

Continuous Query Processing of Spatio-temporal Data Streams in PLACE

Mohamed F. Mokbel, Xiaopeng Xiong, Moustafa A. Hammad¹, and Walid G. Aref *

*Department of Computer Science, Purdue University, West Lafayette, IN
E-mail:{mokbel,xxiong,aref}@cs.purdue.edu*

¹*Department of Computer Science, University of Calgary, Calgary, Alberta
E-mail:hammad@cpsc.ucalgary.ca*

Abstract. The tremendous increase in the use of cellular phones, GPS-like devices, and RFIDs results in highly dynamic environments where objects as well as queries are continuously moving. In this paper, we present a continuous query processor designed specifically for highly dynamic environments (e.g., location-aware environments). We implemented the proposed continuous query processor inside the PLACE server (Pervasive Location-Aware Computing Environments); a scalable location-aware database server developed at Purdue University. The PLACE server extends data streaming management systems to support location-aware environments. These environments are characterized by the wide variety of continuous spatio-temporal queries and the unbounded spatio-temporal streams. The proposed continuous query processor includes: (1) New *incremental* spatio-temporal operators to support a wide variety of continuous spatio-temporal queries, (2) Extended semantics of sliding window queries to deal with spatial sliding windows as well as temporal sliding windows, and (3) A shared-execution framework for scalable execution of a set of concurrent continuous spatio-temporal queries. Experimental evaluation shows promising performance of the continuous query processor of the PLACE server.

Keywords: Spatio-temporal databases, continuous queries, data stream management systems, location-aware services.

1. Introduction

The rapid increase in spatio-temporal applications calls for new query processing techniques to deal with both the spatial and temporal domains. Examples of these applications include location-aware services, traffic monitoring, and enhanced 911 service. Such applications continuously receive data from mobile objects (e.g., moving vehicles in road networks). The streaming nature of incoming spatio-temporal data poses new challenges that require combining and/or modifying the recent advances in both spatio-temporal database systems and data stream management systems.

* This work was supported in part by the National Science Foundation under Grants IIS-0093116, IIS-0209120, and 0010044-CCR.



Recent research efforts for continuous spatio-temporal query processing, e.g., see (Hadjieleftheriou et al., 2003; Iwerks et al., 2003; Kwon et al., 2002; Lazaridis et al., 2002; Lee et al., 2003; Mokbel et al., 2004a; Prabhakar et al., 2002; Saltenis et al., 2000; Tao et al., 2002; Tao et al., 2003a) rely mainly on the ability of storing and indexing spatio-temporal data. Given the dynamic environment of spatio-temporal applications, the main idea is to modify traditional data indices to support frequent updates. Examples of these indices include modified grid structures, e.g., (Gedik and Liu, 2004; Mokbel et al., 2004a), modified B-trees, e.g., (Jensen et al., 2004), modified R-trees, e.g., (Kwon et al., 2002; Lee et al., 2003), and time-parameterized R-trees, e.g., (Saltenis and Jensen, 2002; Saltenis et al., 2000; Tao et al., 2003a). Although these indexing schemes give better support for updates than their counterpart traditional indices, issues of high arrival rates of both objects and queries, infinite source of data, and spatio-temporal streams are overlooked by these approaches. With the notion of spatio-temporal streams, only in-memory algorithms for continuous queries need to be realized.

Numerous research efforts are devoted to stream query processing (e.g., see (Arasu et al., 2004; Abadi et al., 2003; Chandrasekaran et al., 2003; Cranor et al., 2003; Hammad et al., 2004; Motwani et al., 2003; Yao and Gehrke, 2002)). The main focus is to provide the ability to process incoming data streams online against a set of outstanding and continuous queries. However, the spatial and temporal properties of both data and queries are overlooked. A spatio-temporal data stream distinguishes itself from a traditional data stream in the following: (1) Queries as well as data have the ability to change their locations continuously. Thus, the arrival of a new data item (e.g., the location) of an object, say p , at some time t_2 ($t_2 > t_1$) may result in expiring the previous location information of p at time t_1 (*predicate-based* sliding window). This is in contrast to traditional data streams where data is expired only as it becomes old in the system (time-based expiration in sliding window queries). (2) An object may be added to or removed from the answer set of a spatio-temporal query (*positive* and *negative* answers). For example, consider a set of moving vehicles that move in and out of a certain range query. Thus the query answer may be represented *progressively* by a sequence of *positive* and *negative* updates. This is in contrast to traditional queries where only an addition to the query answer is permitted. (3) Due to the mobility of both objects and queries, any delay in processing spatio-temporal queries may result in an obsolete answer. Consider a query that asks about moving objects that lie in a certain region. If the query answer is delayed, the answer may be outdated where objects are continuously changing their loca-

tions. These distinguishing characteristics of spatio-temporal streams require revisiting traditional data stream management systems to have special handling of spatio-temporal streams.

In this paper, we present the PLACE server (**P**ervasive **L**ocation-**A**ware **C**omputing **E**nvironments) (Aref et al., 2003; Mokbel et al., 2004b); a scalable location-aware database server currently being developed at Purdue University. The PLACE server extends the Nile (Hammad et al., 2004) data stream management system to support continuous query processing of spatio-temporal streams. The PLACE server aims to bridge the areas of spatio-temporal databases and data stream management systems. The main idea is to furnish traditional data stream management systems with the basic functionalities that support processing incoming spatio-temporal streams against a set of outstanding continuous spatio-temporal queries. In particular, the continuous query processor of the PLACE server (Mokbel et al., 2004c) has the following distinguishing characteristics:

1. **Predicate-based window queries:** The PLACE continuous query processor extends the processing of continuous sliding window queries beyond time-based and tuple-count windows to accommodate for the so called *predicate-based* window queries. In *predicate-based* window queries, objects are qualified to be part of the window once they satisfy a certain query predicate. Similarly, objects are expired only when they no longer satisfy a certain predicate. *Predicate-based* windows are a generalization of time-based and tuple-count sliding windows.
2. **Incremental evaluation.** The PLACE continuous query processor employs an *incremental* evaluation paradigm by continuously updating the query answer. We distinguish between two types of updates; namely *positive* and *negative* updates (Mokbel et al., 2004a). A positive/negative update indicates that a certain object needs to be added to/removed from the query answer.
3. **Spatio-temporal operators.** The PLACE continuous query processor goes beyond the idea of implementing high level algorithms for continuous spatio-temporal queries. Instead, the PLACE server encapsulates the spatio-temporal query algorithms into a set of primitive spatio-temporal pipelined operators (e.g., INSIDE and k NN operators) that can be part of a larger query plan. Having a set of primitive spatio-temporal operators results in supporting a wide variety of continuous spatio-temporal queries and in having flexible query optimizers where multiple candidate query plans can be produced.

4. **Scalability.** We use a *shared-execution* paradigm as a means of achieving scalability in terms of the number of outstanding continuous spatio-temporal queries.

The rest of the paper is organized as follows. Section 2 highlights the challenges in realizing the continuous query processor of the PLACE server along with the related work of each challenge. In Section 3, we present an overview of the data model and SQL language used by the PLACE server. Section 4 presents *predicate-based* window queries. The *incremental* evaluation of the PLACE server is discussed in Section 5. The scalability in terms of the number of outstanding spatio-temporal queries is addressed in Section 6. The GUI interfaces for the PLACE server and PLACE clients are presented in Section 7. Section 8 presents experimental results that evaluate the performance of the PLACE server. Section 9 highlights the basic features of the Nile data stream query processor in which the PLACE continuous query processor is based on. Finally, Section 10 concludes the paper.

