PREPRINT. In Proceedings of the 4th Workshop on Algorithms Data Structures, volume 955, pages 473-481. Springer Verlag, 1995.

On the Difficulty of Range Searching

Arne Andersson* Kurt Swanson*

Dept. of Computer Science, Lund University,
Box 118, S-221 00 LUND, Sweden

Abstract. The problem of range searching is fundamental and well studied, and a large number of solutions have been suggested in the literature. The only existing non-trivial lower bound that closely matches known upper bounds with respect to time/space tradeoff is given for the pointer machine model. However, the pointer machine prohibits a number of possible and natural operations, such as the use of arrays and bit manipulation. In particular, such operations have proven useful in some special cases such as one-dimensional and rectilinear queries.

In this article, we consider the general problem of (2-dimensional) range reporting allowing arbitrarily convex queries. We show that using a traditional approach, even when incorporating techniques like those used in fusion trees, a (poly-) logarithmic query time can not be achieved unless more than linear space is used. Our arguments are based on a new non-trivial lower bound in a model of computation which, in contrast to the pointer machine model, allows for the use of arrays and bit manipulation. The crucial property of our model, *Layered Partitions*, is that it can be used to describe all known algorithms for processing range queries, as well as many other data structures used to represent multi-dimensional data.

We show that $\Omega\left(\frac{\log n}{\log T(n)}\right)$ partitions must be used to allow queries in $O(T(n)+k)$ time, where k is the number of reported elements, for any growing function $T(n)$. In some special cases, as for rectilinear queries, these partitions may be stored in compressed form, which has been exploited by the M-structure of Chazelle. However, so far there has been no indication that such compression would be feasible in the general case, in which case any algorithm based on our model, and supporting range searching in $O(\log^c n + k)$ time requires $\Omega\left(\frac{n \log n}{\log\log n}\right)$ space. (Note that it may be possible to obtain a better upper bound with an algorithm not adhering to the model of layered partitions.) Hence, we show that removing the restrictions of the pointer machine model does not help in obtaining a significantly improved time/space tradeoff — any solution based on traditional representations of point sets cannot combine linear space and polylogarithmic time.

* email: {arne, kurt}@dna.lth.se

1 Introduction

1.1 Discrepancy between upper and lower bound

The complexity of range searching has been a long standing open question. The problem is fundamental and easy to formulate: given a set of points in a multidimensional space, create a data structure that facilitates reporting of all points inside a given query region. This formulation is frequently called range reporting. Although many attempts have been made, the time-space tradeoff for this problem is still unclear.

Taking a general approach, we note the following trivial lower bounds: Let $F(n)$ be the time required to find *one* point in the query region—or to discover that the region is empty—in a set of n points. Then, an immediate lower bound on time complexity is $\Omega(F(n) + k)$, where k is the number of points to be reported. The same lower bound applies when $F(n)$ is the cost of counting the number of points in the region or computing the weighted sum of a set of weighted points. As an example, by a modification of the proof by Miltersen [13], a lower bound of $F(n) = \Omega(\log^{1/3-o(1)} n)$ can be obtained [12]. (The original bound given in that article is expressed in terms of the size of the universe. This is also the case for similar lower bounds on *existential range searching*, which reports whether or not a given region is empty, given by Miltersen et al [14].) In the same way, the arithmetic lower bound by Chazelle [7] for the problem of *dominance searching* (where the sum of weights of all points dominated by a given point is computed), applies to our problem. However, these bounds are far from tight with respect to the time/space complexity of range reporting.

The only existing non-trivial lower bound that closely matches known upper bounds with respect to time/space tradeoff is given for the *pointer machine* model, where $F(n) = \Omega(\log n)$ even in the one-dimensional case. In this model, it has been shown [6] that $\Theta(n \log n / \log\log n)$ space is necessary and sufficient in order to achieve optimal query time complexity. In fact, not even polylogarithmic time can be achieved with less space.

However, the restrictions of the pointer machine model are not realistic. This weakness is particularly evident for range reporting since it has been explicitly demonstrated that the use of arrays and bit manipulation can help. In such a more general, and more realistic, model we can expect faster searching than $\Omega(\log n)$ [9]. As another possible improvement, we may reduce space requirements by packing information about more than one point into one machine word. This option of storing point sets in compressed form has been utilized by Chazelle [5] in the special case of *rectilinear* range queries; a query cost of $O(\log^4 n + k)$ can be achieved using only $O(n)$ space, and a query cost of $O(\log n + k)$ can be achieved using $O(n \log^\epsilon n)$ space on the RAM model. But what about the general problem?

