DIETI UNIVERSITA' DEGLI STUDI DI NAPOLI FEDERICO II

On Network Throughput Variability in Microsoft Azure Cloud

Valerio Persico, Pietro Marchetta, Alessio Botta, Antonio Pescapè

University of Napoli "Federico II" (Italy)

IEEE GLOBECOM 2015

San Diego, CA, (USA) December 7 2015

(Public) Cloud Paradigm: XaaS

An increasing number of services rely on Public Clouds

- Pay-as-you-go resources and no upfront investments
- Real-time provisioning
- *Autoscaling*

Cloud infrastructures are backed by huge investments from the providers

- Research
- Complex infrastructures

Public-Cloud Network Performance

- About the Cloud…
  *There is NO CLOUD, just other people's computers**

- About the network…
  Without high-performance networks,
  *there would be no such thing as cloud computing***

- What about its performance?

**http://fsfe.org*
***Mogul and Popa, 2012*

Public Cloud Networks

- Taxonomy:
  - Cloud-to-user
  - Inter-datacenter
  - Intra-datacenter

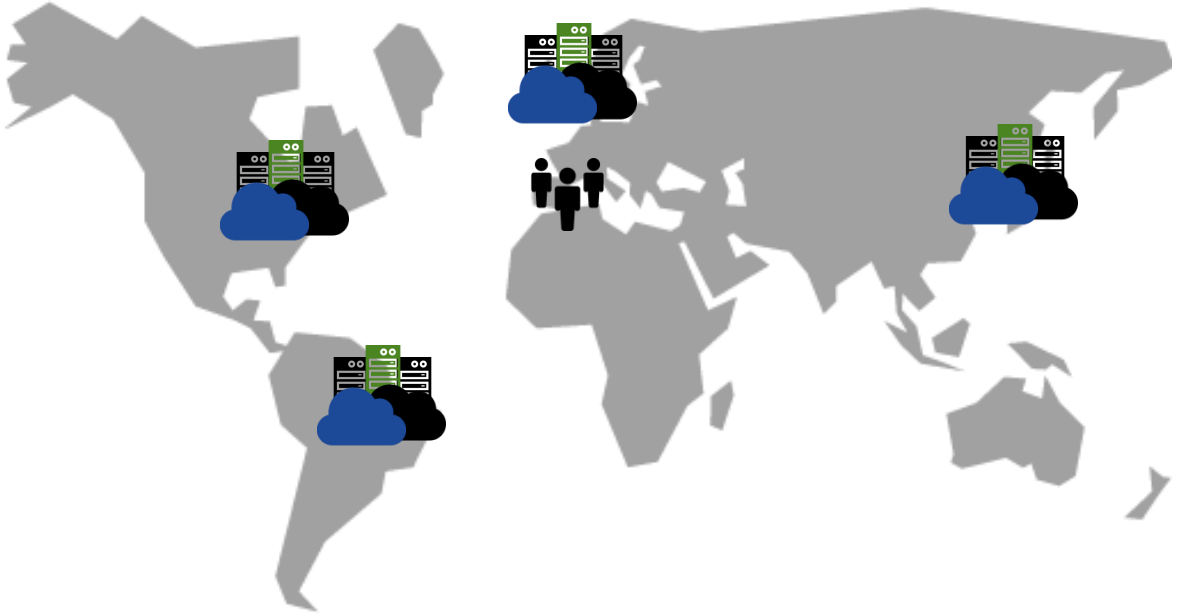

Public Cloud Networks

- Taxonomy:
  - **Cloud-to-user**
  - Inter-datacenter
  - Intra-datacenter

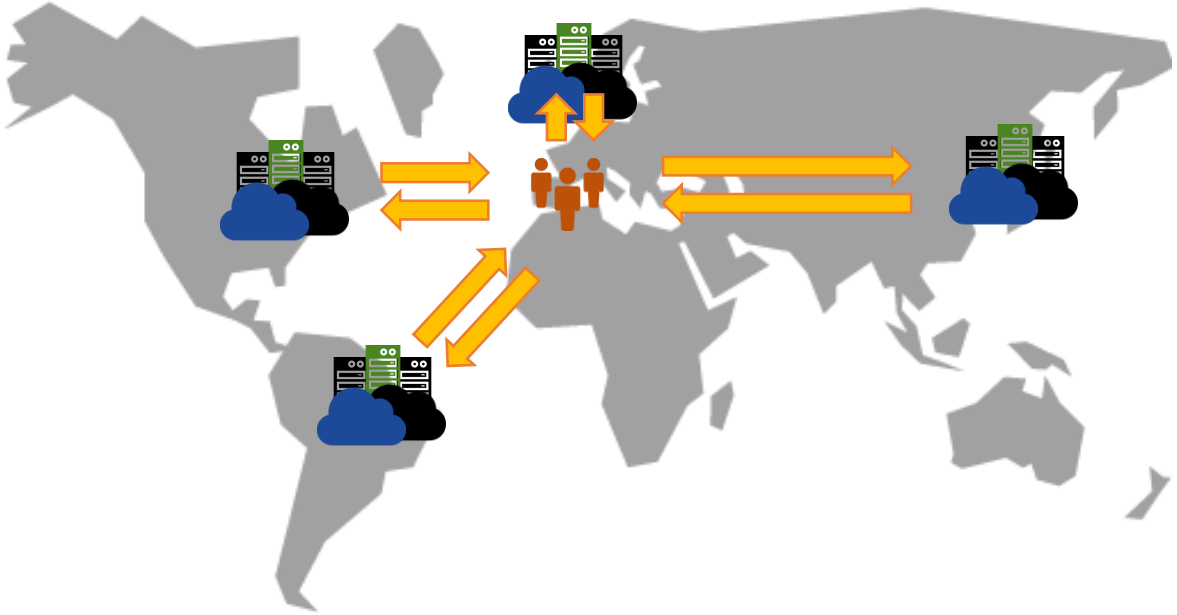

Public Cloud Networks

- Taxonomy:
  - Cloud-to-user
  - **Inter-datacenter**
  - Intra-datacenter

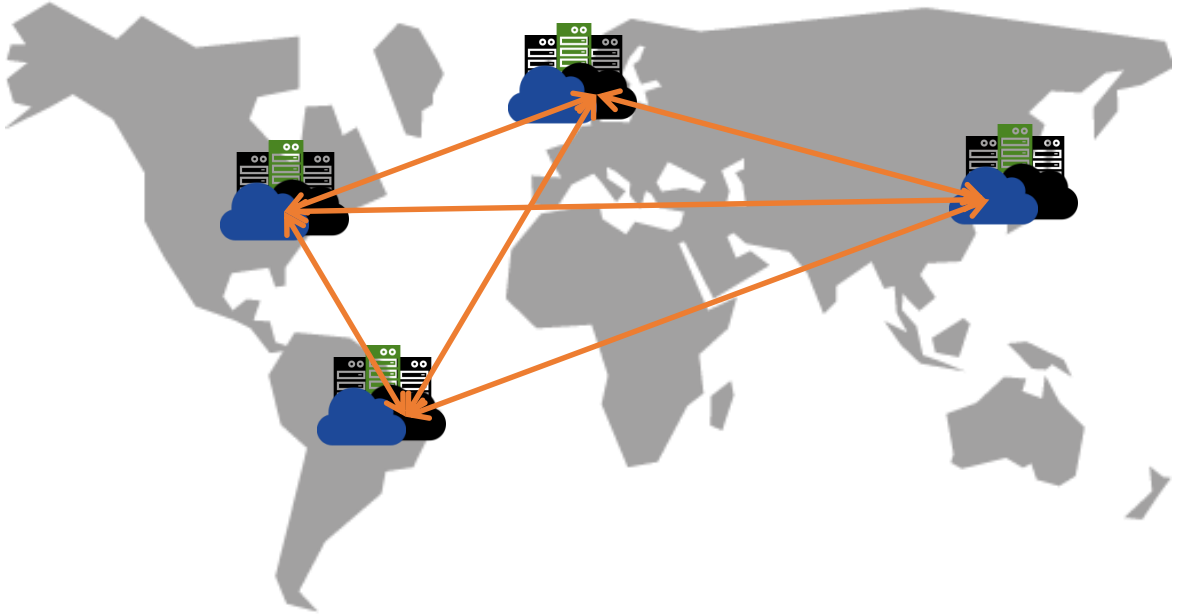

Public Cloud Networks

- Taxonomy:
  - Cloud-to-user
  - Inter-datacenter
  - **Intra-datacenter**

Intra-datacenter network

Why its performance is so important?