2. Challenges and Related Work

In this section, we highlight some challenges in realizing the continuous query processor of the PLACE location-aware server. With each challenge, we summarize the related work in the areas of spatio-temporal databases and data streams. Then, we highlight briefly how the PLACE continuous query processor deals with these challenges.

2.1. CHALLENGE I: MASSIVE SIZE OF INCOMING SPATIO-TEMPORAL DATA STREAMS

Spatio-temporal databases. Existing continuous query processors for spatio-temporal databases assume explicitly that all incoming data can be stored in secondary storage. A wide variety of spatio-temporal access methods (e.g., see (Mokbel et al., 2003b) for a survey) has been introduced to deal with massive sizes of spatio-temporal data. However, with the streaming input, only in-memory algorithms are feasible. There is limited work that exploits the spatial and/or temporal properties of data streams. The spatial properties of data streams are addressed recently in (Cormode and Muthukrishnan, 2003; Hershberger and Suri, 2004) to solve geometric problems, e.g., computing the convex hull (Hershberger and Suri, 2004). In (Sun et al., 2004), spatio-temporal histograms are used as synopses for approximate query processing on spatio-temporal data streams. Up to the authors' knowledge, there

is no existing work that addresses continuous query processing for spatio-temporal streams.

Data stream management systems. A common challenge for all data stream management systems is the infinite size of the incoming data streams. With the inability to store all the incoming data into memory, data stream management systems tend to store only data that is of interest to any outstanding continuous query. Once a stored data becomes out of interest of all outstanding queries, it is expired from the memory leaving its space to a more important incoming tuple. One approach of expiring in-memory data is to use *punctuation* (Tucker et al., 2003). A *punctuated* tuple indicates the expiration of a certain set of stored tuples. To decide whether a certain incoming tuple is important or not, every continuous query is associated with a *historical* window that limits the important tuples to the most recent ones. A *historical* window could be time-based (e.g., the last 1 hour) or tuple-based (e.g., the last 100 tuples). Such queries are termed *sliding-window* queries (e.g., see (Arasu et al., 2004; Babu and Widom, 2001; Chandrasekaran et al., 2003; Golab et al., 2004; Hammad et al., 2003b; Kang et al., 2003; Madden et al., 2002; Srivastava and Widom, 2004)). In *sliding-window* queries, tuples are dropped (expired) from the system in a first-in-first-expire policy. In case of large window sizes, *load shedding* techniques, e.g., (Babcock et al., 2003; Babcock et al., 2004; Tatbul et al., 2003) are utilized to drop some tuples from memory.

The PLACE approach. Trying to deploy the idea of *sliding-window* queries from traditional data streams would result in limiting the functionality of the PLACE server. A wide variety of continuous spatio-temporal queries are considered as NOW queries (Clifford et al., 1997) (i.e., no historical information is needed). Thus, in the PLACE server, we generalize the time-based and tuple-based sliding window queries to *predicate-based* window queries. A tuple is considered important if it satisfies at least one query predicate from all outstanding continuous queries. Once a tuple no longer satisfies any query predicate, it is expired from the server. This implies that tuple expiration is predicate-based rather than first-in-first-expire as in the commonly used time-based or tuple-count data streams. Expiration in *predicate-based* window queries is different from *punctuated* streams (Tucker et al., 2003) in that tuples in *predicate-based* windows are expired one at a time, rather than expiring a set of tuples using a certain *punctuated* tuple.

2.2. CHALLENGE II: CONTINUOUS EVALUATION OF CONTINUOUS QUERIES

Spatio-temporal databases. Most of the existing techniques in spatio-temporal databases abstract the continuous query into a series of snapshot queries executed at regular interval times. Mainly, three different approaches are investigated: (1) The validity of the results (Zhang et al., 2003; Zheng and Lee, 2001). With each query answer, the server returns a *valid time* (Zheng and Lee, 2001) or a *valid region* (Zhang et al., 2003) of the answer. Once the valid time is expired or the client goes out of the valid region, the client resubmits the continuous query for reevaluation. (2) Caching the results. The main idea is to cache the previous result either in the client side (Song and Roussopoulos, 2001) or in the server side (Lazaridis et al., 2002). Previously cached results are used to prune the search for the new results of k -nearest-neighbor queries (Song and Roussopoulos, 2001) and range queries (Lazaridis et al., 2002). (3) Precomputing the result (Lazaridis et al., 2002; Tao et al., 2002). If the trajectory of query movement is known a priori, then by using computational geometry for stationary objects (Tao et al., 2002) or velocity information for moving objects (Lazaridis et al., 2002), we can identify which objects will be nearest-neighbors (Tao et al., 2002) to or within a range (Lazaridis et al., 2002) from the query trajectory. If the trajectory information changes, then the query needs to be reevaluated. Up to the authors' knowledge, our earlier work, SINA (Mokbel et al., 2004a), is the only work that addresses incremental evaluation of continuous queries in spatio-temporal databases.

Data stream management systems. Due to the newly introduced *sliding-window* queries, several algorithms are proposed for each query operator, e.g., window join (Golab et al., 2004; Hammad et al., 2003a; Hammad et al., 2003b; Kang et al., 2003; Srivastava and Widom, 2004) and window aggregates (Arasu and Widom, 2004; Datar et al., 2002). All window algorithms for traditional data streaming provide progressive updates of the query answer.

The PLACE approach. A distinguished characteristic of spatio-temporal streams is that we need to have the ability to *remove* some parts of the query answer (e.g., an object moves out of the range query). This feature is not available in traditional data streams where the query answer is append-only. In the PLACE server, we apply a progressive evaluation paradigm by extending the ideas of SINA (Mokbel et al., 2004a) to be applicable to spatio-temporal streams rather than being tied to disk storage.

2.3. CHALLENGE III: WIDE VARIETY OF CONTINUOUS QUERY TYPES

Spatio-temporal databases. A major challenge for spatio-temporal query processors is the wide variety of spatio-temporal query types (e.g., see (Mokbel et al., 2003a) for a thorough classification of spatio-temporal queries). Thus, it becomes a difficult task to provide a database system with the ability to support all kinds of spatio-temporal queries. The DOMINO database system (Wolfson et al., 2002; Wolfson et al., 1999) is the first attempt to build a database of moving objects on top of existing DBMSs (Wolfson et al., 1998). One main focus of DOMINO is to support new kinds of spatio-temporal attributes and query language for moving objects. However, the query processing issues are not addressed. Other than DOMINO, most of the existing query processing techniques focus on solving special cases of continuous spatio-temporal queries, e.g., (Song and Roussopoulos, 2001; Tao et al., 2002; Zhang et al., 2003; Zheng and Lee, 2001) focus on *moving* queries on *stationary* objects while (Cai et al., 2004; Gedik and Liu, 2004; Hadjieleftheriou et al., 2003; Prabhakar et al., 2002) focus on *stationary* range queries on *moving* objects. Other work focuses on aggregate queries (e.g., see (Hadjieleftheriou et al., 2003; Sun et al., 2004; Tao et al., 2004)), k -NN queries (e.g., see (Iwerks et al., 2003; Song and Roussopoulos, 2001)) and reverse nearest-neighbor queries (Benetis et al., 2002).

