

Deconstructing Amazon EC2 Spot Instance Pricing

Orna Agmon Ben-Yehuda Muli Ben-Yehuda
Assaf Schuster Dan Tsafir

Department of Computer Science
Technion — Israel Institute of Technology

Haifux, 26.9.2011

Leonard Kleinrock's 5 golden guidelines to research

- 1 Conduct the 100-year test—will your work be remembered in 100 years?
- 2 Don't fall in love with your model
- 3 Beware of mindless simulation: Ask the obvious questions
- 4 Understand your own results—Use your intuition
- 5 Look for “Gee, that's funny”

Amazon rent virtual machines with prices which vary according to:

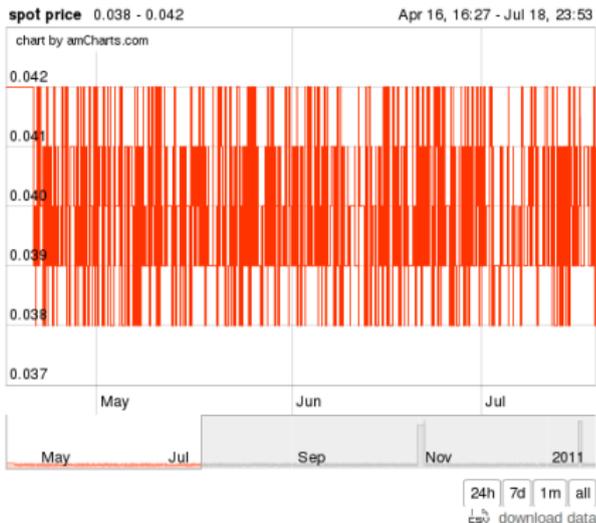
- Instance types
- Regions
- operating systems
- Payment by the hour (a fraction counts as a full hour), unless...
- Commitment level: Reserved, on-demand, spot

cloud exchange

ec2 | ap-southeast-1 | linux | m1.small — \$ 0.038 | 40%

linux	us-w	us-e	eu	ap
m1.s	0.039 41	0.031 36	0.039 41	0.038 40
m1.l	0.162 43	0.125 37	0.162 43	0.165 43
m1.xl	0.315 41	0.241 35	0.304 40	0.325 43
c1.m	0.084 44	0.059 35	0.078 41	0.083 44
c1.xl	0.304 40	0.235 35	0.306 40	0.323 43
m2.xl	0.249 44	0.168 34	0.251 44	0.233 41
m2.2xl	0.537 40	0.440 37	0.563 42	0.545 41
m2.4xl	1.070 40	0.863 36	1.162 43	1.110 41

win	us-w	us-e	eu	ap
m1.s	0.068 52	0.051 43	0.068 57	0.066 55
m1.l	0.280 54	0.200 42	0.273 57	0.257 54
m1.xl	0.542 52	0.385 40	0.543 57	0.542 56
c1.m	0.170 55	0.121 42	0.169 58	0.166 57
c1.xl	0.635 51	0.496 43	0.638 55	0.690 59
m2.xl	0.325 47	0.251 40	0.328 53	0.326 53
m2.2xl	0.707 45	0.553 38	0.717 50	0.753 52
m2.4xl	1.422 45	1.058 37	1.411 49	1.454 50



data: amazon web services — visualization: tim lossen — times shown are UTC

Amazon EC2 Spot instances declaration

“Spot Instances [...] allow customers to bid on unused Amazon EC2 capacity and run those instances for as long as their bid exceeds the current Spot Price. The Spot Price changes periodically *based on supply and demand*, and customers whose bids exceeds it gain access to the available Spot Instances.”

Why sell spot instances?

- Must keep idle machines anyhow, for elasticity reserved and on-demand instances.
- Selling easily-evacuated instances on idle machines can cover the expenses of running them idle
- Idle machines are sunk costs. To start earning, the provider only needs to charge for them the difference between electricity costs of idle and active machines.

What is the EC2 spot instance pricing mechanism?

- Uniform Price
- Sealed Bid
- Market Based
 - What does this mean?

What is the EC2 spot instance pricing mechanism?

- Uniform Price
- Sealed Bid
- Market Based
 - What does this mean?

What is the EC2 spot instance pricing mechanism?

- Uniform Price
- Sealed Bid
- Market Based
 - What does this mean?

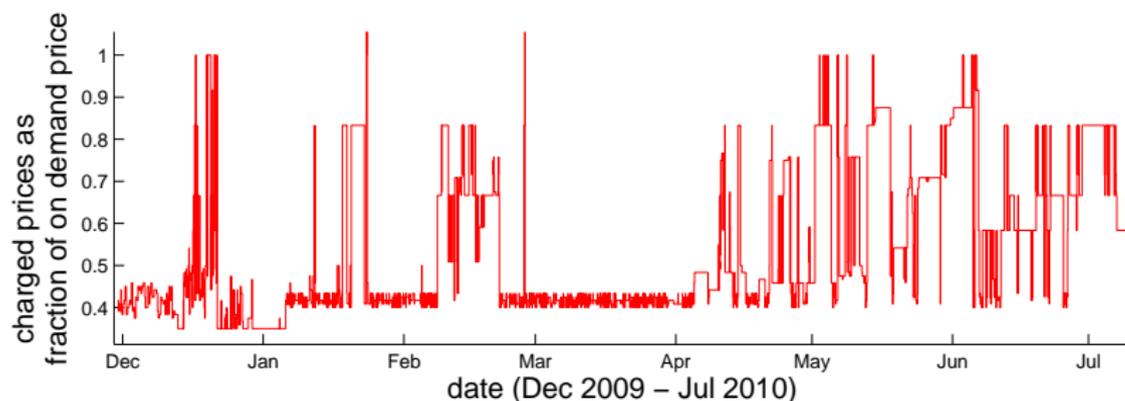
What is the EC2 spot instance pricing mechanism?

- Uniform Price
- Sealed Bid
- Market Based
 - What does this mean?

An example for what Amazon might be doing under this definition

- The provider is the auctioneer and the seller
- Clients bid secretly
- The provider sorts the bids (descending order)
- The provider grants only first N bids.
 - N is limited by the actual supply
 - N may be smaller (retroactive limitation of supply): minimal price, (hidden) reserve price, revenue maximization.
- All bidders who got in pay the highest of the prices which did not (bid $N + 1$) or the minimal price.

