

# On Network Throughput Variability in Microsoft Azure Cloud

Valerio Persico, Pietro Marchetta, Alessio Botta, Antonio Pescapé  
University of Naples Federico II and NM2 srl (Italy)  
{valerio.persico, pietro.marchetta, a.botta, pescape}@unina.it

**Abstract**—The dependence of the industry on cloud-based infrastructures has grown much faster than our understanding of the performance limits and dynamics of these environments. An aspect only marginally analyzed in the past is related to the performance of the intra-cloud network connecting the virtual machines (VMs) deployed in the same data center. The few available works either do not exhaustively describe the adopted methodology or employed different approaches causing the analyses to be hard to replicate, and the results to be hard to compare. In addition, cloud customers can today highly customize their cloud environments while previous works considered only a few of the scenarios in which a customer may operate.

In this paper, we provide an intra-cloud network performance characterization of Microsoft (MS) Azure, a leading provider only preliminary investigated from this angle. We first propose and thoroughly detail a methodology to carry out similar analyses, thus encouraging its replication also in other contexts; then we apply this methodology to characterize the intra-cloud network performance in terms of maximum network throughput. More specifically, we investigate whether and how the achievable throughput between two VMs varies (i) over time; (ii) when the customer operates different decisions on VM size, network configuration, geographic region, and transport protocol; and (iii) when the customer operates the same decisions on these factors. Our analysis aims at addressing the gap existing in the literature by providing the most exhaustive and detailed results about the intra-cloud network performance for MS Azure today available.

## I. INTRODUCTION

An increasing number of services and applications is now delivered through cloud-based infrastructures and a large number of organizations more and more depends on cloud. To drive this increasing demand, cloud providers made huge investments supporting the public cloud paradigm where anyone can take advantage of virtual resources at varying performance and cost. This result is achieved through the adoption of sophisticated virtualization strategies for sharing resources among a great number of largely uncoordinated and mutually untrusted customers. While cloud customers are not aware of the complexity of similar mechanisms, the Quality of the Service (QoS) they perceive is deeply affected by them. Although cloud customers need guarantees on the QoS, cloud providers often provide no or only coarse grained information about the performance a customer should expect from the cloud [4]. A perfect example in this regard is represented by the intra-cloud, high-performance network, i.e. the network connecting resources deployed inside the same data center. All providers grant high-performance network connectivity to their customers but they provide no more than qualitative information about its performance, mainly due to security and commercial

reasons [14], [13], [13], [5]. This limited information strongly weakens the ability of customers to understand whether and how the performance of their applications is impacted by the network [17]. This aspect is much more awkward since the intra-cloud network proved to be a potential bottleneck for data- and computation-intensive applications [16], [17], [7], [9]. In this context, it is not surprising that measuring intra-cloud network performance has very recently attracted interest from the research community [16], [15], [11], [9], [14], [17], [12], [8]. From the pure point of view of the cloud customer, the adoption of *non-cooperative* approaches – that do not rely on the support of the provider – seems the only viable solution to obtain detailed information on the network performance. Hence, recently several services and frameworks providing cloud customers with benchmarks also for cloud network performance have been released [2], [1]. Using these approaches to exhaustively characterize the intra-cloud network performance of a provider is, however, also particularly challenging for a number of reasons including: (i) the high number of possible scenarios in which a customer may operate; (ii) the huge cost of the analysis due to the pay-as-you-go paradigm, exacerbated by the need of repeated measurements to provide meaningful performance statistics; (iii) the continuous evolution of the cloud infrastructures, features, and services; (iv) the strongly limited visibility on the heterogeneous resource sharing and virtualization techniques adopted by cloud providers. For all these reasons, despite the efforts in literature, the overall picture of the intra-cloud network performance in public clouds is scattered and incomplete: only a few scenarios have been taken into account and not all the cloud providers have been deeply investigated.

