Vol.2011-ARC-195 No.19
Vol.2011-OS-117 No.19
2011/4/14

Internet Traffic Characteristics of Virtual Machine on Amazon EC2

CHUNGHAN LEE,†1 HIROTAKE ABE,†2 TOSHIO HIROTSU†3 and KYOJI UMEMURA†4

Cloud computing has attracted a great deal of attention and provides users to acquire computing resources from data centers. Platforms typically do not permit access to actual physical machines but rely on virtualization technologies. Amazon EC2 is a popular cloud service, and provides the platform as an instance to users. However, virtualization technologies can impair communication performance. We found 'oversize packet spacing', which can be caused by CPU scheduling latency, was a major cause of throughput instability on PlanetLab. The oversize packet spacings are anomalies on virtualized network environments. Moreover, large delay variations and unstable TCP/UDP throughput were reported on Amazon EC2. However, they used their monitoring program only and focused on the characteristics among the instances. We choose the instance as a sender node and a native node as a receiver node. To understand Internet traffic characteristics and to observe the anomalies, we simultaneously measured both network throughput and resource states by using monitoring program offered by Amazon. In the results, very few of packet spacings are the anomalies, and it is hard to determine a major cause of throughput instability. Moreover, CPU utilization is stable during the throughput measurements. Finally, we show that packet loss is occurred continuously although RTT is very similar to mean RTT using ping.

1. Introduction

Cloud services emerge as a new paradigm recently, and provide users to acquire computing resources from large scale data centers of service providers, such as Amazon, Google, and Microsoft. Moreover, cloud platforms typically do not permit access to actual physical machines but rely on virtualization technologies for low cost, security, and technical flexibility. Amazon EC2 is one of popular cloud services. It is a component of Amazon web services (AWS)[1], and its web service interface allows users to obtain *Amazon Machine Instance (AMI)* and to run applications on the AMIs. The AMI based Xen is provided to user as the platform. The instances e.g., micro, small, large, and so on, have different resources, such as CPU cores, memory, and storage. Amazon EC2 provides the ability to place the instances in multiple locations. The locations are composed of regions and availability zones. The zones are distinct locations that are engineered to be insulated from failures in other zones and provide inexpensive and low latency network connectivity to other zones in the same region. It is currently available in five regions: US East (Northern Virginia), US West (Northern California), EU (Ireland), and Asia Pacific (Singapore, Tokyo).

†1 Dept. of Electronic and Information Engineering, Toyohashi University of Technology
†2 Cybermedia center, Osaka University
†3 Faculty of Computer and Information Sciences, Hosei University
†4 Dept. of Computer Science and Engineering, Toyohashi University of Technology

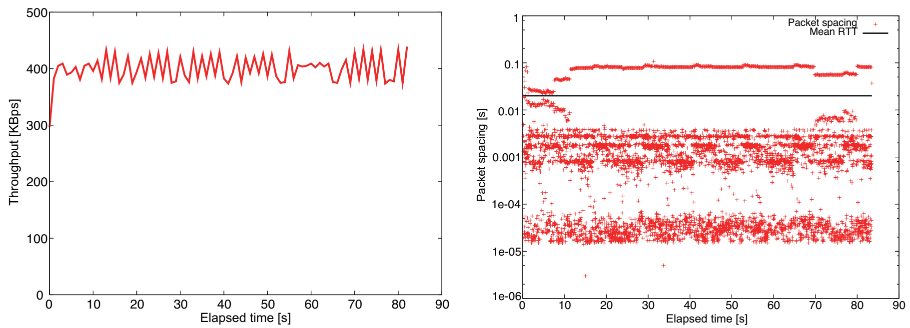

(a) Throughput measurement result (32 MB)　(b) Packet spacing

Fig. 1 Throughput measurement result and packet spacing on PlanetLab.