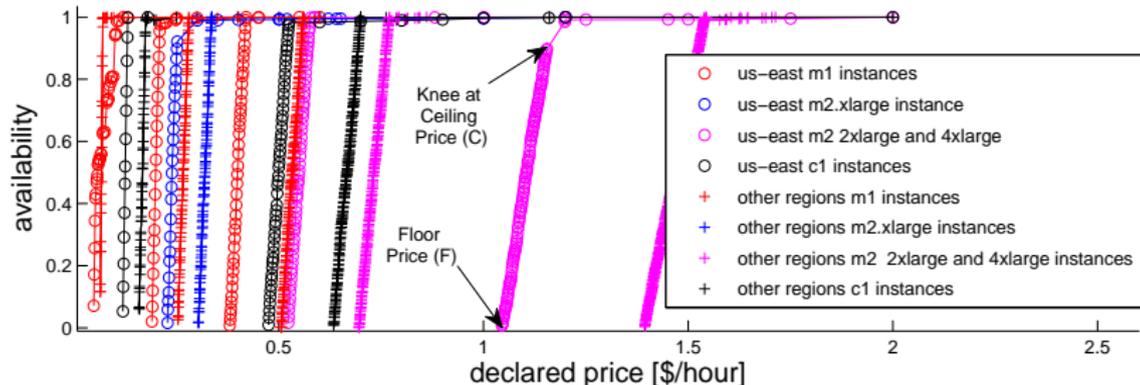
Amazon encourages clients to look at price histories, and bid accordingly.



Availability of bid price

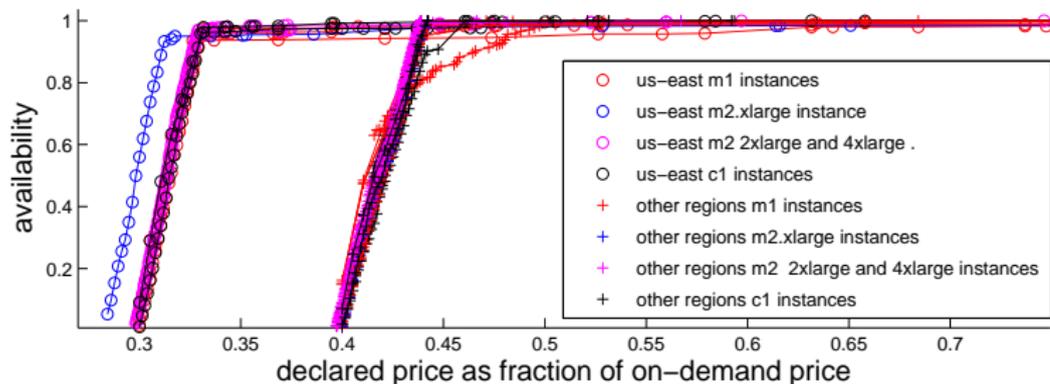
- The total time in which the spot price was above the bid price, divided by the total time.
- The probability that at a uniformly random time, the bid would immediately get in.

Windows instances availability as function of price



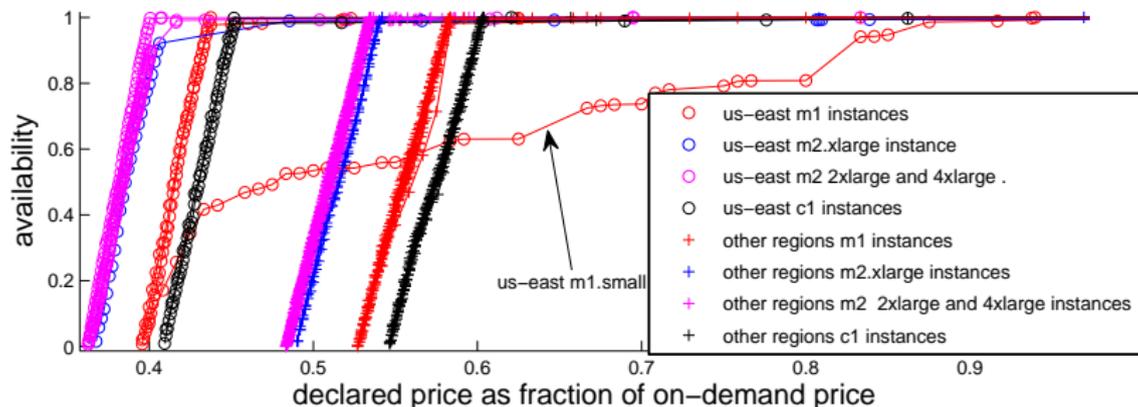
Regular shapes at different prices: straight lines and high knees. Looks similar for Linux.

Linux instances availability as function of normalized price



Two groups of regions. The forest disappears.

Windows instances availability as function of normalized price



A repeating pattern within the two region groups.

Real client bids would have to account for:

- 1 normalized prices turning out **identical for various Linux types** but different for Windows types;
- 2 a **rigid linear connection** between availability and price that turns out identical for different types and regions;
- 3 a **singular region** having a normalized price range different than all the rest (which turn out to have identical ranges);
- 4 normalized prices for Windows instances which differ from one another by identical amounts in each region class, creating the **same pattern for both region classes**.

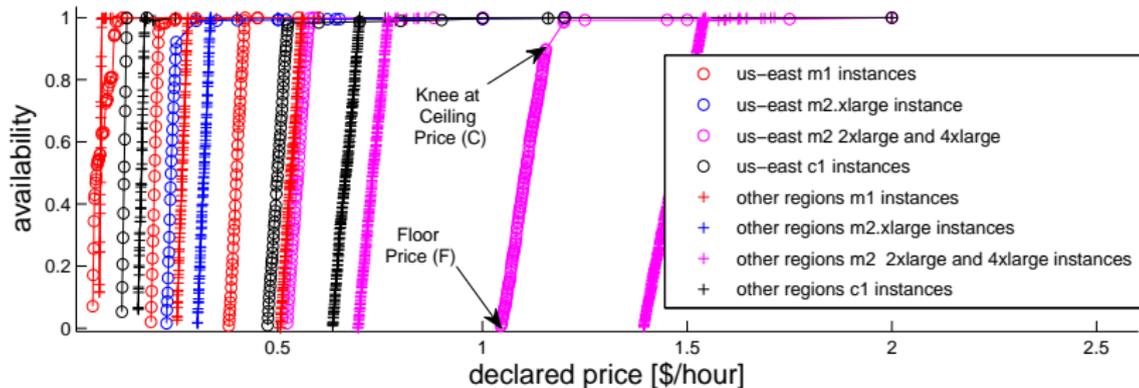
Our hypothesis

- Amazon uses a dynamic algorithm to set a reserve price for the auction, independently of client bids.
- Most of the time, the auction's result is identical to the reserve price.
- Because of that, usually the prices Amazon announces are not market-driven.
- Supported by both Occam's razor and simulations.