1.2 A lower bound in a relevant model

We concentrate on two dimensional range reporting with convex query regions. It should be noted that our lower bound only holds for a (relevant) class of

algorithms to solve range searching, and not the problem itself.

In examining the large set of known data structures, one finds that they all have one property in common: they may all be viewed as representing one or more partitions of the plane, where each partition divides the plane into $\Theta(n)$ convex areas. When a query is made, the answer is given by intersecting the query region with a selected subset of these partitions.

Based on this observation, we define the model of *Layered Partitions*, which can be used to emulate all known solutions to the range searching problem, as well as many other data structures used in computational geometry. Our lower bound on the time-space tradeoff is given in terms of the number of partitions needed in order to achieve a certain query cost. We show that, in order to support queries in time $O(T(n)+k)$, $\Omega(\log n/\log T(n))$ partitions must be represented. Hence, if each partition does require $\Theta(n)$ space (and there is no evidence pointing to the contrary for the general case of convex query regions), then any algorithm based on layered partitions, which supports range reporting in $O(\log^c n + k)$, $c = O(1)$, time requires $\Omega\left(\frac{n \log n}{\log \log n}\right)$ space.

Our model of computation is general in the sense that it allows all kinds of bit-manipulation techniques, such as fusion trees [9], to be used in order to speed up queries, thus avoiding inherent weaknesses in the pointer machine model.

2 Computational model

Our computational model is based on a data structure paradigm which can be used to describe all data structures for the range reporting problem found in the literature, as well as many other data structures used to represent multidimensional data.

A central part of our model is the *partition*. Given a set of points in the plane, we define the "universe" as the smallest enclosing rectangle containing all points in the set. A partition divides the universe in $\Theta(n)$ convex regions, each region contains at most one point.

A data structure in the model represents a set of partitions $P_1, P_2, \ldots$ in the plane. Range queries are processed in the following way:

1. Split the query region into m subregions $R_1, \ldots, R_m$.
2. Associate each R_i with a partition $P(R_i)$.
3. Examine all regions in $P(R_i)$ that intersect R_i and report which points are contained in R_i.

The cost of processing the query is defined as:

$$\sum_{i=1}^{m} \text{number of regions in } P(R_i) \text{ that intersect } R_i$$

Note that we only consider the time required to access information in partitions, and not the time needed to determine how to divide the query region into

subregions, nor the time to determine which partitions to use, nor time spent searching in any ancillary data. Thus differences between the RAM and decision tree models are negated.

We claim that this model covers the classical data structures used to solve this, and similar, problems. Among others, the following data structures can be described as layered partitions: k-d-Tree [15], Multistage direct access (multilevel k-ranges) [2], Filtering search [3], Range Trees [1] (see also [16] and [10]), Quad trees [8], Priority search trees [11], Voronoi diagrams, and M-structures [5] (see also [4]).

As an example, we indicate how to describe, in terms of layered partitions, priority search trees as well as the data structure used in filtering search. In Figure 1, we illustrate how to view a priority search tree as a layered partition (a single layer). In a priority search tree, the point with highest y-coordinate is stored in the root. The rest of the points are divided between the two subtrees according to a *split value* which is also stored in the root. All points whose x-coordinate is less than the split value are stored in the left subtree, the other points are stored in the right subtree. A priority search tree supports range queries where the query region is a rectangle for which the topmost edge in the rectangle is located above all stored points.

In the figure, the universe covers $([0;20],[0;10])$ and the query region (shaded) is $([8;18],[6;10])$. In each tree node, the upper values are split values (not needed for the leaves) and the lower values are point coordinates. The partition created by the tree is illustrated below the tree; each node corresponds to one region. Vertical segments represent split values, and horizontal segments separate each node's point from those of its subtree. The horizontal segments are somewhat arbitrarily drawn, but conform to the nature of priority search trees as well as layered partitions. This particular query rectangle contains five regions, the corresponding nodes are shaded.

