Translating a Regular Grid over a Point Set*

Prosenjit Bose† Marc van Kreveld‡ Anil Maheshwari†
Pat Morin§ Jason Morrison†

Abstract

We consider the problem of translating a (finite or infinite) square grid G over a set S of n points in the plane in order to maximize some objective function. We say that a grid cell is k-occupied if it contains k or more points of S. The main set of problems we study have to do with translating an infinite grid so that the number of k-occupied cells is maximized or minimized. For these problems we obtain running times of the form $O(kn \text{ polylog } n)$. We also consider the problem of translating a finite size grid, with m cells, in order to maximize the number of k-occupied cells. Here we obtain a running time of the form $O(knm \text{ polylog } nm)$.

1 Introduction

In this paper we consider the problem of translating a (finite or infinite) square grid G over a set S of n points in the plane so that some objective function is optimized. For a grid cell C of G we say that C is *k-occupied* if C contains at least k points of S. The *k-occupancy* of G is the number of cells of G that are k-occupied. A cell is simply *occupied* if it is 1-occupied and the *occupancy* denotes the 1-occupancy.

For infinite grids, we consider both minimization and maximization versions of the following problems where G is free to be translated by any vector and k is part of the input:

Problem 1. *Find a translation of G that optimizes the occupancy of G.*

Problem 2. *Find a translation of G that optimizes the k-occupancy of G.*

*This work was partly funded by the Natural Sciences and Engineering Research Council of Canada and the Dutch Organization for Scientific Research (N.W.O).

†School of Computer Science, Carleton University, 1125 Colonel By Drive, Ottawa, Canada, K1S 5B6, {jit,maheshwari,morrison}@scs.carleton.ca .

‡Department of Computer Science, Utrecht University, P.O. Box 80.089, 3508 TB, Utrecht, The Netherlands, marc@cs.uu.nl .

§School of Computer Science, McGill University, 3480 University Street, McConnell Engineering Building., Room 318, Montreal, Quebec, Canada, H3A 2A7, morin@cgm.cs.mcgill.ca .

Problem 3. *Find a translation of G (if one exists) so that no cell of G is k-occupied.*

We also consider the following problems of optimizing k:

Problem 4. *Find the smallest integer k such that there exists a translation of G with k-occupancy 0.*

Problem 5. *Find the largest integer k such that there exists a translation of G with k-occupancy at least r.*

For a finite grid G and a fixed integer k we also consider the following problem:

Problem 6. *Find a translation of G that maximizes k-occupancy.*

This problem arises in several applications. For example, a problem from cartography is that of laying out maps into books, where one wants to avoid book pages with no features or with too many features. In the problem of gridding/interpolation one wants to derive a regular grid from irregular data. In order to minimize the interpolation error one would like to ensure that every grid cell contains at least 1 point to interpolate from. Finally, in the design of hash functions [2, 5], one wants to minimize the maximum number of points contained in any bucket.

The main tools we use to solve these problems are plane sweep [4] and a dynamic data structure for maintaining the maximum clique in an interval graph under insertions and deletions of intervals. Our approach is to reduce the grid translation problem to that of finding a maximally (or minimally) covered point in an arrangement of rectangles and then apply plane sweep and our data structure to find this point.

The previous work most closely related to our results is the work of Lee [9] on finding the maximum clique in a rectangle intersection graph. A rectangle intersection graph $G = (V, E)$ has vertices V, $|V| = n$, such that each vertex represents one of n iso-oriented rectangles. Every edge in E corresponds to a pair of rectangles that intersect and all such intesections are represented in E. Lee also uses plane sweep along with a data structure for maintaining the *size* of a maximum clique in an interval graph. The data structure is based on segment trees [3] and has $O(\log n)$ time for insertions, deletions and queries.

As we were previously unaware of Lee's work, we developed a data structure based on prefix-sum trees that also has $O(\log n)$ running time for all operations. This data structure seems to be more flexible than the segment tree data structure. For example, it is fully dynamic, not requiring that the endpoints of the segments be known in advance. It can also report a value that is contained in the largest number of intervals. It is unclear whether the segment tree data structure can be extended to do either of these without increasing the running time. We conjecture that this data structure will find applications beyond those in this paper.

Several other researchers have considered grid placement problems in which rotations and resizing of the grid cells can take place [1, 2]. In these applications

Problem	Running Time	Memory
Problem 1	$O(n \log n)$	$O(n)$
Problem 2	$O(kn \log n)$	$O(kn)$
Problem 3	$O(kn \log n)$	$O(kn)$
Problem 4	$O(kn \log n \log k)$	$O(kn)$
Problem 5	$O(kn \log n \log k)$	$O(kn)$
Problem 6	$O(kmn \log mn)$	$O(kn)$

Table 1: Summary of results.

the objective function that is minimized is usually the maximum number of points that fall in any given cell. Because of the increased number of degrees of freedom, the running times of the algorithms are quite high (usually close to quadratic in n).

