Temporal Range Exploration of Large Scale Multidimensional Time Series Data

Joseph JaJa

Jusub Kim

Qin Wang

Institute for Advanced Computer Studies
Department of Electrical and Computer Engineering
University of Maryland, College Park, MD 20742
E-mail: {joseph, jusub, qinwang}@umiacs.umd.edu

Abstract—We consider the problem of querying large scale multidimensional time series data to discover events of interest, test and validate hypotheses, or to associate temporal patterns with specific events. This type of data currently dominate most other types of available data, and will very likely become even more prevalent in the future given the current trends in collecting time series of business, scientific, demographic, and simulation data. The ability to explore such collections interactively, even at a coarse level, will be critical in discovering the information and knowledge embedded in such collections. We develop indexing techniques and search algorithms to efficiently handle temporal range value querying of multidimensional time series data. Our indexing uses linear space data structures that enable the handling of queries in I/O time that is essentially the same as that of handling a single time slice, assuming the availability of a logarithmic number of processors as a function of size of the temporal window. A data structure with provably good bounds is also presented for the case when the number of multidimensional objects is relatively small. These techniques improve significantly over standard techniques for either serial or parallel processing, and are evaluated by extensive experimental results that confirm their superior performance.

I. INTRODUCTION

While considerable work has been performed on indexing multidimensional data (see for example [1]), relatively little efforts have been made for developing techniques that specifically deal with time series of multidimensional data. However, such type of data is abundantly available, and is currently being generated at an unprecedented rate in a wide variety of applications that include environmental monitoring, scientific simulations, medical and financial databases, and demographic studies. For example, the remotely sensed data alone is expected to exceed several terabytes per day in the

next couple of years. This type of spatio-temporal data constitutes large scale multidimensional time series data that are currently very hard to manipulate or analyze. Another example are the tens of thousands of weather stations around the world which provide hourly or daily surface data such as precipitation, temperature, winds, pressure, and snowfall. Such data can be used to model and predict short-term and long-term weather patterns or correlate spatio-temporal patterns with phenomena such as storms, hurricanes, or tornados. Similarly, in the stock market, each stock can be characterized by its daily opening price, closing price, and trading volume, and hence the collection of long time series of such data for various stocks can be used to understand short and long term financial trends.

In multidimensional time-series data, we are given a collection of n time series such that each time series describes the evolution of an object (point) in multi-dimensional space as a function of time. A possible approach for exploring such data can be based on determining which objects behave in a certain way over some time window. Such exploration can be used for example to test a hypothesis relating patterns to specific events that happened during that time window or classifying objects based on their behavior within that time window. Since quite often, we will be experimenting with many variations of a pattern to determine appropriate correlations to an event of interest, or experimenting with many variations of the parameters of a certain hypothesis to test its validity, it is critical that each exploration be achieved very quickly preferably on the available large scale multidimensional data without sampling. We focus in this paper on techniques that minimize the overall number of I/O accesses and that are suited for sequential implementation as well as parallel implementation on

(small) clusters of processors.

Current multidimensional access techniques handle two types of multidimensional objects, points and extended objects such as lines, polygons, or polyhedra. In this paper we focus only on multidimensional point data and address the temporal extensions of the orthogonal range value queries, which constitute the most fundamental type of queries for multidimensional data. This type of queries is introduced next.¹

Given N objects each specified by a set of d attributes, let $O_i(l)$ indicate l th attribute value of object i and $D(.,.)$ represent a distance metric between two multidimensional points.

Query 1-1. (Range Query in Multidimensional Space) *Given d value ranges $[a_l, b_l]$, $1 \leq l \leq d$, determine the set of objects that fall within the query rectangle defined by these ranges.*

$RangeQ = \{O_i \mid a_l \leq O_i(l) \leq b_l, \text{ for } \forall l, 1 \leq l \leq d\}.$

For the case of multidimensional time-series data, we are primarily interested in addressing the multidimensional data trends along the time axis. By incorporating the time-interval component, we can extend the above types of queries into two special cases and a more general case.