It is expected that virtualization technologies can affect computation and communication performance. In our previous work[2], we showed 'oversize packet spacing', which can be caused by CPU scheduling latency, was a major cause of throughput instability on PlanetLab[3)4)] even when no significant changes occur in well-known network metrics. The oversize packet spacings are anomalies on virtualized network environments. Fig. 1 shows a throughput measurement result and packet spacing on PlanetLab. In this case, the mean value of RTT was 0.022 s, and this value was similar to RTT using ping. The throughput measurement result was 393 KBps despite the stable RTT. The mean value of packet spacing was 0.008177 s, and the anomalies were the major cause of throughput decrease. Moreover, Wang et al.[5] demonstrated Internet traffic characteristics, such as abnormal large delay variations and unstable TCP/UDP throughput, caused by the impact of virtualization on Amazon EC2. However, they used their monitoring program only. Moreover, their approach would be inappropriate because the under layer of virtual machine is typically block-box. Then, they focused on the characteristics among Amazon EC2 instances only.

The aims of our study are to observe the anomalies and to understand Internet traffic characteristics on Amazon EC2. We chose the micro instance as a sender node and a native node on Toyohashi University of Technology (TUT) as a receiver node to eliminate the impact of virtualization. We simultaneously measured both network throughput and resource states on a node by using Amazon CloudWatch[6] provided by AWS. Then, we investigated the anomalies and well-known network metrics, such as RTT, notify window, and packet loss by using packet-level analysis to understand the Internet traffic characteristics on Amazon EC2. Finally, we presented CPU utilization using Amazon CloudWatch. In this paper, we present the Internet traffic characteristics on Amazon EC2. The contributions of this work are:

- In analysis results, very few of packet spacings are the anomalies, and it is hard to determine a major cause of throughput instability. Our results are different from the results shown by Wang et al.[5].
- Previous works[5][7][8] did not use Amazon CloudWatch but we use it to observe correct resource state. We show that the CPU utilization is stable during the throughput measurements.
- We present that packet loss is occurred continuously during the throughput measurements although RTT is very similar to mean RTT using ping. It is the major cause of throughput instability.

The rest of this paper is organized as follows. First, we describe related work in Section 2. Second, our measurement methodology is explained in Section 3. Next, we present throughput measurement results and discuss analysis results in Section 4. Finally, we conclude the paper with a summary of the main points in Section 5.

2. Related work

2.1 Fluctuation in packet spacing

Lee et al.[2] investigated anomalies on the virtualized network testbed; they found that CPU consumption on a per-slice basis is an important parameter for anomaly estimation. However, measuring CPU availability consumes idle CPUs, decreasing the number of CPUs that can be allocated to users on the node who require CPU resources.

Peterson et al.[9] deployed a packet forwarding overlay between Seattle and Washington, D.C. on the virtualized testbed and compared round trip time (RTT) on both virtualized and non-virtualized hosts between the nodes. The RTT on non-virtualized hosts was nearly constant, while that on the overlay varied widely. This RTT fluctuation was caused by CPU scheduling latency. Although the scheduling latency at a node is a serious problem for network measurements, the relationship between packet spacing fluctuation and scheduling latency was not considered.

Spring et al.[10] showed that the load on a node prevents accurate latency measurement and precise spacing for packet trains. They ran traceroute and tcpdump in parallel to acquire timestamps for the application and kernel levels and showed the differences between these timestamps when sending probes and receiving responses. Moreover, they transmitted packet trains to determine how the CPU load impaired precisely-spaced packets. However, they showed no clear conditions for the type of load.

2.2 Performance analysis of Amazon EC2

Iosup et al.[7] investigated EC2 performance using micro-benchmarks, kernels, and e-Science workloads. They evaluated the instances as a scientific computing platform only. Dejun et al.[8] analyzed EC2 performance for service-oriented applications. They developed a CPU-intensive web application and database read and write-intensive applications to simulate different types of workload pattern. Despite the performance of the instances to be stable, the multiple instances of the same type showed heterogeneous performances. The performance analysis was focused, and they did not discuss traffic characteristics of instances.

To our best knowledge, Wang et al.[5] focused on networking performance among the instances, and demonstrated very different characteristics, such as abnormal large delay variations and unstable TCP/UDP throughput, caused by the impact of virtualization on node. However, they focused on the regions on Amazon EC2 data center only, and did not investigate networking performance between the region and different types of network. In these works, they did not use monitoring programs offered by AWS. Because the under layer on the virtual machine is block-box typically, their approach would be inappropriate for the performance analysis.

3. Measurement methodology

In this section, we introduce throughput measurement method with resource monitoring to observe the anomalies on Amazon EC2.