In comparison, we study a constrained version of the problem where only rigid translation is allowed. With this constraint we are able to consider variations with k-occupancy rather than just 1-occupancy. We obtain the low-order running times and memory requirements shown in Table 1. All results are valid for grid cells of arbitrary aspect ratio and scale. When we do use the idea of square unit size cells in our descriptions it is only for the sake of clarity and no results are affected by the change.

The remainder of this paper is organized as follows: In Section 2 we describe our new data structure for maintaining the maximum clique in an interval graph. In Section 3 we show how to use this data structure to solve Problems 1–6. In Section 4 we summarize and conclude with open problems.

2 The Data Structure

In this section we describe a data structure that maintains a function $f : \mathbb{R} \to \mathbb{Z}$ under the following update and query operations.

1. INSERT(T, a, b): Increase the value of $f(x)$ by 1 for all $x \in [a, b)$.
2. DELETE(T, a, b): Decrease the value of $f(x)$ by 1 for all $x \in [a, b)$.
3. MAX-COVER(): Return $\max\{f(x) : x \in \mathbb{R}\}$.
4. MAX-COVER-WITNESS(): Return a value x^* such that $f(x^*) = \max\{f(x) : x \in \mathbb{R}\}$.

Furthermore, a simple modification of this data structure answers the minimization queries MIN-COVER() and MIN-COVER-WITNESS() without any difficulties.

The data structure we use is a binary search tree T whose root is denoted by r. Each node v of T stores an interval. We denote by INT(v) the interval

stored at v, and by $\textsc{start}(v)$, respectively $\textsc{stop}(v)$, the left, respectively right, endpoint of $\textsc{int}(v)$. For an internal node v, the left, respectively right, child of v is denoted by $\textsc{left}(v)$, respectively $\textsc{right}(v)$. The interval of an internal node v is always $\textsc{int}(v) = \textsc{int}(\textsc{left}(v)) \cup \textsc{int}(\textsc{right}(v))$. Initially, T contains one node r with $\textsc{int}(r) = [-\infty, +\infty)$.

In addition to this information, each node v of T also maintains two values, $\textsc{stab}(v)$ and $\Delta(v)$. Let $K(v) = \max\{f(x) : x \in \textsc{int}(v)\}$ so that $K(r) = \textsc{Max-Cover}()$. Intuitively, we would like value of $\textsc{stab}(v)$ to be $K(v)$. However, maintaining this property would take linear time for each Insert and Delete operation. Instead, we maintain the following less strict invariant.

Invariant 1. *For every node v of T, $K(v) = \textsc{stab}(v) + \sum_{u \in P(v)} \Delta(v)$, where $P(v)$ is the set of nodes (including v) on the path from v to the root of T.*

Invariant 1 has several implications. One of these is that an answer to Max-Cover queries can be given in constant time since $K(r) = \textsc{stab}(r) + \Delta(r)$. Furthermore, for a node v, if $\Delta(u) = 0$ for all nodes u on the path from v to r then $\textsc{stab}(v) = K(v)$.

Our strategy for maintaining the stab information is a lazy one. We try to avoid updating the values of nodes whenever possible. Van Kreveld and Overmars [10] use a similar mechanism in the design of concatenable segment trees. When it becomes necessary, the procedure that we use to update values at internal nodes of T is $\textsc{Push-Delta}(v)$, defined as follows.

$\textsc{Push-Delta}(v)$

1: $\textsc{stab}(v) \leftarrow \textsc{stab}(v) + \Delta(v)$
2: $\Delta(\textsc{left}(v)) \leftarrow \Delta(\textsc{left}(v)) + \Delta(v)$
3: $\Delta(\textsc{right}(v)) \leftarrow \Delta(\textsc{right}(v)) + \Delta(v)$
4: $\Delta(v) \leftarrow 0$

It is clear that, if Invariant 1 holds for node v and its two children before calling $\textsc{Push-Delta}(v)$, then it continues to hold after calling $\textsc{Push-Delta}(v)$.

