

Efficient Evaluation of Continuous Range Queries on Moving Objects

D. V. Kalashnikov S. Prabhakar W. G. Aref S. E. Hambrusch
Department of Computer Sciences
Purdue University
West Lafayette, Indiana 47907
{dvk, sunil, aref, seh}@cs.purdue.edu

Abstract. In this paper we evaluate several in-memory algorithms for efficient and scalable processing of continuous range queries over collections of moving objects. Constant updates to the index are avoided by query indexing. No constraints are imposed on the speed or path of moving objects. We present a detailed analysis of a grid approach which shows the best results for both skewed and uniform data. A sorting based optimization is developed for significantly improving the cache hit ratio. Experimental evaluation establishes that indexing queries using the Grid index yields orders of magnitude better performance than other index structures such as R*-trees.

1 Introduction

The problem of handling different types of queries on moving objects has caught wide attention recently due to the proliferation of GPS and wireless technologies [2][6][13][22]. Using these technologies, it is possible to develop systems where a local server tracks the locations of mobile objects. Mobile objects can report their location to this server through a wireless interface, or the objects can be tracked through ground-based radar or satellites.

As an example, consider a system where aircraft are tracked by radar and their location is reported to the server. Regions of space in which aircraft can be detected by enemy radar or anti-aircraft systems are identified on the server. The server continuously monitors the location of friendly aircraft with respect to these regions and issues alerts as soon as an aircraft is in a detection area. Alternatively, there might be areas where only specially designated aircraft are allowed to be. Such areas can be monitored continuously at the server and detect trespassers. More generally, the location of people and vehicles can be tracked with such systems.

In this paper we address the important problem of evaluating continuous range queries on mobile objects. In contrast to regular queries that are evaluated once, a continuous query remains active over a period of time and has to be continuously evaluated during this time. At any time there will be several continuous queries running at the server. Each of these queries needs to be re-evaluated as the objects move. A major challenge for this problem is repeatedly evaluating all queries within a reasonable amount of time as the numbers of objects being tracked and continuous queries increases.

While the focus of our work is on moving object environments, the techniques presented here can easily be applied in a more general setting. For example, to continuously compute an ϵ -join¹ between a set of relatively fixed points and another set of points that may move arbitrarily. Such continuous queries over ever-changing data are also important for many applications that handle streaming data from sensors and wireless location-based services in Location-commerce (L-commerce) [5][15]. Towards this goal of wider applicability of our techniques, we make no assumptions about the speed and nature of the movement of objects.

Current efforts at evaluating queries over moving objects have focused on the development of disk-based indexes. The problem of scalable, real-time execution of continuous queries is not well suited for disk-based indexing for the

¹ An ϵ -join between two sets of points is defined as all pairs of points, one from each set, such that the distance between the points is less than or equal to ϵ .

following reasons: (i) the need to update the index as objects move; (ii) the need to re-evaluate all queries when any object moves; and (iii) achieving very short execution times for large numbers of moving objects and queries. These factors, combined with the drastically dropping main memory costs makes main memory evaluation highly attractive. The growing importance of main memory based algorithms has been underscored by the Asilomar report [4], which projects that within 10 years main memory sizes will be in the range of terabytes.

The location of a mobile object can be represented in memory as a 2-dimensional point while its other attributes can be stored on disk. One local server is likely to be responsible for handling a limited number of mobile objects (e.g. 1,000,000). As we see later in this paper, all necessary data and auxiliary structures for even large problem sizes can be easily kept in the main memory of an average modern PC.

In order for the solution to be effective it is necessary to efficiently compute the matching between large numbers of objects and queries. While multidimensional indexes tailored for main memory, as proposed in [13], would perform better than disk-oriented structures, the use of an index on the moving objects suffers from the need for constant updating as the objects move – resulting in degraded performance. To avoid this need for constant updating of the index structure and to improve the processing of continuous queries, we propose a very different approach: *Query Indexing*. In contrast to the traditional approach of building an index on the moving objects, we propose to build an index on the queries. This approach is especially suited for evaluating continuous queries over moving objects, since the queries remain active for long periods of time, and objects are constantly moving.

In this paper we investigate several in-memory index structures for efficient and scalable processing of continuous queries. We evaluate not only indexes designed to be used in main memory exclusively but also disk-based indexes adapted and optimized for main memory. Our results show that, contrary to conventional wisdom, using a simple grid-like structure gives the best performance, even when the data is highly skewed. We also propose an effective technique for improving the caching performance. The proposed solutions are extremely efficient, e.g. 100,000 (25,000) continuous queries over 100,000 (1,000,000) objects can be evaluated in as little as 0.288 (0.762) seconds! The use of query indexing is critical for achieving such efficient processing. We also present an analytical evaluation for the optimal grid size. The analysis matches well with the experimental results. A technique for improving the cache-hit rate is developed that achieves a speed up of 100%.

