

New Algorithms for Computer Graphics

Mark H. Overmars

*Department of Computer Science
University of Utrecht
P.O.Box 80.012
3508 TA Utrecht
The Netherlands*

1. Introduction

The area of computational geometry deals with the study of algorithms for problems concerning geometric objects like e.g. lines, polygons, circles, etc. in the plane and in higher dimensional space. Since its introduction in 1976 by Shamos the field has developed rapidly and nowadays there are even special conferences and journals devoted to the topic. A list of publications by Edelsbrunner and van Leeuwen [6] collected in 1982 already contained over 650 papers. And this number has rapidly increased since then.

Clearly, a large number of problems in computer graphics deals with geometric objects as well. Examples are hidden line elimination, windowing problems, intersection problems, etc. Hence, computer graphics can benefit from the techniques developed in computational geometry.

In computational geometry many new sophisticated data structures and algorithms have been designed for dealing with geometric objects. These *geometric data structures* normally work in object space, i.e., objects are stored in a structural way rather than being broken down to e.g. pixels. In computer graphics many data structures work in image space.

Storing and treating objects in object space normally gives much more flexibility. Hence, people in computer graphics should be aware of the existence of such object space data structures.

On the other hand, in many practical situations in graphics objects are taken from some bounded domain or can be rounded to some bounded domain. Hence, it seems useful to study problems from computational geometry in a bounded universe, i.e., on a grid. To state this more formally, let U be some integer. We assume that all points, endpoints of line segments, vertices of polygons, etc. have

integer coordinates between 0 and $U-1$. In other words, the points are chosen from $[0..U-1]^d$, also denoted as U^d where d is the dimension. We call U the *universe size*. For example, when the points are pixels on 1024*1024 screen, $U=1024$. Note that, with appropriate scaling and translation, U^d can approximate any bounded region of the real d -dimensional space. Note that in this framework e.g. line segments are still real line segments that might intersect in some point that is not on the grid. We do not break objects down to pixels.

Computational geometry on a grid has recently been studied by a number of researchers (See e.g. [8, 11-14, 16, 19]). In this paper we will describe a number of the results obtained so far, showing that on a grid many problems from computational geometry can be solved more efficiently and, often, in a much simpler way than in arbitrary real space. This is mainly due to the fact that we can use table look-up rather than searching in a number of applications and that there often is no need for rebalancing structures. Efficient solutions will be presented for problems like windowing, finding intersections between geometric objects and hidden surface elimination.

We do not present hardware solutions for the problem. But, due to the simplicity of the algorithms presented we claim that they can be implemented in hardware in a much more easy way than general object space methods without sacrificing too much of the generality of the solutions. Further study of these new algorithmical techniques might be important for the development of computer graphics and might lead to new approaches, both in hardware and in software.

This paper is organised as follows:

In Sections 2, 3 and 4 we treat two basic problems: sorting and searching. Most algorithms use sorting and searching as basic steps and, hence, it is important to investigate their complexity on a grid. We give efficient, simple methods for *lexicographical sorting* on a multi-dimensional grid and demonstrate some search structures, especially the *vanEmdeBoas tree* [3, 4].

In Section 5 we treat the classical problem of finding the *convex hull* of a set of points on a planar grid. Moreover we look at related problems like finding *maximal elements*. We show that on a grid these problems can be solved in linear time.

In Sections 6, 7 and 8 we treat *intersection problems*. First of all we look at the problem of finding all intersections in a set of axis-parallel rectangles. This is especially relevant in computer graphics applications where objects are surrounded by bounding boxes. Finding intersections between bounding boxes is the first step in finding intersections between the objects themselves. Secondly, we look at intersections in a set of line segments. Finding such intersections can be a first step in hidden line removal algorithms [25].

Most of the contents of this paper has been copied from [20].

2. Sorting

In this section we will shortly recall two well-known results on sorting.

