

## Preface

Nowadays, applications that require a high degree of distribution and loosely-coupled connectivity are ubiquitous in various domains, including scientific databases, bioinformatics, and multimedia retrieval. In all these applications, data is typically voluminous and multidimensional, and support for advanced query operators is required for effective querying and efficient processing. Moreover, the highly distributed setting calls for a system architecture with salient properties, including scalability, fault-tolerance, autonomy, and dynamic participation.

Peer-to-peer (P2P) systems emerge as a powerful model for searching vast amounts of data distributed over independent sources. Each peer stores autonomously its own data and the objective is to support efficient and effective techniques for query processing and advanced data analysis. In this context, a key requirement of query processing is retrieval of the exact and complete result set. To achieve this goal, a scalable framework is presented that relies on data summaries that are distributed and maintained as multidimensional routing indices. Different types of data summaries enable efficient processing of a variety of advanced query operators. This book focuses on queries for similarity search, skyline queries, and top- $k$  queries, and identifies appropriate data summaries, proposes effective indexing methods at local level as well as routing indices at network level, and efficient processing algorithms for each query type.

*Akrivi Vlachou  
Christos Doulkeridis  
Kjetil Nørnvåg  
Yannis Kotidis*

Peer-to-Peer Query Processing over Multidimensional  
Data

Vlachou, A.; Doulkeridis, C.; Norvag, K.; Kotidis, Y.

2012, XII, 84 p. 18 illus., Softcover

ISBN: 978-1-4614-2109-2