2.1 Query Operations

Because of Invariant 1, the $\textsc{Max-Cover}()$ operation can be implemented trivially in constant time since $\textsc{Max-Cover}() = K(r) = \textsc{stab}(r) + \Delta(r)$. The procedure $\textsc{Max-Cover-Witness}()$ requires a little more care, and is implemented by the following algorithm.

```
MAX-COVER-WITNESS()
1: v ← r
2: while v is not a leaf do
3:    PUSH-DELTA(v)
4:    if STAB(LEFT(v)) + Δ(LEFT(v)) > STAB(RIGHT(v)) + Δ(RIGHT(v)) then
5:       v ← LEFT(v)
6:    else
7:       v ← RIGHT(v)
8:    end if
9: end while
10: report any x* ∈ INT(v)
```

To see that this procedure is correct note that, before line 4 is executed, $\Delta(u) = 0$ for all $u \in P(v)$. Therefore, by Invariant 1, $K(u) = \textsc{stab}(u) + \Delta(u)$ for each child u of v. So, during each iteration, v is reset to the child u of v that maximizes $K(u)$. It follows that the algorithm reports a point in $\textsc{int}(v)$ where v is a leaf such that all values in $\textsc{int}(v)$ maximize f. Since the only modifications to T are made by calls to PUSH-DELTA, which maintain Invariant 1, Invariant 1 is still true for all nodes after calling MAX-COVER-WITNESS().

2.2 Update Operations

Insertion and deletion operations are done using one general operation, INSERT-INFINITE-INTERVAL(a, δ), where a is a real number and δ is an integer. INSERT-INFINITE-INTERVAL(a, δ) has the effect of increasing $f(x)$ by δ for all $x \in [a, +\infty)$. It should be clear that INSERT(a, b) and DELETE(a, b) can each be implemented with two calls to INSERT-INFINITE-INTERVAL (using values of $\delta = \pm 1$).

We implement INSERT-INFINITE-INTERVAL by first creating a new leaf w with $\textsc{start}(w) = a$ if T does not already contain such a leaf. We then traverse $P(w)$ from r to w calling PUSH-DELTA for each node we encounter. At the same time, we update the Δ values for nodes that lie to the right of $P(w)$. After doing this, Invariant 1 is maintained for all nodes except those on $P(w)$. A final traversal from w back up to r corrects the STAB and Δ values for all nodes of $P(w)$ thereby restoring Invariant 1.

These ideas are implemented by the following pseudocode.

```
INSERT-INFINITE-INTERVAL(a, δ)
 1: BASIC-INSERT(a)
 2: v ← r
 3: while v is not a leaf do
 4:    PUSH-DELTA(v)
 5:    if a ≥ START(RIGHT(v)) then
 6:       v ← RIGHT(v)
 7:    else
 8:       Δ(RIGHT(v)) ← Δ(RIGHT(v)) + δ
 9:       v ← LEFT(v)
10:    end if
11: end while
12: STAB(v) ← STAB(v) + δ
13: repeat
14:    v ← PARENT(v)
15:    STAB(v) ← max{STAB(LEFT(v)) + Δ(LEFT(v)), STAB(RIGHT(v)) +
       Δ(RIGHT(v))}
16: until v = r
```

We are now considering two different versions of the function f. To avoid confusion, we let f_{a} denote the function f before the operation and f_{b} denote the function f after the operation. Similarly, we refer to Invariant 1.a as Invariant 1 using the function $f = f_{\mathrm{a}}$ and Invariant 1.b as Invariant 1 using $f = f_{\mathrm{b}}$. Finally, $K_{\mathrm{a}}(v)$ and $K_{\mathrm{b}}(v)$ refer to the function $K(v)$ using $f = f_{\mathrm{a}}$ and $f = f_{\mathrm{b}}$, respectively.

The function BASIC-INSERT(a) creates a node w with START$(w) = a$ if one does not already exist and rebalances T. How this is accomplished is the topic of the next section. For now, we can assume that after calling BASIC-INSERT(a), there exists a node v with START$(v) = a$ and Invariant 1 holds for every node in T.

Claim 1. *Let w denote the value of v in Line 12. Then, by the time Line 12 is executed, Invariant 1.b holds for all nodes not on $P(w)$.*

Proof. Suppose x is a node that has an ancestor u which is a node not on $P(w)$ that is a left child of a node on $P(w)$ (Fig. 1.a). Then Invariant 1.a and Invariant 1.b are the same for x. The only changes to nodes in $P(x)$ are caused by calls to PUSH-DELTA, which maintains Invariant 1.a. Therefore Invariant 1.b holds for x.

Suppose therefore that x has an ancestor u that is a node not on $P(w)$ but which is the right child of a node on $P(w)$ (Fig. 1.b). Then $K_{\mathrm{b}}(x) = K_{\mathrm{a}}(x) + \delta$. If the pseudocode omitted Line 8, then Invariant 1.a would hold. However, because we do not omit Line 8 Invariant 1.b holds. □

Claim 2. *After Line 16, Invariant 1.b holds for all nodes in T.*

Proof. After Line 11, we only update STAB values of nodes in $P(w)$, so we can not invalidate Invariant 1.b for any node not on $P(w)$. Therefore, we need only show that Invariant 1.b is maintained for nodes of $P(w)$.

