Distributed Systems Meet Economics: Pricing in the Cloud

Hongyi Wang[†] Qingfeng Jing[‡] Rishan Chen[°] Bingsheng He[†] Zhengping Qian[†] Lidong Zhou[†]
[†]*Microsoft Research Asia* [‡]*Shanghai Jiao Tong University* [°]*Peking University*

Presenter: Rishan Chen
Peking University and Microsoft Research, Asia
June 2010, Boston, MA

Cloud is a distributed system

- System metrics
  - Throughput
  - Latency / response time
  - Failure rate
  - Power consumption, etc.
- As a pay-as-you-go service
  - Two parties connected by the pricing scheme
  - It's all about the money!

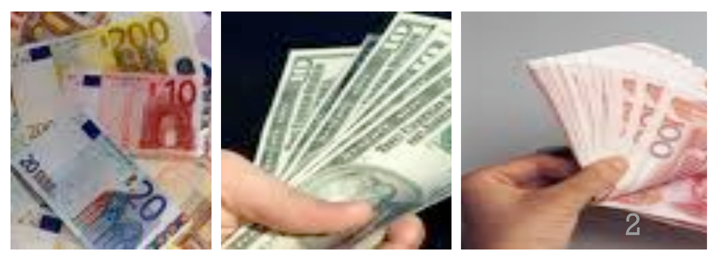

Pricing in the Cloud

- It significantly changes the landscape of system design: **Cost** as an explicit and measurable system metric
  - How both parties optimize their logic
  - Is the pricing fair
  - How does the pricing interplay with the evolving system dynamics
  - How to measure the cost of failures, etc.

Methodology overview

- Approximate a typical workload in current cloud computing
  - `Postmark` (I/O-intensive)
  - `PARSEC`: `Dedup`, `BlackScholes` (CPU-intensive)
  - `Hadoop` (large-scale data processing)
- Complementary approaches for evaluations
  - A black-box approach with Amazon EC2
  - Built a cloud-computing test bed, *Spring*, to perform fully controlled experiments

Preliminary results

- Pricing may give different indices for users and providers for system optimizations (e.g., consolidation)
- System performance variations may lead to pricing fairness issues
- System evolution (e.g., adoption of new hardware like SSD) may affect pricing scheme
- Failures need to be better dealt with regarding to the cost

Highlights of our study /1

- Pricing (profit) versus throughput

Number of concurrent VMs	One VM	Two VMs	Four VMs
Average cost per task ($)	0.004	0.004	0.012
Profit ($)	-0.009	0.002	0.028
Throughput (tasks/h)	28.3	56.4	33.9

Run Postmark continuously and report the number for four tasks; we compare the consolidation of x VMs on a single physical machine.