Our hypothesis—requirements from the dynamic algorithm

- Input: a pricing band—floor price F , ceiling price C (per instance type and region, as fraction of on-demand price)
- Output:
 - dynamic changes of the reserve price such that the availability graph is linear in the $[F,C]$ range.
 - the reserve price never drops below the floor
 - the reserve price never rises above the ceiling
 - the spot price may rise above the ceiling due to market considerations

Ceiling and floor of the pricing band

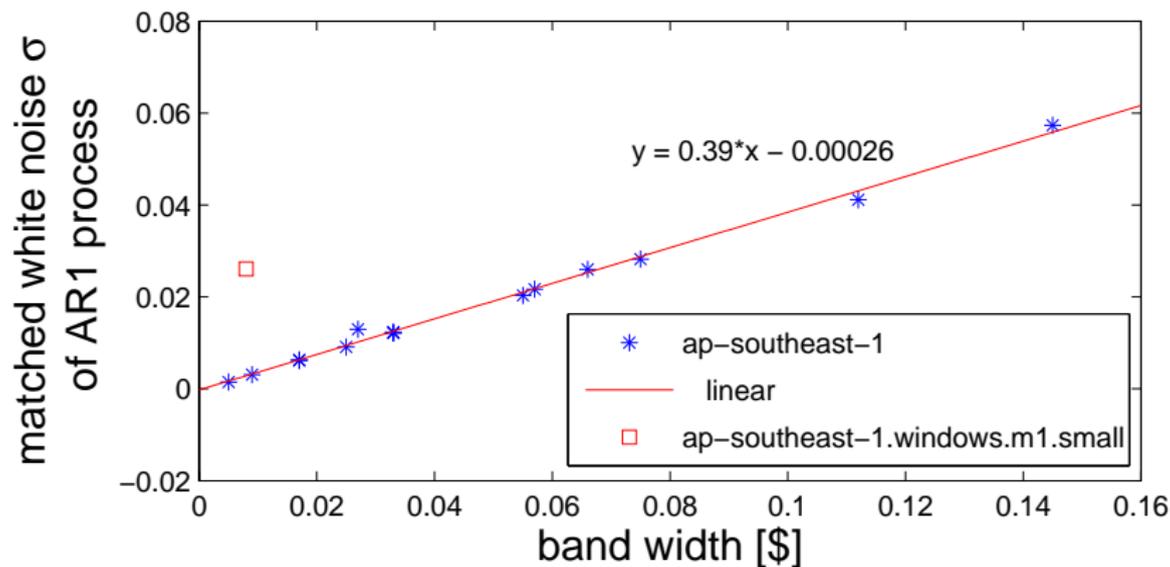


Matching an auto-regressive process $AR(1)$ for ap-southeast.windows types

$$\Delta_j = -a_1 \Delta_{j-1} + \epsilon(\sigma)$$

- Δ_j is the difference of two consequent prices
- $a_1 = 0.7$
- $\epsilon(\sigma)$ - white noise with a standard deviation $\sigma = 0.39(C - F)$
- m1.small matched $a_1 = 0.5, \sigma = 0.5(C - F)$.

Variance of the matched AR(1) process



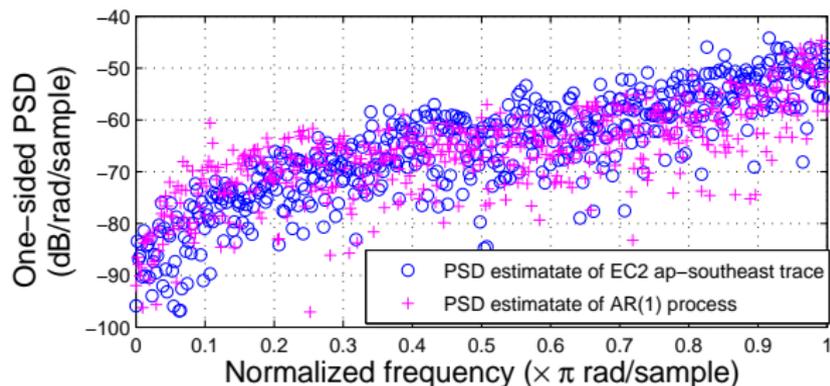
The close fit supports our hypothesis.

Constructing the reserve price algorithm

- Start with $P_0 = F$ and a price change of $\Delta_0 = 0.1(F - C)$ (this works, but not everything does. You can choose other conditions).
- $P_i = P_{i-1} + \Delta_i$
- $\Delta_i = -0.7 \cdot \Delta_{i-1} + \epsilon(0.39 \cdot (C - F))$
- Truncate the process to the range $[F, C]$ by regenerating the white noise component while P_i is outside the $[F, C]$ range or identical to P_{i-1} .
- Round all prices to 0.1 cent.

Is the constructed algorithm consistent with reality?

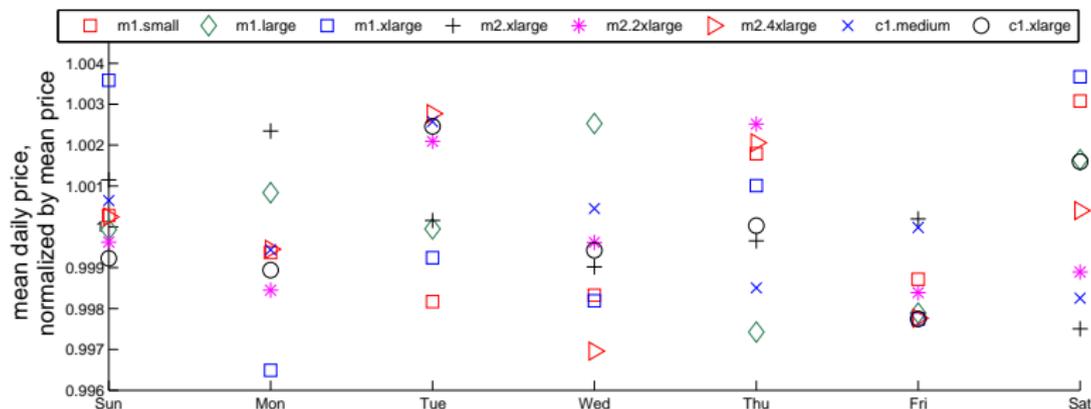
Periodogram: “The modulus-squared of the discrete Fourier transform of the time series (with the appropriate normalization)” —Vaughan and Uttley.