In this paper, we aim at improving the knowledge on the performance of intra-cloud networks by focusing on Microsoft (MS) Azure, one of the leading public cloud provider [10]. Previous works only marginally analyzed the performance of this cloud component for MS Azure by considering very few limited scenarios. To fill this gap, we provide a detailed characterization of the intra-cloud network performance for this provider in terms of maximum achievable throughput. More specifically, we provide the following main contributions: (i) we propose and adopt a thorough methodology to characterize the intra-cloud network performance that takes into account all the aspects causing similar analyses to be extremely complicated; (ii) thanks to 800 hours of experimentation, we investigate whether and how the maximum network throughput achievable between two VMs deployed on MS Azure varies (a) over time; (b) across different scenarios – i.e. when the customer operates different decisions on Virtual Machine (VM)

size, network configuration, geographic region, and transport protocol; and (c) within the same scenario – i.e. when the customer operates the same decisions on these factors. The documented results represent the most detailed and exhaustive study of the intra-cloud network performance of MS Azure today available.

The paper is organized as follows: Sec. II positions the paper against the related work; Sec. III describes the methodology we propose and adopt for this type of analysis; Sec. IV discusses the main results on the achievable network throughput and its variability; finally, Sec. V ends the paper with concluding remarks.

## II. RELATED WORK

Several works in literature analyzed the maximum network throughput achievable in a public-cloud environment by leveraging non-cooperative approaches [16], [15], [11], [9], [14], [12], [8]. However, only few of them analyzed the MS Azure public cloud and investigated the performance of the intra-cloud network. Hill et al. [8] measured the TCP network throughput among 10 VMs of small size reporting throughput values ranging between 80 and 800 Mbps. They also reported great throughput variability over time between fixed VMs. Tudoran et al. [16] considered 20 small- and 20 extra-large-sized VMs reporting that extra-large VMs did not exhibit the high TCP throughput variability exposed by small VMs. Finally, Li et al. [11] reported that the TCP throughput between two VMs placed in the same data center is always close to 800 Mbps.<sup>1</sup> None of these works aimed at characterizing the intra-cloud network performance as a primary goal. All of them only marginally or preliminarily analyzed the network performance by considering a limited number of scenarios in which a cloud customer may operate. Moreover, the methodology adopted in these works is not thoroughly described making the results hard to compare. Finally, compared to the time in which these analyses have been conducted, today MS Azure allows the customers to heavily customize their own cloud environments. Therefore, the impact on the network performance of the decisions operated by the customers during the VM-deployment phase is today largely unknown. To fill all these gaps, we propose and thoroughly describe a methodology, taking into account all the relevant factors having an impact on the network performance (VM size, geographical region, configurations, transport protocol). Thanks to this methodology, we provide the most complete and deepened network throughput characterization available today for the MS Azure cloud.

## III. MEASURING INTRA-CLOUD NETWORK THROUGHPUT

In this section, we first discuss the reference architecture and tools we used in our experimentation (Sec. III-A). Then, we describe and motivate the subset of scenarios of interest we took into account (Sec. III-B).

### A. Reference architecture and tools

In our experiments, we adopt the reference architecture reported in Fig. 1: we aim at measuring the maximum network

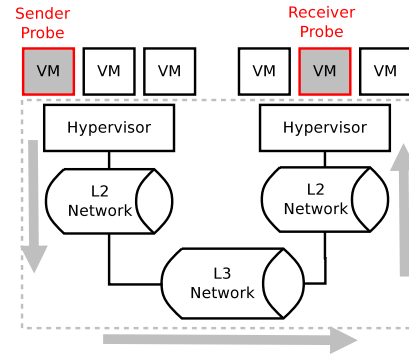


Fig. 1: **Reference architecture.** We measured the network throughput between a pair of VMs deployed on the same data center. L2/L3 devices as well as hypervisors are considered as part of the network.