3.1 Throughput measurement

We choose the micro instance as a sender node to observe the anomalies and to identify the impact of virtualization. The five micro instances are used for the experiments. We choose a native node on TUT as a receiver node to eliminate the impact of virtualization. The native node has an Intel Xeon CPU (4 CPU cores) with a 4 GB RAM and a 1 Gbps network card and runs on Ubuntu 10.04. The micro instance is with 613 MB memory, 1 EC2 compute unit (Intel Xeon), and 10 GB on EBS storage. However, we cannot have information about the physical node. Amazon EC2 uses Xen, while PlanetLab uses V-Server. It is with operating system-level virtualization while the micro instance is with para-virtualization. The instances are in US East - N. Virginia (us-east- 1a). Min, max, and mean RTT between the region on Amazon EC2 and our node are 209.63, 369.782, and 213.126 ms. Our approach to choose the micro instance is appropriate because the experiments use mainly network resources.

A prediction method using pairs of different-sized connections has been previously proposed[11]. This method, which we call *connection pair*, uses a small-sized probe transfer to predict the throughput of a large-sized data transfer. Wolski et al.[11][12] used different-sized pairs of connections and empirically established the basic probe size as 64 KB for the Network Weather Service (NWS). However, they did not consider the appropriate size for probe transfer and network virtualization. We generated various sizes of connection pairs to observe anomalies. In our previous work[13], we found the large-sized one is better than the basic size. In order to identify the validity of the large-sized probe on Amazon EC2, we used both small and large-sized probe for the connection pair. The sizes for the connection pairs are shown in Table 1.

3.2 Resource monitoring

If we observe resource state using sysstat, the information of resource metrics will be incorrect because the under layer is normally black-box. We used Amazon CloudWatch offered by AWS to observe the resource state during the throughput measurement. It correctly provides resource metrics such as CPU utilization, disk I/O, and network traffic per the instance. However, it does not provide memory utilization, the number of the instances on physical node, and so on. Moreover, the minimum granularity of the metrics is limited as 1 minute. The measurement methodology is described in Fig. 2.

Table 1 Probe and data size combinations for connection pairs.

Index	Probe size [KB]	Data size [MB]	Number of connection pairs
I_0	16	8	3
I_1	16	16	3
I_2	32	16	3
I_3	64	16	3
I_4	64	32	3
I_5	512	8	3
I_6	512	16	3
I_7	1024	16	3
I_8	1024	16	3
I_9	2048	32	3

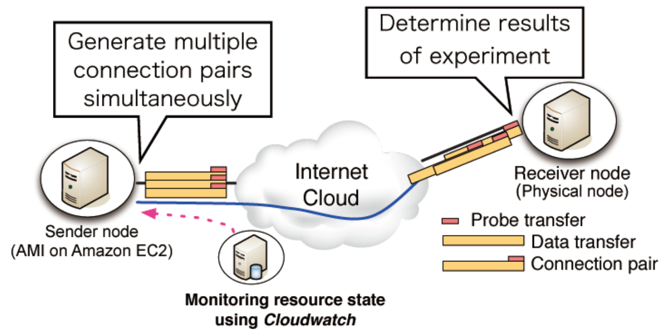

Fig. 2 Connection pair measurement method with resource monitoring.

4. Analysis results

In this section, we describe throughput measurement results, packet-level analysis, and resource state. Our objectives in packet-level analysis are to observe the anomalies and to understand the Internet traffic characteristics. We show CPU utilization to investigate the impact of virtualization to communication performance.

4.1 Throughput measurement results

Although we measured the throughput for the connection pair over 3 days, some of it was not gathered by libpcap error. A dataset consisted of approximately 1,000 connection pairs over 18 hours. Throughput measurement results at indices I_0 and I_5 are shown in Fig. 3. Low network throughput was observed at all the throughput measurement results. While the results where the probe size was 16 KB at index I_0 were clustered at the certain areas, the results where the probe size was 512 KB at index I_5 were widely spread. The results from index I_1 to I_4 were similar to the results at index I_0, and the others were similar to the results at index I_5. The large-sized probe would be appropriate for the connection pair.

