National Science Foundation

Project Title: Efficient Indexing for Spatiotemporal Applications

Vassilis J. Tsotras, PI
Dimitrios Gunopulos, co-PI
Department of Computer Science and Engineering
University of California, Riverside

Contact Information:

Vassilis J. Tsotras
Dimitrios Gunopulos
Department of Computer Science and Engineering
University of California, Riverside, CA 92521
Phone: (909) 787-2888, Fax : (909) 787-4643
Email: tsotras@cs.ucr.edu, http://www.cs.ucr.edu/~tsotras
dg@cs.ucr.edu, http://www.cs.ucr.edu/~dg

WWW Page: http://www.cs.ucr.edu/~dg/spatiotemporal.html

List of Supported Students:

Marios Hadjieleftheriou, Department of Computer Science and Engineering, University of California, Riverside, CA 92521
Dimitrios Papadopoulos, Department of Computer Science and Engineering, University of California, Riverside, CA 92521

Project Award Information:

Award Number: IIS-9907477. Duration: Oct. 1, 1999-Sept.30, 2002
Title: Efficient Indexing for Spatiotemporal Applications

Keywords: spatiotemporal databases, access methods, moving points, neighbor queries.

Project Summary:

Indexing spatiotemporal data is an important problem for many applications (global change, transportation, social and multimedia applications). The goal of this project is to provide efficient access methods for data whose geometry changes over time. Two time-varying spatial attributes are considered, the object position and extent. Based on the rate by which these spatial attributes change, the discrete and continuous spatiotemporal environments are identified. In the discrete environment, spatiotemporal data changes in discrete steps. Efficient ways to answer historical queries on any past state of such spatiotemporal data are examined. In particular, selection, neighbor, aggregate, join and similarity queries are addressed using a "partial persistence" methodology. In the continuous spatiotemporal environment, data changes continuously. Instead of keeping the data position/extent at discrete times (which would result in enormous update/storage requirements) the functions by which this data changes are stored. This introduces the novel problem of indexing functions. Using this approach, selection, neighbor and aggregation queries about future locations of moving objects in one and two dimensions are addressed. The methods used in this project are expected to achieve at least 30% improvement over traditional access methods. The applicability of the completed work reaches multiple settings, including Geographic Information Systems, multimedia databases and transportation systems.

Publications and Products:

Journal/Conference publications:

[1] D. Zhang, A. Markowetz, V. Tsotras, D. Gunopulos, B. Seeger: "Efficient Computation of Temporal Aggregates with Range Predicates," in PODS 2001.
[2] C. Domeniconi, J. Peng & D. Gunopulos: "An adaptive Metric Machine for Pattern Classification," in NIPS 2000.
[3] D. Gunopulos, G. Kollios, V. Tsotras, C. Domeniconi: "Using kernels to approximate multi-dimensional aggregate range queries over real attributes.", Workshop on New Perspectives in Kernel-Based Learning Methods, NIPS 2000.
[4] Shu-Yao Chien, Vassilis J. Tsotras, Carlo Zaniolo: "Version Management of XML Documents: Copy-Based versus Edit-Based Schemes." In RIDE-DM 2001.

