International Journal of Advance Research in Science and Engineering
Vol. No.9, Issue No. 03, March 2020
www.ijarse.com
IJARSE
ISSN 2319 - 8354

ENERGY AND BANDWIDTH CONSTRAINED MULTICAST DATA FORWARDING SCHEME FOR CLOUD DATA CENTERS

S. Indhirani [1] , E. Mythili [1] , S. Umajagadeeswari [1] ,

Dr. S. Radha [2]

[1]B.E-(Final) Computer Science and Engineering , Sengunthar Engineering College (Autonomous), Tiruchengode-638205

[2]Assistant professor – Computer Science and Engineering , Sengunthar Engineering College (Autonomous), Tiruchengode-638205

ABSTRACT

Cloud computing technology shares the resources through the internet. Processors, computing architecture and storage pools are shared as software and infrastructure service. A remote data center provides data and software's using the network bandwidth. Amazon Simple Storage Service (S3) and Amazon Elastic Compute Cloud (EC2) are the popular cloud service providers. The Cloud MP cast is an overlay system loaded in the cloud data centers to support bulk data transfer process. The routing planner component is employed to estimate the most cost efficient plan. The overlay routing topology is configured with the discovered transmission plan. The overlay distribution tree and traffic levels are analyzed to select the data transfer path. The bandwidth cost estimation is carried out with geographical location and pricing policy information. The cloud data center bulk transfer process is handled with the Enhanced Cloud MP cast scheme. Data splitting approach is constructed to support multiple path based data forwarding process.

1.INTRODUCTION

The past few years have witnessed an explosion in the popularity of cloud computing services. A key advantage of cloud computing is the ability to geo-distribute data over multiple data-centers, both to increase availability and reliability as well as provide lower user latency—low latencies are critical to business revenues, for example, Amazon estimated every 100 ms of latency costs 1 percent in sales .

There are numerous examples of research efforts and successful commercial applications that distribute content across the globe to achieve good performance. Netflix is a prominent example of an application that disseminates its contents over a Cloud Service Provider (CSP) infrastructure to offer on-demand media services. In addition to significant reductions in latencies, geo-replication has become important to ensure high availability despite failures—e.g., for disaster recovery purposes. Further, other applications including software distribution, virtual machines cloning, distributed databases, and data warehousing may require geo replication.

2. RELATED WORK

The Internet community has proposed several mechanisms to cut operational expenditures of online service companies many of which may also apply to CSPs. In contrast to these works, our focus is on reducing costs of customers of CSPs.

Prior work has pointed out that the cost of electricity varies with time of day and across locations, and explored moving computation to data-centers that are cheaper at a given time. Similarly, some studies focus on how ISP pro- videos charge wholesale clients for bandwidth . Specifically, researchers have proposed to carry out data backups and other bulk data transfer tasks during off-peak hours when usage is lower. This approach was further improved by splitting backup activity into chunks and leveraging software defined networking (SDN).

Several works have designed bulk data transfer schemes assuming ISPs charge their clients for bandwidth following a 95th percentile usage model. That is, given a time series that represents the bandwidth used by clients, the bill corresponds to the 95th percentile sample regardless of traffic at lower percentiles. These works propose to transfer data when usage is below the 95th percentile as no additional charges would be incurred. Moreover, the authors in showed that not only bulk transfers, but also multimedia content with different quality requirements can be forwarded over an overlay topology so as to minimize the probability of exceeding pre-estimated percentiles.

3. ARCHITECTURE DIAGRAM

CSPs charge their customers for data transfers both to other data-center and the Internet by

IJARSE
ISSN 2319 - 8354

counting the number of GBs each customer transmits per data-center during a month and multiplying by its cost. The cost rates are published in the terms and conditions sheets of each CSP. Table 1 shows Azure and EC2 pricing in $/GB on a monthly basis as of 2014.