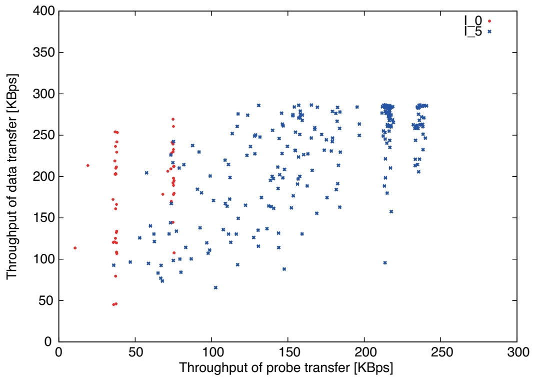

Fig. 3 Connection pair throughput at indices I_0 and I_5.

4.2 Network metrics

First, we analyze the well-known network metrics to understand the Internet traffic characteristics on Amazon EC2. Moreover, we observe packet spacings for reproduction of the anomalies, and compare them to the results with the anomalies on PlanetLab.

4.2.1 RTT and packet loss

The RTT normally fluctuates and network throughput decreases as a result of network congestion on the end-to-end path. Thus, a lot of packets would be dropped in the congestion. These well-known network metrics are important to understand Internet traffic characteristics. Fig. 4 shows RTT of connection pair at index I_8. The throughput measurement results of probe and data were 121.4 KBps and 98.5 KBps respectively. Although the throughput decreases were occurred, the most values of it were stable and the mean values of RTT of them were 0.216 s and 0.220 s. The values are similar to the mean RTT (0.213 s) using ping. Moreover, packet loss rate of the data transfer was 1.304 %. It would be small but the packet loss was occurred continuously during the throughput measurements. Table 2 shows mean values of the network metrics for all the indices. There were no significant changes in the mean RTT, so the network condition did not exhibit any abrupt changes. The mean values at all the indices were higher than that using ping because the sender node had to wait for a timeout. Although the mean values of RTT were not significant changes, there was packet loss continuously. The loss rate at all the indices was less than 10 %, however the packet loss was similar to the above results.

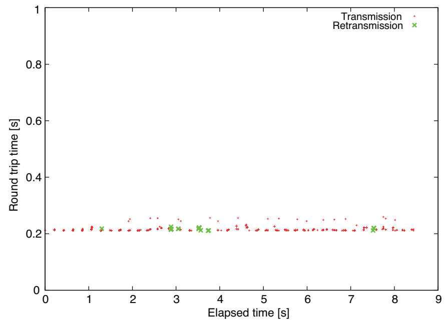

(a) Probe transfer (1 MB)

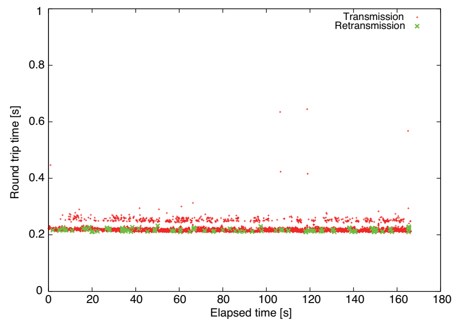

(b) Data transfer (16 MB)

Fig. 4 RTT of connection pair at index I_8.

4.2.2 Packet spacing

Normally, the packet spacing is a very short period over non-virtualized environments. To investigate packet spacing over a non-virtualized environment, we considered that a local environment consists of a router and two native nodes. Each native node has an Intel Pentium 4 processor with a 1 GB RAM and a 100 Mbps network card and runs on Ubuntu 9.04. The router was connected between these nodes via a 100 Mbps link. The packet spacing where data size was 32 MB is shown in Fig. 5. Mean packet spacing was 0.000010 s. There was no abrupt fluctuation in the packet spacing and these values did not affect the throughput measurement. In order to observe the anomalies on Amazon

Table 2 Mean values of well-known network metrics

Index	Mean RTT [s] (Probe)	Mean RTT [s] (Data)	Mean loss rate [%] (Probe)	Mean loss rate [%] (Data)
I_0	0.215	0.218	2.025	1.675
I_1	0.215	0.217	1.711	1.436
I_2	0.215	0.217	1.86	1.462
I_3	0.217	0.218	0.056	1.192
I_4	0.217	0.217	5.397	1.234
I_5	0.216	0.215	1.122	0.581
I_6	0.218	0.217	1.646	0.9
I_7	0.218	0.218	1.775	0.794
I_8	0.216	0.215	0.645	0.42
I_9	0.217	0.216	1.167	0.709

