

Indexing the Trajectories of Moving Objects

Dieter Pfoser
Computer Technology Institute
Aktaiou 11 & Pouloupoulou Str.
11851 Athens, Hellas
pfoser@cti.gr

Abstract

The domain of spatiotemporal applications is a treasure trove of new types of data and queries. In this work, the focus is on a spatiotemporal sub-domain, namely the trajectories of moving point objects. We examine the issues posed by this type of data with respect to indexing and point out existing approaches and research directions. An important aspect of movement is the scenario in which it occurs. Three different scenarios, namely unconstrained movement, constrained movement, and movement in networks are used to categorize various indexing approaches. Each of these scenarios give us different means to either simplify indexing, or to improve the overall query processing performance.

1 Introduction

Several application areas contribute to the growing number of different types of spatiotemporal data. For example, we are currently experiencing rapid technological developments that promise widespread use of on-line mobile personal information appliances [14]. Mobility as a keyword is introduced as a concern to many applications and services. One aspect of mobility is movement and thus the change of location. Applications in this context include emerging ones such as location-based services, but also “classical” ones such as fleet management, dealing with the optimal spatiotemporal distribution of vehicles [1].

Applications such as these warrant the study of indexing mobile objects. In particular, our interest is in *recording the movements of mobile objects, i.e., their trajectories, and indexing those for post-processing (e.g., data mining) purposes*. Thus, we will not concern ourselves with the indexing of the current positions and the predicted movements of objects such as treated in, e.g., [16, 17]. The size and shape of an object is often fixed and of little importance—only its position matters. Thus, the problem becomes one of recording the position of a moving object across time. The movement of an object may then be represented by a trajectory in the three dimensional space composed of two spatial dimension and one time dimension [11].

Typically, access methods are developed only having the data and the queries in mind. However, for trajectory data we can also consider the constraints which the objects in their movement are subjected to. Specifically, we may distinguish among three movement scenarios, namely unconstrained movement (e.g., vessels at sea), constrained movement (e.g., pedestrians), and movement in transportation networks (e.g., trains and, typically,

Copyright 0000 IEEE. Personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution to servers or lists, or to reuse any copyrighted component of this work in other works must be obtained from the IEEE.

Bulletin of the IEEE Computer Society Technical Committee on Data Engineering

cars). As we will see in this work, the latter two scenarios allow us to optimize the query processing technique or to use simpler access methods to index the data.

The outline of the paper is as follows. Section 2 defines the basic concepts and points out the difficulties when indexing trajectories. Sections 3, 4, and 5 point out indexing approaches grouped according to the three movement scenarios. Finally, Section 6 gives conclusions and directions for future research.

2 On Trajectories and Queries

The central question in this work is how we can ease the task of indexing point-object movements. This section gives a brief description of the data, the related queries, and the respective indexing challenges.

2.1 Trajectories

A trajectory is the data we obtain by recording the position of a moving point object across time. Consider the following application context. Optimizing transportation, especially in highly populated and thus congested areas, is a very challenging task that may be supported by an information system. A core application in this context is fleet management [1]. Vehicles equipped with GPS receivers transmit their positions to a central computer using either radio communication links or mobile phones. At the central site, the data is processed and utilized.

To record the movement of an object, we would have to know its position at all times, i.e., on a continuous basis. However, GPS and telecommunications technologies only allow us to sample an object's position, i.e., to obtain the position at discrete instances of time, such as every few seconds. A first approach to represent the movements of objects would be to simply store the position samples. This would mean that we could not answer queries about the objects' movements at times in-between those of the sampled positions. Rather, to obtain the entire movement, we have to interpolate. The simplest approach is to use linear interpolation, as opposed to other methods such as polynomial splines. The sampled positions then become the endpoints of line segments of polylines, and the movement of an object is represented by an entire polyline in 3D space. The solid line in Figure 1(a) represents the movement of an object. Space and time coordinates are combined to form a single coordinate system. The dashed line shows the projection of the movement into the 2D (spatial) plane [10, 11].