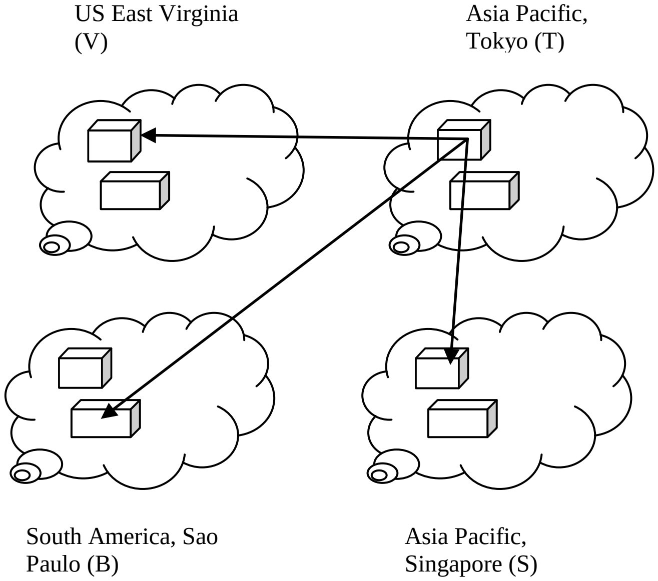

Figure (3.1)

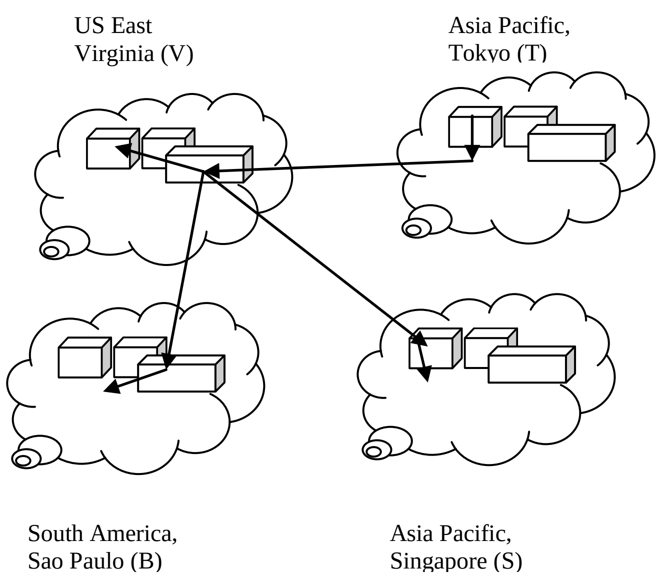

Figure (3.2)

International Journal of Advance Research in Science and Engineering
Vol. No.9, Issue No. 03, March 2020
www.ijarse.com
IJARSE
ISSN 2319 - 8354

4. WORKING PRINCIPLE

CloudMPcast is an overlay system that executes in each data-center of an application deployment, where each application/tenant instantiates its own deployment of CloudMPCast. The key component is a routing planner, which computes the most cost-efficient transmission plan (possibly the trivial solution) while observing the transfer time, for any given application request. Once a transmission plan is generated, CloudMPcast contacts the other instances of CloudMPcast in the data-centers involved in the solution to configure the overlay routing topology.

Each CloudMPcast node monitors the TCP throughput between itself and other data-centers in the deployment. CloudMPcast periodically (every 5 minutes) updates it's inter data-center bandwidth matrix using well-known techniques to inform computations for future transfers. Our measurements on real cloud deployments indicate that inter data-center bandwidth is relatively stable over time-scales longer than individual bulk transfers. However, if significant throughput changes are detected which indicates that an overlay distribution tree previously computed for an on-going transfer is no longer performing well from a transfer time perspective, CloudMPcast aborts the rest of the transfer. The destination data-centers directly contact the source data-center to obtain data not yet received. To reduce system complexity, CloudMPcast does not dynamically create alternate overlay distribution plans for an ongoing bulk transfer.

Since CloudMPcast may forward traffic to hub data-centers, a potential concern is whether this may induce network congestion at those data-centers and wide-area network in general, resulting in a reduction in the real throughput. Since CloudMPcast only manages requests of a given application/tenant, the total number of flows managed by CloudMPcast is small compared to the total number of flows competing for bandwidth in the network core.

