

Technical Report

Department of Computer Science
and Engineering
University of Minnesota
4-192 EECS Building
200 Union Street SE
Minneapolis, MN 55455-0159 USA

TR 97-016

Spatial Databases:
Accomplishments and Research
Needs

By: S. Shekhar, S. Ravada,
A. Fetterer, X. Liu, and C.T. Lu

Spatial Databases: Accomplishments and Research Needs

S. Shekhar, S. Ravada, A. Fetterer, X. Liu, C.T. Lu

Computer Science Department, University of Minnesota

EE/CS 4-192, 200 Union St. SE., Minneapolis, MN 55455

[*shekhar|siva|fetterer|xliu|ctl*u]@cs.umn.edu

<http://www.cs.umn.edu/research/shashi-group/>

Abstract

Spatial databases have been an active area of research for over two decades, addressing the growing data management and analysis needs of spatial applications such as Geographic Information Systems. This research has produced a taxonomy of models for space, spatial data types and operators, spatial query languages and processing strategies, as well as spatial indexes and clustering techniques. However, more research is needed to improve support for network and field data, as well as query processing (e.g. cost models, bulk load). Another important need is to apply the spatial data management accomplishments to newer applications such as data warehouses and multimedia information systems. The objective of this paper is to identify recent accomplishments and the research needs in the near term.

Keywords: Spatial Databases, Multi-Dimensional, Object-Relational, Databases, Geographic Information Systems

1 Introduction

1.1 Spatial Databases

A spatial database [4, 6, 15] management system aims at the effective and efficient management of data related to a space such as the physical world (geography, urban planning, astronomy); parts of living organisms (anatomy of the human body); engineering design (very large scale integrated circuits, the design of an automobile or the molecular structure of a pharmaceutical drug); and conceptual information space (a multi-dimensional decision support system, fluid flow, or an electro-magnetic field).

The field of spatial database research has been an active area of research for over two decades. Its results, e.g. spatial multi-dimensional indexes, are being used in a number of areas. The field of spatial databases can be defined by its accomplishments, and current research is aimed at improving its functionality and its performance. The impetus for improving functionality comes from the needs of existing applications such as Geographic Information Systems (GIS) and Computer Aided Design (CAD), as well as from potential applications such as Multimedia Information System (MMIS), Data Warehousing (DWH) and NASA's Earth Observation System (EOS).

Commercial examples of spatial database management include Informix's spatial datablades (i.e. 2D, 3D, Geodetic), Oracle's Universal server with Spatial Data Option, and ESRI's spatial data engine (SDE). Research prototype examples of spatial database management systems include spatial datablades with Postgres, Geo2, and Paradise. The functionalities provided by these systems include a set of spatial data types such as a point, line-segment, polygon and a set of spatial operations such as inside, intersection, and distance. The spatial types and operations may be made a part of a query language such as SQL, which allows spatial querying when combined with an object-relational database management system [1, 14]. The performance enhancement provided by these systems includes a multi-dimensional spatial index, and algorithms for spatial access methods, spatial range queries and spatial joins. Spatial indexing with concurrency control may be implemented in the object-relational server for performance reasons.

Existing and emerging applications require new functionalities including the modeling of network spaces and continuous fields. The performance needs of emerging applications require not only the management of large data-sets, but also new processing strategies for spatial

set-operations, field operations (e.g. slope), and network analysis (e.g. shortest-path, route-evaluation).

1.2 Related Work and Our Contributions

Recent reports[4, 6, 15] have described the accomplishments of spatial database research and have prioritized research needs. A broad survey of spatial database requirements of and an overview of research results is provided by [15, 4]. Basic modeling requirements for spatial objects such as points, lines, and polygons are given in terms of their geometry, topology and object relationships (topological, directional, metric, network). Requirements are given for other user-level issues such as graphical input and output and query language support. Spatial clustering and indexing techniques such as Grid-files, Z-order, Quad-tree, Kd-trees, R-trees and associated join strategies are described. Finally, an architecture for spatial databases is given in terms of the object-relational model.

