarge

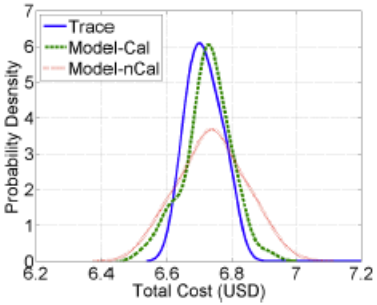

(f) c1.xlarge

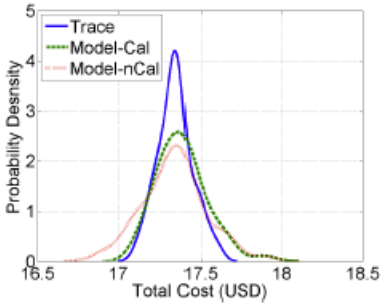

(g) m2.2xlarge

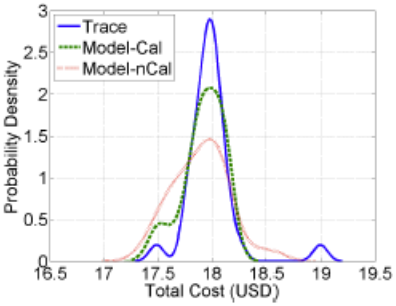

(h) m2.4xlarge

Model Validation (Dec.10-Feb.11)

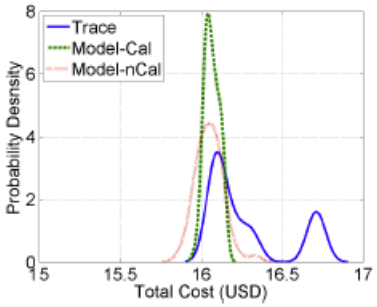

(a) m1.small

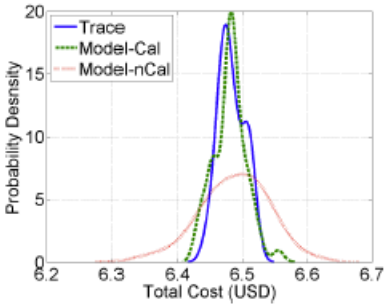

(b) c1.medium

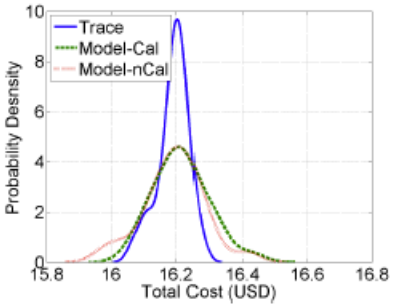

(c) m1.large

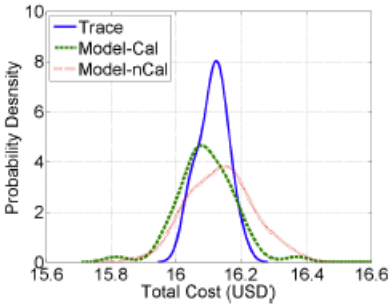

(d) m2.xlarge

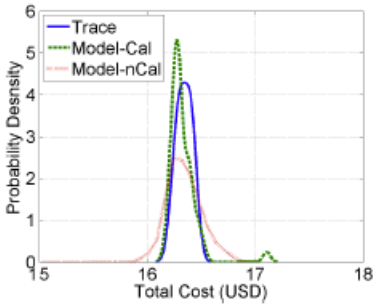

(e) m1.xlarge

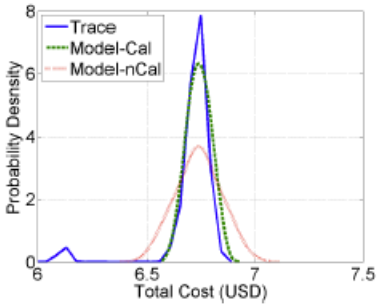

(f) c1.xlarge

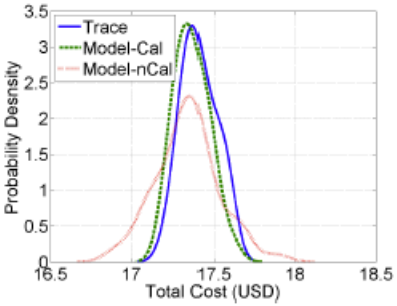

(g) m2.2xlarge

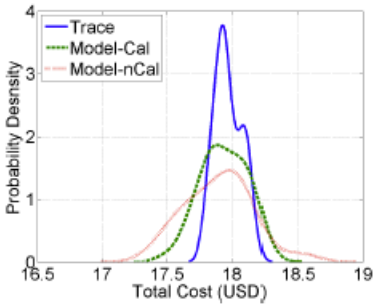

(h) m2.4xlarge

Conclusions

- Discovery Statistical Model for Spot Prices
  - Mixture of Gaussians
- Useful Model for
  - Stochastic Scheduling Algorithms
  - Fault-tolerant Mechanism (check-pointing)
- Future Work
  - Take into account user's bid
  - Cloud Broker to optimize monetary cost

Thank You