Data stream management systems. Existing data stream management systems (e.g., Aurora (Abadi et al., 2003), Telegraph (Chandrasekaran et al., 2003), Gigascope (Cranor et al., 2003), Nile (Hammad et al., 2004), STREAM (Motwani et al., 2003), and Cougar (Yao and Gehrke, 2002)) provide new algorithms for (almost) all traditional query operators. However, there is no special handling and/or optimization for spatio-temporal queries.

The PLACE approach. In the PLACE server, we go beyond the idea of having tailored high-level algorithms for each specific spatio-temporal query. Instead, we furnish existing data stream management systems by a set of primitive spatio-temporal pipeline operators (e.g., the **INSIDE** and **kNN** operators). Spatio-temporal operators are combined with traditional streaming operators to provide the ability of having complex query plans that represent a wide variety of continuous spatio-temporal queries. In addition, the PLACE server provides a uniform framework that is applicable to all mutability combinations of objects and queries, e.g., *moving* queries on *stationary* objects, *stationary* queries on *moving* objects, and *moving* queries on *moving* objects.

2.4. CHALLENGE IV: LARGE NUMBER OF CONCURRENT CONTINUOUS QUERIES

Spatio-temporal databases. Most of the existing spatio-temporal algorithms focus on evaluating only one spatio-temporal query (e.g., (Benetis et al., 2002; Iwerks et al., 2003; Lazaridis et al., 2002; Song and Roussopoulos, 2001; Tao et al., 2002; Tao et al., 2003b; Zhang et al., 2003; Zheng and Lee, 2001)). Optimization techniques for evaluating a set of continuous spatio-temporal queries are addressed recently for centralized (Mokbel et al., 2004a; Prabhakar et al., 2002) and distributed environments (Cai et al., 2004; Gedik and Liu, 2004). In centralized environments, the Q-index (Prabhakar et al., 2002) is presented as an R-tree-like index structure to index the *stationary* queries instead of objects. SINA (Mokbel et al., 2004a) uses a *shared* grid structure to index both objects and queries. Then, evaluating a set of continuous queries is abstracted as a spatial join (using the grid index) between objects and queries. In distributed environments, the main idea of (Cai et al., 2004; Gedik and Liu, 2004) is to ship part of the query processing down to the moving objects, while the server acts mainly as a mediator among moving objects.

Data stream management systems. There is a lot of research in sharing the execution of concurrent continuous queries (e.g., NiagaraCQ (Chen et al., 2002; Chen et al., 2000) and PSoup (Chandrasekaran and Franklin, 2002; Chandrasekaran and Franklin, 2003)). The main idea is to have a shared query plan for all continuous queries. Other forms of sharing at the operator level are investigated, for window join (Golab and Ozsu, 2003; Hammad et al., 2003b) and window aggregate operators (Arasu and Widom, 2004).

The PLACE approach. In the PLACE server, we employ a *shared-execution* paradigm similar to that in NiagaraCQ (Chen et al., 2000) and SINA (Mokbel et al., 2004b). The execution of a set of concurrent continuous spatio-temporal queries is performed as a spatial join between two incoming streams. The first stream represents the streaming objects while the second stream represents the streaming queries.

3. The PLACE Server

In this section, we present the data modelling and SQL language used by the PLACE server.

3.1. DATA MODEL

By subscribing with the PLACE server, moving objects are required to send their location updates periodically to the PLACE server. A location update from the client (moving object) to the server has the format (OID, x, y) , where OID is the object identifier, (x, y) is the location of the moving object in the two-dimensional space. An update is timestamped upon its arrival at the server side. Once an object stops moving (e.g., an object reaches its destination or the cellular phone is shut down) it sends to the server a *disappear* message that indicates that the object is no further moving.

Due to the highly dynamic nature of location-aware environments and the infinite size of incoming spatio-temporal streams, we cannot store all incoming data. Thus, the PLACE server employs a three-level storage hierarchy. First, a subset of the incoming data streams is stored in in-memory buffers. In-memory buffers are associated with the outstanding continuous queries at the server. Each query determines which tuples are needed to be in its buffer and when these tuples are expired, i.e., deleted from the buffer. Second, we keep an in-disk storage that keeps track with only one reading of each moving object and query. Since, we cannot update the disk storage every time we receive an update from moving objects, we sample the input data by choosing every k th reading to flush to disk. Moreover, we cache the readings of moving objects/queries and flush them once to the secondary storage every T time units. Data on the secondary storage are indexed using a simple grid structure (Mokbel et al., 2004a). Third, every $T_{archive}$ time units, we take a snapshot of the in-disk database and flush it to a repository server. The repository server acts as a multi-version structure of the moving objects that supports historical queries. Stationary objects (e.g., gas stations, hospitals, restaurants) are preloaded to the system as relational tables that are infrequently updated.

3.2. EXTENDED SQL SYNTAX

As the PLACE server (Mokbel et al., 2004b) extends both PREDATOR (Seshadri, 1998) and NILE (Hammad et al., 2004), we extend the SQL language provided by both systems to support spatio-temporal operators. Mainly, we add the INSIDE and kNN operators to support continuous range queries and k -nearest-neighbor queries respectively. A continuous query is registered at the PLACE server using the SQL:

```
REGISTER QUERY query_name AS
SELECT select_clause
FROM from_clause
```

WHERE *where_clause*
 INSIDE *inside_clause*
 kNN *knn_clause*
 WINDOW *window_clause*

The REGISTER QUERY statement registers the continuous query at the PLACE server with the *query_name* as its identifier. The *select_clause*, *from_clause*, and *where_clause* are inherited from the PREDATOR (Seshadri, 1998) database management statement. The *window_clause* is inherited from the NILE (Hammad et al., 2004) stream query processor to support continuous sliding window queries (Hammad et al., 2003b). A continuous query is dropped from the system using the SQL DROP QUERY *query_name*.

The *inside_clause* may represent stationary/moving rectangular or circular range queries. Moving queries are tied to *focal* objects. As the *focal* object reports movement update to the server, we update the query region. One *focal* object may issue several queries. A rectangular range query can have one of the following two forms:

- Static range query (x_1, y_1, x_2, y_2) , where (x_1, y_1) and (x_2, y_2) represent the top left and bottom right corners of the rectangular range query.
- Moving rectangular range query $(M', ID, xdist, ydist)$, where M' is a flag to indicate that the query is moving, ID is the identifier of the query *focal* point, and $xdist$ and $ydist$ are the length and width of the query rectangle.

A circular range query has the same syntax except that we define only the radius instead of (x, y) . Similarly, the *knn_clause* for continuous k -nearest-neighbor queries may have one of the following two forms:

- Static k NN query (k, x, y) , where k is the number of the neighbors to be maintained, and (x, y) is the center of the query point.
- Moving k NN query (M', k, ID) , where M' is a flag to indicate that the query is moving, k is the number of neighbors to be maintained, and ID is the identifier of the query *focal* point.

4. Predicate-based Sliding Windows

With the unbounded incoming spatio-temporal streams, it becomes infeasible to store all incoming tuples. However, some input tuples may

be buffered in memory for a limited time. The choice of the stored tuples is mainly query dependent, i.e., we store only the tuples that are of interest to any outstanding continuous query. In addition, there should be a mechanism to expire (delete) some of the stored tuples and replace them with other tuples that become more relevant to the outstanding continuous queries. The PLACE server employs a *predicate-based* window policy, where each query is associated with a certain predicate. Only tuples that satisfy at least one query predicate are stored in memory. Based on *predicate-based* window queries, the PLACE server supports three types of tuple expiration, namely, *temporal* expiration (sliding-window queries), *spatial* expiration, and *predicate-based* expiration. Data expiration in *punctuated streams* (Tucker et al., 2003) can be represented as a special case of *predicate-based* expiration.