Given m snapshots of N d -dimensional objects at time instances $t_1, t_2, \dots, t_m$, let $O_i^j(l)$ denote l th attribute value of object i at time t_j .

Query 2-1. (Temporal Range Query 2-1 for Multidimensional Time-series Data) *Given d value ranges $[a_l, b_l]$, $1 \leq l \leq d$, and a time interval $[t_s, t_e]$, determine the set of objects that fall within the query range values during every time instance that appears in the interval $[t_s, t_e]$.*

$TRangeQ1 = \{O_i \mid a_l \leq O_i^j(l) \leq b_l, \text{ for } \forall l, 1 \leq l \leq d \text{ at } \forall j \text{ time stamps, } t_s \leq t_j \leq t_e\}.$

Query 2-2. (Temporal Range Query 2-2 for Multidimensional Time-series Data) *Given d value ranges $[a_l, b_l]$, $1 \leq l \leq d$, and a time interval $[t_s, t_e]$, determine the set of objects that fall within the query range values at some time instance within $[t_s, t_e]$.*

$TRangeQ2 = \{O_i \mid a_l \leq O_i^j(l) \leq b_l, \text{ for } \forall l, 1 \leq l \leq d \text{ for some } j, t_s \leq t_j \leq t_e\}.$

Query 2-3. (Temporal Range Query 2-3 for Multidimensional Time-series Data) *Given d value ranges $[a_l, b_l]$, $1 \leq l \leq d$, and a time interval $[t_s, t_e]$, determine the set of objects, each of which falling within the query range values in at least a certain fraction p of time stamps during the query time interval.*

$TRangeQ3 = \{O_i \mid a_l \leq O_i^j(l) \leq b_l, \text{ for } \forall l, 1 \leq l \leq d \text{ for at least a fraction } p, (0 \leq p \leq 1), \text{ of time stamps during the time interval } [t_s, t_e]\}.$

In this extended abstract, we focus on Query 2-1 and introduce very efficient strategies to handle such queries. The performance of our techniques is confirmed by extensive simulation results on widely different types of synthetic data. Extensions to Queries 2-2 and 2-3 as well as extensions to the case when the data is noisy are also briefly mentioned in the last section of this extended abstract.

A. Possible Approaches Based on Standard Techniques

There are two straightforward ways to extend existing multidimensional access technique to handle the above queries. The first consists of viewing the multidimensional time series data in $(d + 1)$ dimensions, and use existing techniques to handle the temporal range queries, which amount to treating the temporal dimension just like any other dimension. This implies that object i at time t_j is represented by the coordinates $(O_i^j(1), O_i^j(2), \dots, O_i^j(d), t_j)$ in $(d + 1)$ dimensional space, i.e., the evolution of an object along m time instances is represented by m points in $(d + 1)$ -dimensional space. Such an approach can be couched within the framework explored for *generalized intersection searching* [2], which translates for our case into coloring each point in $(d + 1)$ -dimensional space with its object id. Hence the m points describing the evolution of object i are colored with color i . As a result, the temporal range queries are transformed into determining the distinct colors that appear at a certain frequency within the query rectangle. For example, Query 2-2 amounts to determining the number of distinct colors (and hence object ids) of the points that fall within the $(d + 1)$ -dimensional query rectangle. The best known internal memory algorithms for special cases of this problem appear in [3] but no external memory algorithms are known to the authors best knowledge.

There are two main disadvantages with such an approach. The first is the fact that, for any time window of size w , the search algorithm, based on any technique to solve the orthogonal range value problem, will identify

¹In the remainder of this paper, an object means a multidimensional point.

some subset of the corresponding w points of each object, which fall within the query range values. Hence for our queries, the number of candidate points explored can be arbitrarily larger than the output size (consisting of the indices of the objects that satisfy the query), which is undesirable especially for large time windows. The second disadvantage is the fact that the resulting data structure, say an R-tree[4] or any of its variants [5], [6], [7], [8], cannot be easily handled on a cluster of processors and corresponding parallel search algorithms for such data structures tend to be complex and not necessarily scalable. Our simulation results illustrate the substantial inferior performance of such an approach relative to our new approach, even in the case of serial processing.