Research needed to improve the performance of spatial databases in the context of object-relational databases was listed in [6]. The primary research needs identified were concurrency control techniques for spatial indexing methods, the development of cost models for query strategies, and the development of new spatial join algorithms beyond nested-loop and tree matching.

Many of the research needs identified in [6] have since been addressed. For example, concurrency control techniques for R-trees have been studied [7] in the context of R-link trees. Also, new spatial join strategies using space partitioning have been explored. In this paper, we identify the recent accomplishments in spatial databases as well as current research needs based on publications in journals and conference proceedings and recent commercial trends.

1.3 Scope and Outline

The role of the spatial database component is dependent on the type of database management system (DBMS) involved: relational, object-oriented or object-relational. In this paper, we focus the discussion of spatial databases in the context of the object-relational [1, 14, 13] databases, which provide extensibility to many components of traditional databases to support new application domains. Spatial databases have been one of the most common applications of object-relational databases and have influenced their design a great deal. Object-relational databases allow the inclusion of spatial data-types, spatial operations, and multi-dimensional

indexing systems. This three-layer architectural framework is shown in Figure 1, and it consists of an object-relational database management system, a spatial database, and a spatial application such as a GIS or MMIS. The interface between the application and the spatial data system maps application-specific constructs to the spatial database. The spatial database associates the application requirements to the functionality provided by the DBMS. The interface to the DBMS supports specialized query processing, which in turn supports the core database requirements for achieving acceptable performance.

