

GraCT: A Grammar-based Compressed Index for Trajectory Data [☆]

Nieves R. Brisaboa^a, Adrián Gómez-Brandón^{a,*}, Gonzalo Navarro^b, José R. Paramá^a

^a*Universidade da Coruña, Facultade de Informática, CITIC, Campus de Elviña s/n, 15071 A
Coruña, Spain*

^b*Millennium Institute for Foundational Research on Data (IMFD), Department of Computer Science,
University of Chile, Beauchef 851, Santiago, Chile*

Abstract

We introduce a compressed data structure for the storage of free trajectories of moving objects (such as ships and planes) that efficiently supports various spatio-temporal queries. Our structure, dubbed GraCT, stores the absolute positions of all the objects at regular time intervals (snapshots) using a k^2 -tree, which is a space- and time-efficient version of a region quadtree. Positions between snapshots are represented as logs of relative movements and compressed using Re-Pair, a grammar-based compressor. The nonterminals of this grammar are enhanced with MBR information to enable fast queries.

The GraCT structure of a dataset occupies less than the raw data compressed with a powerful traditional compressor such as *p7zip*. Further, instead of requiring full decompression to access the data like a traditional compressor, GraCT supports direct access to object trajectories or to their position at specific time instants, as well as spatial range and nearest-neighbor queries on time instants and/or time intervals.

Compared to traditional methods for storing and indexing spatio-temporal data, GraCT requires two orders of magnitude less space, and is competitive in query times. In particular, thanks to its compressed representation, the GraCT structure may reside in main memory in situations where any classical uncompressed index must resort to disk, thereby being one or two orders of magnitude faster.

Keywords: Compact data structures, moving objects databases.

1. Introduction

More than two decades after it emerged, the field of moving object databases is still an active area of research. The data collected from the GPS positions of large sets of cars, ships, planes, smartphones and wearable devices, has led to a growing interest in applications that exploit trajectory information, for example with data mining purposes.

[☆]A shorter, preliminary version of this paper appeared in *Proc. SPIRE 2016* [7].

*Corresponding author. Tel. +34981167000 Fax. +34981167160.

Email addresses: brisaboa@udc.es (Nieves R. Brisaboa), adrian.gbrandon@udc.es (Adrián Gómez-Brandón), gnavarro@dcc.uchile.cl (Gonzalo Navarro), jose.parama@udc.es (José R. Paramá)

Preprint submitted to Elsevier

December 28, 2018

As these datasets grow in size and the applications managing moving objects become more sophisticated, more space- and time-efficient storage techniques are required.

A trajectory is the path followed by a moving object through space as a function of time. Due to storage requirements and the limitations of the devices used to acquire the object positions, the continuous movement of an object is usually approximated with discrete samples of spatio-temporal location points. The more samples taken, the more accurate the trajectory. High sampling rates, however, result in large amounts of data, which increase storage, transmission and processing needs. Even when storage, network and processing capacity grows rapidly, the collected data grows even faster, and thus it is necessary to aim for reduced trajectory representations [54].

Traditional methods for compressing trajectories include line generalization (or simplification) techniques, which keep only some of the trajectory points and discard the rest. Obviously, this approach loses some information on the real trajectory. A lossless strategy to obtain compression is the use of *delta compression*, where each new position is stored as the difference with the previous one. This idea exploits the fact that consecutive positions are expected to be closer to each other, and that smaller numbers can be stored using fewer bits. Extracting a whole trajectory with this arrangement is easy. Efficiently accessing the position of an object at a given time, instead, requires sampling some absolute positions at regular time intervals, which introduces a space/time tradeoff.

Another issue is how to index the trajectory data in order to answer queries other than just retrieving a whole trajectory or finding the position of an object at a given time instant. A number of indexes have been proposed since the 90's to handle a rich set of queries on trajectory data. Most indexes were modifications of the R-tree [14] augmented with another dimension to deal with the time. None of those works, however, aim to compress the data. Rather, they are designed to work on disk, which is much slower than the main memory.

With the increasing gap in the access time of main memory versus disk, compressing the trajectories in order to query them in main memory has become an attractive option. Some recent proposals following this trend [10, 48] build on delta compression, coupled with an encoding that favors small numbers. The optimal codes for delta compression can be obtained with a statistical encoder that exploits frequency bias (typically, smaller numbers are more frequent).

In this article we introduce a new way to compress trajectories and demonstrate that it is much more effective than statistical compression on real-world data. Instead of exploiting the higher frequency of smaller differences, we exploit the fact that, in many applications, trajectories tend to be similar to others, wholly or piecewise. We resort to *grammar compression* [18], which is a method from the family of dictionary-based compression methods. Grammar compression generates a context-free grammar that generates (only) a given input sequence. Our input sequence is the concatenation of all the trajectories in differential form. Grammar compression exploits the similarities between trajectories: if many trajectories are similar, then a small grammar can be found that generates the whole dataset.

Our index, called *Grammar-based Compressed Representation of Trajectories (GraCT)*, is an in-memory representation that combines grammar compression of trajectories with additional data structures to support efficient queries of various kinds on the dataset, without the need of decompressing it. On one hand, *snapshots* of the positions of all the objects at regular time intervals are created, and those positions are indexed in a

compressed quadtree representation called k^2 -tree [8]. Such spatial index allows us restricting the set of candidate objects when answering queries. The movements of the objects between two consecutive snapshots are represented in a sequence called the *log of movements*, which is represented in differential form and compressed with Re-Pair, an effective grammar compression method [20]. To avoid decompressing and sequentially traversing the logs to answer queries, the nonterminals of the grammar are enhanced with additional information, most notably the minimum bounding rectangle (MBR) of all the (relative) positions represented by the nonterminal. MBRs allow us skipping whole nonterminal symbols without the need to decompress them, thereby speeding up the traversal of logs. This arrangement enables GraCT to answer *time-slice*, *time-interval* and *k-nearest neighbor* queries efficiently.

GraCT requires discretizing the data in order to obtain compression. When the data must be used with full precision, GraCT can be used as an *in-memory cached index* using a sufficient degree of discretization. The candidate set of answers delivered by GraCT are then read from disk and verified.

We evaluate our techniques on three real-life datasets: ships on the sea, commercial aircrafts, and taxis in a city. We show that the grammar compression of GraCT is highly effective, reducing the raw data to 4%–7% of its original size. This is even less than the space achieved by *p7zip*, a powerful dictionary-based compressor that does not support direct access nor querying. It is also less than half the space used by a baseline index we implement that replaces grammar compression by statistical symbolwise compression. Compared to a traditional method for storing and indexing spatio-temporal data (concretely, the MVR-tree [43]), the compressed representation of GraCT occupies two orders of magnitude less space. In relation to query times, GraCT is up to 4 times faster than the statistically compressed index, except for retrieving whole trajectories, where it is slightly slower. GraCT is also faster than the MVR-tree running in main memory (which, as explained, is feasible for small datasets only) for all queries except on nearest-neighbor queries, on which it can be clearly slower. GraCT is one or two orders of magnitude faster, however, when the MVR-tree cannot be maintained in main memory. Finally, a simple combination with a relational storage of the full-precision trajectory points shows that an in-memory cache based on GraCT is an order of magnitude faster than the MVR-tree operating on disk for the basic time-slice queries.

The paper is structured as follows. Section 2 covers related work. Section 3 introduces the GraCT structure and Section 4 explains in detail the compression of the logs. Sections 5 and 6 describe how the various supported queries are processed in GraCT. Section 7 presents the experimental results. Finally, Section 8 presents our conclusions and directions for future research.

2. Related Work

We describe in this section the previous work on indexing and compressing trajectories. Section 2.1 describes the related work on indexing trajectories, Section 2.2 on compressing trajectories, and Section 2.3 on combining both. Readers willing to go directly to the description of our new index may jump directly to Section 2.4, where the basic concepts of compact data structures are covered.

2.1. Indexing trajectories

Indexes that handle free trajectories of points in the space are also called spatio-temporal indexes. One type spatio-temporal index builds on a classic multidimensional spatial index, in the form of a temporally augmented R-tree [14]. While the R-tree uses two-dimensional Minimum Bounding Rectangles (MBRs) to enclose the spatial objects stored in the database, the 3DR-tree [46] uses Minimum Bounding Boxes (MBBs), where the third dimension is time, to enclose the segments of the trajectories. The problem is that the three-dimensional space covered by an MBB can be large, even in small segments, resulting in high levels of overlap and limited power of discrimination. The STR-Tree [34] is an extension of the 3DR-tree designed to overcome this problem, by modifying the insertion/split strategy that produces the MBBs. The same work also proposes an index called TB-tree, where several segments of a trajectory are bundled into partial trajectories that are inserted as MBBs of an R-tree.

A second kind of index is a versioned R-tree, which creates an R-tree for each timestamp and a B-tree to select the relevant R-trees. Creating an R-tree for each timestamp requires large amounts of space. To overcome this, instead of storing the complete R-tree for each timestamp, these techniques store only the part of the R-tree that is different from the R-tree in the previous timestamp. Examples of indexes of this type are MR-tree [51], HR-tree [29], HR+-tree [42] and MV3R-tree [43].

A third family of indexes, called grid-based indexes, partition the space and build a temporal index for each of the spatial partitions. The Scalable and Efficient Trajectory Index (SETI) [9] divides the spatial dimension into static, non-overlapping cells. Within each cell, the trajectory segments are indexed by time using an R^* -tree. If a segment crosses the boundary between two cells, it is clipped and stored twice in the R^* -trees of both partitions. Other examples from this family include Multi Time Split B-Tree [55], Compressed Start-End tree [49], and PIST [4]. More recently, GCOTraj [52] presented a work in the same line, it also divides the space into cells, and all the subtrajectories in a cell are stored in the same disk page. Disk pages are ordered on disk following a space filling curve to optimize I/O times.

PA-trees [31] use a completely different approach to avoid MBBs. The index is built with linear and order- k polynomial approximations of the trajectories. Of course, the approximation deviates from the real trajectory, thus the index also stores the maximum deviation with respect the actual trajectory to avoid false negatives.

A recent active research line is the management of trajectories in the distributed computing framework. PRADASE [24] and CloST [41] are trajectory storage systems built on top of Hadoop and extended with a spatio-temporal index. MD-HBase [33], R-HBase [15], and GeoMesa [16] use distributed key/value stores. TrajSpark [53] is an in-memory system based on Spark and equipped with a two-level spatio-temporal index. While we do not address distributed computing in this article, we remark that, in general, a compressed storage allows distributing the spatiotemporal data across fewer computers, thereby reducing the communication costs.

The closest predecessor of our GraCT index is SEST-Index [13, 50]. This was the first structure considering snapshots and logs of events, but it did not consider compression and was oriented to storing “events” (objects that appear in or disappear from an area).

2.2. Compressing trajectories

A simple way to reduce the size of a trajectory is to produce a new trajectory that approximates the original one by selectively removing some of its points. This approach, called trajectory simplification, is based on so-called line generalization techniques. This strategy is typically used as a preprocessing step that is applied before the trajectories are compressed or indexed [54]. Therefore, these methods are compatible with other compression and indexing techniques.

The simplest trajectory simplification method is uniform sampling [35], which selects points at regular intervals (e.g., 5th, 10th, 15th, etc.), and discards the rest. Though simple to apply, the resulting trajectory may lose too much precision.

More evolved algorithms analyze the points in order to identify those containing more information, while discarding more redundant surrounding points. One of the best-known algorithms of this type is Douglas-Peucker [11]. Other methods in this family include top-down time ratio [25], reservoir sampling [47], sliding window [17], SQUISH [26], Bounded Quadrant System [23], and OPERB [22].

Trajectory compression techniques based on speed and direction include points in the approximated trajectory provided they represent a significant change in the course of the trajectory. Examples of methods from this family include dead reckoning [44], threshold-guided sampling, and STTrace [35].

Finally, some methods use the knowledge of an underlying network to decide which points are retained (e.g., those at bus stops or at street intersections) [38, 36, 40].

Independently of whether and how trajectory simplification is applied, the resulting trajectories can be compressed. The typical alternative is *delta compression*, which stores the first point in the trajectory in absolute form and all the others as a the difference with the previous point. This applies to all three components (x, y, and timestamp).

Trajic [32] uses a strategy related to delta compression. It uses a predictor of the next point of a trajectory and an encoder of the residuals, that is, the difference between the prediction and the actual point. If the residual is small, the encoder uses few bits.

Finally, a somewhat related research line faces the compression of traffic data over networks, mainly flow at intersections and average speed, using low-dimensional models such as principal component analysis or low-rank approximation [21, 1, 2]. Basically, these techniques approximately represent a matrix as the product of smaller matrices. Similar techniques are used in 3D animated video to compress the trajectories formed by the positions of the vertices of animated synthetic figures [37, 45], and could probably be applied to compress GPS trajectories.

GraCT does not discard points, but it discretizes the positions according to an underlying grid, whose size is parameterized by the needs of the applications (e.g., it makes no sense to store ship positions with a precision of centimeters). It also uses a discrete temporal scale, interpolating the measures so that each trajectory has a point at each time instant. GraCT is unique in using grammar-based compression on the trajectories.

2.3. Trajectory compression and indexing

There are just a few systems combining compression and indexing of trajectories.

TrajStore [10] is similar to SETI [9], as it also slices the trajectories into subtrajectories that fit in the cells into which the space was divided. The difference is that the data of each cell is compressed and that the cells are adapted to the distribution of the data.

TrajStore uses delta compression and a cluster-based compression, which groups several similar trajectories and stores only one of them; thus it is a lossy method. In addition, the cells are spatially indexed by a global quadtree, and each cell has a temporal index of the disk pages that make up the cell.

SharkDB [48] is an in-memory trajectory storage system that combines a column-oriented data structure and compression. SharkDB divides the time into frames of a given length (e.g., one minute) and stores one point for each trajectory within that frame. All of the points within the frame are stored as a column of a column-oriented database management system. To reduce data size, it also uses delta compression. SharkDB uses multithreading to speed up query processing.

PRESS [39] is a system that stores compressed trajectories and supports queries, but the trajectories must be paths in an underlying network (e.g., train stations, bus stops). We only consider free trajectories in this article.

2.4. Compact data structures and compression methods

Compact data structures [30] are a new family of data structures that combine query support with compressed data representation. They often use space close to that of a plain compressed representation, while supporting queries on the data in time competitive with classical data structures. Our index makes use of various compact data structures and compression methods.

2.4.1. Rank and Select

The following two operations over bit sequences (or bitmaps) are basic components of most compact data structures: $rank_b(B, p)$ counts the number of occurrences of bit b in bitmap B up to position p , whereas $select_b(B, n)$ returns the position in bitmap B where the n -th occurrence of bit b is located.

A number of data structures (see [27], for example) allow these operations to be solved in constant time, using $n + o(n)$ bits of total space (in practice, approximately 5% extra space over the original bitmap).

2.4.2. k^2 -tree

The k^2 -tree is a compact data structure originally designed to represent Web graphs within reduced space, allowing them to be navigated directly in compressed form [8]. In general, a k^2 -tree can be used to represent the adjacency matrix of any graph, as well as binary matrices.

Conceptually, the k^2 -tree is a k^2 -ary tree built from a binary matrix by recursively subdividing the matrix into k^2 submatrices of the same size. It starts by subdividing the original $n \times n$ matrix into k^2 submatrices of size n^2/k^2 . The submatrices are ordered from left to right and from top to bottom. Each submatrix generates a child of the root node whose value is 1 if there is at least one 1 in the cells of that submatrix, and 0 otherwise. The subdivision proceeds recursively for each child with value 1 until it reaches a submatrix full of 0s, or until it reaches the cells of the original matrix (i.e., submatrices of size 1×1). Figure 1 shows an example.

Instead of using a pointer-based representation, the k^2 -tree is compactly stored using bitmaps T and L (see Figure 1). T stores all the bits of the k^2 -tree, except those in the last level. The bits are placed according to a level-wise traversal: first the k^2 binary




Figure 1: Example of a binary matrix (left), the associated k^2 -tree conceptual representation (top right), and its compact representation (bottom right), for $k = 2$.




Figure 2: Example of Re-Pair compression.

values of the children of the root node, then the values of the second level, and so on. L stores the last level of the tree, consisting of cell values of the original binary matrix.

Any cell, row, column, or region of the matrix can be obtained efficiently via *rank* and *select* operations over bitmaps T and L . For example, assuming a value of 1 at position p in T , its k^2 children start at position $p_{children} = rank_1(T, p) \cdot k^2$ of T . If the children of a node return a position $p_{children} > |T|$, the actual values of the cells are retrieved by accessing $L[p_{children} - |T|]$. Similarly, the parent of position p in $T : L$ is $q - (q \bmod k^2)$, where $q = select_1(T, \lfloor p/k^2 \rfloor)$, and $q \bmod k^2$ indicates the submatrix of p within that of its parent.