The second straightforward approach would be to build a multidimensional indexing data structure for the d -dimensional points at each time instance, and then sequentially search each of the data structures for each time instance of the query interval. This approach, while easy to implement, can be quite inefficient and will generate, as we proceed along the time axis, many possible candidates most of which will be ruled out by the end of the process. Moreover, while this strategy leads to a fast parallel implementation by analyzing all the time slices in parallel, the number of processors required is proportional to the length of the query interval, as opposed to our strategy that will only require a logarithmic number of such processors and report each proper object only once.

A more involved approach can be based on more sophisticated data structures such as MR-tree[8], Historical R-tree(HR-tree)[9][10], and RT-tree[8]. These data structures focus on reducing the redundancy in a series of R-trees built along the time axis by making use of identical branches in consecutive R-trees. None of these techniques are appropriate for our problem since the only possible strategy seems to involve proceeding sequentially in time through the different temporal versions, which amounts in the worst case to at least the same amount of work as that required by the second straightforward approach.

We should note that special cases of our problem were addressed in [11] in the case of internal memory algorithms.

B. Computational Model

Before proceeding further, we introduce our computational model, which is more or less the standard I/O model used in the literature [12]. This model is defined

by the following parameters: n , the size of the input; M , the size of the internal memory; and B , the size of a disk block. An I/O operation is defined as the transfer of one block of contiguously stored data between disk and internal memory. Hence, scanning n contiguously stored items takes $O(n/B)$ I/O operations.

In the parallel domain, we use a particular version of the Parallel Disk Model (PDM), which is an extension of the above model. In the PDM model, we assume the presence of p processors, each with its own disk drive, such that the p processors are connected through a commodity interconnection network. Each I/O operation by a processor involves the transfer of a block of B consecutive words. For $p = 1$, we obtain the standard I/O model.

Our overall approach consists of building a time hierarchy of R-trees, such that each R-tree is of size $O(Nd/B)$, where N is the number of objects, d is the number of attributes characterizing each object and B is the block size of I/O transfer. Note that in our case the input size is $O(Ndm)$ and hence the input occupies $O(Ndm/B)$ blocks on the disk. Our R-tree is based on the Hilbert space filling curve and is used to handle the rectangle containment problem for Query 2-1 and standard multidimensional points for Queries 2-2 and 2-3. As we will see later, our version of the rectangle containment problem can be solved quite efficiently without having to traverse too many search paths in the R-tree. Each query can be handled by a parallel search of $O(\log w)$ such R-trees, where w is the length of the query time window.

We also present a very efficient solution for the special case when the number of objects is relatively small. The overall structure is an extension of the B-tree and achieves provably optimal bounds whenever $Nd = O(B)$, regardless of the temporal length of the time series.

II. OVERALL STRATEGY FOR TEMPORAL RANGE QUERYING

Our indexing scheme consists of a temporal hierarchical tree organization such that each node has a pointer to an R-tree structure that captures the minimum bounding boxes (in the d -dimensional space) of the objects during the corresponding time interval. Hence each R-tree can be used to handle an arbitrary range value query over the corresponding time interval. The overall organization is shown in Fig 1. The temporal hierarchical tree organization can be used to decompose an arbitrary query time interval into a (short) sequence

of the time segments represented in the tree. Hence this reduces the problem to searching a small number of R-trees, each of size $O(Nd/B)$ blocks. We will show how to index the objects so that the intersection of the candidate objects over these time segments can be obtained very quickly.