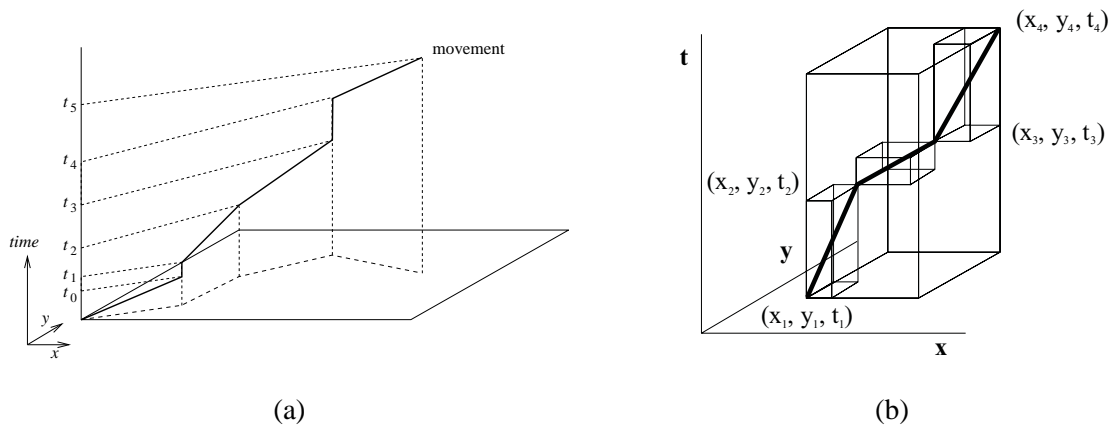


Figure 1: Trajectories: (a) 2+1 dimensional space, (b) approximating trajectories using MBBs

In classical spatial databases only position information is available. In our case, however, we have also derived information, e.g., speed, acceleration, traveled distance, etc. Information is derived from the combination of spatial and temporal data. Further, we do not just index collections of line segments-these are parts of larger,

Query Type		Operation	Signature
Coordinate-based Queries		overlap, inside, etc.	range
Trajectory -based Queries	Topological Queries	enter, leave, cross, bypass	$\text{range} \times \{\text{segments}\}$ $\rightarrow \{\text{segments}\}$
	Navigational Queries	traveled distance, covered area (top or average), speed, heading, parked	$\{\text{segments}\} \rightarrow \text{int}$ $\{\text{segments}\} \rightarrow \text{real}$ $\{\text{segments}\} \rightarrow \text{bool}$

Table 1: Types of spatiotemporal queries

semantically meaningful objects, namely trajectories. These semantic properties of the data are reflected in the types of queries of interest for the data, as will be discussed next.

2.2 Queries

The queries of interest for trajectories comprise what can be conceived as adapted spatial queries, e.g., range queries of the form “Find all objects within a given area at some time during a given time interval” and queries with no spatial counterpart. Here, the so-called trajectory-based queries are classified in *topological* queries, which involve the entire movement of an object (enter, leave, cross, and bypass), and *navigational* queries, which involve derived information, such as speed and heading. Table 1 summarizes the query types. We use the signature-like notation as presented in [2]. The “operation” column lists the operations used for several query types, and the “signature” column presents the involved types. For example, a coordinate-based query uses the inside operation to determine the segments within the specified range. The notation segments simply refers to a set, it does not capture whether this set constitutes one or more trajectories. An important issue in dealing with spatiotemporal queries is to extract information related to (partial) trajectories, e.g., “What were the trajectories of objects after they left Tucson between 7 a.m. and 8 a.m. today, in the next hour?” This type of query is referred to as *combined* query, since we first have to select the respective trajectories (“...left Tucson between 7 a.m. and 8 a.m. today...”) and subsequently the respective portions (“...in the next hour?”). More details on trajectory-related queries can be found in [10, 13].

2.3 Indexing Fundamentals

Trajectories are three-dimensional spatial entities, and we can use spatial access methods to index this data. However, there are difficulties. Trajectories are decomposed into their constituent line segments, which are then indexed. The use of the R-tree [3] is illustrated in Figure 1(b). The R-tree approximates the data objects by Minimum Bounding Boxes (MBBs), here three-dimensional intervals. Approximating segments using MBBs proves to be inefficient. Figure 1(b) shows that we introduce large amounts of “dead space.” It is evident that the MBBs cover large parts of space, whereas the actual space occupied by the trajectory is small. This leads to high overlap and consequently to a small discrimination capability of the index structure. (Note that approximating trajectories with MBBs without prior decomposition into segments just serves to make things worse.)