Theorem 2.1 *Given a set of n keys with a value between 0 and $U - 1$, they can be sorted in time $O(n + U)$ using $O(n + U)$ storage.*

Proof. We make an array of U buckets and initialise them to be empty. Next we put each element in the set in its bucket using the key as array index. Finally we walk along the buckets in order, reporting the key when a bucket is not empty. The bounds follow trivially. \square

On a two dimensional grid it is often important to sort the points *lexicographical*, i.e., $p = (p_1, p_2) < p' = (p_1', p_2')$ if $p_1 < p_1'$ or $p_1 = p_1'$ and $p_2 < p_2'$. To sort a set of points lexicographical we can use a simple version of bucket sort:

Theorem 2.2 *Given a set of n points in U^2 , they can be sorted lexicographical in time $O(n + U)$ using $O(n + U)$ storage.*

Proof. We create two arrays of U buckets, but this time each bucket will contain a list of points. First we sort the points by second coordinate, using one of the arrays of buckets. If more points have the same second coordinate we put them in the list in the corresponding bucket. Next treating the points ordered by second coordinate we put them in the other array of buckets using their first coordinate as an index. When more points have the same first coordinate they will be put in the list in the bucket. Because we treat the points by increasing second coordinate, such a list corresponding to a particular first coordinate will contain the points in the order of their second coordinate. Hence, walking along the buckets and reporting the points in the order in which they appear in the lists will give the points in lexicographic order. The bounds follow trivially. \square

On a d -dimensional grid we can use a similar technique. First we sort all points with respect to the d th coordinate. Next, treating the points in this order, we sort them with respect to the $d - 1$ th coordinate, next on the $d - 2$ th coordinate, etc. This immediately leads to the following result (noting that we need only two arrays of buckets):

Theorem 2.3 *Given a set of n points in U^d , they can be sorted lexicographical in time $O(d(n + U))$ using $O(n + U)$ storage.*

These results are, in fact, not optimal. Kirkpatrick and Reisch [16] have shown that a set of n keys in U can be sorted in time $O(n(1 + \log \log_n U))$ using $O(n + U)$ storage. Their result can easily be extended to sort a set of points in U^d in time $O(dn(1 + \log \log_n U))$ using $O(n + U)$ storage. On the other hand, when U is relatively small, this result is not any better and the method is more complicated.

3. Simple searching

Let V be a set of n keys in U . We often need to store V such that questions like: Is $K \in V$, which keys in V lie in between A and B , what is the rank of a key K in V , etc., can be answered efficiently. In arbitrary real space search trees, like e.g. AVL-trees, are used for answering such questions. On a grid we can solve these problems much more efficiently.

Theorem 3.1 *Given a set V of n keys in U , we can store V using $O(n + U)$ storage such that given a key K we can determine in $O(1)$ time the largest key in V that is $\leq K$.*

Proof. The method is trivial. We make an array KEY of length U . In KEY[i] we store whether $i \in V$ and we store a pointer to the largest $j < i$ such that $j \in V$. Clearly, this takes $O(n + U)$ storage and queries can be answered in time $O(1)$. The structure can be built in time $O(n + U)$. \square

A similar structure can be used to solve the ranking problem that asks how many keys are $\leq K$ in time $O(1)$ and the range searching problem that asks for all keys in $[A : B]$ in time $O(1 + k)$ where k is the number of reported answers.

The structure is static, i.e., it is hard to insert or delete keys in the set V . For dynamic structures, see the next subsection.

The structure works fine when U is not much larger than n but when U is much larger the amount of storage required becomes too large. Fredman et al. [7] describe a technique, based on perfect hashing, that can answer queries of the form: is $K \in V$, in time $O(1)$ using only $O(n)$ storage. Unfortunately, their structure has a very high preprocessing time of $O(n^3 \log U)$. Moreover, as it is based on hashing, it cannot answer questions like: what is the largest key $\leq K$.