Fig. 1. Overall Strategy Based on the Temporal Hierarchy

A. Temporal Hierarchies

The timestamps of the objects can be used to group the objects into a tree hierarchy. In many applications, a natural hierarchy can be defined such as groupings by days, weeks, months and years. If no such natural hierarchy exists, we can use a balanced k -ary tree for some value of k that depends on the application. Once the hierarchy is set, any temporal range can be decomposed into a series of time segments (days, weeks, months, for instance) represented by the nodes in the tree hierarchy. Therefore, a temporal range query can then be answered by issuing a series of temporal range queries on the corresponding time segments. For the remainder of the paper, we take $k = 2$, in which case one can easily show that the number of nodes in the binary tree needed to accommodate any time window is logarithmic in the size of the time window. The query objects consist of the intersection or union of the results at the corresponding nodes, depending on the type of the query. Given any node in the tree hierarchy, the non-temporal data attributes of the objects can be aggregated over a time segment into an auxiliary structure attached to the node corresponding to that time segment. In our case, the minimum and maximum attribute values for each object over the time segment are computed and stored.

B. Rectangle Containment Query

For a fixed time segment, the attributes of each object lie within the rectangle defined by the minimum and the maximum values of each attribute over this

time segment. The handling of Query 2-1 comes down to determining the objects whose rectangles are fully contained in the query rectangle (defined by the range values). Given that we expect the number of objects to be large (with the number of attributes ranging typically between 1 and 20), the data corresponding to each time slice is typically too large to fit in main memory. We will represent the set of rectangles stored at each node of the time hierarchy by an R-tree that makes use of Hilbert space filling curves. Such R-trees seem to work well in practice even for a moderately high number of dimensions. We adapt such a strategy to our problem as follows.

Each leaf of the time hierarchy requires an R-tree to represent the corresponding N d -dimensional points, and hence this can be built using the standard Hilbert R-tree. Each interior node of the time hierarchy requires an R-tree that represents the corresponding N rectangles in d dimensions. Such an R-tree can be built by sorting the rectangles according to their Hilbert indices (properly defined), and proceeding in a sequence of hierarchical groupings as in the standard way to build a Hilbert R-tree. However we attach additional information at each node as follows.

We associate with each internal R-tree node a pointer to an ID list corresponding to the objects stored in its subtree. We make use of this ID list as follows. Suppose that, during the search, a minimum bounding rectangle in the R-Tree is found to be fully enclosed within the query rectangle. This means that all the objects contained inside this minimum bounding rectangle satisfy the query relative to the time window. Hence we can use the list of IDs to determine these objects, without having to proceed further down the R-tree as in the standard search process. In fact, based on the way we constructed the R-tree, we can create an overall list of IDs sorted by their Hilbert indices, and then represent each ID list at a node by two pointers to the overall list. All consecutive objects between these two pointers fall within the Minimum Bounding Rectangle (MBR) of that node. Since the size of object ID list is small, multiple consecutive accesses can be carried out through buffering in memory, which substantially improves our R-Tree search time, almost independent of the number of qualified objects.

C. Query Algorithm

We first focus on the algorithm to handle Query 2-1, assuming a single processor.

We start by determining the minimum number of nodes in the temporal hierarchy required to handle the query time window. Our time hierarchy is a balanced binary tree stored as a linear array in internal memory, and hence this step is quite simple and can be carried out very quickly. At the end of this step, we identify $O(\log m)$ nodes, each with a pointer to an R-tree.

Starting with the node representing the longest time segment in the time query window, we search the corresponding R-tree that was already augmented by the appropriate sorted ID list. The IDs of objects satisfying the query are returned. The same search process can be performed on the remaining nodes.

The output list of IDs is generated incrementally by intersecting the returned ID lists from the various R-tree searches.

We now consider the case when we have a cluster of $p = O(\log m)$ processors. The first step of determining the appropriate nodes of the temporal hierarchy is the same as before. We allocate the handling of the corresponding R-trees in a round robin fashion, assuming that the nodes are sorted by their heights (and hence by the lengths of their time intervals). Hence each processor will then handle its R-trees in the same way as before. The list of p ID lists are then merged easily on a single processor.