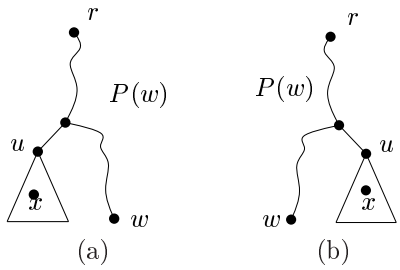

Figure 1: The two cases in the proof of Claim 1.

To show that Invariant 1.b is maintained for a node $v \in P(w)$, we use induction on the distance from v to w to show that, after time Line 15 is executed for v, Invariant 1.b holds for v. For the case $v = w$, we note that, because of the calls to PUSH-DELTA in Line 4, by the time Line 12 is executed, the value of $\text{STAB}(v) + \Delta(v) = K_{\text{a}}(v)$. But $K_{\text{b}}(v) = K_{\text{a}}(v) + \delta$ and is therefore correctly updated in Line 12.

Next we consider a node v at distance greater than 0 from w. Note that every ancestor u of v (including v itself) has $\Delta(u) = 0$ (from Line 4), and Invariant 1.b holds for both children of v (by the inductive hypothesis and Claim 1). This implies that for each child x of v, $K_{\text{b}}(x) = \text{STAB}(x) + \Delta(x)$. Since, by definition, $K_{\text{b}}(v)$ is the maximum of $K_{\text{b}}(\text{LEFT}(x))$ and $K_{\text{b}}(\text{RIGHT}(x))$, Invariant 1.b holds after the execution of Line 15. □

Therefore, INSERT-INFINITE-INTERVAL correctly updates the tree T so as to maintain Invariant 1. A simple examination of the code for INSERT-INFINITE-INTERVAL and MAX-COVER-WITNESS shows that their running time is $O(h)$ where h is the height of the tree. Next we show that this height can be kept to $O(\log n)$.

2.3 Rebalancing

One detail missing from the description of INSERT-INFINITE-INTERVAL is the implementation of BASIC-INSERT. Implementations of this function are described in many algorithms texts (c.f., Cormen *et al* [6]) so we do not go into details here.

Several methods exists for rebalancing binary search trees so that their height is $O(\log n)$. Usually, these methods use *rotations* to rebalance the tree. Which method used is a matter of preference, the important fact is that if we can show how to update the STAB, Δ, and INT, values during a rotation then any of the methods can be used.

A left rotation is illustrated in Fig. 2. We leave it to the reader to verify that the following wrapper, placed around a left rotation, correctly updates the information stored at x and y. Right rotation is implemented symmetrically.

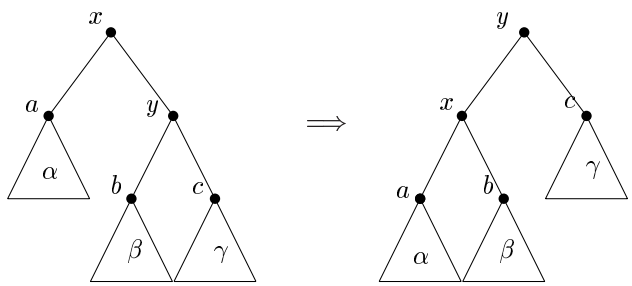

Figure 2: A left rotation

LEFT-ROTATE-WITH-UPDATE(x, y)

1: PUSH-DELTA(x)
2: PUSH-DELTA(y)
3: LEFT-ROTATE(x, y)
4: $\text{INT}(x) \leftarrow \text{INT}(a) \cup \text{INT}(b)$
5: $\text{STAB}(x) \leftarrow \max\{\text{STAB}(a) + \Delta(a), \text{STAB}(b) + \Delta(b)\}$
6: $\text{STAB}(y) \leftarrow \max\{\text{STAB}(x), \text{STAB}(c) + \Delta(c)\}$

We have just proven the following result.

Theorem 1. *A set of n intervals can be stored in a data structure requiring $O(n)$ space and supporting the operations,* INSERT, DELETE, MAX-COVER, *and* MAX-COVER-WITNESS *in $O(\log n)$ time per operation.*

3 Applications

In this section we show how the data structure from the previous section can be applied to solve a number of problems related to placing a regular square grid G over a set $S = \{p_1, \ldots, p_n\}$ of n points in the plane.

3.1 Infinite Grids

Throughout this section we assume that G is an infinite grid consisting of unit squares and that one of G's vertices is located at the origin. It is clear that all translations of G are equivalent up to additions of integers to the x and y components. Therefore we need only consider translations whose x and y components are in the range $[0, 1)$ so we represent translations as points in the square $[0, 1) \times [0, 1)$.