In filtering search, the data structure is organized in $\Theta(\log n/\log\log n)$ layers, each layer may be described as consisting of three partitions. At the ith layer, we have the following three partitions:

- The first partition divides the universe into $\Theta(\log^i n)$ vertical segments, each segment containing $\Theta(n/\log^i n)$ points. Each segment is in turn divided into horizontal segments, each segment containing one point.
- The second partition divides the universe into the same $\Theta(\log^i n)$ vertical segments. This time, however, each segment is represented as a priority search tree, based on the left edge of the segment. A priority search tree may in turn be described as a partition as previously shown.
- The third partition is similar to the second one, with the difference that the priority search trees are based on the right edge of their segments.

3 Lower bound

Our lower bound is based upon a specific example. We describe a simple layout pattern for points in the plane, and prove that, according to the above described

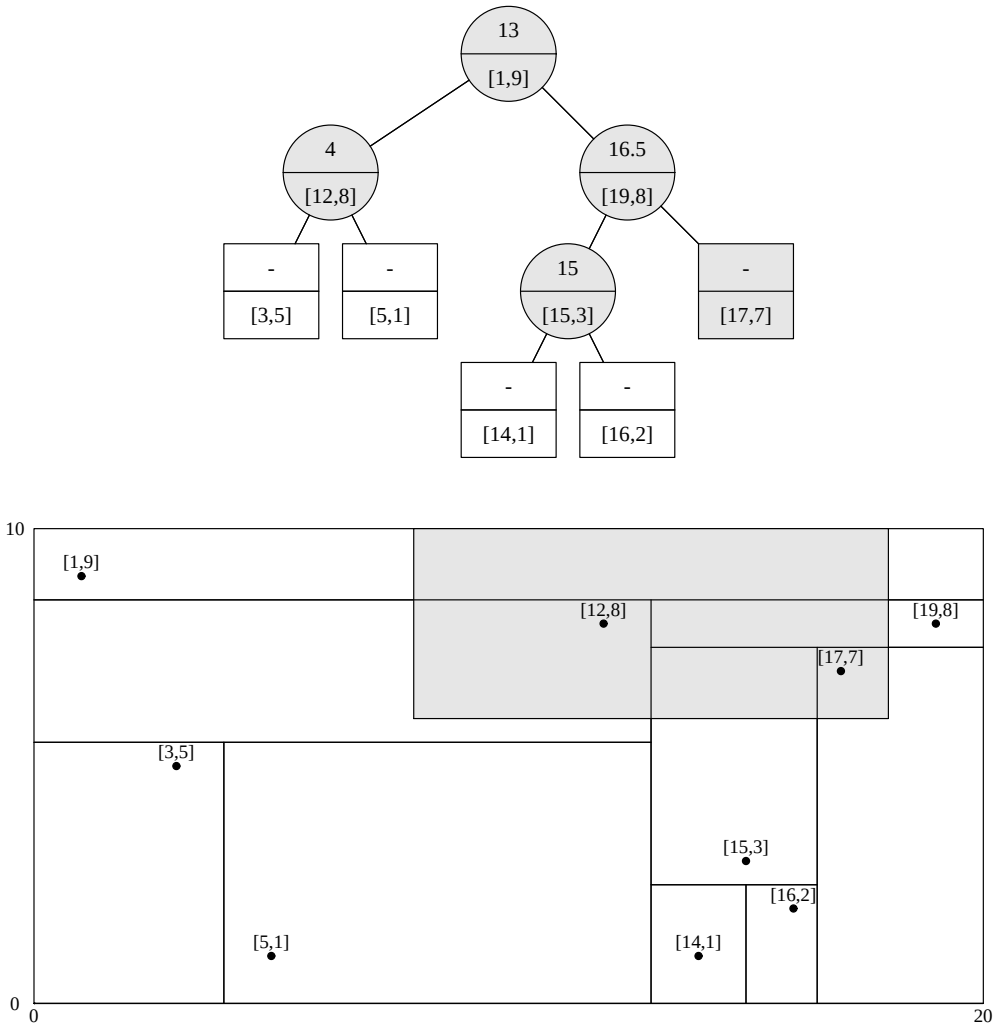

Fig. 1. Searching in a priority search tree

model of computation, any algorithm that solves range searching in $O(T(n)+k)$ time, must represent $\Omega\left(\frac{\log n}{\log T(n)}\right)$ partitions.

For the sake of simplicity, we shall describe this layout in terms of rectangular queries. However, our lower bound construction can easily be extended to a set of convex non-rectangular queries. We indicate how to perform this extension below.