throughput achievable between a pair of VMs deployed by the same cloud customer on the same data center. For our experiments, these VMs have been instrumented with a standard operating system and all the necessary network measurement and diagnostic tools we need for estimating the network throughput. As depicted in the figure, the traffic generated by a VM sender normally traverses the hypervisor layer at the sender side first. Then it flows through L2/L3 devices and middleboxes composing the intra-cloud high performance network. Finally, the traffic reaches the hypervisor at the receiver side before being delivered to the receiver VM. Note that as a normal cloud customer, we are not aware of the specific location of the sender and receiver VMs, which may also be hosted and managed by the same hypervisor. Furthermore, we adopt a black box approach for which the L2/L3 devices as well as the hypervisors are considered as part of the network connecting sender and receiver VM. Basically, we consider as *network* all the logical and physical components interposed between the two VMs. We use the network measurement tool named *nuttcp* to measure the network throughput between the VMs: this tool allows us to measure the raw TCP or UDP application layer throughput by transferring memory buffers from the sender to the receiver VM.

### B. Identifying scenarios of interest

In MS Azure, a customer can highly customize her cloud environment by acting on a predefined set of *factors* during the deployment phase. We aim at assessing whether these factors have an impact on the network throughput measured between pairs of VMs and quantifying it. In our analyses, we consider the following factors: VM type and size, geographical region, configuration, transport protocol, and VM placing and relocation. We clarify in the following how we selected specific values for each factor: the combination of these values determines the scenarios of interest we considered in this paper.

**VM type and size.** When deploying a VM on the cloud, a customer can choose its *type* and *size*. The VM type – also referred to as *series* in the Azure documentation [3] – indicates a family of VMs optimized for a given task (storage, computation, etc.). Once the type is selected, the customer can decide the size of the VM to further specify storage and computation capabilities. Both the VM type and size influence the hourly cost of the VM. In our experiments, we used *general purpose* VMs. These VMs provide a good balance between

<sup>1</sup>Note that the authors used labels instead of names to identify the different providers. We inferred MS Azure performance among their results by looking at the different geographical regions of the data centers.

TABLE I: Selected sizes for VM of A series (Mar 2015 [3]).  
\*: prices may vary across regions.

VM Size	Acronym	CPU cores	RAM (GB)	Max data disks	Approx. Monthly Cost* (\$)
medium	A2	2	3.5	4	from 89 to 119
large	A3	4	7	8	from 179 to 238
xlarge	A4	8	14	16	from 357 to 476

CPU and memory, making them a good choice for many applications [3]. We also considered VMs of three different sizes named in the Azure documentation as A2, A3 and A4; we simply refer to them as *medium*, *large* and *xlarge*, respectively. Tab. I reports more details on the VMs adopted in our analyses. Note how MS Azure provides details only about RAM, CPU, and storage capabilities: neither quantitative nor qualitative description of the expected network performance is provided to the customer.

**Region.** A MS Azure customer can deploy her VMs in data centers spread in 17 geographically distributed regions. The choice is typically made according to technical and legal requirements. In this paper, we aim at measuring the network throughput between VMs placed inside the same region. Since this operation is extremely costly due to the pay-as-you-go paradigm, we were forced to focus on a subset of regions. To obtain a representative picture of the intra-cloud network performance of the cloud provider, we selected four regions in different continents: West US (*California*), North Europe (*Ireland*), Southeast Asia (*Singapore*), Brazil South (*Sao Paulo*). These regions serve customers from distinct zones and may be consequently subjected to different workloads. They also became available at different points in time, thus potentially leveraging different technologies. Only few works in literature investigated the intra-cloud network performance in multiple regions. For MS Azure, this aspect has been investigated only in Li *et al.* [11] where maximum throughput variability across the regions is reported.

**Configuration.** During the deployment process, the customer can deploy her VMs in the same *affinity group* or *virtual network (VNET)*. According to the MS Azure documentation [3], the affinity group is a way to “group cloud services by proximity to each other in the Azure data center in order to achieve optimal performance”. VNETs, instead, are a way to allow services and VMs to “communicate securely with each other”. Placing a VM in a given affinity group or VNET is optional and comes with no additional cost for the customer. During our experimentation, creating a VNET also required the specification of an affinity group while this requirement has been now removed. In our analysis, we considered the following three *configurations*: (i) the two communicating VMs are deployed in the same (cloud-only) VNET (hereafter simply  $CFG_1$ ); (ii) the two VMs are deployed in the same affinity group ( $CFG_2$ ); (iii) the customer provides no preferences ( $CFG_3$ ).