4.1. TEMPORAL EXPIRATION

Temporal expiration is used commonly to support continuous sliding-window queries in data streams. A sliding window query involves a *historical* time window w . Any object that has a timestamp within the current sliding window of any outstanding query Q is buffered in-memory with the associated buffer of Q . Stored tuples follow strictly a first-in-first-expire policy.

An example for a *historical* sliding window query submitted to the PLACE server is: Q_1 : “*Continuously, report the number of cars that passed by region R in the last hour*”.

```
SELECT COUNT(ObjectID)
FROM MovingObjects
WHERE type = Car
INSIDE R
WINDOW 1 hour
```

Notice that Q_1 buffers all incoming tuples during the previous hour. A tuple is expired (i.e., deleted from the query buffer) once it goes out of the sliding time window (i.e., when the car tuple becomes more than one hour old).

4.2. SPATIAL EXPIRATION

NOW queries are more common in spatio-temporal data stream applications than historical queries. For example, monitoring applications are concerned with the actual current answer not the accumulated historical one. Such applications cannot be realized using only temporal expiration. Thus, the PLACE server introduces a new type of expiration that depends on the *spatial* location of the moving objects instead of

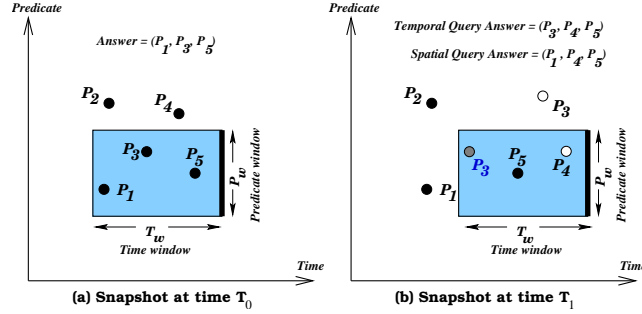


Figure 1. Temporal and spatial expiration.

their timestamps. An incoming tuple, say o , is stored in the in-memory buffer associated with a query Q only if o satisfies the *spatial* window (e.g., region) of Q . A stored tuple is expired only when it steps out of the *spatial* window.

An example of a spatial expiration query is: Q_2 : “Continuously, report the number of cars in a certain area.” Notice that unlike Q_1 , in Q_2 , we are concerned about the actual current number of cars not the number of cars in the recent history. The SQL of Q_2 is similar to that of Q_1 with only the removal of the window statement.

Figure 1 gives the difference between a temporal-expiration and a spatial-expiration in a range query. Figure 1a gives a snapshot of the database at time T_0 . The vertical **bold** line represents the spatial predicate window in the one-dimensional space (along the y axis). The shaded rectangular area represents the temporal window along the time dimension (the x axis). A query with a temporal window is interested in objects that lie in the shaded area while a query with a spatial window is interested in the actual objects that lie in the spatial region regardless of their timestamps. The answer to both the temporal and spatial window queries at time T_0 is (P_1, P_3, P_5) . Figure 1b gives a snapshot of the database at time T_1 . Only objects P_3 and P_4 change their locations (the new locations are plotted as white circle while the old location of P_3 is plotted as a gray circle). For the temporal query, although P_1 still satisfies the query spatial predicate, P_1 is discarded from the server. Also, although P_3 becomes out of the query predicate, the old value of P_3 has not yet expired. Thus, the temporal query answer at T_1 would be (P_3, P_4, P_5) . However, the answer of the spatial window query would be (P_1, P_4, P_5) .

4.3. PREDICATE-BASED EXPIRATION

Predicate-based expiration generalizes both the temporal- and spatial-based expiration methods. An incoming tuple is stored in-memory only

when it satisfies a certain query predicate. A predicate could be as simple as a temporal or a spatial predicate. Due to the nature of spatio-temporal streams, other forms of predicates may arise. For example, consider the query Q_3 : “For each moving object, continuously report the elapsed time between each two consecutive readings”. Such a query contains a self-join where objects from the stream of moving objects are self-joined based on the object identifier. The query buffer needs to maintain only the latest reading of each moving object. Once the reading of a certain object is reported, the previous reading is expired. *Punctuation-based* expiration (Tucker et al., 2003) can be considered as a special case of *predicate-based* when we consider the predicate as the arrival of a certain artificial tuple (the *punctuated* tuple).

5. Incremental Evaluation

To avoid reevaluating continuous spatio-temporal queries, we employ an *incremental evaluation* paradigm in the PLACE continuous query processor. The main idea is to only report the changes of the answer from the last evaluation time. By employing *incremental evaluation*, the PLACE server achieves the following goals: (1) Fast query evaluation, since we compute only the updates of the answer not the whole answer. (2) In a typical location-aware server, query results are sent to the users via satellite servers (Acharya et al., 1996; Hambrusch et al., 2001). Thus, limiting the amount of transmitted data to the updates only rather than the whole query answer saves in network bandwidth. (3) When encapsulating incremental algorithms into physical pipelined query operators, limiting the tuples that go through the whole query pipeline to only the updates reduces the flow in the pipeline. Thus, efficient query processing is achieved.

To realize the incremental evaluation processing in the PLACE server, we go through three main steps. First, we define the high level concept of incremental updates, by defining two types of updates; *positive* and *negative* updates (Mokbel, 2004; Mokbel et al., 2004a). Second, we encapsulate the processing of incremental algorithms into pipelined query operators. Third, we modify traditional pipelined query operators (e.g., distinct and join) to deal with the concept of negative tuples (Hammad et al., 2003c).

5.1. POSITIVE/NEGATIVE UPDATES

Incremental evaluation is achieved through updating the previous query answer. Mainly, we distinguish between two types of updates; *positive*

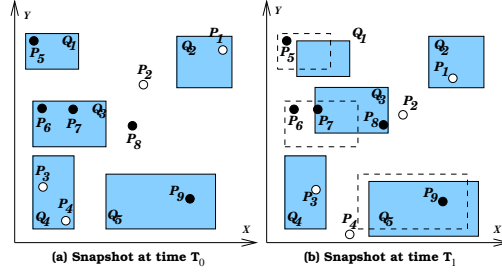


Figure 2. Incremental evaluation of range queries

updates and *negative* updates. A positive/negative update indicates that a certain object needs to be added to/removed from the query answer. A query answer is represented in the form $(QID, OList)$, where QID is the query identifier and $OList$ is the query answer. The PLACE server continuously updates the query answer with updates of the form (QID, \pm, OID) where \pm indicates the type of the update and OID is the object identifier.

Figure 2 gives an example of applying the concept of positive/negative updates on a set of continuous range queries. The snapshot of the database at time T_0 is given in Figure 2a with nine moving objects, p_1 to p_9 , and five continuous range queries, Q_1 to Q_5 . The answer of the queries at time T_0 is represented as (Q_1, P_5) , (Q_2, P_1) , (Q_3, P_6, P_7) , (Q_4, P_3, P_4) , and (Q_5, P_9) . At time T_1 (Figure 2b), only the objects p_1, p_2, p_3 , and p_4 and the queries Q_1, Q_3 , and Q_5 change their locations. As a result, the PLACE server reports the following updates: $(Q_1, -P_5)$, $(Q_3, -P_6)$, $(Q_3, +P_8)$, and $(Q_4, -p_4)$.

