A Tour Of Spatial Databases

Draft[1]
August, 2001

Authors:
Shashi Shekhar
Sanjay Chawla

Contact:
Department of Computer Science
University of Minnesota
200 Union St. SE, 4 - 192 EE/CS Bldg.
Minneapolis, MN 55455
email: {shekhar, chawla}@cs.umn.edu
web: http://www.cs.umn.edu/research/shashi-group/

Draft date: 08/29/2001
To be published by Prentice Hall

[1] Copies only for review. Limit distribution.

Contents

List of Figures

List of Tables

Preface

Over the years it has become evident in many areas of computer applications that the functionality of database management systems has to be enlarged to include spatially referenced data. The study of spatial database management systems (SDBMS) is a step in the direction of providing models and algorithms for the efficient handling of data related to space.

Spatial databases have now been an active area of research for over two decades. Their results, e.g., spatial multidimensional indexing, are being used in many different areas. The principle impetus for research in SDBMs comes from the needs of existing applications such as geographical information systems (GIS) and computer aided design (CAD), as well as potential applications such as multimedia information systems, data warehousing and NASA's earth observation system. These spatial applications have over one million existing users.

Major players in the commercial database industry have products specifically designed to handle spatial data. These products include the Spatial Data Engine (SDE), by Environment Systems Research Institute (ESRI); as well as Spatial Datablades for object-relational database servers from many vendors including Oracle, IBM and Informix. Research prototypes include Postgres, Geo2, and Paradise. The functionality provided by these systems includes a set of spatial data types such as the point, line, and polygon, and a set of spatial operations, including intersection, enclosure and distance. An industry-wide standard set of spatial data types and operations has been developed by the Open Geographic Information Systems (OGIS) consortium. The spatial types and operations can be made a part of an object-relational query language such as SQL3. The performance enhancement provided by these systems includes a multidimensional spatial index and algorithms for spatial access methods, spatial range queries, and spatial joins.

The integration of spatial data into traditonal databases amounts to resolving many nontrivial issues at various levels. They range from deep ontological questions about the modeling of space, for example whether it should be field-based or object-based, thus paralleling the wave-particle duality in physics, to more mundane but important issues about file management. These diverse topics make research in SDBMS truly interdisciplinary.

Let us use the example of a country dataset to highlight the special needs of spatial databases. A country has at least one nonspatial datum, its name, and one spatial datum, its boundary. There is no ambiguity about storing or representing its name, but unfortunately it is not true for its boundary. Assuming that the boundary is represented as a collection of straight lines, we need to include a spatial data type *line* and the companion types *point* and *region* in the database system to facilitate spatial queries on the object *country*. These new data types need to be manipulated and composed according to some fixed rules leading to the creation of a spatial algebra. Since spatial data is inherently visual and usually voluminous, database systems have to be augmented to provide visual query processing and special spatial indexing . Other important database issues like concurrency control, bulk loading, storage, and security have to be revisited and fine-tuned to build an effective spatial database management system.

This book evolved from the class notes of a graduate course on Scientific Databases (Csci 8705) in University of Minnesota. Researchers and students both within and outside the Computer Science Department found the course very useful and applicable to their work. Despite the good response and high level of interest in the topic, no textbook available in the market was able to meet the interdisciplinary needs of the audience. A recent book by [Scholl et al., 2001] focusses on traditional topics related to query languages and access methods while leaving out current topics such as spatial networks (e.g. roadmaps) and data mining for spatial patterns. A recent monograph by [Adam and Gangopadhyay, 1997] confines itself to cataloging research papers on database issues in GIS, with little reference to the industrial state-of-the-art, Another [Worboys, 1995] also focuses on GIS and has only two chapters devoted to database issues. Many of these book did not provide adequate instuctional support, e.g. questions and problems at the end of each chapter to allow students to assess their understanding of the main concepts. Not suprisingly, our colleagues in academia working in databases, parallel computing, multimedia information, civil and mechanical engineering and forestry have expressed a strong desire for a comprehensive text on spatial databases. Industry professionals involved in software development for GIS and CAD/CAM have also made several requests for information on spatial databases in a collected form.

As a first step towards developing this book we completed a survey paper, "Spatial databases: Accomplishments and Research Needs," for IEEE Transactions on Knowledge and Data Engineering. We noticed that the research literature in Computer Science was skewed with numerous publications on some topics (e.g. spatial indexes, spatial join algorithms) and relatively scarce publications on many other important topics (e.g. conceptual modeling of spatial data). We looked to GIS industry as well as GIS researchers outside Computer Science for ideas in these areas for the relevant chapters in the book.

Being on sabbatical for the academic year 1997-98 facilitated initial work on the book. We completed a draft of many chapters by expanding on the course notes of CSci 8705 based on research literature. First draft of the book was used in database courses at the University of Minnesota. Subsequently the book was revised using feedback from reviewers, collagues and students.

