

INFORMATIK BERICHTE

355 – 11/2009

Spatiotemporal Pattern Queries

Mahmoud Attia Sakr, Ralf Hartmut Güting



**Fakultät für Mathematik und Informatik
Postfach 940
D-58084 Hagen**

Spatiotemporal Pattern Queries

Mahmoud A.Sakr ^{#1,*2}, Ralf H.Güting ^{#1}

^{#1}*Database Systems for New Applications, FernUniversität in Hagen
58084 Hagen, Germany*

^{*2}*Faculty of Computer and Information Sciences, University of Ain Shams
Cairo, Egypt*

¹mahmoud.sakr@fernuni-hagen.de

²rhg@fernuni-hagen.de

November 16, 2009

Abstract

Spatiotemporal pattern queries allow for querying moving objects by their movement profiles. That is, one can specify for example temporal order constraints on the fulfillment of predicates on moving objects. We propose a complete design for spatiotemporal pattern queries in the context of spatiotemporal DBMSs. The design builds on the well established concept of *lifted predicates*. Hence, unlike previous approaches, it is not restricted to specific sets of predicates. It can express a wide range of spatiotemporal patterns including various types of predicates such as kNN, range, metric, topological, set operations, aggregations, distance, direction, and boolean operations. Our design covers the language integration in SQL, the evaluation of the queries, and the integration with the query optimizer. We also propose a simple language for defining temporal constraints. The approach allows for queries that were never available. We provide a complete implementation in C++ and Prolog in the context of the SECONDO platform. The implementation is made publicly available online as a SECONDO Plugin. We have also included in the Plugin automatic scripts for repeating the experiments in this paper.

1 Introduction

Moving objects are objects that change their position and/ or extent with time. Having the trajectories of these objects stored in a suitable database system allows for issuing spatiotemporal queries. One can query, for example, for animals which crossed a certain lake during a certain time or for the total length of a car trajectory inside a certain zone.

Spatiotemporal pattern queries provide a more complex query framework for moving objects. In particular, they specify a temporal order among a set of spatiotemporal predicates. They are used to filter a set of trajectories and pass only those which fulfill a certain relative or exact temporal ordering of predicates during their movement.

The term Spatiotemporal Patterns (STP) is used in the literature with different meanings. In the database literature, the research on STP goes in two directions, namely data mining and database query. Although there is no clear cut between both areas, it happened that the literature clusters itself into these two classes.

From the data mining perspective, STPs refer to the collective movement behavior, also called group patterns. The methods analyze simultaneous movements and the interaction between objects (e.g. patterns like leadership, play, fighting, migration, trend-setting, ... etc). The research in this direction aims at developing a toolbox of data mining algorithms and visual analytic techniques for movement analysis. For example, algorithms for the flock, leadership, convergence and encounter patterns are presented in

[13]. More systematically, Dodge et al. [8], presented a classification of the movement patterns. They gave examples for different classes and referred to relevant mining algorithms in the literature. The next logical step is to build tools that can, more generally, identify classes of patterns.

In this paper, we are interested in the other perspective. It is focused on the development of query methods that can be added to extensible spatiotemporal DBMS. These query methods should allow the user to query for database objects that depict individual spatiotemporal patterns. In this sense, every object/tuple can individually answer the user query. This is the intuitive way to query the moving objects by their movement profiles.

The pattern is a set of spatiotemporal predicates that are related to each other by temporal constraints. For example, suppose predicates P , Q , and R that can hold over a time interval or a single instant. We would like to be able to express spatiotemporal pattern conditions like the following:

- P then (later) Q then R .
- P ending before 8:30 then Q for no more than 1 hour.
- $(Q$ then $R)$ during P .

The predicates P , Q , R , etc. might be of the form:

- Vehicle X is on road W .
- The extent of the storm area Y is larger than 4 square kms.
- The speed of air plane Z is between 400 and 500 km/h.

The literature in this perspective refers to spatiotemporal patterns using different names:

- They are called *spatiotemporal patterns* in [9] and [17].
- In [20] they are called *inverse elementary queries*. They are *inverse* because the user already knows how to describe the pattern and wants to retrieve objects depicting the pattern. This is opposite to looking for a *frequent pattern*. They are *elementary* in the sense that one trajectory can individually answer the query. This is opposite to the *synoptic queries* that target collective patterns similar to the data mining approach.
- In [21] they are called *trajectory based queries* because they rely on sequential information in the trajectory. This is in contrast to *coordinate based queries* that concentrate on a single part of the trajectory.

We call them spatiotemporal patterns (STP). So far, there exist only few proposals for handling STP queries. A language based approach is proposed by Mouza and Rigaux [19]. They discretize the spatial space into labeled zones (e.g. a spatial grid), and the temporal dimension into constant-size intervals. The trajectories are strings of labels that show the visiting order of the cells. The patterns are represented as regular expressions. Thanks to this representation, the problem of matching a spatiotemporal pattern is reduced to matching a regular expression against a set of strings. Hadjieleftheriou et. al. [17] propose efficient algorithms that use a specialized index structure to evaluate STP queries consisting of spatial and nearest neighbor predicates. Another discussion by Erwig [9] outlines some ideas to extend the spatiotemporal predicates [11] towards spatiotemporal patterns. We discuss the details and limitations of these approaches in more detail in Section 6.

Our contributions are the following:

- We propose a new approach for answering STP queries that is based on a very general and powerful class of predicates, the so-called lifted predicates [16]. They are very powerful as they are simply the time dependent version of arbitrary static predicates. Instead of returning a *bool* value (like standard predicates) they return a *moving(bool)* (a boolean function of time). Our approach allows one to formulate temporal constraints on the results of arbitrary expressions returning such moving booleans. Formulating STP queries over lifted predicates allows for a wide range of queries that are not addressed before.
- Thanks to the clean design, the proposed approach can be easily extended to support more complex patterns. In Section 5, we describe an extension that further increases the expressive power.
- In contrast to previous work we are able to actually integrate STP queries into the query optimizer. Obviously for an efficient execution of pattern queries on large databases the use of indexes is mandatory. In Section 7 we consider how STP queries can be mapped by the query optimizer to efficient index accesses.
- We propose a simple language for describing the relationship between two intervals (e.g. Allen’s operators). The language makes it easier, from the user point of view, to express interval relations without the need to memorize their names.
- We provide the complete implementation for the proposed design in the context of the SECONDO platform [4]. The implementation is made publicly available as a SECONDO Plugin and can be downloaded from the Plugin web site [1]. Parallel to this paper, we have written a user manual describing how to install and run our spatiotemporal pattern algebra within a SECONDO system.
- We also provide automatic scripts to repeat the experiments in this paper. The scripts are installed during the installation of the Plugin. In Section 11 we describe the detailed procedure to repeat the experiments. The scripts, together with the well documented source code provided in the Plugin, allow the readers to explore our approach, further elaborate on it, and compare with other future approaches.

The rest of this paper is organized as follows. Section 2 gives a brief background about the problem domain and recalls some necessary definitions from previous work in moving objects databases. In Section 3, we outline the proposed approach. Section 4 formalizes the spatiotemporal pattern predicate as a constraint satisfaction problem. Then we describe the evaluation algorithms. In Section 5, the basic spatiotemporal pattern predicate is extended into a more expressive version. Section 6 reviews the previous related work. In Section 7 we show how to integrate our approach seamlessly with the query optimizers. Section 8 is dedicated to the technical aspects of the implementation in the SECONDO framework. The experimental evaluation is shown in Section 9. In Section 10, we demonstrate two application examples that emphasize the expressive power of our approach. Section 11 and the Appendices at the end of the paper describe the experimental repeatability. Finally we conclude in Section 12.

2 Moving Objects Databases

Moving objects can be abstracted as geometries that change their position and/or extent with time. In previous work [16], [12], and [7], a model for representing and querying moving objects is proposed. The work is based on abstract data types (ADT). The *moving* type constructor is used to construct the moving counterpart of every static data type. Moving geometries are represented using three abstractions; *moving(point)*, *moving(region)* and *moving(line)*. Simple data types (e.g. *integer*, *bool*, *real*) are also mapped to *moving* types. In the *abstract model* [16], moving objects are modeled as temporal functions that map time to geometry or value. For example, moving points are modeled as curves in the 3D space (i.e. time to the 2D space).


In [12] a discrete data model implementing the abstract model is defined. For all data types in the abstract model, corresponding *discrete* types whose domains are defined in terms of finite representations are introduced. In the discrete model, moving types are represented by the sliced representation as units.

Definition 1 A data type $\text{moving}(\alpha)$ is a temporally ordered sequence of units. Every unit is a pair $(I, \text{Instant} \rightarrow \alpha)$. The semantic of a unit is that at any time instant during the interval I , the value of the instance can be calculated from the temporal function $\text{Instant} \rightarrow \alpha$. Units are not allowed to temporally overlap, yet gaps are possible (i.e. periods during which the value of the object is undefined). \square

The $\text{moving}(\text{point})$, for example, is modeled as a temporally ordered list of units. Every unit is a pair that consists of a time interval and a linear function in time. The semantics of a unit is that the position of the point at any time instance within the interval is obtained by evaluating the temporal function. This is illustrated in Figure 1.

Figure 1: The sliced representation of a $\text{moving}(\text{point})$

["2003-11-20-06:06" "2003-11-20-06:06:08.692"[, (16229.0 1252.0), (16673.0 1387.0)
["2003-11-20-06:06:08.692" "2003-11-20-06:06:24.776"[, (16673.0 1387.0), (16266.0 1672.0)
["2003-11-20-06:06:24.776" "2003-11-20-06:06:32.264"[, (16266.0 1672.0), (16444.0 1818.0)
["2003-11-20-06:06:32.264" "2003-11-20-06:06:39.139"[, (16444.0 1818.0), (16144.0 2227.0)



The model offers a large number of operations that fall in three classes:

1. Static operations over the non-moving types. Examples are the topological predicates, set operations and aggregations.
2. Spatiotemporal operations offered for the temporal types (e.g. trajectory of a $\text{moving}(\text{point})$, area of a $\text{moving}(\text{region})$).
3. Lifted operations offered for combinations of moving and non-moving types. Basically they are time dependent versions of the static operations.

Lifted operations are obtained by a mechanism called *temporal lifting*. All the static operations defined for non-moving types are uniformly and consistently made applicable to the corresponding moving types. For example, a static predicate and its corresponding lifted predicate are defined as follows.

Definition 2 A *static predicate* is a function with the signature

$$P_1 \times \dots \times P_n \rightarrow \text{bool}$$

where P_i is any static data type (e.g. *integer*, *point*, *region*). \square

Example: FernUni *inside* Hagen.

Definition 3 A *lifted predicate* is a function with the signature

$$P_1 \times \dots \times P_k \times \uparrow P_{k+1} \times \dots \times \uparrow P_n \rightarrow \uparrow \text{bool}$$

where \uparrow is the *moving* type constructor. It can be applied to any static data type and returns its moving counterpart. A lifted predicate is obtained by allowing one or more of the parameters of a static predicate to be of a *moving* data type. Consequently, the return type is a $\text{moving}(\text{bool})$, also denoted *mbool*. \square

Example: Train_RE1206 *inside* Hagen.

The large number of operations in the second and third classes of operations allow for spatiotemporal queries that involve an arbitrary set of spatiotemporal conditions. They fall short, however, of expressing a relative temporal order of the conditions. Spatiotemporal pattern predicates (STPP), that we propose in this paper, allow for expressing such temporal order. The importance of spatiotemporal patterns for many fields of application is illustrated in [9].