[5] Shu-Yao Chien, Vassilis J. Tsotras, Carlo Zaniolo: "Version Management of XML Documents". In WebDB (Informal Proceedings) 2000: 75-80.
[6] C. Domeniconi, J. Peng and D. Gunopulos: "Adaptive Metric Nearest Neighbor Classification." In Computer Vision & Pattern Recognition 2000.
[7] George Kollios, Dimitrios Gunopulos, Vassilis J. Tsotras: "Nearest Neighbor Queries in a Mobile Environment", in Proceedings of Spatio-Temporal Database Management International Workshop, STDBM'99, Edinburgh, Scotland, September 10-11, 1999, LNCS 1678, pp: 119-134.
[8] Dimitrios Gunopulos, George Kollios, Vassilis J. Tsotras: "All-Pairs Nearest Neighbors in a Mobile Environment", 7th Hellenic Conference on Informatics, Ioannina, Greece, August 26-29, 1999. Proceedings appear in *Advances in Informatics*, edited by D.I. Fotiadis, S.D. Nikolopoulos, World Scientific Publishers, 1999.
[9] George Kollios, Dimitrios Gunopulos, Vassilis Tsotras: "Indexing Animated Objects", in Proceedings of 5th International Workshop on Multimedia Information Systems, MIS '99, Indian Wells, CA, October 21-23, 1999.
[10] George Kollios, Dimitrios Gunopulos, Vassilis Tsotras, Alex Delis, Marios Hadjieleftheriou: "Indexing Animated Objects Using Spatiotemporal Access Methods", to appear in IEEE Transactions of Knowledge and Data Engineering.
[11] Dimitrios Gunopulos, George Kollios, Vassilis J. Tsotras, Carlotta Domeniconi: "Approximating Multi-Dimensional Aggregate Range Queries over real Attributes", in 2000 ACM SIGMOD , Texas, May 14-19, 2000.

Book:
Yannis Manolopoulos, Yannis Theodoridis, Vassilis J. Tsotras, *Advanced Database Indexing*, Kluwer Academic Publishers, Boston, November 1999, 312 pages, ISBN 0-7923-7716-8.

Project Impact:
Dr. Kollios graduated with a Ph.D. degree in Summer 2000, and is currently an Assistant Professor in the Dept. of Computer Science, Boston University. Prof. Tsotras was his advisor. There are two Research Assistants currently working on this project, Mr Hadjieleftheriou and Mr. Papadopoulos. Both are 2nd year Ph.D. students in the Univ. of California, Riverside. Parts of this research are used in a Database Graduate Seminar course (CS267) taught at the University of California, Riverside (Winter 2000). For this project we have established a cooperation with ESRI (Environmental Systems Research Institute), the prime provider of GIS. Our contact at ESRI is Dr. Erik Hoel. Through this cooperation we expect to have access to real spatiotemporal datasets/applications. We also cooperate with the Department of Defense.

Goals, Objectives, and Targeted Activities:
The objective of this proposal is to design efficient access methods for addressing spatiotemporal queries. A spatiotemporal query specifies spatial/temporal predicates and retrieves all objects that satisfy them. A spatial predicate is defined in terms of a point or an extent while a temporal predicate can involve a time instant or a time interval. In particular we are interested in:
(a) selection queries: "find all objects contained in a given area *S* at a given time *t*,"
(b) neighbor queries: "find which object became the closest to a given point *s* during time interval *T*," or, "find the 5 closest ambulances to an accident position in the next 10 minutes,"
(c) aggregate queries: "find how many objects passed through area *S* during time interval *T*," or, "find the fastest object that will pass through area *S* in the next 5 minutes from now,"
(d) join queries: "given two spatiotemporal relations *R*1 and *R*2, find pairs of objects whose extents intersected during the time interval *T*," or "find pairs of planes that will come closer than 1 mile in the next 5 minutes,"
(e) similarity queries: "given an area *S* find the time instants when there were more than 10 objects in *S*", or, "find objects that moved similarly to the movement of a given object *x* over an interval *T*."

During the first year of this project we have worked mainly on selection, neighbor and join queries. In particular, reference [7] above examines how to efficeintly support neighbor queries in a database of moving objects. Reference [8] addresses the all-pairs nearest neighbor problem which is also related to spatiotemporal joins. In [9] and [10] we provide efficient indexing for selection queries. The application comes from an animated movie, which can also be seen as a spatiotemporal evolution (where the frame sequence corresponds to time evolution). Many of the spatiotemporal applications use data whose values range over the domain of reals. In [11] we examine ways to approximate range queries over attributes from real domains. Recently we worked on aggregate and similarity spatiotemporal queries, as well as more complex selection and join queries. In [1] we consider aggregate spatiotemporal queries. [2], [3] and [6] present extentions on classification algorithms and data approximation techniques that take advantage of the spatial properties of the data. [4] and [5] consider the problem of efficiently maintaining the structure of documents that changes over time.

Project References:

This is a new project; data and publications will be posted as they become available through the project's web page.