The remainder of this paper is organized as follows. In Section 2 we present related work. Section 3 describes the problem of continuous query processing, *Query Indexing*, and the index structures considered. We also present an effective technique for improving the cache hit rate. Section 4 presents the experimental results and Section 5 concludes the paper.

2 Related work

The growing importance of moving object environments is reflected in the recent body of work addressing issues such as indexing, uncertainty management, broadcasting, and models for spatio-temporal data. Optimization of disk-based index structures has been explored recently for B⁺-trees [21] and multidimensional indexes [13]. Both studies investigate the redesign of the nodes in order to improve cache performance. Neither study addresses the problem of executing continuous queries or the constant movement of objects (changes to data). The goal of our technique is to efficiently and continuously re-generate the mapping between moving objects and queries. Our technique makes no assumptions about the future positions of objects. It is also not necessary for objects to move according to well-behaved patterns as in [22]. The problem of scalable, efficient computation of continuous range queries over moving objects is ideally suited for main memory evaluation. To the best of our knowledge no existing work addresses the main memory execution of multiple concurrent queries on moving objects as proposed in the following sections.

Indexing techniques for moving objects have been proposed in the literature, e.g., [3], [16] index the histories, or trajectories, of the positions of moving objects, while [22] indexes the current and anticipated future positions of the moving objects. In [14], trajectories are mapped to points in a higher-dimensional space which are then indexed. In [22], objects are indexed in their native environment with the index structure being parameterized with velocity vectors so that the index can be viewed at future times. This is achieved by assuming that an object will remain at the same speed and in the same direction until an update is received from the object.

Uncertainty in the positions of the objects is dealt with by controlling the update frequency [17][28], where objects report their positions and velocity vectors when their actual positions deviate from what they have previously reported by some threshold. Tayeb et. al. [26] use quadtrees to index the trajectories of one-dimensional moving

points. Kollios [14] et. al. map moving objects and their velocities into points and store the points in a kD-tree. Pfoser et. al [18][19] index the past trajectories of moving objects that are presented as connected line segments. The problem of answering a range query for a collection of moving objects is addressed in [2] through the use of indexing schemes using external range trees. [27] and [29] consider the management of collections of moving points in the plane by describing the current and expected positions of each point in the future. They address how often to update the locations of the points to balance the costs of updates against imprecision in the point positions. Issues relating to location dependent database querying are addressed in [24]. Broadcast of data becomes an important technique for scalable communication in the mobile environment. Efficient broadcast techniques are proposed in [1][8][10][12]. Spatio-temporal database models to support moving objects, spatio-temporal types and supporting operations have been developed in [6][7]. An excellent review of multidimensional index structures including grid-like and Quad-Tree based structures can be found in [26].

3 Continuous Query Evaluation

3.1 Updating Object Locations

The issue of obtaining the updated locations of objects is independent of the technique used for evaluating the queries. Since the focus of this research is on the efficient evaluation of queries, we assume that updated location information is available at the server, without considering how exactly it was made available. Below we briefly discuss the common assumptions made in order to reduce communication.

The most common assumption is that each object moves on a straight line path with a constant speed, and updates the server with its direction of movement and speed when they change. A similar assumption is that objects are moving with constant speed on a known road. A mutual feature of these assumptions is that for each moving object the server can determine its location based upon a formula. In our experiments new locations of objects are generated at the beginning of each cycle. While some index structures for moving objects rely upon restricted models of movement (e.g. [25]), *Query Indexing* allows objects to move arbitrarily. Therefore, the objects can move anywhere in the domain, but the overall object distribution chosen for the experiment is maintained (uniform, skewed, etc.).

3.2 Query Indexing

The problem of continuous query evaluation is: *Given a set of queries and a set of moving objects, continuously determine the set of objects that are contained within each query.*

Clearly, with a large number of queries and moving objects, it is infeasible to re-evaluate each query whenever any object moves. A more practical approach is to re-evaluate all queries periodically taking into account the latest positions of all objects. In order for the results to be useful, a short re-evaluation period is desired. The goal of the query evaluation techniques is therefore to re-evaluate all queries in as short a time as possible.

