

This material is presented to ensure timely dissemination of scholarly and technical work. Copyright and all rights therein are retained by authors or by other copyright holders. All persons copying this information are expected to adhere to the terms and constraints invoked by each author's copyright. In most cases, these works may not be reposted without the explicit permission of the copyright holder.

Letter from the Special Issue Editor

Aspects of key computing and communication hardware technologies continue to improve rapidly, some at sustained exponential rates. These developments, including advances in geo-positioning, contribute to making research in spatio-temporal data management more relevant than ever.

As the field of data management is maturing, emphasis will be increasingly on rigor. For example, it becomes increasingly important that new contributions be based on the growing body of existing contributions. As other examples, prototype implementation and rigorous experimental studies will become increasingly important.

The contributions in this issue further state of the art in spatio-temporal query processing, but do so indirectly. They do not propose new query processing techniques—instead, their focus is to contribute to the infrastructure for conducting research in spatio-temporal query processing. The term *infrastructure* is interpreted broadly, thus covering aspects such as publicly available query processing toolkits and implementations of query processing techniques; real data, synthetic-data generators, and benchmarks; standards; and surveys of research contributions.

This issue's first paper, by Kornacker et al., presents *amdb*, a graphical design tool for access methods that is built on top of the so-called Generalized Search Tree abstraction (see the coverage of the GiST indexing toolkit in the sixth paper). An analysis framework, complete with performance metrics and support for visualization and debugging, aids the designers of an access method in studying and thus improving their access method. In the second paper, Cammert et al. cover the *eXtensible and fleXible Library (XXL)* for efficient query processing that is being developed at University of Marburg. XXL offers infrastructure that makes it easier to implement advanced query processing functionality, it offers a framework for meaningful comparisons of access methods, and it aims to serve as a repository for query processing techniques and use-cases.

When experimentally evaluating query processing techniques, real as well as synthetic data sets are important. The former aid in ensuring that a technique under study is subjected to realistic conditions. However, real data sets may not be available; further, a single real data set is likely to capture only a specific type of use. In contrast, synthetic data generators allow the generation of data sets with specific properties, thus making it possible to subject a technique to a wide variety of conditions.

In the third paper, Brinkhoff considers the generation of data sets intended for the testing of query processing techniques to do with "moving objects." He covers his own Network-based Generator and Kaufman et al.'s City Simulator, both of which assume that the object movement, from which the generated data result, is constrained to a transportation network. The fourth paper, by Nascimento et al., covers three other data generators for moving objects, *GSTD*, *G-TERD*, and *Oporto*, which do not constrain movement to a network. *GSTD* generates moving-point and moving-rectangle data. *G-TERD* produces sequences of raster images. Being the most elaborate data generator of the three, it is covered in detail in the fifth paper, by Manolopoulos et al. *Oporto* generates data corresponding to fishing-at-sea scenarios. Nascimento et al. also cover several real data sets.

The sixth paper presents a survey of spatio-temporal access methods—methods that index the spatial aspect together with only the past, with only the current time, and with the current time and the future. In this paper, Mokbel et al. cover almost 30 methods. (Note also the survey by Agarwal and Procopiuc in last year's June issue of the Bulletin.) Mokbel et al. also cover two indexing toolkits: *GiST*, which concerns B-tree and R-tree like bounding-region trees, and *SP-GiST*, which concerns space-partitioning trees.

In the last paper, Schmidt and Jensen cover standards and standardization efforts of general relevance to spatio-temporal query processing, and of particular relevance to spatio-temporal data exchange.

It is my hope that this issue will be a useful reference to the spatio-temporal data management research community and will help move spatio-temporal query processing research in the right, rigorous, and experimental direction.

Christian S. Jensen
Department of Computer Science
Aalborg University, Denmark