Our design builds on the concept of lifted predicates, hence we can easily leverage a considerable part of the available infrastructure.

3 Spatiotemporal Pattern Predicates

In this section we describe the proposed model. We do so using a series of definitions with examples.

Definition 4 A *predicate alias* is a query level unique identifier that identifies a predicate. \square

Example: Train_RE1206 inside Hagen *as* *RE_Hagen_Predicate*.

We compose the spatiotemporal pattern predicate using a set of lifted predicates. Predicate aliases are needed to refer to the different lifted predicates in a user query.

Definition 5 A *temporal connector* is an infix binary predicate/constraint that accepts two *mbool* parameters and enforces a certain temporal arrangement between the pairs of their units. The operation has the signature

$$mbool \times mbool \rightarrow bool$$

\square

Assume the expression $P \odot Q$ where P and Q are lifted predicates each returning an *mbool* value. Let the set of time intervals of the units during which P is true be called P^{true} and similarly Q^{true} for Q . The temporal connector \odot is evaluated by calculating the Cartesian product of P^{true} and Q^{true} , then applying the temporal constraint on every pair of time intervals. The temporal connector is satisfied if one or more pairs satisfy the constraint and we call such a pair a *supported assignment*. Temporal connectors can be simple or vectors as shown in Definitions 6, 7.

Definition 6 *Simple temporal connectors* are temporal connectors that enforce only one interval relationship. The set of simple temporal connectors is inspired from the 13 Allen's operators [5] with the addition that the intervals may degenerate into time instants. Hence 26 simple temporal connectors are possible. \square

We use a simple language for writing the simple connectors. The letters *aa* denote the begin and end time instants of the left hand side argument. Similarly *bb* are the begin and end of the right hand side. The order of letters describes the constraint, that is, a sequence *ab* means $a < b$. The dot symbol denotes the equality constraint, hence, the sequence *a.b* means $a = b$. Table 1 lists the 26 possible temporal constraints with their graphical illustration.

A temporal connector can alternatively be written as a vector of simple temporal connectors.

Definition 7 A *vector temporal connector* is a set of simple temporal connectors. Vectors are interpreted as the disjunction of their constituent simple temporal connectors. Hence 2^{26} vector temporal connectors are possible. In the following we use the operator *vec* as a tool for constructing these vectors (e.g. `vec(aabb, abab, a.bab)`). \square

Now we define the spatiotemporal pattern predicate.

Table 1: Simple temporal connectors

Connector	Illustration	Connector	Illustration
Both arguments are intervals (Allen's operators)			
aabb	aaaa bbbb	abba	aaaaaaaa bbbb
bbaa	aaaa bbbb	a.bab	aaaa bbbbbbbb
aa.bb	aaaa bbbb	a.bba	aaaaaaaa bbbb
bb.aa	aaaa bbbb	baa.b	aaaa bbbbbbbb
abab	aaaa bbbb	aba.b	aaaaaa bbbb
baba	aaaa bbbb	a.ba.b	aaaa bbbb
baab	aaaa bbbbbbbb		
The LHS argument is an instant			
a.ab	a bbbb	bb.a.a	a bbbb
a.a.bb	a bbbb	bba.a	a bbbb
ba.ab	a bbbb		
The RHS argument is an instant			
b.baa	aaaa b	aa.b.b	aaaa b
b.b.aa	aaaa b	aab.b	aaaa b
ab.ba	aaaa b		
Both arguments are instants			
a.ab.b	a b	b.ba.a	a b
a.a.b.b	a b		

Definition 8 A spatiotemporal pattern predicate (STPP) is a triple $\langle t, L, C \rangle$ where t is a tuple containing at least one moving object, L is a set of aliased lifted predicates that apply to the moving object in t and C is a set of binary constraints. Every binary constraint in C is in the form $L_i \odot L_j$ where \odot is a temporal connector (simple or vector). The predicate is fulfilled if and only if the evaluations of the set of aliased predicates L fulfill all the constraints in C . In SQL, we use the operator *pattern* to write spatiotemporal pattern predicates. \square

Example: A query for possible bank robbers may look for the cars which entered a gas station, kept close to the bank for a while, then drove away fast. The query may be written as follows:

```

SELECT c.licencenumber
FROM cars c, landmark l
WHERE l.type = "gas station" and
  pattern([ c.trip inside l.region as gas,
    distance(c.trip, bank) < 50.0 as bnk,
    speed(c.trip) > 100000 as leaving],
    [gas vec(aabb) bnk, bnk vec(abab, aa.bb, aabb) leaving])

```

The predicate *pattern* in this query gets, as input, single tuples from the set of tuples generated by the *SELECT FROM* clauses. This is already the first parameter to the STPP. Every tuple contains the attribute *trip*, a *moving(point)* that stores the car's trajectory. The STPP includes three lifted predicates with aliases *gas*, *bnk*, and *leaving*, each of which returns a *moving(bool)*. Two constraints apply to the lifted predicates; a simple temporal connector between the first and the second predicate, and a vector temporal connector between the second and third predicate. The first constraint states that the car came close to the bank after it has left the gas station. The second constraint is a bit more tricky. We wish to say that the car left the bank area quickly. This means that the car started fast, or may have started normally and then sped up after a while. Therefore we use a vector temporal connector to state all possibilities.

For syntactic elegance, we allow for defining names for the temporal connectors. Using the *let* statement, it is possible to write

```
let then = vec(abab, aa.bb, aabb);
let later = vec(aabb);

SELECT c.licencenumber
FROM cars c, landmark l
WHERE l.type = "gas station" and
      pattern([c.trip inside l.region as gas,
               distance(c.trip, bank) < 50.0 as bnk,
               speed(c.trip) > 100000 as leaving],
              [gas later bnk, bnk then leaving])
```

4 Evaluating Spatiotemporal Pattern Predicates

The spatiotemporal pattern predicate can be modeled as a Constraint Satisfaction Problem (CSP).

Definition 9 Formally, a *constraint satisfaction problem* is defined as a triple $\langle X, D, C \rangle$, where X is a set of variables, D is a set of initial domains and C is a set of constraints. Each variable $X_i \in X$ has a non-empty domain $D_i \in D$. CSP algorithms remove values from the domains during evaluation once it is discovered that the values cannot be part of a solution. Each constraint involves a subset of variables and specifies the allowable combinations of values for this subset. An assignment for a subset of variables is *supported* if it satisfies all constraints. A solution to the CSP is in turn a *supported assignment* of all variables. \square

The definition of CSP maps to the definition of spatiotemporal pattern predicates. The variables map to the lifted predicates and their evaluations are the variable domains. The temporal connectors are the constraints, hence only binary constraints (i.e. constraints involving exactly two variables) exist.

A CSP having only binary constraints is called *binary CSP* and can be represented graphically in a *constraints graph*. The nodes of the graph are the variables and the links are the binary constraints. Two nodes are linked if they share a constraint. The neighborhood of a variable in the constraints graph are all variables that are directly linked to it. The spatiotemporal pattern predicate is fulfilled if and only if its corresponding CSP has at least one supported assignment.

CSPs are usually solved using variants of the backtracking algorithm. The algorithm is a depth-first tree search that starts with an empty list of assigned variables and recursively tries to find a solution (i.e. a supported assignments of all variables). In every call, backtracking adds a new variable to its list and tries all the possible assignments. If an assignment is supported, a new recursive call is made. Otherwise the algorithm backtracks to the last assigned variable. The algorithm runs in exponential time and space.

Constraint propagation methods [6] (also called local consistency methods) can reduce the domains before backtracking to improve the performance. Examples are the ARC Consistency and Neighborhood Inverse Consistency (NIC) algorithms. They detect and remove some values from the variable domains that cannot be part of a solution. Local consistency algorithms do not guarantee backtrack-free search.

To have the nice property of backtrack-free search one would need to enforce n -consistency (equivalent to global consistency), which is again exponential in time and space.

The solvers for CSPs assume that the domains of the variables are known in advance. This is, however, a precondition that we wish to avoid. In STPP, calculating the domain of a variable is equivalent to evaluating the corresponding lifted predicate. Since this can be expensive, we wish to delay the evaluation of the domains.

The proposed algorithm *Solve Pattern* tries to solve the sub-CSP of $k - 1$ variables (CSP_{k-1}) first and then to extend it to CSP_k . Therefore, an early stop is possible if a solution to the CSP_{k-1} cannot be found.

It uses three data structures: the SA list (for Supported Assignments), the Agenda and the Constraint Graph. The Agenda keeps a list of variables that are not yet consumed by the algorithm. One variable from the Agenda is consumed in every iteration. Every supported assignment in the SA list is a solution for the sub-CSP consisting of the variables that have been evaluated so far. In iteration k there are $k - 1$ previously evaluated variables and one newly evaluated variable (X_k with domain D_k). Every entry in SA at this iteration is a solution for the CSP_{k-1} . To extend the SA, the Cartesian product of SA and D_k is calculated. Then only the entries that constitute a solution for CSP_k are kept in SA. CSP_k is constructed using the consumed variables and their corresponding constraints in the constraint graph.

Algorithm Solve Pattern

input: variables, constraints

output: whether the CSP consistent or not

1. Clear SA, Agenda and Constraint Graph
2. Add all variables to Agenda
3. Add all constraints to the Constraint Graph
4. WHILE Agenda not empty
 - (a) Pick a variable X_i from the Agenda
 - (b) Calculate the variable domain D_i (i.e. evaluate the corresponding lifted predicate)
 - (c) Extend SA with D_i
 - (d) IF SA is empty return NotConsistent
5. return Consistent

Algorithm Extend

input: i, D_i ; the index and the domain of the newly evaluated variable

1. IF SA is empty
 - (a) FOREACH interval I in D_i
 - i. INSERT a new row sa in SA having sa[i]= I and undefined for all other variables
- ELSE
 - (a) set SA = the Cartesian product $SA \times D_i$
 - (b) Construct the subgraph CSP_k that involves the variables in SA from the Constraint Graph.
 - (c) FOREACH row sa in SA
 - i. IF sa does not satisfy the CSP_k , remove sa from SA

The methodology for picking the variables from the Agenda has a big effect on the run time. The best method will choose the variables so that inconsistencies are detected soon. For example, suppose an STPP having four predicates with aliases u , v , w , and x . The constraints are $u \text{ vec}(abab) x$, $v \text{ later } x$, and $w \text{ vec}(bb.a.a) x$. If the variables are picked in sequential order u , v , w , then x , the space and time costs are the maximum. Since u , v , and w are not connected by any constraints, the SA is populated by the Cartesian product of their domains in the first three iterations. The actual filter to SA starts in the fourth iteration after x is picked.

The function that picks the variables from the Agenda chooses the variables according to their *connectivity rank* in the Constraint Graph. The connectivity rank of a variable is the summation of its individual connectivities in the Constraint Graph. If a given variable is connected to an Agenda variable with a constraint, it gets 0.5 connectivity score for this constraint. This means that evaluating this variable contributes 50% in evaluating the constraint because the other variable is still not evaluated. If the other variable in the constraint is a non-Agenda variable (i.e. a variable that is already evaluated), the connectivity score is 1. Back again to the example, in the first iteration, the variables u , v , and w have connectivity ranks of 0.5, whereas x has 1.5. Therefore, x is picked in the first iteration. In the second iteration u , v , and w have equal connectivity ranks of 1, so the algorithm picks any of them.