The close fit supports our hypothesis.

Is the AR(1) process natural or artificial?

The weekly average of ap-southeast.windows types:



A natural process is expected to have a significant weekly cycle.

Is the AR(1) process partly natural or artificial?

- Partly natural : some real bids within band, some reserve prices.
- A partly natural process is expected to have a mean price above mid-range.
 - The mean price is 98%-100% of the mid-range.
 - Many clients already noted that bidding inside the band is not cost effective.

Is the AR(1) process partly natural or artificial?

- Partly natural : some real bids within band, some reserve prices.
- A partly natural process is expected to have a mean price above mid-range.
 - The mean price is 98%-100% of the mid-range.
 - Many clients already noted that bidding inside the band is not cost effective.

Is the AR(1) process partly natural or artificial?

- Partly natural : some real bids within band, some reserve prices.
- A partly natural process is expected to have a mean price above mid-range.
 - The mean price is 98%-100% of the mid-range.
 - Many clients already noted that bidding inside the band is not cost effective.

Is the AR(1) process partly natural or artificial?

- Partly natural : some real bids within band, some reserve prices.
- A partly natural process is expected to have a mean price above mid-range.
 - The mean price is 98%-100% of the mid-range.
 - Many clients already noted that bidding inside the band is not cost effective.

Are traces as a whole natural or artificial?

- 98% of the time, prices are within the band.

If our hypothesis is correct, then:

- 98% of the time spot prices carry little information about real client bids!
- Researchers *cannot* learn from spot prices about client valuations for products, nor about supply and demand
- The spot price is *not necessarily* a market clearing price.

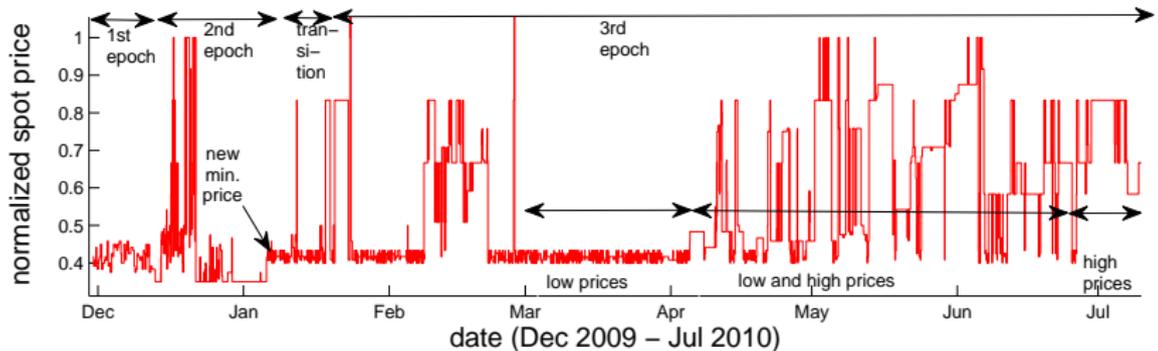
If our hypothesis is correct, then:

- 98% of the time spot prices carry little information about real client bids!
- Researchers *cannot* learn from spot prices about client valuations for products, nor about supply and demand
- The spot price is *not necessarily* a market clearing price.

If our hypothesis is correct, then:

- 98% of the time spot prices carry little information about real client bids!
- Researchers *cannot* learn from spot prices about client valuations for products, nor about supply and demand
- The spot price is *not necessarily* a market clearing price.

Pricing epochs



Qualitative changes define epochs

- 11/2009—internal usage or beta testers (epoch 1)
- 12/2009—announcement (epoch 2)
- 01/2010—appearance of pricing bands (epoch 3, starting in a transfer period)
- 07/2010—change of timing algorithm in us-east
- 02/2011—another change of the timing algorithm in us-east

Qualitative changes define epochs

- 11/2009—internal usage or beta testers (epoch 1)
- 12/2009—announcement (epoch 2)
- 01/2010—appearance of pricing bands (epoch 3, starting in a transfer period)
- 07/2010—change of timing algorithm in us-east
- 02/2011—another change of the timing algorithm in us-east

Qualitative changes define epochs

- 11/2009—internal usage or beta testers (epoch 1)
- 12/2009—announcement (epoch 2)
- 01/2010—appearance of pricing bands (epoch 3, starting in a transfer period)
- 07/2010—change of timing algorithm in us-east
- 02/2011—another change of the timing algorithm in us-east

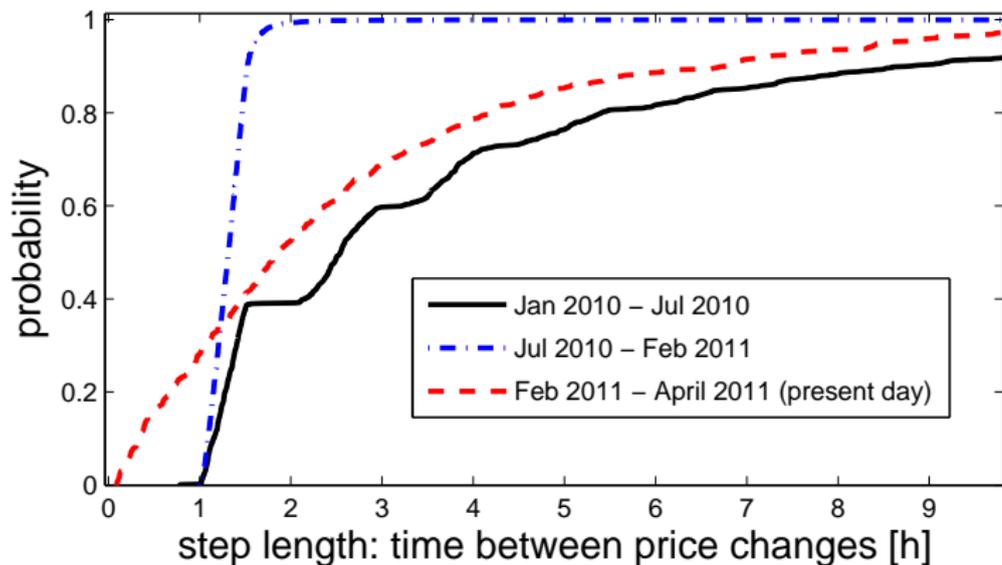
Qualitative changes define epochs

- 11/2009—internal usage or beta testers (epoch 1)
- 12/2009—announcement (epoch 2)
- 01/2010—appearance of pricing bands (epoch 3, starting in a transfer period)
- 07/2010—change of timing algorithm in us-east
- 02/2011—another change of the timing algorithm in us-east