The naïve approach to this problem is to compare each object against each query in every period. A more effective solution seems to be the traditional approach of building an index (such as an R-tree) on the objects to speed up the queries. Although for regular queries this would result in improvements in performance over the naïve approach, it suffers a major drawback for the moving objects environment: the index needs to be continuously updated as object positions change. Maintaining such an index on mobile data proves to be quite a challenging task [22].

We propose a novel solution to this problem: *Query Indexing*. Instead of building an index on the moving objects (which would require frequent updating), create an index on the more stable queries [20]. Any spatial index structure can be used to build the query index.

The evaluation of continuous queries in each cycle proceeds as follows. For object P , let $P.qset$ denote the set of all queries in which P is contained. For query Q , let $Q.pset$ denote the set of objects contained in Q . The goal is to compute the set $Q.pset$ for each query Q based upon the current locations of the objects by the end of each cycle. Because of the lack of speed limitations $Q.pset$ and $P.qset$ can be completely different from one cycle to the next. Consequently, incremental solutions are infeasible. In each cycle, we first use the query index to compute $P.qset$ for each object P . Next, for each query Q in $P.qset$, we add P to $Q.pset$.

Some important consequences of indexing the queries instead of the data should be noted. Firstly, the index needs no modification unless there is a change to the set of queries – a relatively less frequent event in comparison to changes to object locations since we are dealing with continuous queries. Secondly, the location of an object can change greatly from one cycle to the next without having any impact on the performance. In other words, there is no restriction on the nature of movement of the objects. This is not true for object indexing techniques such as [25] which rely upon certain assumptions about the movement of objects.

Next we discuss the feasibility of in-memory query indexing by evaluating different types of indexes for queries. Clearly, if we are unable to select an appropriate index, the time needed to complete one cycle can be large (e.g. 1 min) and the approach would be unacceptable. Below we briefly discuss indexing techniques for building a query index in main memory. In Section 4 we evaluate the performance of these alternative indexes.

3.3 Indexing Techniques

We consider the following five well-known index structures: R*-Tree, R-Tree, CR-Tree, Quad-Tree, and Grid. The R-tree and R*-tree index structures are designed to be disk-based structures. This is principally achieved by choosing the node size to be a multiple of disk page size. The CR-Tree, on the other hand, is a variant of R-trees optimized for main memory. All of these indexes were implemented for main memory use. In order to make a fair comparison, we did not choose large node sizes for these trees. Instead, we experimentally determined the best choice of node size for main-memory evaluation. All three indexes (R*-Tree, R-Tree, and CR-Tree) showed best performance when the number of entries per node was chosen to be five. This value was used for all experiments.

Details of the CR-Tree are described in [13]. The main idea is to make R-Tree cache-conscious by compressing MBRs. This is achieved by using so-called Quantized Relative Minimum Bounding Rectangles (QRMBR). Other well-known minor optimizations have also been proposed in the paper. We implemented the CR-Tree based upon the main idea of QRMBRs without the other optimizations.

Because many variations exist, we describe the Grid index as it is used here for query indexing. The Grid index is a two-dimensional array of “cells”. Each cell represents a region of space generated by partitioning the domain using a uniform grid. Figure 3-1 shows an example of a Grid. Throughout the paper, we assume that the domain is normalized to the unit square.

In this example, the domain is divided into a 10x10 grid of 100 cells, each of size 0.1 x 0.1. Since we have a uniform grid, given the coordinates of an object, it is easy to calculate the cell that it falls under. The cell coordinates for a point $z(x, y)$ are easily computed as:

$$\text{cell_x} = (\text{int}) x / hX;$$

$$\text{cell_y} = (\text{int}) y / hY;$$

where x and y are coordinates of the point we are interested in, hX and hY are the horizontal and vertical size of each cell (e.g., in our example $hX=hY=0.1$). Applying the formula above we can determine the cell that a point belongs to in $O(1)$ time. Each cell contains two lists that are identified as *full* and *part* (see Figure 3-1). The *full* list of a cell contains *pointers* to all the queries that fully cover the cell. The *part* list of each cell contains *pointers* to all the queries that only partially cover the cell.

The choice of data structures for the *full* and *part* lists is critical for performance. We implemented these lists as dynamic arrays rather than lists. A dynamic array is a standard data structure for arrays whose size can be adjusted dynamically. Using dynamic arrays instead of lists improves the performance by roughly 40%.