5. RESULT AND DISCUSSION

The Cloud Service Provider (CSP) and Cloud Data Centers (CDC) are deployed with the support of the cloud infrastructure. The data centers, replicas and regions detail are collected from the user. The cloud infrastructure is deployed with specified data centers and replicas. The data centers and replicas are assigned with relevant regions to ensure better data delivery to the nodes

The historical data are maintained under the traced traffic log data files. The historical data values are analyzed to estimate the traffic patterns. The traffic patterns are derived for each time slots. The request frequency is measured in each time slot to indicate the traffic levels

The traffic prediction process is initiated for all regions. The region level traffic and total traffic levels are used in the forecasting process. The regions are categorized with normal and high traffic levels. The bulk data transfer operations are initiated with reference to the traffic conditions

The network bandwidth cost is managed by the Internet Service Provider (ISP).The static pricing policy and dynamic pricing policy models are used by the Internet Service Providers (ISP).The dynamic pricing policy includes the homogeneous, heterogeneous and discount based pricing mechanisms. The homogeneous cost model charges the same cost level for all conditions

The bandwidth scheduling is applied with the current request traffic levels. The data transmission and pause time intervals are allotted in the bandwidth scheduling process. The CloudMPcast and Enhanced CloudMPcast schemes uses separate bandwidth scheduling mechanisms. The data splitting process is initiated to identify the multiple path for the data transmission process

International Journal of Advance Research in Science and Engineering
Vol. No.9, Issue No. 03, March 2020
www.ijarse.com
IJARSE
ISSN 2319 - 8354