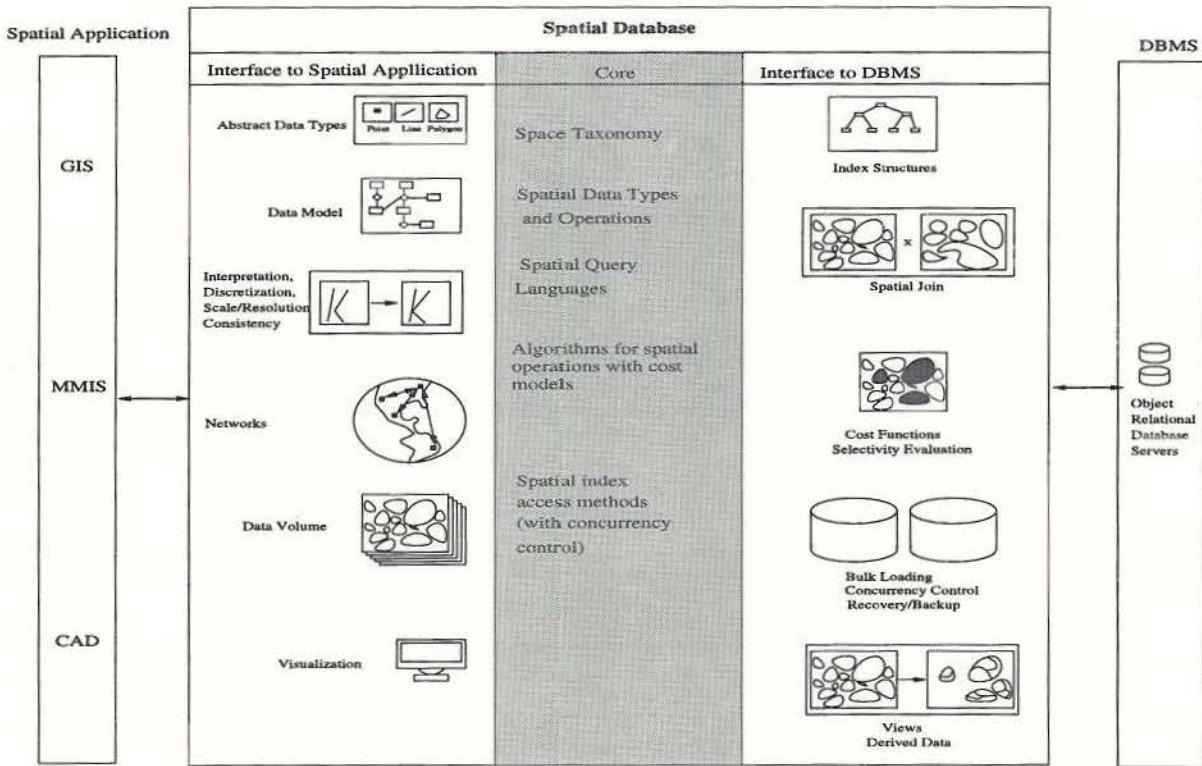


Figure 1: 3-layer architecture

Emerging trends such as world-wide-web interfaces, multimedia data, and image processing are likely to impact the data sharing and analysis needs of spatial databases. Scaling up to large datasets requires new research in many areas beyond spatial databases, including research on file-systems, device-drivers for tertiary storage, computer networks, and visualization software and algorithms related to graphics and computational geometry. This paper does not explore those issues.

The remainder of the paper is organized as follows: Section 2 describes the recent advances in

spatial databases. Section 3 states the research needs for spatial databases. Section 4 highlights our conclusions and motivates exploration of applications whose needs are not currently met by spatial databases.

2 Accomplishments

Research into spatial databases has mainly focused on developing a space taxonomy, spatial data models, spatial query languages and processing strategies, and spatial access methods. This section lists recent important accomplishments, not only for the current applications of spatial databases, but also for the emerging database problems that have spatial dimensions.

2.1 Space Taxonomy

Space is a framework to formalize specific relationships among a set of objects. Depending on the relationships of interest, different models of space such as set-based space, topological space, euclidean space, metric space and network space can be used. Set-based space uses the basic notion of elements, element-equality, sets and membership to formalize the set relationships such as set-equality, subset, union, cardinality, relation, function, and convexity. Relational and object-relational databases use this model of space.

Topological space uses the basic notion of a neighborhood and points to formalize the extended object relationships such as boundary, interior, open, closed, within, connected, and overlaps, which are invariant under elastic deformation. Combinatorial topological space formalizes relationships such as Euler's formula ($\#faces + \#vertices - \#edges = 1$ for planar configuration). Network space is a form of topological space in which the connectivity property among nodes formalizes graph properties such as connectivity, iso-morphism, shortest-path, and planarity.

Euclidean coordinatized space uses the notion of a coordinate system to transform spatial properties and relationships to properties of tuples of real numbers. Metric spaces formalize the distance relationships using positive symmetric functions that obey the triangle inequality. Many multidimensional applications use euclidean coordinatized space with metrics such as distance.

2.2 Spatial Data Model

A spatial data model [9, 15] is a type of data-abstraction that hides the details of data-storage. There are two common models of spatial information: field-based and object-based. The field-based model treats spatial information such as altitude, rainfall and temperature as a collection of spatial functions transforming a space-partition to an attribute domain. The object-based model treats the information space as if it is populated by discrete, identifiable, spatially-referenced entities. The operations on spatial objects include *distance* and *boundary*. The operations on fields include local, focal, and zonal operations, as shown in Table 1. The fields may be continuous, differentiable, discrete, and isotropic or anisotropic, with positive or negative auto-correlation. Certain field operations (slope or interpolation) assume certain field properties (differentiable or positive auto-correlation).

An implementation of a spatial data model in the context of object-relational databases consists of a set of spatial data types and the operations on those types. Much work has been done over the last decade on the design of spatial data types, and consensus is slowly emerging via standardization efforts like OGIS [3]. Figure 3 illustrates the spatial data-type hierarchy of OGIS [3] which consists of points, curves, surfaces and solids. The spatial data types and operations can be embedded in an object-relational query language to implement a spatial query language.