2.4.3. Re-Pair

Re-Pair [20] is a grammar-based compression method. Given a sequence of integers I (called *terminals*), the method proceeds as follows: (1) it obtains the most frequent pair of integers ab in I ; (2) it adds rule $s \rightarrow ab$ to a dictionary R , where s is a new symbol not present in I (called a *non-terminal*); (3) every occurrence of ab in I is replaced by s , and (4) steps 1-3 are repeated until all pairs in I appear only once (see Figure 2). The resulting sequence after compressing I is called C . Every symbol in C represents a phrase (a sequence of one or more of the integers in I). If the length of the represented phrase is 1, then the phrase consists of an original (terminal) symbol; otherwise, the phrase is represented by a new (non-terminal) symbol. Re-Pair can be implemented in linear time and a phrase may be recursively expanded in optimal time (i.e., proportional to its length).

2.4.4. DACs

Directly Addressable Codes (DACs) [5] are a variable-length encoding scheme for sequences of integers that supports fast direct access to any given position in a sequence. In other words, it allows the i^{th} integer to be decoded without the need to decompress the preceding integers. If the sequence of integers has many small numbers and few large ones, then DACs obtain a very compact representation.

Given a sequence of integers $X = x_1, x_2, \dots, x_n$, DACs take the binary representation of that sequence and rearrange it into a level-shaped structure as follows: the first level, B_1 , contains the first (least significant) n_1 bits of the binary representation of each integer. A bitmap C_1 is added to indicate whether the binary representation of each integer requires more than n_1 bits (1) or not (0). In the second level, B_2 stores the next n_2 bits of the integers with a value of 1 in B_1 . A bitmap C_2 marks the integers that need more than $n_1 + n_2$ bits, and so on. This process is repeated for as many levels as needed. The number of levels ℓ and the number n_l of bits at each level l , with $1 \leq l \leq \ell$, is calculated in order to maximize compression. Each value x_i is then retrieved using less than ℓ rank operations on the bitmaps C_l and extracting chunks from the arrays B_l .

2.4.5. Permutations

A permutation of the integers $[1..n] = \{1, \dots, n\}$ is a reordering of the values $\{1, \dots, n\}$. The two main operations over the reordered sequence are $\pi(i)$ and $\pi^{-1}(j)$. The first yields the number at position $i \in [1..n]$ in the sequence, while the second identifies the position in the sequence where $j \in [1..n]$ appears.

The simplest way to store a permutation is an array $P[1..n]$, where $P[i] = \pi(i)$, so that $\pi(i)$ is answered in constant time. In order to answer $\pi^{-1}(j)$ efficiently, we can double the space by storing a second array. We use instead an intermediate option that uses $(1 + \epsilon)n$ instead of $2n$ cells and answers $\pi^{-1}(j)$ queries in time $O(1/\epsilon)$ [28].

3. The GraCT Index

3.1. Discretization

GraCT represents moving objects that follow free trajectories in space. The method assumes that the positions of all the objects are synchronized and stored at regular time instants (e.g., every minute). Since GPS devices may report the positions of objects at times that do not match the required time instants, the positions in the desired time instants are obtained from the raw timestamped positions using interpolation.

A raster model is used to represent the space, which is divided into cells of a fixed size, and objects are assumed to fit in one cell. The size of the cells and the length of the period between represented time instants are parameters that can be adapted to the specific domain. Obviously, the shorter is the length of the period and the smaller is the size of the cell, the more accurate will be the trajectory data, but less compression will be achieved. This is a tradeoff between space usage and precision in the storage.

Alternatively, we may regard this discretization as a way to obtain a more compact in-memory index, while the full data is maintained on disk and accessed only to filter out false-positives from a usually small set of candidate answers. We will also experiment with this variant of the index, showing that this is better than directly operating on the full data on disk.

3.2. Snapshots

Every d time instants, GraCT uses a data structure based on k^2 -trees to represent the absolute positions of all the objects. These data structures are called *snapshots*. The distance d between snapshots is another parameter of the system, which trades space for query time. Between two consecutive snapshots, the trajectory of each moving object is represented as a *log*, which is an array of movements represented in differential form.

Let \mathcal{S}_h denote the snapshot representing the position of all the objects at time instant k . The components of \mathcal{S}_h are:

- The *time instant* represented by the snapshot.
- A k^2 -tree storing the positions in the space (i.e., the cells of the raster) where there are objects.
- An array of integers *perm* storing the identifiers of the objects at each cell in the space where there are objects.
- A bitmap Q , which is an auxiliary data structure of *perm* that is used to determine the correspondence between positions in the space and the positions of *perm*.

The k^2 -tree represents a binary matrix where a cell set to 1 indicates that the cell contains one or more objects. However, we need to know what objects are in the cell set to 1. Each 1 in the binary matrix corresponds to a bit set to 1 in bitmap L of the k^2 -tree. The list of object identifiers corresponding to each of those bits set to 1 is stored in *perm*, where the object identifiers are sorted according to their order of appearance in L . This array turns out to be a permutation of all the object identifiers. Bitmap Q is aligned with *perm*; a 0 at a given position of Q means that the object identifier aligned in *perm* is the last object in the cell, while a 1 signals that the next object identifier in *perm* is in the same cell.

Figure 3 shows an example. The left side shows the matrix representing the space and the objects placed in the cells (these cells are the same containing a 1 in Figure 1). The four arrays on the right are the actual data structures that represent the snapshot information. Arrays T and L are the same as those in Figure 1 and mark the positions having objects. The objects in those positions are stored in arrays Q and *perm*. For instance, the object identifiers corresponding to the second 1 in L (which is at position 5 in L and corresponds to the cell at position (9,5)) are stored starting at position 3 in *perm*, since the first 0 (signaling the end of the entries corresponding to the first 1 in L) is at position 2. In order to find out how many objects there are in the corresponding cell, we access Q starting at position 3 and search for the first 0 after that position, which is at position 4; this shows that there are two objects in the inspected cell. By accessing positions 3 and 4 in *perm*, we obtain the object identifiers 4 and 5. The object identifiers corresponding to the third 1 in L start at position 5 in *perm*, and so on.

With these structures used to represent the absolute positions of all moving objects at the snapshots, we can answer two types of queries in time $O(\log_k s + \text{output})$, where $s \times s$ is the grid size and *output* is the size of the query output:

- *Find the objects in a given cell*: First, we traverse the tree from the root until we reach position n in L corresponding to that cell. Next, we count the number of 1s




Figure 3: Position of objects in the space (left) and the arrays (right) that make up a snapshot representing that situation.

in the array of leaves L until position n ; this gives us the number of leaves with objects up to the n -th leaf, $x = \text{rank}_1(L, n)$. Then we calculate the position of the $(x - 1)$ -th 0 in Q , which indicates the last bit of the previous leaf with objects, and we add 1 to obtain the first position in our leaf. That is, $p = \text{select}_0(Q, x - 1) + 1$ is the position in perm of the first object identifier corresponding to the searched position. From p , we read all the object identifiers aligned with 1s in Q until we reach a 0, which signals the last object identifier in that leaf.

- *Find the position of a given object in the space.* First, by using the π^{-1} operation (see Section 2.4.5), we obtain the position k in perm of the searched object. Next, we have to find the leaf in L corresponding to the k -th position of perm . For this, we calculate the number of leaves with objects before the object in position k of perm ; that is, we count the number of 0s until the position before k , $y = \text{rank}_0(Q, k - 1)$. Then we find in L the position of the $(y + 1)$ -th 1; that is, $\text{select}_1(L, y + 1)$. With that position in L , we can traverse the k^2 -tree upwards in order to obtain the position in the space of that cell, and thus the position of the object.

3.3. Log of relative movements

The movements of objects between two consecutive snapshots are recorded in sequences of consecutive object positions called *logs*, with one separate sequence per object. Logs actually represent the positions in relative form with respect to the previous one, where the first absolute position can be obtained from the preceding snapshot. Therefore, logs can be thought of as storing the consecutive *movements* of the objects along time. As we will see soon, this model must be extended to consider cases where objects appear and disappear.

Objects change their positions along the two Cartesian axes, so every movement in the log can be described with two integers. In order to save space, our method encodes the two values using a single positive integer. The cells around the actual position of an object are enumerated following a spiral in which the origin is the initial position of the object, as shown in Figure 4 (left). As an example, assume that an object moves one




Figure 4: Encoding movements.

cell to the East and one cell to the North with respect to the previous known position. Instead of encoding the movement as the pair (1,1), it is encoded as an 8. Figure 4 (right) shows the trajectory of an object starting at cell (0,2). Each number indicates a movement between two consecutive time instants. Since most relative movements involve short distances, this technique usually produces a sequence of small numbers.

Between two consecutive snapshots \mathcal{S}_h and \mathcal{S}_{h+d} , each object is represented by a log, $\mathcal{L}_{h,h+d}(id)$, where id is the identifier of the object.

Figure 5 presents a situation in which there is a snapshot every eight time instants. Time instants t_0 and t_8 are represented with snapshots, while the time instants in the intervals $[t_1, t_8]$ and $[t_9, t_{16}]$ are represented with logs. Note that, although t_8 is represented with a snapshot, it is also included in the preceding logs. This duplicity allows us traversing the logs in both directions, which accelerates certain queries, as seen later. The meaning of the other components of the figure is also explained in the sequel.

3.3.1. Managing appearances and disappearances

Because the model does not explicitly store the time instant of each recorded position, a mechanism is required to signal which objects are active at each time instant and to determine to which time instant each log entry corresponds.

For this purpose, the log manages three events: when an object starts emitting its positions, when it stops, and when it stops emitting signals for periods of time. In fact, the three cases are homogeneously managed by the following events, each identified in the log with a special codeword:

- *Disappearances (D)*. The occurrence of codeword D in $\mathcal{L}_{h,h+d}(id)$ means that object id stopped emitting its position at the corresponding time instant of that portion of the log, and that id did not restart emitting (at least) until the time instant of the next snapshot (\mathcal{S}_{h+d}).

To enable backward traversal of the logs, each disappearance must store the time instant and the absolute coordinates of the object by the time it disappeared.

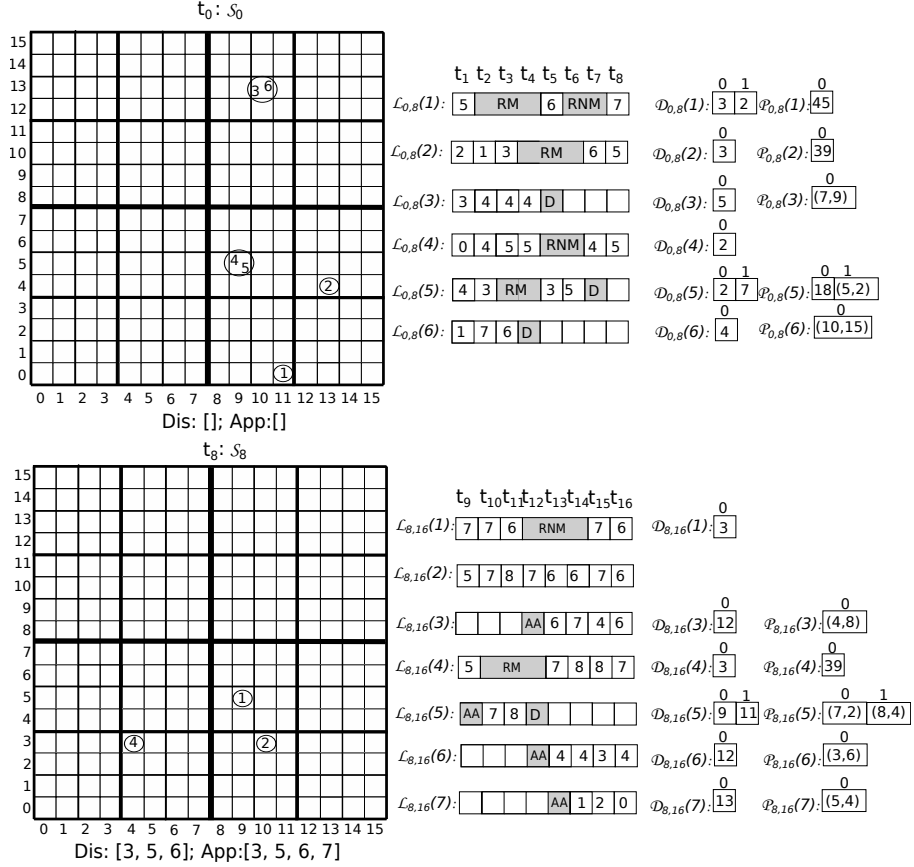


Figure 5: Snapshots and logs.

- *Absolute appearance (AA)*. The occurrence of codeword AA in $\mathcal{L}_{h,h+d}(id)$ means that id started to emit its position in the corresponding time instant of that portion of the log, and that id had not emitted its positions (at least) since the previous snapshot (\mathcal{S}_h). AA may signal the first appearance of an object or its reappearance after having disappeared in a preceding log.

To enable forward traversal of the logs, an absolute appearance stores the time instant and the absolute coordinates where the object appears. Note that, even if the object is reappearing, it might do so far away from the position where it had disappeared in a previous log.

- *Relative disappearance*. The occurrence of a codeword indicating a relative disappearance in $\mathcal{L}_{h,h+d}(id)$ in an entry corresponding to time instant t_i means that id stopped emitting positions at that time instant, and that id restarted emitting its positions in a time instant between t_i and the time instant of the next snapshot (\mathcal{S}_{h+d}). In other words, a relative disappearance occurs when the object disappears and reappears within the same log.

A relative disappearance must store the number of time instants in which id was not emitting its positions, and possibly a movement, depending on the type of relative disappearance:

- *Relative disappearance with movement (RM)* means that id appears in a different position with respect to its last known position. This requires the log to store the relative movement (using the spiral encoding) with respect to that last known position.
- *Relative disappearance without movement (RNM)* means that id appears in the same position where the signal was lost, therefore there is no need to store a movement.

The extra information is stored in two auxiliary arrays associated with each log $\mathcal{L}_{h,h+d}(id)$:

- Array $\mathcal{D}_{h,h+d}(id)$ stores:
 - in the case of relative disappearances, the time the disappearance lasted;
 - in the case of an absolute appearance or a disappearance, the absolute time instant of that event.
- Array $\mathcal{P}_{h,h+d}(id)$ stores:
 - in the case of relative disappearances with movement, the relative movement with respect to the last known position;
 - in the case of an absolute appearance or a disappearance, the absolute position of that event.

Figure 5 shows two snapshots; \mathcal{S}_0 is the snapshot of Figure 3. The right-hand side shows the logs representing the movements of objects in the time instants after those snapshots, thus representing the positions of objects between time instants t_0 and t_{16} . Object 1 disappears at t_2 and then reappears at t_5 (see $\mathcal{L}_{0,8}(1)$). Since the time instants occur between the same pair of snapshots, this is a relative disappearance. The number of time instants that the disappearance lasts is stored in array $\mathcal{D}_{0,8}(1)$. In the first position of $\mathcal{D}_{0,8}(1)$ there is a 3, which means that the first disappearance of $\mathcal{L}_{0,8}(1)$ lasted three time instants. In the same portion of the log, there is another relative disappearance between time instants t_6 and t_7 , therefore in the second entry of array $\mathcal{D}_{0,8}(1)$ there is a 2, indicating that the disappearance lasted two time instants. The relative disappearance that starts at t_2 is of type *RM*, so the log must store the position where object 1 reappears. Array $\mathcal{P}_{0,8}(1)$ stores this position using spiral encoding with respect to the last known position. In the example, a 45 means that object 1 reappeared three cells to the North.

In the entry in $\mathcal{L}_{0,8}(3)$ corresponding to time instant t_5 , codeword *D* means that object 3 disappears at that time instant and does not reappear in $\mathcal{L}_{0,8}(3)$. Therefore, $\mathcal{D}_{0,8}(3)$ stores a 5 in its first position (there are no prior relative reappearances), indicating that the object disappeared at t_5 , and in $\mathcal{P}_{0,8}(3)$, the only entry stores the coordinates (7,9), which denote that the object was in that position when it disappeared. Object 3 reappears in $\mathcal{L}_{9,16}(3)$, in the entry corresponding to t_{12} . This reappearance is signaled

- $s \rightarrow ab$ is the normal rule of Re-Pair,
- $\#t$ is the number of time instants covered by the rule,
- (x, y) are the relative coordinates of the final position of the object after the application of the rule (that is, the total displacement in both coordinates), and
- MBR is the minimum bounding rectangle enclosing the movements of the rule. MBR is represented as $(x1, y1, x2, y2)$, where $(x1, y1)$ corresponds to the bottom-left corner and $(x2, y2)$ to the top-right corner.

For example, the rules in Figure 2 are enriched as follows. The first rule of R is $W \rightarrow 2, 9, 2, (3,0), (0, -1, 3, 0)$: 2 and 9 are the substituted symbols; the next 2 indicates that the rule represents a sequence of two movements; $(3,0)$ indicates the position of the object after the application of the rule if we start at $(0,0)$, and the last four values are two corners (bottom-left and top-right) defining a rectangle that encloses all the movements encoded by the rule.