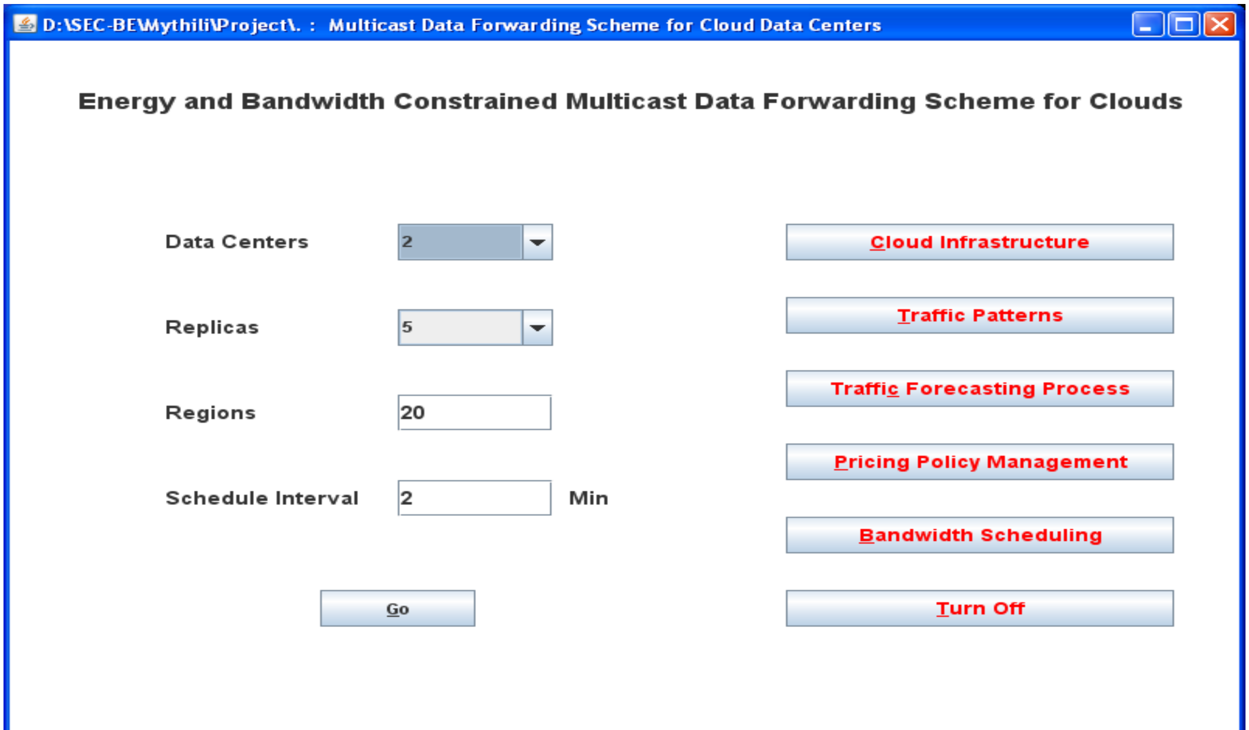

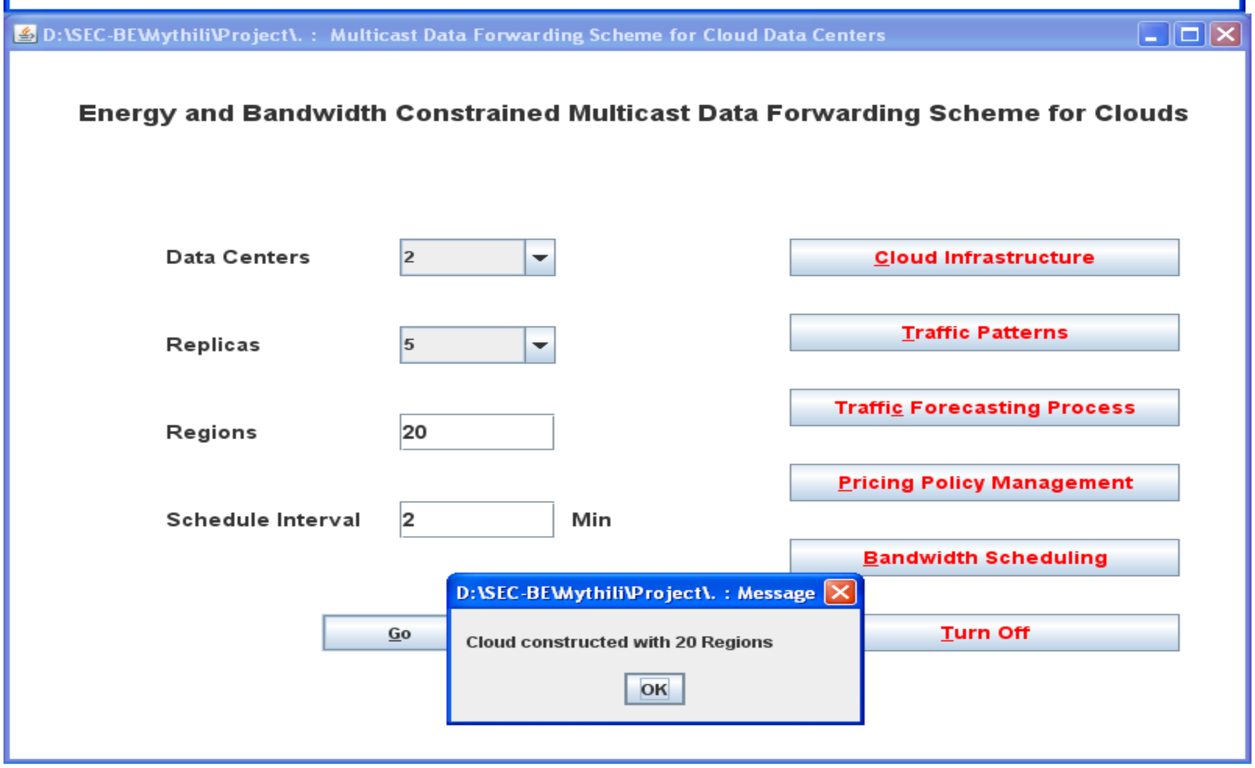

International Journal of Advance Research in Science and Engineering
Vol. No.9, Issue No. 03, March 2020
www.ijarse.com
IJARSE
ISSN 2319 - 8354