**L4 transport protocol.** In our experiments, we measured both UDP and TCP network throughput. UDP is useful to analyze the performance of the raw IP traffic. Indeed, UDP adds no closed loop-control, leaving the complete control on the traffic to the application that generates it, no matter of the state of the network. On the other hand, TCP throughput,




Fig. 3: **An instance of long-term campaign.** The maximum intra-cloud network throughput between two medium VMs deployed on Singapore shows a slight variability in each 5-minute-long experiments over 24 hours. However, the average/median value per experiment is strongly stable over time.

which is governed by flow and congestion control, makes the generated traffic to be subjected to the status of the network path and provides information on the performance of the numerous TCP-based applications. In traditional scenarios, UDP typically achieves higher network throughput than TCP [6].

**VM placing and relocation.** A cloud customer does not have neither visibility nor influence on which specific physical machine in the data center hosts her VMs. In our analyses, we also obtained empirical evidences suggesting the presence of mechanisms nullifying traditional topology-discovery strategies such as ICMP filtering. These mechanisms prevented us to gain knowledge about the relative position of the VMs through the adoption of tools like traceroute and ping. The only way we have to indirectly evaluate the impact of VM placement is to *relocate* the VMs multiple times, i.e., we ask the cloud provider to terminate the VMs and recreate them from scratch. By measuring each time the network throughput, we can derive a more exhaustive picture of the network throughput a generic cloud customer may receive when deploying on the cloud two VMs with given characteristics.

The proposed methodology is general enough to be employed for monitoring also other relevant network performance indices such as latency, jitter, and packet loss.

#### IV. EXPERIMENTAL RESULTS

By applying the methodology described in Sec. III, we performed several experimental campaigns to measure the network throughput achievable in the MS Azure cloud. In this section, we discuss the main results obtained with 800 hours of experimentation.

##### A. Network throughput variability over time.

The first aspect we investigated is whether and how the network throughput varies over time once the two communicating VMs are deployed on the cloud. With this analysis, we aim at verifying whether the performance of the network slice assigned by the provider to a customer dynamically changes. We discovered that the maximum intra-cloud network throughput is very stable on average, while a slightly higher variability is observed when the instantaneous values are taken into account. To carry out this analysis over large periods of time, we performed a series of 24-hour-long campaigns. In each campaign, (i) we created and deployed two same-sized




Fig. 2: **TCP instantaneous throughput variability over time.** Deploying larger VMs led to higher variability (2a). Sao Paulo and Ireland regions (2b) as well as  $CFG_3$  (2c) exposed to lower variability.

VMs in a region and monitored them for 24 hours; (ii) every hour the sender VM transferred synthetic traffic to the receiver VM at the highest possible rate with `nattcp` for 5 minutes; (iii) for each 5-minute-long experiment, we registered the network throughput achieved. We launched multiple campaigns by considering all the possible combinations of VM sizes, regions, configurations, and L4 protocols. To quantify how much the network throughput varied, we rely on the Coefficient of Variation (CoV), defined as the ratio of the standard deviation to the mean: a very low CoV implies that the standard deviation is negligible compared to the mean, thus underling a low throughput variability. In *all* the long-lasting campaigns we performed, the estimated CoV was always below 0.01, i.e. the average network throughput between two fixed VMs computed over 5-minute-long experiments in 24-hours of experimentation proved to be very stable. Hence, no daily pattern was observed. This happened independently from the VM size, the cloud region, the configuration, or the L4 protocol. An instance of 24-hour campaign for two medium-sized VMs deployed in the Singapore region is reported in Fig. 3: note how the average/median values do not significantly vary over time. We also investigated the variability of throughput instantaneous values (1-second granularity). For each experiment we performed, the CoV computed over the instantaneous values was always lower than 0.1, i.e. the instantaneous values only slightly fluctuate around the average. In case of TCP traffic, the factors seem having a non-negligible impact on the network throughput variability over time as shown in Fig. 2. In this case, we noticed how (i) the variability over time was slightly higher for larger VM sizes – see Fig. 2a; (ii) deploying medium-sized VMs in regions such as Singapore and California exposed to higher throughput variability compared to San Paulo and Ireland – see Fig. 2b; (iii) a higher variability affected medium-sized VMs placed in the same VNET ( $CFG_1$ ) or affinity group ( $CFG_2$ ) compared to the case in which we let the cloud provider taking a decision about it ( $CFG_3$ ) – see Fig. 2c. The similar factor impact was not observed in case of UDP traffic.