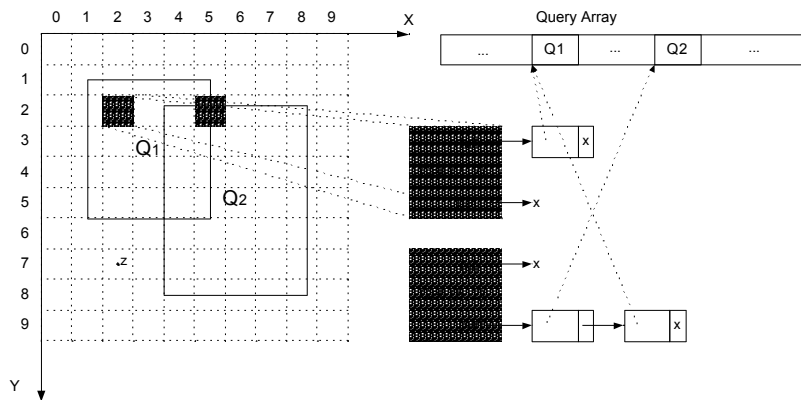


Figure 3-1 Example of Grid

An analytical solution for the appropriate choice of grid size is presented in Section 3.4. As will be seen in the experimental section, this simple one-level grid index outperforms all other structures for uniform as well as skewed data. However, for the case of highly skewed data (e.g. roughly half the queries fall within one cell), the *full* and *part* grow too large. Such situations are easily handled by switching to a two-tier grid. If any of the lists grows beyond a threshold value, the grid index converts to a directory grid and a few secondary grids. The directory grid is used to determine which secondary grid to use. Each directory grid cell points to a secondary grid. The secondary grid is used in the same way as the one-level grid. While this idea of generating an extra layer can be applied as many times as is necessitated by the data distribution, it is unlikely in practice. Consider for example, that the domain of interest represents a 1000-kilometer by 1000-kilometer region. With a 1000 x 1000 grid, a cell of the two-tier grid corresponds a square of side 1 meter – it is very unlikely that there will be very many objects or queries in such a small region in practice.

The astute reader will observe that the Grid index and the quad-tree are closely related – both are space partitioning and split a region if it is overfull. There is however, a subtle difference: the Grid index as described is a height-balanced structure. In contrast a quad-tree is not necessarily height-balanced. The advantage for the Grid is therefore that it avoids conditional (“if”) branches in its search algorithm.

Another advantage of the Grid is that it typically has far more cells per level. A quad-tree therefore tends to be very deep, especially for skewed data. We expect that a quad-tree like structure that has more cells per level would perform better than a quad-tree. In order to test this hypothesis, we also consider what we call a 32-Tree. The 32-tree is identical to a quad-tree, except that it divides a cell using a 32x32 grid, unlike quad-tree’s 2x2 grid. As with the Grid index, pointers to children are used instead of keeping an array of pointers to children. In order to further improve performance, we implemented the following optimization: In addition to leaf nodes, internal nodes can also have an associated *full* list. Only leaf nodes have a *part* list. A *full* list contains all queries that fully cover the bounding rectangle (BR) of the node, but do not fully cover the BR of its parent. Adding a rectangle (or region) to a node proceed as follows. If the rectangle fully covers the BR, it is added to the *full* list and the algorithm stops for that node. If this is a leaf node and there is space in the *part* list, the rectangle is added to *part list*. Otherwise the set of all relevant children is determined and the procedure is applied to each of them.

Storing *full* lists in non-leaf nodes has two advantages. One advantage is saving of space: without such lists, when a query fully covers a node’s BR it would be duplicated in all the node’s children. A second advantage is that it has the potential to speed up point queries. If a point query is falls within the BR of a node then it is relevant to all queries in the *full* list of this node – no further checks are needed. Leaf node split is based on the *part* list size only. While there are many more optimizations possible some of which can be found in [26], we did not explore these further. The purpose of studying the 32-Tree is to establish the generality and flexibility of the grid-based approach.

3.4 Choice of Cell Size

We now present an analysis of appropriate choice for the tile size for the Grid index in the context of main-memory query indexing. Consider the case where m square with side q queries uniformly distributed on $[0,1] \times [0,1]$ are added to index, see Figure 3-2. Let c denote the side of each cell. Then q can be represented as $i \cdot c + x$, where $x \in [0, c)$ and i is an integer. We now analyze the average number of cells partly covered by a query. Without loss of generality let us consider the case where the top-left corner of query Q is somewhere within cell $(0,0)$ and $q > c$.

It can be verified that if the corner is inside Set0 then Q is present in $4 \cdot i$ part lists, for Set1 this number is $4 \cdot i + 2$, and for Set2 it is $4 \cdot i + 4$. Assuming uniform distribution, the probability that the corner is inside Set0 is $(c-x)^2/c^2$, inside Set1 is $2x(c-x)/c^2$, and inside Set2 is x^2/c^2 .