Qualitative changes define epochs

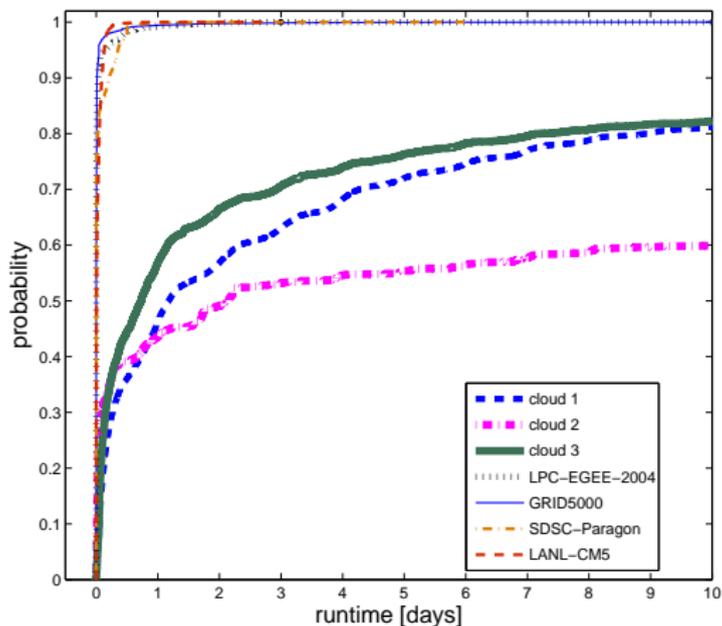
- 11/2009—internal usage or beta testers (epoch 1)
- 12/2009—announcement (epoch 2)
- 01/2010—appearance of pricing bands (epoch 3, starting in a transfer period)
- 07/2010—change of timing algorithm in us-east
- 02/2011—another change of the timing algorithm in us-east

Price Changing Timing (us-east)

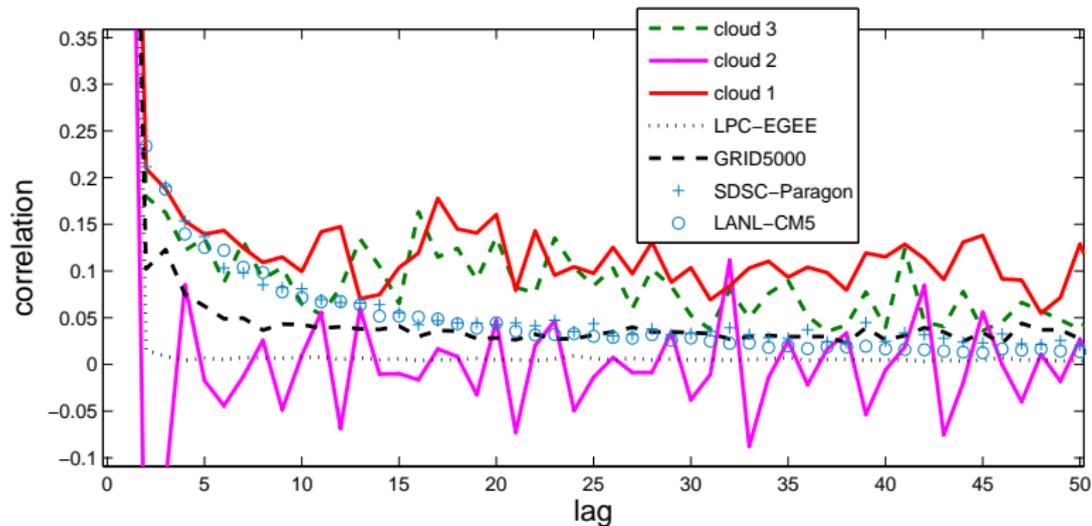


- Workload traces of large systems. Truncated to tasks longer than 10 minutes, shorter than 24 hours.
- Grid: LPC-EGEE, a cluster of a large grid.
- Clouds: 3 internal clouds of a Fortune-500 company.
- What are the characteristics of the workloads? Are they good cloud workloads?

Instance Duration CDF



Auto-correlation of Instance Inter-Arrival time



We know almost nothing about client bids. We test three models, to show that the qualitative results are insensitive to the model.

- Pareto distribution (a widely applicable economic distribution) with a minimal value of 0.4, and a Pareto index of 2, a reasonable value for income distribution.
- $\mathcal{N}(0.7, 0.3^2)$, truncated at 0.4.
- A linear mapping from runtimes to $(0.4, 1]$, which reflects client aversion to having long-running instances terminated.