- 75% of the cloud traffic (e.g., multi-tier applications, scientific computation, etc.)
- Bottleneck for computation
- Its variability can severely compromise customer experience

Intra-datacenter IP traffic*

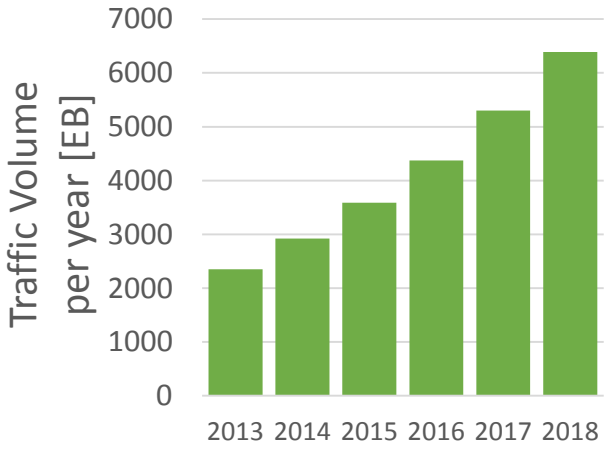

*Cisco Global Cloud Index: Forecast and Methodology

Monitoring cloud networks through a *non-cooperative* approach

- Monitoring (and benchmarking) the public-cloud network without relying on information restricted to the provider

- Purpose
  - Validation of the (poor) information supplied by the provider
  - Augmented view to support services and applications

- In this work
  1. We propose a methodology
  2. We focus on Microsoft Azure intra-datacenter network
  3. We characterize the performance in terms of network throughput and its variability

What we know (as general customers)

- All cloud providers provide (high-performance) network connectivity to customer VMs
- Plenty of prior work aimed at various specific approaches to sharing network resources among customers and providing cloud network guarantees
- Only *qualitative* information disclosed by providers (at most)

What we do **not** know (as general customers)

- Providers seldom make any promise about network performance
  - Customers suffer from highly-variable, unpredictable network performance
- What is the optimization goal of the provider?
  - Saving datacenter power consumption?
  - Guaranteeing better performance to specific sets of users?
- Datacenter topology and virtual machine (VM) location are kept hidden

Reference architecture

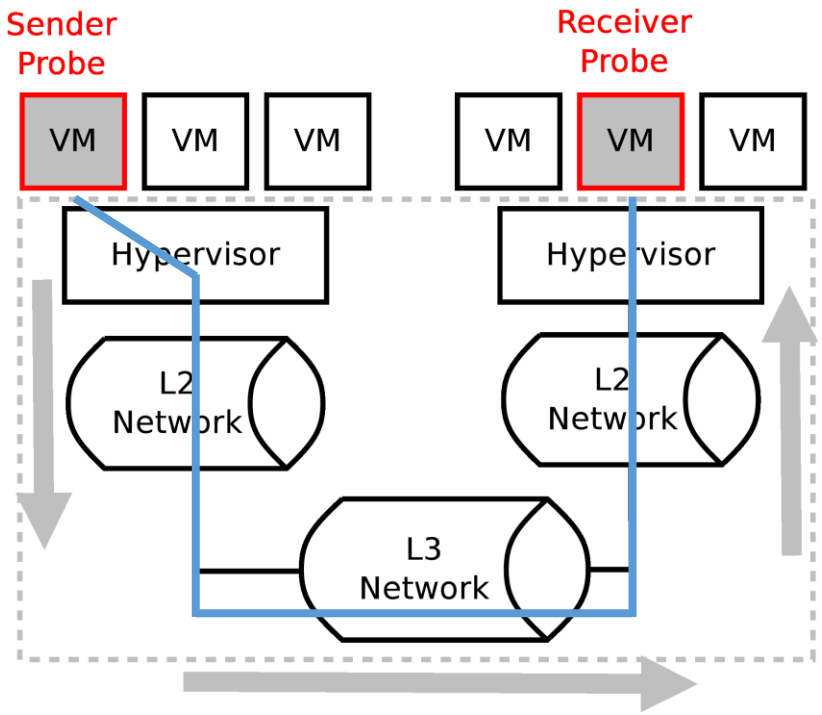

Intra-Cloud Network

- Generation of Synthetic Traffic
- From a Sender Probe to a Receiver Probe
- Black-Box approach