5.2. SPATIO-TEMPORAL INCREMENTAL PIPELINED OPERATORS

Two alternative approaches can be utilized in implementing spatio-temporal algorithms inside the PLACE server: using SQL *table functions* (Reinwald and Pirahesh, 1998) or by encapsulating the algorithms in physical query operators. Since there is no straightforward method for pushing query predicates into table functions (Reinwald et al., 1999), the performances is limited and the approach does not give enough flexibility in optimizing the issued queries. In the PLACE server we encapsulate our algorithms inside physical pipelined query operators that can be part of a query execution plan. By having pipelined query operators, we achieve three goals: (1) Spatio-temporal operators can be combined with other operators (e.g., distinct, aggregate, and join operators) to support incremental evaluation for a wide variety of continuous spatio-temporal queries. (2) Pushing spatio-temporal operators deep in the query execution plan reduces the number of tuples




Figure 3. Shared-execution of continuous queries.

in the query pipeline. This reduction comes from the fact that spatio-temporal operators act as filters to the above operators. (3) Flexibility in the query optimizer where multiple candidate execution plans can be produced.

The main idea of spatio-temporal operators is to keep track of the recently reported answer of each query Q in a query buffer termed $Q.Answer$. Then, for each newly incoming tuple P , we perform two tests: Test I: Is P part of the previously reported $Q.Answer$? Test II: Does P qualify to be part of the current answer? Based on the results of the two tests, we distinguish among four cases:

- **Case I:** P is part of $Q.Answer$ and P still qualify to be part of the current answer. As we process only the updates of the previously reported result, P will not be processed.
- **Case II:** P is part of $Q.Answer$, however, P does not qualify to be part of the answer anymore. In this case, we report a *negative* update P^- to the above query operator. The *negative* update indicates that P is *spatially* expired from the answer.
- **Case III:** P is not part of $Q.Answer$, however, P qualifies to be part of the current answer. In this case, we report a *positive* update to the above query operator.
- **Case IV:** P is not part of $Q.Answer$ and P still does not qualify to be part of the current answer. In this case, P has no effect on Q .

6. Scalability

The PLACE continuous query processor exploits a *shared-execution* paradigm (Mokbel et al., 2003a; Mokbel et al., 2004a; Xiong et al.,

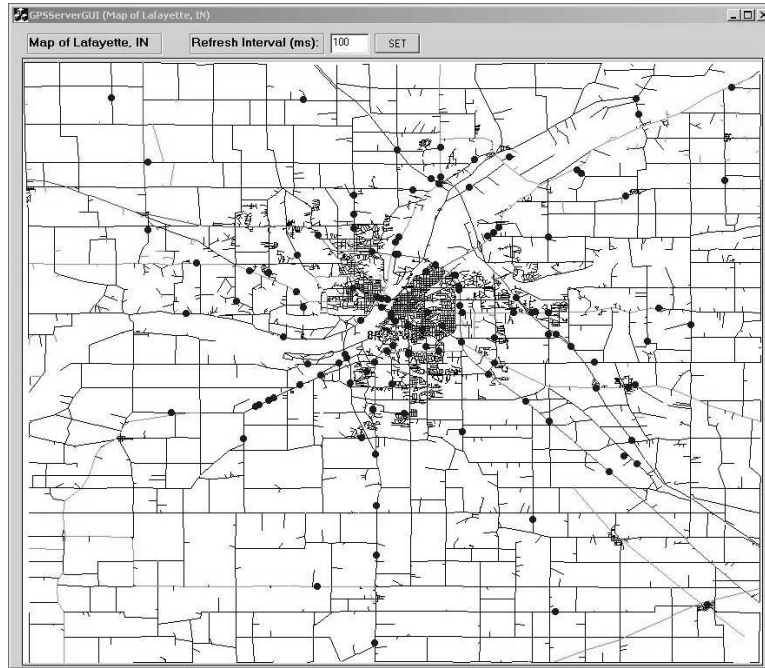


Figure 4. Server GUI.

2004) as a means for achieving scalability in terms of the number of concurrently executing continuous spatio-temporal queries. The main idea is to group similar queries in a query table. Then, the evaluation of a set of continuous queries is modelled as a spatial join between moving objects and moving queries. Similar ideas of shared-execution have been exploited in the NiagaraCQ (Chen et al., 2000) for web queries and PSoup (Chandrasekaran and Franklin, 2002; Chandrasekaran and Franklin, 2003) for streaming queries.

Figure 3a gives the execution plans of two simple continuous spatio-temporal queries, Q_1 : “Find the objects inside region R_1 ”, and Q_2 : “Find the objects inside region R_2 ”. With *shared-execution*, we have the execution plan of Figure 3b. Shared-execution for a collection of spatio-temporal range queries can be expressed in the PLACE server by issuing the following continuous query:

```

SELECT Q.ID, O.ID
FROM QueryTable Q, ObjectTable O
WHERE O.location inside Q.region

```

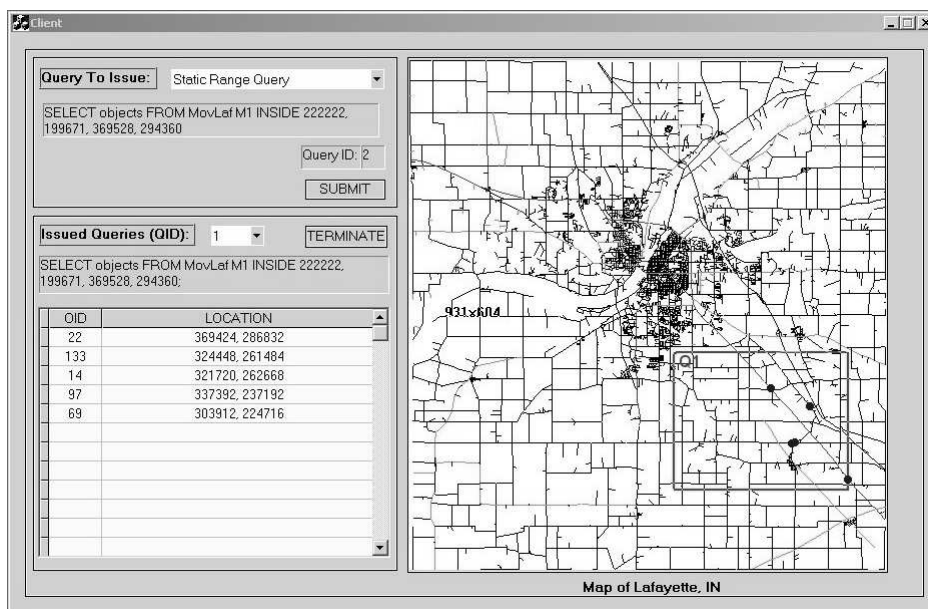


Figure 5. Client GUI.

7. User interface in PLACE

Figures 4 and 5 give snapshots of the server and client graphical user interface (GUI) of PLACE, respectively. The server GUI displays all moving objects on the map¹. The client GUI simulates a client end-device used by the users. Users can choose the type of query from a list of available query types (stationary/moving queries and stationary/moving k -nearest-neighbor queries). The spatial region of the query can be determined using the map of the area of interest. By pressing the *submit* button, the client translates the query into SQL language and transmits it to the PLACE server. The result appears both in the list of Figure 5 and as moving objects on the map. A client can see only, on its map, the objects that belong to its issued query.

8. Performance Evaluation

In this section, we present experiments that show the promising performance of the continuous query processor in the PLACE server. We use the *Network-based Generator of Moving Objects* (Brinkhoff, 2002) to generate a set of 100K moving objects and 100K moving queries.