Area Background:

A spatiotemporal database system manages data whose geometry changes over time. There are many applications that create spatiotemporal data, including global change (as in climate or land cover changes), transportation (traffic surveillance data, intelligent transportation systems), social (demographic, health, etc.), and multimedia (animated movies) applications. An important temporal feature of spatiotemporal data is the rate by which spatial attributes change. We distinguish between *discrete* and *continuous* changes. Traditional databases assume that data stored in the database remains constant until explicitly modified through an update. For example, if a price field is $5, it remains $5 until explicitly updated. While this model serves well many applications where data changes in discrete steps, it is not appropriate for applications with continuously changing data (the dynamic attributes of [SWCD97, WXCJ98]). Consider a database keeping the position of moving objects (like automobiles). The primary goal of this database is to correctly represent the real world while objects move. Continuous updating about each object's position leads to serious performance overhead. Updating the database only at given time instants limits query accuracy. A better approach would be to represent the position of each moving object as a function of time; then object positions change as time proceeds without the need of explicit updates.

While spatiotemporal applications and data abound the field has only recently attracted the efforts of database researchers. While many temporal [ST99] and spatial indexes exist, very few works have addressed the combination [SY91, TOW98, VTS98, KGT99]. Most of research has concentrated to spatiotemporal database models and query languages [EGSV98, CR99, W94, E93, SWCD97, WXCJ98]. [BGH97] examines spatiotemporal structures for main memory. Due to the temporal component, spatiotemporal databases need to manage large amounts of data accumulated over long periods of time (historical data). It is thus important to develop efficient access methods (indices) to access such databases.

Area References:

[BGH97] J. Basch, L. Guibas and J. Hershberger, "Data Structures for Mobile Data", Proc. ACM-SIAM SODA, pp. 747-756, Jan. 1997.

[CR99] J. Chomicki, P. Revesz, "A Geometric Framework for Specifying Spatiotemporal Objects", Proc. 6th International Workshop on Time Representation and Reasoning, May 1999.

[E93] M.J. Egenhofer, "What's Special about Spatial Database Requirements for Vehicle Navigation in Geographic Space", Proc. SIGMOD Conference, 398-402 1993.

[EGSV98] M. Erwig, R.H. Guting, M. Schneider and M. Vazirgianis, "Spatio-temporal Data Types: An Approach to Modeling and Querying Moving Objects in Databases", ACM GIS Symposium, pp. 131-136, 1998.

[KGT99] G. Kollios, D. Gunopulos, V.J. Tsotras, "On Indexing Mobile Objects", in Proc. ACM PODS, 1999.

[ST99] B. Salzberg, V.J. Tsotras, "A Comparison of Access Methods for Time-Evolving Data", ACM Computing Surveys, June 1999.

[SY91] S. Shekhar and T.A. Yang, "Motion in a Geographical Database System", Proc. of 2nd SSD, pp. 339-357, 1991.

[SWCD97] A. P. Sistla, O. Wolfson, S. Chamberlain, S. Dao, "Modeling and Querying Moving Objects", Proc. IEEE ICDE, pp. 422-432, Apr. 1997.

[TOW98] J. Tayeb, O. Ulusoy, O. Wolfson, "A Quadtree-Based Dynamic Attribute Indexing Method", The Computer Journal, Vol. 41, No. 3, pp. 185-200, 1998.

[VTS98] M. Vazirgiannis, Y.Theodoridis, T.K. Sellis, "Spatio-Temporal Composition and Indexing for Large Multimedia Applications", Multimedia Systems, 6(4): 284-298, 1998.

[W94] M. Worboys, "A Unified Model for Spatial and Temporal Information", The Computer Journal, 37(1): 36-34, 1994.

[WXCJ98] O. Wolfson, B. Xu, S. Chamberlain, L. Jiang, "Moving Objects Databases: Issues and Solutions", Proc. SSDBM, pp. 111-122, Jul. 1998.

Potential Related Projects:

We would like to cooperate with colleagues working on spatial, geospatial, temporal and pictorial databases.