Reference architecture

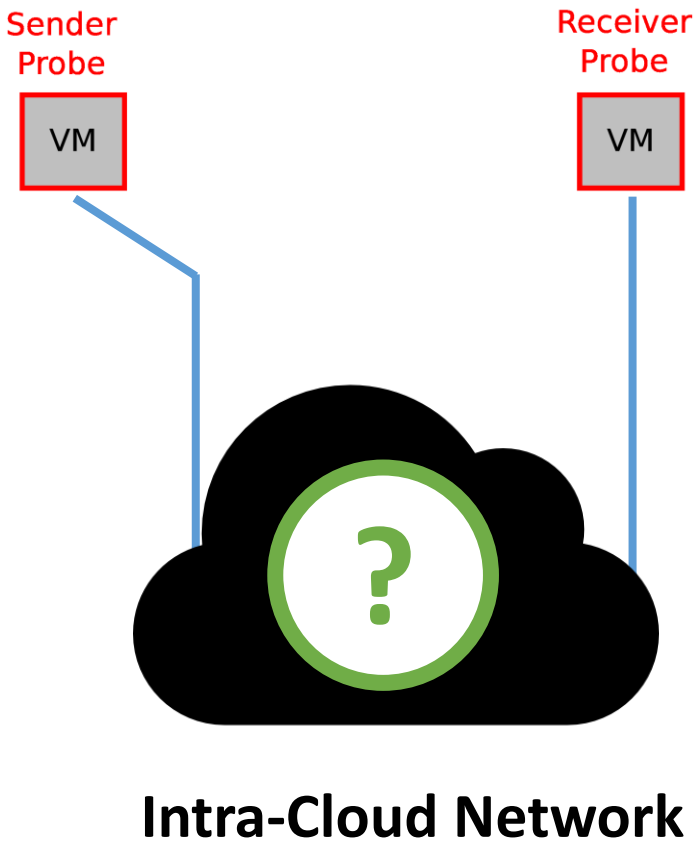

- Generation of Synthetic Traffic
- From a Sender Probe to a Receiver Probe
- Black-Box approach

Factors to identify scenarios

- Region
  - California (US)
  - Ireland (EU)
  - Singapore (ASIA)
  - Sao Paulo (BRA)
- VM type and size

  General purpose
  - Medium (M)
  - Large(L)
  - ExtraLarge (XL)
- Configuration
  - Same VNET (VN)
  - Same Affinity Group (AG)
  - None (NO)
- Transport protocol
  - TCP
  - UDP

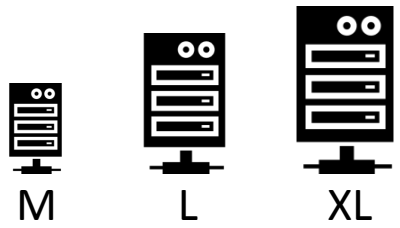

Factors to identify scenarios

- Region
  - California (US)
  - Ireland (EU)
  - Singapore (ASIA)
  - Sao Paulo (BRA)
- VM type and size

  General purpose
  - Medium (M)
  - Large(L)
  - ExtraLarge (XL)
- Configuration
  - Same VNET (VN)
  - Same Affinity Group (AG)
  - None (NO)
- Transport protocol
  - TCP
  - UDP

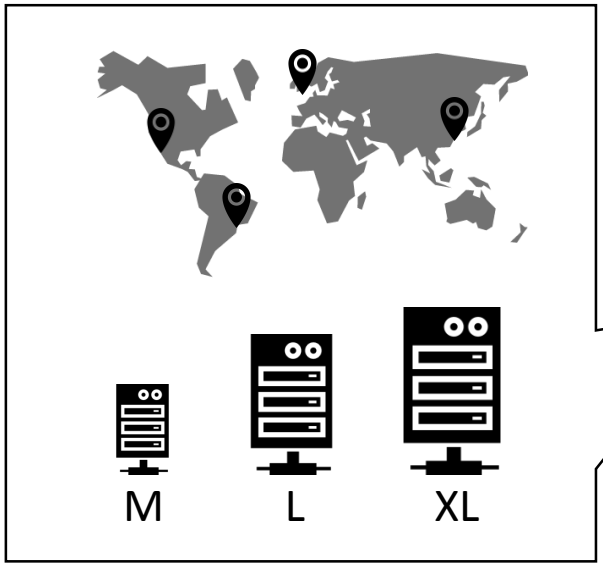

Directly impact costs

Experimental details and results

- 800-hour-long experimental campaign
- Intra-datacenter throughput
- Throughput Variability
  1. over time
  2. across different scenarios
  3. in the same scenario