With this additional information, most of the non-terminal symbols of a compressed log file \mathcal{L} do not need to be decompressed in order to obtain the position of an object at a given time instant between two snapshots. Consider, for example, the compressed log $\mathcal{L} = Z, Y, 9, Y, X, Z, X$ of Figure 6 (this is the array C of Figure 2). Assume that we want to find the position of the object at the 5th time instant, that is, $(7, 2)$. From the preceding snapshot we obtain that the absolute position of the object at the beginning of the log is $(0,1)$. Next, we inspect the log \mathcal{L} from the beginning. The first value is a Z . The enriched rule indicates that this symbol represents four time instants, after which the object is displaced six columns to the East while remaining in the same row. Since it started at $(0, 1)$ the object will be at position $(6, 1)$ at t_4 . Since our target time instant is $t_5 > t_4$, we can skip the whole nonterminal Z . The next symbol is Y , which lasts two time instants. This would take us to t_6 , which surpasses our target time instant. Only in the last step of the search, therefore, we have to decompress the rule and process its components: $Y \rightarrow 8\ 7$. The 8 is a terminal symbol that lasts one time instant, and thus it is enough to reach our target time instant. An 8 moves one column to the East and one to the North, which, applied to the current position $(6, 1)$, takes us to the correct position $(7, 2)$.

The MBR component also helps speeding up other queries, as seen later.

These additional elements that enrich the rules are compressed with DACs. To obtain better compression, the times of all of the rules are compressed using one DAC, separately from the three pairs of coordinates of all of the rules, which are compressed using a separate DAC. Figure 7 shows the resulting data structures that make up the GraCT index for our running example in Figure 5.

Finally, as a summary, Figure 8 shows a flow chart of the whole process to create a GraCT data structure.

5. Query Preliminaries

This section presents a series of definitions that will be used later.




Figure 7: GraCT structures.

Definition 5.1. Let t_b and t_e be two time instants where $t_b \leq t_e$. Then $CT_O(t_b, t_e) = (m_1, m_2, \dots, m_j)$ represents the compressed trajectory (that is, the movements) of object O between time instants t_b and t_e . Each m_i can be:

- $S_{(x,y)}$, an absolute position (x, y) obtained from a snapshot.
- NT/T , a grammar symbol encoding the movements of O , where NT denotes a non-terminal and T denotes a terminal.
- $D_{t,(x,y)}$, a disappearance code, where t is the time instant of the disappearance and (x, y) is the object position at that time instant.
- $RNM_{\#}$, a reappearance without movement code, where $\#$ is the duration of the disappearance.
- $RM_{\#,value}$, a reappearance with movement code, where $\#$ is the duration of the disappearance and value is the spiral code representing the movement.
- $AA_{t,(x,y)}$, an absolute appearance code, where t is the time instant of the appearance and (x, y) is the object position at that time instant.

If object O is missing at t_b , then m_1 will refer to the first time instant after t_b containing information about the position of O . In the same way, if the object is missing

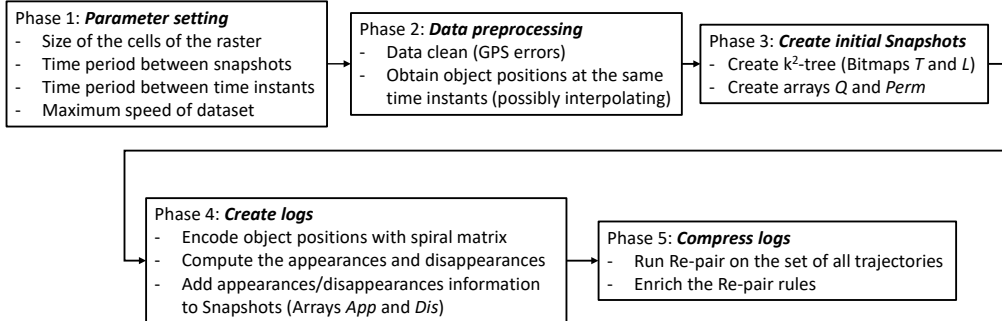


Figure 8: Flow chart of the process to create a GraCT index.

at t_e , then m_j will be the movement of O corresponding to the closest time instant with information about the position of O prior to t_e .

The sequence $CT_O(t_b, t_e)$ is obtained by extracting the portions of the logs and the positions at the snapshots of O covering the time interval $[t_b, t_e]$.

For example, in our running example in Figures 5 and 7, for object O_6 , the compressed trajectory between 0 and 13 is $CT_{O_6}(0, 13) = (\mathcal{S}_{(10,13)}, 1, V, D_{4,(10,15)}, AA_{12,(3,6)}, Y)$. Note that the last code (Y) covers time instants 13 and 14, since it is a non-terminal symbol. In contrast, in $CT_{O_5}(11, 15) = (AA_{11,(7,2)}, W, D_{13,(8,4)})$. The last code ($D_{13,(8,4)}$) corresponds to time instant 13, since there are no positions after that time instant in the requested time interval.

Definition 5.2. If m_i is an element of $CT_O(t_b, t_e)$, p_c is the absolute position of object O at time instant t_c , and $t_b \leq t_c \leq t_e$, then operation $moveJ(p_c, t_c, t_e, m_i)$ returns either:

- $\langle t_{nc}, p_{nc} \rangle$, where p_{nc} is the position resulting from the application of the movements encoded in m_i to p_c , and $t_{nc} \leq t_e$ is the time instant of that position, or
- $\langle t_e, p_e \rangle$, where p_e is the position at t_e , if m_i encodes movements surpassing t_e .

The name $moveJ$ is the abbreviation of “move jump”. If m_i is a non-terminal code, it uses $m_i.(x, y)$, the relative coordinates of the enriched information included in the definition of the non-terminal symbol m_i , which, when applied to the previous position, gives the new position of the object after the application of all the movements encoded by m_i . For example, $moveJ((9, 5), t_1, t_3, Z)$ uses the relative (x, y) coordinates of the definition of $Z \rightarrow 4, 5, 2, (-2, -1), (-2, -1, 0, 0)$, that is $(-2, -1)$, which, when applied to $(9, 5)$, obtain the new position $(7, 4)$. Note that the process never obtains the position at t_2 but jumps from $\langle t_1, (9, 5) \rangle$ to $\langle t_3, (7, 4) \rangle$.

Note also that if we issue $moveJ((9, 5), t_1, t_2, Z)$, where the application of Z would produce the position at $t_3 > t_2 = t_e$, the operation has to decompress Z into its components (4 and 5) and apply them one by one until it reaches t_e . In this case, therefore, only code 4 is applied, to move the object to $(8, 4)$. The decompression of non-terminal codes in $moveJ$ only occurs in this situation.

In this example we fully decompress the nonterminal that exceeds t_e , but in general we continue advancing recursively on the terminals and nonterminals of the right-hand

$CT_O(t_b, t_e)$	compressed trajectory of O between time instants t_b and t_e
$moveJ(p_c, t_c, t_e, m_i)$	returns the pair $\langle t_{nc}, p_{nc} \rangle$ resulting from applying the movements in m_i to position p_c , but if that application produces a position surpassing t_e , then the operation returns $\langle t_e, p_e \rangle$.
$RmoveJ(p_c, t_b, t_c, m_i)$	returns the pair $\langle t_{nc}, p_{nc} \rangle$ resulting from undoing the movements in m_i to position p_c , but if that application produces a position previous to t_b , then the operation returns $\langle t_b, p_b \rangle$.
$moveS(p_c, t_c, t_e, m_i)$	returns $\langle (t_{c+1}, p_{c+1}), \dots, (t_h, p_h) \rangle$, the list of time instants/positions resulting from applying m_i to p_c , except if m_i encodes movements that surpass t_e , in which case $h = e$.
$ER_r(t_b, t_e)$	expansion of region r from t_b to t_e using the maximum speed
$ER_O(t_b, t_e)$	expanded region of object O from t_b to t_e using the maximum speed
$\overline{p_i p_j}$	distance in cells of the shortest trajectory from point p_i to point p_j
\overline{pr}	distance in cells of the shortest trajectory from point p to region r

Table 1: Summary of definitions.

side of its rule. That is, if the application of the next nonterminal Z would exceed t_e , and the rule associated with Z is $Z \rightarrow Z_1 Z_2$, then we replace Z by $Z_1 Z_2$ and try again. This means that, if applying Z_1 also exceeds t_e we will continue recursively with Z_1 , and otherwise we will process the whole Z_1 in constant time and continue recursively with Z_2 . Therefore, if we are indexing t time instants in total and use a balanced grammar (i.e., of height $O(\log t)$), then $moveJ$ can be implemented in time $O(1)$ if m_i does not exceed t_e and $O(\log t)$ if it does.

Definition 5.3. *If m_i is an element of $CT_O(t_b, t_e)$, p_c is the absolute position of object O at time instant t_c , and $t_b \leq t_c \leq t_e$, then $RmoveJ(p_c, t_b, t_c, m_i)$ returns either:*

- $\langle t_{nc}, p_{nc} \rangle$, where p_{nc} is the position resulting from undoing the movements encoded in m_i to p_c , and t_{nc} is the time instant of that position, or
- $\langle t_b, p_b \rangle$, where p_b is the position at t_b , if m_i encodes movements that obtain a time instant earlier than t_b .

This is a reverse version of $moveJ$, and can be implemented within the same time complexity.

Definition 5.4. *If m_i is an element of $CT_O(t_b, t_e)$, p_c is the absolute position of object O at time instant t_c , and $t_b \leq t_c \leq t_e$, then $moveS(p_c, t_c, t_e, m_i)$ returns the list of pairs $\langle (t_{c+1}, p_{c+1}), \dots, (t_h, p_h) \rangle$ where $t_h \leq t_e$, which is the list of time instants/positions resulting from applying m_i to p_c . If m_i encodes movements that surpass t_e , however, the decompression ends at t_e .*

The name $moveS$ is the abbreviation of “*move step by step*”. Unlike $moveJ$, if m_i is a non-terminal code, $moveS$ decompresses the symbol until it obtains terminal codes, then applies those terminal symbols to p_c and returns the positions of each symbol.

For example, $moveS((9, 5), t_1, t_3, Z)$ decompresses Z to obtain the terminal codes 4 and 5, and thus returns $\langle t_2, (8, 4) \rangle, \langle t_3, (7, 4) \rangle$. In contrast to $moveJ$, $moveS$ does obtain the position at t_2 . This operation can be implemented in time $O(h - c + \log t)$ on a balanced grammar.

Definition 5.5. *Let $r = [x_1, y_1] \times [x_2, y_2]$ be a rectangular region in the two-dimensional space. Let t_b and t_e be two time instants, where $t_b < t_e$, and let M_{sp} be the maximum speed of the dataset. Then $ER_r(t_b, t_e)$ denotes the expanded region of r from t_b to t_e using M_{sp} : $ER_r(t_b, t_e) = [x_1 - \delta, y_1 - \delta] \times [x_2 + \delta, y_2 + \delta]$, where $\delta = M_{sp} \cdot (t_e - t_b)$.*

The maximum speed of the dataset is the maximum speed of any object during the period of time covered by the dataset. $ER_r(t_b, t_e)$ is the rectangle resulting from expanding r in all directions (North, South, West and East) so that any object within r at t_b is bound to be within $ER_r(t_b, t_e)$ at t_e . In other words, to obtain $ER_r(t_b, t_e)$, r is enlarged by $M_{sp} \cdot (t_e - t_b)$ cells in all four directions.

Definition 5.6. Let $t_b < t_e$ be two time instants, and let M_{sp} be the maximum speed of the dataset. Then $ER_O(t_b, t_e)$ denotes the expanded region of object O from t_b to t_e using the maximum speed M_{sp} . That is, if O is at cell (x, y) , then $ER_O(t_b, t_e) = ER_{r_{(x,y)}}(t_b, t_e)$, where $r_{(x,y)} = [x, y] \times [x, y]$.

That is, $ER_O(t_b, t_e)$ delimits the area where object O is bounded to belong at t_e even if running at maximum speed M_{sp} in any direction, starting at time instant t_b .

Definition 5.7. Given a point p_i and a point p_j , the distance between p_i and p_j ($\overline{p_i p_j}$) is the length of the shortest trajectory (in number of cells) from p_i to p_j , using Euclidean distance.

Definition 5.8. Given a point p and a region r , the distance between r and p (\overline{pr}) is the length of the shortest trajectory (in number of cells) between $r = [x_1, x_2] \times [y_1, y_2]$ and $p = (x, y)$. This is zero if $x_1 \leq x \leq x_2$ and $y_1 \leq y \leq y_2$. Otherwise, it is $\min(|x - x_1|, |x - x_2|)$ if $y_1 \leq y \leq y_2$ and $\min(|y - y_1|, |y - y_2|)$ if $x_1 \leq x \leq x_2$. Otherwise, it is the minimum Euclidean distance from p to the four corners of r .

Table 1 summarizes these definitions.

6. Supporting Queries

We now describe how each of the supported queries are answered using the structures of GraCT.

6.1. Obtain the position of an object in a given time instant

Algorithm 1 shows the pseudocode to answer this query. The algorithm first identifies a time instant close to t_q where the position of the object is represented in absolute form and, from that time instant, applies the relative movements in order to obtain the absolute position of the object at time instant t_q .

The first option to find that time instant is the snapshot \mathcal{S}_h that is nearest to the queried time instant (line 1 of the algorithm). If t_q coincides with the time instant of \mathcal{S}_h , then the position of the object in the snapshot is returned (line 28; the position could be *null* if the object is not present in \mathcal{S}_h). The other possible cases are handled in lines 4-15 (when t_q follows the snapshot) and 16-27 (when t_q precedes the snapshot).

In the first case, t_c stores the time instant and p_c the absolute position at any step of the log traversal. In line 4, t_c and p_c store the information of the object at \mathcal{S}_h . If the object is missing at \mathcal{S}_h (i.e., $p_c = \text{null}$), line 6 checks if O_{id} appears in the immediate next portion of the log ($\mathcal{L}_{h,h+d}(O_{id})$), as signaled by the presence of O_{id} in $\mathcal{S}_h.App$. In line 7, the algorithm obtains from the first entry of $\mathcal{D}_{h,h+d}(O_{id})$ the time instant of the appearance of O_{id} . If the appearance is prior to or occurs exactly at t_q , then the

Algorithm 1: SearchObject(O_{id}, t_q)

```
1  $\mathcal{S}_h \leftarrow \text{nearestSnapshot}(t_q)$ ;  
2  $p_c \leftarrow \mathcal{S}_h.\text{findObject}(O_{id})$ ; // Finds position  $O_{id}$  in the closest snapshot  
3  $t_c \leftarrow \mathcal{S}_h.\text{time}$ ;  
4 if  $t_c < t_q$  then // Move forward from preceding snapshot  
5   if  $p_c = \text{null}$  then //  $O_{id}$  is not in the snapshot  
6     if  $O_{id} \in \mathcal{S}_h.\text{App}$  then // Appears in  $\mathcal{L}_{h,h+d}(O_{id})$   
7        $t_c \leftarrow \mathcal{D}_{h,h+d}(O_{id})[0]$ ; // Obtain the time of appearance  
8       if  $t_c \leq t_q$  then  
9          $p_c \leftarrow \mathcal{P}_{h,h+d}(O_{id})[0]$ ; // Obtain the location of appearance  
10      else return null;  
11     else return null;  
12    $T \leftarrow CT_{O_{id}}(t_{c+1}, t_q)$ ; // The trajectory, as a stack (first on top)  
13   while  $T$  is not empty do  
14      $m \leftarrow T.\text{pop}$ ;  
15      $\langle t_c, p_c \rangle \leftarrow \text{moveJ}(p_c, t_c, t_q, m)$ ;  
16 else if  $t_c > t_q$  then // Move backward from following snapshot  
17   if  $p_c = \text{null}$  then //  $O_{id}$  is not in the snapshot  
18     if  $O_{id} \in \mathcal{S}_h.\text{Dis}$  then // Disappears in  $\mathcal{L}_{h-d,h}(O_{id})$   
19        $t_c \leftarrow \mathcal{D}_{h-d,h}(O_{id})[\text{last}]$ ; // Obtain the time of disappearance  
20       if  $t_c \geq t_q$  then  
21          $p_c \leftarrow \mathcal{P}_{h-d,h}(O_{id})[\text{last}]$ ; // Obtain the location of disappearance  
22       else return null;  
23     else return null;  
24    $T \leftarrow CT_{O_{id}}(t_q, t_{c-1})$ ; // The trajectory, as a stack (last on top)  
25   while  $T$  is not empty do  
26      $m \leftarrow T.\text{pop}$ ;  
27      $\langle t_c, p_c \rangle \leftarrow \text{RmoveJ}(p_c, t_q, t_c, m)$ ;  
28 if  $t_c = t_q$  then return  $p_c$ ;  
29 else return null;
```

algorithm obtains the absolute position from $\mathcal{P}_{h,h+d}(O_{id})$. Otherwise, the object does not have a position at t_q and the algorithm returns *null*.