This variable picking methodology tries to maximize the number of evaluated constraints in every iteration with the hope that they filter the SA list and detect inconsistencies as soon as possible.

The time cost of the *Solve Pattern* algorithm is

$$\sum_{i=1}^n \prod_{k=1}^i d_k \times e_k$$

where n is the number of variables, d_k is the number of values in the domain of the k^{th} variable and e_k is the number of constraints in CSP_k . The storage cost is

$$\sum_{i=1}^n \prod_{k=1}^i d_k$$

The algorithm runs in $O(ed^n)$ and takes $O(d^n)$ space.

The exponential time and space costs are not prohibitive in this case. This is because the calculations done within the iterations are simple comparisons of time instances. Moreover, the number of variables in an STP query is expected to be less than 8 in the normal case. The *Solve Pattern* algorithm is more focused on minimizing the number of evaluated lifted predicates (statement 4.b of the algorithm). The cost of evaluating the lifted predicates varies, but it is expected to be expensive because the evaluation usually requires retrieving and processing the complete trajectory of the moving object. The run time analysis of many lifted predicates is illustrated in [7].

5 Extending the Definition of STPP

Back to the example of bank robbers, a sharp eyed reader will notice that the provided SQL statement can retrieve undesired tuples. Suppose that long enough trajectories are kept in the database. A car that entered a gas station in one day, passed close to the bank in the next day, and in a third day sped up will be part of the result. To avoid this, we would like to constrain the period between leaving the gas station till speeding up to be at most 1 hour.

Indeed the proposed design is flexible so that such an extension is easy to integrate. The idea is that after the STPP is evaluated, the SA data structure contains all the supported assignments. As illustrated before, a supported assignment assigns an interval to each lifted predicate during which it is satisfied. At the same time the interval values of all variables satisfy all the constraints in the STPP. Now that we

know the time intervals, we can impose more constraints on them. For example, we state that the period between leaving the gas station (first predicate) till speeding up (third predicate) must be at most 1 hour.

To implement the extension, two changes are required:

1. Change step 5 in the *Solve Pattern* algorithm to return SA.
2. Extend the definition of STPP (Definition 8) by Definition 10.

Definition 10 An *extended spatiotemporal pattern predicate*, denoted *patternex* in SQL, is a quadruple $\langle t, P, C, f \rangle$ where $\langle t, P, C \rangle$ is an STPP and f is a boolean expression that filters the list of supported assignments SA after solving the STPP. \square

The processing of extended STPP is done in two parts that both must succeed. The first part processes the triple $\langle t, P, C \rangle$ as described above. The second part, which is processed only after the success of the first part, evaluates the boolean expression. Hence, conditions on the list SA are possible.

Syntactically, the user is provided with two functions *start* and *end*. They accept a predicate alias and return the start and end of the assigned interval. Note that SA may include several supported assignments. The expression f is tested iteratively against each entry in SA till it is true; otherwise, the extended STPP fails.

Example: The SQL for the bank robbers example is rewritten as follows:

```
SELECT c.licencenumber
FROM cars c, landmark l
WHERE l.type = "gas station" and
      patternex([c.trip inside l.region as gas,
                 distance(c.trip, bank) < 50.0 as bnk,
                 speed(c.trip) > 100000 as leaving],
               [gas later bnk, bnk then leaving],
               start(leaving) - end(gas) < 1)
```

More complex conditions can be issued in the second part. The time intervals can be used, for example, to retrieve parts from the moving object trajectory to express additional spatial conditions. For example, the query for possible bank robbers may more specifically look for the cars which entered a gas station, made a round or more surrounding the bank, then drove away fast. To check that the car made a round surrounding the bank, a possible solution is to check the part of the car trajectory close to the bank for self intersection. The query may be written as follows

```
SELECT c.licencenumber
FROM cars c, landmark l
WHERE l.type = "gas station" and
      patternex([c.trip inside l.region as gas,
                 distance(c.trip, bank) < 50.0 as bnk,
                 speed(c.trip) > 100000 as leaving],
               [gas later bnk, bnk then leaving],
               isSelfIntersecting(
                 trajectoryPart(c.trip, start(bnk), end(bnk))) and
                 (start(leaving) - end(bnk)) < 1)
```

where *trajectoryPart* computes the spatial trajectory of the moving object between two time instants and *isSelfIntersecting* checks a line for self intersection.

6 Related Work

Although our technical development is not yet finished, we discuss related work already at this point of the paper. This is because none of the related work addresses issues of query optimization and system integration that we study below.

A theory and a design for spatiotemporal pattern queries, although important, are not yet well established. Only few proposals exist. In [19], a model that relies on a discrete view of the spatiotemporal space is presented. The 2D space is partitioned in a finite set of user defined partitions, called zones. The time axis is partitioned into constant-sized intervals. Every spatial partition (zone) is given a label. The trajectories are then represented as strings of labels. If the moving object entered in a zone a , for example, the character a is appended to its trajectory. If the same object moved to zone b and stayed there for three time units then the string bbb is appended and so on. The pattern is composed in the user query as a formal expression, which is then evaluated using efficient string matching techniques.

Their approach is not general in the sense that the space and time have to be partitioned. The partitioning depends on the intended application and has to be done in advance. Moreover, only patterns that describe the change of location of moving points can be expressed. The approach leaves back all other kinds of predicates (e.g. topological, metric, distance) as well as other types of moving objects (e.g. moving regions).

In [17], an index structure and efficient algorithms to evaluate STP queries that consist of spatial and neighborhood predicates is presented. They addressed the problem of conjoint neighborhood queries (e.g. find all objects that were as close as possible to A at time T_1 then were as close as possible to B at time T_2). The two NN conditions in this query have to be evaluated conjointly. In other words, an object which minimizes the sum of the two distances at the two time points is the answer of this query.

Again the approach addresses only limited kinds of predicates, and handles moving points only. Moreover, it is not extensible in the context of systems. The evaluation of the predicates (lifted predicates in our case) is an integral part of the evaluation of the STPP. Hence, the algorithms have to be extended for every predicate. An extensible design for a DBMS requires separating the evaluation of the predicates from the evaluation of the STPP. This is not the case in the approach of [17]. On the other hand, the evaluation of the predicates within the STPP allows for more efficient evaluation. It allows also for the conjoint neighborhood queries that are not possible in our approach.

The series of publications [10], [11], [9], and [22] provide a concrete formalism for spatiotemporal predicates and developments. A spatiotemporal development is a composite structure built as an alternating sequence of spatiotemporal and spatial predicates, and they are themselves spatiotemporal predicates. They describe the change, wrt. time, in the spatial relationship between two moving objects. Consider, for example, a moving point MP and a moving region MR . The development MP *Crosses* MR is defined as:

`Crosses= Disjoint meet Inside meet Disjoint`

where *meet* is a spatial predicate that yields true when its two arguments touch each other, and *Disjoint* is a spatiotemporal predicate that yields true when its two arguments are always spatially disjoint. The spatiotemporal predicates, denoted by being capitalized, differ from the spatial predicates in that, the former hold at time intervals while the later hold at instants. Spatiotemporal developments consider two spatiotemporal objects and precisely describe the change in their topological relationship. In this sense, they can be thought of as describing *micro STP*.

Complex spatiotemporal developments can be defined by means of regular expressions over other spatial and spatiotemporal predicates [10], [11]. To evaluate a development for a pair of moving objects, one has to find a connected alternating sequence of time intervals and instants during which the constituting sequence of spatiotemporal and spatial predicates hold. If such partitioning of time cannot be found, the development yields false.

In contrast, our approach does not partition the moving objects. We evaluate the predicates for the complete movement, by means of lifted predicates, then evaluate the STPP over the resulting moving booleans. This difference in formalizing STP predicates has the following consequences:

1. The spatiotemporal developments by Erwig and Schneider [10], [11] are capable of describing the change in the topological relationship between two moving objects (i.e. topological predicates). They cannot handle other kinds of predicates.

2. A pattern that is described by a spatiotemporal development is restricted to two moving objects. A natural way of describing the movement pattern of an object would involve its interactions with many other objects in the spatiotemporal space. This is inherent to our approach because the lifted predicate can be independently parameterized.
3. The formalism of spatiotemporal developments requires that the constituting predicates are fulfilled in a connected sequence of time intervals and instants. Our approach can more freely express all the possible interval relationships.
4. Spatiotemporal developments, in contrast to our approach, can not express patterns involving non-spatial moving objects (e.g. *moving(integer)*).

The first and third limitations of spatiotemporal predicates were partially discussed by Erwig in [9]. He outlined some ideas (without proposing a full fledged design) to generalize the spatiotemporal developments towards spatiotemporal patterns. The last point shows another kind of patterns that is not addressed by the three reviewed approaches, that is temporal patterns. It is possible to use a *moving(integer)*, for example, to encode the T-shirt number of the player who possesses the ball in a soccer game. We would like to be able to express patterns on this (e.g. find all the attacks where player 10 passes to 4, then 4 passes to 20). Although such purely temporal patterns (i.e. patterns not involving moving objects) are not in the focus of our design, they are a nice result of the proposed modular two steps design. Since the lifted predicate (first step) can process temporal moving objects, the STP predicate (second step) leverages this capability for free.

Furthermore, the extended STP predicate in Section 5 increases the expressive power by making it possible for the user to access the fulfillment times of the lifted predicates.

In the other hand, our approach can not currently express repetitions and alternations in the pattern, in contrast to Mouza and Rigaux [19], and Erwig and Schneider [11]. Difficulties are in the language integration in SQL, and the modification of our CSP-based evaluation algorithm. We address this limitation in the future work.

7 Optimizing Spatiotemporal Pattern Predicates

In Section 4 we explained the evaluation of the spatiotemporal pattern predicate. The proposed algorithm is efficient because it avoids the unnecessary evaluation of lifted predicates. In the context of large-scale DBMS, this is not enough. Obviously for an efficient execution of pattern queries on large databases the use of indexes is mandatory. It should be triggered by the query optimizer during the creation of the executable plans.

In this section, we demonstrate a generic procedure for integrating the STPP with query optimizers. We do not assume a specific optimizer or optimization technique. The optimizer is however required to have some basic features that will probably be available in any query optimizer. In the following subsection, we describe these basic assumptions.

7.1 Query Optimization

A typical query optimizer contains two basic modules; the *rewriter* and the *planner* [18]. The rewriter uses some heuristics to transform a query into another equivalent query that is, hopefully, more efficient or easier to handle in further optimization phases. The planner creates for the user query (or the rewritten version) the set of possible *execution plans* (possibly restricted to some classes of plans). Finally it applies a selection methodology (e.g. cost based) to select the best plan.

We assume that the query optimizer contains the rewriter and the planner modules. We also assume that it supports the data types and operations on moving objects, in SQL predicates as described in [16] and [12].

7.2 Query Optimization for Spatiotemporal Pattern Predicates

One observation that we like to make clear is that the STPP itself does not process database objects directly. Instead, the first operation applied is the evaluation of the lifted predicates that compose the STPP. The idea, hence, is to design a general framework for optimizing the lifted predicates within the STPP. This framework should trigger the optimizer to use the available indexes for the currently supported lifted predicates as well as for those that might be added in the future. It should utilize the common index structures. Although specialized indexes, as in [17], can achieve higher performance, the overhead of maintaining them within a system is high and they only serve specific purposes, which makes them unfavorable in the context of systems.