According to these results, when a customer deploys two VMs on Azure, the communication between them can achieve a maximum throughput which on average is not subjected to significant changes over time, no matter the size, the region, the configuration, or the transport protocol. At the same time, instantaneous throughput oscillations exist and are larger for larger VM sizes, for specific regions, and in specific configurations. Since the average network throughput does not significantly change over time, measuring the throughput between two VMs even over a short period of time provides a

good indication of the throughput that will be achieved in the future between the same VMs. We leveraged this finding to design the experimental campaigns discussed in the following.

### B. Network throughput variability across scenarios.

The second aspect we investigated is whether and how the network throughput varies in different scenarios, i.e. when the customer operates different decisions during the deployment phase. To this end, we performed 5.5K short campaigns. In each campaign, (i) we created from scratch and deployed two same-sized VMs; (ii) we measured the network throughput between them for 5 minutes and registered the average throughput achieved with both the L4 protocols; finally (iii) we terminated the VMs. We considered again all the possible combinations of sizes, regions, configurations, and L4 protocols. We performed multiple campaigns for each combination, collecting multiple network throughput samples. Note that, as already described in Sec. IV-A, the network throughput measured over 5 minutes represents a good indication of the performance that will be achieved in the next 24 hours.

Tab. IIa and IIb provide the average value computed over all the network throughput samples obtained in each scenario for TCP and UDP traffic, respectively. For the first time in literature, we provide a breakdown of the network throughput in Azure on such a large number of factors. A closer look at these results suggests the following findings: first of all (i) all the factors have a non-negligible impact on the average network throughput – different values can be observed by reading these tables by rows (i.e. same region, but different sizes and configurations) and by columns (i.e. same VM size and configuration, but different regions); (ii) differently from what happens in traditional environments, UDP traffic does not clearly achieve better performance than TCP, although these experiments were performed in the same conditions and closely in time; (iii)  $CFG_1$  and  $CFG_2$  typically guarantee higher network throughput than  $CFG_3$ , with only few exceptions such as the case of xlarge-sized VMs deployed in Singapore and California; (iv) in general, deploying VMs of larger size leads to an improvement also from the network performance point of view, although this result is not clearly acknowledged in the Azure documentation; also interestingly, this does not happen always: in both California and Ireland, when  $CFG_1$  or  $CFG_2$  is in place, medium VMs receive network throughput very similar to large VMs; (v) the network throughput significantly varies across the regions, i.e. VMs of a given size and in a particular configuration may receive

TABLE II: Average network throughput achievable in different scenarios [Mbps]. Underlined values refer to scenarios exposing high variability. We considered as high variable the scenarios associated to a standard deviation larger than 100 Mbps. For the remaining scenarios the standard deviation was found to be always smaller than 20 Mbps.

(a) TCP traffic.

