A Shapley Value Approach for Cost Allocation in the Cloud

Qizhen Zhang, Haoran Wang, *Yang Chen*, Tao Qin, Ying Yan, *Thomas Moscibroda*
{*yachen*, *moscitho*}@microsoft.com
Microsoft Research

SoCC 2015

Bin Packing in the Cloud

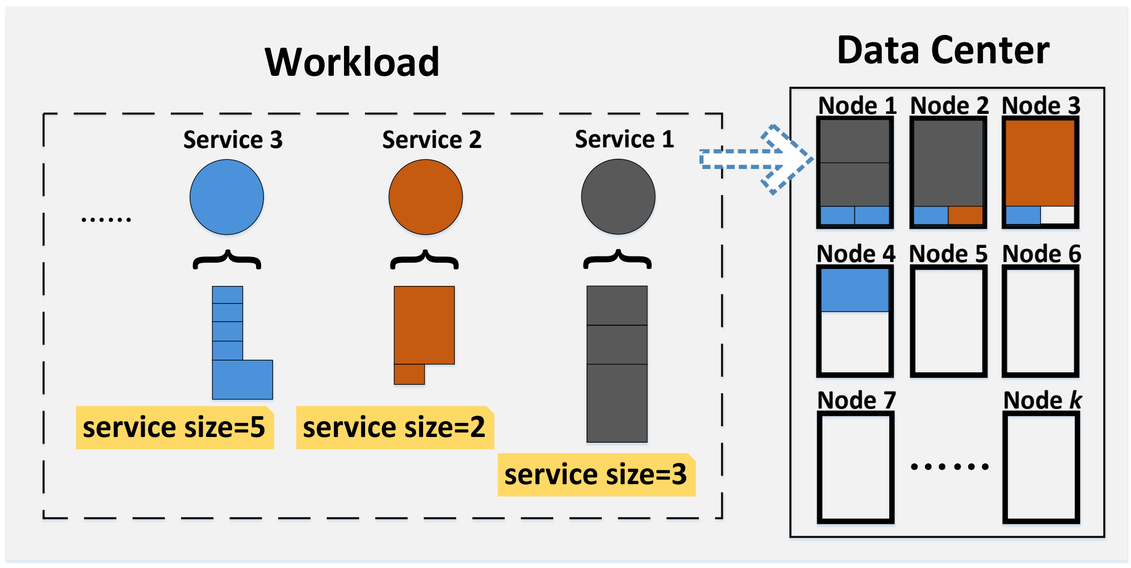

Resource allocation in the cloud can be seen as bin packing problem. A *workload* is a sequence of *services*, each of which is constituted of several *instances*, i.e. virtual machines.