Other trajectory indexing problems include trajectory preservation and skewed data growth. As for the first problem, spatial indices tend to group segments into nodes according to spatial proximity. However, in the case of trajectories, it is beneficial to some queries if the index preserves trajectories, i.e., to group segments according to their trajectory and only then according to proximity (cf. [13]). The second problem refers to the fact that trajectory data grows mostly in the temporal dimension. The spatial dimensions are fixed, e.g., the city limits. Exploiting this property of the data can further increase query performance.

2.4 Movement Scenarios

Depending on the particular objects and applications under consideration, the movements of the objects may be subject to constraints. Specifically, we may distinguish among three movement scenarios, namely unconstrained movement (e.g., vessels at sea), constrained movement (e.g., pedestrians), and movement in transportation networks (e.g., trains and, typically, cars). Figures 2(a)-(c) exemplifies these three scenarios. In each case, within the bounding square white illustrates space in which movement can occur, whereas black signifies the opposite. Unconstrained movement is the scenario mostly asserted in work on spatiotemporal access methods. However, it hardly represents reality. Constrained movement and movement in networks represent similar movement scenarios. The former assumes that there exist spatial objects that constrain the movement, e.g., when considering the movement of cars, houses, lakes, parks, etc. (examples of the black boxes in Figure 2(b)) are examples of such constraints. Movement in networks is then merely an abstraction of constrained movement. Here, one is only interested in positions of objects with respect to the network and not with respect to a two-dimensional reference system. For example, we may expect that many applications will be interested only in the positions of cars with respect to the road network, rather than in their absolute coordinates. The movement effectively occurs in a different space than for the first two scenarios.

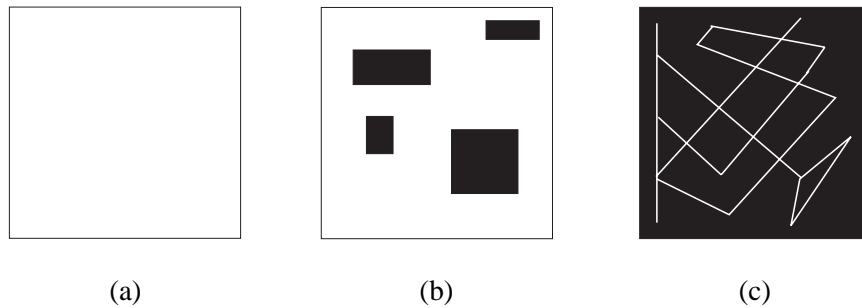


Figure 2: Movement scenarios: (a) unconstrained, (b) constrained, and (c) movement in a network

Section 2.3 illustrates the indexing challenge related to trajectories. In exploiting auxiliary information such as infrastructure and networks, we can improve query performance and simplify indexing. The following three sections explore indexing approaches related to three movement scenarios presented in this section.

3 On Indexing Unconstrained Movement

Unconstrained movement is conceptually the simplest case for trajectory indexing. In the following, we describe several access methods that are tailored to the requirements of the data and the queries.

3.1 The TB-tree

The TB-tree [13] is an access method that considers the particularities of the data, namely trajectory preservation and temporal growth, and aims at efficiently processing related types of queries.

An underlying assumption for the R-tree is that all inserted geometries are independent. In our context this translates to all line segments being independent. However, line segments are part of trajectories and the R-tree only implicitly maintains this knowledge. With the TB-tree, we aim for an access method that strictly preserves trajectories. As a consequence, each leaf node in the index should only contain segments belonging to one trajectory. Thus, the index becomes actually a *trajectory bundle*. This approach is only feasible by making some concessions to the most important R-tree property, namely node overlap, or spatial discrimination. As a drawback, spatially close line segments from different trajectories will be stored in different nodes. This,

in turn, increases the node overlap, decreases the space discrimination, and, thus, increases the classical range query cost. The TB-tree trades space discrimination for trajectory preservation.