EC2, we investigate packet spacings for the connection pair, and compare them to the results, where the probe size was 64 KB and the data size was 32 MB, that experiences throughput instability caused by the anomalies between two sites on PlanetLab. Fig. 6 shows mean packet spacing and network throughput for the connection pair at all the indices on Amazon EC2 and PlanetLab. The mean values on PlanetLab were larger than the others, and their patterns were very similar to Fig. 1. Most values of mean packet spacing on Amazon EC2 were stable like the result on local environment. However, some of it was increased, and similar to the results on PlanetLab. To identify the reproduction of the anomalies, we present that the case where the mean packet spacings of probe and data were 0.003620 s and 0.000415 s at index I_8. Fig. 7 shows packet spacing of outgoing transfers at the case. The throughput measurement results of probe and data was 193.9 KBps and 251 KBps respectively. Due to the values where packet spacings of probe and data were 0.469552 s and 0.475770 s, the mean packet spacing of outgoing transfers was increased. In the other results, very few of packet spacings were the anomalies, and it is hard to determine the major cause of throughput instability. Our results are different from the results shown by Wang et al.[5]. Typically, the fluctuation in packet spacing can be caused by flow control in TCP mechanism. Packet will be not sent immediately after ACK packet. We observe notify window to identify the effect of flow control. Fig. 8 shows packet spacing and notify window of outgoing transfers at the above case. The window is saturated to up maximum size. There is no impact of flow control. We will clarify a cause of very few of the anomalies as a future work.

To summarize the Internet traffic characteristics on Amazon EC2, very few of packet spacings were the anomalies and they were not the major cause of throughput instability. Packets were continuously dropped during the throughput measurements although RTT was very similar to mean RTT using ping. It was the major cause of throughput instability. In the data centers, tens of nodes per rack are connected via top of rack switches that connect to high degree aggregation switches[14]. The previous works[14][15] presented impairments among flows in the data center. Our implications of throughput instability are that variant types of many flows co-exist in the data center and available buffer size in switches would be insufficient. These would affect the measurement results.

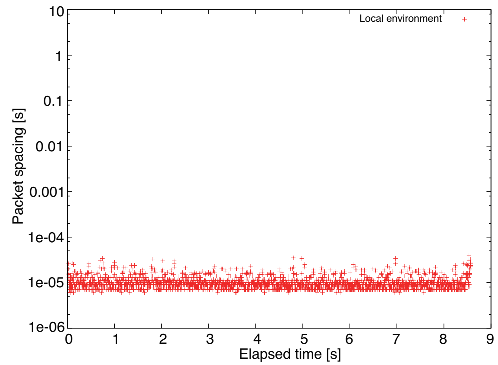

Fig. 5 Packet spacing on local environment.

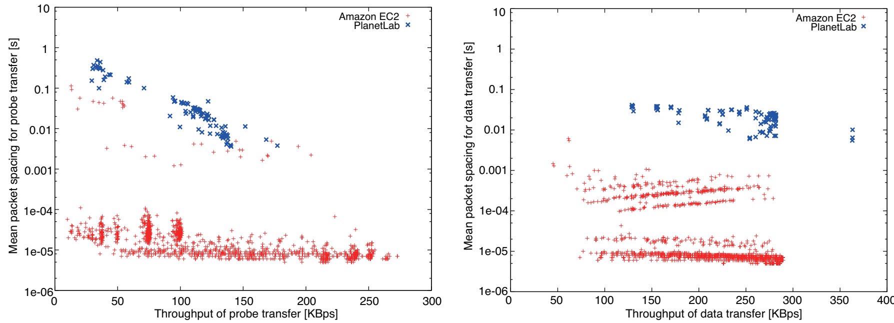

(a) Probe transfer

(b) Data transfer

Fig. 6 Mean packet spacing and throughput at all indices on Amazon EC2 and PlanetLab.

