Scale-Out Beyond Map-Reduce

Raghu Ramakrishnan
Microsoft
One Microsoft Way, Redmond WA 98052
raghu@microsoft.com

CISL Team Members
Microsoft
One Microsoft Way, Redmond WA 98052
emailed@microsoft.com

Categories and Subject Descriptors
D.1.3 [Concurrent Programming], H.2.3 [Languages] H.2.4 [Systems], I.2.5 [Programming Languages and Software], I.2.6 [Learning], I.5 [Pattern Recognition]

General Terms
Algorithms, Management, Languages.

Keywords
Analytics, Big Data, data science, Hadoop, YARN, REEF, Map-Reduce, SQL, Machine Learning, scale-out.

1. ABSTRACT
The amount and variety of data being collected in the enterprise is growing at a staggering pace. The default now is to capture and store any and all data, in anticipation of potential future strategic value, and vast amounts of data are being generated by instrumenting key customer and systems touch points. Until recently, data was gathered for well-defined objectives such as auditing, forensics, reporting and line-of-business operations; now, exploratory and predictive analysis is becoming ubiquitous. These differences in data heterogeneity, scale and usage are leading to a new generation of data management and analytic systems, where the emphasis is on supporting a wide range of large datasets to be stored uniformly and analyzed seamlessly using whatever techniques are most appropriate, including traditional tools like SQL and BI and newer tools, e.g., for machine learning. These new systems are necessarily based on *scale-out* architectures for both storage and computation. The terms Big Data and data science are often used to refer to this class of systems and applications.

Hadoop has become a key building block in the new generation of scale-out systems. Early versions of analytic tools over Hadoop, such as Hive [1] and Pig [2] for SQL-like queries, were implemented by translation into Map-Reduce computations. This approach has inherent limitations, and the emergence of resource managers such as YARN [3] and Mesos [4] has opened the door for newer analytic tools to bypass the Map-Reduce layer. This trend is especially significant for iterative computations such as graph analytics and machine learning, for which Map-Reduce is widely recognized to be a poor fit. In fact, the website of the machine learning toolkit Apache Mahout [5] explicitly warns about the slow performance of some of the algorithms on Hadoop.

In this talk, I will examine this architectural trend, and argue that resource managers are a first step in re-factoring the early implementations of Map-Reduce, and that more work is needed if we wish to support a variety of analytic tools on a common scale-out computational fabric. I will then present REEF, which runs on top of resource managers like YARN and provides support for task monitoring and restart, data movement and communications, and distributed state management. Finally, I will illustrate the value of using REEF to implement iterative algorithms for graph analytics and machine learning.

KDD'13, August 11–14, 2013, Chicago, Illinois, USA.
ACM 978-1-4503-2174-7/13/08.

2. Short Bio
Raghu Ramakrishnan heads the Cloud and Information Services Lab (CISL) in the Data Platforms Group at Microsoft. From 1987 to 2006, he was a professor at University of Wisconsin-Madison, where he wrote the widely-used text "Database Management Systems" and led a wide range of research projects in database systems (e.g., the CORAL deductive database, the DEVise data visualization tool, SQL extensions to handle sequence data) and data mining (scalable clustering, mining over data streams). In 1999, he founded QUIQ, a company that introduced a cloud-based question-answering service. He joined Yahoo! in 2006 as a Yahoo! Fellow, and over the next six years served as Chief Scientist for the Audience (portal), Cloud and Search divisions, driving content recommendation algorithms (CORE), cloud data stores (PNUTS), and semantic search ("Web of Things"). Ramakrishnan has received several awards, including the ACM SIGKDD Innovations Award, the SIGMOD 10-year Test-of-Time Award, the IIT Madras Distinguished Alumnus Award, and the Packard Fellowship in Science and Engineering. He is a Fellow of the ACM and IEEE.

3. ACKNOWLEDGMENTS
This work was informed by interactions with ML team members in CISL and MSR, Beysim Sezgin and S. Sudarshan from the AIP team at Microsoft, and Dave Campbell in DPG.

4. REFERENCES
1. The Apache Software Foundation, "Apache Hive," 05 2013. [Online]. Available: http://hive.apache.org/.
2. The Apache Foundation, "Apache Pig," 11 12 2012. [Online]. Available: http://pig.apache.org/.
3. The Apache Project, "Apache Hadoop NextGen MapReduce (YARN)," The Apache Project, [Online]. Available: http://hadoop.apache.org/docs/r0.23.0/hadoop-yarn/hadoop-yarn-site/YARN.html.
4. B. Hindman, A. Konwinski, M. Zaharia and I. Stoica, "A Common Substrate for Cluster Computing," in *HotCloud*, 2009.
5. The Apache Software Foundation, "Apache Mahout," 05 2013. [Online]. Available: http://mahout.apache.org/.