The idea is to add each of the lifted predicates, in a modified form, as an extra *standard predicate* to the query, that is, a predicate returning a boolean value. The standard predicate is chosen according to the lifted predicate, so that the fulfillment of the standard predicate implies that the lifted predicate is fulfilled at least once. This is done during query rewriting. The additional standard predicates in the rewritten query trigger the planner to use the available indexes. To illustrate the idea, the following query shows how the bank robbers query in Section 3 is rewritten.

```
SELECT c.licencenumber
FROM cars c, landmark l
WHERE l.type = "gas station" and
      pattern([c.trip inside l.region as gas,
              distance(c.trip, bank) < 50.0 as bnk,
              speed(c.trip) > 100000 as leaving],
            [gas later bnk, bnk then leaving])
and
c.trip passes l.region and
sometimes(distance(c.trip, bank) < 50.0) and
sometimes(speed(c.trip) > 100000)
```

The three lifted predicates in the STPP x inside y , $\text{distance}(x, y) < z$, and $\text{speed}(x) < y$ are mapped to the standard predicates x passes y , $\text{sometimes}(\text{distance}(x, y) < z)$, and $\text{sometimes}(\text{speed}(x) < y)$, respectively. Here *sometimes(.)* is a predicate that accepts a *moving(bool)* and yields true if the argument ever assumes true during its lifetime, otherwise false. Each of the standard predicates ensures that the corresponding lifted predicate is fulfilled at least once, a necessary but not sufficient condition for the *pattern* predicate to be fulfilled. Clearly, the rewritten query is equivalent to the original query.

The choice of the standard predicate depends on the type of the lifted predicate and the types of the arguments. For example, the lifted spatial range predicates (i.e. the spatial projection can be described by a box) are mapped into the *passes* standard predicate. The *passes* predicate [16], in this example, is fulfilled if the car *c.trip* ever passed the gas station *l.region*. If *passes* fails, then we know that *inside* is never true and that *pattern* will also fail. The planner should have for the added *passes* predicate already some optimization rule available (e.g. use a spatial R-tree index when available). In Section 9.2.2 we show an optimized query written in the SECONDO executable language.

To generalize this solution, we define a table of mappings between the lifted predicates (or groups of them) and the standard predicates. Clearly, this mapping is extensible for the lifted predicates that can be introduced in the future. The mapping for the set of lifted predicates proposed in [16] is shown in Table 2.

For the lifted spatial range predicates, they map into *passes* and the available translation rules for *passes* do the rest. The $\text{distance}(x, y) < z$ is conceptually equivalent to a lifted spatial range predicate, where the spatial range is the minimum bounding box of the static argument extended by z in every side. Other types of lifted predicates are mapped into *sometimes*. We need to provide translation rules that translate *sometimes(.)* into index lookups. For every type of lifted predicates, one such translation rule is required. For example, the *sometimes(Pred)*, where *Pred* is a lifted left range predicate, searches for a

Table 2: Mapping lifted predicates into standard predicates.

Lifted Predicates	Type	Standard Predicates
$\sigma = \alpha$ $\underline{mpoint} \times \underline{point} \rightarrow \underline{mbool}$ $\underline{mregion} \times \underline{region} \rightarrow \underline{mbool}$ $\sigma \text{ inside } \alpha$ $\underline{mpoint} \times \underline{region} \rightarrow \underline{mbool}$ $\underline{mpoint} \times \underline{points} \rightarrow \underline{mbool}$ $\underline{mpoint} \times \underline{line} \rightarrow \underline{mbool}$ $\underline{mregion} \times \underline{region} \rightarrow \underline{mbool}$ $\underline{mregion} \times \underline{points} \rightarrow \underline{mbool}$ $\underline{mregion} \times \underline{line} \rightarrow \underline{mbool}$ $\sigma \text{ intersects } \alpha$ $\underline{mregion} \times \underline{points} \rightarrow \underline{mbool}$ $\underline{mregion} \times \underline{region} \rightarrow \underline{mbool}$ $\underline{mregion} \times \underline{line} \rightarrow \underline{mbool}$	lifted spatial range	$\sigma \text{ passes } \alpha$
$\sigma = \alpha$ $\underline{mint} \times \underline{int} \rightarrow \underline{mbool}$ $\underline{mbool} \times \underline{bool} \rightarrow \underline{mbool}$ $\underline{mstring} \times \underline{string} \rightarrow \underline{mbool}$ $\underline{mreal} \times \underline{real} \rightarrow \underline{mbool}$	lifted equality	sometimes ($\sigma = \alpha$)
$\sigma \leq \alpha, \sigma < \alpha$ $\underline{mint} \times \underline{int} \rightarrow \underline{mbool}$ $\underline{mbool} \times \underline{bool} \rightarrow \underline{mbool}$ $\underline{mstring} \times \underline{string} \rightarrow \underline{mbool}$ $\underline{mreal} \times \underline{real} \rightarrow \underline{mbool}$	lifted left range	sometimes ($\sigma \leq \alpha$), sometimes ($\sigma < \alpha$)
$\sigma \geq \alpha, \sigma > \alpha$ $\underline{mint} \times \underline{int} \rightarrow \underline{mbool}$ $\underline{mbool} \times \underline{bool} \rightarrow \underline{mbool}$ $\underline{mstring} \times \underline{string} \rightarrow \underline{mbool}$ $\underline{mreal} \times \underline{real} \rightarrow \underline{mbool}$	lifted right range	sometimes ($\sigma \geq \alpha$), sometimes ($\sigma > \alpha$)
$\text{distance}(\sigma, \alpha) < \text{threshold}$ $\underline{mpoint} \times \underline{region} \rightarrow \underline{mreal}$ $\underline{mpoint} \times \underline{point} \rightarrow \underline{mreal}$ $\underline{mregion} \times \underline{point} \rightarrow \underline{mreal}$ $\underline{mregion} \times \underline{region} \rightarrow \underline{mreal}$	lifted spatial range	$\sigma \text{ passes } \text{enlargeRect}(\text{bbox}(\alpha), \text{threshold}, \text{threshold})$
Other lifted predicates, P		sometimes (P)

B-tree defined on the units of the moving object, and performs a left range search in the B-tree. We show examples for these translation rules within SECONDO in Section 8.2.

This two steps optimization helps to develop a general framework for optimizing the *sometimes*(.) predicate, which may also appear directly in the user queries. Note that we can alternatively rewrite all lifted predicates into *sometimes*(.), and provide translation rules accordingly. It remains an implementation decision, which approach to use.

8 The Implementation in SECONDO

SECONDO [4], [14], [15] is an extensible DBMS platform that does not presume a specific database model. Rather it is open for new database model implementations. For example, it should be possible to implement relational, object-oriented, spatial, temporal, or XML models.

SECONDO consists of three loosely coupled modules: the kernel, GUI and query optimizer. The kernel includes the command manager, query processor, algebra manager and storage manager. The

kernel may be extended by algebra modules. In an algebra module one can define new data types and/or new operations. The integration of the new types and/or operations in the query language is then achieved by adding syntax rules to the command manager.

The SECONDO kernel accepts queries in a special syntax called SECONDO *executable language*. The SQL-like syntax is provided by the optimizer. For more information about SECONDO modules see [4] and [3]. For more information about extending SECONDO see the documentation on [2].

If it is the case that a new data type needs a special graphical user interface (GUI) for display, the SECONDO GUI module is also extensible by adding viewer modules. Several viewers exist that can display different data types. Moving objects, for example, are animated in the *Hoes* viewer with a time slider to navigate forwards and backwards.

A large part of the moving objects database model presented in [16], [12], [7], that we also assume in the paper, is realized in SECONDO. That is, the current SECONDO version 2.9.1 includes the algebra modules, the viewer modules, and the optimizer support for moving objects. In the following subsections, we describe the implementation of our STPP in SECONDO 2.9.1. This implementation is available as a SECONDO Plugin as explained in Section 11.

8.1 Extending the Kernel

We have implemented the STPP in the SECONDO kernel in a new algebra module called *STPatternAlgebra*. The algebra contains:

1. One data type *stvector*. The class represents the temporal connectors. Simple temporal connectors are treated as a special case of vector temporal connectors (i.e. vectors having only one element). The SECONDO operator *vec* is used to create an *stvector* instance. The operator accepts a set of strings from Table 1, and constructs the *stvector* instance accordingly.

Example: `vec("aabb" , "a.abb" , "a.a.bb")`.

2. The *stconstraint* operator. The operator represents a constraint within the STPP. The signature of the operator is

$$\underline{string} \times \underline{string} \times \underline{stvector} \rightarrow \underline{bool}$$

The first and second parameters are the aliases for two lifted predicates.

Example: `stconstraint("predicate1" , "predicate2" , vec("a.a.bb"))`.

3. The *stpattern* operator. The operator implements the STPP, Section 4. It has the signature

$$\underline{tuple} \times \underline{AliasedPredicateList} \times \underline{ConstraintList} \rightarrow \underline{bool}$$

where the *AliasedPredicateList* is a list of lifted predicates, each of which has an alias, and the *ConstraintList* is a list of the *stconstraint* operators.

4. The *stpatternex* operator. The operator implements the extended STPP, Section 5. It has the signature

$$\underline{tuple} \times \underline{AliasedPredicateList} \times \underline{ConstraintList} \times \underline{bool} \rightarrow \underline{bool}$$

5. The *start* and the *end* operators described in Section 5. They accept a *string* representing a predicate alias and return the start or end of the corresponding time interval in the *SA* list. The operators have the signature

$$\underline{string} \rightarrow \underline{instant}$$

Using these operators, the query for bank robbers can be written in SECONDO executable language as follows:

```

query cars feed {c}
landmark feed {l}
  filter[.type_l = "gas station"]
product
filter[.
  stpatternex[gas: .trip_c inside .region_l,
    bnk: distance(.trip_c, bank) < 50.0,
    leaving: speed(.trip_c) > 100000;
  stconstraint("gas", "bnk", vec("aabb")),
  stconstraint("bnk", "leaving", vec("abab", "aa.bb", "aabb"))];
  duration2real(start("leaving") - end("gas")) < (1/24) ]]
consume

```

where *feed* is a postfix operator that scans a relation sequentially and converts it into a stream of tuples. The query performs a cross product between the tuples of the *cars* relation and the tuples of *landmark* relation that has the value “gas station” in their *type* attribute. The resulting tuple stream after the cross product is filtered using the extended STP predicate *stpatternex*. Finally, the *consume* operator converts the resulting tuple stream into a relation, so that it can be displayed.

8.2 Extending the Optimizer

The SECONDO optimizer is written in Prolog. It implements an SQL-like query language which is translated into an optimized query in SECONDO executable language. The SECONDO optimizer includes a separate rewriting module that can be switched on and off by setting the optimizer options. The planner implements a novel cost based optimization algorithm which is based on *shortest path search in a predicate order graph*. The predicate order graph (POG) is a weighted graph whose nodes represent sets of evaluated predicates and whose edges represent predicates, containing all possible orders of predicates. For each predicate edge from node x to node y , so-called plan edges are added that represent possible evaluation methods for this predicate. Every complete path via plan edges in the POG from the bottom-most node (i.e. zero evaluated predicates) till the top-most node (i.e. all predicates evaluated) represents a different execution plan. Different paths/execution plans represent different orderings of the predicates and different evaluation methods. The plan edges of the graph are weighted by their estimated costs, which in turn are based on given selectivities. Selectivities of predicates are either retrieved from prerecorded values, or estimated by sending selection or join queries on small samples of the involved relations to the SECONDO kernel and reading the cardinality of the results. The algorithm is described in more detail in [15] as well as in the SECONDO programmers guide [2].