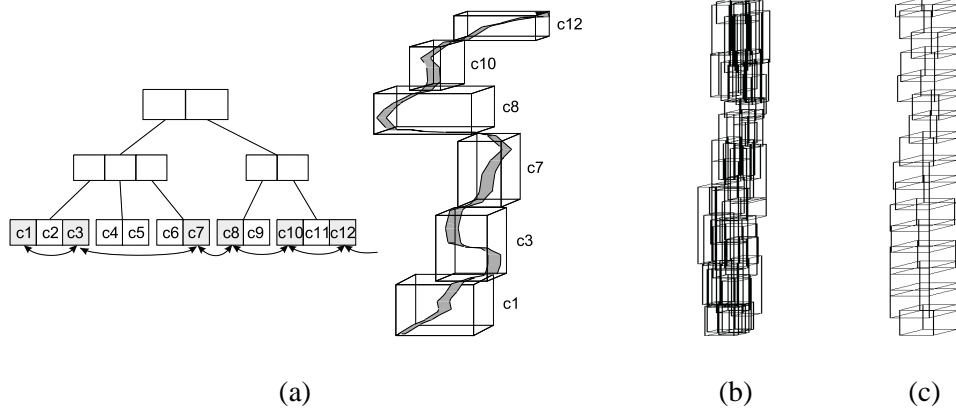


Figure 3: TB-tree: (a) structure, non-leaf level MBB structure (b) R-tree and (c) TB-tree

Employing the above constraints, the TB-tree structure consists of a set of leaf nodes, each containing a partial trajectory, organized in a tree hierarchy. For query processing, it is necessary to be able to retrieve segments based on their trajectory identifier. Thus, leaf nodes containing segments of the same trajectory are linked together by a doubly-linked list data structure. This preserves trajectory evolution and overall improves especially the performance of processing trajectory-based and combined queries. Figure 3(a) shows a part of a TB-tree structure and a trajectory illustrating the overall data structure. For clarity, the trajectory is drawn as a band rather than a line. The trajectory symbolized by the grey band is fragmented across six nodes, c1, c3, etc. In the TB-tree these leaf nodes are connected through the doubly-linked list.

Figures 3(b) and (c) give an impression of the quality of the non-leaf level arrangement of bounding boxes in the R- and the TB-tree, respectively. The R-tree grouping is dominated by purely spatial characteristics such as proximity, ignoring the temporal growth of the data. The TB-tree structure is dominated by trajectory preservation. Since the data is inserted with a growing time horizon, the MBBs exhibit a temporal clustering.

The TB-tree structure is not the only approach to the indexing of trajectories. The following section gives a brief overview of what other techniques exist.

3.2 Other Approaches

Nasciemento et al. [8] adopt the trajectory model as described in Section 2.1 and investigate the suitability of multidimensional access methods for trajectory data. In their work, they compare the performance of various R-tree versions for different range query loads.

In [7], Hadjileftheriou et al. define an approach to reduce the dead space introduced by MBB approximations of trajectories (cf. Section 2.3) by introducing “artificial object updates.” They effectively manipulate the partitioning of a trajectory into segments. A partially persistent tree structure [5, 18] is used to index the data. This approach generalizes a previous work [4] in which it was assumed that the objects move with a linear function of time, whereas in [7] more complex functions are permitted.

Porkaew et al. [15, 6] examine the indexing of trajectories in Native Space (cf. Section 2.1) vs. Parametric Space. In Parametric Space, segments of trajectories are represented in terms of a location and a motion vector. In their experimental evaluation, the authors use an R-tree structure as an index for both representations.

Other works further propose access methods that go beyond trajectories towards the indexing of moving spatial shapes in general. These approaches however do not always consider continuous but only discrete changes, e.g., Tzouramanis et al. [19] employ overlapping linear quad-trees to index discretely changing spatial objects.

Tao and Papadias [18] propose a method called MV3R-tree to index the past locations of moving shapes. Their method combines multi-version B-trees and 3D R-trees to process timestamp and interval queries.

4 On Constrained Movement

The indexing approaches in the previous section focus on trajectory data and related queries. In this section, we assert that the movement of objects is constrained by elements termed *infrastructure*. In terms of data, infrastructure represents “black-out” areas for movement and, thus, there are no trajectories where there are infrastructure elements. However, in the index, approximations of the data are used, which introduce dead space and, consequently, the areas covered by infrastructure are not empty. This will incur unnecessary search in the index as well as produce a certain number of falsely reported answers, which must subsequently be eliminated. Both lead to extra I/O operations and thus negatively affect the performance.

To eliminate this extra I/O, we can use infrastructure in a pre-processing step, i.e., why should we look for objects, where there cannot be any? The strategy we choose is to query the infrastructure to save on querying the trajectory data. Overall, this will turn out to be favorable, since the number of infrastructure elements can be assumed to be very small compared to the trajectory data. Further, the trajectory data is growing with time, whereas the size of the infrastructure data remains more or less constant.