2.3 Spatial Query Processing

The efficient processing of spatial queries requires both efficient representation and efficient algorithms. Common representations of spatial data in an object model include spaghetti, the node-arc-area(NAA) model, the doubly-connected-edge-list (DCEL), and boundary representation, some of which are shown in Figure 2 using entity-relationship diagrams. The NAA model differentiates between the topological concepts (node, arc, areas) and the embedding space (points, lines, areas). The spaghetti-ring and DCEL focus on the topological concepts. The representation of the field data model includes a regular tessellation (triangular, square, hexagonal grid), as well as triangular irregular networks (TIN).

The spatial queries[2] shown in Table 2 are often processed using *filter* and *refine* techniques. Approximate geometry such as the minimal orthogonal bounding rectangle of an extended spatial object is first used to quickly filter out many irrelevant objects. Exact geometry is then used for the remaining spatial objects to complete the processing. Strategies for range-queries

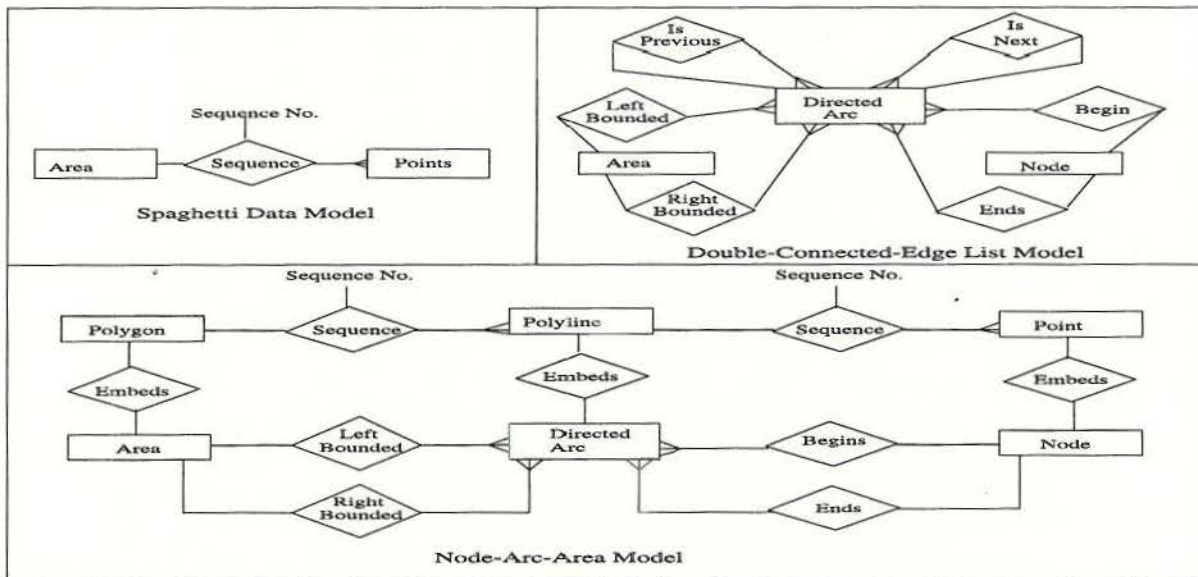


Figure 2: Entity Relationship Diagrams for Common Representations of Spatial Data

include a scan and index-search in conjunction with the plane-sweep algorithm. Strategies for the spatial-join include the nested loop, tree matching, and space partitioning. To speed up computation for large spatial objects (it is common for polygons to have 1000 edges), object indices are used in extended filtering. Strategies such as object approximation and tree matching originated in spatial-databases, and can potentially be applied in other domains.