Therefore, on the average each query will end up in $avg = [4i(c-x)^2 + (4i+2)2x(c-x) + (4i+4)x^2]/c^2$ part lists. This formula simplifies to $avg = 4c[x+i \cdot c]/c^2 = 4q/c$.

Correspondingly, m queries will end up in $4qm/c$ part lists. Given the fact that total number of cells can be computed as $1/c^2$, each cell will have part list of size $4qmc$ on the average.

For our implementation, the difference in time needed to process a cell when its part list is empty vs. when its part list has size one is very small. By choosing c such that $4qmc=1$, that is $c=1/4qm$, on the average, the size of a cell's part list will be one. The case where $q < c$ can be analyzed similarly. The final formula for choosing c is:

$$c = \frac{1}{4qm}, \text{ if } q > \frac{1}{2\sqrt{m}}; \text{ and } c = \frac{1}{\sqrt{m}} - q, \text{ otherwise.}$$

For our example in Figure 4-7. we have $m=25,000$; $q=0.01$. Since $\frac{1}{2\sqrt{m}} = 0.001 < q$, the first formula should be used, i.e.

$c=1/4qm=0.001$. This means that grid should be of size 1000x1000 and larger grids will not produce better performance. We study the impact of tile size in the experimental section and show that the results match the analytical prediction.




Figure 3-2. Choosing grid size.

3.5 Improving the Cache Hit Rate

The performance of main-memory algorithms is greatly affected by cache hit rates. In this section we describe an optimization that can drastically improve cache hit rates (and consequently the overall performance) for the query indexing approach.

In each cycle the processing involves searching the index structure for each objects current location in order to determine the queries that cover the object's current location.

For each object, its cell is computed, and the *full* and *part* lists of this cell are accessed. The algorithm simply processes objects in sequential order in the array. Consider the example shown in **Figure 3-3**. The order in which the objects appear in the array is shown on the left of the figure in the "Unsorted Point Array". Note that we use the terms object and point interchangeably. In this example, the lists pointed to by Cell (0,0) will be accessed for processing object P_2 and then later for processing object P_n . Since several other points are processed after P_2 and before P_n , it is very likely that the lists for Cell(0,0) will not be in the cache when P_n is processed – resulting in a cache miss.




Figure 3-3 Example of sorted and un-sorted Object Arrays

If we expect that objects will maintain their locality, then we can improve cache-hit rates by altering the order of processing the objects. If we re-order objects in the array such that objects that are close together in our 2D domain are also close together in the object array, as in the array on the right labeled "Sorted Point Array" shown in **Figure 3-3**. With this ordering, object P_2 will be analyzed first and therefore Cell(0, 0) and its lists will be processed. Then, object P_n will be analyzed and Cell(0, 0) and its lists will be processed again. In this situation everything relevant to Cell(0, 0) is likely to remain in the CPU cache after the first processing and will be reused from the cache during the second processing. The speed up effect is also achieved because objects that are close together are more likely to be covered by the same queries than objects that are far apart, thus queries are more likely to be retrieved from the cache rather than from main memory.

Sorting the objects to ensure that objects that are close to each other are also close in the array order can easily be achieved. One possible approach is to group objects by cells – i.e. all objects that fall under each cell are placed adjacently in the array and thus processed together. The objects grouped by cell can then be placed in the array using a row-major or column-major ordering for the cells. Although this will be effective, the benefit of the sorting is lost if the object moves out of its current cell and enters an adjacent cell that is not close by in the ordering used for the cells (e.g. object moves to adjacent cell in next row and row major ordering is used). We propose an alternative approach: Order the points using any of the well-known space filling curves such as Z-order or Hilbert curve. We choose to use a sorting based on the Z-order. Z-sorting significantly improves the performance of the main memory algorithm, as will be seen in the experiment section.

It is important to understand that the use of this technique does not require that objects have to preserve their locality. The only effect of sorting the objects according to their earlier positions is to alter the order in which objects are processed in each cycle. The objects are still free to move arbitrarily. Of course, the *effectiveness* of this technique relies upon objects maintaining their locality over a period of time. If it turns out that objects do not maintain their locality then we are, on the average, no worse than the situation in which we do not sort. Thus, for the case where objects preserve locality sorting the objects based upon their location at some time can be beneficial. It should also be noted that the exact position used for each object is not important. Thus the sorting can be carried out infrequently (say once a day).