We arrange the point set to create classes of empty rectangles $R_1, R_2, R_3, ...$ We denote the entire set of rectangles as $R = R_1 \cup R_2 \cup R_3 \cup ...$ The class R_i is constructed recursively in the following way (see Figure 2):

1. Initially $i = 1$ and $R = R_1$ which contains one rectangle.
2. Set $i = i + 1$.
3. Make $T^3(n)$ copies of the entire set of rectangles R and place them on a horizontal line. This implies that the number of rectangles of class $R_j, (\forall j \mid 1 \leq j < i)$ increase by a factor of $T^3(n)$.
4. Add a class of rectangles R_i as $T^{3(i-1)}(n)$ long rectangles, evenly spaced over the universe from top to bottom, each one intersecting all copies made in the previous step (as shown in Figure 2).
5. Make the holes (rectangular regions between intersecting rectangles) small enough so that their total area is negligible. One point is placed in each hole.
6. If there are points left, go to 2.

We assume, without loss of generality, that the number of points, n, exactly fills the last class, without any remaining points. (We note that this provides for an infinite number of constructions).

We now indicate how this construction can be rephrased for a set of convex and non-rectangular queries. To do this, we observe that we have left small gaps between all rectangles. Hence, instead of placing our points on strict lines, we may arrange them in a slightly more irregular pattern, creating convex, but not necessarily rectangular, empty query regions.

Lemma 1. *The number of classes of rectangles that can be created by the above described process is* $\Omega\left(\frac{\log n}{\log T(n)}\right)$, *given* n *points.*

Proof. Let $S(i)$ equal the number of points required to create i classes of rectangles. When creating the ith class, we create $T^3(n)$ copies of everything that we had before, the number of points needed for this purpose is $S(i-1)T^3(n)$. Next, we add $T^{3(i-1)}(n)$ rectangles. Each of these rectangles crosses $T^3(n)$ rectangles in R_{i-1}, creating $(T^{3(i-1)}(n)-1)(T^3(n)-1)$ new holes for level R_i, and we add one point in each hole, which gives us:

$$S(i) = S(i-1)T^3(n) + (T^{3(i-1)}(n)-1)(T^3(n)-1) \tag{1}$$

We first prove that $S(i) < T^{4i}(n)$ by induction. As $S(1) = 0$, then $S(2) = (T^{3\cdot 1}(n)-1)(T^3(n)-1) < T^{4\cdot 2}(n)$. By induction assume $S(i-1) < T^{4(i-1)}(n)$, then, for $T(n) > 1$ and $i > 2$:

