



Aalto University

Exploiting Hardware Heterogeneity within the Same Instance Type of Amazon EC2

Zhonghong Ou, Hao Zhuang, Jukka K. Nurminen, Antti Ylä-Jääski

Aalto University, Finland

Pan Hui

Deutsch Telekom Laboratories, Berlin, Germany

Outline

- Motivation
- Experimental configuration
- Experimental results
 - Micro-benchmark
 - Application benchmark
- Cost analysis
- Conclusion & future work

Motivation

- Cloud computing attracts attention because of
 - Pay-as-you-go
 - Theoretically unlimited resource
 - Reduced Capital Expenditure (CAPEX) and Operating Expense (OPEX)
 - And more...
- Amazon EC2
 - Introduced in 2006
 - Provisioning various categories of instances, diversified types of instances within the same category
- Hardware heterogeneity likely from
 - Hardware upgrade and replacement
- Research problems
 - Homogeneous vs. heterogeneous?
 - Performance variation?

Experimental configurations

- CPUID
 - Non-trapping instruction
- Confirmed with
 - `cat /proc/cpuinfo`
- Collected info from Amazon EC2 for two periods of time
 - Apr. – Jul. 2011
 - Jan. – Mar. 2012
 - 200 instances collected for each instance type
- Micro-benchmark
 - CPU performance: UnixBench
 - Memory performance: Redis
 - Disk performance: Dbench
- Application benchmark
 - Httperf

Hardware information

Table 1: Hardware configuration

Instance type	CPU model	%(2011)	%(2012)
<i>m1.small</i>	E5507	45%	12%
	E5430	34%	38%
	E5645	3%	30%
	2218HE	18%	20%
<i>m1.large</i>	E5507	58%	40%
	E5430	29%	17%
	E5645	5%	42%
	2218HE	4%	1%
	270	4%	-
<i>m1.xlarge</i>	E5507	31%	6%
	E5430	27%	46%
	E5645	40%	48%
	270	2%	-


Released:
 E5507: Q1'10
 E5430: Q4'07
 E5645: Q1'10
 2218HE: Q3'06

Newer processor models replace older ones progressively
 Hardware info varies significantly among different availability zones

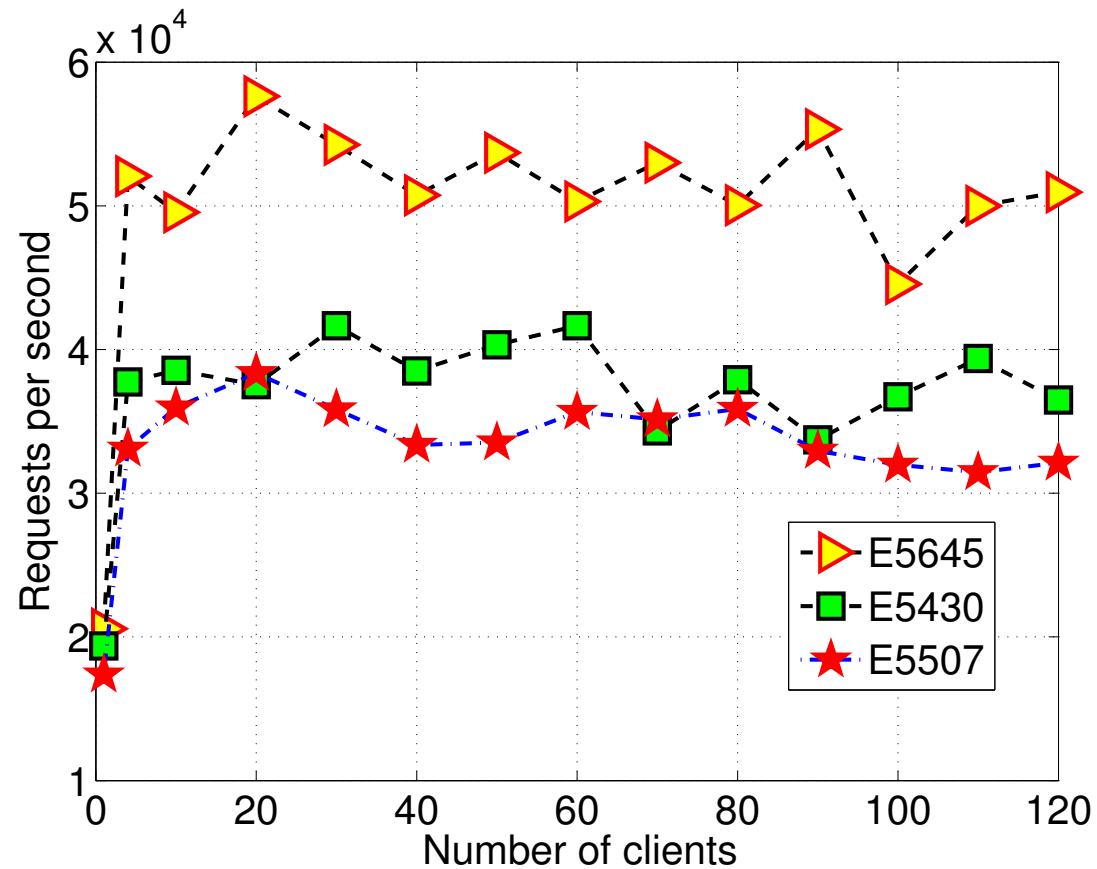
Outline

- Motivation
- Experimental configuration
- **Experimental results**
 - Micro-benchmark
 - CPU performance: UnixBench
 - Memory performance: Redis
 - Disk performance: Dbench
 - Application benchmark
 - Httpperf
- **Cost analysis**
- **Conclusion & future work**