We believe the following features are unique aspects of this book:

- The aim of this book is to provide a comprehensive overview of management systems. It covers current topics, e.g. spatial networks, spatial data mining in addition to traditional topics, e.g. query languages, indexing, query processing,
- A set of questions and problems is provided in each chapter to allow readers to test understanding of core concepts, to apply core concepts in new application domain and to think beyond the material presented in the chapter. Additional instruction aids (e.g. laboratories, lecture notes) are planned on the course web-site.
- A concerted effort has been made to "look beyond GIS." Techniques from SDBMS are finding applications in many diverse areas including multimedia information, CAD/CAM, astronomy, meterology, molecular biology, and computational mechanics.
- In each chapter, we have tried to bring out the object-relational database framework,

which is a clear trend in commercial database applications. This framework allows spatial databases to reuse relational database facilties when possible, while extending relational database facilities when needed.

- We will use spatial data types and operations specified by standards (e.g. OGIS) to illustrate common spatial database queries. These standards can be incorporated into an object-relational query language like SQL-3.

Chapter Organization

The book is divided into eight chapters, each one an important subarea within spatial databases. We introduce the field of spatial databeses in Chapter 1. In Chapter 2, we focus on spatial data models and introduce the field vs. object dichotomy and its implications for database design. Chapter 3 discusses the neccessary enhancements required to make traditional query languages compatible with spatial databases. We provide an extensive discussion of various proposals to extend SQL with spatial capabilities. Spatial databases deal with extraordinarily large amounts of data, and it is essential for DBMS to provide sophisticated storage, compression and indexing methods to enhance the performance of query processing. Spatial data storage and indexing schemes were covered in Chapter 4. ¿From query languages and indexing, we move on to query processing and optimization in Chapter 5. Here we discover that many standard techniques from traditional databases have to be abandoned or drastically modified in order to be applicable in a spatial context. We also introduce the filter-refine paradigm for spatial query processing. In chapter 6, we will show how spatial database technology is being applied to spatial networks. In this chapter we will also cover network data models and query languages. Chapter 7 covers the emerging field of spatial data mining. In this chapter we expose the readers to the concept of spatial dependency that is prevelent in spatial data sets, and show how this can be modeled and incorporated into data mining process. Finally in Chapter 8, we discuss emerging trends in spatial databases.

Acknowledgements

We have received help from many people and we are extremely grateful to them. This book would not have started without encouragement from Prof. Vipin Kumar, Computer Science Department, University of Minnesota; and Dr. Jack Dangermond, President, Environmental Systems Research Institute (ESRI). This book benefitted a great deal from access to ESRI researchers and products. We are also greatful to Dr. Siva Ravada (Oracle Corporation) and Dr. Robert Uleman (Illustra Inc.) for help with understanding of their spatial datablades. We are thankful to Alan Apt and his wonderful staff at Prentice Hall for helping with the process of book-writing and constant encouragement. The presentation improved a great deal from the comments from the anonymous reviewers and we are grateful to them.

Our research related to spatial database has been supported by many organizations including the United Nations Development Programme, the National Science Foundation, the

National Aeronautics and Space Agency, the Army Research Laboratories, the US Department of Agriculture, the Federal Highway Administration, the US Department of Transportation, the Minnesota Department of Transportation, the Center for Urban and Regional Affairs, and the Computing Devices International. Many of the research project resulted in surveys of research literature as well in development of innovative techniques for problems related to spatial databases.

Special thanks to the members of spatial database research group in Computer Science department at the University of Minnesota. They contributed in many different ways including literature surveys, development of examples and figures, insights into various methods as well in developing proper problem formulations and innovative solution methods. We are extremely grateful to Vatsavai Ranga Raju for careful review and multiple improvements to the earlier drafts. We also thank the students in different offerings of Csci 8701, and Csci 8705 for working with the earlier drafts of the books and providing helpful suggestions towards revising the material.

We benefitted from discussions with many other people over the years. They include Marvin Bauer, Yvan Bedard, Paul Bolstad, Nick Bourbakis, Thomas Burk, John Carlis, Jai Chakrapani, Vladimir Cherkassky, Douglas Chub, William Craig, Max Donath, Phil Emmerman, Max Egenhofer, Michael Goodchild, Ralf Hartmut Gueting, Oliver Gunther, John Gurney, Jia-Wei Han, Ravi Janardan, George Karypis, Hans-Peter Kriegel, Robert McMaster, Robert Pierre, Shamkant Navathe, Raymond Ng, Hanan Samet, Paul Schrater, Jaideep Srivastava, Benjamin Wah, Kyu-Young Whang, and Michael Worboys.