3.1.1 Optimizing 1-Occupancy

In this section, we consider the problem of optimizing the 1-occupancy, i.e., optimizing the number of grid cells occupied by at least 1 point of S. We

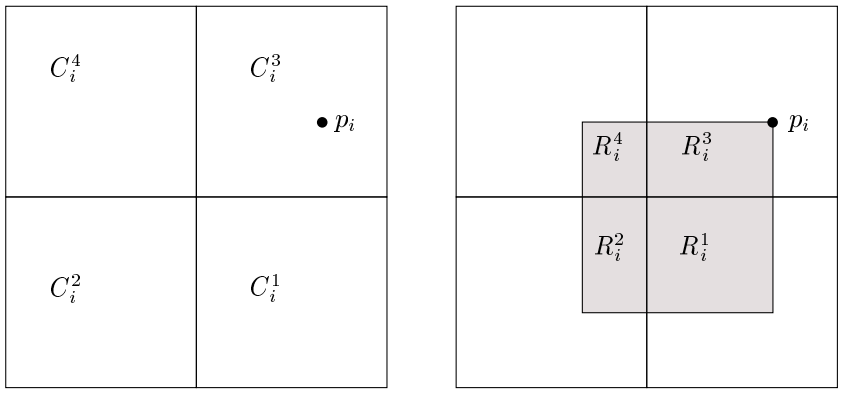

Figure 3: The point p_i will lie in one of four cells after the translation (left) which partitions the set of possible translations into four rectangles (right).

focus on the maximization problem since the modifications required to solve the minimization problem are obvious.

Since we only consider translation of at most 1 unit in the x and y directions, each point $p_i \in S$ must occupy 1 of 4 cells of G after translation, call these C_i^1, C_i^2, C_i^3, and C_i^4. This partitions the set of possible translations into four rectangles R_i^1, R_i^2, R_i^3, and R_i^4 that contain the top-left, top-right, bottom-left, and bottom-right corners of $[0,1) \times [0,1)$, respectively (see Fig. 3). We call R_i^j an *occupation rectangle* for cell C_i^j.

For a cell C, let the *k-occupied set* of C be the set of translations of G for which C is k-occupied. It follows that the k-occupied set of C is the set of points contained in k or more occupation rectangles of C.

Lemma 1. *The 1-occupied set of a cell C can be computed in $O(s \log s)$ time, where s is the number of occupation rectangles of C.*

Proof. It is clear from the above discussion that the 1-occupied set of C is the union of the occupation rectangles of C. In general, the union of s rectangles has complexity $O(s^2)$, however occupation rectangles have the property that they are contained in the unit square, and each contains at least 1 corner of the unit square.

Let G_1 be the set of occupation rectangles that contain the top-left or top-right corner of the unit square and let G_2 be the set of occupation rectangles that contain the bottom-left or bottom-right corner of the unit square. Then the boundary of $\cup G_1$ is given by the lower envelope E_1 of the bottom edges of rectangles in G_1 and the boundary of $\cup G_2$ is given by the upper envelope E_2 of the top edges of rectangles in G_2. Since both of these are envelopes of at most s parallel line segments they each have at most $2s$ vertices and can be computed in $O(s \log s)$ time [8].

E_1 and E_2 are both x-monotone and the space above E_1 is $\cup G_1$ and the

space above E_2 is $\cup G_2$. Any vertex of $\cup(G_1 \cup G_2)$ is either a vertex of G_1 or G_2 or an intersection point of E_1 and E_2. Since two monotone polygonal chains of length $O(s)$ intersect at most $O(s)$ times, the boundary of the 1-occupied set of C is of size $O(s)$ and can be obtained in an additional $O(s)$ time by merging E_1 and E_2. □

To determine a translation that maximizes 1-occupancy we proceed as follows: For each point $p_i \in S$ we compute the four occupation rectangles generated by p_i in $O(1)$ time using the floor function. Using an $O(n \log n)$ time sorting algorithm, we then gather the occupation rectangles of each cell C and compute the 1-occupied set for C in $O(s \log s)$ time, where s is the number of occupation rectangles of C. Since each point in S generates at most 4 occupation rectangles, this step takes $O(n \log n)$ time for all cells.

Finally, we partition the 1-occupied regions into rectangles. The problem of computing the translation that maximizes the number of occupied cells then becomes the problem of finding the point contained in the largest number of rectangles, i.e., finding the maximum clique in a rectangle intersection graph. This problem can be solved in $O(n \log n)$ time using plane sweep along with the data structure of Section 2. Refer to Lee [9] for details of the plane sweep.