Region	Medium VMs			Large VMs			Xlarge VMs		
	CFG <sub>1</sub>	CFG <sub>2</sub>	CFG <sub>3</sub>	CFG <sub>1</sub>	CFG <sub>2</sub>	CFG <sub>3</sub>	CFG <sub>1</sub>	CFG <sub>2</sub>	CFG <sub>3</sub>
California	761.5	760.1	<u>380.4</u>	757.8	772.1	<u>741.0</u>	943.0	950.8	<u>1121.6</u>
Ireland	<u>733.6</u>	<u>727.9</u>	<u>331.2</u>	749.3	749.0	<u>512.3</u>	946.5	926.5	<u>810.3</u>
Sao Paulo	<u>255.6</u>	<u>256.7</u>	<u>237.0</u>	734.0	758.9	736.3	946.6	944.5	931.3
Singapore	<u>578.4</u>	<u>587.3</u>	<u>450.2</u>	731.9	751.5	<u>802.0</u>	942.3	952.2	<u>1112.9</u>

(b) UDP traffic.

Region	Medium VMs			Large VMs			Xlarge VMs		
	CFG <sub>1</sub>	CFG <sub>2</sub>	CFG <sub>3</sub>	CFG <sub>1</sub>	CFG <sub>2</sub>	CFG <sub>3</sub>	CFG <sub>1</sub>	CFG <sub>2</sub>	CFG <sub>3</sub>
California	750.0	746.4	<u>354.6</u>	754.9	758.6	<u>610.9</u>	936.0	925.8	<u>738.9</u>
Ireland	<u>734.4</u>	<u>723.6</u>	<u>315.5</u>	768.0	760.0	<u>435.8</u>	882.3	845.9	<u>686.4</u>
Sao Paulo	<u>250.2</u>	<u>256.4</u>	<u>231.6</u>	743.9	753.2	737.8	858.3	863.9	859.7
Singapore	<u>565.2</u>	<u>568.2</u>	<u>437.1</u>	740.1	747.8	<u>687.0</u>	900.6	903.7	<u>703.2</u>

higher throughput in specific regions. This is true especially for medium VMs for which the throughput in Sao Paulo proved to be much lower than what can be achieved in other regions.


When similar analyses were conducted in the past, only few factors were under the control of the Azure customers [11], [8], [16]. Nowadays, instead, the customer can highly customize her cloud environment by acting on a wide set of factors. In this new context, we demonstrated that considering one or few combinations of factor values may lead to an incomplete image of the network performance a customer may experience. Since considering all the possible combinations is extremely costly both in terms of money and time, researchers willing to conduct similar analyses must very carefully design and describe the adopted methodology. This is essential not only to cross-validate and replicate the reported findings but also to understand their scope and validity. The methodology we proposed can be used to face the complexity of a similar task. For the Azure customers, these results clearly demonstrate that the decisions operated during the VM-deployment phase are very important at least from the network performance point of view. In this regard, Tab. IIa and Tab. IIb may support those customers willing to optimize their network performance by selecting the best combinations of VM sizes, regions, and configurations. For instance, we clearly demonstrated that it is convenient for the customers to deploy always the VMs inside the same VNET ( $CFG_1$ ) or the same affinity group ( $CFG_2$ ) since these configurations provided higher network throughput at no additional cost: this gain is more remarkable in specific regions such as California and Ireland.

### C. Network throughput variability inside each scenario.

In some of the considered scenarios, repeated measurements provided very different values of throughput. These scenarios are the ones for which Tab. II reports underlined values. In these cases, even if the customer operates exactly the same decisions during the deployment phase, the provider may grant very different levels of network throughput for her. This result is particularly interesting since according to Sec. IV-A, an *unlucky* customer receiving low network performance should not expect any significant improvement over time. Some extreme cases where multiple experiments