¹ The map in Figures 4 and 5 is for the Greater Lafayette area, Indiana, USA.




Figure 6. A) Moving objects (%). B) Query size

The output of the generator is a set of moving objects that move on the road network of a given city. We choose some points randomly and consider them as centers of square range queries.

8.1. SIZE OF INCREMENTAL ANSWER

Figure 6 compares between the size of the incremental answer returned by utilizing the incremental approach and the size of the complete answer. The location-aware server buffers the received updates from moving objects and queries and evaluates the queries every 5 seconds. Figure 6a gives the effect of the number of moving objects that reported a change of location within the last 5 seconds. The size of the complete answer is constant and is orders of magnitude of the size of the worst-case incremental answer. In Figure 6b, the query side length varies from 0.01 to 0.02. The size of the complete answer increases dramatically to up to seven times that of the incremental result. The saving in the answer size affects directly the communication cost from the server to the clients.

8.2. PIPELINED SPATIO-TEMPORAL OPERATORS

In this section, we compare the implementation of spatio-temporal algorithms at the application-level (e.g., table functions) with the encapsulation of the spatio-temporal algorithms inside query operators.

8.2.1. Pipeline with a Select Operator

Consider the query Q : “Continuously report all trucks that are within $MyArea$ ”. $MyArea$ can be either a stationary or moving range query. A high-level implementation of this query has only a selection operator that selects only the “trucks”. Then, a high-level algorithm implementation would take the selection output and produce incrementally the query result. However, an encapsulation of the `INSIDE` algorithm into a physical operator allows for more flexible plans.




Figure 7. A. Query SQL and pipeline. B. Pipelined operator with Select.

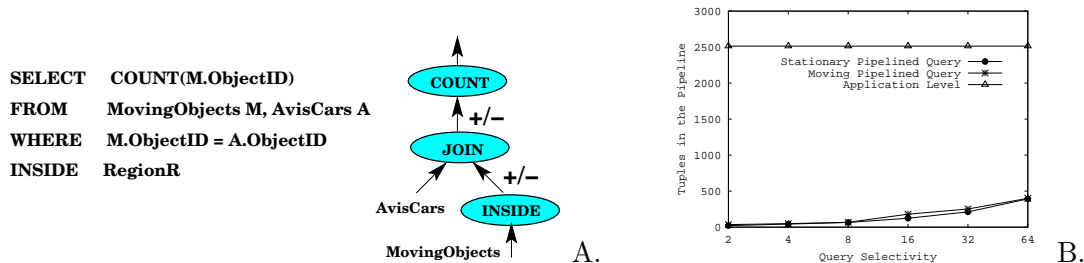


Figure 8. A) Query SQL and pipeline. B) Pipelined operator with Join.

Figure 7 compares the high-level implementation of the above query with pipelined operators for both stationary and moving queries. The selectivity of the queries varies from 2% to 64%. The selectivity of the selection operator is 5%. Our measure of comparison is the number of tuples that go through the query evaluation pipeline. When algorithms are implemented at the application level, the performance is not affected by the selectivity. However, when `INSIDE` is pushed before the *selection*, it acts as a filter for the query evaluation pipeline, thus, limiting the tuples through the pipeline to only the incremental updates. With `INSIDE` selectivity less than 32%, pushing `INSIDE` before the selection affects the performance greatly.

8.2.2. Pipeline with a Join Operator

In this section, we consider a more complex query plan that contains a *join* operator. Consider the query Q : “Continuously report moving objects that belong to my favorite set of objects and lie within MyArea”. A high-level implementation would probe a streaming database engine to join all moving objects with my favorite set of objects. Then, the output of the join is sent to the algorithm for further processing. However, with the `INSIDE` operator, we can have a query evaluation plan as that of Figure 8a where the `INSIDE` operator is pushed below the *Join* operator.

Figure 8 compares the high-level implementation of the above query with the pipelined `INSIDE` operator for both stationary and moving queries. The selectivity of the queries varies from 2% to 64%. Unlike the case of *selection* operators, there is a dramatic increase in performance (around ten orders of magnitude) when `INSIDE` is implemented as a pipelined operator. The main reason in this dramatic gain in performance is the high overhead incurred when evaluating the *join* operation. Thus, the `INSIDE` operator filters out the input tuples and limits the input to the *join* operator to only the incremental *positive* and *negative* updates.

9. Supporting Negative Tuples in Nile

The PLACE location-aware server is built by modifying and augmenting the Nile data stream management (Hammad et al., 2004) system to support scalable execution of continuous queries, moving queries, and spatio-temporal operators. The underlying Nile query processor uses the *negative tuple* approach (Hammad et al., 2003c) to handle continuous time-based sliding-window queries. The basic idea is to introduce a new operator, termed `W-Expire`, that is responsible for *expiring* old tuples. For each *expired* tuple T , the `W-Expire` operator emits a negative tuple T^- that has the same value as that of T , however it is tagged with a *negative* sign. By having the `W-Expire` operator at the bottom of the pipeline, all other operators do not need to be aware of any timing issues of the incoming tuples.

Figure 9 gives a typical query pipeline in the Nile query processor. Two incoming streams (A and B) are joined together. Then, the output is joined with a relational table R . The output of the latter join is sent to the `DISTINCT` and `COUNT` operators. Two `W-Expire` operators, one for each stream, are used at the bottom of the pipeline plan. All the above operators need to have the ability to recognize and process the incoming *negative* tuples.

9.1. NEGATIVE TUPLES IN TRADITIONAL OPERATORS

The Nile query processor (Hammad et al., 2004) modifies traditional stream operators to recognize and process the newly introduced *negative* tuples. The modifications are mainly the following:

- *Selection* and *Join* operators handle *negative* tuples in the same way as the regular *positive* tuples. The only difference is that the output will be in the form of a *negative* tuple.




Figure 9. A query pipelined execution plan in the Nile stream query processor.

- *Aggregates* update their aggregate functions by considering the received *negative* tuple.
- The *Distinct* operator reports a *negative* tuple at the output only if the corresponding *positive* tuple is in the recently reported result.

For more detail about handling the *negative* tuples in various query operators, the reader is referred to (Hammad et al., 2003c).

10. Conclusion

In this paper, we presented the continuous query processor of the PLACE (Pervasive Location-Aware Computing Environments) server; a database server for location-aware environments currently developed at Purdue University. The PLACE server extends both the PREDATOR database management system and the NILE stream query processor to deal with unbounded spatio-temporal streams. In addition to the temporal tuple expiration defined in sliding-window queries, we maintain other forms of tuple expirations (e.g., spatial expiration). To efficiently handle large number of continuous queries, we employ an incremental evaluation paradigm that contains: (1) Defining the concept of *positive* and *negative* updates, (2) Encapsulating the algorithms for incremental processing into pipelined spatio-temporal operators, and (3) Modifying traditional query operators (e.g., distinct and join) to deal with the *negative* updates that comes from the spatio-temporal operators. Shared-execution is employed by the continuous query processor as a means of achieving scalability in terms of the number of concurrently continuous queries. Experimental results show the promising performance of the PLACE continuous query processor.