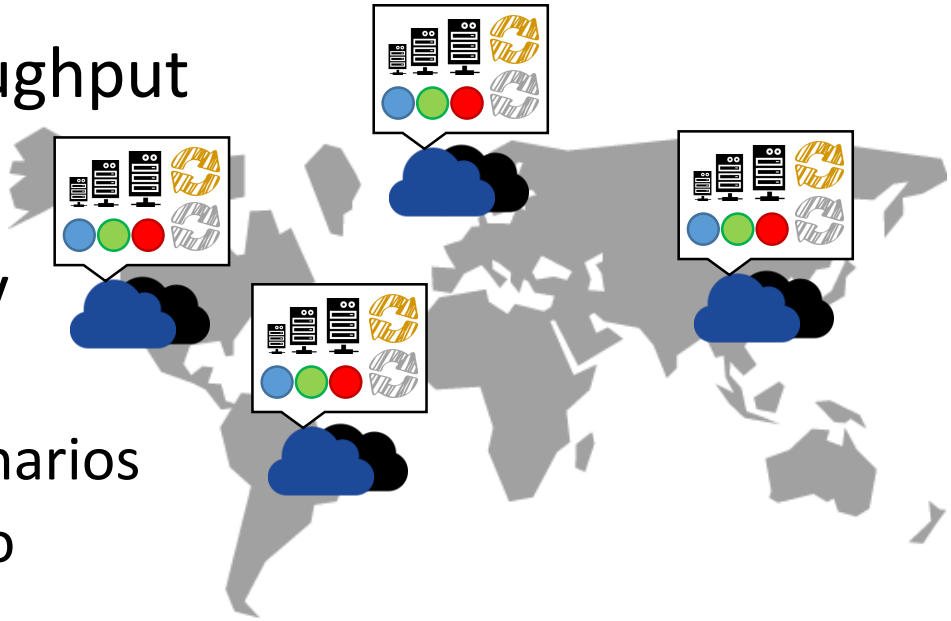

Variability over time

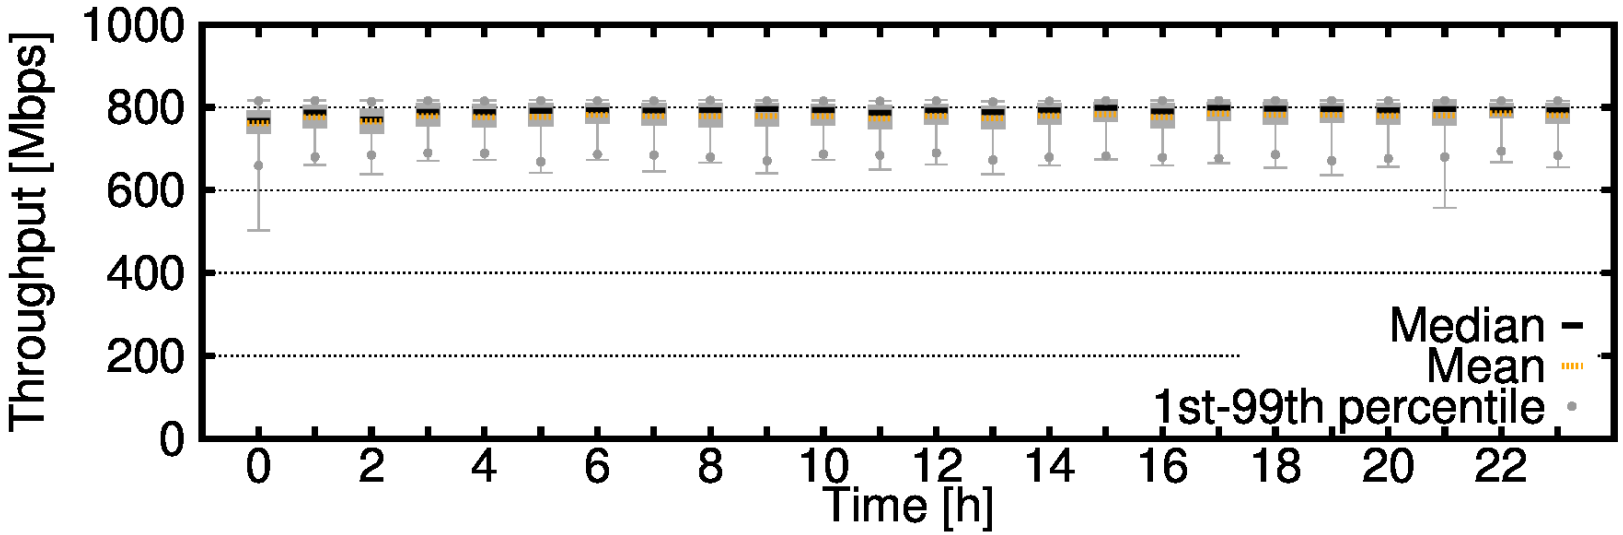

Among different experiments (average over 5-minute-long experiments)

- CoV (*Coefficient of Variation*) < 0.1
- What about the variability inside the same experiment?