Workload Cost in the Cloud

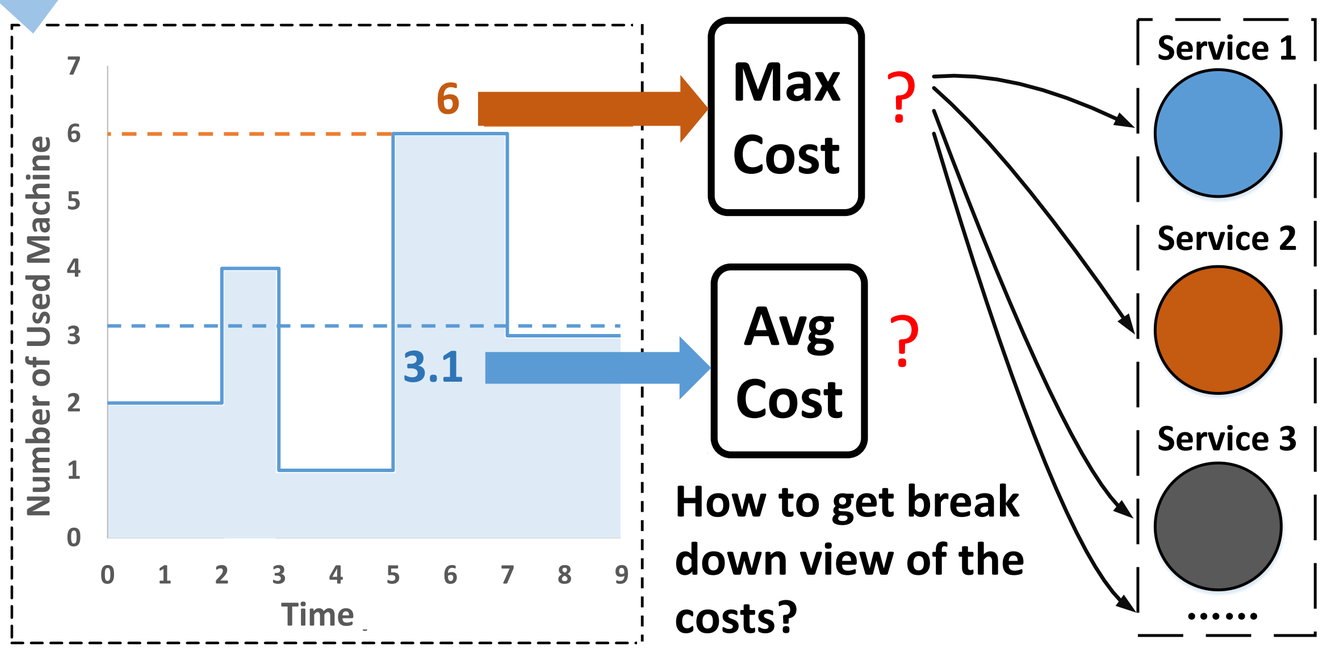

What is the cost of a service? → Cost allocation in the cloud: Allocate total cost to all services in a fair way.

Pricing in the Cloud

Microsoft Azure

INSTANCE	CORES	RAM	PRICE
A1	1	1.75 GB	$0.09/hr (~$67/mo)
A2	2	3.5 GB	$0.18/hr (~$134/mo)
A3	4	7 GB	$0.36/hr (~$268/mo)
A4	8	14 GB	$0.72/hr (~$536/mo)

amazon web services™

	vCPU	Memory (GiB)	Linux/UNIX Usage
m3.medium	1	3.75	$0.067 per Hour
m3.large	2	7.5	$0.133 per Hour
m3.xlarge	4	15	$0.266 per Hour
m3.2xlarge	8	30	$0.532 per Hour

Current cloud pricing is linear in resource size. Is the cost incurred by a service really linear to its size?

Cost Allocation based on Shapley Value

Marginal Cost: $\phi_i(v) = v(\mathbb{N}) - v(\mathbb{N}\backslash\{i\})$

Shapley value takes all possible coalitions into account: $O(2^{|\mathbb{N}|})$

Groups of similar services form coalitions

Shapley Value is a classic concept, widely applied to cost allocation problems in coalition games. The goal is to fairly allocate surplus or cost to each player in the game.

$$\phi_i(v) = \sum_{\mathbb{S}\subseteq\mathbb{N}\backslash\{i\}} \frac{|\mathbb{S}|!\,(|\mathbb{N}| - |\mathbb{S}| - 1)!}{|\mathbb{N}|!}\left(v(\mathbb{S}\cup\{i\}) - v(\mathbb{S})\right)$$

