

Introduction to the EDBT 2006 Special Section

This issue contains three substantially expanded versions of papers that were presented at the 10th International Conference on Extending Database Technology (EDBT), held in Munich, Germany on March 26–31, 2006.

The article “Physical Design Refinement: The ‘Merge-Reduce’ Approach”, by Bruno and Chaudhuri, addresses a very practical issue, physical database design, taking into account both indexes and materialized views. It formalizes the process and offers various techniques that can be used to modify an initial, DBA-provided physical configuration so that it satisfies important constraints that have been specified.

The article “Unified Framework for Fast Exact and Approximate Search in Dissimilarity Spaces”, by Skopal, takes a very novel approach to similarity searching and provides some nontrivial results and general algorithms for using metric indexes to index objects based on nonmetric similarity measures.

Finally, the article “Constructing a Generic Natural Language Interface for an XML Database”, by Li, Yang, and Jagadish, is one of the few articles in recent times that deals with the very important problem of providing a natural language query interface to a database, an XML database in this case. It describes how natural language requests are translated into XQuery expressions and verifies the validity of the approach with experiments involving real users.

I wish to thank the authors for submitting their articles, and associate editor Christian Jensen and Editor-in-Chief Richard Snodgrass for their work in putting together this special section of TODS on EDBT 2006.

YANNIS IOANNIDIS
EDBT 2006 Program Committee Chair