D. Time and Space Complexity

The space requirement of our overall data structure is clearly linear in the input size since we have $O(m)$ R-trees, each is built on N rectangles in d -dimensional space. Handling a query on a single processor will require the search through at $O(\log m)$ R-trees, each of size $O(Nd)$. For the case when we have a cluster of $p = O(\log m)$ processors, our I/O complexity reduces to $O(\log m/p)$ R-tree searches, and hence if $p = \Theta(\log m)$, the I/O complexity is the same as that for handling a single time slice, which is clearly the very best one can hope for. Notice that our strategy is still valid if we substitute a better data structure for handling the rectangle containment problem in d -dimensional space.

III. EFFICIENT STRATEGIES FOR A SMALL NUMBER OF OBJECTS

In this section, we consider the case when the number of objects is relatively small and provide a provably good solution to Query 2-1. Our solution uses a B+-tree, whose interior nodes have been augmented with auxiliary information. To illustrate the idea, we start with

Fig. 2. The minmaxB+-tree data structure. Each key in an interior node holds the minimum and the maximum values of each attribute of the corresponding child node.

the case when $Nd = O(B)$, where B is the block size. Let $k = \max\{2, \lfloor \frac{B}{2Nd} \rfloor\}$. We build a k -ary tree on the ordered time stamps such that each leaf contains the data corresponding to all the objects for k consecutive time stamps, all stored in at most two consecutive blocks. For an interior node, we store in addition to the key splitters, the minimum and maximum values of each attribute of each object, over the time interval corresponding to each child of the node. We call such a structure the minmaxB+-tree, which is illustrated in Fig. 2 for the case of a single object. It is easy to verify that each node is of size $O(B)$, the height of the tree is $O(\log_k m)$, and the overall size is linear.

The handling of a Query 2-1 proceeds as follows:

1. Determine the lowest common ancestor Q of the leaves of the minmaxB+-tree corresponding to the start and end time stamps of the query temporal window $[t_s, t_e]$. Note that this can be done in internal memory without any I/O accesses.

2. Let the i th and j th children of Q be those children that are on the paths from Q to t_s and t_e respectively. We begin by determining, for each key between i and j , the objects whose minimum and maximum values fall within the query range values. These objects are now potential candidates for our query. We proceed to check whether there are some objects that also satisfy the query bounds at i and j . These can be immediately declared as satisfying the query. We then proceed with the remaining candidates, from Q along the paths leading to t_s and t_e .

3. At each interior node along the paths to t_s and t_e , we refine the list of candidates by checking whether all the attributes' minimum and maximum values of the children that have time stamps greater than t_s or less than t_e lie within the query range values. If necessary, we repeat this step until we reach the leaf nodes.

4. At the two leaf nodes to which t_s and t_e belong, we complete the refinement of the list of the candidates thereby generating a list of the objects that satisfy the query.

Notice that the I/O complexity is proportional to the height of the tree, and hence it is of $O(\log_k m)$.

For larger values of N , we use a balanced binary tree built on the time stamps $t_1 < t_2 < \dots < t_m$ such that each interior node has the minimum and maximum values of each attribute of each object during the time interval represented by each child. These values are stored consecutively in $O(Nd/B)$ blocks. We proceed as before, except each step will take $O(Nd/B)$ I/O steps. Clearly this is only efficient for small values of N .

We will present experimental results comparing this approach to the simple approach of using the standard B+-tree for a single object (which clearly favors the latter approach). Even for this case, our approach is shown to be significantly superior especially for large query time windows. In the full paper, we will provide detailed experimental results comparing our minmaxB+-tree approach to our scheme for the general case that uses R-trees. Note that the minmaxB+-tree has provably optimal bounds whenever $Nd \leq cB$, for some constant c , and the main practical issue is to determine the highest value of c for which the approach of this section is superior. We estimate that this approach will indeed be superior whenever the number of objects is in the hundreds. Otherwise, our general approach is superior.