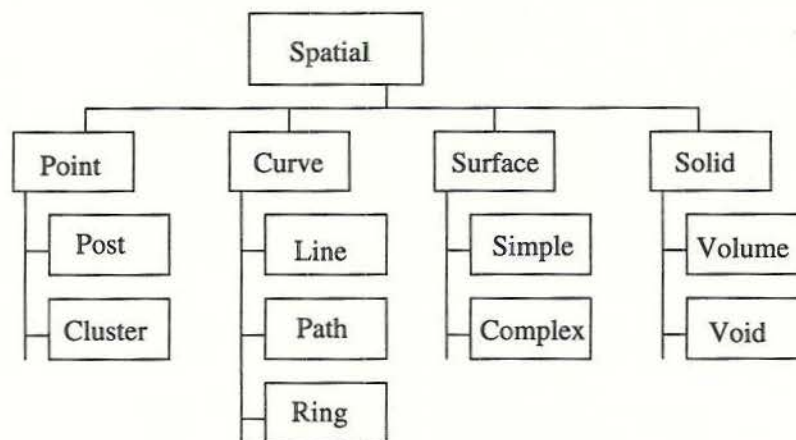


Figure 3: Spatial Data Type Hierarchy

Data model	Operator Group	Operation
Vector Object	Set-Oriented	equals, is a member of, is empty, is a subset of, is disjoint from, intersection, union, difference, cardinality
	Topological	boundary, interior, closure, meets, overlaps, is inside, covers, connected, components, extremes, is within
	Metric	distance, bearing/angle, length, area, perimeter, centroid
	Network	successors, ancestors, connected, shortest-path
	Dynamic	translate, rotate, scale, shear, split, merge
Raster field	Local	Point-wise sums, differences, maximums, means, etc
	Focal	slope, aspect, weighted average of neighborhood
	Zonal	sum or mean or maximum of field values in each zone

Table 1: A Sample of Spatial Operations

2.4 Spatial File Organization and Indices

The physical design of a spatial database optimizes the instructions to storage devices for performing common operations on spatial data files. File designs for secondary storage include clustering methods as well as spatial hashing methods. Simple spatial clustering is based on the notion of space-filling curves, which can be used to linearize a space, as illustrated in Figure 4. Metric clustering techniques use the notion of distance to group nearest neighbors together in a metric space. Topological clustering methods like connectivity-clustered access methods [10] use the min-cut partitioning of a graph representation to efficiently support graph traversal operations. The physical organization of files can be supplemented with indices, which are data-structures to improve the performance of search operations.

Classical one-dimensional indices such as the B^+ tree can be used for spatial data by linearizing a multi-dimensional space using a space-filling curve such as the Z-order (see Figure 4). A large number of spatial indices [8] have been explored for multi-dimensional euclidean space. Representative indices for point objects include Grid files, Point-Quad-Trees, and Kd-trees. Representative indices for extended objects include the R-tree family, the Field tree, Cell tree, BSP tree, and Balanced and Nested grid files.

Concurrency control for spatial access methods [7] is provided by the R-link tree, which is a variant of the R-tree with additional sibling pointers that allow tracking of modifications. Concurrency is provided during operations such as search, insert, and delete. The R-link tree




Figure 4: Space-Filling Curves to Linearize a Multidimensional Space

is also recoverable in a write-ahead logging environment.

Single Table Queries	
<i>Grouping</i>	Recode all land with silty soil to silt-loam soil
<i>Isolate</i>	Select all land owned by Steve Steiner
<i>Classify</i>	If the population density is less than 100 people / sq. mi., land is acceptable
<i>Scale</i>	Change all measurement to the metric system
<i>Rank</i>	If the road is an Interstate, assign it code 1; if the road is a state or US Highway, assign it code 2; otherwise assign it code 3
<i>Evaluate</i>	If the road code is 1, then assign it Interstate; if the road code is 2, then assign it Main Artery; if the road code is 3, assign it Local Road
<i>Rescale</i>	Apply a function to the population density
Multi-Table Queries	
<i>Attribute Join</i>	Join the Forest layer with the layer containing forest-cover codes
<i>Zonal</i>	Produce a new map showing state populations given county population
<i>Registration</i>	Align two layers to a common grid reference
<i>Spatial Join</i>	Overlay the land-use and vegetation layers to produce a new layer