Theorem 2. *Given n points $S = \{p_1, \ldots, p_n\}$ and an infinite grid G it is possible to report a translation of G which minimizes or maximizes the occupancy of G in $O(n \log n)$ time and $O(n)$ space.*

3.1.2 Optimizing k-Occupancy

Next we consider the problem of optimizing the k-occupancy. As before, the algorithm for maximizing k-occupancy is easily translated into an algorithm for minimizing k-occupancy. Therefore, we focus on the maximization problem.

We use the same general approach of computing, for each cell C, the k-occupied set of C. By partitioning these sets into rectangles and then applying the algorithm for maximum clique in a rectangle intersection graph we obtain our result. However, the difficulty comes from the fact that the k-occupied set of a cell has a richer combinatorial structure than its 1-occupied set.

To compute the k-occupied set of a cell C having s occupation rectangles, we partition the occupation rectangles of C into four sets S_1, S_2, S_3 and S_4, depending on which of the four corners of the unit square they contain (see Fig. 4). Rectangles that contain more than one corner of the unit square are assigned arbitrarily to one of their possible sets.

For each set S_i we then compute an arrangement A_i that partitions the unit square into $k+1$ regions $R_i^0, \ldots, R_i^{k-1}, R_i^{\geq k}$ where all points in R_i^x are contained in x occupation rectangles (see the top-right of Fig. 4). This is done by a vertical line sweep over the rectangles in S_i starting at the corner of the unit square contained in all elements of S_i. The sweep line status is maintained by a balanced-binary tree and we construct the arrangement by keeping track of the k-smallest elements in the tree. The only events occur when rectangles terminate, at which point an element is removed from the tree, causing at most

k of the $k+1$-smallest elements to change rank. Each event can therefore be handled in $O(k + \log s)$ time and the resulting arrangement has complexity $O(ks)$.

Each A_i consists of k x-y monotone chains (sometimes calls staircases), where the directions of monotonicity are different for each. Because the chains are monotone, the intersection of any horizontal or vertical line segment with all chains consists of at most $4k$ points and line segments. Since the chains consists of a total of $O(s)$ line segments, this implies that the total number of intersections between all chains is $O(sk)$, and these intersections can be computed in $O(sk \log s)$ time using plane sweep [4]. This results in an arrangement A of orthogonal line segments that has $O(sk)$ faces.

For each face F of A, we can examine the faces of A_1, A_2, A_3 and A_4 that contain F. Since each face F' of A_i is labelled as being contained in $0, \ldots k-1$ or $\geq k$ occupation rectangles of S_i, we can determine whether the points in F are contained in k or more occupation rectangles of C. Thus we can determine whether for each face F of A, whether or not F is part of the k-occupied set. This step can be performed in $O(sk \log s)$ time using point location or, if we are a bit more careful, in $O(sk)$ time by doing a depth-first traversal of the faces of A.

Thus, for each cell C, we can find the k-occupied set of C in $O(sk)$ time, where s is the number of occupation rectangles of C. As in the previous section, once the the k-occupied sets have been computed for each cell, a plane sweep over all sets yields an optimal translation of G in an additional $O(kn \log n)$ time.

Theorem 3. *Given n points $S = \{p_1, \ldots, p_n\}$, an integer $1 \leq k \leq n$, and an infinite grid G it is possible to report a translation of G which minimizes or maximizes the k-occupancy of G in $O(kn \log n)$ time and $O(kn)$ space.*

3.1.3 Finding a Translation with k-Occupancy 0

We note that the problem of finding a translation with k-occupancy 0 is solvable using the algorithm of the previous section. The translation with the minimum k-occupancy is found and if that k-occupancy is 0 then the solution is found otherwise none exists.

Theorem 4. *Given n points $S = \{p_1, \ldots, p_n\}$, an integer $1 \leq k \leq n$, and an infinite grid G it is possible to decide if there exists a translation of G such that no cell of G is k-occupied and report such a translation if one exists in $O(kn \log n)$ time and $O(kn)$ space.*

3.1.4 Optimizing k

To solve Problems 4 and 5, we can perform a "doubling search" on the value of k. Initially, we guess k to be 1 and repeatedly double the value of k. For each value of k we use Theorem 3 or Theorem 4 to test if the value of k has become too large. If it has, we stop and perform binary search on the range $(k/2, k)$ to find the optimal value of k. The additional cost of this binary search is a $\log k$ factor in the running time.