IV. EXPERIMENTAL RESULTS

The experimental results reported here are for a single processor and make use of synthetic data sets generated by a software package developed at the University of Maryland. We can specify N , d , m , and one of 4 possible distributions, and the software will return the corresponding data set. The processor used is a Sun Ultrasparc 359MHz, and each of the disks used is a 9GB disk, 7200RPM, with I/O throughput of 15MB/s. We start by presenting our main results of the general case, followed by a brief description of the case when we have a small number of objects.

A. The General Case

For the general case, we took $N = 10,000$ objects, each with a time series of length $m = 10,000$. The value of d was set to 2, 4, 8, and the distributions used were the Gaussian and the uniform distributions. Our time hierarchy is a balanced binary tree stored in an integer array of length 19,999. For a given input, the performance is measured in terms of the number of blocks interchanged between the disk and the internal memory, and in terms of the query time. Given that the performance will depend on the query time window and the query range values, we describe our results in two parts, the first focusing on the query performance when the size of the time window changes, and the second focusing on the query performance when the attribute range values change. We also compare our results to those obtained by incorporating the time dimension as an additional dimension using a single Hilbert R-tree, and to the sequential scan over the time window, where we use a Hilbert R-tree for each time instance. The comparison shows substantial gains achieved by our scheme over these two schemes, even for $p = 1$. With more processors, our scheme will achieve even stronger performance relative to the standard approaches.

1) Performance as a Function of Time Window Length: We conducted a series of tests based on queries with fixed spatial range values but increasing time windows from a single time slice to a time window of size 1000. The results on the 4-D dataset are shown in Fig. 3 through Fig. 5, using Gaussian distributions for the first two dimensions and uniform distributions for the remaining two. Note that each R-Tree takes 256 pages, each page of size 2K. In each case, we measure the number of accessed page blocks, query time, and the number of R-trees searched. For these graphs, the starting time used was 5001, but the results of other possible starting times are very similar.

These results clearly indicate that the number of pages accessed and the number of R-trees searched follow each a logarithmic pattern as a function of the query time window size. The relative variations from a logarithmic curve fit are small. Another immediate result is that the query time follows the same logarithmic pattern.

Fig. 6 and Fig. 7 are the performance results of using a single R-tree with the same series of queries on the same input. Notice that as the time window size increases, the query rectangle overlaps with more MBRs in the single R-tree, and hence the search algorithm explores many

Fig. 3. Number of Page Block Accesses vs. Query Time Window Size

Fig. 5. Number of Qualified Objects vs. Query Time Window Size

Fig. 4. Query Time vs. Query Time Window Size

Fig. 6. Number of Accessed Page Blocks vs. Query Time Window Size When Using a Single R-Tree

more paths in the R-tree. This observation is confirmed by the experimental results, which show a substantial increase in the number of pages accessed and in the query time as the temporal window size increases.

2) *Performance As a Function of the Query Range Values:* We now fix the temporal window size and change the spatial range values gradually until nearly all the objects are included. The variations of the spatial range values are expressed in terms of the percentage of the query range values rectangle out of the MBR that includes almost all the objects. The experimental results describing the number of accessed pages, query time, queried R-Trees, and number of qualified objects, for certain query settings, are summarized in the Fig. 8 through Fig. 11. Other averaged results are summarized in Table I .

Since the query time interval is fixed (at a size

of 100), the maximum number of the R-trees to be searched is also fixed (at most 9). However the search algorithm may explore fewer R-trees, since the search will stop if the R-tree corresponding to any node does not have any solutions. Notice that we start our search with the node covering the largest time segment, and hence we have a good probability of stopping early if there are no qualifying objects. We change the size of the range values rectangle from 40% to 100% of the MBR containing almost all the objects. The graph of the number of accessed page blocks is similar to the number of queried R-Trees except when the query rectangle gets larger and hence more objects get qualified in which case the number of page blocks decreases substantially while the number of queried R-trees increases slightly until it reaches its maximum value. This implies that, with larger query range rectangles, fewer nodes of the Hilbert R-

Fig. 7. Query Time vs. Query Time Window Size When Using a Single R-Tree

Fig. 9. Query Time vs. Query Range Rectangle as a Percentage of the Overall Bounding Rectangle

