

Guest Editorial: Special Section on the International Conference on Data Engineering

Isabel F. Cruz, Elena Ferrari, *Fellow, IEEE*, and Yufei Tao

THE 30th IEEE International Conference on Data Engineering (ICDE) took place in Chicago, IL, USA, from March 31 to April 4, 2014. The conference received 446 paper submissions for the research track, 33 submissions for the industrial track, and 65 demo proposals. The research program featured 89 papers, the industrial program 10 papers, and the demonstration program 28 demos. The conference program also included six seminar tutorials and one panel.

This special section is comprised of the long versions of five exceptional papers selected from the research program. These papers were highly rated by the research program committee, and have been extended into the current versions with new technical results. They have gone through another stringent review process to meet the high quality standards of the *IEEE Transactions on Knowledge and Data Engineering*. These papers together cover a range of active topics in today's database research, as we introduce below.

The paper "Scaling HTM-Supported Database Transactions to Many Cores"—whose short version won the best-paper award at the conference—by Leis, Kemper, and Neumann explores an innovative way to deploy hardware transactional memory (HTM) for efficient concurrency control on database transactions. The authors propose several fundamental synchronization techniques that significantly facilitate the implementation of concurrent HTM-based algorithms. Their experiments exhibit nearly lock-free execution of transactions in a variety of practical scenarios.

The paper "Top- k Preferences in High Dimensions" by Yu, Agarwal, and Yang revisits top- k queries and reverse top- k queries, both of which are among the most extensively studied problems in the database area. The authors develop new algorithms with an impressive fusion of theory and practice. An important feature of their solutions is a new way to formalize *sparsity*, motivated by the observation that a multidimensional object often has non-zero weights on some, but not all, dimensions. This formalization is of independent interests, and may find additional use on other problems.

The paper "Answering Pattern Queries Using Views" by Wang, Xin, Fan, Wenfei, and Wu, investigates the novel issue of answering graph pattern queries using views.

- I.F. Cruz is with the University of Illinois at Chicago, Chicago, IL 60607-7053. E-mail: ifc@cs.uic.edu.
- E. Ferrari is with the University of Insubria, Italy. E-mail: elena.ferrari@uninsubria.it.
- Y. Tao is with the Chinese University of Hong Kong, Hong Kong. E-mail: taoyf@cse.cuhk.edu.hk.

For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org, and reference the Digital Object Identifier below. Digital Object Identifier no. 10.1109/TKDE.2015.2495958

For this purpose, the authors propose a notion of pattern containment and efficient algorithms to answer graph pattern queries, based on this characterization. They also study the complexity of determining minimal and maximum containment of pattern queries and show the effectiveness of the proposed algorithms on large real-life graph data.

The paper "Merlin: Exploratory Analysis with Imprecise Queries" by Qarabaqi and Riedewald addresses key challenges in the exploratory search of large databases. The Merlin system supports imprecise queries, provides feedback about the sensitivity of the result to changes of individual conditions, and suggests new conditions. Since dealing with probabilities and expectations is computationally expensive, Merlin can trade result accuracy for faster response time based on a user-controlled real-time threshold.

The paper "Venus: Scalable Real-time Spatial Queries on Microblogs with Adaptive Load Shedding" by Magdy, Mokbel, Elnikety, Nath, and He extends the state of the art of spatio-temporal real-time querying of hundreds of millions of incoming microblogs with high arrival rates in main-memory, while providing fast query responses and efficient memory utilization. The Venus system supports an efficient in-memory spatio-temporal index, a scalable query processor, and a comprehensive suite of memory optimization techniques.

The authors thank all of the authors for meticulously revising their papers to address the comments raised by the conference and journal reviewers. Their appreciation also goes to the program committee members of ICDE'14 for their hard work in assessing the submissions. They also cherish this opportunity to thank all of the parties involved to make the conference a tremendous success. Last but not least, we sincerely hope that you will enjoy this special section.

Isabel F. Cruz
Elena Ferrari
Yufei Tao
Guest Editors



Isabel F. Cruz received the PhD degree from the University of Toronto. She is a professor of computer science at the University of Illinois at Chicago. Among her honors, she received the US National Science Foundation (NSF) CAREER Award and she was a University of Illinois Faculty scholar. She has published more than 120 peer-reviewed papers in databases, Semantic Web, and information integration and visualization, receiving three best paper awards. She has served as program committee chair and conference chair for leading conferences, including IEEE ICDE, ISWC, and ACM SIGSPATIAL GIS.



Elena Ferrari is a full professor of computer science at the University of Insubria, Italy, and a scientific director in the K&SM Research Center. Her research activities are related to access control, privacy, and trust. In 2009, she received the IEEE Computer Society's Technical Achievement Award for "outstanding and innovative contributions to secure data management." She received a Google Award in 2010, and an IBM Faculty Award in 2014. She is a fellow of the IEEE and an ACM distinguished scientist.



Yufei Tao is a full professor at the Chinese University of Hong Kong. From 2011 to 2013, he was simultaneously a visiting professor, under the World Class University program of the Korean government, in the Division of Web Science and Technology, Korea Advanced Institute of Science and Technology (KAIST), Korea. He served as an associate editor of the *ACM Transactions on Database Systems (TODS)*, and of the *IEEE Transactions on Knowledge and Data Engineering (TKDE)*. He received two SIGMOD Best

Paper Awards (in 2013 and 2015, respectively), and a Hong Kong Young Scientist Award in 2002.

▷ **For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.**