in the same scenario exposed completely different values of network throughput are reported in Fig. 4. Note how the network throughput is quantized around few different values: although less marked, this behavior was qualitatively observed also for the other affected scenarios not detailed here due to space constraints. Fig. 4a reports the network throughput measured in Ireland across different campaigns where we adopted  $CFG_3$  and two VMs of the same size: medium VMs received 200Mbps, 450Mbps, and more than 650Mbps in 70%, 10% and 20% of the campaigns, respectively. Hence, different customers deploying medium VMs on Ireland with  $CFG_3$  may benefit from very different throughput values. A similar behavior has been qualitatively observed also for larger VMs. Interestingly, medium VMs received a network throughput higher than 75% of the campaigns involving large VMs for about 30% of the campaigns. Similarly, in 20% of the cases, large VMs achieved higher network throughput than 85% of the campaigns involving xlarge VMs: adopting larger VMs not always leads to higher network performance. Other interesting cases are reported in Fig. 4b showing how a customer deploying medium VMs with  $CFG_1$  in a given region may benefit from one of two very different network throughput values. In Ireland and California, the customer receives almost always high network throughput (760 Mbps). In Sao Paulo, instead, the customer often receives low network throughput (200 Mbps) but in 10% of the cases, lucky customers can take advantage of 760 Mbps. In the light of the observed variability, we also investigated the minimum achievable throughput a customer should expect from Azure when she adopts VMs of a given size. Tab. III reports the 1<sup>st</sup> percentile of the network throughput samples we registered across all the experiments for each VM size and for both L4 protocols: empirical evidences suggest that independently from the cloud region or the configuration, a customer adopting larger VMs at higher fee can definitely achieve higher network performance in the worst case.

According to these results, an Azure customer must be aware that even when the same decisions are operated during the deployment process, the network throughput received may significantly vary. To this regard, one approach a customer may employ to achieve high network performance is to (a.) select



(a) Region: Ireland; Configuration:  $CFG_3$ ; L4: TCP.

(b) Size: medium; Configuration:  $CFG_1$ ; L4: TCP.

Fig. 4: **Network throughput variability inside the scenarios.** A customer operating the same decisions during the deployment phase may receive very different network throughput.

the desired combination of factor values during the deployment phase according to her needs, costs and expected network performance; (b.) deploy the VMs and measure the maximum network throughput with classic network diagnostic tools such as the ones we adopted for this paper; (c.) compare the obtained throughput with the reference values we reported in Tab. II. When the measured throughput is much lower than the reference value, there are no possibilities for the customer to receive a higher throughput over time since the throughput proved to be very stable. In this case, a good option for the customer is to (i) terminate the VMs, (ii) create and deploy the VMs from scratch by operating the same decisions on the factor values; and (iii) measure the throughput again to check whether it is satisfactory or not. For instance, it is convenient for a customer deploying two medium-sized VMs in  $CFG_1$  on Singapore to terminate and recreate them when she measures only 200 Mbps of achievable throughput: by relocating the VMs, she has high chance to get 750 Mbps of throughput (see Fig. 4). Note that all these operations should be performed during the deployment phase, i.e. before the customer applications or services are made active.

## V. CONCLUSION

In this paper, by performing active measurements and by selecting the interesting scenarios as presented in Sec. III, we characterize the performance of the intra-cloud network in terms of the maximum achievable throughput in MS Azure. Thanks to 800 hours of experimentation, we discovered that – once a pair of VMs has been deployed – the maximum throughput between them is very stable over time. Moreover, experimental results showed that the factors under the direct control of a general customer (such as VM size, configuration, and region) may heavily influence the expected performance. Although general trends can be spotted (e.g. larger VMs perform better than smaller on average, or putting two VMs in the same affinity group allows to reach better performance), we also observed non-negligible variability in our results. For instance, in some scenarios, even when all the factors under the direct control of a customer are kept fixed, the maximum achievable throughput across different experiments (i.e., measured between different VM pairs) may significantly vary: the measured values polarized around a set of few very different throughput values. These findings allowed us to derive guidelines for both researchers aiming at characterizing the intra-cloud network performance in a public cloud environment and customers willing to obtain the best network performance.

TABLE III: Minimum throughput guaranteed by the provider [Mbps].