Variability over time

Inside the same experiment (1-second-samples)

- CoV is always lower than 0.2
- Some factors impact variability more than others

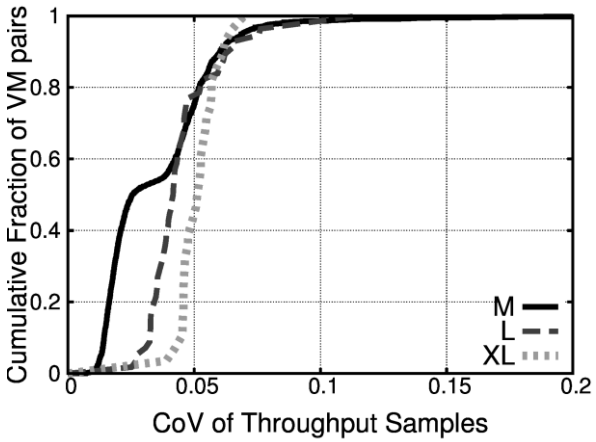

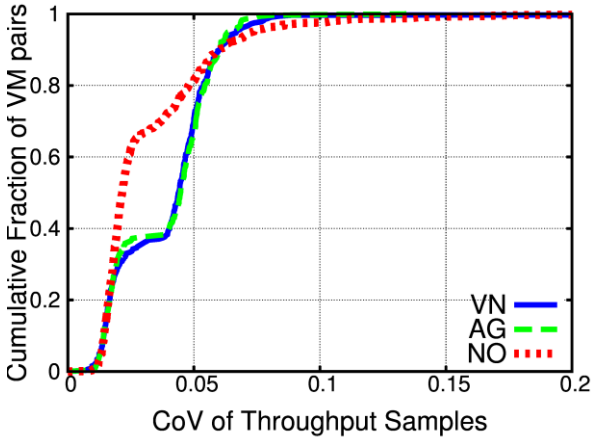

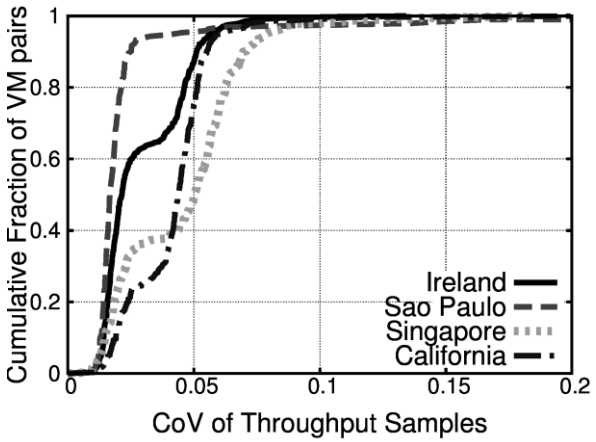

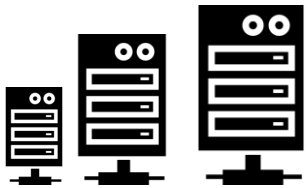

Variability across scenarios

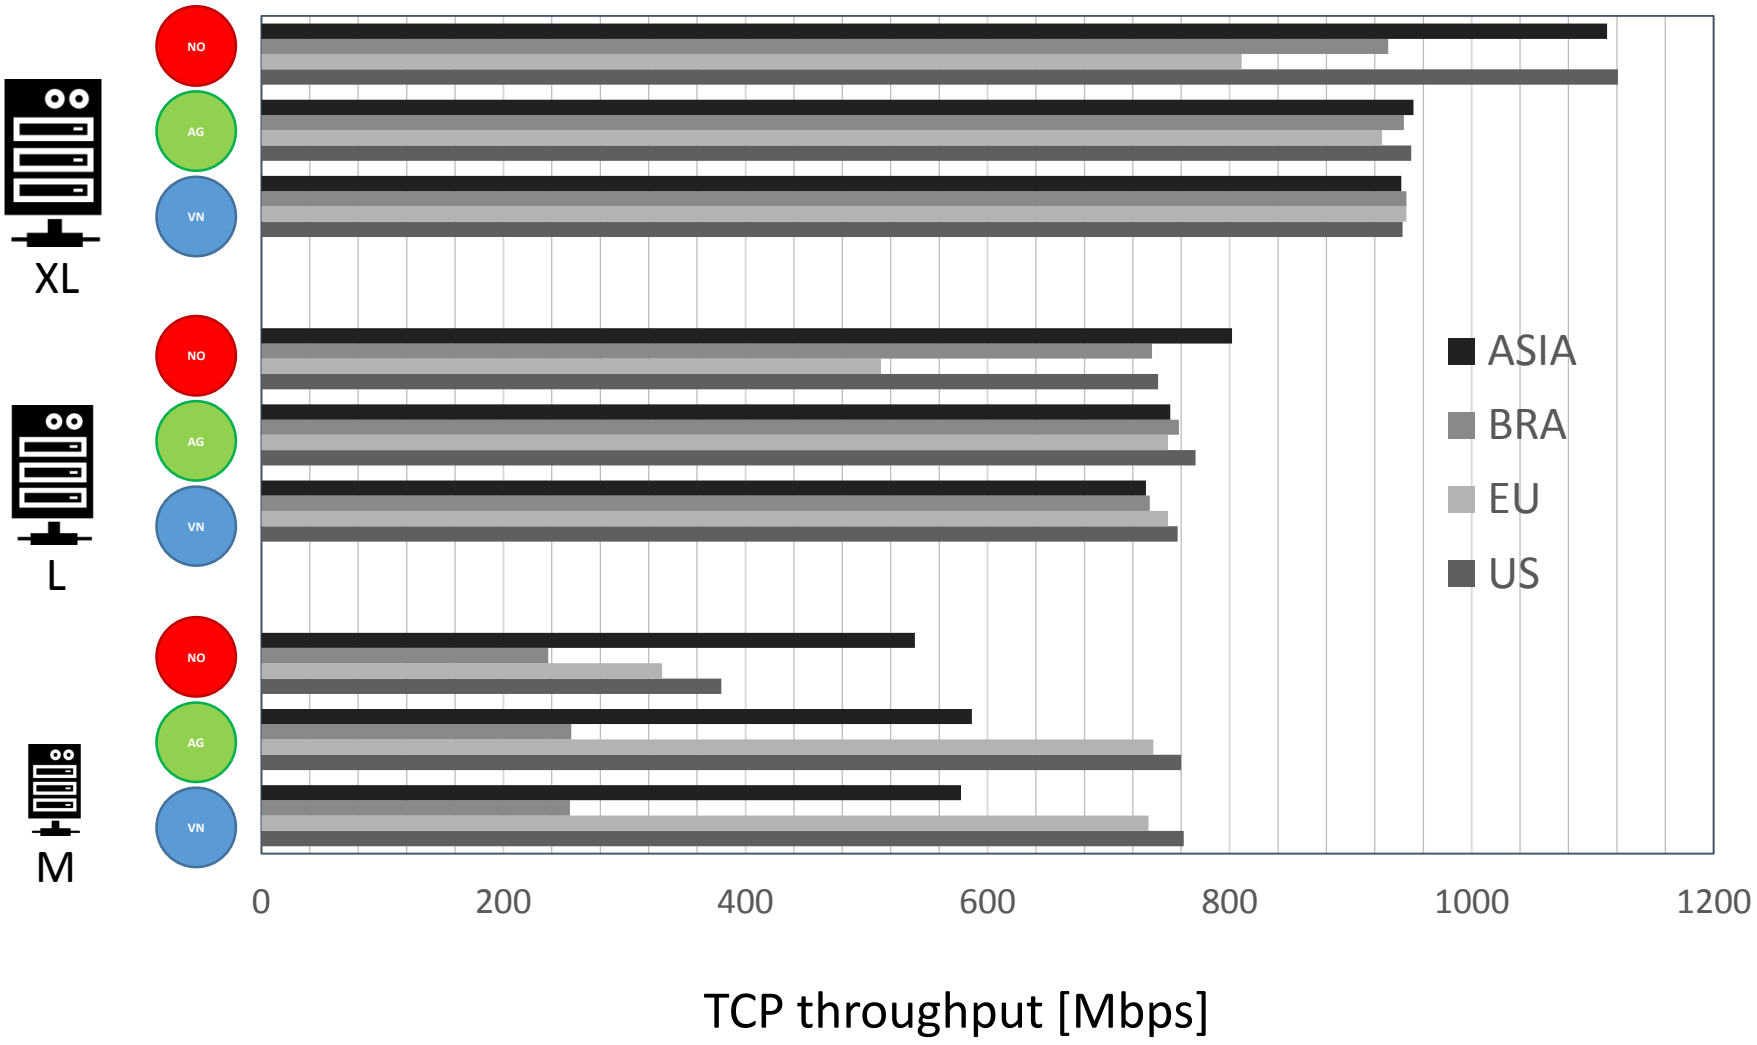

Variability in the same scenario

- Repeated experiments may provide very different values

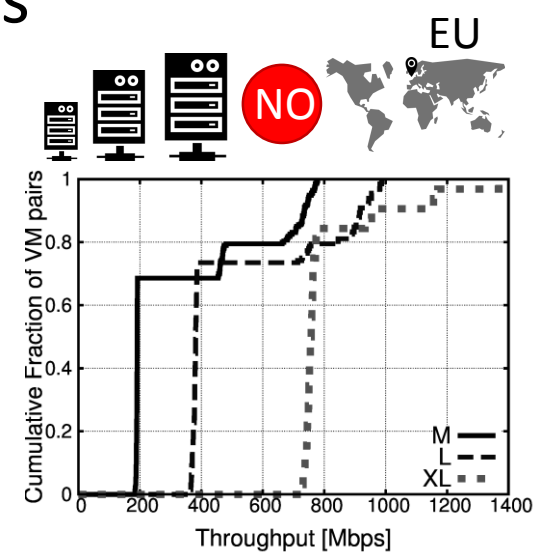