Once we have the absolute position of the object, lines 12-15 traverse the compressed trajectory using *moveJ* operations from time instant t_c (either the closest snapshot or an absolute appearance) to t_q . If the algorithm is able to reach t_q , the position is returned.

Lines 16-27 perform the same process as lines 4-15 but in reverse, for the case where t_q occurs before the snapshot.

If there are t time instants, n objects, the grid is of size $s \times s$, k is the arity of our k^2 -tree, d is the number of time instants between two snapshots, and our grammar is balanced, then this query takes time $O(d + \log t + \log_k s + \log n)$.

6.2. Obtain the trajectory of an object between two time instants

Algorithm 2 shows the pseudocode to answer this query. It first obtains a time instant preceding t_b with an absolute position of the object, and then it simply follows the trajectory until it reaches t_e . Lines 1-12 are responsible for finding a close preceding time instant with an absolute value, lines 13-17 then advance until reaching t_b , and finally lines 18-26 collect the trajectory from t_b to t_e .

Line 1 obtains the preceding snapshot closest to t_b (\mathcal{S}_h). If the snapshot has information about the position of the object, p_c stores that position and t_c is the time instant

Algorithm 2: Search Trajectory (O_{id}, t_b, t_e)

```
1  $\mathcal{S}_h \leftarrow \text{nearestPreviousSnapshot}(t_b)$ ; // The closest snapshot preceding  $t_b$ 
2  $p_c \leftarrow \mathcal{S}_h.\text{findObject}(O_{id})$ ; // The position in the preceding snapshot
3  $t_c \leftarrow \mathcal{S}_h.\text{time}$ ;
4 while  $p_c = \text{null}$  and  $t_c \leq t_e$  do
5   if  $O_{id} \in \mathcal{S}_h.\text{App}$  then // Appears before next snapshot
6      $t_c \leftarrow \mathcal{D}_{h,h+d}(O_{id})[0]$ ;
7     if  $t_c \leq t_e$  then  $p_c \leftarrow \mathcal{P}_{h,h+d}(O_{id})[0]$ ;
8   else // Move to next snapshot
9      $\mathcal{S}_h \leftarrow \mathcal{S}_h.\text{nextSnapshot}()$ ;
10     $p_c \leftarrow \mathcal{S}_h.\text{findObject}(O_{id})$ ;
11     $t_c \leftarrow \mathcal{S}_h.\text{time}$ ;
12 if  $t_c > t_e$  then return null; // Does not appear before  $t_e$ 
13 if  $t_c < t_b$  then // Skip until reaching  $t_b$ 
14    $T \leftarrow CT_{O_{id}}(t_c, t_{b-1})$ ; // The trajectory, as a stack (first on top)
15   while  $|T| > 1$  do
16      $m \leftarrow T.\text{pop}$ ;
17      $\langle t_c, p_c \rangle \leftarrow \text{moveJ}(p_c, t_c, t_{b-1}, m)$ ;
18  $\text{Answer} \leftarrow \langle \rangle$ ;
19  $T \leftarrow CT_{O_{id}}(t_{c+1}, t_e)$ ; // The trajectory, as a stack (first on top)
20 while  $T$  is not empty do
21    $m \leftarrow T.\text{pop}$ ;
22   foreach  $\langle t_j, p_j \rangle$  in  $\text{moveS}(p_c, t_c, t_e, m)$  do
23     if  $t_j \geq t_b$  then  $\text{Answer.add}(\langle t_j, p_j \rangle)$ ;
24    $t_c \leftarrow t_j$ ;
25    $p_c \leftarrow p_j$ ;
26 return Answer
```

of the snapshot. In lines 4-12, the algorithm deals with the case in which the object is not in \mathcal{S}_h . If the object appears in the next portion of the log ($O_{id} \in \mathcal{S}_h.\text{App}$), then the algorithm obtains from $\mathcal{D}_{h,h+d}(O_{id})[0]$ the time instant of that appearance (t_c); if that time instant does not exceed t_e , it also obtains the absolute position, which is stored in p_c . Otherwise, the algorithm moves to the next snapshot, until finding an absolute position or exceeding the time instant t_e .

If the flow reaches line 13, the algorithm has found a suitable timestamp $t_c \leq t_b$ for which the absolute position of the object (p_c) is known. If $t_c < t_b$, lines 13-17 skip entries of the log(s) until $t_c = t_b$ and p_c is the position of the object at t_b . Finally, lines 18-26 traverse from instants t_b to t_e , this time collecting all the positions in a sequence Answer , which is finally returned.

Compared to the complexity of Algorithm 1, we have an additional cost of $O((t_e - t_b + 1)(1 + (\log n + \log_k s)/d))$, which can be made the optimal $O(t_e - t_b + 1)$ by choosing a sufficiently large $d = \Omega(\log n + \log_k s)$. This is, however, a worst-case complexity that ignores the fact that the log is compressed: on average, the term $O((\log n + \log_k s)(t_e - t_b + 1)/d)$ must be multiplied by the compression ratio.

6.3. Time slice query

Given a time instant t_q and a window rectangle $r = [x_1, x_2] \times [y_1, y_2]$ of the space, this query returns the objects lying within r at time t_q , and their positions. Algorithm 3 shows the pseudocode.

The process starts by obtaining the snapshot nearest to t_q , \mathcal{S}_h . The algorithm considers three cases. Firstly, if t_q is the time instant of \mathcal{S}_h , the algorithm just needs to access the snapshot in order to obtain the objects within r . The second case occurs when t_q is between the snapshots \mathcal{S}_h and \mathcal{S}_{h+d} . In this case, the algorithm has to follow the trajectory of the objects from \mathcal{S}_h until time t_q , to see if they end up within r . Instead of following all of the objects in the snapshot, the algorithm tracks only those within $ER_r(\mathcal{S}_h.time, t_q)$, ignoring objects that have no chance of being in the answer. The algorithm must also include in the tracked objects those that are not present in \mathcal{S}_h but appear at a later time instant, not exceeding t_q , and are capable of reaching r at t_q .

Once this list of candidates has been obtained, the algorithm follows the log to obtain the trajectory of the tracked objects. However, the algorithm prunes the tracking further as it processes the log: a candidate object may follow a direction that takes it away from region r , so the algorithm rechecks the condition after applying each code of the compressed trajectory, discarding the objects as soon as they become incapable of reaching r at time t_q .

The third case is analogous, and occurs when t_q is between \mathcal{S}_{h-d} and \mathcal{S}_h .

To illustrate the idea using the running example shown in Figures 5 and 7, Figure 9 outlines the process of a time slice query where $t_q = t_{10}$ and $r = [7, 10] \times [3, 4]$. Since the target time instant is t_{10} , the closest snapshot is \mathcal{S}_8 . Figure 9(a) shows the target region r and the extended region $ER_r(t_8, t_{10})$ superimposed. In this example, $M_{sp} = 1$ (that is, an object can only move one position per time instant), so $ER_r(t_8, t_{10})$ is r enlarged by two cells in all directions.

Line 1 obtains the snapshot closest to t_{10} (\mathcal{S}_8). Since t_q is not the time instant of the snapshot, the algorithm retrieves in line 4 all the objects in the extended region $ER_r(\mathcal{S}_h.time, t_q) = [5, 12] \times [1, 6]$, which are stored in the variable *Candidates*; in our example, these are the objects O_2 and O_1 . Next, in lines 7-11, for each object in $\mathcal{S}_h.App$, the algorithm obtains the time and position of its appearance (t_c and p_c); if that time is t_q or earlier, and that position is within $ER_r(t_c, t_{10})$, then the object is added to the candidates to be tracked. In our example, objects O_3 , O_6 and O_7 appear after t_{10} , therefore they are discarded. Instead, O_5 appears at t_9 , so it is added to the list of candidates, because it is also within $ER_r(t_9, t_{10})$ (see Figure 9(b)). The resulting list of candidates in our example comprises O_1 , O_2 and O_5 .

In lines 12-20, the algorithm follows the movements of each candidate object using the log, until it reaches t_q or until the object can be discarded. There are two ways to preempt the log traversal of an object. One is that we reach a nonterminal that already covers t_q while its MBR is disjoint with r ; this is discarded in lines 16-17. If this does not happen, the log entry is processed in line 18 and we consider the second condition: now, at this point, we can ensure that the object will not reach the region r . This is verified in line 19. Finally, line 20 includes the object in the result if we reach t_q and the p_c is in r . Lines 22-35 process the logs backwards from the following snapshot.

Assuming we first process O_1 , Step 1 of Table 2 shows the initial state of the algorithm. The algorithm traverses the trajectory of O_1 and checks if it can reach r at t_q after applying each m_i of $CT_{O_1}(t_9, t_{10}) = \langle 7, V \rangle$. The state after applying the first code (7) is shown in Step 2. The algorithm checks if the new position, (9, 6), lies inside the extended region corresponding to $t_c = t_9$. As shown in Figure 9(b), O_1 is outside the region, so the object is discarded.

The algorithm now processes O_2 . Its initial position, $p_c = (10, 3)$, and its time instant,

Algorithm 3: Time Slice (r, t_q)

```
1  $\mathcal{S}_h \leftarrow \text{nearestSnapshot}(t_q)$ ;  
2  $t_c \leftarrow \mathcal{S}_h.\text{time}$ ;  
3 if  $t_c = t_q$  then return  $\mathcal{S}_h.\text{findObjectsInRegion}(r)$  ;  
4  $\text{Candidates} \leftarrow \mathcal{S}_h.\text{findObjectsInRegion}(ER_r(t_c, t_q))$ ;  
5  $\text{Answer} \leftarrow \langle \rangle$ ;  
6 if  $t_c < t_q$  then  
7   foreach  $O_j \in \mathcal{S}_h.\text{App}$  do  
8      $t_c \leftarrow \mathcal{D}_{h,h+d}(O_j)[0]$ ;  
9      $p_c \leftarrow \mathcal{P}_{h,h+d}(O_j)[0]$ ;  
10    if  $t_c \leq t_q$  and  $p_c \in ER_r(t_c, t_q)$  then  
11       $\text{Candidates.add}(\langle O_j, t_c, p_c \rangle)$   
12    foreach  $\langle O_j, t_c, p_c \rangle \in \text{Candidates}$  do  
13       $T \leftarrow CT_{O_j}(t_{c+1}, t_q)$ ; // The trajectory, as a stack (first on top)  
14      while  $T$  is not empty do  
15         $m \leftarrow T.\text{pop}$ ;  
16        if  $m = NT_{\text{value}}$  and  $t_c + \text{value.}\#t \geq t_q$  and  $(p_c + \text{value.MBR}) \cap r = \emptyset$  then  
17          break // The nonterminal is disjoint with  $r$ , skip to next object  
18           $\langle t_c, p_c \rangle \leftarrow \text{moveJ}(p_c, t_c, t_q, m)$ ;  
19          if  $p_c \notin ER_r(t_c, t_q)$  then break; //  $O_j$  cannot reach  $r$ , skip to next object;  
20          if  $p_c \in r$  and  $t_c = t_q$  then Answer.add $(\langle O_j, p_c \rangle)$ ;  
21 else //  $t_c > t_q$   
22   foreach  $O_j \in \mathcal{S}_h.\text{Dis}$  do  
23      $t_c \leftarrow \mathcal{D}_{h-d,h}(O_j)[\text{last}]$ ;  
24      $p_c \leftarrow \mathcal{P}_{h-d,h}(O_j)[\text{last}]$ ;  
25     if  $t_c \geq t_q$  and  $p_c \in ER_r(t_q, t_c)$  then  
26        $\text{Candidates.add}(\langle O_j, t_c, p_c \rangle)$   
27     foreach  $\langle O_j, t_c, p_c \rangle \in \text{Candidates}$  do  
28        $T \leftarrow CT_{O_j}(t_q, t_{c-1})$ ; // The reversed trajectory, as a stack (last on top)  
29       while  $T$  is not empty do  
30          $m \leftarrow T.\text{pop}$ ;  
31         if  $m = NT_{\text{value}}$  and  $t_c - \text{value.}\#t \leq t_q$  and  $(p_c - \text{value.MBR}) \cap r = \emptyset$  then  
32           break // The nonterminal is disjoint with  $r$ , skip to next object  
33            $\langle t_c, p_c \rangle \leftarrow \text{RmoveJ}(p_c, t_c, t_q, m)$ ;  
34           if  $p_c \notin ER_r(t_q, t_c)$  then break; //  $O_j$  cannot reach  $r$ , skip to next object;  
35           if  $p_c \in r$  and  $t_c = t_q$  then Answer.add $(\langle O_j, p_c \rangle)$ ;  
36 return  $\text{Answer}$ ;
```

$t_c = t_8$, are obtained from the candidate list; this state is shown in Step 3. The state of Step 4 is obtained by applying the first code of $CT_{O_2}(t_9, t_{10}) = \langle 5, W \rangle$. Since $(9, 3)$ is within $ER_r(t_9, t_{10})$, as shown by Figure 9(b), the algorithm applies the second code of the trajectory: $m_i = W$ (a non-terminal symbol). The algorithm checks that the application of the non-terminal symbol reaches or exceeds t_q ($t_c + \text{value.}\#t \geq t_q$). In our example, $t_c = t_9$ and $W.\#t = 2$, so the rule would lead us to t_{11} . The MBR of W , however, is $[0, 1] \times [1, 2]$, which added to the current position $(9, 3)$ is $[9, 10] \times [4, 5]$. This is not disjoint with r , so O_2 cannot be discarded. The algorithm then has to decompress W , to determine if the object is within r after applying the part of W that spans exactly up to t_q . Since $\text{moveJ}((9, 3), t_9, t_{10}, W)$ yields that the position of O_2 at t_{10} is $(9, 4)$, as shown in Step 5, the object O_2 is reported.

Finally, Step 6 shows that we start processing O_5 at t_9 , where it appears. The case of this object is similar to O_2 , and it also becomes reported in Step 7.

	O_j	t_c	i	m_i	p_c	Answer
Step 1	O_1	t_8			(9,5)	
Step 2	O_1	t_9	0	7	(9,6)	
Step 3	O_2	t_8			(10,3)	
Step 4	O_2	t_9	0	5	(9,3)	
Step 5	O_2	t_{10}	1	W	(9,4)	O_2
Step 6	O_5	t_9			(7,2)	O_2
Step 7	O_5	t_{10}	0	W	(7,3)	O_2, O_5

Table 2: Trace of Algorithm 3 with the example of Figure 9.




Figure 9: Representation of a time-slice query (t_9 and t_{10} are not represented with snapshots).

An average-case analysis of the time complexity of this query can be made assuming that the objects distribute uniformly in the space and without taking advantage of the compression of the log. Let the query region r be of dimensions $\ell_x \times \ell_y$, and the query instant be at distance t from its closest snapshot. The spatial query on the k^2 -tree, for a range of size $(\ell_x + 2t) \times (\ell_y + 2t)$ and retrieving c candidates, takes $O(\ell_x + \ell_y + t + (c + 1)k \log_k s)$ time [30, Sec. 10.2.1]. On average, $c = \frac{(\ell_x + 2t)(\ell_y + 2t)}{\ell_x \ell_y} \cdot o$, where o is the actual output size. Further, we traverse all the objects that are within ranges of size $(\ell_x + 2t) \times (\ell_y + 2t)$ (at the snapshot time) decreasing by one until size $\ell_x \times \ell_y$ (t instants later or earlier). The sum of all those costs is, on average, $o \cdot \sum_{i=0}^t \frac{(\ell_x + 2i)(\ell_y + 2i)}{\ell_x \ell_y} = O(o \cdot (t + t^2(\ell_x + \ell_y + t)) / (\ell_x \ell_y))$. In total, the query has a fixed cost of $O(\ell_x + \ell_y + t + k \log_k s)$ plus an average cost of $O(k \log_k s + t + (t(\ell_x + \ell_y) \log_k s + t^2(\ell_x + \ell_y + t + k \log_k s)) / (\ell_x \ell_y))$ per object reported.

The analysis shows that, for this query, it is especially worthwhile to start from the closest snapshot, be it preceding or following t_q , so that t is on average $d/4$. In addition to having to process fewer log entries, the distance t between $\mathcal{S}_h.time$ and t_q directly impacts the size of the snapshot area from where we have to collect the candidate objects, and therefore the average number of candidates. The analysis also shows that the effect of t is reduced when the query area is larger.

