Guest Editors' Notes

Welcome to the March 2025 issue of the ACM SIGMOD Record, a special issue presenting papers and technical perspectives covering a selection of the best conference papers from 2024 in data management. These papers have been given the **2024 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, the papers 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 Angela Bonifati, Rada Chirkova, Alfons Kemper, Samuel Madden, Kenneth Ross (chair) and Ke Yi. We solicited articles from PODS 2024, SIGMOD 2024, VLDB 2024, ICDE 2024, EDBT 2024, and ICDT 2024, as well as from community nominations. We used a careful review process, in which each nominated paper was discussed in a virtual meeting. Papers with conflict of interest were discussed in the absence of the conflicting committee members. This year, ten articles were finally selected as 2024 Research Highlight Award winners.

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 in the March 2025 issue 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 2024 research highlight award winners include:

1) Theoretical results for estimating the output size of a join query ("Join Size Bounds using Lp Norms on Degree Sequences");
2) A technique to dynamically partition data in a manner that reveals nothing about the history of operations on the data ("History-Independent Dynamic Partitioning: Operation-Order Privacy in Ordered Data Structures");
3) A new type of synopsis for high volume data streams that efficiently supports queries on dynamically chosen attributes ("OmniSketch: Efficient Multi-Dimensional High-Velocity Stream Analytics with Arbitrary Predicates");
4) A way to compose different components of a query processing pipeline using partial evaluation and a unified exchange format ("BOSS - An Architecture for Database Kernel Composition");
5) A method to ensure consistency across multiple database caches ("CausalMesh: A Causal Cache for Stateful Serverless Computing");
6) A comprehensive analysis of semantic and implementation issues for views in graph databases ("Implementation Strategies for Views over Property Graphs");

7) Theoretical and experimental results showing how to maintain a sample of a join result over streaming data (“Reservoir Sampling over Joins”);
8) A formalism for defining relational query patterns that intuitively capture query structure (“On The Reasonable Effectiveness of Relational Diagrams”);
9) A system for automatically cleaning data in files of comma-separated values (“TASHEEH: Repairing Row-Structure in Raw CSV Files”);
10) An application of large language models to database tuning, where the system reads the manuals and recommends tuning parameters (“GPTuner: A Manual-Reading Database Tuning System via GPT-Guided Bayesian Optimization”);

On behalf of the SIGMOD Record Editorial Board, we hope that you enjoy reading the March 2025 issue of the SIGMOD Record!

Angela Bonifati Rada Chirkova Alfons Kemper
Samuel Madden Kenneth Ross Ke Yi

March 2025