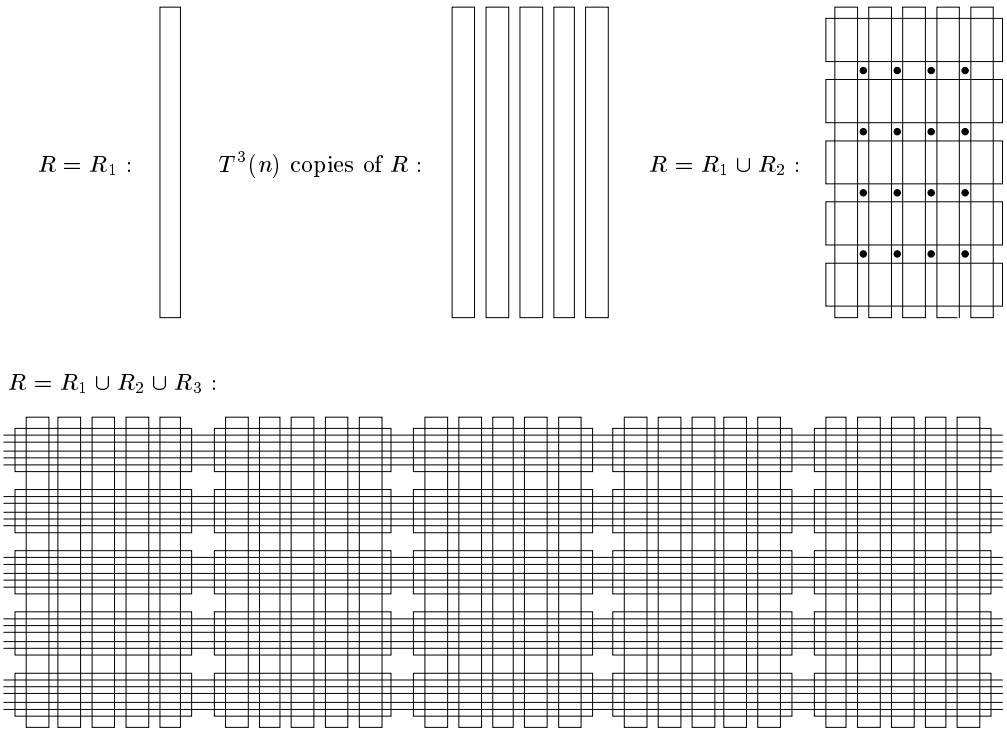

Fig. 2. The growth of R, for $T^3(n) = 5$, yielding 16 holes for $R = R_1 \cup R_2$, and 176 total holes for $R = R_1 \cup R_2 \cup R_3$. (The 25 rectangles in R_3 are shown as lines. Only points in R_2 are drawn.)

$$S(i) < T^{4(i-1)}(n)T^3(n) + (T^{3(i-1)}(n) - 1)(T^3(n) - 1) \tag{2}$$
$$< T^{4i-1}(n) + T^{3i}(n) - T^{3(i-1)}(n) - T^3(n) + 1 < T^{4i}(n) \tag{3}$$

We now set $n = S(i)$ and solve for i:

$$n = S(i) < T^{4i}(n) \tag{4}$$
$$\log n < 4i \log T(n) \tag{5}$$
$$\frac{\log n}{\log T(n)} < 4i \tag{6}$$
$$i = \Omega\left(\frac{\log n}{\log T(n)}\right) \tag{7}$$

Next, we study how many partitions must be maintained in order to perform range queries efficiently. Intuitively, our goal is to show that there must be at least as many partitions as there are classes of empty rectangles.

We say that a rectangle used in a layered partition is *q-sized* if both its height and width is at least $1/T(n)$ times the height and width of a rectangle in R_q.

Lemma 2. *A q-sized rectangle used in a layered partition data structure is not s-sized if $q \neq s$.*

Proof. The proof follows from the way our R_q's are constructed and the fact that no rectangle used in a partition may contain more than one point.

Without loss of generality, assume $q < s$. The height of a q-sized rectangle is greater than $1/T(n)$ times the height of a rectangle in R_q, by definition. The height of a q-sized rectangle is greater than $T^2(n)$ times the height of a rectangle in R_s due to the recursive construction of R_q and R_s.

Assume a q-sized rectangle used in a layered partition is at least $1/T(n)$ times the width of a rectangle in R_s. Then the q-sized rectangle must cross over at least $T^2(n)$ rectangles in R_{s-1}. By construction, there are $\Omega(T^3(n))$ points between two rectangles in R_q. Therefore, due to its height, the q-sized rectangle must enclose at least $T^2(n)$ points in between each pair of rectangles it crosses in R_{s-1}. Thus the q-sized rectangle contains at least $T^2(n)(T^2(n)-1) = T^4(n) - T^2(n)$ points. It can thereby not be used in a layered partition. This contradicts the assumption on its width, and thus, by definition, it cannot be s-sized.

Lemma 3. *For any solution to the range query problem, the following must hold: for each q, all but at most $1/T(n)$-th of the universe must be covered by q-sized rectangles represented by the solution.*

Proof. Assume that we chose to use an arbitrary rectangle in R_q as a query region. Then, in order to perform our query in time $T(n)$ our data structure must contain rectangles such that any rectangle in R_q can be covered by at most $T(n)$ rectangles.

By the construction of the example, all but a negligible part of the universe is covered by each set R_q. It thus suffices to show that for each rectangle in R_q, only $1/T(n)$-th of the rectangle may not be covered by q-sized rectangles in the solution. From Lemma 2 and the definition of q-sized, it follows that any s-sized rectangle, $(s \neq q)$, covers less than $\frac{1}{T^3(n)}$ of the rectangle. Thus, $\Theta(T(n))$ non-q-sized rectangles cover $O(\frac{1}{T^2(n)})$ of the rectangle. Therefore, the remainder of the rectangle must be covered by q-sized rectangles. If not, the cost of a query would exceed $O(T(n))$.