Region	TCP			UDP		
	medium	large	xlarge	medium	large	xlarge
California	186.6	374.6	929.5	171.2	350.8	658.0
Ireland	185.4	364.1	728.5	170.5	327.5	468.8
Sao Paulo	185.1	707.5	907.1	175.0	708.1	842.6
Singapore	186.0	718.1	935.0	174.9	519.1	614.9

## ACKNOWLEDGEMENTS

This work is partially funded by the MIUR projects: PLATINO ( $PON01_01007$ ), SMART HEALTH ( $PON04a2_C$ ) and SIRIO ( $PON01_02425$ ), and art. 11 DM 593/2000 for NM2 srl.

## REFERENCES

- [1] CloudHarmony. <https://cloudharmony.com/>.
- [2] Perfkit. <https://github.com/GoogleCloudPlatform/PerfKitBenchmarker/>.
- [3] Microsoft azure, Mar. 2015.
- [4] G. Aceto, A. Botta, W. de Donato, and A. Pescapé. Cloud monitoring: A survey. *Computer Networks*, 57(9):2093 – 2115, 2013.
- [5] H. Ballani, P. Costa, T. Karagiannis, and A. Rowstron. Towards predictable datacenter networks. In *Proceedings of the ACM SIGCOMM 2011 Conference*, SIGCOMM ’11, pages 242–253, New York, NY, USA, 2011. ACM.
- [6] A. Botta, D. Emma, A. Pescapé, and G. Ventre. Systematic performance modeling and characterization of heterogeneous ip networks. In *Parallel and Distributed Systems, 2005. Proceedings. 11th International Conference on*, volume 2, pages 120–124. IEEE, 2005.
- [7] A. Greenberg, J. R. Hamilton, N. Jain, S. Kandula, C. Kim, P. Lahiri, D. A. Maltz, P. Patel, and S. Sengupta. V12: A scalable and flexible data center network. *SIGCOMM Comput. Commun. Rev.*, 39(4):51–62, Aug. 2009.
- [8] Z. Hill, J. Li, M. Mao, A. Ruiz-Alvarez, and M. Humphrey. Early observations on the performance of windows azure. In *Proceedings of the 19th ACM International Symposium on High Performance Distributed Computing*, pages 367–376. ACM, 2010.
- [9] K. LaCurts, S. Deng, A. Goyal, and H. Balakrishnan. Choreo: Network-aware task placement for cloud applications. In *Proceedings of the 2013 Conference on Internet Measurement Conference*, IMC ’13, pages 191–204, New York, NY, USA, 2013. ACM.
- [10] L. Leong et al. Gartner, magic quadrant for cloud infrastructure as a service, 2014.
- [11] A. Li, X. Yang, S. Kandula, and M. Zhang. Cloudcmp: Comparing public cloud providers. In *Proceedings of the 10th ACM SIGCOMM Conference on Internet Measurement*, IMC ’10, pages 1–14, New York, NY, USA, 2010. ACM.
- [12] D. Mangot. Measuring ec2 system performance, 2009.
- [13] J. C. Mogul and L. Popa. What we talk about when we talk about cloud network performance. *SIGCOMM Comput. Commun. Rev.*, 42(5):44–48, Sept. 2012.
- [14] C. Raiciu, M. Ionescu, and D. Niculescu. Opening up black box networks with clouddtalk. In *Proc. 4th USENIX Conference on Hot Topics in Cloud Computing*, page 6, 2012.
- [15] J. Schad, J. Dittrich, and J.-A. Quiané-Ruiz. Runtime measurements in the cloud: Observing, analyzing, and reducing variance. *Proc. VLDB Endow.*, 3(1-2):460–471, Sept. 2010.
- [16] R. Tudoran, A. Costan, G. Antoniu, and L. Bougé. A performance evaluation of azure and nimbus clouds for scientific applications. In *Proceedings of the 2nd International Workshop on Cloud Computing Platforms*, page 4. ACM, 2012.
- [17] G. Wang and T. Ng. The impact of virtualization on network performance of amazon ec2 data center. In *INFOCOM, 2010 Proceedings IEEE*, pages 1–9, March 2010.