Simulator Event-Driven Loop

Events are:

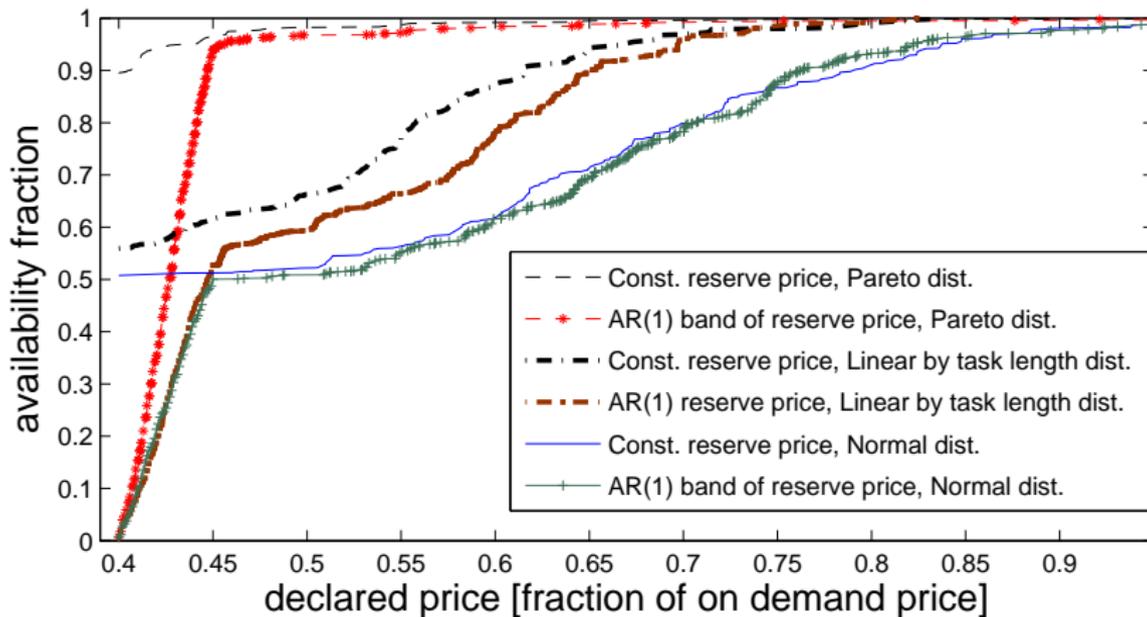
- Instance submission
- Instance termination
- price changes due to:
 - A scheduled change
 - A waiting instance with a bid higher than the spot price.

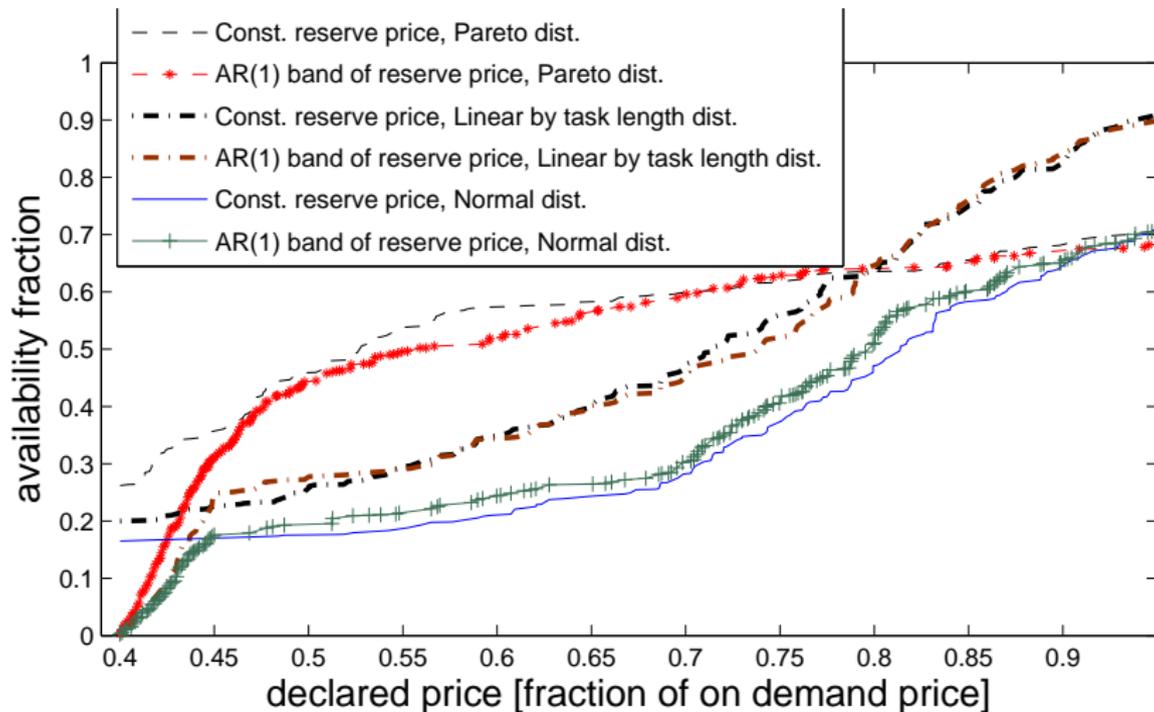
System:

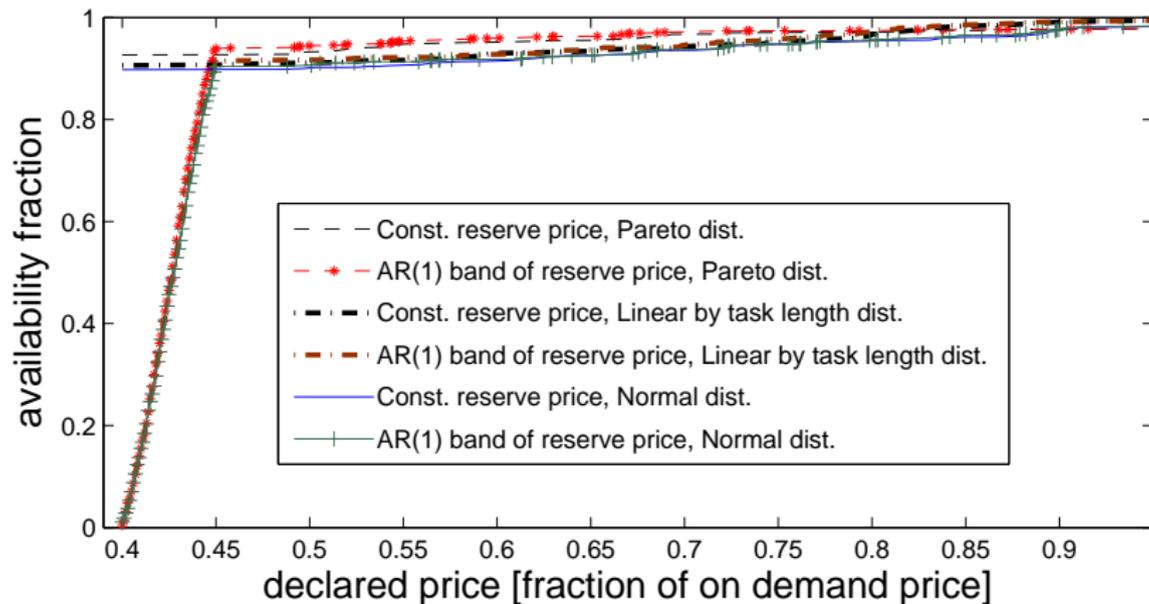
according to the respective maximal values in each trace.