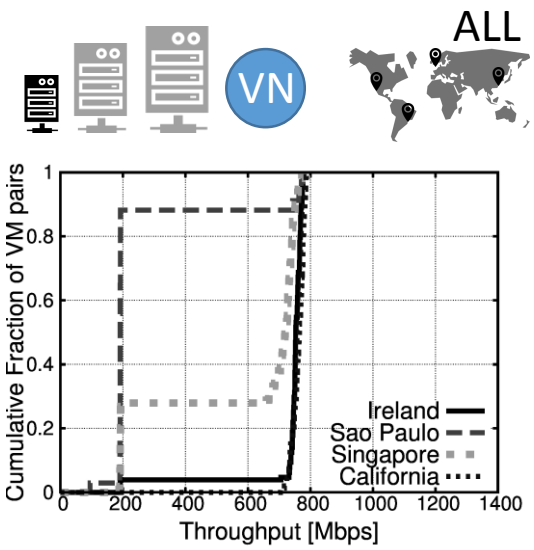

- Considering the absence of variability over time, *unlucky* customers should not expect any significant improvement

Minimum throughput guaranteed

	M	L	XL
US	186.6	374.6	929.5
EU	185.4	364.1	728.5
BRA	185.1	707.5	907.1
ASIA	186.0	718.1	935.0

1st percentile [Mbps]

Conclusion

- We propose a characterization of the achievable throughput of the intra-datacenter network for MS Azure through non-cooperative approaches

- Network throughput is stable over time