Table 2: Typical Spatial Queries from GIS

2.5 Other Accomplishments

Spatial applications like NASA's Earth Observation System(EOS) have some of the largest data sets encountered in any application to date. This has prompted new research in database-file design for storage on tertiary storage devices such as juke-boxes. Representative results include those from the Sequoia 2000 project[12]. High-performance spatial applications such as flight simulators with geographic accuracy have triggered the development of new parallel formalizations for the range query and the spatial join query, including declustering methods

and dynamic-load balancing techniques for multi-dimensional spatial data[11]. Other interesting developments include hierarchical algorithms for shortest path computation [5] and view materialization.

3 Research Needs

Spatial databases are being used for an increasing number of new applications, such as Intelligent Transportation Systems, NASA's Earth Observation System, Multimedia Information Systems (MMIS) and Data Warehouses. This section lists representative research needs.

3.1 Space Taxonomy

Many spatial applications manipulate continuous spaces of different scales and with different levels of discretization. A sequence of operations on discretized data can lead to growing errors similar to the ones introduced by finite-precision arithmetic on numbers. There are preliminary results [4] on the use of discrete basis and bounding errors with peg-board semantics. Another related problem concerns interpolation to estimate the continuous field from a discretization. Negative spatial auto-correlation makes interpolation error-prone. Further work is needed on frameworks to formalize the discretization process, its associated errors, and interpolation.

3.2 Spatial Data Model

Spatial data models have been developed for topological, metric and coordinatized euclidean space. There has been very little work towards developing data models, data types (e.g. node, edge, path), and a kernel set of operations (e.g. get-successors, shortest path) for network space, despite their critical role in applications like transportation and utility management (telephone, gas, electric).

Similarly, there is a need for developing the field data model towards a field-based query language. Operations on fields will be needed to help derive new information such as land-cover classification; the fields involved include temperature, texture, and water content, and are obtained through imaging in different bands such as other infrared, visible bands, or microwave.

3.3 Spatial Query Processing

Many open research areas exist at the logical level of query processing, including query-cost modeling, strategies for bulk loading as well as queries related to fields and networks. Cost models are used to rank and select the promising processing strategies, given a spatial query and a spatial data set. Traditional cost models may not be accurate in estimating the cost of strategies for spatial operations, due to the distance metric as well as the semantic gap between relational operators and spatial operation. Cost models are needed to estimate the selectivity of spatial search and join operations towards comparison of execution-costs of alternative processing strategies for spatial operations during query optimization. Preliminary work in the context of the R-tree, tree-matching join, and fractal-models is promising, but more work is needed.

Many processing strategies have been developed for range queries and spatial join queries using the overlap predicate. However, there is a need to develop and evaluate strategies for many other frequent queries such as those in Table 3. These include queries on objects using predicates other than overlap, queries on fields such as slope analysis as well as queries on networks such as shortest path to a set of destinations. Bulk loading strategies for spatial data also need further study.

Buffer	Find the areas 500 ft. from power lines
Voronize	Classify households as to which supermarket they are closest to
Neighborhood	Determine slope based on elevation
Network	Find the shortest path from the warehouse to all delivery stops
Allocation	Where is the best place to build a new restaurant
Transformation	Triangulate a layer based on elevation
Bulk Load	Load a spatial data file into the database
Raster ↔ Vector	Convert between raster and vector representations

Table 3: Difficult Spatial Queries from GIS

3.4 Spatial File Organization and Indices: Physical Level

Many file organizations and indices with distance metrics have been developed for coordinatized euclidean space. However, little work has been done on file clustering and indices for network spaces such as road-maps and telephone networks. Further work is needed, both to characterize the access patterns of the graph algorithms that underlie network operations and to design access methods.

