



Guest Editorial: Special issue on VLDB 2019

Fatma Özcan¹ · Lei Chen²

© Springer-Verlag GmbH Germany, part of Springer Nature 2020

This is the special issue of best papers selected from VLDB 2019 Conference, which was held in Los Angeles, California, from August 26 to August 30, 2019. VLDB 2019 covered many aspects of data management and analytics, including data integration, cloud databases, distributed transactions, query processing and optimization, crowdsourcing, graph analytics, scalable machine learning, and distributed systems. VLDB 2019 received 677 research submissions. Out of 677 submissions includes 587 regular research papers, 8 vision papers, 31 innovation systems and applications papers, and 51 experimental and analysis papers. The Review Board, Associate Editors, and Editors-in-Chief of PVLDB volume 12 have worked very hard and have selected 128 papers to be presented at VLDB 2019, with an acceptance rate of 18.9%.

Based on the recommendation from the associate editors and the program chairs, seven outstanding papers were selected as best paper candidates from the accepted ones. The Best Paper Selection Committee, consisting of M. Tamer Özsu (chair), Divesh Srivastava, Ashraf Abounaga, Georgia Koutrika, Fatma Özcan, and Lei Chen, reviewed all the seven papers thoroughly and selected the best paper awards for the conference. We invited the authors of the seven selected papers to submit an extended version of their papers, and five of them submitted to this special issue. The reviewers for the manuscripts submitted to the journal were a mix of those who had originally reviewed the conference versions, as well as additional experts who reviewed only the extended submissions. After two rounds of reviewing, all five papers were accepted for publication in this issue, covering a diverse spectrum of topics ranging from distributed consistency and concurrency control, cloud storage scheme to data explanation tools, and subjective database querying.

In the paper “Autoscaling Tiered Cloud Storage in Anna”, C. Wu, V. Sreekanti and J. M. Hellerstein proposed a novel

solution to extend, Anna, a distributed key-value store into an autoscaling, multi-tier service for the cloud. The goals of traditional cloud storage services lead to poor cost-performance trade-offs for applications. To improve performance, developers are inhibited by two key types of barriers, cost-performance barriers and static deployment barriers. As a distributed key-value store, Anna was initially developed based on a fully shared-nothing, thread-per-core architecture with background gossip protocol across the cores and nodes. The authors have extended Anna by incorporating the following new designs: multi-master selective replication, a vertical tiering of storage layers and elasticity of each tier, making it to dynamically adjust configuration and match resources to workloads, thus, overcome the cost-performance and static deployment barriers. The experimental results demonstrate that the extended Anna outperforms AWS ElasticCache and Masstree by up to an order of magnitude and DynamoDB by more than two orders of magnitude in efficiency.

Motivated by the applicability and extensibility limitations of the current explanation engines which are designed as standalone data processing tools that do not interoperate with traditional, SQL-based analytics workflows, in the paper “DIFF: A Relational Interface for Large-Scale Data Explanation” F. Abuzaid, P. Kraft, S. Suri, E. Gan, E. Xu, A. Shenoy, A. Ananthanarayan, J. Sheu, E. Meijer, X. Wu, J. Naughton, P. Bailis, and M. Zaharia proposed the DIFF operator, which unifies the core functionality of existing explanation engines with declarative relational query processing to capture the same semantics of the existing engines together with the production use cases in industry. In the extended version, they have introduced an ANTI DIFF operator and designed new logical and physical query optimizations for DIFF. The experimental results demonstrated that DIFF can outperform existing explanation engines by up to an order of magnitude.

The paper “Interactive Checks for Coordination Avoidance” by M. Whittaker and J. M. Hellerstein proposed to use invariant confluence to combine the best of strong consistency and weak consistency to achieve both easy to reason coordination as well as high performance. However, the challenge is how to determine whether or not an object is invariant

✉ Lei Chen
leichen@cse.ust.hk

¹ IBM Research - Almaden, San Jose, USA

² Hong Kong University of Science and Technology, Kowloon, Hong Kong

confluent. In this paper, they first defined conditions making a commonly used sufficient condition for invariant confluence both necessary and sufficient. After that, based on this sufficient condition, they have designed a decision process for a general-purpose interactive invariant confluence and a novel sufficient condition that can be checked automatically. After that, they introduced a generalization of invariant confluence to replicate non-invariant confluent objects with a small amount of coordination. They implemented a prototype system, Lucky, which can efficiently handle common real-world workloads. Segmented invariant confluent replication can deliver up to an order of magnitude more throughput than linearizable replication for low contention workloads.

The paper “Gossip-Based Visibility Control for High Performance Geo-Distributed Transactions” by Hua Fan and Wojciech Golab pointed out to huge concurrency control and data replication overheads for ACID transactions under conflicts across globally distributed data. To address this problem, the authors have proposed a novel distributed protocol, called Ocean Vista, to guarantee strict serializability. Specifically, they adopt a multi-version protocol to track visibility using version watermarks and achieved correct visibility decisions using efficient gossip. They improved transaction efficiency under high cross data center network delays through gossiping the watermarks, which enabled asynchronous transaction processing and acknowledging transaction visibility in batches in the concurrency control and replication protocols. The experiments in a multidata center cloud environment well demonstrated that Ocean Vista outperforms a leading distributed transaction processing engine (TAPIR) more than 10-fold in terms of peak throughput with reasonable additional latency for gossip, and a more restricted transaction model.

The last paper of this special issue is “Querying Subjective Data” by Y. Li, A. Feng, J. Li, S. Chen, S. Mumick, A. Halevy, V. Li and W.-C. Tan. In this paper, the authors have proposed a very interesting study, querying subjective databases. Online users currently can search products only through objective attributes. However, more often, they want to find products through searching subjective opinions left by other users. To support search over subjective experiences, users’ experience should be modeled and stored in the database for future queries. In this paper, the authors introduced OpineDB, a subjective database system which can model and store subjective experiences. They also proposed solutions to translate subjective queries by matching the user query phrases to the subjective database schema. Their experimental study on real hotel and restaurant review data clearly demonstrated the advantages of integrating experiential conditions into user queries.

Finally, we thank all the authors for their tremendous efforts in extending their conference manuscripts significantly. We also would like to thank all the reviewers, including the previous reviewers in the conference and the new reviewers, for their insightful suggestions to help the authors improve the quality of their papers. Finally, we reiterate our profound gratitude to the Best Papers Committee of VLDB 2019 and AEs for their efforts on selecting these papers from a large and competitive high quality publications.

PVLDB Volume 12 Editors-in-Chief

VLDB 2019 Program Committee Chairs

Publisher’s Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.