- Several factors under the direct control of the customer may influence the perceived performance
- Customers can derive deployment and usage guidelines
  - Performance prediction
  - Performance enhancement
  - Cost reduction

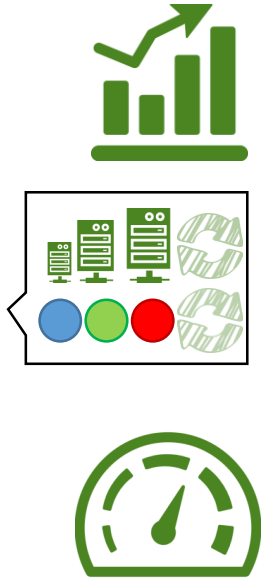

Questions?

valerio.persico@unina.it

http://wpage.unina.it/valerio.persico

DIETI UNIVERSITÀ DEGLI STUDI DI NAPOLI FEDERICO II

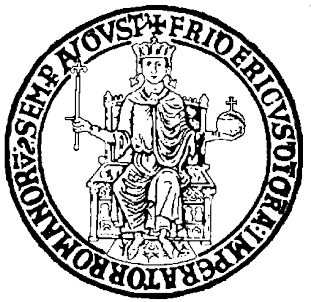

Icons made by Freepik from www.flaticon.com are licensed under CC BY 3.0