The R-link tree is among the few approaches available for concurrency control on the R-tree. New approaches are needed for concurrency-control techniques for other spatial indices. The data volume of emerging spatial applications such as NASA's EOS is among the highest of any database application. Sequoia 2000 provides an approach towards tertiary storage files and indices. Other approaches based to manage databases on tertiary storage need to be investigated.

3.5 Other

Other research needs include benchmarking, work-flow modeling, and the visual presentation of results. The Sequoia 2000 [12] benchmark characterizes the data and queries in Earth Science applications. The performance of loading data, raster queries, spatial selection, spatial joins, and recursion is addressed in 11 benchmark queries. Similar benchmarks are needed to characterize the spatial data management needs of other applications such as GIS, DWH, and transportation.

The work-flow in some spatial applications such as GIS is based on manipulating layers to produce new, derived layers. Typically, the layers are combined in a tree-based manner, starting with a large number of source layers and producing new layers until a final result layer is produced. Information about dependence among layers is useful for change propagation if the source layers are modified.

Spatial databases may require a different type of concurrency support than needed by traditional databases. For example, transactions in traditional systems tend to be short (on the order of seconds). However, in spatial databases, these transactions can last up to a couple of hours for editing and browsing. Similarly, recovery and backup issues may also change, as the spatial objects tend to be large (a few megabytes) when compared to their counterparts in traditional systems. There is a need to characterize the work flow of spatial applications.

Many spatial applications present results visually in the form of maps which consist of graphic images, 3D displays, and animations. They also allow users to query the visual representation by pointing to the visual representation using devices like a mouse or a pen. Further work is needed to explore the impact on database performance of both querying by pointing and visually presenting the results.

4 Summary and Discussion

Table 4 shows statistics of recent research publications in four representative conferences including the Symposium on Spatial Databases, ACM SIGMOD, VLDB, and IEEE Data Engineering. While publication count¹ is not always the most informative measure of research accomplishments or needs, it does provide an overview of research interests. This table brings out the fact that many research needs identified in this paper have very few publications. These research needs include spatial graphs, field based modeling, cost modeling and concurrency control for query processing techniques, and discretization and error propagation. Similarly, many of the major accomplishments identified in this survey have a large number of publications and have seen standardization and industrial use. These accomplishments include object-based data modeling, spatial data types, filter and refine techniques for query processing, and spatial indexing.

Many of these techniques are being used in an increasing number of spatial (multi-dimensional) applications such as GIS, CAD, and EOS. We believe that other emerging multi-dimensional applications such as multimedia information systems will use these methods to solve problems such as searching and indexing spatial content. We illustrate the possibilities in the context of multi-media information systems with text, audio and video data over the world-wide-web.

Multimedia data has a spatial content which can be queried using the same spatial operators that have become popular in geographic information systems. The spatial operator *inside of* can be applied to text to locate sentences that contain the word "multimedia". Audio is often broken into channels with each channel containing input from a different source; for instance, trumpet, guitar, and voice. These channels are analogous to layers in GIS and can be manipulated similarly. A spatial join could determine all of the locations where the input from both piano and voice is over a certain decibel threshold.

Web data is also highly spatial. Hyper-links connect different web sites together in the same way that spatial graph databases store transportation and utility data. Index servers could use spatial graph operators to answer queries such as: "What is the shortest hyper-path from my homepage to a web page containing a picture of a red bicycle, where the picture is to the left of the text".