Our extension to the optimizer has three major parts: query rewriting, operator description, and translation rules. In the query rewriting, we choose to rewrite all the lifted predicates into *sometimes(.)*. This is because an accurate rewriting based on the mapping in Table 2 requires that we know the data types of the arguments. The SECONDO optimizer knows the data types only after query rewriting is done.

Following are the Prolog rules that do the rewriting:

```

inferPatternPredicates([], []).

inferPatternPredicates([Pred|Preds],
  [sometimes(Pred)|Preds2]):-
  assert(removefilter(sometimes(Pred))),
  inferPatternPredicates(Preds,Preds2).

```

where the *inferPatternPredicate* accepts the list of the lifted predicates within the STPP as a first argument, and yields the a list of rewritten predicates in the second argument. The additional *sometimes(.)*

predicates are kept in the table `removefilter(.)`, so that it is possible to exclude them from the executable plan afterwards.

In the operator descriptions, we annotated the lifted predicates by their types (e.g. lifted left range) as in Table 2. Then we provided translation rules for *sometimes(.)* for every type of lifted predicates. Following is an example for such a rule:

```
indexselectLifted(arg(N), Pred ) =>
  gettuples(rdup(sort(windowintersectsS(
    dbobject(IndexName), BBox))), rel(Name, *))
:-
  Pred =..[Op, Arg1, Arg2],
  ((Arg1 = attr(_, _, _), Attr= Arg1) ;
   (Arg2 = attr(_, _, _), Attr= Arg2)),
  argument(N, rel(Name, )),
  getTypeTree(Arg1, _, [_, _, T1]),
  getTypeTree(Arg2, _, [_, _, T2]),
  isLiftedSpatialRangePred(Op, [T1, T2]),
  (
    ( memberchk(T1, [rect, rect2, region, point, line, points, sline]),
      BBox= bbox(Arg1)
    );
    ( memberchk(T2, [rect, rect2, region, point, line, points, sline]),
      BBox= bbox(Arg2)
    )
  ),
  hasIndex(rel(Name, _), Attr, DCindex, spatial(rtree, unit)),
  dcName2externalName(DCindex, IndexName).
```

where this rule translates the *lifted spatial range* predicates into an R-tree window query, as indicated in the rule header. The `=>` operator can be read as *translates into*. It means that the expression to the right is the translation of the expression to the left, if the conditions in the rule body hold. The body of the rule starts by inferring the types of the arguments of the lifted predicate within the *sometimes(.)*. Then it uses them to make sure that the predicate is of the type *lifted spatial range*. Finally, it checks whether a spatial R-tree index on the involved relation and attribute is available in the catalog. It tries to find a spatial R-tree built on the units of the moving object. Similar translation rules are provided for other types of indexes. The optimized query in Section 9.2.2 shows the effect of these translation rules.

9 Experimental Evaluation

We proceed with an experimental evaluation of the proposed technique. The intention is to give an insight into the performance. It is clear that the runtime of an STP predicate depends on the number and types of the lifted predicates. Therefore, we show two experiments. The first measures only the overhead of evaluating the spatiotemporal pattern predicate. That is, we set the time of evaluating the lifted predicates to negligible values.

In the second experiment, we generate random STP predicates with varying numbers of lifted predicates and constraints and measure the run time of the queries. The experiment also evaluates the optimization of STPP. Every query is run twice; once without invoking the optimizer, and another time with the optimizer being invoked.

The experiments use the *berlintest* database that is available with the free distribution of *SECONDO*. The experiments are run on a *SECONDO* platform installed on a Linux machine. The machine is a Pentium-4 dual-core 3.0 GHz processor with 2 GBytes main memory.

9.1 The Overhead of Evaluating STPP

To perform the first experiment, we add two operators to *SECONDO*; *randommbool* and *passmbool*. The operator *randommbool* accepts an *instant* and creates an *mbool* object whose definition time starts at the given time instant, and consists of a random number of units. The operator *passmbool* mimics a lifted predicate. It accepts the name of an *mbool* database object, loads the object and returns it. More details are given below.

9.1.1 Preparing the Data

This section describes how the test data for the first experiment is created. The *randommbool* operator is used to create a set of 30 random *mbool* instances and store them as database objects. The operator creates *mbool* objects with a random number of units varying between 0 and 20. The first unit starts at the time instant provided in the argument. Every unit has a random duration between 2 and 50000 milliseconds. The value of the first unit is randomly set to *true* or *false*. The value of every other unit is the negation of its preceding unit. Hence, the minimal representation requirement [12] of the moving types in *SECONDO* is met. That is, adjacent units can not be further merged because they have different values.

The 30 *mbool* objects are created by calling `randommbool(now())` 30 consecutive times. This increases the probability that the definition times of the objects temporally overlap.

9.1.2 Generating the Queries

The queries of the first experiment are selection queries consisting of one filter condition in the form of an STPP. The queries are generated with different experimental settings, that is, different numbers of lifted predicates and constraints in the STPP. The number of lifted predicates varies between 2 and 8. The number of constraints varies between 1 and 16. The queries are not generated for every combination. For example, it does not make sense to generate STPP with 2 lifted predicates and 10 constraints. For N lifted predicates, the number of constraints varies between $N - 1$ and $2N$. The rationale of this is that, if the number of constraints is less than $N - 1$, then the constraint network can not be complete (i.e. some predicates are not referenced within constraints). On the other hand, having more than $2N$ constraints increases the probability of contradicting constraints. For every experimental setting, 100 random queries are evaluated and the average run time is recorded.

A query with 3 lifted predicates and 2 constraints, for example, looks like:

```
query thousand feed
  filter[
    stpattern[a: passmbool(mb5),
              b: passmbool(mb13),
              c: passmbool(mb3);
    stconstraint("b", "a", later),
    stconstraint("b", "c", vec("abab")) ]]
count
```

where `query thousand feed` streams the *thousand* relation, which contains 1000 tuples. For every tuple, the STPP *stpattern* is evaluated. Note that the predicate does not depend on the tuples. That is, the same predicate is executed 1000 times in the query. This is to minimize the effect of the time taken by *SECONDO* to prepare for query execution. The lifted predicates are all in the form of `passmbool(X)`, where X is one of the 30 stored random *mbool* objects.

The constraints are generated so that the constraint graph is complete. We start by initializing a set called *connected* having one randomly selected alias. For every constraint, the two aliases are randomly chosen from the set of aliases in the query, so that at least one of them belongs to the set *connected*. The other alias is added to the set *connected* if it was not already a member. After the required number of

constraints is generated, we check the completeness of the graph. If it is not complete, the process is repeated till we get a connected graph. The temporal connector for every constraint is randomly chosen from a set containing 31 temporal connectors namely, the 26 simple temporal connectors in Table 1 and 5 vector temporal connectors (later, follows, immediately, meanwhile, and then) (shown in Appendix A).

Before running the queries, we query for the 30 *mbool* objects so that they are loaded into the database buffer. The measured run times should, hence, show the overhead of evaluating the STPP in SECONDO because other costs are made negligible.

9.1.3 Results

The results are shown in Figure 2. The number of lifted predicates is denoted as N . Increasing the number of lifted predicates and constraints in the STPP does not have a great effect on the run time. This is a direct result of the early pruning strategy in the *Solve Pattern* algorithm. The results show that the evaluation of STPP is efficient in terms of run time.




Figure 2: The overhead of evaluating STPPs

9.2 STPP with Optimization

The second experiment is intended to evaluate the run time of STP queries. It also evaluates the effect of the proposed optimization. Unlike the first experiment, the STPPs in this experiment contain lifted predicates. We generate 10 random queries for every experimental setting and record the average run time. Every query is run twice; without being optimized, and after optimization.

9.2.1 Preparing the Data

The queries use the *Trains20* relation. It is generated by replicating the tuples of the *Trains* relation in the *berlintest* database 20 times. The *Trains* relation was created by simulating the underground trains of the city Berlin. The simulation is based on the real train schedules and the real underground network of Berlin. The simulated period is about 4 hours in one day. The schema of *Trains20* is similar to *Trains* with the additional attribute *Serial*:

Trains20[*Serial*: *int*, *Id*: *int*, *Line*: *int*, *Up*: *bool*, *Trip*: *mpoint*]

where *Trip* is an *mpoint* representing the trajectory of the train. The relation contains 11240 tuples and has a disk size of 158 MB. To evaluate the optimizer, a spatial R-tree index called *Trains20_Trip_sptuni* is built on the units of the *Trip* attribute. A set of 300 points is also created to be used in the queries. The points represent geometries of the top 300 tuples in the *Restaurants* relation in the *berlintest* database.

9.2.2 Generating the Queries

The queries are generated in the same way as in the first experiment. In this experiment, however, we use actual lifted predicates instead of *passmbool*. Every lifted predicate in the STPP is randomly chosen from

1. $\text{distance}(\text{trip}, \text{randomPoint}) < \text{randomDistance}$.
2. $\text{speed}(\text{trip}) > \text{randomSpeed}$.

where *randomPoint* is a *point* object selected randomly from the 300 restaurant points, *randomDistance* ranges between 0 and 50, and *randomSpeed* ranges between 0 and 30. The $\text{distance}(\cdot, \cdot) < \cdot$ is a sample for the lifted predicates that can be mapped into index access, so that we can evaluate the optimizer. While the queries in the first experiment are created directly in the SECONDO executable language, they are created here in SECONDO SQL. It is an SQL-like syntax that looks similar to the standard SQL, but obeys Prolog rules. The main differences are that everything is written in lower case, and lists are placed within square brackets.

Here is one query example from the generated queries:

```
SELECT count(*)
FROM trains20
WHERE pattern([ distance(trip, point170) < 18.0 as a,
                speed(trip) > 11.0 as b],
              [stconstraint("a", "b", vec("b.ba.a"))])
```

where *pattern* is the SQL operator equivalent to *stpattern* in the executable language. The rewritten version of the query as generated by the rewriting module of the SECONDO optimizer is:

```
SELECT count(*)
FROM trains20
WHERE [ pattern([ distance(trip, point170) < 18.0 as a,
                        speed(trip) > 11.0 as b],
                  [stconstraint("a", "b", vec("b.ba.a"))]),
        sometimes(distance(trip, point170) < 18.0),
        sometimes(speed(trip) > 11.0)]
```

Finally, the optimal execution plan is:

```
Trains20_Trip_sptuni
windowintersectsS[ enlargeRect(bbox(point170), 18.0, 18.0)]
sort rdup Trains20  gettuples
  filter[sometimes((distance(.Trip,point170) < 18.0))]
  {0.00480288, 1.69712}
  project[Trip]
  filter[. stpattern[ a: (distance(.Trip, point170) < 18.0),
                    b: (speed(.Trip) > 11.0);
                    stconstraint("a", "b", vec("b.ba.a"))]]
  {0.00480288, 1.49038}
  filter[sometimes((speed(.Trip) > 11.0))]
  {0.883731, 1.48077}
count
```

where the predicates are placed within the *filter[]* operator, which means that they belong to the *where* clause in SQL. The rewriter generates for the two lifted predicates in the original query two standard *sometimes* predicates. The predicate *sometimes*($\text{distance}(\cdot, \cdot) < \cdot$) is handled by the optimizer as a special kind of range predicate. Since the optimizer can find the spatial R-tree index that we created, it is used. The index access part in the query is:

```
Trains20-Trip_sptuni windowintersectsS[enlargeRect(., ., .)]
```

This part expands the minimum bounding box of *point170* by the distance threshold value 18.0. The enlarged box is intersected with the R-tree to get the candidate tuple id's. The rest of the query retrieves the data of the candidate tuples and performs the query. The pairs of numbers between the curly brackets do not affect the semantics of the query. They are estimated predicate selectivities and run time statistics used to help estimate the query execution progress.