The general principle is to decompose a given query window based on the infrastructure contained in it. The intuition is to segment the parts of the query window not occupied by infrastructure into well-shaped rectangles, i.e., as square as possible. In Figure 4 where few but large infrastructure elements are shown as black rectangles, the possible outcome of such a segmentation process is shown as white rectangles.

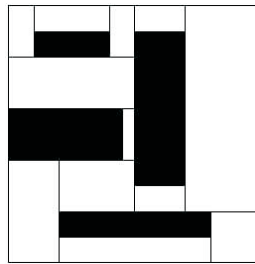


Figure 4: A query window segmentation example

The query windows resulting from this segmentation are subsequently used to query the trajectory data index (instead of the large query window ranging of infrastructure). In [12] the conditions under which this approach is favorable are established.

5 On Movement in Networks

In many applications movement is modeled to occur in a network rather than its native two-dimensional space. The reason is simplicity. When dealing with network-constrained movement, one is not interested in spatial extents, e.g., the thickness of the road, or the absolute position of the object in terms of its x- and y-coordinates, but rather in relative positions with respect to the network, e.g., Kilometer 21 of Highway 101.

The space as defined by a network is quite different from the Euclidean space that the network is embedded into. Intuitively, the dimensionality of the networked-constrained space is lower than that of the space it is embedded in. In the literature, the term 1.5 dimensional has been used.

Modeling movement with respect to a network simplifies the trajectory data obtained. The two spatial dimensions are essentially reduced to one. Figures 5(a) and (b) illustrate this principle by showing the same

trajectory in a three-dimensional and a two-dimensional space, respectively. A two-dimensional network is reduced to a set of one-dimensional segments and the trajectory data is mapped accordingly.

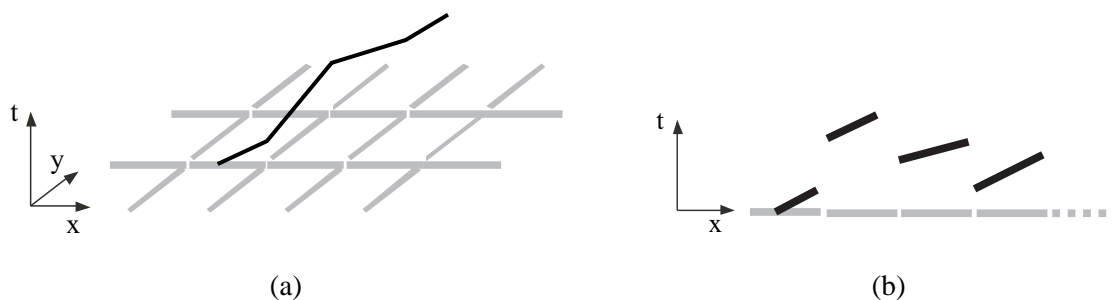


Figure 5: Network movement: trajectories in (a) 3D and (b) 2D space

Lowering the dimensionality of the data reduces the indexing challenge considerably. Off-the-shelf database management systems typically do not offer three-dimensional indexes and thus do not contend with trajectories. Although it is desirable to design new access methods for new types of data, it may not be attractive in a short or medium term. For example, it took a dozen years before the R-tree found its way into some commercial database products [9]. Depending on the type of data, it can be beneficial to use existing access methods by transforming the data. Considering movement in a network is such a transformation. We can use a simple access method such as a two-dimensional R-tree and this, in turn, allows for an easy integration of the new type of data into commercial database management systems.

6 Conclusions and Future Work

Spatiotemporal data is emerging from a broad range of applications. In this work, we present selected methods for the indexing of trajectories, a type of data that stems from recording the movement of point objects. The existing indexing approaches are grouped according to *three movement scenarios*, constrained movement, unconstrained movement, and movement in networks. Each of these scenarios can aid the processing of spatiotemporal queries in different ways. Unconstrained movement is the typical showcase for the definition of new access methods. Constrained movement allows us to reduce the extent of query windows. Movement in networks reduces the dimensionality of the data and thus of the index.

Directions for future work can either be to define more efficient and/or more specialized access methods, or to satisfy existing needs arising from real applications [1]. This can be achieved by trying to handle spatiotemporal data with our current knowledge in connection with the means available from off-the-shelf database management systems [20].

Acknowledgements

The author would like to thank his co-authors on this subject, Christian S. Jensen and Yannis Theodoridis.

Part of this research was supported by the CHOROCHRONOS project, funded by the European Commission DG XII, contract no. ERBFMRX-CT96-0056, and the Nykredit Corporation.