Willard [26] designed a structure, based on [7], that can answer such questions in time $O(\log \log U)$ using $O(n)$ storage. Unfortunately, also his structure has a very high preprocessing time.

4. The vanEmdeBoas tree

The vanEmdeBoas tree [3, 4] is an efficient dynamic data structure for storing a set of keys in a one-dimensional grid U , as long as U is not too large. Here we will describe a variation of the structure that obtains exactly the same time and space bounds but is easier to describe and implement. We will call it a vanEmde-Boas tree anyhow.

Let V be the set of points and let U be the universe. We construct a complete balanced binary tree T of depth $\log U$ on the universe. (For convenience we assume that $U=2^i$ for some i .) Now you can think of this tree as having the following form. It consists of a top tree T_0 of depth $\left\lceil \frac{\log U}{2} \right\rceil$ having about \sqrt{U} leaves and about \sqrt{U} bottom trees T_1, T_2, \dots also of size about \sqrt{U} . T_0 is a tree with a universe of size \sqrt{U} being the indexes of the bottom tree. Also the bottom trees have a universe of size about \sqrt{U} . Each of the trees T_i is in the same way subdivided in a top tree $T_{i,0}$ and bottom trees $T_{i,1}, T_{i,2}, \dots$, each of size $\sqrt[4]{U}$, etc. See figure 1 for the partition obtained.




Figure 1: The partition of the tree.

Keys will be stored in the tree T in the following way. We split V in subsets V_1, V_2, \dots corresponding to the subtrees T_1, T_2, \dots . Each V_i we store in T_i in the following way. At the root of T_i we store the size of V_i . If V_i is empty, nothing else will be stored in T_i (although the subtree will be there). If V_i contains 1 or 2 keys, we store them at the root of T_i . If V_i contains more than 2 keys we store at the root the largest and smallest key in V_i , the other keys we store in T_i recursively in the same way. (I.e., we split V_i in subsets $V_{i,1}, V_{i,2}, \dots$ corresponding to $T_{i,1}, T_{i,2}, \dots$, etc.) When V_i is not empty, we insert i in T_0 . Finally, we link all keys in T in a double linked list such that, once we find a key, we can always find its two neighbors.

Clearly, the main T has $O(U)$ nodes and at each node at most 2 keys and the size are stored. Hence, the structure uses $O(U)$ storage. (Note that it will require $\Omega(U)$ storage, even when V is very small.)

The vanEmdeBoas tree is used for so-called neighbor queries that ask for some neighbor of a given query key K (i.e., the first key larger or smaller, or equal to K). Of course, because we linked the keys in a double linked list, once we have a neighbor we can always in $O(1)$ time find the other neighbor. This also allows us to find e.g. all keys between K_1 and K_2 in time $O(1)$ per answer once we found a neighbor of K_1 .

To perform a neighbor query with key K we proceed in the following way. Determine the tree T_i K comes in. Because T is a complete binary tree this can be done in time $O(1)$. When T_i contains 1 or 2 keys (indicated with the root) one of the 1 or 2 keys stored at the root must be a neighbor and we are done. When T_i contains more than 2 keys, we first check whether one of the two keys stored at the root is a neighbor. (This can be done in $O(1)$ time by looking at them and their neighbors in the double linked list.) If so we are done. Otherwise the neighbor must be stored in T_i and we continue our search there recursively in the same way. (I.e., we look at the appropriate $T_{i,j}$, etc.) When T_i is empty, we search for a neighbor j of i in T_0 . This means that T_j is the first non-empty subtree to the left or to the right of T_i . Assume $j < i$. In that case the largest key stored at the root of T_j is a neighbor of K . Similar, when $j > i$, the smallest key stored at T_j is a neighbor.

To estimate the time required, note that after $O(1)$ work we are left with a subtree of size \sqrt{U} . Hence, if $Q(U)$ denotes the query time over a tree of size U , we get $Q(U) = O(1) + Q(\sqrt{U})$, which leads to $Q(U) = O(\log \log U)$.

