Special section on large-scale analytics

Wolfgang Lehner · Michael J. Franklin

Published online: 6 September 2012
© Springer-Verlag 2012

Big Data is no longer exclusively the domain of big organizations. Companies, collaborations, and organizations of all types and sizes are increasingly faced with the need to analyze and make sense of large and growing collections of data. Solving the challenge of large-scale analytics requires innovation across the spectrum of data management: Large volumes of data have to be acquired, processed, stored, and eventually reclaimed. Complex statistical procedures must be applied to those large data sets. Transactional guarantees are required to provide a consistent picture with operational systems. Metadata must be maintained to provide the context of the underlying raw data for later analysis. These challenges must be faced independently and together in order to establish a scalable, affordable, and flexible large-scale analytics infrastructure.

This special section focuses on conceptual and systems-architecture issues in this emerging area. The three selected papers present recent efforts that push the envelope of novel schemes for large-scale analytics and provide a deeper understanding and assessment of the current state of the art.

The first paper of the special section focuses on optimizing the well-known MapReduce processing paradigm. The proposed approach exploits the fact that MapReduce clusters support multiple, concurrently running jobs often accessing the same set of data. The paper entitled “On the optimization of schedules for MapReduce workloads in the presence of shared scans” improves traditional MapReduce processing

with two main contributions. First, the authors introduce the technique of *cyclic piggybacking* to implement shared scans and reduce to overall access cost. Then, given the ability to share scans, the paper addresses the optimization of job scheduling to exploit the shared scans in an optimal fashion. The paper then presents the CIRCUMFLEX scheduler, a generalized version of the FLEX scheduler allowing a wide variety of different cost metrics to be used. Insights into the implementation of these techniques and results of simulation and real benchmark experiments are given as well.

The second paper addresses one of the hot topics in database-centric research in the context of the MapReduce data processing paradigm by trying to bridge the gap between both worlds. In their article entitled “SCOPE: parallel databases meet MapReduce,” the authors carefully layout the design principles and implementation aspects of the SCOPE system. The general goal of the Scope system is to provide an efficient and flexible platform for very large-scale/massive data analytics by providing an easy to program interface supported by sophisticated optimization to achieve high performance and reliability. The paper outlines the compilation of a query from SQL-like query specification via the optimizer, which turns the query into a data-flow graph for the distributed query engine. The paper also outlines the core principles of the execution environment, discussing the role of a job manager orchestrating the cluster resources and generating the schedule for the different jobs to provide fault tolerant behavior. In summary, the paper gives a superb view into the details of a large-scale platform used daily for various analytics and data mining activities at Microsoft.

The third paper in this special section, entitled “GBASE: an efficient analysis platform for large graphs,” tackles some key challenges in the context of very large graphs. Graph-structured data are commonplace in domains such as social network analytics, logistics, or even security.

W. Lehner (✉)
Technische Universität Dresden,
Nöthnitzer Str. 46, 01187 Dresden, Germany
e-mail: wolfgang.lehner@tu-dresden.de

M. J. Franklin
Algorithms, Machines and People Lab, University of California,
465 Soda Hall #1776, Berkeley, CA 94720, USA
e-mail: franklin@cs.berkeley.edu

In contrast to classical table-like structures, graph structures are more entity-focused, implying that individual nodes may have different types and may have specific relationships to other nodes. Running large analytical queries on top of graphs impacts the requirements for different types of operators and storage structures. This paper outlines the approach taken within the GBASE project. The authors carefully layout storage and compression structures used and present primitive graph operations allowing complex analytical queries through composition. Extensive experi-

mental results are given to underline the efficiency of the approach taken.

Taken together, these papers provide a strong point of reference for future research and teaching on this crucial topic. We thank the authors of all of the papers submitted to this special section as well as the diligent reviewers who provided perspective, advice, and keen assessments of the submissions in the context of this fast-moving field. We look forward to continue the innovation and leadership of the VLDB community in the “Big Data” analytics space.