References

- [1] V. Delis, D. Pfoser, Y. Theodoridis, and N. Tryfona. Movement mining in fleet management applications. Project Proposal, General Secretariat of Science and Technology, Athens, Hellas, 2001.

- [2] R. Güting, M. Böhlen, M. Erwig, C. S. Jensen, N. Lorentzos, M. Schneider, and M. Vazirgiannis. A foundation for representing and querying moving objects. *ACM Transactions on Database Systems*, 25(1):1–42, 2000.
- [3] A. Guttman. R-trees: a dynamic index structure for spatial searching. In *Proceedings of the ACM SIGMOD Conference on Management of Data*, pages 47–57, 1984.
- [4] G. Kollios, D. Gunopulos, V. Tsotras, A. Delis, and M. Hadjieleftheriou. Indexing animated objects using spatio-temporal access methods. *IEEE Transactions on Knowledge and Data Engineering*, pages 742–777, September 2001.
- [5] A. Kumar, V. Tsotras, and C. Faloutsos. Designing access methods for bitemporal databases. *IEEE Transactions on Knowledge and Data Engineering*, 10(1):1–20, 1998.
- [6] I. Lazaridis, K. Porkaew, and S. Mehrotra. Dynamic queries over mobile objects. In *Proceedings of the 8th International Conference on Extending Database Technology*, pages 269–286, 2002.
- [7] Hadjieleftheriou M, G. Kollios, V. J. Tsotras, and D. Gunopulos. Efficient indexing of spatiotemporal objects. In *Proceedings of the 8th International Conference on Extending Database Technology*, pages 251–268, 2002.
- [8] M. Nascimento, J. R. O. Silva, and Y. Theodoridis. Evaluation of access structures for discretely moving points. In *Proceedings of the International Workshop on Spatio-Temporal Database Management*, pages 171–188, 1999.
- [9] Oracle Corporation. Oracle spatial cartridge developers guide and reference, release 9.0.1., Web site: <http://technet.oracle.com>, 2001.
- [10] D. Pfoser. *Issues in the Management of Moving Point Objects*. PhD thesis, Department of Computer Science, Aalborg University, Denmark, 2000.
- [11] D. Pfoser and C. S. Jensen. Capturing the uncertainty of moving-object representations. In *Proceedings of the 6th International Symposium on Spatial Databases*, pages 111–132, 1999.
- [12] D. Pfoser and C. S. Jensen. Querying the trajectories of on-line mobile objects. In *Proceedings of the 2nd ACM International Workshop on Data Engineering for Wireless and Mobile Access*, pages 66–73, 2001.
- [13] D. Pfoser, C. S. Jensen, and Y. Theodoridis. Novel approaches to the indexing of moving object trajectories. In *Proceedings of the 26th International Conference on Very Large Databases*, pages 395–406, 2000.
- [14] E. Pitoura. DB-Globe: an IST-FET research project on global computing. Web site: <http://softsys.cs.uoi.gr/dbglobe/>, 2002.
- [15] K. Porkaew, I. Lazaridis, and S. Mehrotra. Querying mobile objects in spatio-temporal databases. In *Proceedings of the 7th International Symposium on Advances in Spatial and Temporal Databases*, pages 55–78, 2001.
- [16] S. Saltenis and C. S. Jensen. Indexing of moving objects for location-based services. In *Proceedings of the 18th International Conference on Data Engineering*, to appear, 2002.
- [17] S. Saltenis, C. S. Jensen, S. Leutenegger, and M. A. Lopez. Indexing the positions of continuously moving objects. In *Proceedings of the ACM SIGMOD Conference on Management of Data*, pages 331–342, 2000.
- [18] Y. Tao and D. Papadias. Mv3R-tree: a spatiotemporal access method for timestamp and interval queries. In *Proceedings of the 27th International Conference on Very Large Databases*, pages 431–440, 2001.
- [19] T. Tzouramanis, M. Vassilakopoulos, and Y. Manolopoulos. Overlapping linear quadtrees and spatio-temporal query processing. *The Computer Journal*, 43(4):325–343, 2000.
- [20] M. Vazirgiannis and O. Wolfson. A spatiotemporal model and language for moving objects on road networks. In *Proceedings of the 7th International Symposium on Advances in Spatial and Temporal Databases*, pages 20–35, 2001.