References

- Abadi, D. J., D. Carney, U. Cetintemel, M. Cherniack, C. Convey, S. Lee, M. Stonebraker, N. Tatbul, and S. B. Zdonik: 2003, 'Aurora: A New Model and Architecture for Data Stream Management'. *VLDB Journal* **12**(2).
- Acharya, S., M. J. Franklin, and S. B. Zdonik: 1996, 'Disseminating Updates on Broadcast Disks'. In: *Proceedings of the International Conference on Very Large Data Bases, VLDB*.
- Arasu, A., B. Babcock, S. Babu, J. Cieslewicz, M. Datar, K. Ito, R. Motwani, U. Srivastava, and J. Widom: 2004, 'STREAM: The Stanford Data Stream Management System'.
- Arasu, A. and J. Widom: 2004, 'Resource Sharing in Continuous Sliding-Window Aggregates'. In: *Proceedings of the International Conference on Very Large Data Bases, VLDB*.
- Aref, W. G., S. E. Hambruch, and S. Prabhakar: 2003, 'Pervasive Location Aware Computing Environments (PLACE)'. <http://www.cs.purdue.edu/place/>.
- Babcock, B., M. Datar, and R. Motwani: 2003, 'Sampling from a moving window over streaming data'. In: *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms, SODA*.
- Babcock, B., M. Datar, and R. Motwani: 2004, 'Load Shedding for Aggregation Queries over Data Streams'. In: *Proceedings of the International Conference on Data Engineering, ICDE*.
- Babu, S. and J. Widom: 2001, 'Continuous Queries over Data Streams'. *SIGMOD Record* **30**(3).
- Benetis, R., C. S. Jensen, G. Karciuskas, and S. Saltenis: 2002, 'Nearest Neighbor and Reverse Nearest Neighbor Queries for Moving Objects'. In: *Proceedings of the International Database Engineering and Applications Symposium, IDEAS*.
- Brinkhoff, T.: 2002, 'A Framework for Generating Network-Based Moving Objects'. *GeoInformatica* **6**(2).
- Cai, Y., K. A. Hua, and G. Cao: 2004, 'Processing Range-Monitoring Queries on Heterogeneous Mobile Objects'. In: *Mobile Data Management, MDM*.
- Chandrasekaran, S., O. Cooper, A. Deshpande, M. J. Franklin, J. M. Hellerstein, W. Hong, S. Krishnamurthy, S. Madden, V. Raman, F. Reiss, and M. A. Shah: 2003, 'TelegraphCQ: Continuous Dataflow Processing for an Uncertain World'. In: *Proceedings of the International Conference on Innovative Data Systems Research, CIDR*.
- Chandrasekaran, S. and M. J. Franklin: 2002, 'Streaming Queries over Streaming Data'. In: *Proceedings of the International Conference on Very Large Data Bases, VLDB*.
- Chandrasekaran, S. and M. J. Franklin: 2003, 'PSoup: a system for streaming queries over streaming data'. *VLDB Journal* **12**(2), 140–156.
- Chen, J., D. J. DeWitt, and J. F. Naughton: 2002, 'Design and Evaluation of Alternative Selection Placement Strategies in Optimizing Continuous Queries'. In: *Proceedings of the International Conference on Data Engineering, ICDE*.
- Chen, J., D. J. DeWitt, F. Tian, and Y. Wang: 2000, 'NiagaraCQ: A Scalable Continuous Query System for Internet Databases'. In: *Proceedings of the ACM International Conference on Management of Data, SIGMOD*.
- Clifford, J., C. E. Dyreson, T. Isakowitz, C. S. Jensen, and R. T. Snodgrass: 1997, 'On the Semantics of "Now" in Databases'. *ACM Transactions on Database Systems*, *TODS* **22**(2).

- Cormode, G. and S. Muthukrishnan: 2003, 'Radial Histograms for Spatial Streams'. Technical Report DIMACS TR: 2003-11, Rutgers University.
- Cranor, C., T. Johnson, O. Spataschek, and V. Shkapenyuk: 2003, 'Gigascope: a Stream Database for Network Applications'. In: *Proceedings of the ACM International Conference on Management of Data, SIGMOD*.
- Datar, M., A. Gionis, P. Indyk, and R. Motwani: 2002, 'Maintaining stream statistics over sliding windows'. In: *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms, SODA*.
- Gedik, B. and L. Liu: 2004, 'MobiEyes: Distributed Processing of Continuously Moving Queries on Moving Objects in a Mobile System'. In: *Proceedings of the International Conference on Extending Database Technology, EDBT*.
- Golab, L., S. Garg, and M. T. Ozsü: 2004, 'On Indexing Sliding Windows over Online Data Streams'. In: *Proceedings of the International Conference on Extending Database Technology, EDBT*.
- Golab, L. and M. T. Ozsü: 2003, 'Processing Sliding Window Multi-Joins in Continuous Queries over Data Streams'. In: *Proceedings of the International Conference on Very Large Data Bases, VLDB*.
- Hadjieleftheriou, M., G. Kollios, D. Gunopulos, and V. J. Tsotras: 2003, 'On-Line Discovery of Dense Areas in Spatio-temporal Databases'. In: *Proceedings of the International Symposium on Advances in Spatial and Temporal Databases, SSTD*.
- Hambrusch, S. E., C.-M. Liu, W. G. Aref, and S. Prabhakar: 2001, 'Query Processing in Broadcasted Spatial Index Trees'. In: *Proceedings of the International Symposium on Advances in Spatial and Temporal Databases, SSTD*.
- Hammad, M. A., W. G. Aref, and A. K. Elmagarmid: 2003a, 'Stream Window Join: Tracking Moving Objects in Sensor-Network Databases'. In: *Proceedings of the International Conference on Scientific and Statistical Database Management, SSDBM*.
- Hammad, M. A., M. J. Franklin, W. G. Aref, and A. K. Elmagarmid: 2003b, 'Scheduling for shared window joins over data streams'. In: *Proceedings of the International Conference on Very Large Data Bases, VLDB*.
- Hammad, M. A., T. M. Ghanem, W. G. Aref, A. K. Elmagarmid, and M. F. Mokbel: 2003c, 'Efficient Pipelined Execution of Sliding-Window Queries Over Data Streams'. Technical Report TR CSD-03-035, Purdue University Department of Computer Sciences.
- Hammad, M. A., M. F. Mokbel, M. H. Ali, W. G. Aref, A. C. Catlin, A. K. Elmagarmid, M. Eltabakh, M. G. Elfeky, T. M. Ghanem, R. Gwadera, I. F. Ilyas, M. Marzouk, and X. Xiong: 2004, 'Nile: A Query Processing Engine for Data Streams (Demo)'. In: *Proceedings of the International Conference on Data Engineering, ICDE*.
- Hershberger, J. and S. Suri: 2004, 'Adaptive Sampling for Geometric Problems over Data Streams'. In: *Proceedings of the ACM Symposium on Principles of Database Systems, PODS*.
- Iwerks, G. S., H. Samet, and K. Smith: 2003, 'Continuous K-Nearest Neighbor Queries for Continuously Moving Points with Updates'. In: *Proceedings of the International Conference on Very Large Data Bases, VLDB*.
- Jensen, C. S., D. Lin, and B. C. Ooi: 2004, 'Query and Update Efficient B+-Tree Based Indexing of Moving Objects'. In: *Proceedings of the International Conference on Very Large Data Bases, VLDB*.