9.2.3 Results

In Figure 3, the chart to the left shows the average run times of the non-optimized STP queries. The chart to the right shows the average run times of their optimized counterparts. The N is again the number of lifted predicates. The run times of the optimized STPP are very promising.

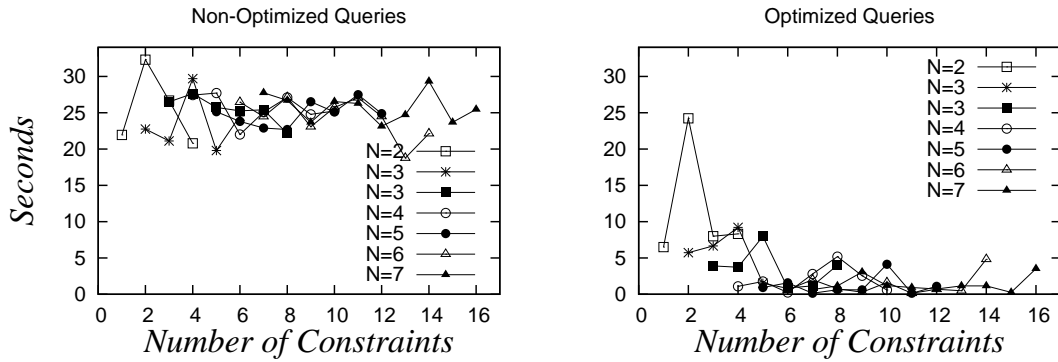


Figure 3: The run times for STP queries on the *Trains20* relation

The high peak in the optimized queries chart at $N = 2$ and *Number of Constraints* = 2 is because it happened that five of the ten generated queries have only *speed(.) < .* predicates. Since the *sometimes(speed(.) < .)* predicate does not map into index access, the average run time for this experimental setting is close to the non-optimized version.

10 Application Examples

To illustrate the expressive power of the proposed approach, we present in the following two subsections more examples for STP queries. Section 10.1 demonstrates a scenario called *Finding Ali*. It is about a kid called Ali, who moves on the street network of Cairo (the capital of Egypt). He makes several trips riding in several cars. We want to query for these cars using their movement profiles.

In Section 10.2, we demonstrate example queries that the reader can try himself/herself in *SECONDO*. The queries are based on the *berlintest* database, that is available with the *SECONDO* distribution. Unlike the first application, the queries are not linked to a single scenario. Hence we can demonstrate STP queries that involve moving points, moving regions, and many kinds of lifted operations.

10.1 Finding Ali

We assume that the road network of Cairo is observed for one month and that the complete trajectories of the cars are stored in the database. The queries assume the following schema:

- Car[PlatesNumber: *string*, Trip: *mpoint*] where Trip is the complete trajectory of the car for the whole observation period.

- Landmark[Name: *string*, Type: *string*, Location: *point*]
- Heliopolis: A *region* object marking the boundary of the district *Heliopolis* where Ali lives.
- AliHome: A *point* object marking Ali's home.
- FamilyHome: A *point* object marking the house of the father's family.
- SportsClub: A *region* object marking the boundary of the sports club in which Ali is a member.

10.1.1 The Go-to-school Trips With the School Bus

The bus starts at the school at 6:00 am - 6:30 am, enters the district Heliopolis at 6:45 am - 7:00 am, stops near Ali's home, picks Ali, exits Heliopolis at 7:45 am - 8:00 am, then goes back to school.

This query can be written without a spatiotemporal pattern predicate. The spatiotemporal window of every predicate is known. It can be expressed as a conjunction of 5 spatiotemporal range predicates (Bus inside School at the time interval [6:00, 6:30] AND Bus inside Heliopolis at the time interval [6:45, 7] ...). We include this as an example of spatiotemporal pattern queries that can be expressed without STPP.

10.1.2 The Evening Trips With Grandfather

Starting from Ali's home, the grandfather drives Ali to the sports club. They stop at the sports club for at least two hours. After the club they go by car to buy some bread, then back home.

```
SELECT c.PlatesNumber
FROM Car c, Landmark l
WHERE l.Type like("%Bakery%") and
  patternex([distance(c.Trip, AliHome) < 20.0 as AtHome,
    c.Trip inside SportsClub as AtClub,
    distance(c.Trip, l.Location) < 20.0 as AtBakery,
    distance(c.Trip, AliHome) < 20.0 as BackHome],
    [AtHome later AtClub,
    AtClub later AtBakery,
    AtBakery later BackHome],
    end("AtClub") - start("AtClub") >= 2.0 and
    daypart(AtHome) = daypart(BackHome))
```

In this query, the extended STPP is used to state that they stayed at least two hours in the sports club and that the whole pattern occurred in one day. Another note is that the query uses the predicate `distance(c.Trip, AliHome) < 20.0` twice with two different aliases. The two aliases are needed to write the constraints. It is the responsibility of the query optimizer to detect this common predicate (i.e. using common sub-expression optimization techniques) and evaluate it only once.

10.1.3 The Weekend Trips With Mother

The mother starts from Ali's home, drives only in main roads, stops near a shopping mall for at most 4 hours then back home. The trip to the mall takes more than 1.5 times the estimated time because the mother uses only main roads. In Cairo it is easier to drive in main roads but they have high traffic.

```
SELECT c.PlatesNumber
FROM Car c, Landmark l
WHERE l.Type like("%Mall%") and
  patternex([distance(c.Trip, AliHome) < 20.0 as AtHome,
    distance(c.Trip, l.Location) < 40.0 as AtMall,
    distance(c.Trip, AliHome) < 20.0 as BackHome],
    [AtHome later AtMall,
```

```

    AtMall later BackHome],
end("AtMall") - start("AtMall") <= 4.0 and
(start("AtMall") - end("AtHome") >
1.5 * EstimatedDriveTime(l.location, AliHome) ))

```

where we assume for simplicity that *EstimatedDriveTime* is a function that computes the normal period that a drive between two places takes. It may do so by finding the shortest path and multiply by the average driving speed.

10.2 The Berlittest Example

In this example, we use the database *berlittest*, more specifically, the *Trains* relation and three newly added relations with the following schemas:

```

SnowStorms[Serial: int, Storm: mregion]
TrainsMeet[Line: int, Uptrip: mpoint, Downtrip: mpoint, Stations: points]
TrainsDelay[Id: int, Line: int, Actual: mpoint, Schedule: mpoint]

```

The *SnowStorms* relation contains 72 tuples, each of which contains a moving region, representing a snow storm that moves over Berlin. The *TrainsMeet* relation is generated from the *Trains* relation. The tuples contain all possible combinations of two trains that belong to the same line and move in opposite directions. The *Stations* attribute represents the train stations of the associated line. The *TrainsDelay* relation is also generated from the *Trains* relation. Each tuple contains the original *Trip* attribute (renamed into *Schedule*), and a delayed copy of it with delays of around 30 minutes. The scripts for creating the three relation and for executing the example queries are available for download as will be explained in Appendix D.

Table 3 lists the lifted operations used within the queries. We have designed the queries so that they illustrate the expressive power of our approach by using various lifted operations to compose complex pattern queries. The table shows only the operator signatures that are used in the queries. The complete list of valid signatures is in [16].

10.2.1 Find the snow storms that passed over the train station *mehringdamm* with speed greater than 40 km/h.

```

SELECT *
FROM   snowstorms
WHERE  pattern([not(isempty(storm at mehringdamm)) as pred1,
               speed(rough_center(storm)) > 40.0 as pred2],
               [stconstraint("pred1", "pred2", together)])

```

where *together* is a vector temporal connector that yields true if the two predicates happen simultaneously.

10.2.2 Find the snow storms that could increase their area over 1/4 square km during the first traversed 5 km.

```

SELECT *
FROM   snowstorms
WHERE  pattern(
        [distancetraversed(rough_center(storm)) <= 5000.0 as pred1,
         area(storm) > 250000.0 as pred2],
        [stconstraint("pred1", "pred2", meanwhile)])

```


10.2.3 Find the trains whose up and down trips meet inside one of the train stations.

```

SELECT  *
FROM    trainsmeet
WHERE   pattern(
        [not(isempty(intersection(uptrip, downtrip))) as pred1,
         uptrip inside stations as pred2 ],
        [stconstraint("pred1","pred2", together)])
ORDERBY line

```

10.2.4 Find the trains that encountered a delay of more than 30 minutes after passing through the snow storm *msnow*.

```

SELECT *
FROM    trainsdelay
WHERE   pattern([not(delay(actual, schedule) > 1800.0) as pred1,

```

Table 3: Lifted Operations

Operation	Signature	Type	Meaning
at	$\underline{mregion} \times \underline{point} \rightarrow \underline{mpoint}$	topological operation	computes a moving point that exist whenever the point argument is inside the moving region argument.
isempty	$\underline{mpoint} \rightarrow \underline{mbool}$	set operation	true whenever the argument is defined.
not	$\underline{mbool} \rightarrow \underline{mbool}$	boolean operation	logical negation.
rough_center	$\underline{mregion} \rightarrow \underline{mpoint}$	aggregation	aggregates the moving region into a moving point that represents its center of gravity.
speed	$\underline{mpoint} \rightarrow \underline{mreal}$	metric property	the metric speed of the moving point.
distancetraversed	$\underline{mpoint} \rightarrow \underline{mreal}$	metric property	the distance that the moving point traversed since the start of its definition time.
area	$\underline{mregion} \rightarrow \underline{mreal}$	metric property	the area of the moving region.
intersection	$\underline{mpoint} \times \underline{mpoint} \rightarrow \underline{mpoint}$	set operation	computes the common parts of the two arguments.
inside	$\underline{mpoint} \times \underline{mregion} \rightarrow \underline{mbool}$ $\underline{mpoint} \times \underline{points} \rightarrow \underline{mbool}$	spatial range predicate	true whenever the <i>mpoint</i> is contained in the <i>mregion</i> , or passes some of the <i>points</i> .
delay	$\underline{mpoint} \times \underline{mpoint} \rightarrow \underline{mreal}$	metric operation	considers the first argument <i>actual</i> , and the second <i>schedule movement</i> and computes the delay of the actual movement in seconds.
=	$\underline{mpoint} \times \underline{point} \rightarrow \underline{mbool}$	spatial range predicate	true whenever the moving point passes the point.
xangle	$\underline{mpoint} \rightarrow \underline{mreal}$	direction	the angle (in degrees) between x-axis and the tangent of the moving point.
and	$\underline{mbool} \times \underline{mbool} \rightarrow \underline{mbool}$	boolean operation	logical and.
<, <=, >, >=	$\underline{mreal} \times \underline{real} \rightarrow \underline{mbool}$	left/right range predicate	true in the time intervals during which the comparison holds.

```

    actual inside msnow as pred2,
    delay(actual, schedule) > 1800.0 as pred3 ],
    [stconstraint("pred1", "pred2", vec("abab", "aba.b", "abba")),
    stconstraint("pred2", "pred3",
        vec("abab", "aba.b", "abba", "aa.bb", "aabb"))])

```

10.2.5 Find the trains that are always heading north-west after passing *mehringdamm*.

```

SELECT *
FROM   trains
WHERE  patternex([trip = mehringdamm as pred1,
    ndefunit(((xangle(trip) >= 90.0) and
        (xangle(trip) <=180.0)), int2bool(1)) as pred2],
    [stconstraint("pred1", "pred2", then)],
    (((start("pred2")- end("pred1")) < create_duration(0, 120000))
    and
    ((inst(final(trip)) - end("pred2")) < create_duration(0, 15000))))

```

where we use the *ndefunit* operator in this query to replace the undefined periods within the *mbool* by *true* units. This is because the *xangle*¹ operator yields undefined during the train stops in the stations. In other words, *pred2* is true whenever the train is not heading other than north-west. The query restricts the results to the trains which started heading north at most 2 minutes after passing *mehringdamm* and remained so till at least 15 seconds before the end of the trip. These time margins are used to cut out small noisy parts in the data, so that the query yields results.