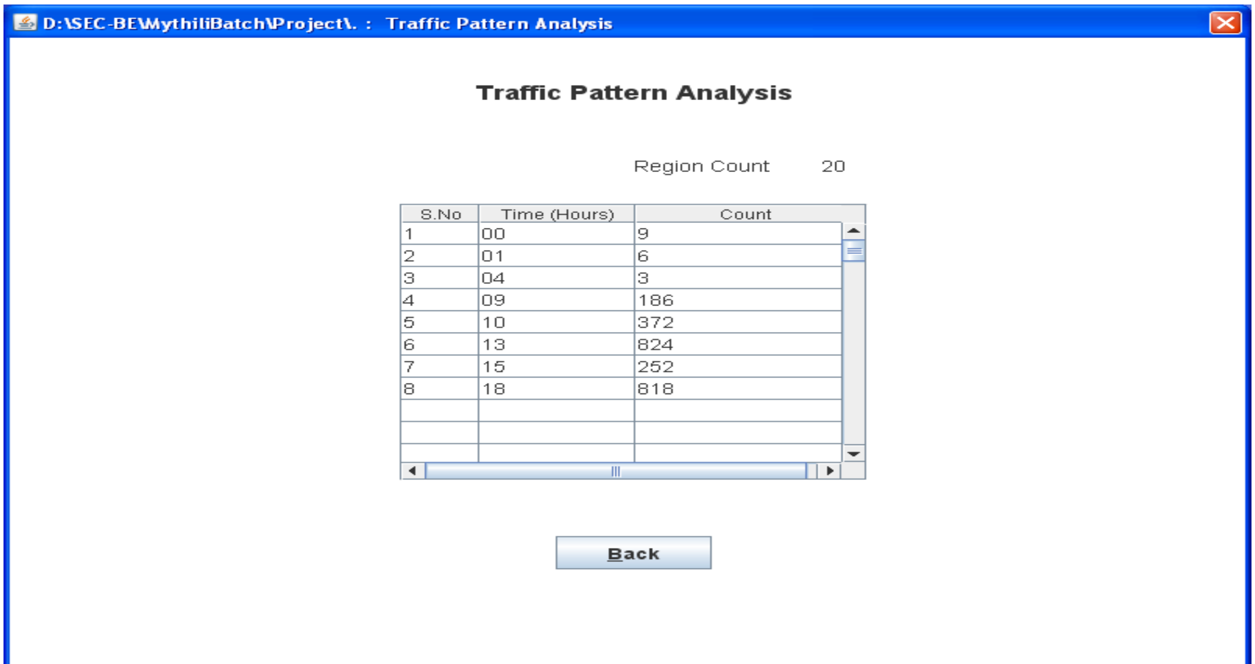

S.No	Time (Hours)	Count
1	00	9
2	01	6
3	04	3
4	09	186
5	10	372
6	13	824
7	15	252
8	18	818

