

Guest Editor's Notes

Welcome to the March 2022 issue of the ACM SIGMOD Record!

The new year of 2022 begins with a special issue on the **2021 ACM SIGMOD Research Highlights Awards**. These are awards for the database community to showcase a set of research projects that exemplify core database research. In particular, each of these projects addresses an important problem, represents a definitive milestone in solving the problem, and has the potential of significant impact. These awards also aim to make the selected works widely known in the database community, to our industry partners, and to the broader ACM community.

The awards committee and editorial board included Marcelo Arenas, Rada Chirkova, Wim Martens, Kenneth Ross, and Jun Yang. We solicited articles from PODS 2021, SIGMOD 2021, VLDB 2021, ICDE 2021, EDBT 2021, and ICDT 2021, as well as from community nominations. Through a careful review process ten articles were finally selected as 2021 Research Highlights. The authors of each article worked closely with the editorial board to rewrite the article into a compact 8-page format and improved it to appeal to the broad data-management community. In addition, each research highlight is accompanied by a one-page technical perspective written by an expert on the topic presented in the article. The technical perspective provides the reader with an overview of the background, the motivation, and the key innovation of the featured research highlight, as well as its scientific and practical significance.

The 2021 research highlights cover a broad set of topics, including (a) a significant step forward for query optimization that combines traditional wisdom with neural networks (“Bao: Making Learned Query Optimization Practical”); (b) a new approach that makes it easier for data processing systems to exploit high-speed networks (“DFI: The Data Flow Interface for High-Speed Networks”); (c) a foundation for future cloud information systems that aim at balancing performance and consistency (“FoundationDB: A Distributed Key Value Store”); (d) a new framework that learns deep contextualized representations on relational Web tables (“TURL: Table Understanding through Representation Learning”); (e) a highly scalable approach for embedding attributed networks on a single server (“No PANE, No Gain: Scaling Attributed Network Embedding in a Single Server”); (f) key innovations for bipartite matching that significantly improve the computation of the assignment costs (“Bipartite Matching: What to do in the Real World When Computing Assignment Costs Dominates Finding the Optimal Assignment”); (g) a system that combines the advantages of imperative and functional dataflow systems (“Imperative or Functional Control Flow Handling: Why not the Best of Both Worlds?”); (h) new insights for estimating quantiles in streams with strong theoretical guarantees, which are already available in the Apache Dataskeches library (“Relative Error Streaming Quantiles”); (i) characterizations and answers to foundational questions on the consistency problem for relations under bag semantics (“Structure and Complexity of Bag Consistency”); and (j) an interesting new connection between counting the number of models of a CNF formula and the distinct elements in a data stream (“Model Counting Meets Distinct Elements in a Data Stream”).

On behalf of the SIGMOD Record Editorial Board, I hope that you enjoy reading the March 2022 issue of the SIGMOD Record!

Wim Martens

March 2022