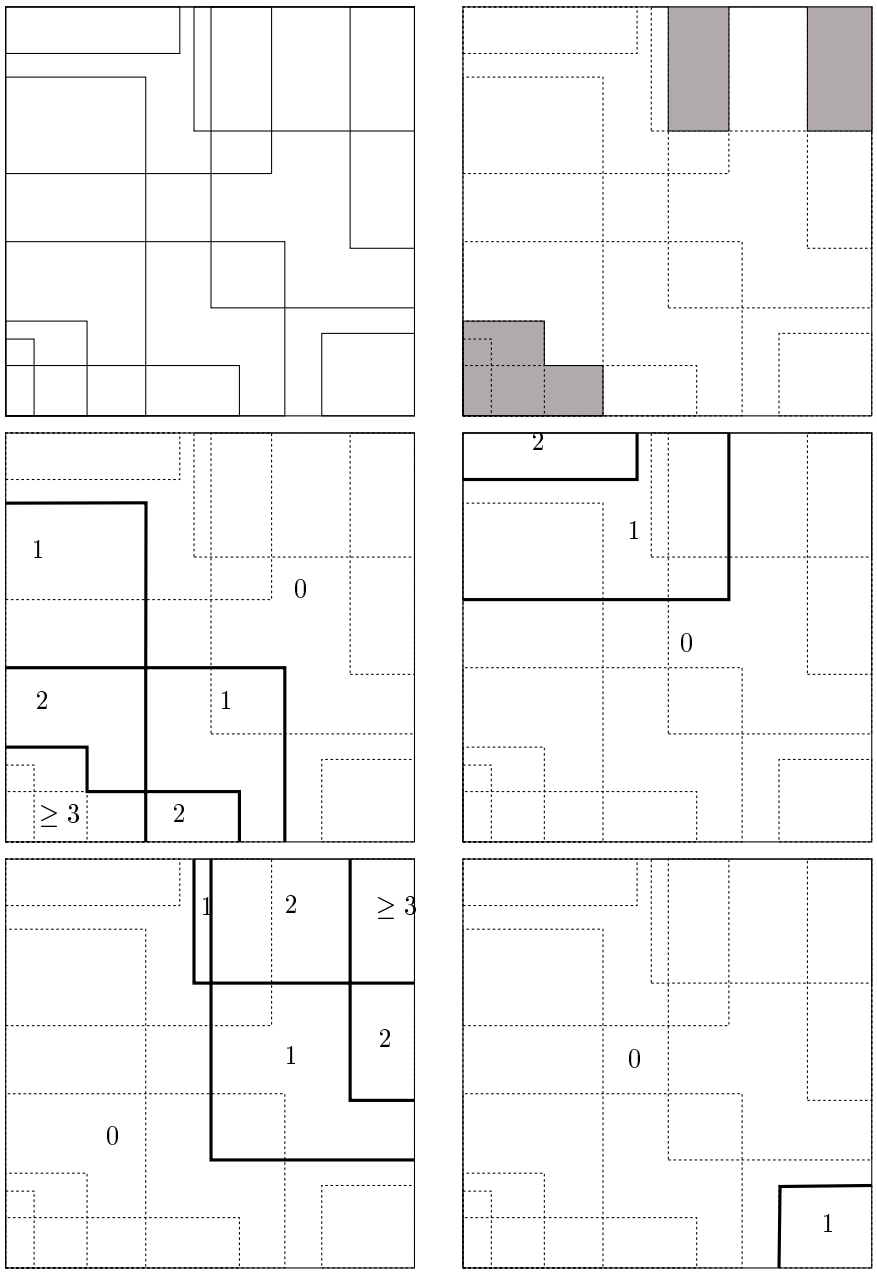

Figure 4: The original set of occupation rectangles for C, the k-occupied set of C and the four arrangements $A_1, \ldots, A_4$ for $k = 3$.

Theorem 5. *Given n points $S = \{p_1, \ldots, p_n\}$ and an infinite grid G it is possible to report the smallest integer k for which there exists a translation with k-occupancy equal to 0 in $O(kn \log n \log k)$ time and $O(kn)$ space. Additionally, given an arbitrary non-negative integer r, it is possible to report the largest integer k for which there exists a translation with k-occupancy at least r in the same time and space.*

3.2 Placing a Finite Grid

Next we turn to the problem of placing an $m_r \times m_c$ grid G over the point set S so that the k-occupancy of G is optimized. This appears to be a very different problem, since now all translations are not equivalent up to integer additions to the x and y components. It should come as no surprise then that the grid size $m_r \times m_c = m \geq 1$ becomes a factor in the running time.

For each point $p_i \in S$, we transform p_i into a unit square s_i whose bottom right corner is p_i. Thus, if we place a grid square with its top-left corner inside of s_i then that square will be occupied. Next, we compute the region R contained in k or more s_i's. This is achieved, as in Section 3.1.2, by overlaying an infinite grid and calculating the k-occupancy region for each occupied cell. The region R is simply the union of each of these disjoint k-occupancy regions. It follows that if we place a unit square with its top-left corner in R, then that square is k-occupied.