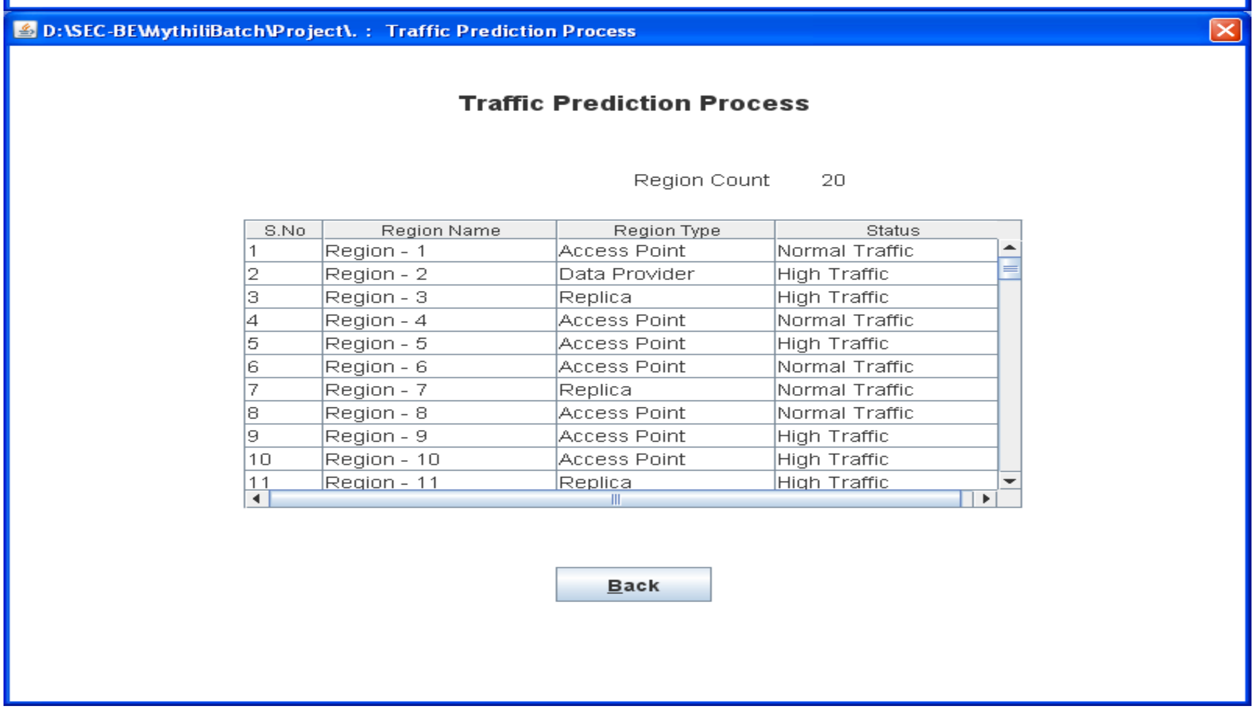

S.No	Region Name	Region Type	Status
1	Region - 1	Access Point	Normal Traffic
2	Region - 2	Data Provider	High Traffic
3	Region - 3	Replica	High Traffic
4	Region - 4	Access Point	Normal Traffic
5	Region - 5	Access Point	High Traffic
6	Region - 6	Access Point	Normal Traffic
7	Region - 7	Replica	Normal Traffic
8	Region - 8	Access Point	Normal Traffic
9	Region - 9	Access Point	High Traffic
10	Region - 10	Access Point	High Traffic
11	Region - 11	Replica	High Traffic

International Journal of Advance Research in Science and Engineering
Vol. No.9, Issue No. 03, March 2020
www.ijarse.com
IJARSE
ISSN 2319 - 8354

6. CONCLUSION

We have presented CloudMPcast, a systematic approach to constructing multi data-center distribution trees for bulk transfers. CloudMPcast optimizes dollar costs of distribution taking public cloud charging models into account, while ensuring end-to-end data transfer times are not affected.

Extensive evaluations of CloudMPcast leveraging an extensive set of inter-data-center bandwidth and latency measurements from both Azure and EC2 have shown significant benefits. Cost savings range from 10 to 60 percent across a wide variety of scenarios, which translates to mil- lions of dollars a year. Further, the results also point to the critical importance of exploiting both volume discounts and pricing heterogeneity across a variety of settings. This ensures savings can be achieved even when only one class of discount is applicable—e.g., rack space only offers volume discounts. When both discounts are applicable, considering them together provides even better results

REFERENCES

[1] High Scalability. Latency is everywhere and it costs you Sales- how to crush it [Online]. Available: http://highscalability.com/ latency-everywhere-and-it-costs-you-sales-how-crush-it/, 2009.

[2] B. Urgaonkar, G. Pacifici, P. Shenoy, M. Spreitzer, and A. Tantawi, "An analytical model for Multi-tier Internet services and its applications," in *Proc. ACM SIGMETRICS Int. Conf. Meas. Model. Comput. Syst.*, 2005, pp. 291–302.

[3] Z. Wu and H. V. Madhyastha, "Understanding the latency bene- fits of Multi-cloud web service deployments," *SIGCOMM Comput. Commun. Rev.*, vol. 43, no. 2, pp. 13–20, 2013.