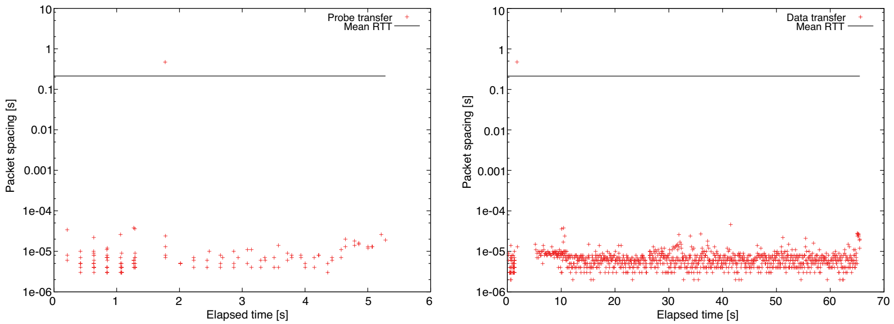

(a) Probe transfer (1 MB) (b) Data transfer (16 MB)

Fig. 7 Packet spacing of outgoing transfers at index I_8.

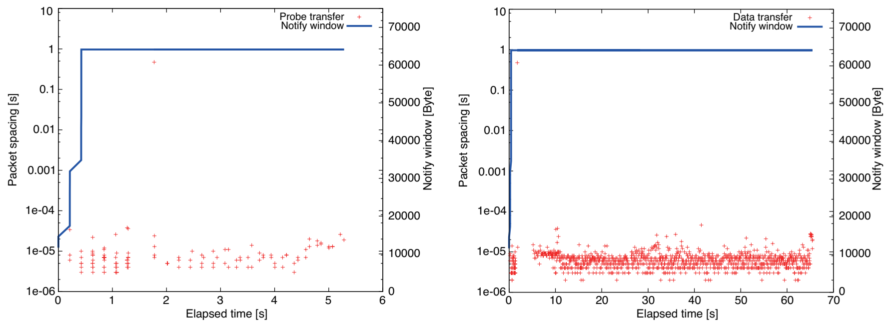

(a) Probe transfer (1 MB) (b) Data transfer (16 MB)

Fig. 8 Packet spacing and notify window of outgoing transfers at index I_8.

4.3 CPU utilization

Because the packet spacings were stable, the CPU utilization will be low. Fig. 9 shows CPU utilization at index I_0 and I_5. Most of values are less than 20 %. In the experiments, the CPUs were not busy, and the CPUs were fully allocated to the instances. The others were similar to the above cases. In the data centers, the CPU utilization would be low due to resource policy and the structure of data center. Again, very few of anomalies were occurred on Amazon EC2, and they were not the major cause of throughput decreases.

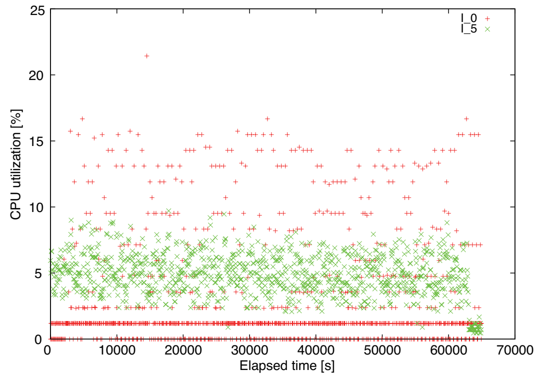

Fig. 9 CPU utilization on Amazon EC2 (index I_0 and I_5)

5. Conclusion

We measured throughput for the connection pair between the micro instance on Amazon EC2 and the native node on TUT, and observed correct resource state by using Amazon CloudWatch. Analysis results showed that very few of packet spacings are the anomalies, and it is hard to determine the major cause of throughput instability. Our results are different from the results shown by Wang et al.[5]. Moreover, packet loss was occurred continuously during the throughput measurements although RTT was very similar to mean RTT using ping. It was the major cause of throughput instability. In the data center, variant types of many flows would co-exist, and they would affect their communication performance. The CPU utilization was less than 20 % during the measurements. The CPUs were not busy, and the CPUs were fully allocated to the instances. Thus, there was no CPU scheduling latency. In the data center, the CPU utilization would be low due to resource policy and the structure of data center.