6.4. Time interval query

Time-interval queries return the objects that were within r at any time instant during time interval $[t_b, t_e]$. An easy solution to this query is to use a time-slice algorithm

Algorithm 4: Time interval (r, t_b, t_e)

```
1  $\mathcal{S}_h \leftarrow \text{nearestPreviousSnapshot}(t_b)$ ;  
2  $\text{Answer} \leftarrow \langle \rangle$ ;  
3 while  $\mathcal{S}_{h+d}.\text{time} < t_e$  do  
4   |  $\text{ProcessPortion}(r, t_b, t_e, \mathcal{S}_h, \text{Answer})$ ;  
5   |  $\mathcal{S}_h \leftarrow \mathcal{S}_h.\text{nextSnapshot}()$ ;  
6 return  $\text{Answer}$ 
```

with two small modifications: (i) use the expanded region with respect to t_e , (ii) when the processed time instant (t_c) is within $[t_b, t_e]$ and the object is inside r , report and stop tracking it. Our actual solution speeds up this basic procedure by skipping whole nonterminals whenever possible.

We split the query interval $[t_b, t_e]$ into the log portions it spans, and treat each portion separately, accumulating the answers. Algorithm 4 calls *ProcessPortion* (Algorithm 5) for each portion of the log overlapped by $[t_b, t_e]$.

ProcessPortion is a modification of the time-slice algorithm. In lines 3-10, it collects the candidates to verify from the snapshot and the appearances that follow, avoiding the objects already reported. Lines 11-26 treat the trajectory of each candidate. A novelty with respect to time-slice queries is in lines 18-19: since we only report the objects and not their positions, we can directly report an object if the MBR of its next nonterminal is fully contained in the query region r . Otherwise, we must expand the nonterminal (lines 20-21) unless its MBR is disjoint with r , in which case we can process it as a whole, even if it exceeds t_{last} . Note that, since we manage the expansion of the nonterminals in line 21, line 23 always treats a complete symbol of the trajectory, knowing that it cannot exceed t_{last} (or that, if it exceeds t_{last} , it is because its MBR is disjoint with r). Finally, lines 24-26 process a single step, adding the objects found inside r and skipping them when they are shown to be unable to reach r within the time left until t_{last} .

We have omitted other possible speedups for simplicity. For example, let $r = [x_1, x_2] \times [y_1, y_2]$ and $pc + \text{value.MBR} = [x'_1, x'_2] \times [y'_1, y'_2]$. Then, in line 18, we can also directly add O_j to the answers if $t_c + \text{value.}\#t \leq t_e$ and, either $x_1 \leq x'_1 \leq x'_2 \leq x_2$ and $[y_1, y_2] \cap [y'_1, y'_2] \neq \emptyset$, or if $y_1 \leq y'_1 \leq y'_2 \leq y_2$ and $[x_1, x_2] \cap [x'_1, x'_2] \neq \emptyset$. In both cases, we know that O_j must be inside r at some moment not exceeding t_e .

The analysis for these queries is similar to that for time-slice queries, now taking t as the distance between t_e and the closest snapshot preceding t_b (and ignoring the help of intermediate snapshots). While omitted for simplicity, we can also process this query from the snapshot following t_e with the aim of reducing t .

6.5. Nearest neighbor queries

Nearest neighbor (*knn*) queries return the K objects nearest to a given query point p_q at a given time instant t_q . To solve this query, we obtain the candidate objects from the snapshot (\mathcal{S}_h) that is nearest to t_q , follow their trajectory up to t_q , and retain the K candidates that are closest to p_q at time t_q . The problem is that, unlike in time-slice queries, we cannot easily determine which objects cannot be part of the solution. This condition changes dynamically as we find better and better candidates for the answer set: given the K closest objects to p_q found up to now, the distance from p_q to the K -th closest object in t_q can be used to compute an upper bound to the distance between

Algorithm 5: ProcessPortion($r, t_b, t_e, \mathcal{S}_h, Answer$)

```

1  $t_c \leftarrow \mathcal{S}_h.time$ ;
2  $t_{last} \leftarrow \min(t_e, \mathcal{S}_{h+d}.time)$ ;
3  $Candidates \leftarrow \langle \rangle$ ; // The candidates to be considered, as a list
4 foreach  $\langle O_j, t_c, p_c \rangle \in \mathcal{S}_h.findObjectsInRegion(ER_r(t_c, t_{last}))$  do
5 |   if  $O_j \notin Answer$  then  $Candidates.add(\langle O_j, t_c, p_c \rangle)$ ;
6 foreach  $O_j \in \mathcal{S}_h.App$  do
7 |    $t_c \leftarrow \mathcal{D}_{h,h+d}(O_j)[0]$ ;
8 |    $p_c \leftarrow \mathcal{P}_{h,h+d}(O_j)[0]$ ;
9 |   if  $t_c \leq t_{last}$  and  $p_c \in ER_r(t_c, t_{last})$  and  $O_j \notin Answer$  then
10 | |    $Candidates.add(\langle O_j, t_c, p_c \rangle)$ ;
11 foreach  $\langle O_j, t_c, p_c \rangle \in Candidates$  do
12 |   if  $p_c \in r$  and  $t_c \in [t_b, t_e]$  then  $Answer.add(O_j)$ ;
13 |   else
14 | |    $T \leftarrow CT_{O_j}(t_{c+1}, t_{last})$ ; // The trajectory, as a stack (first on top)
15 | |   while  $t_c < t_{last}$  and  $T$  is not empty do
16 | | |    $m \leftarrow T.pop$ ;
17 | | |   if  $m = NT_{value}$  then
18 | | | |   if  $(p_c + value.MBR) \subseteq r$  then
19 | | | | |    $Answer.add(O_j)$ ; break; // Include object and skip
20 | | | |   else if  $(p_c + value.MBR) \cap r \neq \emptyset$  and  $value \rightarrow v_1 v_2$  then
21 | | | | |    $T.push(v_2)$ ;  $T.push(v_1)$ ; // Replace nonterminal
22 | | |   else
23 | | | |    $\langle t_c, p_c \rangle \leftarrow moveJ(p_c, t_c, \mathcal{S}_{h+d}.time, m)$ ;
24 | | | |   if  $p_c \in r$  and  $t_c \in [t_b, t_e]$  then
25 | | | | |    $Answer.add(O_j)$ ; break; // Include object and skip
26 | | | |   else if  $t_c < t_{last}$  and  $p_c \notin ER_r(t_c, t_{last})$  then break; // Skip candidate ;

```

a candidate and p_q in \mathcal{S}_h , so that if the candidate is farther away, it has no chance of getting closer than the current K -th closest object at t_q .

Since this upper bound becomes more restrictive as we improve the answer with closer objects, it is important to obtain objects close to p_q as soon as possible, as this allows us discarding other objects earlier. For this purpose, we will use the hierarchical space partitioning induced by the k^2 -tree of \mathcal{S}_h , and will traverse the regions prioritized by their potential relevance to the query (i.e., those closer to t_q will be processed first).

The traversal is carried out with a priority queue, which stores k^2 -tree regions r prioritized by increasing values of $d(r) = \overline{p_q r}$. Initially the whole grid, corresponding to the k^2 -tree root, is inserted in the priority queue, and then we iteratively extract the most promising node v (i.e., covering the region r with minimum $d(r)$ value). Let $\{r_0, r_1, \dots, r_{k^2-1}\}$ be the subregions of r induced by the children of v . The nonempty subregions r_i corresponding to children v_i of v are then reinserted in the queue with priority $d(r_i) = \overline{p_q r_i}$. When the children of the region r we extract are leaves of the k^2 -tree, the corresponding objects O_j are processed in a different way.

Those candidate objects O_j we find are inserted in a second priority queue, using a similar ordering on $\overline{p_q p_c}$, where p_c is the position of O_j ; we also record the time instant t_c of the snapshot \mathcal{S}_h associated with the objects. Initially, we accept in this second queue every object that does not disappear in its log by time instant t_q , until we have K elements. From this point, objects O_j are inserted only if at time t_q they can get closer to p_q than the K -th object closest to p_q in the second queue. More precisely, let O , which was known to be at position p in time instant t , be the K -th object closest to p_q in the

Algorithm 6: ObtainCandidates($\mathcal{S}_h, K, t_q, p_q$)

```

1  $v \leftarrow \mathcal{S}_h.root$ ; // The root of the  $k^2$ -tree of the grid,  $v.r$  is its region
2  $Q_R \leftarrow \emptyset$ ; // Priority queue of pairs (region,dist), sorted by increasing dist
3  $Q_C \leftarrow \emptyset$ ; // Priority queue of tuples (obj,time,pos,dist), sorted by increasing dist
4  $Q_R.add(\langle v, 0 \rangle)$ ; // Insert the whole region as the first pair
5  $d_{max} \leftarrow +\infty$ ; // Infinite until obtaining  $K$  candidates
6  $t_c \leftarrow \mathcal{S}_h.time$ ;
7 while  $Q_R$  is not empty do
8    $\langle v, dr \rangle \leftarrow Q_R.extractMin$ ; // Extract minimum distance from queue
9    $d_{min} \leftarrow dr - M_{sp} \cdot (t_q - t_c)$ ;
10  if  $d_{min} \geq d_{max}$  then break; // Best region cannot beat  $K$ -th candidate
11  if children of  $v$  are not leaves then
12    foreach  $v'$  nonempty child of  $v$  do
13       $Q_R.add(\langle v', p_q v'.r \rangle)$ ;
14  else
15    foreach  $O_j$  inside  $v.r$  do
16      if  $O_j$  is present in the log  $\mathcal{L}_{h,h+d}$  by time  $t_q$  then
17         $p_c \leftarrow \mathcal{S}_h.findObject(O_j)$ ; // Fast to compute from  $v$ 
18         $d_{min} \leftarrow \overline{p_q p_c} - M_{sp} \cdot (t_q - t_c)$ ;
19        if  $d_{min} < d_{max}$  then
20           $Q_C.add(\langle O_j, t_c, p_c, d_{min} \rangle)$ ;
21          if  $|Q_C| \geq K$  then
22            Let  $\langle O, t, p, d \rangle$  be the  $K$ -th smallest distance in  $Q_C$ ;
23             $d_{max} \leftarrow \overline{p_q p} + M_{sp} \cdot (t_q - t)$ 
24 foreach  $O_j \in \mathcal{S}_h.App$  do // Finally, add candidates that appear in the log
25   if  $O_j$  is present in the log  $\mathcal{L}_{h,h+d}$  by time  $t_q$  then
26      $p_c \leftarrow \mathcal{P}_{h,h+d}(O_j)[0]$ ;  $t_c \leftarrow \mathcal{D}_{h,h+d}(O_j)[0]$ ;
27      $d_{min} \leftarrow \overline{p_q p_c} - M_{sp} \cdot (t_q - t_c)$ ;
28     if  $d_{min} < d_{max}$  then
29        $Q_C.add(\langle O_j, t_c, p_c, d_{min} \rangle)$ ;
30       if  $|Q_C| \geq K$  then
31         Let  $\langle O, t, p, d \rangle$  be the  $K$ -th smallest distance in  $Q_C$ ;
32          $d_{max} \leftarrow \overline{p_q p} + M_{sp} \cdot (t_q - t)$ 
33 return  $Q_C$ 

```

priority queue (this is easily maintained by keeping the K closest objects in a separate queue prioritized by decreasing distance). Then the maximum distance to p_q it can reach at time t_q is $d_{max} = \overline{p_q p} + M_{sp} \cdot (t_q - t)$. Let O_j be at position p_c in time instant t_c . Then, at instant t_q , O_j cannot be closer to p_q than $d_{min}(O_j) = \overline{p_q p_c} - M_{sp} \cdot (t_q - t_c)$. Therefore, we insert O_j as a candidate only if $d_{min}(O_j) < d_{max}$. When we populate the queue of candidates, we must also consider the objects O_j in $\mathcal{S}_h.App$ appearing at times $t_c \leq t_q$ (and not disappearing before t_q). We prioritize the candidates O_j using $d_{min}(O_j)$.

This condition to insert an object as a candidate gives also a criterion to stop the processing of the priority queue of regions. If we extract region r from the queue, no object from r can be, at time instant t_q , at a distance to p_q below $d_{min}(r) = \overline{p_q r} - M_{sp} \cdot (t_q - t_c)$. Therefore, if $d_{min}(r) \geq d_{max}$, there is no need to consider any point from region r , nor from any region yet to be extracted from the first priority queue. We can, at that point, discard the first queue and start processing our queue of candidates. Algorithm 6 gives the code (for simplicity, we consider only the nearest snapshot preceding t_q).

Let us assume that $k = 2$, $M_{sp} = 1$, $K = 1$, $t_c = t_8$, $t_q = t_9$, $p_q = (10, 0)$, and the grid \mathcal{S}_8 of Figure 10. We number the regions with a sequence of subindices in $\{0, 1, 2, 3\}$

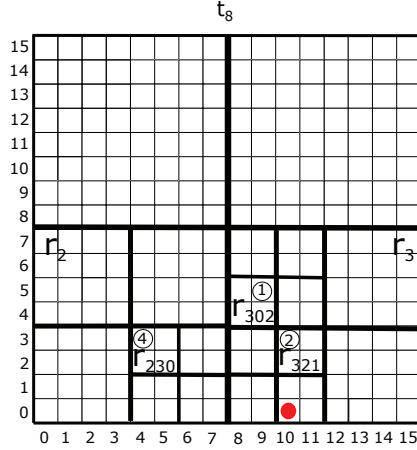


Figure 10: Example of regions of the space inserted in the queue.

considering the consecutive levels. The children of the root at r_0 to r_3 , the children of r_0 are r_{00} to r_{03} , and so on. From the root, we insert r_2 and r_3 in the queue (r_0 and r_1 are empty). We then extract r_3 , the closest one to p_q (indeed, containing it), and reinsert r_{30} and r_{32} . The next most promising region is r_{32} , from whose children we only reinsert r_{321} , the nonempty one. From the current set, $\{r_2, r_{30}, r_{321}\}$, the most promising one is r_{321} , at distance $\overline{p_q r_{321}} = 2$. So we extract it and insert object O_2 as our first candidate, at distance $\overline{p_q O_2} = 3$. Thus, at time t_9 , O_2 could be at maximum distance $\overline{p_q O_2} = 4$. Since the next most promising area, r_2 , is at distance $\overline{p_q r_2} = 3$, it is still worth considering, so we reinsert its nonempty child r_{23} . This is still more promising than r_{30} , so we extract it and reinsert r_{230} , its nonempty child. Now our candidate regions are $\{r_{30}, r_{230}\}$, and the most promising is r_{30} . We extract it and reinsert r_{302} , its nonempty child, which becomes the first region to extract. Since $\overline{p_q r_{302}} = \sqrt{17} \approx 4.1$, a point inside r_{302} might reach distance ≈ 3.1 to p_q at t_9 , and thus it might still be interesting. We then extract it and obtain O_1 , which is at distance $\overline{p_q O_1} = \sqrt{26} \approx 5.1$, and then can get as close as distance ≈ 4.1 in t_9 . There is no point in considering O_1 , then, because it cannot get closer to p_q than O_2 . Similarly, there is no point in considering the remaining region r_{230} (nor any other region that would have remained to process), because it is at distance $\overline{p_q r_{230}} = \sqrt{29} \approx 5.4$. In this case, we only have one candidate.

From the priority queue of candidates, the algorithm extracts the first one (with position p_c , at time t_c , with minimum d_{min}) and calculates the next position of the object using *moveJ*. The object is then reinserted in the queue with its new position p'_c and time instant t'_c , and with d_{min} recomputed accordingly. As we iterate, the candidates advance in their log $\mathcal{L}_{h,h+d}$, with the most promising objects advancing sooner towards t_q . At some point, we will start extracting objects with time $t_c = t_q$. For those objects, d_{min} is their exact distance to p_q at time t_q . We store them in a third priority queue, of maximum size K and sorted by decreasing d_{min} . When we have K answers in this third queue, the insertion of a new result may displace the currently K -th result (the first in the priority queue). Further, we can use the value $d_{min}(O)$ of the K -th result O to preempt the processing of the queue of candidates: if for the candidate O_j with the

Algorithm 7: Knn(K, p_q, t_q)

```

1  $\mathcal{S}_h \leftarrow \text{findNearestPreviousSnapshot}(t_q)$ ;
2  $Q_C \leftarrow \text{ObtainCandidates}(\mathcal{S}_h, K, t_q, p_q)$ ;
3  $Q_F \leftarrow \emptyset$ ; // Priority queue of pairs (obj,dist) sorted by decreasing dist, max size  $K$ 
4  $d_{max} \leftarrow +\infty$ ; // Infinite until obtaining  $K$  results
5 if  $\mathcal{S}_h.time = t_q$  then return first  $K$  objects of  $Q_C$  ;
6 while  $Q_C$  is not empty do
7    $\langle O_j, t_c, p_c, d_{min} \rangle \leftarrow Q_C.extractMin$ ;
8   if  $|Q_F| \geq K$  and  $d_{min} \geq d_{max}$  then break; // Best candidate cannot beat  $K$ -th result
9   if  $t_c = t_q$  then // We have a result
10     $Q_F.add(\langle O_j, d_{min} \rangle)$ ; //  $d_{min}$  is  $\overline{p_q p_c}$ 
11    if  $|Q_F| > K$  then  $Q_F.extractMax$ ; // maximum size is  $K$ , discard one;
12    if  $|Q_F| \geq K$  then  $\langle O, d_{max} \rangle \leftarrow Q_F.inspectMax$ ; // update  $d_{max}$ ,  $O$  is dummy;
13  else // Advance one position in the log; we have a pointer to  $CT_{O_j}$ 
14     $\langle t_c, p_c \rangle \leftarrow \text{moveJ}(p_c, t_c, t_q, CT_{O_j}(t_{c+1}, t_q)[0])$ ;
15     $d_{min} \leftarrow \overline{p_q p_c} - M_{sp} \cdot (t_q - t_c)$ ;
16     $Q_C.add(\langle O_j, t_c, p_c, d_{min} \rangle)$ ;
17 return  $Q_F$ 

```

minimum d_{min} value it holds that $d_{min}(O_j) \geq d_{min}(O)$, then there is no chance that O_j (nor any other candidate object remaining in the second priority queue) alters the current set of results. We can then stop the processing and return the results in the third queue as the definitive answer.