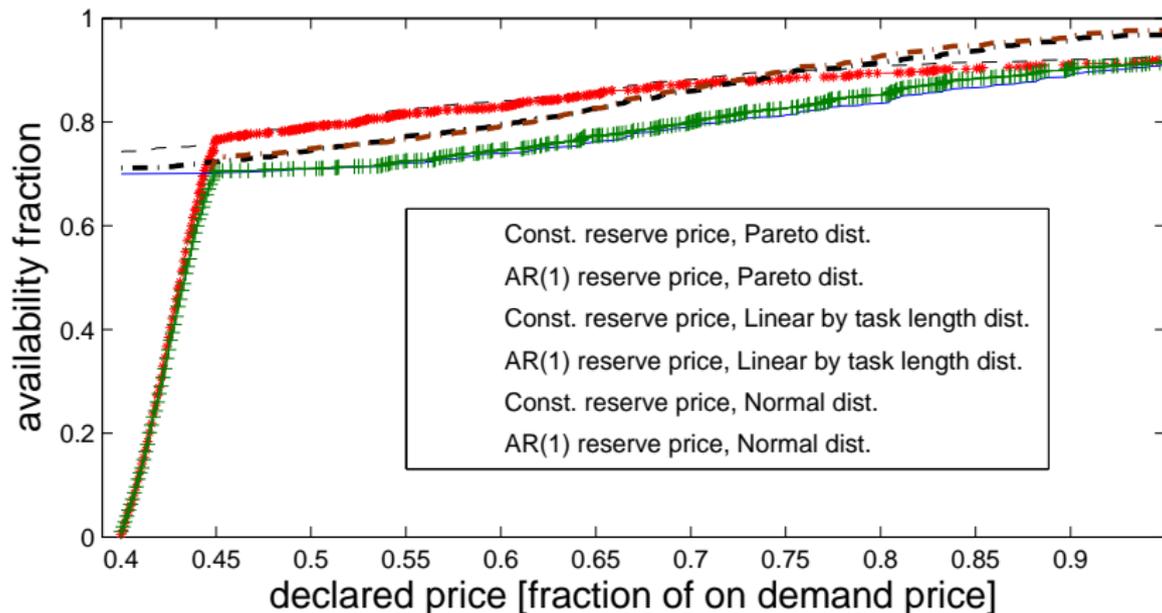
Ending:

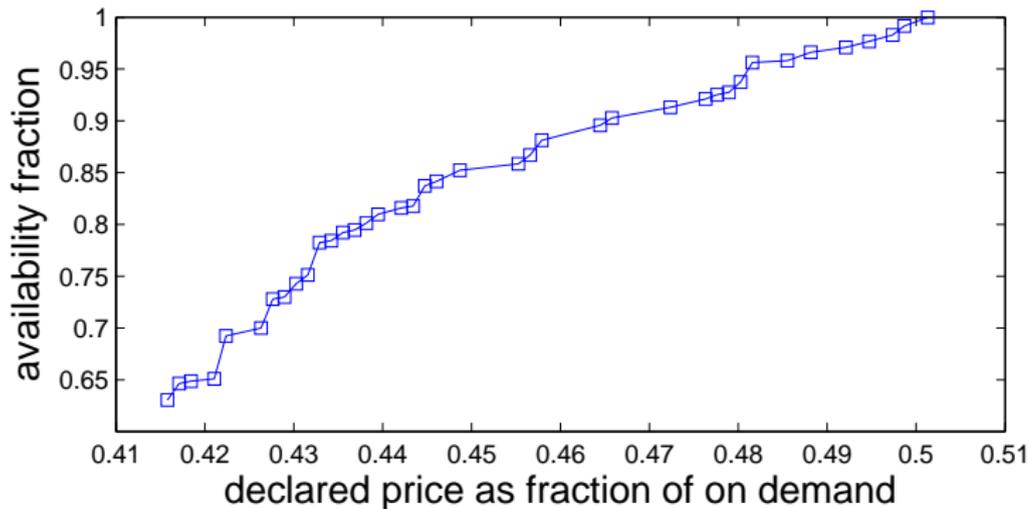
when the last input-trace job had been submitted.