[4] P. Shankaranarayanan, A. Sivakumar, S. Rao, and M. Tawarma- lani, "Performance sensitive replication in geo-distributed cloud datastores," in *Proc. IEEE/IFIP Int. Conf. Dependable Syst. Netw.*, 2014, pp. 240–251.

[5] Z. Wu, M. Butkiewicz, D. Perkins, E. Katz-Bassett, and H. V. Madhyastha, "SPANStore: Cost-effective geo-replicated storage spanning multiple cloud services," in *Proc. ACM Symp. Oper. Syst. Principles*, 2013, pp. 292–308.

[6] Forrester Research. The future of data center wide-area network- ing [Online]. Available: http://www.forrester.com, 2010.

[7] J. Hamilton's blog. Inter-datacenter replication & Geo-redundancy [Online]. Available: http://perspectives.mvdirona.com, 2010.

[8] Y. Chen, S. Jain, V. K. Adhikari, Z. li Zhang, and K. Xu, "A first look at Inter-data center traffic characteristics via Yahoo! data- sets," in *Proc. IEEE INFOCOM*, 2011, pp. 1620–1628.

[9] E. Zohar, I. Cidon, and O. Mokryn, "The power of prediction: Cloud bandwidth and cost reduction," in *Proc. ACM SIGCOMM*, 2011, pp. 86–97.

[10] N. Laoutaris, M. Sirivianos, X. Yang, and P. Rodriguez, "Inter- datacenter bulk transfers with Net stitcher," in *Proc. ACM SIGCOMM*, 2011, pp. 74–85.

[11] T. Nandagopal and K. P. N. Puttaswamy, "Lowering inter-data- center bandwidth costs via bulk data scheduling," in *Proc. Int. Symp. Cluster, Cloud Grid Comput.*, 2012, pp. 244–251.

[12] Y. Feng, B. Li, and B. Li, "Jet way: Minimizing costs on inter-data- center video traffic," in *Proc. ACM Multimedia Conf.*, 2012, pp. 259–268.

[13] K. He, A. Fisher, L. Wang, A. Gember, A. Akella, and T. Ristenpart, "Next stop, the cloud: Understanding modern web service deployment in EC2 and Azure," in *Proc. ACM Internet Meas. Conf.*, 2013, pp. 177–190.

[14] A. Qureshi, R. Weber, H. Balakrishnan, J. Guttag, and B. Maggs, "Cutting the electric bill for Internet-scale systems," in *Proc. ACM SIGCOMM*, 2009, pp. 123–134.

[15] L. Gyarmati, R. Stanojevic, M. Sirivianos, and N. Laoutaris, "Sharing the cost of backbone networks: Cui bono?" in *Proc. ACM Internet Meas. Conf.*, 2012, pp. 509–522.

[16] Y. Wu, Z. Zhang, C. Wu, C. Guo, Z. Li, and F. Lau, "Orchestrating bulk data transfers across geo-distributed datacenters," *IEEE Trans. Cloud Comput.*, To appear, doi: 10.1109/TCC.2015.2389842, 2015.

[17] Z. Ou, H. Zhuang, A. Lukyanenko, J. Nurminen, P. Hui, V. Mazalov, and A. Yla-Jaaski, "Is the same instance type created equal? exploiting heterogeneity of public clouds," *IEEE Trans. Cloud Computing.*, vol. 1, no. 2, pp. 201–214, Jul.–Dec. 2013.

[18] G. Wang and T. S. Eugene Ng, “The impact of virtualization on network performance of Amazon EC2 data center,” in *Proc. IEEE INFOCOM*, 2010, pp. 1–9.

[19] C. Guo, G. Lu, H. J. Wang, S. Yang, C. Kong, P. Sun, W. Wu, and Y. Zhang, “Second net: A data center network virtualization architecture with bandwidth guarantees,” in *Proc. ACM Co-NEXT Conf.*, 2010, p. 15.

[20] Y. Xu, Z. Musgrave, B. Noble, and M. Bailey, “Bobtail: Avoiding long tails in the cloud,” in *Proc. USENIX Conf. Net. Syst. Design Implementation*, 2013, pp. 329–342.