

# Special Issue on Graph Processing: Techniques and Applications

Xuemin Lin<sup>1</sup> · Jeffrey Xu Yu<sup>2</sup>

Received: 21 February 2017 / Accepted: 22 February 2017  
© The Author(s) 2017. This article is published with open access at Springerlink.com

Given the existence of large graphs in many real applications including the web, online social networks, RDF, etc., there is a need to find new approaches to efficiently process such large graphs. This special issue presents high-quality research ideas related to the graph processing from both techniques and applications.

The first paper of this issue: *Big Graph Analyses: From Queries to Dependencies and Association Rules* by Wenfei Fan and Chunming Hu, is a position paper which presents an overview of the recent advances in the study of big graphs, from theory to systems to applications. The theory of bounded evaluability is introduced based on which querying big graphs can be done by a bounded amount of the data.

The second paper of this issue: *Graph-based RDF Data Management* by Lei Zou and M. Tamer Özsu, gives an overview of the systems they have developed over years including gStore, gStore-D, and gAnswer to process RDF data and discusses the design philosophy behind the systems developed.

The third paper of this issue: *Distance-aware Selective Online Query Processing over Large Distributed Graphs* by Xiaofei Zhang and Lei Chen, focuses on online selective graph queries based on the shortest path distance semantics and discusses how to eliminate redundant graph traversal by utilizing a distance-aware index constructed offline.

The fourth paper of this issue: *Graph Partitioning for Distributed Graph Processing* by Makoto Onizuka, Toshimasa Fujimori, and Hiroaki Shiokawa, presents effective graph partitioning techniques to achieve low communication cost and good load balance among computational nodes at the same time in a distributed system.

The fifth paper of this issue: *Efficient Breadth-First Search on Massively Parallel and Distributed Memory Machines* by Koji Ueno, Toyotaro Suzumura, Naoya Maruyama, Katsuki Fujisawa, and Satoshi Matsuoka, focuses on one of the most fundamental graph algorithm breadth-first search (BFS) and discusses new methods to compute BFS for graphs with one trillion vertices using supercomputers. The benchmark of Graph 500 is introduced, and the results of their algorithms are presented.

The sixth paper of this issue: *Investigating TSP Heuristics for Location-Based Services* by Weihuang Huang and Jeffrey Xu Yu, investigates 22 heuristics and shows the best heuristics in terms of accuracy and efficiency using benchmarks, real datasets, and synthetic datasets.

We would like to thank all those who contributed to this special issue, Lizhou Zhou, Jianzhong Li, and Elisa Bertino, the managing editor, and editor-in-chiefs, for their guidance, the reviewers for their thorough comments that help in enhancing the quality of the papers, and the authors for submitting their papers to this special issue.

---

✉ Jeffrey Xu Yu  
yu@se.cuhk.edu.hk  
Xuemin Lin  
lxue@cse.unsw.edu.au

<sup>1</sup> The University of New South Wales, Sydney, Australia

<sup>2</sup> The Chinese University of Hong Kong, Hong Kong, China

**Open Access** This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.