Editor's Notes

Welcome to the December 2019 issue of the ACM SIGMOD Record!

This issue starts with the Database Principles column featuring an article by Cuenca Grau, Horrocks, Kaminski, Kostylev, and Motik that discusses Limit Datalog, query language that could be appropriate for data-analysis applications. The language extends the well-known declarative language Datalog in ways that make the result powerful enough to naturally capture important data-analysis tasks on complex data sets, while keeping the reasoning no harder than in the case of Datalog. The importance of the language being declarative is that in specifying queries, data analysts can focus on the goals to be achieved in the analysis, rather than on the procedural implementations of the needed tasks. The article details the construction and complexity considerations in designing the language, and supplies illustrations via examples motivated by practical applications. The authors also discuss how the language can be used as a basis for understanding the expressive power of key data-analysis constructs, and outline directions of future work.

The Surveys column features an article by Zhang, Zhang, Wu, Johns, and He that studies the problem of improving the utilization of modern hardware in data-stream processing systems (DSPSs). The article discusses the challenges of optimizing system latency and throughput toward achieving real-time data analytics on large-scale streaming data, in presence of underutilization of the available hardware resources. Toward addressing the challenges, the authors present a systematic study of recent works in the field, organized along the axes of computational optimization, stream I/O optimization, and query deployment. The article also summarizes how hardware-conscious optimization techniques mitigate the gap between DSPSs and the requirements of real-time stream processing that were formulated by Stonebraker and colleagues in their 2005 SIGMOD Record paper. The article concludes by formulating major open issues and proposing directions of future research in the area.

The Research Articles column presents an article by Papadakis, Tsekouras, Thanos, Giannakopoulos, Palpanas, and Koubarakis that introduces the Java gEneric Data Integration (JedAI) toolkit. The focus of the toolkit is on addressing two major challenges in end-to-end entity resolution: development of extensible open-source tools and provision of solutions that apply not only to structured, but also to semi- or even unstructured data. The article details the composition and functionalities of JedAI, outlines its user-friendly interface, and reports the results of the experimental performance comparisons of the toolkit with three state-of-the-art packages described in the literature, using six real data sets. The JedAI code is available from the authors of the article.

The Distinguished Profiles column features Anastasia Ailamaki, professor at the Swiss Federal Institute of Technology (EPFL), and former professor at Carnegie Mellon. Natassa is an ACM Fellow and a Sloan Fellow; she has received the European Young Investigator Award, the 2018 Nemitsas Prize in Computer Science from the President of the Republic of Cyprus, and the 2019 Edgar F. Codd Innovation Award from ACM SIGMOD, as well as ten Best Paper awards for her research papers. Her Ph.D. is from the University of Wisconsin, Madison. In this interview, Natassa talks about her research in two very different systems areas, about merit-based hiring, and about how she approaches working with students and staff in her large research group. She also discusses challenges and tradeoffs in hardware/software codesign, including potential opportunities for both database researchers and database users. Natassa outlines ideas for emerging applications, such as just-in-time access to data, and shares her views on research in energy efficiency for database-management

systems. She also provides insights on a range of other topics, including the additional things that she could work on if she had the time, and gives advice on database research and mentorship.

This issue features two reports. First, it is our pleasure to include the 2018 Seattle Report on database research. Every few years a group of database researchers meets to discuss the state of database research, its impact in practice, and important new directions. This report summarizes the discussion and conclusions of the ninth such meeting, held October 9-10, 2018 in Seattle. The meeting identified the changes that have taken place in the last five plus years, for reasons that include recent technological breakthroughs in a number of areas, including machine learning and artificial intelligence, the rise of interdisciplinary data science, and new trends in the hardware landscape. These changes have given rise to an unprecedented opportunity for the database-research community to have a transformative impact on today's world. The report details specific research challenges and opportunities in the areas of data science, data governance, cloud services, and database engines. The article also outlines community challenges, including those particular to end-to-end solutions in the hands of users, the data-science software ecosystem, and the impact on university campuses stemming from the rise of data science on the curricula. The second article in the Reports column, by Manghi, Candela, Lazzeri, and Silvello, focuses on the Italian Research Conference on Digital Libraries (IRCDL), annual Italian forum for research topics revolving around digital libraries. The 2019 IRCDL conference took place in Pisa; its theme was "Digital Libraries: Supporting Open Science." Science is increasingly becoming digital, and research results are no longer just traditional scientific publications, but are represented by digital artifacts as well, including data sets, software, and experiments. In this landscape, Digital Libraries are an integral part of the evolution of the research outputs. The 2019 IRCDL conference focused on the topics related to findability, preservation, interlinking, and reuse of the research products via Digital Libraries. The article summarizes the ideas presented and discussed in the conference, and outlines key observations and emerging research directions.

On behalf of the SIGMOD Record Editorial board, I hope that you enjoy reading the December 2019 issue of the SIGMOD Record!

Your submissions to the SIGMOD Record are welcome via the submission site:
https://mc.manuscriptcentral.com/sigmodrecord

Prior to submission, please read the Editorial Policy on the SIGMOD Record's website:
https://sigmodrecord.org/sigmod-record-editorial-policy/

Rada Chirkova

December 2019

Past SIGMOD Record Editors:

Yanlei Diao (2014-2019)	Ioana Manolescu (2009-2013)	Alexandros Labrinidis (2007–2009)
Mario Nascimento (2005–2007)	Ling Liu (2000–2004)	Michael Franklin (1996–2000)
Jennifer Widom (1995–1996)	Arie Segev (1989–1995)	Margaret H. Dunham (1986–1988)
Jon D. Clark (1984–1985)	Thomas J. Cook (1981–1983)	Douglas S. Kerr (1976-1978)
Randall Rustin (1974-1975)	Daniel O'Connell (1971–1973)	Harrison R. Morse (1969)