

Datacenters as Computers: Google Engineering & Database Research Perspectives

Shivakumar Venkataraman
Google Inc.
1600 Amphitheater Parkway
Mountain View, CA 94043
shivav@google.com

Divyakant Agrawal^{*}
Department of Computer Science
University of California at Santa Barbara
Santa Barbara, CA 93106
agrawal@cs.ucsb.edu

ABSTRACT

In this collaborative keynote address, we will share Google's experience in building a scalable data infrastructure that leverages datacenters for managing Google's advertising data over the last decade. In order to support the massive online advertising platform at Google, the data infrastructure must simultaneously support both transactional and analytical workloads. The focus of this talk will be to highlight how the datacenter architecture and the cloud computing paradigm has enabled us to manage the exponential growth in data volumes and user queries, make our services highly available and fault tolerant to massive datacenter outages, and deliver results with very low latencies. We note that other Internet companies have also undergone similar growth in data volumes and user queries. In fact, this phenomenon has resulted in at least two new terms in the technology lexicon: *big data* and *cloud computing*. Cloud computing (and datacenters) have been largely responsible for scaling the data volumes from terabytes range just a few years ago to now reaching in the exabyte range over the next couple of years. Delivering solutions at this scale that are fault-tolerant, latency sensitive, and highly available requires a combination of research advances with engineering ingenuity at Google and elsewhere. Next, we will try to answer the following question: is a datacenter just another (very large) computer? Or, does it fundamentally change the design principles for data-centric applications and systems. We will conclude with some of the unique research challenges that need to be addressed in order to sustain continuous growth in data volumes while supporting high throughput and low latencies.

^{*}On leave from UCSB: Visiting Scientist, Google Inc. (2013-14); Director of Research, Qatar Computing Research Institute (2014-15).

This work is licensed under the Creative Commons Attribution-NonCommercial-NoDerivs 3.0 Unported License. To view a copy of this license, visit <http://creativecommons.org/licenses/by-nc-nd/3.0/>. Obtain permission prior to any use beyond those covered by the license. Contact copyright holder by emailing info@vldb.org. Articles from this volume were invited to present their results at the 40th International Conference on Very Large Data Bases, September 1st - 5th 2014, Hangzhou, China. *Proceedings of the VLDB Endowment*, Vol. 7, No. 13
Copyright 2014 VLDB Endowment 2150-8097/14/08.

Speaker Biographies

Shivakumar Venkataraman is Vice President of Engineering for Google's Advertising Infrastructure and Payments Systems. He received his BS in Computer Science from IIT, Madras in 1990 and received his MS and PhD in Computer Science from University of Wisconsin at Madison in 1991 and 1996 respectively. From 1996 to 2000, he worked as an Advisory Software Engineer with IBM working on the development of IBM's federated query optimizers and associated technologies. After leaving IBM in 2000, he worked with Cohera Corporation, PeopleSoft, Required Technologies, and AdeSoft. He also served as a Visiting Faculty member at UC Berkeley in 2002. He has been with Google since 2003. At Google, Dr. Venkataraman is recognized for the vision in the development of critical technologies for databases: scalable distributed database management system F1 [3, 4], scalable data warehousing solution Mesa [2], and scalable log-processing system Photon [1].

Divyakant Agrawal is a Professor of Computer Science and the Director of Engineering Computing Infrastructure at the University of California at Santa Barbara. His research expertise is in the areas of database systems, distributed computing, data warehousing, and large-scale information systems. From January 2006 through December 2007, Dr. Agrawal served as VP of Data Solutions and Advertising Systems at the Internet Search Company ASK.com (his work at ASK.com resulted in a Technology Patent). Dr. Agrawal has also served as a Visiting Senior Research Scientist at the NEC Laboratories of America in Cupertino, CA from 1997 to 2009 (his work at NEC resulted in several Technology Patents). During his professional career, Dr. Agrawal has served on numerous Program Committees of International Conferences, Symposia, and Workshops and served as an editor of the journal of Distributed and Parallel Databases (1993-2008), the VLDB journal (2003-2008), and IEEE Transactions on Knowledge and Data Engineering (2012-2014). He currently serves as the Editor-in-Chief of Distributed and Parallel Databases and is on the editorial boards of the ACM Transactions on Database Systems, ACM Transactions on Spatial Algorithms and Systems, and ACM Books. He also serves on the Board of Trustees of the VLDB Endowment and served on the Executive Committee of ACM Special Interest Group SIGSPATIAL from 2011-14. Dr. Agrawal's research philosophy is to develop data management solutions that are theoretically sound and are relevant in practice. He has authored more than 350 research manuscripts on wide range of topics related to data

management and distributed systems. He has advised more than 35 Doctoral students during his academic career and received the 2011 Outstanding Graduate Mentor Award from the Academic Senate at UC Santa Barbara. Dr. Agrawal is an ACM Distinguished Scientist (2010), an ACM Fellow (2012), and an IEEE Fellow (2012). His current interests are in the area of scalable data management and data analysis in cloud computing environments, security and privacy of data in the cloud, and scalable analytics over social networks data and social media. In 2013-14, he was on a sabbatical leave from UCSB serving as a Visiting Scientist in the Advertising Infrastructure Group at Google, Inc. in Mountain View, CA. In 2014-15, he will be on leave from UCSB and will serve as a Director of Research in Data Analytics at Qatar Computing Research Institute.

1. REFERENCES

- [1] R. Ananthanarayanan, V. Basker, S. Das, A. Gupta, H. Jiang, T. Qiu, A. Reznichenko, D. Ryabkov, M. Singh, and S. Venkataraman. Photon: Fault-tolerant and scalable joining of continuous data streams. In *SIGMOD Conference*, pages 577–588, 2013.
- [2] A. Gupta, F. Yang, J. Govig, A. Kirsch, K. Chan, K. Lai, S. Wu, S. G. Dhoot, A. R. Kumar, A. Agiwal, S. Bhansali, M. Hong, J. Cameron, M. Siddiqi, D. Jones, J. Shute, A. Gubarev, S. Venkataraman, and D. Agrawal. Mesa: Geo-Replicated, Near Real-Time, Scalable Data Warehousing. *PVLDB*, 2014. To appear.
- [3] J. Shute, M. Oancea, S. Ellner, B. Handy, E. Rollins, B. Samwel, R. Vingralek, C. Whipkey, X. Chen, B. Jegerlehner, K. Littlefield, and P. Tong. F1: The fault-tolerant distributed RDBMS supporting Google’s Ad business. In *SIGMOD Conference*, pages 777–778, 2012.
- [4] J. Shute, R. Vingralek, B. Samwel, B. Handy, C. Whipkey, E. Rollins, M. Oancea, K. Littlefield, D. Menestrina, S. Ellner, J. Cieslewicz, I. Rae, T. Stancescu, and H. Apte. F1: A Distributed SQL Database That Scales. *PVLDB*, 6(11):1068–1079, 2013.