A video database such as a movie server can take advantage of techniques developed for

¹A complete bibliography can be found at <http://www.cs.umn.edu/research/shashi-group>

Subject	Category	Publication Volume
Spatial Data Model & Query Language	Discretization, Error Control	*
	Euclidean Coordinatized Space	***
	Topological Space	*
	Network Space	*
	Field Space	*
Spatial Query Processing	Representation	*
	Strategies for Range Query	**
	Strategies for Spatial Join	***
	Strategies for Field Operations	**
	Strategies for Network Operations	*
	Cost Modeling	*
Spatial File Organization & Indices	Euclidean Coordinatized Space	****
	Network Space	*
	Topological Space	*
	Concurrency Control	*
Others	Parallel Processing	**
	Distributed Spatial Data	*
	Benchmark	*
	Work flow, Transaction	*
	Visual Presentation	**
	Data Mining, Similarity Search	*
	Spatio-temporal	*
	Architecture	**

Table 4: Recent Research in Spatial Databases

spatial databases. Consider the movie *Toy Story*: each frame contains spatial content with objects interacting in topological relationships. For instance, Buzz Lightyear could be above the trees when he is flying and frames in the movie could be queried based on those relationships. For example, if you cannot remember when in the movie an important event occurred, but you can remember that Buzz Lightyear was in front of a tree, you would be able to query the movie using that relationship to determine when in the movie that event took place. Such queries exploit the topological relationships inherent in all tangible objects.

Acknowledgments

This work is sponsored in part by the Army High Performance Computing Research Center under the auspices of the Department of the Army, Army Research Laboratory cooperative agreement number DAAH04-95-2-0003/contract number DAAH04-95-C-0008, the content of which does not necessarily reflect the position or the policy of the government, and no official endorsement should be inferred. This work was also supported in part by NSF grant #9631539. Thanks to Professor Jaideep Srivastava for technical commentary and to Christiane McCarthy for helping to improve the readability of the paper.

References

- [1] D. Chamberlin. *Using the New DB2: IBM's Object Relational System*. Morgan Kaufmann, 1997.
- [2] N. Chrisman. *Exploring Geographic Information Systems*. John Wiley and Sons, 1997.
- [3] K. Gardels. Open GIS and On-Line Environmental Libraries. *ACM SIGMOD Record*, 26(1):32-38, March 1997.
- [4] R.H. Gutting. An Introduction to Spatial Database Systems. *VLDB*, 3:357-399, 1994.
- [5] N. Jing, Y. Huang, and E. Rudensteiner. Hierarchical Optimization of Optimal Path Finding for Transportation Applications. In *Proc. Intl. Conf. on Information and Knowledge Management*. ACM, November 1996.
- [6] W. Kim, J. Garza, and A. Kesin. Spatial Data Management in Database Systems. pages 1-13. 3rd Intl. Symp. on Advances in Spatial Databases, Springer Verlag, 1993.
- [7] M. Kornacker and D. Banks. High-Concurrency Locking in R-Trees. In *Proc. Intl. Conf. on VLDB*, 1995.
- [8] H. Samet. *The Design and Analysis of Spatial Data Structures*. Addison-Wesley, 1990.
- [9] S. Shekhar, M. Coyle, D-R. Liu, B. Goyal, and S. Sarkar. Data Models in Geographic Information Systems. *Communication of the ACM*, 40(4), 1997.
- [10] S. Shekhar and D-R. Liu. A Connectivity-Clustered Access Method for Networks and Network Computation. *IEEE TKDE*, 9(1):102-, January 1997.
- [11] S. Shekhar, S. Ravada, V. Kumar, D. Chubb, and G. Turner. Parallelizing a GIS on a Shared Address Space Architecture. *IEEE Computer*, 29(12), December 1996.
- [12] M. Stonebreaker, J. Frew, and J. Dozier. The Sequoia 2000 Storage Benchmark. In *SIGMOD*. ACM, 1993.
- [13] M. Stonebreaker and G. Kennitz. POSTGRES Next-Generation Database Management System. *Communication of the ACM*, 34(10):78-92, 1993.

- [14] M. Stonebreaker and D. Moore. *Object Relational DBMSs: The Next Great Wave*. Morgan Kaufmann, 1997.
- [15] M.F. Worboys. *GIS: A Computing Perspective*. Taylor and Francis, 1995.