To perform an insertion of a key K we proceed in the following way. First we search for a neighbor K' of K in T . Next, determine the subtree T_i K must come in. Increase the size of T_i with 1. If T_i is empty store K at its root and insert i in T_0 . If T_i contains 1 key just store K at its root. If T_i contains 2 or more keys, let A and B (with $A < B$) be the keys stored at its root. If $K < A$ store K at the root and insert A in T_i . If $K > B$ store K here and insert B in T_i . Otherwise insert K in T_i . Finally, using K' insert K in the double linked list.

To delete a key K determine the subtree T_i K is in. Decrease the size of T_i . If the subtree contains 1 key, it must be K and K is stored at the root. Remove K here and remove i from T_0 . If the subtree contains 2 keys, K is stored at the root and we can just remove it. Otherwise, if K is not stored at the root of T_i , delete K from T_i . If K is stored at the root, locate its neighbor in T_i (there must be such a neighbor). This can be done in time $O(1)$ using the double linked list. Replace K by its neighbor, and delete the neighbor from T_i . Finally, delete K from the double linked list.

The insertion and deletion time is bounded by $O(\log \log U)$ following the same arguments as in the query time. We conclude:

Theorem 4.1 *Let V be a set of keys in U . There exists a dynamic structure for storing V , requiring $O(U)$ storage such that insertions, deletions and neighbor queries take time $O(\log \log U)$.*

A number of improvements over the vanEmdeBoas tree have been suggested. Johnson [10] describes a modification to reduce the initialisation time and to perform some queue operations more efficiently. Willard [27] and Karlsson [11] describe dynamic structures on a grid that use only $O(n)$ storage. To obtain this lower space requirement, the update and query time of their methods is increased to $O(\sqrt{\log U})$. We won't go into details of their solutions because the methods are more complicated and unnecessary in the case U is of the same order as n .

It should be noted that, although the method might look complicated, it is in fact very simple to implement. Because the tree is a complete binary tree, it can easily be implemented in the form of an array. Because keys are never stored at the leaves this array will have length $U-1$. Calculating where a particular node is in the tree is a very simple operation on the bits of the key we are searching for. Hence, implementing the structure in hardware (for some fixed U) should be reasonably easy.

5. Convex hulls and related problems

As an example of how the fact that our points have integer coordinates can help us in deriving efficient algorithms we will consider the *convex hull* problem. Given a set of points (or a polygon), the convex hull is the smallest convex polygon containing all the points (or the polygon). A convex hull of a polygon is often a good first estimation for computing e.g. intersections or visibility. Computing the convex hull of a set of n points can be done in $O(n \log n)$ time (see e.g. [1, 9, 17, 23]). For a simple polygon there exists a quite complicated method for computing the convex hull in linear time. We will show that on a grid there exists a very simple linear time solution to both problems. We will only treat the problem of finding the convex hull of a set of points. The convex hull of a polygon is the same as the convex hull of the set of vertices of the polygon.

Theorem 5.1 *Given a set V of n points in U^2 , the convex hull of V can be computed in time $O(n + U)$.*

Proof. The method is based on [9]. First we sort all points by x -coordinate. This can be done in $O(n + U)$ time using Theorem 2.1. Clearly, the leftmost point p_l and the rightmost point p_r belong to the convex hull. The convex hull now consists of an upper chain UC that runs from p_l to p_r having all points below it, and a lower chain LC having all points above it. (See figure 2.)

We will only show how to compute UC. LC can be computed in exactly the same way. We will describe the method without much details. For detail see [9]. We walk from p_l to p_r . We will take care that when we are at some point p_i we found a convex arc from p_l to p_i with all points in between below it. Assume we have reached some point p_i . Now look at the angle $p_{i-1} p_i p_{i+1}$. If this angle is convex we can continue with

$i := i + 1$. Otherwise, p_i will not belong to the upper convex hull and can be removed. We backup and continue with $i := i - 1$. It is easy to see that is this way, in $O(n)$ time, we find the upper convex chain. \square




Figure 2: Convex chains.

