

Editorial

This new issue opens the ninth volume of JIDM. It brings seven articles, one article regularly submitted to the Journal (Regular Papers Section) and six extended versions of best papers from the 32nd Brazilian Symposium on Databases (SBBD 2017 Section).

The article presented in Regular papers Section, entitled Mind your dependencies for semantic query optimization, is authored by E. H. M. Pena, E. Falk, J. A. Meira and E. C. de Almeida. It addresses the problem of semantic query optimization. In this sense, the authors propose a data-driven and query-aware tool, denoted focused dependency selector (FDSel), to find functional dependencies (FDs), relevant for speeding up semantic query execution. In order to reduce the amount of select FDs, information on database workloads are used.

Section *SBBD 2017* includes extended and revised versions of papers selected from the 32nd Brazilian Symposium on Databases. SBBD is the official database conference of the Brazilian Computer Society (SBC) and the largest venue in Latin America for presentation and discussion of research results in the domain of information and data management. The first three articles of the section are extended versions of the best full papers, and the following three are extended versions of the best short papers presented at the symposium. The section opens with the article entitled “Cutting-edge relational graph data management with Edge-k : From one to multiple edges in the same row”, by Lucas C. Scabora, Paulo H. Oliveira, Gabriel Spadon, Daniel S. Kaster, Jose F. Rodrigues-Jr, Agma J. M. Traina, and Caetano Traina Junior. It proposes a technique to store graph data in a k-edge table. It groups k edges originating from the same vertex in a single row. Experimental results show that k-edge outperforms the traditional single edge-table as well as Neo4j when executing Single Source Shortest Path (SSSP) queries. The article “Beyond hit-or-miss: a comparative study of synopses for similarity searching”, by Marcos V. N. Bedo, Daniel de Oliveira, Agma J. M. Traina, and Caetano Traina Jr., tackles the problem of similarity queries performance. To this end, a study of different probability density functions, called synopses, to represent distance distributions is presented. Moreover, a model to determine the cost of range and neighborhood queries is proposed, and results with synopses that better estimate selectivity, radii, I/O, and CPU costs are presented. The article “STACY: Strength of Ties Automatic-Classifer over the Years”, by Michele A. Brandão, Pedro O. S. Vaz de Melo, and Mirella M. Moro, proposes an algorithm to determine how strong the connections among people in a social network are. The algorithm automatically determines the class of such connections, considering how the network evolves over time. Experimental results on academic networks involving co-authorship are presented. The first article extended from a short paper is entitled “Tie Strength Metrics to Rank Pairs of Developers from GitHub”, from Natércia A. Batista, Guilherme A. Sousa, Michele A. Brandão, Ana Paula C. da Silva, and Mirella M. Moro. It analyzes the collaborative network of social coding in order to determine the strength of ties between pairs of software developers. The impact of several properties for characterizing these ties are analyzed, as well as the effectiveness of aggregation methods. The results are compared to the ranking of developers from Github. The paper “Towards an Empirical Evaluation of Scientific Data Indexing and Querying”, by Thaylon Guedes, Vítor Silva, José Camata, Marcos V. N. Bedo, Marta Mattoso, and Daniel De Oliveira, considers the problem of querying results of scientific simulations that are not loaded into a database management system, but stored in their original format. A case study in the domain of Computational Fluid Dynamics for sediment deposition is presented to compare two approaches: PostgresRAW routines for adaptive query processing, and FastBit methods for raw data file indexing and querying. The section closes with the article entitled “LABAREDA: A Predictive and Elastic Load Balancing Service for Cloud-Replicated Databases” by Carlos S. S. Marinho, Leonardo

O. Moreira, Emanuel F. Coutinho, José S. Costa Filho, Flávio R. C. Sousa, and Javam C. Machado. It proposes a load balancing service that helps cloud providers to satisfy their SLA requirements. The service architecture contains a prediction module, for detecting possible SLA violations, and an elasticity module, that allocates new database replicas whenever such a violation is identified. The experimental analysis considers two prediction models: Autoregressive Integrated Moving Average (ARIMA) and Exponential Moving Average (EMA), and shows that both present similar accuracy.

We would like to thank everyone who contributed to this edition of JIDM, particularly reviewers for their valuable comments and authors for their contributions and hard work in preparing their final manuscripts.

Angelo Brayner
Editor-in-Chief

Maristela Holanda
Associate Editor

Carmem Hara
SBBG Guest Editor