4 Experiments

In this section we present the performance results for the index structures. Each index structure was implemented and tested – not simulated. The results report the actual times for the execution of the various algorithms. First we describe the parameters of the experiments, followed by the results and discussion.

In all our experiments we used a 1 GHz Pentium III machine with 2GB of memory. The machine has 32K of level 1 cache (16K for instructions and 16K for data) and 256K level 2 cache. Moving objects were represented as points distributed on the unit square $[0,1] \times [0,1]$. The number of objects ranges from 100,000 to 1,000,000. Range-queries were represented as squares with sides 0.01. Experiments with other sizes of queries yielded similar results and are thus omitted. For distributions of objects and queries in the domain we considered the following cases:

- 1) **Uniform:** Objects and queries are uniformly distributed.
- 2) **Skewed:** The objects and queries are distributed among five clusters. Within each cluster objects and data are distributed normally with a standard deviation of 0.05 for objects and 0.1 for queries.
- 3) **Hyper-skewed:** Half of the objects (queries) are distributed uniformly on $[0,1] \times [0,1]$, the other half on $[0,0.001] \times [0,0.001]$. Queries in $[0,0.001] \times [0,0.001]$ are squares with sides 0.00001 to avoid excessive selectivity.

We consider the skewed case to be most representative. The hyper-skewed case represents a pathological situation designed to study the performance of the schemes under extreme skew. In the majority of our experiments the Grid was chosen to consist of 1000x1000 cells. The testing proceeds as follows: First, queries and objects are generated and put into arrays. Then the index is initialized and the queries are added to it. Then in each cycle, we first update the locations followed by an evaluation of the query results.

4.1 Comparing efficiency of indexes




Figure 4-1




Figure 4-2

Figure 4-1 shows the results for various combinations of number of objects and queries with uniform distribution. The y-axis gives the processing time for one cycle in seconds for each experiment. Figure 4-2 shows similar results for the skewed case.

Each cycle consists of two steps: *moving* objects (i.e., determining current object locations) and *evaluation/processing*. From Figure 4-2 for the case of 100,000 objects and 100,000 queries we can see that the evaluation step for Grid takes 0.628 seconds. Altering object locations takes 0.15 seconds for 100K objects and 1.5 seconds for 1M objects on the average. Thus the length of each cycle is just 0.778 seconds on the average.

The Grid index gives the best performance in all these cases. While the superior performance of Grid for the uniform case is expected, the case for skewed data is surprising. For all experiments the Grid index consisted of only a single level.

Figure 4-3 shows the results for the hyper-skewed case. For the hyper-skewed case, a second level grid is required for the cell containing a large number of queries. It is interesting to see that the Grid index once again outperforms the other schemes. There is a significant difference in performance of Grid and the other approaches for all three distributions. For example, with 1,000,000 objects and 25,000 queries, Grid evaluates all queries in 1.724 seconds as compared to 33.2 seconds for the R-Tree, and 8.5 seconds for the Quad Tree. This extremely fast evaluation implies that with the Grid index, the cycle time is very small – in other words, we can re-compute the set of objects contained in each query every 3.2 seconds or faster (1.7 seconds for the evaluation step + 1.5 seconds for updating the locations of objects). This establishes the feasibility of in-memory query indexing for managing continuous queries.

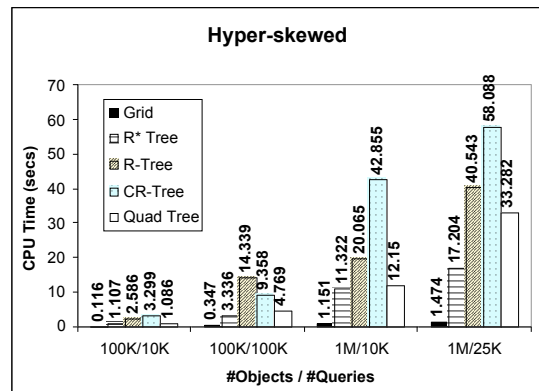


Figure 4-3

4.2 32-Tree index

It can be seen that the quad-tree performs better than R-Tree like data structures for skewed cases, but worse for the majority of the uniform and hyper-skewed cases.

