

Guest Editor's Notes

Welcome to the March 2021 issue of the ACM SIGMOD Record!

The new year of 2021 begins with a special issue on the **2020 ACM SIGMOD Research Highlight Award**. This is an award for the database community to showcase a set of research projects that exemplify core database research. In particular, these projects address an important problem, represent a definitive milestone in solving the problem, and have the potential of significant impact. This award also aims to make the selected works widely known in the database community, to our industry partners, and to the broader ACM community.

The award committee and editorial board included Marcelo Arenas, Rada Chirkova, Wim Martens, Jun Yang, and Divesh Srivastava. We solicited articles from PODS 2020, SIGMOD 2020, VLDB 2020, ICDE 2020, EDBT 2020, and ICDT 2020, as well as from community nominations. Through a careful review process ten articles were finally selected as 2020 Research Highlights. The authors of each article worked closely with an associate editor 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 2020 research highlights cover a broad set of topics, including (a) a substantial advance on the question of reliability of streaming algorithms when the input to the algorithm might depend on the algorithm's earlier output ("A Framework for Adversarially Robust Streaming Algorithms"); (b) key innovations in processing database transactions driven by new high-bandwidth communication technologies that are becoming common in data centers ("Chiller: Contention-centric Transaction Execution and Data Partitioning for Modern Networks"); (c) a versatile framework for automatic extraction of benchmarks and their distributed execution and performance monitoring ("DIAMetrics: Benchmarking Query Engines at Scale"); (d) non-trivial results on finding a 2-approximate solution to the problem of directed densest subgraph discovery using a new concept of $[x, y]$ -core ("Efficient Directed Densest Subgraph Discovery"); (e) insightful work on data structures for similarity search that support algorithmic fairness ("Fair Near Neighbor Search via Sampling"); (f) a general abstraction, along with advanced interfaces, focusing on effective visual search ("From Sketching to Natural Language: Expressive Visual Querying for Accelerating Insight"); (g) practical improvements in hash tables for string processing in real system implementations ("Optimistically Compressed Hash Tables & Strings in the USSR"); (h) answers to foundational questions of the representation and querying of infinite probabilistic databases ("Probabilistic Data with Continuous Distributions"); (i) an elegant quantitative approach to defining and computing explanation scores, based on the Shapley value of cooperative games ("Query Games in Databases"); and (j) design and evaluation of a persistent dynamic hash table using Optane byte-addressable non-volatile memory ("Scaling Dynamic Hash Tables on Real Persistent Memory").

On behalf of the SIGMOD Record Editorial Board, I hope that you enjoy reading the March 2021 issue of the SIGMOD Record!

Divesh Srivastava

March 2021