11 System Use and Experimental Repeatability

The implementation of the described approach is made available as a Plugin for the SECONDO system. It can be downloaded from the Plugin web site [1]. The *User Manual* (also available on the Plugin web site) describes how to install and run the Plugin. We have also made available the scripts for running the experiments and the *Berlintest* application example so that the results are repeatable.

Before running the scripts of the experiments, you need to install:

1. The SECONDO system version 2.9.1 or later². A brief installation guide is given in the *Plugin User Manual* on [1], and a detailed guide is given in the *SECONDO User Manual* [3].
2. The Spatiotemporal Pattern Queries Plugin (STPatterns) as described in [1].

11.1 Repeating the First Experiment

During the installation of the STPattern Plugin, two files are copied to the SECONDO bin directory \$SECONDO_BUILD_DIR/bin. These two files *Expr1Script.sec* and *STPQExpr1Query.csv* (described in Appendix A) automate the repeatability of the first experiment in this paper. The experiment can then be run as follows:

1. Run `SecondoTTYNT` (i.e. in a shell, go to \$SECONDO_BUILD_DIR/bin and write `SecondoTTYNT`).
2. Make sure that the *berlintest* database is restored (i.e. at the SECONDO prompt, write `list databases` and make sure that *berlintest* database is in the list). Otherwise, restore it by writing

¹The *xangle* operator is a corrected copy of the SECONDO *mdirection* operator. It is presented only for the sake of this example. In the SECONDO versions newer than 2.9.1, the *mdirection* operator works fine.

²Since our optimizer extension wraps around the standard optimizer implementation, you may get different optimization results in later SECONDO versions. The described results in this paper are obtained from version 2.9.1

```
restore database berlintest from berlintest
```

at the **SECONDO** prompt (press <return> twice).

3. Execute the script by writing `@Expr1Script.sec` at the **SECONDO** prompt. The script creates the required database objects and executes the experiment queries. This may take half an hour depending on your machine.

Executing the script creates a **SECONDO** relation *STPQExpr1Result* in the *berlintest* database, which stores the experimental results. Its schema is shown in Table 4.

Table 4: The schema of the *STPQExpr1Result* relation

Attribute	Meaning	Example
no	A serial number for the query.	0
queryText	The query text.	thousand feed filter [.stpattern[a:passmbool(mb10), b:passmbool(mb30); stconstraint("a", "b", vec("aa.b.b"))]] count
numPreds	The number of the lifted predicates in the STPP.	2
numConstraints	The number of the constraints in the STPP.	1
ElapsedTimeReal	The measured response time, in seconds, for this query.	0.171932
ElapsedTimeCPU	The measured CPU time, in seconds, for this query	0.16

The experimental results are also saved to a comma separated file *STPQExpr1Result.csv* in the **SECONDO** bin directory. The file has a similar structure as the table *STPQExpr1Result*.

11.2 Repeating the Second Experiment

Repeating the second experiments is also automated by script files that are copied to the **SECONDO** directories during the installation of the STPattern Plugin. For the second experiment, two script files are used; the `$SECONDO_BUILD_DIR/bin/Expr2Script.sec` file creates the necessary database objects, and the `$SECONDO_BUILD_DIR/Optimizer/expr2Queries.pl` executes the queries. The *Expr2Script.sec* file is described in Appendix B, and the *expr2Queries.pl* in Appendix C. The experiment is repeated as follows:

1. Run `SecondoTTYNT`.
2. Make sure that the *berlintest* database is restored, otherwise, restore it.
3. Execute the *Expr2Script.sec* by writing `@Expr2Script.sec` at the **SECONDO** prompt. This creates the necessary database objects.
4. Quit `SecondoTTYNT` (i.e. write `quit` at the **SECONDO** prompt), go to the **SECONDO** optimizer folder `$SECONDO_BUILD_DIR/Optimizer` and write `SecondoPL`. This starts the **SECONDO** optimizer user interface in the single user mode.
5. Write `consult(expr2Queries).` to let Prolog interpret the script file *expr2Queries.pl*.
6. Open the *berlintest* database (i.e. write `open database berlintest.`).

7. Write `runSTPQExpr2DisableOptimization.` to run the queries without enabling the optimization of the STPP, or `runSTPQExpr2EnableOptimization.` to run the queries with the optimization of the STPP being enabled. This can take more than an hour.

The results are saved to the comma separated files *Expr2StatsDO.csv* and *Expr2QueriesDO.csv* in the `SECONDO` optimizer folder if the STPP optimization is disabled. If it is enabled, the results are saved to the files *Expr2StatsEO.csv* and *Expr2QueriesEO.csv*.

The files *Expr2StatsDO.csv* and *Expr2StatsEO.csv* show the run times. They include the columns described in Table 5.

Table 5: The schemas of the *Expr2StatsDO.csv* and *Expr2StatsEO.csv* files

Attribute	Meaning	Example
NumberOfPredicates	The number of the lifted predicates in the STPP.	2
NumberOfConstraints	The number of the constraints in the STPP.	1
Serial	A serial for the query in the range [0,9]. The serial is repeated with every experimental setup	1
ExecTime	The measured response time, in milliseconds, for this query.	443

The files *Expr2QueriesDO.csv* and *Expr2QueriesEO.csv* have a similar structure. They exclude the *ExecTime* attribute and have two more attributes; the *SQL* attribute which stores the SQL-like query, and the *ExecutablePlan* which stores the execution plan generated by the Optimizer.

12 Conclusions

We propose a novel approach for spatiotemporal pattern queries. It combines efficiency, expressiveness and a clean concept. It builds on other moving objects database concepts. Therefore, it is convenient in the context of spatiotemporal DBMSs. Unlike the previous approaches, it is integrated with query optimizers. We also propose an algorithm for evaluating the constraint satisfaction problems, that is customized to fit the efficient evaluation of the spatiotemporal pattern predicates. In the paper, we demonstrate two application examples to emphasize the expressive power of our approach. Our work is completely implemented in the `SECONDO` platform. The implementation and the scripts for experimental repeatability are available on the Web. The experimental evaluation shows that the run times are reasonable. As future work, we intend to support spatiotemporal patterns that involve repetitions and alternations.

References

- [1] `SECONDO` plugins.
http://dna.fernuni-hagen.de/secondo.html/start_content_plugins.html.
- [2] `SECONDO` programmer’s guide.
<http://dna.fernuni-hagen.de/secondo.html/files/programmersguide.pdf>.
- [3] `SECONDO` user manual.
<http://dna.fernuni-hagen.de/secondo.html/files/secondomanual.pdf>.
- [4] `SECONDO` web site.
<http://dna.fernuni-hagen.de/secondo.html/>.
- [5] James F. Allen. Maintaining knowledge about temporal intervals. *Commun. ACM*, 26(11):832–843, 1983.

- [6] Christian Bessiere. *Handbook of Constraint Programming*, chapter 3. Elsevier, 2006.
- [7] José Antonio Cotelo Lema, Luca Forlizzi, Ralf Hartmut Güting, Enrico Nardelli, and Markus Schneider. Algorithms for moving objects databases. *Comput. J.*, 46(6):680–712, 2003.
- [8] Somayeh Dodge, Robert Weibel, and Anna-Katharina Lautenschütz. Towards a taxonomy of movement patterns. *Information Visualization*, 7(3):240–252, 2008.
- [9] Martin Erwig. *Spatio-Temporal Databases (ed. Caluwe, De)*, chapter 2, pages 29–54. Springer-Verlag New York, Inc., 2004.
- [10] Martin Erwig and Markus Schneider. Developments in spatio-temporal query languages. In *DEXA '99: Proceedings of the 10th International Workshop on Database & Expert Systems Applications*, page 441, Washington, DC, USA, 1999. IEEE Computer Society.
- [11] Martin Erwig and Markus Schneider. Spatio-temporal predicates. *IEEE Trans. on Knowl. and Data Eng.*, 14(4):881–901, 2002.
- [12] Luca Forlizzi, Ralf Hartmut Güting, Enrico Nardelli, and Markus Schneider. A data model and data structures for moving objects databases. In *SIGMOD '00: Proceedings of the 2000 ACM SIGMOD International Conference on Management of Data*, pages 319–330, New York, NY, USA, 2000. ACM.
- [13] Joachim Gudmundsson, Marc van Kreveld, and Bettina Speckmann. Efficient detection of motion patterns in spatio-temporal data sets. In *GIS '04: Proceedings of the 12th annual ACM International Workshop on Geographic Information Systems*, pages 250–257, New York, NY, USA, 2004. ACM.
- [14] Ralf Hartmut Güting, Victor Almeida, Dirk Ansorge, Thomas Behr, Zhiming Ding, Thomas Höse, Frank Hoffmann, Markus Spiekermann, and Ulrich Telle. SECONDO: An extensible DBMS platform for research prototyping and teaching. In *ICDE '05: Proceedings of the 21st International Conference on Data Engineering*, pages 1115–1116, Washington, DC, USA, 2005. IEEE Computer Society.
- [15] Ralf Hartmut Güting, Thomas Behr, Victor Almeida, Zhiming Ding, Frank Hoffmann, and Markus Spiekermann. SECONDO: An extensible DBMS architecture and prototype. Technical Report Informatik-Report 313, FernUniversität Hagen, March 2004.
- [16] Ralf Hartmut Güting, Michael H. Böhlen, Martin Erwig, Christian S. Jensen, Nikos A. Lorentzos, Markus Schneider, and Michalis Vazirgiannis. A foundation for representing and querying moving objects. *ACM Trans. Database Syst.*, 25(1):1–42, 2000.
- [17] Marios Hadjieleftheriou, George Kollios, Petko Bakalov, and Vassilis J. Tsotras. Complex spatio-temporal pattern queries. In *VLDB '05: Proceedings of the 31st International Conference on Very Large Data Bases*, pages 877–888. VLDB Endowment, 2005.
- [18] Yannis E. Ioannidis. Query optimization. *ACM Comput. Surv.*, 28(1):121–123, 1996.
- [19] Cédric Mouza and Philippe Rigaux. Mobility patterns. *Geoinformatica*, 9(4):297–319, 2005.
- [20] Mirco Nanni, Bart Kuijpers, Christine Körner, Michael May, and Dino Pedreschi. *Mobility, Data Mining and Privacy*, chapter 10. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 2008.
- [21] Dieter Pfoser, Christian S. Jensen, and Yannis Theodoridis. Novel approaches in query processing for moving object trajectories. In *VLDB '00: Proceedings of the 26th International Conference on Very Large Data Bases*, pages 395–406, San Francisco, CA, USA, 2000. Morgan Kaufmann Publishers Inc.