In fact, the only place where we use the property that the points lie on a grid is when sorting the points. After the sorting the algorithm runs in $O(n)$ time independent of the fact whether the points lie on a grid.

Theorem 5.2 *A convex polygon (hull) on n points on a grid can be stored using $O(U)$ storage such that for a given point p one can determine in $O(1)$ time whether p lies inside or outside the polygon.*

Proof. Scanning from left to right over the polygon, we store for each vertical grid line the two intersection points of this line with the polygon. To search with a point p (not necessarily on the grid) we determine the two vertical grid lines in between which p lies (or the one on which p lies). With the intersection points stored there we can easily determine whether p lies inside or outside the polygon. This clearly takes $O(1)$ time. \square

The structure above can trivially be built in time $O(n + U)$.

As a second example let us consider the *maximal elements* problem. Given a set V of n points in the plane, a point $p = (p_1, p_2)$ in V is called maximal if there is no point in V with x -coordinate larger than p_1 and y -coordinate larger than p_2 . See figure 3.

To find the maximal elements we sort the points by decreasing x -coordinate. Now we walk along the points in this order, keeping track of the point with largest y -coordinate found so far. Let y this y -coordinate. In the beginning $y = -\infty$. When we come at a point $p = (p_1, p_2)$ we check whether $p_2 > y$. If so we report p as a




Figure 3: Maximal elements.

maximal element and set $y = p_2$. And we continue with the next point. This can easily be done in $O(1)$ per point. Hence, the time for sorting dominates the run time of the algorithm.

Theorem 5.3 *Given a set of n point in U^2 , the maximal elements problem can be solved in time $O(n + U)$ using $O(n + U)$ storage.*

When U is large, the results in this section are not very good. Using the sorting method of [16] we can improve the bounds in such cases. See also [15].

6. The interval trie

Let V be a set of n intervals in U , that might intersect or overlap. We want to store V such that we can perform *stabbing queries*, that ask for those intervals containing a given query point, efficiently. When the intervals have arbitrary real coordinates, this problem can be solved using a so-called *interval tree* using $O(n)$ storage in a query time of $O(k + \log n)$ where k is the number of reported intervals (see e.g. [24] for a description). In this subsection we will describe a solution on a grid, called the *interval trie*, that has slightly better asymptotic time complexity but might behave much better when the intervals are relatively small.

Let $F(U) = \log^\epsilon U$ for some $\epsilon > 0$. We split the set V in subsets V_1, V_2, \dots , where V_i contains all intervals of length between $F(U)^{i-1}$ and $F(U)^i$. Clearly, there are at most $O(\log U / \log F(U)) = O(\log U / \log \log U)$ sets V_i . We will treat each V_i separately. For V_i we divide the universe in equal parts of size $F(U)^{i-1}$. Hence, each interval in V_i has its left endpoint in one part, overlaps at most $F(U)$ parts and has its right endpoint in another part. See figure 4. No interval can have its two endpoints in the same part. For each part we have three lists: B

containing all the left endpoints, E containing all the right endpoints and O containing all the overlapping intervals. See figure 5. Both B and E we store as vanEmdeBoas trees, making it possible to insert and delete endpoints in time $O(\log \log U)$. O we store as a double linked list. Moreover, in B we store for each interval where it is stored in O lists.




Figure 4: Parts overlapped by an interval.




Figure 5: Intervals stored at a part.

To perform a stabbing query with point p we query each V_i separately. In V_i we determine the part p lies in in time $O(1)$. We report all intervals in the corresponding O list. In B we start at the leftmost endpoint. If it lies left of p we report it and continue with the next endpoint. If not there are no answers in B . Similar in E . It is easy to see that in this way we find all answers in V_i in time $O(1)$ plus the number of answers found. As we have to query $O(\log U / \log \log U)$ sets the query time is bounded by $O(k + \log U / \log \log U)$ where k is the number of answers found.