Fig. 8. Number of Accessed Page Blocks vs. Query Range Values Rectangle as a Percentage of the Overall Bounding Rectangle

Fig. 10. Number of Queried R-Trees vs. Query Range Rectangle as a Percentage of the Overall Bounding Rectangle

trees are accessed and no further exploration is needed since the MBRs of these nodes are fully contained in the query rectangle. This in particular illustrates the effectiveness of using the links to the ID lists. The query time is consistent with the number of page blocks accessed and not the number of R-trees queried as we expect, except for a few spikes when a large number of new blocks are being accessed, which mostly occurs due to additional R-Trees being queried. Because of disk buffering, once these blocks are accessed, further access to them is much faster. Hence, using our temporal hierarchy, further exploration of spatial attributes on the data with a fixed time window becomes extremely fast. Note however that a sequential access of time slices along the query temporal window cannot possibly exploit such disk locality for moderate to large window sizes and

neither can the scheme based on a single R-tree. In fact, the performance of a single R-tree scheme deteriorates significantly by increasing the query rectangle range values for a fixed window size.

B. The case of a Small Number of Objects

We now consider the case of a small number of objects. We present experimental results comparing the performance of our modified B+-tree to the traditional B+-tree for the case of a single object, which clearly favors the latter. Our test data consists of 200,000 time stamps and four attributes at each time stamp, where the values of each attribute are generated using a gaussian distribution. A sample of the data generated is shown in Figure 12. Both the original B+-tree and the new minmaxB+ tree were built on the same data set. The

Fig. 11. Number of Qualified Objects vs. Query Range Rectangle as a Percentage as a Percentage of the Overall Bounding Rectangle

Dimension	4		8	
Interval length	100	1000	100	1000
Ave. Query Time (ms)	160	275	200	350
Ave. Queried R-Trees	6	10	6	10

TABLE I

AVERAGE QUERY TIME AND NUMBER OF QUERIED R-TREES AT VARIOUS QUERY INTERVAL SITUATIONS UNDER RANGE CHANGING

block size was set to 8 KB and both versions of the B+-tree were of height 3.

The primary benefit of our proposed data structure is that query response time does not depend on the query time window size. Figure 13 shows the relationship between the query response time and query time window size. As seen in Fig. 13, the query time of the traditional B+-tree increases as the query time window increases. The same figure shows that the query response time of our proposed minmaxB+-tree is around 20 msec regardless of the size of the query time window. Note that the average number of blocks accessed using the minmaxB+-tree is between just 1 and 3 regardless of the query time window size, while the number of blocks accessed increases from 3 to 498 as the query time window increases when using the original B+-tree.

We should note that the size of our data structure and the construction time are almost the same as those for the B+ tree. This is due to the fact that we only added some extra information only to the nonleaf nodes which are typically a small part of the total indexing structure. For example, in this experiment, our indexing structure was larger by 38KB compared to the original B+-tree

Fig. 12. A sample synthetic time series data used in the experiment.

Fig. 13. Response time versus query time window size, with query range fixed at 80%. Both the original B+-tree and the minmaxB+-tree are built on 200,000 time stamps and have three levels.

that was of size 4MB.

In the full paper, we will present more detailed experimental results of our indexing structure as the number of objects increases up to a few hundreds. These results will illustrate the performance as a function of the query time window size and the size of the range values rectangle, and will be compared with the corresponding results for the general case.

V. OTHER TYPES OF TEMPORAL RANGE QUERIES

The previous sections have described in some detail the handling of the temporal range query 2-1. However the overall strategy can be applied to the other two types of queries mentioned in the introduction. We use the temporal hierarchical trees as before, with a pointer from each node to an R-tree. However we relax our earlier definition of the rectangles induced by the minimum and maximum values of each attribute over the corresponding time interval. Let's consider for example the handling of Query 2-1 in the case of noisy data.