We then make m translated copies $R_1, \ldots, R_m$ of R, where R_i is R translated by the vector $(-i \bmod m_c, \lfloor i/m_c \rfloor)$. It follows that if we place the top-left corner of a $m_r \times m_c$ grid G at a point that is contained in x of the R_i then the number of k-occupied cells of G is x. Thus, we have reduced the problem of placing G to a problem of finding a point contained in the largest number of regions R_i. Since each R_i is a set of orthogonal polygons they can be decomposed into rectangles thereby reducing the problem to maximum-clique in a rectangle intersection graph.

In summary, we compute R in $O(kn \log n)$ time, partition R into rectangles and create $R_1, \ldots, R_m$. This results in a set of rectangles of size $O(mkn)$. We then solve the maximum-clique problem in the resulting set of rectangles in time $O(mkn \log mn)$. To reduce the space requirements we note that $R_1, \ldots, R_m$ need not be computed explicitly. Indeed, in the priority queue that implements the plane sweep of the maximum clique algorithm, each edge of R need appear only once at any given time. At all times during the sweep, an edge in the queue represents each of the m_c copies with the same y-coordinate. To initialize the queue the left and upper most copy of each edge is inserted into the queue. When an edge is swept over it generates the appropriate m_c simultaneous events, one for each grid column. After it has been swept the edge is translated downward and reinserted into the queue to represent the next set of m_c events. This translation is performed only m_r times for each edge thus providing the m copies of each edge.

Theorem 6. *Given n points $S = \{p_1, \ldots, p_n\}$, an integer $1 \leq k \leq n$, and a*

finite grid G with m cells it is possible to report a translation of G with maximal k-occupancy in $O(mkn \log mn)$ time and $O(kn)$ space.

4 Conclusions

We have studied problems related to placing a regular square grid over a set of points in the plane in order to optimize some function of the k-occupancy. For small values of k most of our algorithms have near-linear running times.

To solve these problems we have introduced a new lazy data structure for maintaining a point in the maximum-clique of an interval graph under insertions and deletions of intervals. This data structure is fully dynamic as it doesn't require the interval endpoints to be known in advance.

The data structure is also rather general. Besides storing intervals, it could be used to, e.g., store piecewise polynomial functions and return a value x that maximizes or minimizes the sum of all the functions evaluated at x. One could imagine using such a data structure with plane-sweep to perform operations like maintaining the minimum in sums of functions, or combining it with the persistence paradigm [7] to get an efficient representation of sums of functions.

Open Problem 1. *Find additional applications for our interval data structure.*

Another obvious open problem is that of improving the running time of the algorithm for Problem 6. In particular, it would be nice to produce an algorithm which is output sensitive to the size of the k-occupied set.

Open Problem 2. *Find an output sensitive algorithm which computes the boundary of a k-occupied region.*

References

[1] P. K. Agarwal, B. K. Bhattacharya, and S. Sen. Output-sensitive algorithms for uniform partitions of points. In A. Aggarwal and C. P. Rangan, editors, *Proceedings of the 10th International Symposium on Algorithms and Computation, ISAAC*, volume 1741 of *Lecture Notes in Computer Science*, pages 403–414. Springer, December 1999.

[2] T. Asano and T. Tokuyama. Algorithms for projecting points to give the most uniform distribution with applications to hashing. *Algorithmica*, 9(6):572–590, June 1993.

[3] J. L. Bentley. Solutions to Klee's rectangle problems. Technical report, Carnegie-Mellon University, 1977.

[4] J. L. Bentley and T. A. Ottmann. Algorithms for reporting and counting geometric intersections. *IEEE Transactions on Computers*, C-28:643–647, September 1979.

[5] D. Comer and M. J. O'Donnell. Geometric problems with applications to hasing. *SIAM Journal on Computing*, 11(2):217–226, May 1982.

[6] T. H. Cormen, C. E. Leiserson, and R. L. Rivest. *Introduction to Algorithms*. MIT Press, Cambridge, 1990.

[7] J. R. Driscoll, N. Sarnak, D. D. Sleator, and R. E. Tarjan. Making data structures persistent. *Journal of Computer and System Sciences*, 38(1):86–124, February 1989.

[8] J. Hershberger. Finding the upper envelope of n line segments in $O(n \log n)$ time. *Information Processing Letters*, 33(4):169–174, December 1989.

[9] D. T. Lee. Maximum clique problem of rectangle graphs. In F. P. Preparata, editor, *Advances in Computing Research*, pages 91–107. JAI Press, 1983.

[10] M. J. van Kreveld and M. H. Overmars. Concatenable segment trees (extended abstract). In B. Monien and R. Cori, editors, *Proceedings of the 6th Annual Symposium on Theoretical Aspects of Computer Science*, volume 349 of *Lecture Notes in Computer Science*, pages 493–504. Springer, February 1989.