- [22] Markus Schneider. Evaluation of spatio-temporal predicates on moving objects. In *ICDE '05: Proceedings of the 21st International Conference on Data Engineering*, pages 516–517, Washington, DC, USA, 2005. IEEE Computer Society.

A The Expr1Script.sec File

This is a commented version of the *Expr1Script.sec* script.

The script runs the first experiment with minimal user interaction. The experiment, as described in Section 9.1, is intended to evaluate the execution overhead of the STPP. This script first creates the required database objects, then executes the queries and logs the run times.

```
close database;
open database berlintest;

let mb1 = randommbool(now());
:
let mb30 = randommbool(now());
```

The commands open the database *berlintest* and creates 30 random *mbool* objects with the names *mb1... mb30*. These objects are needed for the queries. The *randommbool* operator works as described in Section 9.1.1.

```
let later = vec("aabb", "a.abb", "aab.b", "a.ab.b");
let follows = vec("aa.bb", "a.a.bb", "aa.b.b", "a.a.b.b");
let immediately = vec("a.bab", "a.bba", ...
let meanwhile = vec( ...
let then = vec( ...
```

The five vector temporal connectors are used in the queries as examples for vector temporal connectors. They are used together with the 26 simple temporal connectors to generate the queries.

```
let STPQExpr1Query=
[const rel(tuple([no:int, queryText: text,
numPreds: int, numConstraints: int])) value ()]
csvimport['STPQExpr1Query.csv', 0, "", "$"] consume;
```

The query imports the experiment queries from the comma separated file *STPQExpr1Query.csv* and stores them in a *SECONDO* relation called *STPQExpr1Query*. The *[const . value .]* operator tells the *csvimport* operator the schema of the relation, which is shown in Table 6.

Table 6: The schemas of the *STPQExpr1Query.csv* file and the *STPQExpr1Query* table

Attribute	Meaning
no	A serial for the query in the range [0, 4899].
queryText	The query statement written in <i>SECONDO</i> executable language.
numPreds	The number of the lifted predicates in the STPP.
numConstraints	The number of the constraints in the STPP.

The file contains 4900 queries that were randomly generated as described in Section 9.1.2. The queries represent 49 experimental settings, each of which have 100 queries. The following query executes them and logs the results in the relation *STPQExpr1Result*:

```
let STPQExpr1Result =
STPQExpr1Query feed
loopjoin[fun(queryTuple: TUPLE)
evaluate(attr(queryTuple, queryText))
project[ElapsedTimeReal, ElapsedTimeCPU]]
consume;
```

This query can take half an hour depending on your machine. You can query the results relation in any of the *SECONDO* user interfaces [3] and create aggregations for the charts. Additionally, the following query exports the relation to the comma separated file *STPQExpr1Result.csv* in the *SECONDO* bin directory.

```
query STPQExpr1Result feed
  projectextend[; Serial: .no,
    NumberOfPredicates: .numPreds,
    NumberOfConstraints: .numConstraints,
    ResponseTime: .ElapsedTimeReal,
    CPUTime: .ElapsedTimeCPU]
csvexport['STPQExpr1Result.csv', FALSE, TRUE]
count
```

NOTE: We encourage the reader to get information about the *SECONDO* operators by using the built-in operator descriptions. For example, to get help on the operator *csvimport*, write the following query at the *SECONDO* prompt:

```
query SEC2OPERATORINFO feed
  filter[.Name contains "csvimport"]
consume
```

B The Expr2Script.sec File

This is a commented version for the *Expr2Script.sec* script.

The script is used to generate the data required for running the second experiment in this paper without executing the queries. The queries need to be executed in the *SecondoPL* environment afterwards.

```
close database;
open database berlintest;

let RestaurantsNumbered =
  Restaurants feed addcounter[no, 1] head[300] consume;
let point1 =
  RestaurantsNumbered feed filter[.no = 1] extract[geoData];
:
let point300 =
  RestaurantsNumbered feed filter[.no = 300] extract[geoData];
delete RestaurantsNumbered;
```

First, the commands open the database *berlintest*. The geometries of the first 300 restaurants in the *Restaurants* table are then copied to point objects (point1... point300) to be used in the queries.

```
let later = vec("aabb", "a.abb", "aab.b", "a.ab.b");
let follows = vec("aa.bb", "a.a.bb", "aa.b.b", "a.a.b.b");
let immediately = vec("a.bab", "a.bba", ...
let meanwhile = vec( ...
let then = vec( ...
```

The five vector temporal connectors, that are also created in *Expr1Script.sec*, are included here so that the two experiments can be run independently.

```
let Trains20 = thousand feed head[20] Trains feed product consume;
```

This query creates the *Trains20* relation by replicating the tuples of the *Trains* relation 20 times. In the following query, we create an index on the *Trains20* relation to test the proposed STPP optimization.

The index is a spatial R-tree on the units of the *Trip* attribute. Instead of indexing the complete movement, the index is built on the units (i.e. a bounding box is computed for every unit in the Trip). This is done so that the bounding boxes better approximate the moving point.

```
let Trains20_Trip_sptuni =
  Trains20 feed
  projectextend[Trip; TID: tupleid(.)]
  projectextendstream[TID; MBR: units(.Trip)
    use[fun(U: upoint) bbox2d(U) ]]
  sortby[MBR asc]
  bulkloadrtree[MBR];
```

C The expr2Queries.pl File

This Prolog file is used to run the queries of the second experiment and log the execution times. It defines four prolog predicates:

1. runSTPQExpr2DisableOptimization/0: switches off STPP optimization by setting the optimizer options, and executes the queries.
2. runSTPQExpr2EnableOptimization/0: switches on STPP optimization, and executes the queries.
3. executeSQL/4: helper predicate for executing queries.
4. runSTPQExpr2/4: the facts table that stores the queries. The file contains 490 such facts, 10 queries for each of the 49 experimental settings. The queries are randomly generated as described in Section 9.2.2. For every query, the fact also stores its serial, number of lifted predicates, and number of constraints.

D Running the Berlittest Application Example

To execute the queries in the berlittest example, you need first to run the script *BerlittestScript.sec* from the SecondoTTYNT prompt. The script is installed within the STPattern Plugin. You also need to have the berlittest database restored in your system. The script file creates the required database objects but it doesn't execute the queries. It first defines some temporal connectors:

```
close database;
open database berlittest;
let later= vec("aabb", "a.abb", "aab.b", "a.ab.b");
let follows= vec(...
let immediately= vec(...
let meanwhile= vec(...
let then= vec(...
let together= vec(...
```

Then it restores the *SnowStorms* relation from the *SnowStorms* file in the *SECONDO/bin* directory, which is installed with the Plugin.

```
restore SnowStorms from SnowStorms;
```

The following command creates the relation *TrainsMeet*, that is used in the example in Section 10.2.3. Every tuple in the relation is a different combination of an up train, down train of the same line, and the stations where the train line stops.

```

let TrainsMeet =
  Trains feedproject[Line, Trip, Up] {t2} filter[.Up_t2 = FALSE]
  Trains feedproject[Line, Trip, Up] {t1} filter[.Up_t1 = TRUE]
  hashjoin[Line_t2 , Line_t1 , 99997]
  extend[Line: .Line_t1, Uptrip: .Trip_t1, Downtrip: .Trip_t2,
    Stations: ((breakpoints(.Trip_t1, create_duration(0,5000) )
      union val(initial(.Trip_t1)))
      union val(final(.Trip_t1)))])
  project[Line, Uptrip, Downtrip, Stations]
  consume;

```

Next we create the relation *TrainsDelay*, used in the example in Section 10.2.4. Every tuple has a *schedule* and an *actual* moving point. The *schedule* movement is a copy from the *Trip* attribute in the *Trains* relation. The actual movement should have delays of about half an hour. We shift the *Trip* 1795 seconds forward, and apply a random positive or negative delay up to 10 seconds to the result. This creates actual movements with random delays between 29:45 and 30:05 minutes.

```

let TrainsDelay=
  Trains feed
  extend[Schedule: .Trip,
    Actual: randomdelay(
      .Trip translate[create_duration(0, 1795000) , 0.0, 0.0],
      create_duration(0,10000) ) ]
  project[Id, Line, Actual, Schedule]
  consume;

```

After running the *BerlintestScript.sec* script, use the *Javagui* to execute the queries. It is the graphical user interface for SECONDO. To launch it:

1. Start the SECONDO kernel in server mode, the optimizer server, and the GUI:
 In a new shell, go to `$SECONDO_BUILD_DIR/bin`, and type `SecondoMonitor -s`.
 In a new shell, go to `$SECONDO_BUILD_DIR/Optimizer`, and type `StartOptServer`.
 In a new shell, go to `$SECONDO_BUILD_DIR/Javagui`, and type `sgui`. The Javagui will start and connect to both the kernel and the optimization server.
2. Open the database. In the Javagui type:
`open database berlittest.`
3. Set the optimizer options. The SECONDO optimizer maintains a list of options that controls the optimization. The examples in this paper require the options *improvedcosts*, *determinePredSig*, *autoSamples*, *rewriteInference*, *rtreeIndexRules*, and *autosave*. To set each of these options, type in the Javagui:
`optimizer setOption(option)`
4. View the underlying network. Type:
`select * from ubahn` to display the underground trains network.
`select * from trains` to display the moving trains. Use the slider to view the results.
 Select the last query in the top-right panel and press `hide` to hide the trains.
`select * from snowstorms` to display the moving snow storms.
`hide the snow storms.`
5. Type the example queries as in Section 10.2, and make sure to type everything in lower case.

Verzeichnis der zuletzt erschienenen Informatik-Berichte

- [339] Beierle, Chr., Kern-Isberner, G. (Eds.): Dynamics of Knowledge and Belief -
Workshop at the 30th Annual German Conference on Artificial Intelligence, KI-
2007
- [340] Düntgen, Chr., Behr, Th., Güting R. H.: BerlinMOD: A Benchmark for Moving Object
Databases
- [341] Saatz, I.: Unterstützung des didaktisch-methodischen Designs durch einen
Softwareassistenten im e-Learning
- [342] Hönig, C. U.:
Optimales Task-Graph-Scheduling für homogene und heterogene Zielsysteme
- [343] Güting, R. H.:
Operator-Based Query Progress Estimation
- [344] Behr, Th., Güting R. H.:
User Defined Topological Predicates in Database Systems
- [345] vor der Brück, T.; Helbig, H.; Leveling, J.:
The Readability Checker Delite Technical Report
- [346] vor der Brück:
Application of Machine Learning Algorithms for Automatic Knowledge Acquisition
and Readability Analysis Technical Report
- [347] Fechner, B.:
Dynamische Fehlererkennungs- und –behebungsmechanismen für verlässliche
Mikroprozessoren
- [348] Brattka, V., Dillhage, R., Grubba, T., Klutsch, Angela.:
CCA 2008 - Fifth International Conference on Computability and Complexity in
Analysis
- [349] Osterloh, A.:
A Lower Bound for Oblivious Dimensional Routing
- [350] Osterloh, A., Keller, J.:
Das GCA-Modell im Vergleich zum PRAM-Modell
- [351] Fechner, B.:
GPUs for Dependability
- [352] Güting, R. H., Behr, T., Xu, J.:
Efficient k -Nearest Neighbor Search on Moving Object Trajectories
- [353] Bauer, A., Dillhage, R., Hertling, P., Ko K.I., Rettinger, R.:
CCA 2009 Sixth International Conference on Computability and Complexity in
Analysis
- [354] Beierle, C., Kern-Isberner G.
Relational Approaches to Knowledge Representation and Learning