Recall that each Hilbert R-tree corresponds to a fixed time interval, say $[t_1, t_2]$, and is built to represent a set of rectangles, one per object which is determined by the

minimum and maximum values of each of its attributes over $[t_1, t_2]$. During preprocessing, we can use good estimates of the noise for each dimension to shrink each rectangle accordingly along all the dimensions. We call the resulting rectangles the *core subrectangles* and use them to build the R-trees as before. The search process through each R-tree is the same as before. Clearly all the qualified objects will be in the final output list. However we may have a few additional objects that almost satisfy the query, but we expect the number of such additional objects to be extremely small assuming the availability of good noise estimates. We plan to undertake a detailed experimental study of such an approach to fully understand the overall performance on noisy data.

VI. CONCLUSION

We considered in this paper the problem of temporal range value querying of large scale multidimensional time series data. We developed a strategy that reduces the overall problem to handling $O(\log m)$ versions of a single time slice, which can be handled in parallel with almost no communication overhead. We reported on some of our extensive experimental results that clearly illustrate the superior performance of our strategy and its potential for enabling interactive querying of such data.

REFERENCES

- [1] V. Gaede and O. Günther, "Multidimensional access methods," *ACM Computing Surveys*, vol. 30, no. 2, pp. 170–231, 1998. [Online]. Available: citeseer.nj.nec.com/gaede97multidimensional.html
- [2] R. Janardan and M. Lopez, "Generalized intersection searching problems," *International Journal of Computational Geometry & Applications*, vol. 3, no. 1, pp. 39–69, 1993.
- [3] Q. Shi and J. JaJa, "Optimal and near-optimal algorithms for generalized intersection reporting on pointer machines," in *Technical Report CS-TR-4542, Institute for Advanced Computer Studies, University of Maryland*, 2003.
- [4] A. Guttman, "R-trees: a dynamic index structure for spatial searching," *ACMSIGMOD*, pp. 47–57, 1984.
- [5] T. K. Sellis, N. Roussopoulos, and C. Faloutsos, "The r+-tree: A dynamic index for multi-dimensional objects," in *The VLDB Journal*, 1987, pp. 507–518. [Online]. Available: citeseer.nj.nec.com/sellis87rtree.html
- [6] I. Kamel and C. Faloutsos, "Hilbert R-tree: An Improved R-tree using Fractals," in *Proceedings of the Twentieth International Conference on Very Large Databases*, Santiago, Chile, 1994, pp. 500–509. [Online]. Available: citeseer.nj.nec.com/kamel94hilbert.html
- [7] N. Beckmann, H.-P. Kriegel, R. Schneider, and B. Seeger, "The r*-tree: an efficient and robust access method for points and rectangles," in *Proceedings of the 1990 ACM SIGMOD international conference on Management of data*. ACM Press, 1990, pp. 322–331.
- [8] X. X. J. Han and W. Lu, "RT-tree: An Improved R-tree Index Structure for Spatiotemporal Databases," in *Proceedings of the 4th International Symposium on Spatial Data Handling*, 1990, pp. 1040–1049.
- [9] M. Nascimento and J. Silva, "Towards Historical Rtrees," in *Proceedings of ACM Symposium on Applied Computing (ACM-SAC)*, 1998, pp. 235–240.
- [10] Y. Tao and D. Papadias, "Efficient Historical R-trees," in *Proceedings of the 13th IEEE Conference on Scientific and Statistical Database Management(SSDBM)*, Fairfax Virginia, 2001, pp. 223–232. [Online]. Available: <http://www.cs.ust.hk/faculty/dimitris/PAPERS/ssdbm01.pdf>
- [11] Q. Shi and J. JaJa, "Fast Algorithms for a Class of Temporal Range Queries," in *Proceedings of Workshop on Algorithms and Data Structures*, 2003, pp. 91–102.
- [12] J. S. Vitter, "External memory algorithms and data structures: dealing with massive data," *ACM Computing Surveys*, vol. 33, no. 2, pp. 209–271, 2001. [Online]. Available: citeseer.nj.nec.com/vitter00external.html