- Kang, J., J. F. Naughton, and S. Viglas: 2003, 'Evaluating Window Joins over Unbounded Streams'. In: *Proceedings of the International Conference on Data Engineering, ICDE*.
- Kwon, D., S. Lee, and S. Lee: 2002, 'Indexing the Current Positions of Moving Objects Using the Lazy Update R-tree'. In: *Mobile Data Management, MDM*.
- Lazaridis, I., K. Porkaew, and S. Mehrotra: 2002, 'Dynamic Queries over Mobile Objects'. In: *Proceedings of the International Conference on Extending Database Technology, EDBT*.
- Lee, M.-L., W. Hsu, C. S. Jensen, and K. L. Teo: 2003, 'Supporting Frequent Updates in R-Trees: A Bottom-Up Approach'. In: *Proceedings of the International Conference on Very Large Data Bases, VLDB*.
- Madden, S., M. Shah, J. M. Hellerstein, and V. Raman: 2002, 'Continuously adaptive continuous queries over streams'. In: *Proceedings of the ACM International Conference on Management of Data, SIGMOD*.
- Mokbel, M. F.: 2004, 'Continuous Query Processing in Spatio-temporal Databases'. In: *Proceedings of the ICDE/EDBT PhD Workshop*.
- Mokbel, M. F., W. G. Aref, S. E. Hambrusch, and S. Prabhakar: 2003a, 'Towards Scalable Location-aware Services: Requirements and Research Issues'. In: *Proceedings of the ACM workshop on Advances in Geographic Information Systems, ACM GIS*.
- Mokbel, M. F., T. M. Ghanem, and W. G. Aref: 2003b, 'Spatio-temporal Access Methods'. *IEEE Data Engineering Bulletin* **26**(2).
- Mokbel, M. F., X. Xiong, and W. G. Aref: 2004a, 'SINA: Scalable Incremental Processing of Continuous Queries in Spatio-temporal Databases'. In: *Proceedings of the ACM International Conference on Management of Data, SIGMOD*.
- Mokbel, M. F., X. Xiong, W. G. Aref, S. Hambrusch, S. Prabhakar, and M. Hammad: 2004b, 'PLACE: A Query Processor for Handling Real-time Spatio-temporal Data Streams (Demo)'. In: *Proceedings of the International Conference on Very Large Data Bases, VLDB*.
- Mokbel, M. F., X. Xiong, M. A. Hammad, and W. G. Aref: 2004c, 'Continuous Query Processing of Spatio-temporal Data Streams in PLACE'. In: *Proceedings of the second workshop on Spatio-Temporal Database Management, STDBM*.
- Motwani, R., J. Widom, A. Arasu, B. Babcock, S. Babu, M. Datar, G. S. Manku, C. Olston, J. Rosenstein, and R. Varma: 2003, 'Query Processing, Approximation, and Resource Management in a Data Stream Management System'. In: *Proceedings of the International Conference on Innovative Data Systems Research, CIDR*.
- Prabhakar, S., Y. Xia, D. V. Kalashnikov, W. G. Aref, and S. E. Hambrusch: 2002, 'Query Indexing and Velocity Constrained Indexing: Scalable Techniques for Continuous Queries on Moving Objects'. *IEEE Trans. on Computers* **51**(10).
- Reinwald, B. and H. Pirahesh: 1998, 'SQL Open Heterogeneous Data Access'. In: *Proceedings of the ACM International Conference on Management of Data, SIGMOD*.
- Reinwald, B., H. Pirahesh, G. Krishnamoorthy, G. Lapis, B. T. Tran, and S. Vora: 1999, 'Heterogeneous Query Processing through SQL Table Functions'. In: *Proceedings of the International Conference on Data Engineering, ICDE*.
- Saltenis, S. and C. S. Jensen: 2002, 'Indexing of Moving Objects for Location-Based Services'. In: *Proceedings of the International Conference on Data Engineering, ICDE*.

- Saltenis, S., C. S. Jensen, S. T. Leutenegger, and M. A. Lopez: 2000, 'Indexing the Positions of Continuously Moving Objects'. In: *Proceedings of the ACM International Conference on Management of Data, SIGMOD*.
- Seshadri, P.: 1998, 'Predator: A Resource for Database Research'. *SIGMOD Record* **27**(1), 16–20.
- Song, Z. and N. Roussopoulos: 2001, 'K-Nearest Neighbor Search for Moving Query Point'. In: *Proceedings of the International Symposium on Advances in Spatial and Temporal Databases, SSTD*.
- Srivastava, U. and J. Widom: 2004, 'Memory-Limited Execution of Windowed Stream Joins'. In: *Proceedings of the International Conference on Very Large Data Bases, VLDB*.
- Sun, J., D. Papadias, Y. Tao, and B. Liu: 2004, 'Querying about the Past, the Present and the Future in Spatio-Temporal Databases'. In: *Proceedings of the International Conference on Data Engineering, ICDE*.
- Tao, Y., G. Kollios, J. Considine, F. Li, and D. Papadias: 2004, 'Spatio-Temporal Aggregation Using Sketches'. In: *Proceedings of the International Conference on Data Engineering, ICDE*.
- Tao, Y., D. Papadias, and Q. Shen: 2002, 'Continuous Nearest Neighbor Search'. In: *Proceedings of the International Conference on Very Large Data Bases, VLDB*.
- Tao, Y., D. Papadias, and J. Sun: 2003a, 'The TPR*-Tree: An Optimized Spatio-temporal Access Method for Predictive Queries'. In: *Proceedings of the International Conference on Very Large Data Bases, VLDB*.
- Tao, Y., J. Sun, and D. Papadias: 2003b, 'Analysis of Predictive Spatio-Temporal Queries'. *ACM Transactions on Database Systems*, *TODS* **28**(4).
- Tatbul, N., U. Cetintemel, S. B. Zdonik, M. Cherniack, and M. Stonebraker: 2003, 'Load Shedding in a Data Stream Manager'. In: *Proceedings of the International Conference on Very Large Data Bases, VLDB*.
- Tucker, P. A., D. Maier, T. Sheard, and L. Fegaras: 2003, 'Exploiting Punctuation Semantics in Continuous Data Streams'. *IEEE Transactions on Knowledge and Data Engineering, TKDE* **15**(3), 555–568.
- Wolfson, O., H. Cao, H. Lin, G. Trajcevski, F. Zhang, and N. Rische: 2002, 'Management of Dynamic Location Information in DOMINO (Demo)'. In: *Proceedings of the International Conference on Extending Database Technology, EDBT*.
- Wolfson, O., A. P. Sistla, B. Xu, J. Zhou, and S. Chamberlain: 1999, 'DOMINO: Databases fOr MovINg Objects tracking (Demo)'. In: *Proceedings of the ACM International Conference on Management of Data, SIGMOD*.
- Wolfson, O., B. Xu, S. Chamberlain, and L. Jiang: 1998, 'Moving Objects Databases: Issues and Solutions'. In: *Proceedings of the International Conference on Scientific and Statistical Database Management, SSDBM*.
- Xiong, X., M. F. Mokbel, W. G. Aref, S. Hambrusch, and S. Prabhakar: 2004, 'Scalable Spatio-temporal Continuous Query Processing for Location-aware Services'. In: *Proceedings of the International Conference on Scientific and Statistical Database Management, SSDBM*.
- Yao, Y. and J. Gehrke: 2002, 'The Cougar Approach to In-Network Query Processing in Sensor Networks'. *SIGMOD Record* **31**(3).
- Zhang, J., M. Zhu, D. Papadias, Y. Tao, and D. L. Lee: 2003, 'Location-based Spatial Queries'. In: *Proceedings of the ACM International Conference on Management of Data, SIGMOD*.
- Zheng, B. and D. L. Lee: 2001, 'Semantic Caching in Location-Dependent Query Processing'. In: *Proceedings of the International Symposium on Advances in Spatial and Temporal Databases, SSTD*.