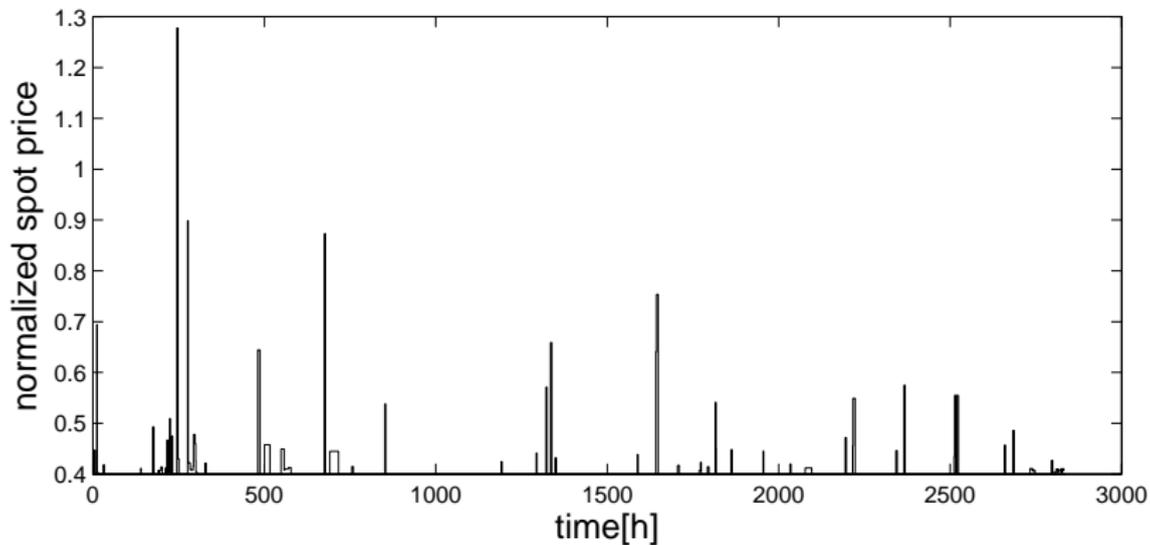




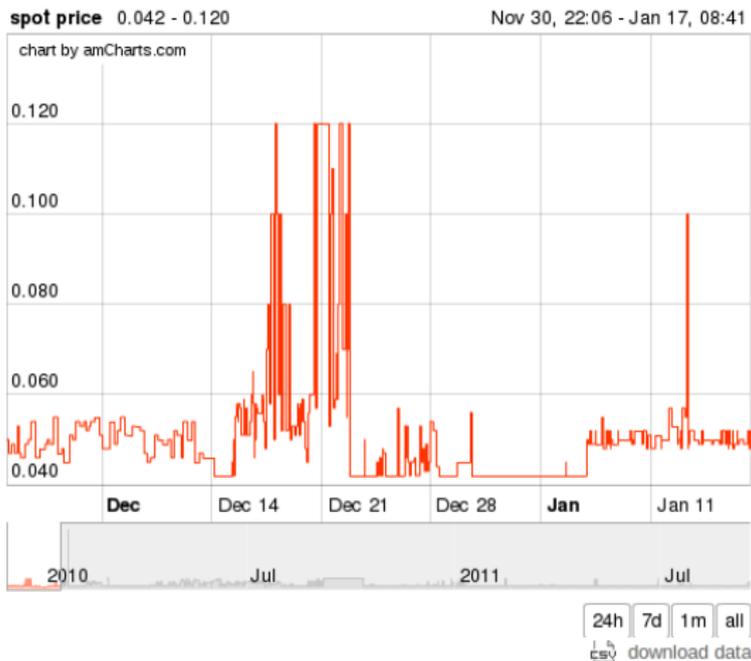




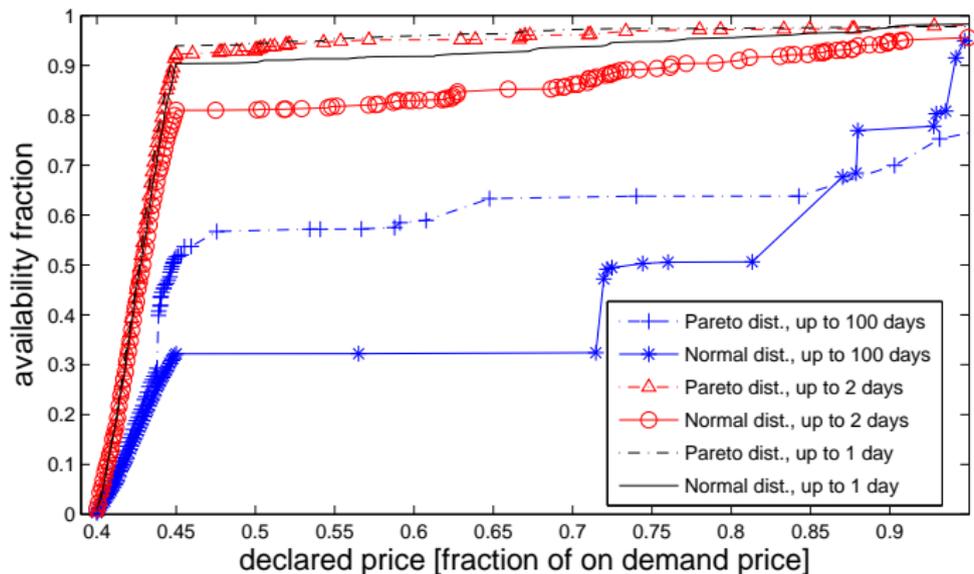
LPC-EGEE History



Reminder - Epoch 2 History



Impact of Truncation of Long Instances



Conclusions From Simulations

- Simulation results show a knee and a linear segment, consistent with traces.
- Existence of simulation knee and linear segment are insensitive to client bidding.
- Existence of simulation knee and linear segment are insensitive to instance length truncation. If our hypothesis is correct, then the EC2 workload is consistent with being characterized by relatively short instances.

Why minimal/reserve price?

- Prevent cannibalization of main offers
- Prevent selling at a loss
- Still allow selling of idle capacity

To publish or not to publish? Minimal or reserve?

- A reserve price can be dynamically changed, no obligation to inform clients.

Dynamic is better than constant reserve price

- Maintains an impression of constant change, thus preventing clients from becoming complacent. Forces clients to
 - Bid higher (unlike using a constant floor)
 - or tolerate sudden unavailability (throws the burden of dealing with elasticity on low-grade SLA clients).
- Clears queues of low bids within the band (unlike using a constant ceiling).

Secret dynamic reserve price is better than known

- Masks times of low demand, price inactivity (impression of false activity)
- Possibly drives up provider stock
- A large band could mask high demand and low supply (illusion of infinite elasticity)
 - (but when the band is small, like in EC2, it is an indication of low demand)

Analyzing Spot Price Traces

- Wee 2011
- Javadi and Buyya 2011

Using Spot Price Traces for Client Strategy Evaluation

- Andrzejak, Kondo and Yi $\times 3$ - minimize costs while meeting an SLA, schedule checkpoints and migrations.
- Mattess, Vecchiola and Buyya 2010—managing peak loads in scientific workloads.
- Chohan et al. 2010—characterizes instance type performance by CDF of time the instance holds at a price.

Using Spot Price Traces to Learn About the Market

- Zhang et al. 2011—assumed EC2 traces reflect real bids.
- Chen et al. 2011—assumed EC2 traces represent market clearing prices.

Conclusions

- It is likely that Amazon sets spot prices using a AR(1) (hidden) reserve price. While prices above the band are probably market driven (2% of the time), the rest of the 98% of the time the spot instance is probably just the reserve price.
- 98% of the time, EC2 traces do not necessarily represent a clearing price and do not necessarily represent real bids.
- EC2 price traces provide more information about amazon than about its clients.
- Many features (minimal price, band width, change timing) are artificial, and may change at will. This divides the traces to epochs.
- This understanding is important to
 - Clients
 - Providers
 - Researchers

Questions?

Contact us at:

{[ladypine](#),[muli](#),[assaf](#),[dan](#) } at cs.technion.ac.il

Thank You!