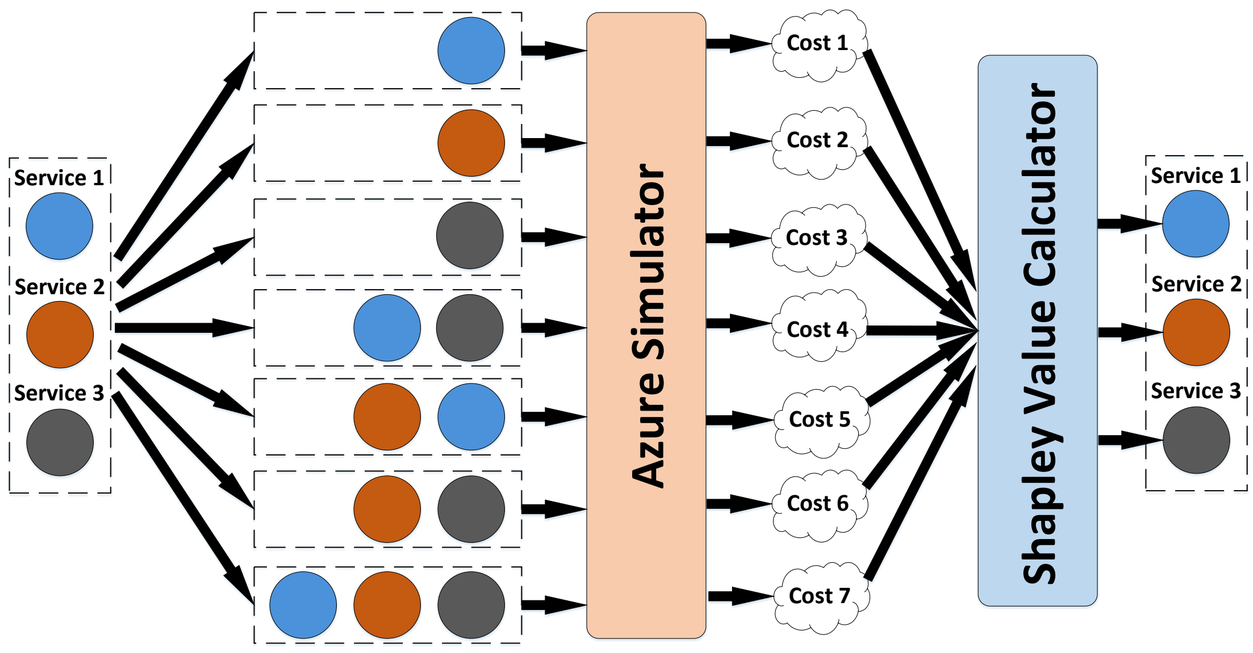

Experimental Results

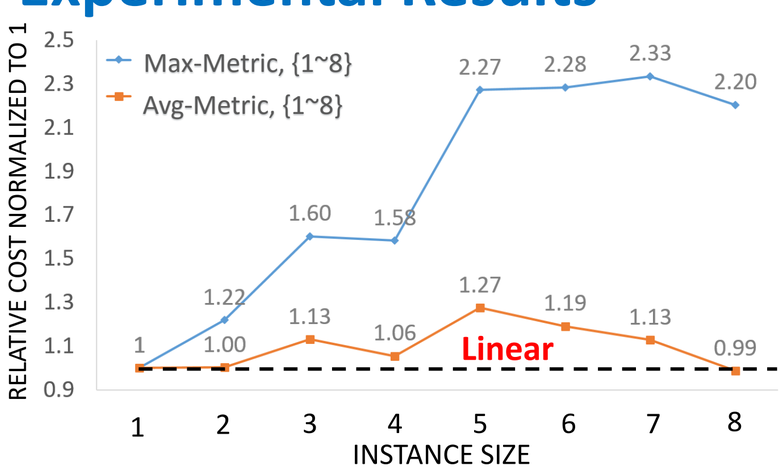

Service costs are non-linear in their instance size.

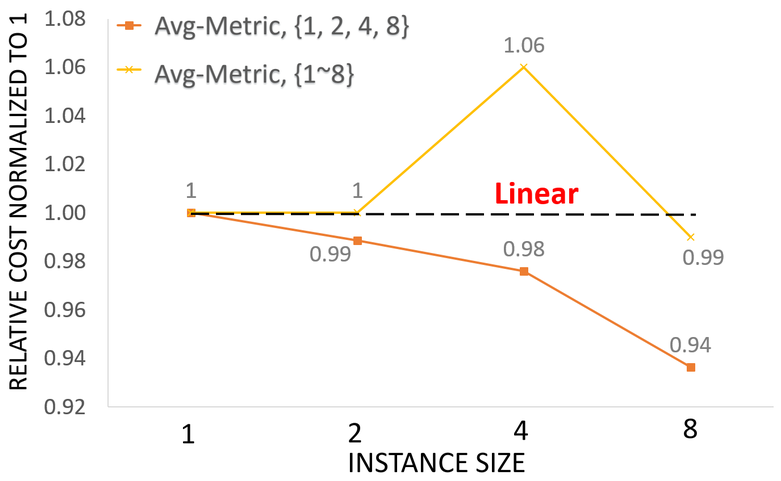

Instance type composition influences cost allocation. Instance types with low fragmentation incur less cost.

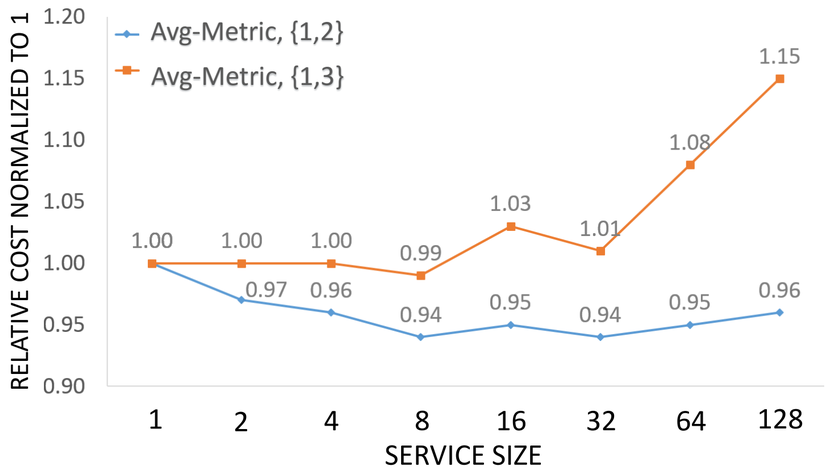

Service size also impacts costs. Large services are more expensive.

Summary

- **Systematic approach for cost allocation in the cloud using Shapley Value.**
- **Derive cost incurred by a specific type of service.**
- **Study impact on cost of workload distribution, service size, VM types, fault-tolerance constraints, etc…**