In future work, we will investigate causes of packet loss and very few of the anomalies clearly, and gather datasets on the different instances in the other regions for the better understand of Internet traffic characteristics on Amazon EC2, and extend our work with additional analysis of the other services on AWS.

Acknowledgments This work was supported in part by JSPS KAKENHI (22500072 and 22700029) and Global COE Program "Frontiers of Intelligent Sensing" from the Ministry of Education, Culture, Sports, Science and Technology, Japan.

References

1) AmazonWebServices(AWS): http://aws.amazon.com/.
2) Lee, C., Abe, H., Hirotsu, T. and kyoji Umemura: Analysis of Anomalies on a Virtualized Network Testbed, *CIT '10: Proceedings of the 10th IEEE International Conference on Computer and Information Technology*, Washington, DC, USA, IEEE Computer Society, pp. 297–304 (2010).
3) PlanetLab: https://www.planet-lab.org/.
4) Bavier, A., Bowman, M., Chun, B., Culler, D., Karlin, S., Muir, S., Peterson, L., Roscoe, T., Spalink, T. and Wawrzoniak, M.: Operating system support for planetary-scale network services, *NSDI'04: Proceedings of the 1st conference on Symposium on Networked Systems Design and Implementation*, Berkeley, CA, USA, USENIX Association, pp.19–19 (2004).
5) Wang, G. and Ng, T. S.E.: The impact of virtualization on network performance of amazon EC2 data center, *INFOCOM'10: Proceedings of the 29th conference on Information communications*, Piscataway, NJ, USA, IEEE Press, pp.1163–1171 (2010).
6) AmazonCloudWatch: http://aws.amazon.com/cloudwatch/.
7) Iosup, R., Ostermann, S., Yigitbasi, N., Prodan, R., Fahringer, T. and Epema, D.: An early performance analysis of cloud computing services for scientific computing, *TU Delft, Tech. Rep., Dec 2008, [Online] Available* (2008).
8) Dejun, J., Pierre, G. and Chi, C.-H.: EC2 performance analysis for resource provisioning of service-oriented applications, *Proceedings of the 2009 international conference on Service-oriented computing*, ICSOC/ServiceWave'09, Berlin, Heidelberg, Springer-Verlag, pp.197–207 (2009).
9) Peterson, L., Bavier, A., Fiuczynski, M.E. and Muir, S.: Experiences building PlanetLab, *OSDI '06: Proceedings of the 7th symposium on Operating systems design and implementation*, Berkeley, CA, USA, USENIX Association, pp.351–366 (2006).
10) Spring, N., Peterson, L., Bavier, A. and Pai, V.: Using PlanetLab for network research: myths, realities, and best practices, *SIGOPS Oper. Syst. Rev.*, Vol.40, No.1, pp.17–24 (2006).
11) Wolski, R., Spring, N. and Hayes, J.: The Network Weather Service: A Distributed Resource Performance Forecasting Service for Metacomputing, *Journal of Future Generation Computing Systems*, Vol.15, pp.757–768 (1999).
12) Swany, M. and Wolski, R.: Multivariate resource performance forecasting in the network weather service, *Supercomputing '02: Proceedings of the 2002 ACM/IEEE conference on Supercomputing*, Los Alamitos, CA, USA, IEEE Computer Society Press, pp.1–10 (2002).
13) Lee, C., Abe, H., Hirotsu, T. and Umemura, K.: Evaluation of the Throughput Measurement Method using Connection Pairs, *WIT '09: In Proceedings of The Tenth Workshop on Internet Technology*, Japan Society for Software Science and Technology (JSSST) (2009).
14) Kandula, S., Sengupta, S., Greenberg, A., Patel, P. and Chaiken, R.: The Nature of Datacenter Traffic: Measurements and Analysis, *Proceedings of the 9th ACM SIGCOMM conference on Internet measurement conference*, IMC '09, New York, NY, USA, ACM, pp. 202–208 (2009).
15) Alizadeh, M., Greenberg, A., Maltz, D.A., Padhye, J., Patel, P., Prabhakar, B., Sengupta, S. and Sridharan, M.: Data center TCP (DCTCP), *Proceedings of the ACM SIGCOMM 2010 conference on SIGCOMM*, SIGCOMM '10, New York, NY, USA, ACM, pp.63–74 (2010).