Algorithm 7 shows the pseudocode of the *knn* query. To illustrate it, the example from Figures 5 and 7 is used to obtain the nearest object ($K = 1$) with respect to location $p_q = (10, 0)$ at $t_q = t_9$. After obtaining the previous closest snapshot to t_9 (\mathcal{S}_8), Algorithm 7 issues the call $\text{ObtainCandidates}(\mathcal{S}_8, 1, t_9)$ to obtain the candidates from the snapshot. As we have already seen, the traversal of regions yields only one candidate, O_2 , with minimum distance $d_{min} = \overline{p_q O_2} - M_{sp} \cdot (t_q - t_c) = 2$. Therefore, the queue of candidates contains only $Q_C = \{\langle O_2, t_8, (10, 3), 2 \rangle\}$. However, objects O_3, O_5, O_6 , and O_7 appear in the log after \mathcal{S}_8 . From those, only O_5 appears before or at t_9 , and thus it can be added to Q_F . Indeed, O_5 appears at t_9 , in position $(7, 2)$. Its minimum (and exact) distance to p_q in t_q is $d_{min} = \overline{p_q(7, 2)} = \sqrt{13} \approx 3.6$. Since $d_{max} = 4$, this object cannot yet be discarded, so the final queue of candidates is $Q_C = \{\langle O_2, t_8, (10, 3), 2 \rangle, \langle O_5, t_9, (7, 2), 3.6 \rangle\}$.

Once Algorithm 7 has all the candidates, it starts processing them in lines 6-16. It first extracts $\langle O_2, t_8, (10, 3), 2 \rangle$. Since t_8 is not the target time, the algorithm considers $CT_{O_2}(t_9, t_9) = \langle 5 \rangle$ and applies moveJ on the first element, 5. This converts t_8 into t_9 and $(10, 3)$ into $(9, 3)$. It then recomputes $d_{min} = \sqrt{10} \approx 3.2$ and reinserts O_2 into Q_C with these new values. The queue of candidates now contains $Q_C = \{\langle O_2, t_9, (9, 3), 3.2 \rangle, \langle O_5, t_9, (7, 2), 3.6 \rangle\}$. The next best candidate is again O_2 , but now its time is t_9 , the same of the query, and this means that we have our first result, $Q_F = \{\langle O_2, 3.2 \rangle\}$. Since now $|Q_F| \geq K$, we can compute $d_{max} = 3.2$. In the next step, when we extract O_5 from Q_C , we have that its d_{min} value, 3.6, is larger than d_{max} , and thus disregard it. We have exhausted Q_C and the answer is in Q_F .

7. Experimental Evaluation

In this section we experimentally evaluate the space and time performance of GraCT and compare it with related work. It has been very hard to obtain implementations of previous indexes. Further, compression-only systems developed for trajectories like Trajic [32] and SQUISH [26] did not compress well on our data, reducing the space of a plain text representation but not the space of a binary representation of the integer coordinates. We have, nevertheless, compared GraCT with two baselines:

scdcCT: This is a structure similar to GraCT with a different representation of the log structure. The sequence of spiral codes that differentially encode the movements is compressed with a statistical zero-order compressor, namely byte-oriented (s, c)-Dense Codes [6] (SCDC) optimized to $s = 183$ and $c = 73$. SCDC combines good compression ratios with fast decoding, thereby supporting fast scanning of the logs. We use scdcCT as a representative of the various systems that use statistical compression, instead of the grammar-based compression used by GraCT. We aim to demonstrate that grammar compression is more effective, and that the ability of GraCT to skip over whole nonterminal symbols makes it faster than statistical encoders that, no matter how fast, must process each movement individually. Note that including sampled absolute values at regular intervals to speed up the processing of logs in scdcCT is equivalent in space to shortening the sampling period d of the snapshots.

MVR-tree: This is the only classic spatio-temporal structure for which we could obtain code. The complete structure is the MV3R-tree [43], which combines an MVR-tree with an auxiliary 3DR-tree, but only the MVR-tree part was available. The 3DR-tree is used to speed up the most basic queries, like obtaining the position of an object at a given time instant or retrieve its trajectory. We aim to show that a classic structure uses orders of magnitude more space than GraCT, without being much faster (actually, the MVR-tree is slower for some queries). We used the MVR-tree implementation provided by the `spatialindex` library¹. This is a disk-based structure, but we make it run in main memory for a direct comparison with GraCT. We also compare both structures on disk, as explained later.

We implemented GraCT and scdcCT in C++, using components from the SDSL library² [12]. The experiments were run on an Intel[®] Core[™] i7-3820 CPU @ 3.60GHz (4 cores) with 10MB of cache and 64GB of RAM, running Ubuntu 12.04.5 LTS with kernel 3.2.0-115 (64 bits), gcc version 4.6.4 with `-O9` optimization. The code for GraCT and scdcCT is publicly available at <https://gitlab.lbd.org.es/adriangbrandon/ct.git>.

We used an implementation for supporting rank operations that adds 6.25% space on top of the bit sequence to accelerate queries. The sample value $1/\epsilon$ of the permutations was set to 5, which increases the size of the data structure by 20% but yields fast π^{-1} queries. Our grammar compressor was the Re-Pair implementation provided by G. Navarro³, using the balanced version. The extra information associated with the

¹<http://libspatialindex.github.io>

²<https://github.com/simongog/sdsl-lite>

³<http://www.dcc.uchile.cl/gnavarro/software/repair.tgz>

Re-Pair nonterminals was encoded using DACs with an unlimited number of levels and without a predefined chunk size. Arrays \mathcal{P} and \mathcal{D} , in contrast, are represented using DACs configured to use chunks of one byte and with a maximum of two levels.

7.1. Datasets

We used real-world and pseudo-real data from three sources.

- **Ships**, a real dataset storing the movements of 3,654 boats sailing in UTM Zone 10 over one month in 2014, obtained from MarineCadastre⁴.
- **Planes**, a real dataset composed by trajectories of aircrafts from 30 different airlines and flights between 30 European airports. This data was obtained from the OpenSky Network⁵; we discarded the altitude component and retained latitude and longitude.
- **Taxis**, a larger synthetic dataset simulating trajectories of taxis in New York City during 2013. We download the data from *NYC Taxis: A Day in the life*⁶. Since this includes only the start and end location of each taxi trip, we computed the trajectory as the shortest route between these two points.

GraCT stores positions in a discrete grid, therefore each position emitted by an object is discretized into a matrix. We chose cell sizes that seem appropriate in each application, in the sense that more detail is probably useless: 50×50 meters in **Ships**, 5000×5000 meters in **Planes** and 10×10 meters in **Taxis**. In addition, our structure deals with object positions at regular timestamps, but in the real datasets, the frequency of signals varied. For this reason, the signals were preprocessed and discretized, in order to obtain the position of each object at regular timestamps: 1 minute for **Ships** and 15 seconds for **Planes** and **Taxis**.

After this normalization step, some location signals were observed to be incorrect, because they implied object movements at an extremely high speed (this kind of errors is not rare in GPS measurements). We set a maximum speed for each dataset, 800km/h for **Planes** and 234km/h for the other datasets. If the signal of an object at a timestamp implies exceeding the maximum speed, that signal is deleted. This amounts to deleting 0.01% of the signals from **Ships** and 0.83% from **Planes**. In addition, the frequency of signals is variable. For example, in **Ships** the frequency of signals from most boats was regular when they were sailing, but it stopped for long periods, typically when they were at a port. This produces disappearances and appearances, which amount to 0.21% of the entries in **Ships** and 0.96% in **Planes**. Finally, since our setup requires the position of each object at regular timestamps, we interpolated the two closest signals if they were less than 15 time instants apart, since an object cannot move far during a short period of time; otherwise we consider it a disappearance/appearance.

Each dataset is stored in a plain text file comprising four columns: *object identifier*, *time instant*, *coordinate x* and *coordinate y*. Each value in the columns is stored as a

⁴<http://marinecadastre.gov/ais>

⁵<https://opensky-network.org>

⁶<http://chriswhong.github.io/nyctaxi>

	Ships	Planes	Taxis
Total objects	3,654	2,263	1,283
Total points	44,304,802	35,315,697	2,524,018,668
Max x	12,719	3,313	1,090,125
Max y	368,186	1,148	1,160,194
Max $time$	44,642	172,547	2,102,879
Size Plain	947.35 MB	671.02 MB	59,263.57 MB
Size Bin	380.27 MB	303.12 MB	26,478.01 MB
Size <i>p7zip</i>	36.14 MB	35.82 MB	3,075.93 MB

Table 3: Datasets and their dimensions.

string. In order to obtain a better compression comparison, the data was alternatively stored in binary form, with the minimum number of bytes. For example, in **Ships**, two bytes were used to represent object identifiers (max value 3,654), two for the time instant column (max value 44,642), two for the x-axis (max value 12,719) and three for the y-axis (max value 368,186). Table 3 shows the resulting sizes.

To illustrate how compressible the data is, the plain data files were compressed with *p7zip*⁷ with default settings. As it can be seen, *p7zip* compresses the data to 9.5%–11.8% of its binary representation.

7.2. Compression

The first experiment compares the compression ratio of GraCT with scdcCT on the two real datasets, **Ships** and **Planes**. GraCT and scdcCT data structures were built on both datasets using snapshot intervals of $d = 120, 240, 360$ and 720 time instants. The construction time for the complete structure was approximately one minute.

Table 4 shows the results of compression. The first row gives the size of both structures for each snapshot period d , whereas the next two rows separate the space due to the snapshots (i.e., the k^2 -trees) and to the compressed logs. It can be seen that the space of the snapshots is much lower than that of the logs, even with the highest sampling rate. Still, the compression of the logs themselves improves for larger d . In this aspect, Re-Pair improves faster than SCDC: the quotient of the log space with $d = 720$ versus $d = 120$ is 0.68 on **Ships** and 0.47 on **Planes**, whereas with scdcCT they are 0.78 and 0.60, respectively.

The last two rows give the compression ratios obtained with respect to the plain and binary representations, respectively. The compression achieved by scdcCT with $d = 720$, 13%–15% of the binary representations, shows that statistical compression of the differences is competitive with state-of-the-art compression like the one offered by *p7zip*. The grammar-compression used by GraCT, on the other hand, obtains less than half the space of scdcCT, and less than 60% of the space used by *p7zip*.

We remark that, within this space, GraCT not only represents the trajectory data, but it offers efficient access and queries on it. Devoided of the extra information it stores on the nonterminals to speed up queries, the space used by the grammar compression of

⁷<http://p7zip.sourceforge.net>

Index		GraCT				scdcCT			
Period		120	240	360	720	120	240	360	720
Ships	Size	34.72	26.92	24.27	21.58	65.93	56.02	52.65	49.28
	Snapshot	4.05	2.13	1.43	0.74	4.05	2.13	1.43	0.74
	Log	30.66	24.79	22.84	20.85	61.88	53.89	51.22	48.55
	Ratio (plain)	3.66%	2.84%	2.56%	2.28%	6.96%	5.91%	5.56%	5.20%
	Ratio (bin)	9.13%	7.08%	6.38%	5.68%	17.34%	14.73%	13.85%	12.96%
Planes	Size	49.24	32.23	26.55	20.83	82.61	59.80	52.21	44.56
	Snapshot	3.85	1.95	1.31	0.68	3.85	1.95	1.31	0.68
	Log	41.06	28.08	23.75	19.38	74.43	55.66	49.41	44.56
	Ratio (plain)	7.34%	4.80%	3.96%	3.10%	12.31%	8.91%	7.78%	6.64%
	Ratio (bin)	16.25%	10.63%	8.76%	6.87%	27.25%	19.73%	17.23%	14.70%

Table 4: Structure sizes (in MB) and compression ratios.

the logs would halve or so. This shows that grammar compression is much more efficient than statistical compression at exploiting the redundancy of real-world trajectory data.

7.3. Query performance

This experiment studies the response times of GraCT and scdcCT, with different snapshot periods, on the real datasets. Figure 11 shows the average time (in milliseconds) per query for the following seven types of queries:

- *object*: this query obtains the position of a given object at a given time instant t_q . We averaged over 20,000 queries for random objects and time instants.
- *trajectory*: this query returns the trajectory followed by an object during an interval $[t_b, t_e]$, where $t_e - t_b$ is fixed at 2,000 time instants. We averaged over 10,000 queries for random objects and time instants t_b .
- *time-slice S*: this query obtains the identifiers and positions of the objects lying within a small region (40×40 cells) at a given time instant. We averaged over 1,000 queries for random region positions and time instants.
- *time-slice L*: this query obtains the identifiers and positions of the objects lying within a larger region (320×320 cells). We averaged over 1,000 queries for random region positions and time instants.
- *time-interval S*: this query obtains the objects present in a small region (40×40 cells) at any time instant between t_b and t_e , where the interval size is $t_e - t_b = 100$ instants. We averaged over 1,000 queries with random region positions and instants t_b .
- *time-interval L*: this query obtains the objects present in a larger region (320×320 cells) over a longer interval of size $t_e - t_b = 500$ instants. We averaged over 1,000 queries with random region positions and instants t_b .

- *knn*: this query obtains the K nearest neighbors to a given position at a given time instant, where K is a random value between 1 and 50. We averaged over 1,000 queries with random objects and time instants.

7.3.1. Object

As seen in Figure 11(a), GraCT outperforms scdcCT in this query, answering within 2–4 microseconds. The logs in GraCT have far fewer entries than those in scdcCT, and in many cases the extra information stored on nonterminals allows processing each GraCT log entry in constant time, just as a simple scdcCT log entry. Further, as shown before, Re-Pair compression improves faster as the logs increase, which makes GraCT times less sensitive than scdcCT to the increased distance between snapshots. The improvements of GraCT over scdcCT range from 5.9% with a snapshot period of $d = 120$ to 46.3% with $d = 720$ on **Ships**, and from 12.3% to 100.4% on **Planes**.

Figure 12 shows the relative time spent on the k^2 -tree of the snapshot versus on the logs for the different queries. Since the times spent on the snapshots are identical, it follows that GraCT processes the logs way faster than scdcCT. This demonstrates that, provided with the proper enhanced data, the use of grammar compression also yields good query times, outperforming the statistical encoding of trajectories because it processes fewer entries at about the same speed per entry.

7.3.2. Trajectory

In trajectory queries, as shown in Figure 11(b), GraCT also outperforms scdcCT, but the difference in response times is small: 12%–20% on **Ships** and 8%–13% on **Planes**. The majority of the time in this query is spent in decompressing the trajectory (of 2,000 entries), which is done at about the same rate in statistical (scdcCT) and grammar-based (GraCT) compression: 25–45 nanoseconds per decompressed datum on **Ships** and 15–20 on **Planes**. Since the trajectories to recover are significantly longer than the snapshot periods, snapshots mostly disrupt the processing of the logs, which explains why times improve with fewer snapshots.

Figure 12 shows the case of shorter trajectories, of lengths 100 (**Trajectory 100**) and 1000 (**Trajectory 1000**), on **Ships**. As expected, more time is spent on the logs to retrieve longer trajectories or with longer sampling periods.

7.3.3. Time-slice S and time-slice L

In time-slice queries, the difference in performance between GraCT and scdcCT becomes much more significant. As shown in Figures 11(c) and 11(d), GraCT performs 2–4 times faster than scdcCT in both datasets, running each query in less than a millisecond.

This query can be regarded as running an *object* query on each of the candidate objects, but as shown in Figure 12, the algorithm spends more time working on the log than for *object* queries, because the k^2 -tree amortizes its traversal cost when retrieving many points. Thus both indexes spend more time in the traversal of the log, where GraCT outperforms scdcCT, and less on the k^2 -tree, where they perform similarly. As a result, GraCT can be up to 4 times faster than scdcCT on *time-slice* queries, whereas in *object* queries it was up to twice as fast.

