

Editor's Notes

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

This issue starts with the Database Principles column featuring an article by Rahul and Tao on designing top-k indexes. The design goal is to address the problem of retrieving only the best k records, in those cases where users do not need to examine all the answers. Clearly, focusing on retrieving only up to a predefined number of the highest-ranked query answers has the potential to improve the efficiency of query processing. In the space of potential mechanisms for preventing the query processor from accessing the lower-ranked answers, the article guides the reader through the state of the art on designing top-k indexes that provide strong theoretical guarantees for both time and space, and are also efficiently updatable. Powerful results detailed in the article show that the problem of designing top-k indexes is no harder than certain natural related problems. For practitioners, this implies that solutions for those related problems can be used as black boxes in the design of top-k indexes with desired expected complexity. The article also discusses open problems, and provides references on related areas.

The Surveys column features an article by Pierri and Ceri that discusses the issue of false news on social media. The article provides a comprehensive study of recent algorithmic advances in detecting, characterizing, and mitigating false news on popular social-media platforms, and also discusses emerging approaches. The discussion centers on 2017-18 results, and includes detailed comparative descriptions of the most promising ideas, methods, and approaches. The authors also outline potential interventions, while also bringing into the picture ethical concerns about censorship. The article provides an extensive bibliography, as well as pointers to, and detailed comparative descriptions of, reference data sets.

The Systems and Prototypes column features an article by Scherzinger, which reports on a course titled "Modern database concepts" taught in summer 2018 at OTH Regensburg. The goals of the intensive course included teaching ideas behind engines such as Hive, as well as design decisions regarding query-language constructs. The article focuses on the hands-on project that was offered as an option to the students in the course. The project objective was to build miniHive, an SQL-on-Hadoop engine for compiling SQL queries into MapReduce workflows. The article describes the scope of the project milestones, the required self-study materials, and some coding challenges. It concludes with observations about the student outcomes as a result of their having taken the project; the discussion is based on the students' self reporting, as well as on their performance in the course. For those interested in the details, complete course materials can be made available by the author on request.

The Distinguished Profiles column features Richard Hipp, winner of the 2017 SIGMOD Systems Award and of the 2005 Google O'Reilly Open Source Award for SQLite. Richard has his own consulting firm, Hwaci; his Ph.D. is from Duke University. In this interview, Richard talks about SQLite, the most widely deployed database engine in the world. He discusses the reasons for the popularity of SQLite, the technical challenges in creating it, the aviation-grade testing involved in the design process, and the ongoing projects. He also shares why working on SQLite has been a dream job, and what things he could work on in the future. Richard discusses potential advantages of consulting careers for graduates, and gives advice on topics that database researchers could find rewarding to work on.

The Reports column features an article by Ailamaki and colleagues that describes the submission-evaluation process for the SIGMOD 2019 research track. The article discusses actionable goals concerning reviews and submissions, and describes in detail the infrastructure deployed to address those. In the SIGMOD 2019 research track, many remarkable things were done to scale the submission-evaluation work while ensuring fairness, including ingenious use of features of the chosen tools, as well as optimization approaches based on integer programming for assigning reviewers to submissions. The article discusses the submission-evaluation pipeline and provides details on the specifics of the needed tuning, on the experience of working with the chosen tools, and on the feedback from authors. The article goes on to discuss the success metrics, summarizes the overall rewarding outcomes, and also suggests some points for improvement.

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

Your submissions to the SIGMOD Record are welcome via the submission site:

<http://sigmod.hosting.acm.org/record>

Prior to submission, please read the Editorial Policy on the SIGMOD Record's website:

<https://sigmodrecord.org/sigmod-record-editorial-policy/>

Yanlei Diao and Rada Chirkova

June 2019

Past SIGMOD Record Editors:

Ioana Manolescu (2009-2013)
Ling Liu (2000-2004)
Arie Segev (1989-1995)
Thomas J. Cook (1981-1983)
Daniel O'Connell (1971-1973)

Alexandros Labrinidis (2007-2009)
Michael Franklin (1996-2000)
Margaret H. Dunham (1986-1988)
Douglas S. Kerr (1976-1978)
Harrison R. Morse (1969)

Mario Nascimento (2005-2007)
Jennifer Widom (1995-1996)
Jon D. Clark (1984-1985)
Randall Rustin (1974-1975)