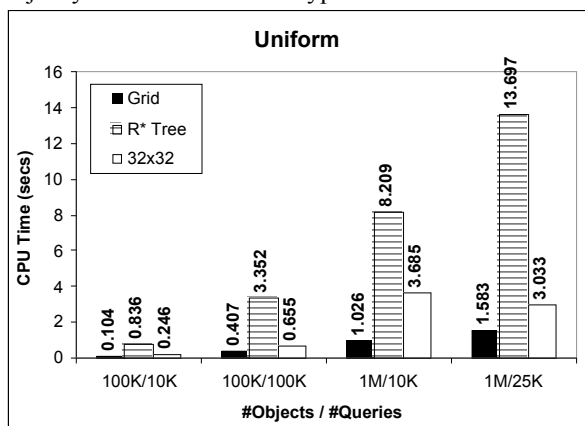


Figure 4-4




Figure 4-5

The problem with a quad-tree for hyper-skewed case is that it has a large height. This suggests that if it were able to zoom faster it would be a better index than R*-Tree. We test this hypothesis by evaluating the 32-Tree which is similar to the quad-tree except that it has more divisions at each node.

The performance of the 32-Tree along with that for the Grid and R*-Tree for uniform, skewed, and hyper-skewed data is shown in Figures 4-4, 4-5, and 4-6 respectively. As can be seen from the graphs our hypothesis is true: the performance of the 32-Tree lies between that of the R*-Tree and the Grid.




Figure 4-6

4.3 Choice of Grid Size

In this experiment we study the impact of the number of cells in the grid. The analysis in Section 3.4 predicted that a choice of a 1000x1000 grid is optimal. Figure 4-7. presents the processing time needed with grid sizes 100x100, 1000x1000, and 2000x2000 cells. As can be seen, increasing the number of cells has the effect of reducing the average number of queries for a cell thereby reducing the processing time. There is a substantial increase in performance as we move from 100x100 cells to 1000x1000 cells. The increase is minor when we move from 1000x1000 to 2000x2000 cells for our case of 1M objects and 25K queries. This behavior corroborates the analytical results.




Figure 4-7. Impact of grid size on processing time.




Figure 4-8. Effectiveness of Z-Sorting.

4.4 Z-Sort Optimization

Figure 4-8 illustrates the effect of the Z-Sort technique on evaluation time for ideally Z-sorted data. Z-sorting reorders the data such that objects that are close together tend to be processed close together. When processing each object in the array from the beginning to the end, objects that close to each other will tend to reuse information stored in the cache rather than retrieving it from main-memory. From the results, we see that sorting objects improves the performance by roughly 50%.

4.5 Time to add/remove queries to/from Grid

We now study the efficiency of modifying the Grid index. The results in Figure 4-9 show how long it takes to add and remove queries to/from an existing index that already contains some queries. Although modifications to queries are expected to be rare, we see that adding or removing queries is done very efficiently with the Grid. For example, the 100% bar shows that 100% of 25K queries can be added or deleted in only 2.408 seconds. The decision whether to add or delete a query at a particular step is made with probability of 0.5 for each query. Therefore we see that even significant changes to the query set can be effectively handled by the Grid approach.




Figure 4-9 Adding and deleting queries.

5 Summary

In this paper we presented a Query-Index approach for in-memory evaluation of continuous range queries on moving objects. We established that the proposed approach is in fact a very efficient solution even if there are no limits on object speed or nature of movement – a common restriction made is similar research. We presented results for seven different in-memory spatial indexes. The Grid approach showed the best result even for the skewed case. A technique of sorting the objects to improve the cache hit-ratio was presented. The performance of the Grid index was roughly doubled with this optimization. An analysis for selecting optimal grid size and experimental validation was presented. We also showed that even though the set of continuous queries is to remain almost unchanged, nevertheless Grid can very efficiently add or remove large numbers of queries. Overall, indexing the queries using the Grid index gives orders of magnitude better performance than other index structures such as R*-trees.