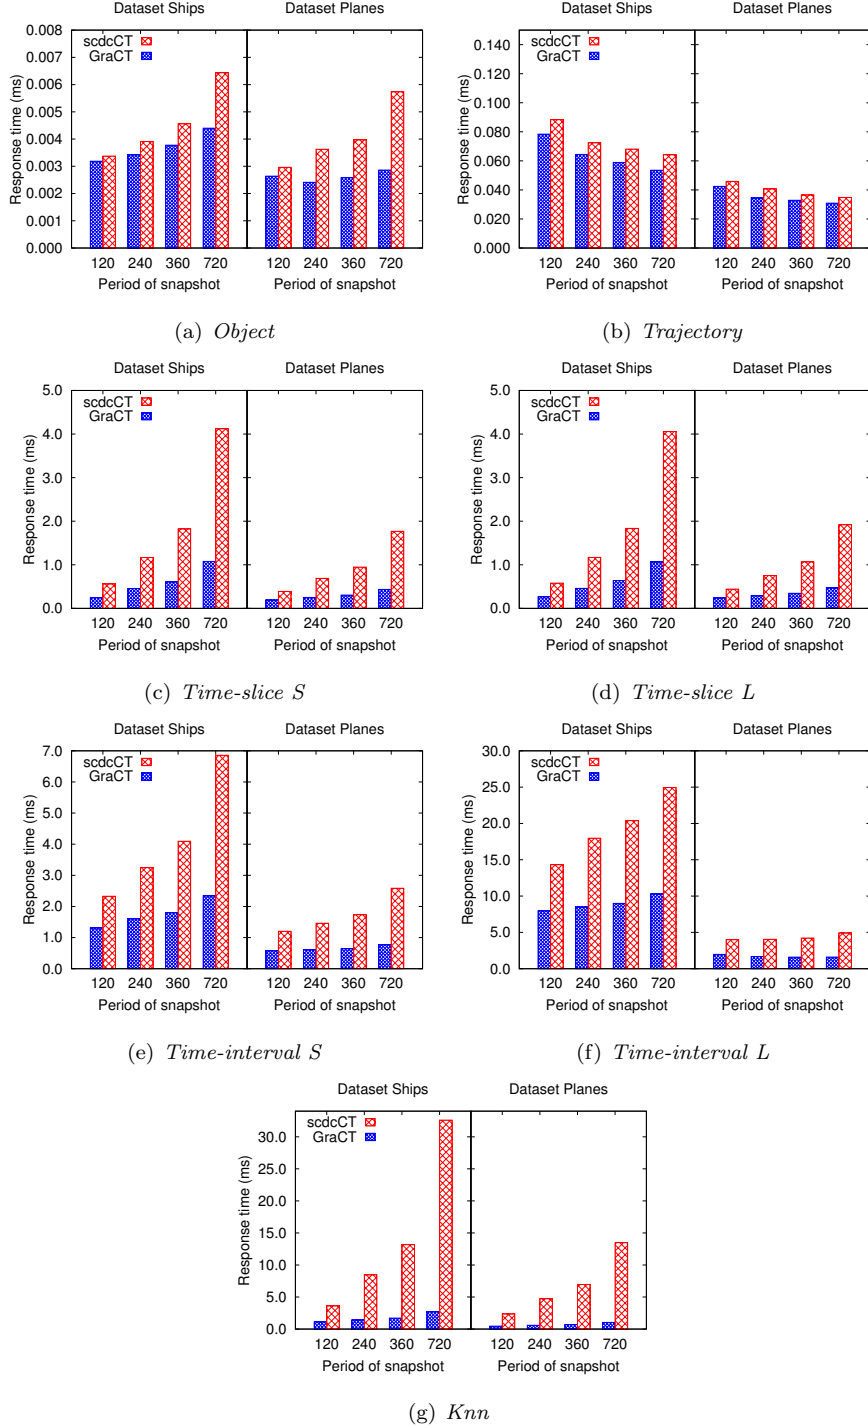


Figure 11: Response time for different queries, in milliseconds per query.




Figure 12: Time spent on the snapshot versus the logs for different queries.

Period	120		240		360		720	
	MBR	no MBR	MBR	no MBR	MBR	no MBR	MBR	no MBR
<i>Slice S</i>	0.25	0.25	0.46	0.47	0.61	0.65	1.08	1.13
<i>Slice L</i>	0.27	0.27	0.46	0.47	0.64	0.67	1.06	1.11
<i>Interval S</i>	1.31	7.40	1.60	12.66	1.80	17.39	2.34	29.68
<i>Interval L</i>	7.96	39.23	8.50	50.84	9.00	59.43	10.32	76.67

Table 5: The impact of using MBRs to preempt time-slice and time-interval queries, in milliseconds.

7.3.4. Time-interval S and time-interval L

As observed in Figures 11(e) and 11(f), time-interval queries are very similar to time-slice queries with respect to the comparison between GraCT and scdcCT, although this time GraCT is only up to 3 times faster than scdcCT. On longer time intervals, the difference is smaller because both structures may have to decompress the movements along the interval $[t_b, t_e]$. The difference is larger than on a *trajectory* query, however, because GraCT can avoid decompressing nonterminals if their MBRs are disjoint or contained in the query range.

Table 5 evaluates the improvement we obtain due to the use of MBRs of nonterminals on *Ships*. As it can be seen, the use of MBRs has a relatively small impact on time-slice queries (below 5%), but a very significant one on time-interval queries: increasing with the sampling period, the speedup factor goes from 5.6 to 12.7 on *time-interval S* and from 4.9 to 7.4 on *time-interval L* queries. Without using MBRs, indeed, GraCT would be slower than scdcCT on time-interval queries.

7.3.5. Knn

Nearest neighbor queries also require to process the logs of many candidate objects. However, the process done on each log entry is significantly heavier than on previous queries: objects are extracted from and reinserted in priority queues, their d_{min} distances are recalculated, etc. In this scenario, the fewer number of entries processed by the grammar-compressed logs of GraCT, which can skip whole nonterminals, is even more valuable. As a result, GraCT is up to an order of magnitude faster than scdcCT for




Figure 13: Evolution of compression ratio and query times as the dataset grows. Dataset sizes are shown in logscale.

these queries, answering them in less than 3 milliseconds.

Figure 12 shows that the prioritized traversal of the k^2 -tree consumes a significant part of the query time on GraCT, but on scdcCT it is counterweighed by the more expensive processing of the log.

7.4. Scalability

In order to evaluate the scalability of our structure in compression and query times, we create five prefixes of **Taxis** with sizes 500 MB, 1,024 MB, 5,120 MB, 10,240 MB, and 61,240 MB. In each longer prefix, more objects are incorporated to the dataset. We built GraCT and scdcCT on each dataset, with the period of snapshots set to 720. Over the GraCT indexes, we ran queries with the same settings of Section 7.3.

As shown in Figure 13(a), the compression ratio of GraCT improves as the dataset grows, from 9% to 4%. This is a consequence of exploiting repetitiveness via grammar compression. As the number of objects and trajectories of the same kind increases, finding similar trajectories is more likely. Instead, as it can be seen, statistical compression is much less sensitive to repetitiveness, improving only from 12% to 10%. This shows another advantage of grammar compression when handling very large datasets.

With the growth of the data, on the other hand, the depth of the grammar used by GraCT also grows, logarithmically in case of being a balanced grammar. This increases the cost of navigating inside a nonterminal, for example. This logarithmic growth is apparent in the *object* queries (Figure 13(b); since the x axis is logarithmic, a logarithm looks as a straight line), though in *trajectory* queries (Figure 13(c)) the time is amortized over the 2,000 positions output (those are decompressed in constant time each from the grammar, with only a logarithmic additive penalty for the whole query). The other queries grow linearly because, as more objects are added on a fixed grid, the output size and/or number of candidates to consider grows proportionally to the data size. This is apparent in Figure 13(d), where the query times divided by dataset size are essentially constant.

7.5. Comparison with a classical uncompressed index

As anticipated, we compare GraCT with the MVR-tree [43], a classic spatio-temporal structure based on the R-tree. The MVR-tree comprises multiple versions of the same data, each associated with a different interval of time, and an R-tree representing the positions of the objects during that interval. The intervals of the versions are disjoint. The MVR-tree exploits the repetition of nodes along consecutive R-trees, by sharing common nodes between versions. It is important to note that this spatio-temporal structure is efficient at solving time-slice, time-interval, and knn queries only, which are solved by identifying versions that overlap with the queried time interval. The algorithm traverses the R-tree of each such version, following the nodes involved in the spatial query range. Queries *object* and *trajectory*, instead, are costly for the MVR-tree. These are speeded up in the full structure, the MV3R-tree, by using an auxiliary 3DR-tree structure [43] (which requires even more space). We therefore omit these queries in our comparison.

The MVR-tree was configured to run in main memory. Both structures were built on the **Ships** and **Planes** datasets described in Section 7.1, and GraCT was configured with the same snapshot distances used in the previous experiments (120, 240, 360 and 720). The size of the MVR-tree in **Ships** was 12.16 GB and in **Planes** it was 11.72 GB, while the maximum-space configuration of GraCT used 36.40MB and 51.63MB, respectively, two orders of magnitude less (more precisely, GraCT uses 342 times less space on **Ships** and 232 times less space on **Planes**). The MVR-tree is also 30–40 times larger than the raw data represented in binary form.

We now study if this gigantic difference in space induces a significant impact on query time performance. Figure 14 shows the times of both structures averaged over 1,000 random queries of each type: *time-slice S*, *time-slice L*, *time-interval S*, *time-interval L*, and *knn* queries. Instead of the sampling $d = 720$, we include a new one, $d = 60$, where GraCT uses 50.70 MB on **Ships** and 83.14 MB on **Planes**. This is still 245 and 144 times smaller than the MVR-tree, respectively, and is sufficient for GraCT to be as fast as MVR-tree on most queries.

Time-slice queries are the most efficient ones for the MVR-tree. Since they retrieve objects within a given region at a single time instant only, the MVR-tree needs to traverse just one R-tree. When the traversal of the tree reaches the leaves, it can retrieve more than one object per leaf, which is faster than traversing one log of GraCT per candidate. Figures 14(a) and 14(b) show that, even in this case, the maximum-space configuration of GraCT matches the performance of the MVR-tree on **Ships**, and outperforms in on

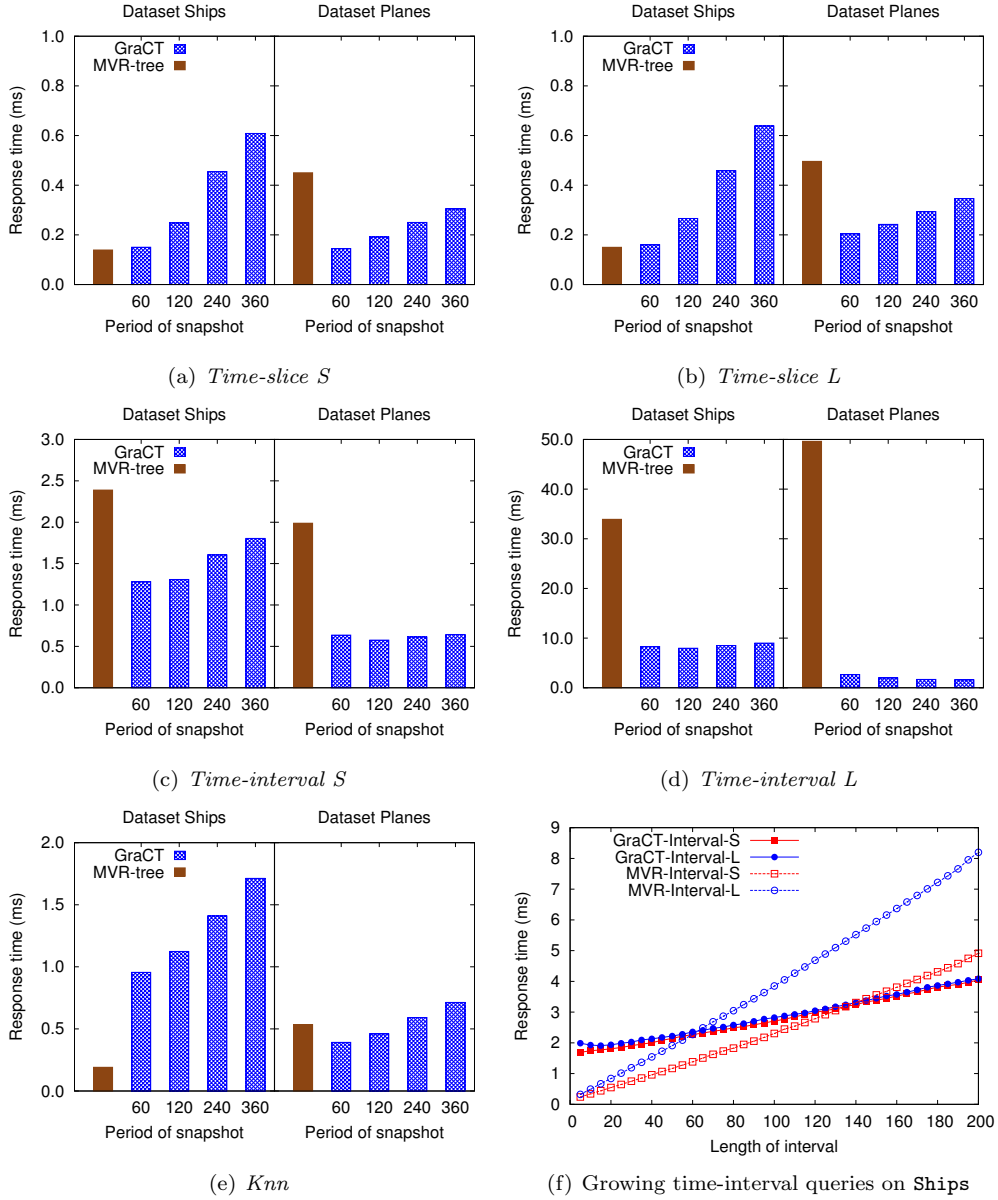


Figure 14: Query time comparison of GraCT with the MVR-tree, running in main memory.

Planes. It is likely that the higher locality of reference when accessing a structure that is two orders of magnitude smaller plays in favor of GraCT.

This advantage of the MVR-tree vanishes on time-interval queries. In an interval of time, the MVR-tree may have to traverse a number of R-trees, which makes it much slower and always loses to GraCT, as seen in Figures 14(c) and 14(d). Figure 14(f)

studies the turning point on **Ships**, by varying the length of the queried time interval, using the smallest space configuration of GraCT (snapshot distance 720). As the time span of the queries increases, both structures slow down. The MVR-tree times, however, increase much faster, becoming slower than GraCT when the time interval exceeds 140 time instants on *time-interval S* and 60 on *interval L*. GraCT is thus less sensitive to the length of the queried interval, outperforming MVR-tree already in short time intervals even on its minimum-space configuration.

Nearest neighbor queries refer again to a single time instant, and therefore the MVR-tree can handle them within a single R-tree, even if it usually has to traverse many more nodes than for a time-slice query. In this query, GraCT also has to traverse a large number of k^2 -tree nodes and verify many candidates because, unlike the MVR-tree, it does not have a snapshot of the precise query time instant. As it can be seen in Figure 14(e), the MVR-tree clearly outperforms GraCT on **Ships**, though not on **Planes**, where GraCT always improves and the MVR-tree always degrades with respect to **Ships**.

7.6. Disk and full precision

With the minimum-space configuration (snapshot sampling of 720), GraCT uses nearly 600 times less space than the MVR-tree. We can then use GraCT to index in main memory datasets for which the MVR-tree would exceed by far the available memory, and would then have to operate on disk. As we show next, even a time-slice query (where the MVR-tree most prominently excels) requires around 50 milliseconds to be solved by the MVR-tree on disk, which is 50 times slower than the time required by GraCT in main memory. If we use a snapshot sampling of 60, then GraCT becomes more than 300 times faster, and is still 150–250 times smaller than the MVR-tree.

The compression of GraCT is achieved, however, by means of discretizing the data and therefore losing precision. In most real systems, full precision is anyway unnecessary. For example, using eight decimals for the GPS position of an object discriminates micrometres, which is much finer than what most physical devices can acquire, and only useful in very specialized fields like tectonic plate mapping.

Still, if full precision is required, GraCT can be used as an in-memory *cached index* of the full-precision data stored in disk, exploiting its very small main memory footprint. To answer a query, GraCT is first used to obtain a list of *candidate* answers, which satisfy the query up to where the precision of GraCT can discriminate. Then, in a second step, those candidate answers are verified by reading the full-precision data from disk.