Theorem 4. *In order to search in $O(T(n)+k)$ time, where k is the number of elements in the query region, $\Omega\left(\frac{\log n}{\log T(n)}\right)$ partitions are required.*

Proof. It follows from lemma 3 that we need to store at least one partition for each separate class of rectangles in the construction, yielding $\Omega\left(\frac{\log n}{\log T(n)}\right)$ partitions.

The theorem yields the following corollary.

Corollary 5. *If each partition requires $\Theta(n)$ space, any algorithm based on layered partitions and supporting range reporting in $O(\log^c n + k)$, $c = O(1)$, time requires $\Omega\left(\frac{n \log n}{\log \log n}\right)$ space.*

4 Conclusion

We feel that our new lower bound provides new insight on the difficulty of range searching. Our model of computation, layered partitions, captures the inherent properties of a large class of data structures. In this model, we have shown that it is not possible to perform range queries in $O(\log^c n + k)$ time and linear space, unless partitions can be stored in compressed form. However, it seems infeasible to be able to apply the bit-encoding technique of M-structures [5] to other than orthogonal rectilinear queries, since the compression heavily exploits the fact that the queries are rectangular. When compression is not feasible, any algorithm based on layered partitions and supporting range reporting in $O(\log^c n + k)$, $c = O(1)$, time requires $\Omega\left(\frac{n \log n}{\log \log n}\right)$ space. Thus we show that no superior upper bound can be achieved using traditional methods and data structures for range searching in the general case.

One might also imagine that a possible way to achieve a better tradeoff between space and time would be to combine some classical data structure, such as range trees, with some sophisticated search method, such as the one used in fusion trees [9]. The number of possibilities seems to be very large and we believe that many researchers have tried methods like this. In this article, we have shown that such an approach would not be fruitful.

Acknowledgments

We would like to thank the referees and Dr. Ola Petersson for many insightful comments.

References

1. J. L. Bentley. Decomposable searching problems. *Inform. Process. Lett.*, 8:244–251, 1979.
2. J. L. Bentley and H. A. Maurer. Efficient worst-case data structures for range searching. *Acta Inform.*, 13:155–168, 1980.
3. B. Chazelle. Filtering search: a new approach to query-answering. In *Proc.* 24th *Annu. IEEE Sympos. Found. Comput. Sci.*, pages 122–132, 1983.
4. B. Chazelle. Slimming down search structures: A functional approach to algorithm design. In *Proc.* 26th *Annu. IEEE Sympos. Found. Comput. Sci.*, pages 165–174, 1985.

5. B. Chazelle. A functional approach to data structures and its use in multidimensional searching. *SIAM J. Comput.*, 17:427–462, 1988.
6. B. Chazelle. Lower bounds for orthogonal range searching, I: the reporting case. *J. ACM*, 37:200–212, 1990.
7. B. Chazelle. Lower bounds for orthogonal range searching: II. the arithmetic model. *Journal of the ACM*, 37(3):439–463, July 1990.
8. R. A. Finkel and J. L. Bentley. Quad trees: a data structure for retrieval on composite keys. *Acta Inform.*, 4:1–9, 1974.
9. M. L. Fredman and D. E. Willard. Blasting through the information theoretic barrier with fusion trees. In *Proc.* 22nd *Annu. ACM Sympos. Theory Comput.*, pages 1–7, 1990.
10. G. S. Lueker. A data structure for orthogonal range queries. In *Proc.* 19th *Annu. IEEE Sympos. Found. Comput. Sci.*, pages 28–34, 1978.
11. E. M. McCreight. Priority search trees. *SIAM J. Comput.*, 14:257–276, 1985.
12. P. B. Miltersen. Personal communication.
13. P. B. Miltersen. Lower bounds for union-split-find related problems on random access machines. In *Proc.* 26th *Ann. ACM STOC*, pages 625–634, 1994.
14. P. B. Miltersen, N. Nisan, S. Safra, and A. Wigderson. On data structures and asymmetric communication complexity. In *Proc.* 27th *Annu. ACM Sympos. Theory Comput.*, 1995. To appear.
15. J. B. Saxe and J. L. Bentley. Transforming static data structures to dynamic structures. In *Proc.* 20th *Annu. IEEE Sympos. Found. Comput. Sci.*, pages 148–168, 1979.
16. D. E. Willard. A new time complexity for orthogonal range queries. In *Proc.* 20th *Allerton Conf. Commun. Control Comput.*, pages 462–471, 1982.

This article was processed using the LaTeX macro package with LLNCS style