To insert or delete an interval we determine the set V_i the interval belongs to (looking at its length). We insert or delete the interval in the appropriate B and E structure. Moreover we have to add it to at most $F(U)$ O -lists. It is easy to see that this takes time $O(\log \log U + F(U)) = O(\log^6 U)$.

Theorem 6.1 *Given a set of n intervals on a grid U , for each $\epsilon > 0$ there exists a dynamic structure for solving the stabbing problem, that can be updated in time $O(\log^\epsilon n)$, with a query time of $O(k + \log U / \log \log U)$, where k is the number of reported answers.*

Although this is only a minor improvement in query time, in practical situations the structure might perform much better. When all intervals are short, only for small i , V_i will contain intervals. As the query time is, in fact, bounded by the number of non-empty sets V_i in such situations it will be much lower.

7. Intersections between rectangles

Intersection problems are very important in computer graphics. For example, many hidden-surface and hidden-line removal algorithm are based on finding intersections of geometric objects. In this section we will consider intersections between axis-parallel rectangles. Axis-parallel rectangles occur especially when working with bounding boxes. Only those objects whose bounding boxes intersect need to be considered further. In the next section we treat the more general problem of finding all intersections in a set of line segments.

To solve the intersection problem for axis-parallel rectangles efficiently on a grid we will use the vanEmdeBoas tree and the interval trie described above.

We will first consider the following subproblem: Let V be a set of n horizontal and vertical line segments, report all intersections between them. To solve this problem we use a plane sweep approach. We sweep a scanline from left to right, stopping at each endpoint of a line segment. With the scanline we keep a vanEmdeBoas tree T , storing all horizontal line segments that are intersected by the scanline.

So we start by sorting all endpoint of the line segments in left to right order. When more endpoint have the same x -coordinate we first store the left endpoints of horizontal line segments, then the vertical line segments with this x -coordinate and finally the right endpoints of horizontal line segments. We initialise T to be empty.

For each endpoint p belonging to a segment s we do the following: If p is the left endpoint of s we search with p in the tree to see whether it lies on any other horizontal line segment. If so we report it. Next we insert s in the vanEmdeBoas tree. If p is a right endpoint we delete s from T . Finally, if s is a vertical line segment $(x, y_1)(x, y_2)$ we search with y_1 in T and walk along the list of stored segments, reporting all segments that lie between y_1 and y_2 .

It is easy to see that the operations take time $O(\log \log U)$ each (plus the number of answers reported).

Theorem 7.1 *Given a set V of n horizontal and vertical line segments in U^2 , all k intersections between them can found in time $O(k + U + n \log \log U)$.*

Two axis-parallel rectangles intersect if either they intersect in the boundary, or one lies completely inside the other. Given a set V of rectangles we can find those pairs that intersect in the boundary by considering each rectangle as four line segments and finding the intersections between the line segments in the way described above. So we are left with the problem of finding all pairs where one is completely contained in the other. Clearly, if a rectangle is completely contained in the other, its leftbottom corner lies in the other rectangle. So we can as well formulate the problem as follows: Given a set of points and a set of axis-parallel rectangles, report all pairs (p, R) where p lies in R .