As a proof-of-concept of the feasibility of this approach, we designed a simple arrangement for time-slice queries, where the MVR-tree performs best. We use the GraCT data structures of the **Ships** dataset with 720 time instants between snapshots. In addition, we store the full-precision data in a relational table of PostgreSQL DBMS v9.1.24.⁸ The table stores tuples $\langle object, time, x, y \rangle$ to represent all the trajectories of all the objects. To efficiently verify the answers to time-slice queries, a hash index is set on the first two columns (together). For simplicity, we still assume that the time instants are discretized; a B-tree can be used for the second column in case they are real numbers as well.

As a baseline, we use the MVR-tree data structure on disk. GraCT, including the full-resolution data and the hash indexes, occupies 4.18 GB, whereas the MVR-tree occupies 78.61 GB. That is, GraCT still requires 18 times less disk space than the MVR-tree.

⁸<https://www.postgresql.org>

	Slice S		Slice L	
	GraCT+DBMS	MVR-tree	GraCT+DBMS	MVR-tree
Cold state	53.59	59.24	313.79	45.76
Warm state	3.57	45.01	5.39	43.64

Table 6: Response times of *time-slice* queries on disk, in milliseconds.

To emulate a scenario where the full-resolution data does not fit in main memory, we configured our machine to use 2GB for the kernel, so that neither the MVR-tree nor the PostgreSQL table can be completely cached by the operating system or the DBMS.

The experiment begins in cold state, with all the data structures residing on disk and the caches of the operating system and the DBMS completely empty. The first row of Table 6 shows the average time of 1000 executions of the first *time-slice* query in cold state. That query requires to completely load the GraCT data structures into main memory. In the case of the MVR-tree, the execution only loads the nodes of the R-tree traversed by the query. The second row shows the times resulting from running 1000 *time-slice* queries after running the first one, that is, in warm state with the GraCT index already loaded in memory. We tested independently the second, third, and so on, seeing no significant differences.

Even in cold state, GraCT is 10% faster than the MVR-tree for *time-slice S* queries, but once in warm state, the response time of GraCT drops to 3.57 milliseconds, whereas the MVR-tree requires around 45 milliseconds, that is, about 12 times slower. GraCT is much slower than the MVR-tree for *time-slice L* queries in cold state, on the other hand, but again, once we enter in warm state, GraCT becomes about 8 times faster. That is, on regime, a disk-based system with full precision using GraCT as an in-memory cached index is an order of magnitude faster than the MVR-tree and uses much less disk space.

Other queries, like *time-interval*, require a more sophisticated arrangement to use GraCT as an in-memory cache. For example, we can store the consecutive points of each trajectory in contiguous form, indexed by time with a B-tree, and a hash-indexed table of objects pointing to their corresponding trajectories. This enables an efficient traversal of the trajectories of candidate objects in the time intervals where GraCT estimates that they may intersect the query range. Those arrangements also support the simpler queries *object* and *trajectory*. Our experiment suggests that it is much more efficient to verify potential answers on disk than to find them using a disk-based structure.

8. Conclusions

This article introduces GraCT, a compact data structure for the representation of moving objects. In our experiments on real-world trajectory datasets, GraCT compresses the data to 4%–7% of its raw representation size. Within this space, GraCT supports not only direct access to the data, but also various queries looking for objects in spatio-temporal windows. The trajectory data is manipulated directly in compressed form, with only small portions requiring decompression in order to answer queries. The space also includes the storage of spatial indexes representing snapshots at regular time intervals, which speeds up queries. The main algorithmic novelty of GraCT is the representation of the logs of movements between snapshots using grammar-based compression. On

our real-life datasets, grammar-based compression is shown to capture the repetitiveness arising from recurrent trajectories, outperforming statistical compression by a factor of 2 and improving as the datasets become larger. A second novelty is to enrich the grammar by associating information with the nonterminals, such as their relative MBR. This enables GraCT to traverse the logs very fast, without decompressing nonterminals in most cases, thereby answering queries within a few milliseconds. A comparison with the MVR-tree, a classical index, shows that GraCT is two orders of magnitude smaller and still competitive in query times when both run in main memory. When the MVR-tree cannot fit in main memory, GraCT is one or two orders of magnitude faster. This includes an arrangement where GraCT is used as a low-precision main-memory cached index that must verify its result on the disk, where the full-precision data is kept.

Future work in this area will explore other compression mechanisms that exploit repetitiveness, looking for ways to process the trajectory data without decompressing it. Some alternatives are the use of block trees [3] or relative Lempel-Ziv [19]. We also plan to go further in the direction of using GraCT as a main-memory cached index, extending our arrangement to handle all the queries considered in this article.

Acknowledgments

This work has been funded by the European Union Horizon 2020 Marie Skłodowska-Curie Action Fund [grant agreement EU H2020 MSCA RISE BIRDS: 690941]; Ministerio de Economía y Competitividad (PGE and FEDER) [grant number TIN2016-78011-C4-1-R], Centro para el desarrollo Tecnológico e Industrial Programa CIEN 2014 (co-founded with FEDER) [grant number ITC-20151247], Ministerio de Educación y Formación Profesional (FPU program) [grant number FPU16/02914]; Xunta de Galicia (co-founded with FEDER) [grant numbers ED431C 2017/58; ED431G/01]; the Chilean National Science and Technology Development Fund (Fondecyt Grant 1-170048); and the Millennium Institute for Foundational Research on Data (IMFD), Chile.

References

- [1] Asif, M. T., Kannan, S., Dauwels, J., Jaillet, P., 2013. Data compression techniques for urban traffic data. In: Proc. IEEE Symposium on Computational Intelligence in Vehicles and Transportation Systems (CIVTS). pp. 44–49.
- [2] Asif, M. T., Srinivasan, K., Mitrovic, N., Dauwels, J., Jaillet, P., 2015. Near-lossless compression for large traffic networks. *IEEE Transactions on Intelligent Transportation Systems* 16 (4), 1817–1826.
- [3] Belazzougui, D., Gagie, T., Gawrychowski, P., Kärkkäinen, J., Ordóñez, A., Puglisi, S. J., Tabei, Y., 2015. Queries on LZ-bounded encodings. In: Proc. 25th Data Compression Conference (DCC). pp. 83–92.
- [4] Botea, V., Mallett, D., Nascimento, M. A., Sander, J., 2008. Pist: An efficient and practical indexing technique for historical spatio-temporal point data. *GeoInformatica* 12 (2), 143–168.
- [5] Brisaboa, N., Ladra, S., Navarro, G., 2013. DACs: Bringing direct access to variable-length codes. *Information Processing and Management* 49 (1), 392–404.
- [6] Brisaboa, N. R., Fariña, A., Navarro, G., Paramá, J. R., 2007. Lightweight natural language text compression. *Information Retrieval* 10 (1), 1–33.
- [7] Brisaboa, N. R., Gómez-Brandón, A., Navarro, G., Paramá, J. R., 2016. GraCT: A grammar based compressed representation of trajectories. In: Proc. 23rd International Symposium on String Processing and Information Retrieval (SPIRE). pp. 218–230.
- [8] Brisaboa, N. R., Ladra, S., Navarro, G., 2014. Compact representation of web graphs with extended functionality. *Information Systems* 39 (1), 152–174.

- [9] Chakka, V. P., Everspaugh, A., Patel, J. M., 2003. Indexing large trajectory data sets with SETI. In: Proc. Conference on Innovative Data Systems Research (CIDR).
- [10] Cudre-Mauroux, P., Wu, E., Madden, S., 2010. Trajstore: An adaptive storage system for very large trajectory data sets. In: Proc. 26th IEEE International Conference on Data Engineering (ICDE). pp. 109–120.
- [11] Douglas, D. H., Peucker, T. K., 1973. Algorithms for the reduction of the number of points required to represent a line or its caricature. *The Canadian Cartographer* 10 (2), 112–122.
- [12] Gog, S., Beller, T., Moffat, A., Petri, M., 2014. From theory to practice: Plug and play with succinct data structures. In: Proc. 13th International Symposium on Experimental Algorithms (SEA). pp. 326–337.
- [13] Gutiérrez, G., Navarro, G., Rodríguez, A., González, A., Orellana, J., 2005. A spatio-temporal access method based on snapshots and events. In: Proc. 13th ACM International Symposium on Advances in Geographic Information Systems (GIS). pp. 115–124.
- [14] Guttman, A., 1984. R-trees: A dynamic index structure for spatial searching. In: Proc. ACM International Conference on Management of Data (SIGMOD). pp. 47–57.
- [15] Huang, S., Wang, B., Zhu, J., Wang, G., Yu, G., 2014. R-HBase: A multi-dimensional indexing framework for cloud computing environment. In: Proc. IEEE International Conference on Data Mining Workshop. pp. 569–574.
- [16] Hughes, J. N., Annex, A., Eichelberger, C. N., Fox, A., Hulbert, A., Ronquest, M., 2015. Geomesa: A distributed architecture for spatio-temporal fusion. In: Proc. SPIE. Vol. 9473.
- [17] Keogh, E., Chu, S., Hart, D., Pazzani, M., 2001. An online algorithm for segmenting time series. In: Proceedings 2001 IEEE International Conference on Data Mining. pp. 289–296.
- [18] Kieffer, J. C., Yang, E.-H., 2000. Grammar-based codes: A new class of universal lossless source codes. *IEEE Transactions on Information Theory* 46 (3), 737–754.
- [19] Kuruppu, S., Puglisi, S. J., Zobel, J., 2010. Relative Lempel-Ziv compression of genomes for large-scale storage and retrieval. In: Proc. 17th International Symposium on String Processing and Information Retrieval (SPIRE). LNCS 6393. pp. 201–206.
- [20] Larsson, N. J., Moffat, A., 2000. Off-line dictionary-based compression. *Proceedings of the IEEE* 88 (11), 1722–1732.
- [21] Li, Q., Jianming, H., Yi, Z., 2007. A flow volumes data compression approach for traffic network based on principal component analysis. In: Proc. IEEE Intelligent Transportation Systems Conference. pp. 125–130.
- [22] Lin, X., Ma, S., Zhang, H., Wo, T., Huai, J., Mar. 2017. One-pass error bounded trajectory simplification. *Proceedings of the VLDB Endowment* 10 (7), 841–852.
- [23] Liu, J., Zhao, K., Sommer, P., Shang, S., Kusy, B., Jurdak, R., 2015. Bounded quadrant system: Error-bounded trajectory compression on the go. In: Proc. 31st IEEE International Conference on Data Engineering (ICDE). pp. 987–998.
- [24] Ma, Q., Yang, B., Qian, W., Zhou, A., 2009. Query processing of massive trajectory data based on mapreduce. In: Proc. 1st International Workshop on Cloud Data Management (CloudDB). pp. 9–16.
- [25] Meratnia, N., de By, R. A., 2004. Spatiotemporal compression techniques for moving point objects. In: Proc. 9th International Conference on Extending Database Technology, (EDBT). pp. 765–782.
- [26] Muckell, J., Olsen, P. W., Hwang, J.-H., Lawson, C. T., Ravi, S. S., 2014. Compression of trajectory data: A comprehensive evaluation and new approach. *GeoInformatica* 18 (3), 435–460.
- [27] Munro, J. I., 1996. Tables. In: Proc. 16th Conference Foundations of Software Technology and Theoretical Computer Science (FSTTCS). pp. 37–42.
- [28] Munro, J. I., Raman, R., Raman, V., Rao, S., 2012. Succinct representations of permutations and functions. *Theoretical Computer Science* 438, 74–88.
- [29] Nascimento, M. A., Silva, J. R. O., 1998. Towards historical R-trees. In: Proc. ACM Symposium on Applied Computing (SAC). pp. 235–240.
- [30] Navarro, G., 2016. *Compact Data Structures – A practical approach*. Cambridge University Press.
- [31] Ni, J., Ravishankar, C. V., 2007. Indexing spatio-temporal trajectories with efficient polynomial approximations. *IEEE Transactions on Knowledge and Data Engineering* 19 (5), 663–678.
- [32] Nibali, A., He, Z., 2015. Trajic: An effective compression system for trajectory data. *IEEE Transactions on Knowledge and Data Engineering* 27 (11), 3138–3151.
- [33] Nishimura, S., Das, S., Agrawal, D., El Abbadi, A., 2013. MD-HBase: design and implementation of an elastic data infrastructure for cloud-scale location services. *Distributed and Parallel Databases* 31 (2), 289–319.
- [34] Pfoser, D., Jensen, C. S., Theodoridis, Y., 2000. Novel approaches to the indexing of moving object

- trajectories. In: Proc. 26th International Conference on Very Large Data Bases (VLDB). pp. 395–406.
- [35] Potamias, M., Patroumpas, K., Sellis, T., 2006. Sampling trajectory streams with spatiotemporal criteria. In: Proc. 18th International Conference on Scientific and Statistical Database Management (SSDBM). pp. 275–284.
- [36] Richter, K.-F., Schmid, F., Laube, P., 2012. Semantic trajectory compression: Representing urban movement in a nutshell. *Journal of Spatial Information Science* 4 (1), 3–30.
- [37] Sattler, M., Sarlette, R., Klein, R., 2005. Simple and efficient compression of animation sequences. In: Proc. ACM SIGGRAPH/Eurographics Symposium on Computer Animation (SCA). pp. 209–217.
- [38] Schmid, F., Richter, K.-F., Laube, P., 2009. Semantic trajectory compression. In: Proc. 11th International Symposium on Spatial and Temporal Databases (SSTD). pp. 411–416.
- [39] Song, R., Sun, W., Zheng, B., Zheng, Y., May 2014. Press: A novel framework of trajectory compression in road networks. *Proceedings of the VLDB Endowment* 7 (9), 661–672.
- [40] Ta, N., Li, G., Chen, B., Feng, J., 2016. Semantic-aware trajectory compression with urban road network. In: Proc. 17th International Conference (WAIM), Part I. pp. 124–136.
- [41] Tan, H., Luo, W., Ni, L. M., 2012. CloST: A hadoop-based storage system for big spatio-temporal data analytics. In: Proc. 21st ACM International Conference on Information and Knowledge Management (CIKM). pp. 2139–2143.
- [42] Tao, Y., Papadias, D., 2001. Efficient historical R-trees. In: Proc. International Conference on Scientific and Statistical Database Management (SSDBM). pp. 223–232.
- [43] Tao, Y., Papadias, D., 2001. MV3R-tree: A spatio-temporal access method for timestamp and interval queries. In: Proc. 27th International Conference on Very Large Data Bases (VLDB). pp. 431–440.
- [44] Trajcevski, G., Cao, H., Scheuermann, P., Wolfson, O., Vaccaro, D., 2006. On-line data reduction and the quality of history in moving objects databases. In: Proc. 5th ACM International Workshop on Data Engineering for Wireless and Mobile Access. pp. 19–26.
- [45] Váša, L., Marras, S., Hormann, K., Brunnett, G., 2014. Compressing dynamic meshes with geometric Laplacians. *Computer Graphics Forum* 33 (2), 145–154.
- [46] Vazirgiannis, M., Theodoridis, Y., Sellis, T. K., 1998. Spatio-temporal composition and indexing for large multimedia applications. *ACM Multimedia Systems Journal* 6 (4), 284–298.
- [47] Vitter, J. S., 1985. Random sampling with a reservoir. *ACM Transactions on Mathematical Software* 11 (1).
- [48] Wang, H., Zheng, K., Xu, J., Zheng, B., Zhou, X., Sadiq, S., 2014. SharkDB: An in-memory column-oriented trajectory storage. In: Proc. 23rd ACM International Conference on Conference on Information and Knowledge Management (CIKM). pp. 1409–1418.
- [49] Wang, L., Zheng, Y., Xie, X., Ma, W.-Y., 2008. A flexible spatio-temporal indexing scheme for large-scale GPS track retrieval. In: Proc. International Conference on Mobile Data Management (MDM). pp. 1–8.
- [50] Worboys, M. F., 2005. Event-oriented approaches to geographic phenomena. *International Journal of Geographical Information Science* 19 (1), 1–28.
- [51] Xu, X., Han, J., Lu, W., 1990. RT-tree: An improved R-tree index structure for spatiotemporal databases. In: Proc. 4th International Symposium on Spatial Data Handling. Vol. 2. pp. 1040–1049.
- [52] Yang, S., He, Z., Chen, Y.-P. P., 2018. GCOTraj: A storage approach for historical trajectory data sets using grid cells ordering. *Information Sciences* 459, 1 – 19.
- [53] Zhang, Z., Jin, C., Mao, J., Yang, X., Zhou, A., 2017. TrajSpark: A scalable and efficient in-memory management system for big trajectory data. In: Proc. 1st International Joint Conference APWeb-WAIM, Part I. pp. 11–26.
- [54] Zheng, Y., Zhou, X. (Eds.), 2011. *Computing with Spatial Trajectories*. Springer.
- [55] Zhou, P., Zhang, D., Salzberg, B., Cooperman, G., Kollios, G., 2005. Close pair queries in moving object databases. In: Proc. 13th Annual ACM International Workshop on Geographic Information Systems (GIS). pp. 2–11.