6 References

- [1] S. Acharya, R. Alonso, M. J. Franklin, and S. B. Zdonik. Broadcast disks: Data management for asymmetric communications environments. In Proceedings of SIGMOD, pages 199-210, 22-25 May 1995.
- [2] P. K. Agarwal, L. Arge, and J. Erickson. Indexing moving points. In Proc. 2000 ACM Symposium on Principles of Database Systems (PODS), 2000.
- [3] A. Becker, S. Gschwind, T. Ohler, B. Seeger, and P. Widmayer. An asymptotically optimal multiversion B-tree. The VLDB Journal, 5(4):264-275, December 1996.
- [4] P. A. Bernstein et. al. The Asilomar Report on Database Research. SIGMOD Record, 27(4), pages 74-80, 1998.
- [5] US Wireless Corp. The market potential of the wireless location industry. <http://www.uswcorp.com/>
- [6] L. Forlizzi, R. H. Guting, E. Nardelli, and M. Scheider. A data model and data structures for moving objects databases. In Proc. of ACM SIGMOD 2000.
- [7] R.H. Guting, M.H.Bohlen, M. Erwig, C.S. Jensen, N.A. Lorentzos, M. Schneider, and M. Vazirgiannis. A foundation for representing and querying moving objects. ACM Transactions on Database Systems, 2000.
- [8] S. E. Hambrusch, C.-M. Liu, W. Aref, and S. Prabhakar. Query Processing in Broadcasted Spatial Index Trees. In 7th International Symposium on Spatial and Temporal Databases (SSTD 2001), July 2001.
- [9] S. Hambrusch., S. Prabhakar, W. Aref, D. Kalashnikov, and Y. Xia, PLACE project, <http://www.cs.purdue.edu/place/>
- [10] Q. Hu, W.-C. Lee, and D. L. Lee. Power conservative multi-attribute queries on data broadcast. In Proceedings of the International Conference on Data Engineering (ICDE), pages 157--166, 2000.
- [11] Q. Hu, W.-C. Lee, and D. L. Lee. A hybrid index technique for power efficient data broadcast. Distributed and Parallel Databases, 9(2):151-177, 2001.
- [12] T. Imielinski, S. Viswanathan, and B. R. Badrinath. Energy efficient indexing on air. In Proc. of SIGMOD 1994.
- [13] K. Kim, S. K. Cha, and K. Kwon. Optimizing multidimensional index trees for main memory access. In Proc. of SIGMOD 2000.
- [14] G. Kollios, D. Gunopulos, and V.J. Tsotras. On indexing mobile objects. In Proc. ACM Symp. on Principles of Database Systems (PODS), June 1999.
- [15] H. Koshima and J. Hoshen. Personal locator services emerge. IEEE Spectrum, 37(2):41-48, 2000.
- [16] A. Kumar, V. J. Tsotras, and C. Faloutsos. Designing access methods for bitemporal databases. 10(1):1-20, 1998.
- [17] P. Pfoser and C. S. Jensen. Capturing the uncertainty of moving-objects representations. In Proceedings of the SSDBM Conf., pages 123-132, 1999.
- [18] P. Pfoser, C.S. Jensen, and Y. Theodoridis. Novel approaches in query processing for moving objects. In Proc. of the 26th Int. Conf. on Very Large Databases (VLDB), Cairo, Egypt, 2000.
- [19] P. Pfoser, Y. Theodoridis, and C.S. Jensen. Indexing trajectories of moving point objects. Technical Report Chorochronos Tech. Rep. CH-99-3, June 1999.
- [20] S. Prabhakar, Y. Xia, D. Kalashnikov, W. Aref, and S. Hambrusch. Query Indexing and Velocity Constrained Indexing: Scalable Techniques for Continuous Queries on Moving Objects. TR 01-012, Purdue University, July 2001.
- [21] J. Rao and K. A. Ross. Making B⁺-Trees Cache Conscious in Main Memory. In Proc. ACM SIGMOD Int. Conf. on Management of Data, pages 475-486, 2000.
- [22] S. Saltenis, C. Jensen, S. Leutenegger, and M. Lopez. Indexing the position of continuously moving objects. In Proceedings of ACM SIGMOD Conference, May 2000.
- [23] H. Samet. The Design and Analysis of Spatial Data Structures. Addison-Wesley, 1990.
- [24] A. Y. Seydim, M. H. Dunham, and V. Kumar. Location dependent query processing. In Proc. of the 2nd ACM international workshop on Data engineering for wireless and mobile access, pp 47-53, 2001
- [25] A. Prasad Sistla, Ouri Wolfson, Sam Chamberlain, and Son Dao. Modeling and querying moving objects. In Proc. of the 14th Int. Conference on Data Engineering (ICDE'97), pages 422-432, 1997.
- [26] Jamel Tayeb, Ozgur Ulusoy, and Ouri Wolfson. A quadtree-based dynamic attribute indexing method. The Computer Journal, 41(3):185-200, 1998.
- [27] O. Wolfson, S. Chamberlain, S. Dao, L. Jiang, and G. Mendez. Cost and imprecision in modeling the position of moving objects. In Proceedings of the Fourteenth International Conference on Data Engineering (ICDE), 1998.
- [28] O. Wolfson, P. A. Sistla, S. Chamberlain, and Y. Yesha. Updating and querying databases that track mobile units. Distributed and Parallel Databases, 7(3):257-387, 1999.
- [29] Ouri Wolfson, Bo Xu, Sam Chamberlain, and L. Jiang. Moving objects databases: Issues and solutions. In Proceedings of the SSDBM Conf., pages 111-122, 1998.