Highlights of our study /2

- Optimizing for cost versus optimizing traditional system metrics

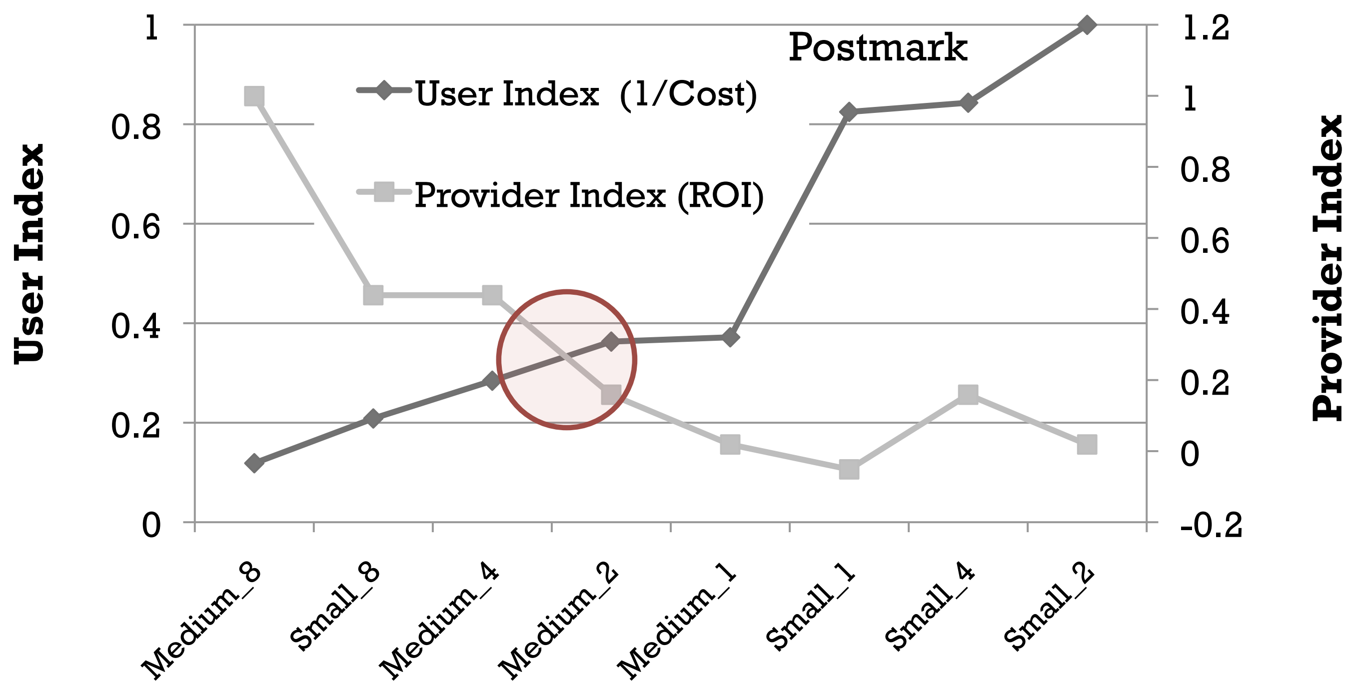

Highlights of our study /3

- Pricing fairness: performance variation

	Postmark	Dedup	BlackScholes
cv	9.1%	11.0%	3.9%
$maxDiff$	40.1%	38.8%	12.6%

Table 8: Variation of different runs on EC2

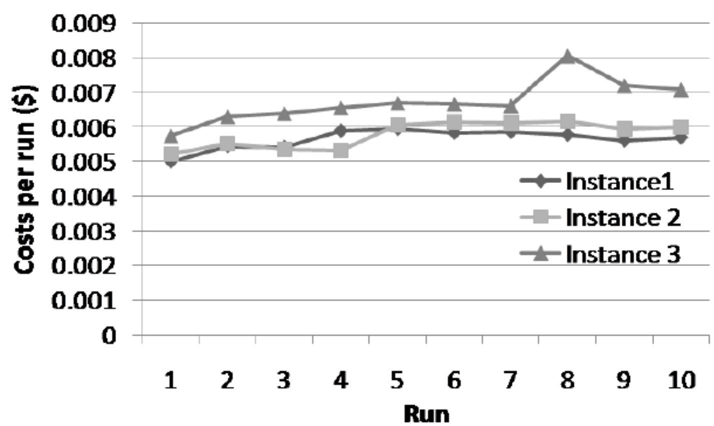

Figure 2: Variations among three instances (Postmark)

Open questions

- What are good properties for a pricing scheme?
- How do users and providers adapt the system design to evolving and even hybrid pricing schemes?
- How is the pricing scheme adapted to the evolving system dynamics and (new) technologies?
- How to deal with failures' cost regarding to the pricing?

Related work

- Other pricing schemes
  - Bilateral
  - Amazon EC2 Spot Instances: Enable you to bid for unused Amazon EC2 capacity
    - Navraj Chohan, et al., *See Spot Run: Using Spot Instances for MapReduce Workflows*, June 2010
  - Microsoft SQL Azure: Make pricing more scalable and more predictable
- Distributed computing w/ Economics
  - Jim Gray, *Distributed Computing Economics*, March 2003
  - Ang Li, et al., *CloudCmp: Shopping for a Cloud Made Easy*, June 2010

Summary

- Pricing is an important bridge between users and providers
- It significantly changes the dynamics in system design
- The interplay between economics and system design can be a fruitful research direction