CPU performance: UnixBench



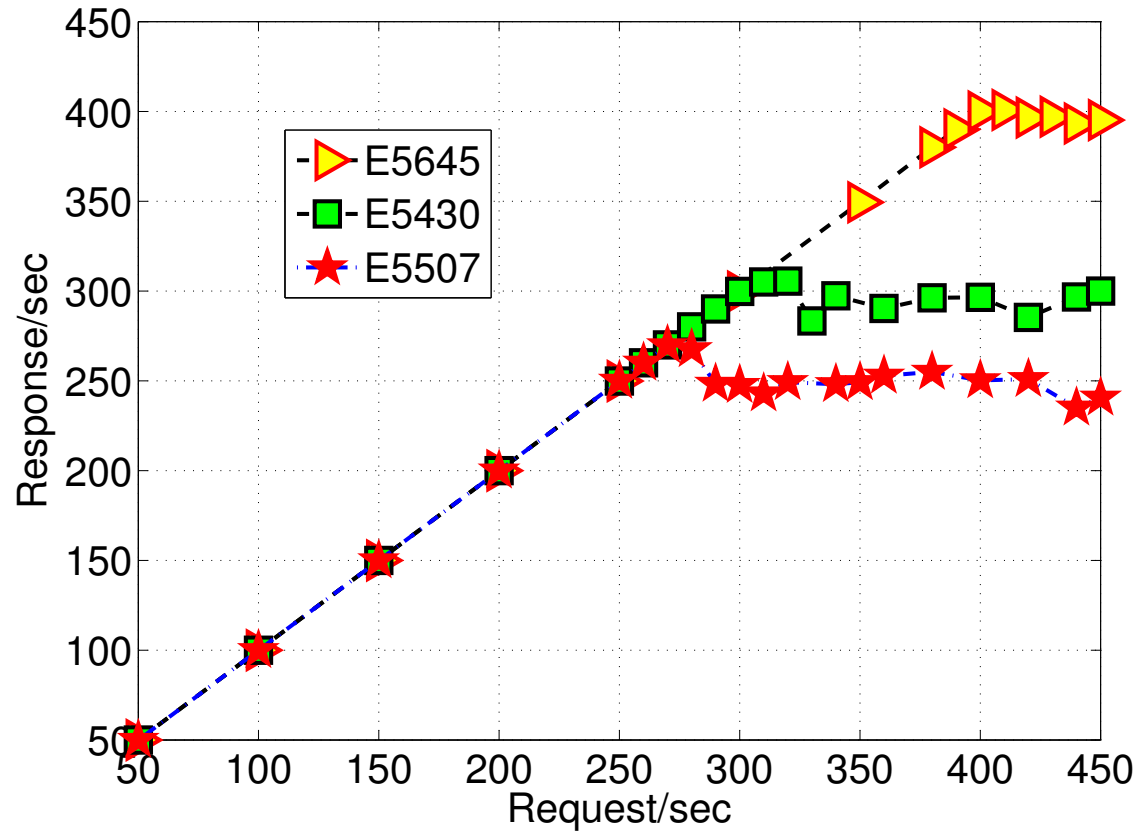
Memory performance: Redis



Disk performance: Dbench

- Shows similar results as UnixBench and Redis
- E5645 is approximately 1.25 times better than E5507
- E5430 is comparable to E5507

Application performance: Httpperf




E5645 is 1.6 times as efficient as E5507

E5430 is 1.2 times as E5507

Cost analysis

- Saving money by seeking for better performing instances, simply using “trial-and-failure” method
 - Applying for instances randomly;
 - Checking if performing well;
 - If not, drop and apply for new ones.

Cost analysis (cont.d)



P : probability of the better-performing instance in the overall instances

Conclusion

- Amazon EC2 uses diversified hardware to host the same type of instance.
- The hardware diversity results in performance variation.
- In general, the variation between the fast instances and slow instances can reach 40%. In some applications, the variation can even approach up to 60%.
- By selecting fast instances within the same instance type, Amazon EC2 users can acquire up to 30% of cost saving, if the fast instances have a relatively low probability.

Future work

- Investigating the scheduling mechanism and analyze its impact on the performance of Amazon EC2 instances.
- Building homogeneous platform from heterogeneous hardware environment, taking into consideration:
 - Processor
 - Memory
 - Disk
 - Networking