To solve this subproblem, we again use a plane sweep approach, sweeping a scanline from left to right over the grid, but this time we store with the scanline an interval trie T (as described in the previous section) storing the projections of the rectangles intersecting the scanline. We sort all points and left and right boundaries of the rectangles by x -coordinate. When more have the same x -coordinate we store first the left boundaries, then the points, and finally the right boundaries. We initialise T to be empty. If we reach a left boundary, we insert the interval in the interval trie. If we reach a right boundary, we delete it. Finally, if we meet a point, we perform a stabbing query with the point. Using the results from the previous section, noting that $\log^6 U = O(\log U / \log \log U)$, we conclude

Theorem 7.2 *Given a set of n points and axis-parallel rectangles in U^2 , all k pairs (p, R) with p contained in R can be found in time $O(k + U + n \log U / \log \log U)$.*

Combining Theorems 7.1 and 7.2 we obtain:

Theorem 7.3 *Given a set of n axis-parallel rectangles in U^2 , all k intersecting pairs can be found in time $O(k + U + n \log U / \log \log U)$.*

As noted in Section 6, the time will, in practise, often be much better.

8. Intersections between line segments

In this subsection we determine intersections among n line segments which are arbitrary oriented, but still are defined by endpoints at integer coordinates between 0 and $U - 1$. Until recently, the best algorithm to find all intersections required $O((k + n) \log n)$ time where k is the number of reported intersections (see Bentley and Ottmann [2] or Preparata and Shamos [24]). Chazelle [5] was the first to provide a solution with a running time of the form $O(k + f(n))$, where $f(n)$ is a subquadratic function of n . In [5] $f(n) = n \log^2 n / \log \log n$.

The line segment intersection problem is important in a number of Computer Graphics applications. For example, Schmitt [25] uses the line segment intersection problem as an important step in a hidden line removal algorithm. The complexity of the line segment intersection problem is dominant in his solution. Hence, any

improvement on the intersection problem improves his hidden surface removal algorithm. Also other hidden surface removal algorithms use line segment intersection as an important step.

We will show that, in the case the line segments have endpoints on an integer grid, all intersections can be found in time $O(k + U + n \log n)$ using $O(n + U)$ space. To this end we will adapt the technique in [2] to work on a grid, replacing search trees by table lookup.

For each x -coordinate x we store three unordered bags of points. $BEGIN(x)$ contains all line segments with a left endpoint with x -coordinate x . $END(x)$ contains all line segments with a right endpoint with x -coordinate x . $INTER(x)$ contains all intersections found so far with x -coordinate in the interval $[x, x + 1)$ (intersections can have coordinates in between grid points). When a line segment is vertical with x -coordinate x , we store it in both the $BEGIN(x)$ and the $END(x)$ bag.

To initialise the structures, each line segment is stored in the appropriate $BEGIN$ and END bag (by using table lookup) and all $INTER$ bags are made empty.

Next we move a scanline from left to right stopping at each x -coordinate. With the scanline we keep a balanced binary search tree T (e.g. an AVL-tree) storing all line segments that are intersected by the scanline in the order of intersection in the internal nodes. We will take care that the $INTER$ bags always contain exactly the intersections between neighboring segments in T . At each stop position x we perform the following three actions:

1. Insert all line segments that are stored in $BEGIN(x)$ in the tree T . For each such line segment s we locate its two neighbors s_1 and s_2 in T . We compute the intersection between s and s_1 and (if any) insert it in the appropriate $INTER$ bag. Similar, we compute the intersection between s and s_2 and store it. Finally, we compute the intersection between s_1 and s_2 . Because these segments are no longer neighbors, we have to remove the intersection from the $INTER$ bag (if there was an intersection). Keeping with each two neighbors in T a pointer to their intersection in the $INTER$ bag this intersection can be removed in time $O(1)$. It is easy to see that the total operation takes time $O(\log n)$ per segment inserted.
2. Delete all line segments that are stored in $END(x)$ from T . For each such segment s we locate its two neighbors s_1 and s_2 in T . As these segments now become neighbors we compute their intersection and, if it exists and lies to the right of x we store it in the appropriate $INTER$ bag. This takes time $O(\log n)$ per segment deleted.
3. For each intersection in $INTER(x)$ we do the following. Let the intersection be between s_1 and s_2 . s_1 and s_2 must be neighbors in T . Let s_0 be the other neighbor of s_1 and s_3 be the other neighbor of s_2 . See figure 6 for the situation. First of all we report the intersection. Secondly, we switch s_1 and s_2 in T . As s_0 and s_1 are no longer neighbors, we remove their intersection (if any)

from the appropriate INTER bag. Similar for s_2 and s_3 . We also remove the intersection between s_1 and s_2 . As s_0 and s_2 have now become neighbors we compute their intersection and, if it exists and HAS NOT BEEN TREATED BEFORE, put it in the appropriate INTER bag (which might be $\text{INTER}(x)$). Similar for s_1 and s_3 . When we keep a pointer from each intersection to the appropriate place in T and the other way round, this total operation can be carried out in time $O(1)$ per intersection.




Figure 6: Intersection of line segments.

The main problem with the above method is that intersections between x and $x + 1$ are NOT treated in left to right order but in the order in which they appear in $\text{INTER}(x)$. Therefore it is very important that we only insert intersections in the INTER bags if they have not been treated before. Otherwise the algorithm could loop forever, reporting the same intersections over and over again. To avoid this we have to check whether the intersection has indeed not been treated before. To this end we use an idea of Myers [18]. Assume we want to add an intersection between s_0 and s_2 to $\text{INTER}(x)$ (when it has to be added to another INTER bag this is always correct). Now look at the order in which s_0 and s_2 have to be at position $x + 1$. If they already are in this order we must have treated their intersection before. Otherwise we can insert it in $\text{INTER}(x)$. This test will take only $O(1)$ time.

We still have to prove that the method correctly reports all intersection

between x and $x + 1$ and that T will have the right shape afterwards. This follows from the following two observations: (i) When an intersection is reported it does exist and has not been reported before. (ii) When $\text{INTER}(x)$ is empty this means that each two neighbors in T either do not intersect between x and $x + 1$ or that their intersection has been treated. For details see [13].

Theorem 8.1 *Given a set of n line segments on U^2 , we can report all intersections between line segments in time $O(k + U + n \log n)$ using space $O(n + U)$, where k is the number of reported intersections.*

Proof. The tree T uses only $O(n)$ storage. The bags need storage $O(n + U)$ because there are $3U$ bag that store $2n$ endpoints and at most n intersections.

The query time follows from the above observations. Each insertion and deletion takes time $O(\log n)$, each intersection takes time $O(1)$ and we have to spend at least $O(1)$ time for each bag. \square

When U is large, this method is not very good. In [13] a second method is presented that takes time $O((k + n) \log \log U + n \log n)$.

9. Conclusions and extensions

In this paper we have demonstrated that techniques from computational geometry, combined with the property that objects have integer coordinates, can lead to simple but efficient solutions for many geometric problems that are related to computer graphics. We believe that these techniques should be applicable in computer graphics applications and can lead to powerful, simple methods that are also very efficient.

Although the presentation of the algorithms does not go into implementation details we think that, because of the simplicity of the solutions, it should not be much work to come to efficient implementations. We realise that in practise, problems normally are not as simple as stated in this paper. Still we believe that the techniques presented will be useful.

It might be an interesting area of further research whether special hardware could be constructed that deals with problems on a grid. For example it should be possible to construct hardware implementations of bucket sorting and of the vanEmdeBoas tree for some fixed U . Such hardware would be independent of screen resolution and could considerably speed up the execution of the algorithms proposed in this paper.

This paper only lists a number of example problems with their solutions. Similar methods can be used to solve other geometric problems on a grid as well. For example, Karlsson and Munro [12] consider the problem of finding nearest neighbors on a grid and Müller [19] discusses the point location problem on a grid.

In Overmars [21,22] the range searching problem on a two-dimensional grid is solved in a (theoretically) very efficient way. Also multi-dimensional problems can be treated in an efficient way. See e.g. Karlsson and Overmars [14] and Gabow e.a. [8].

Theoretical, the techniques presented are not optimal. For example, as mentioned, sorting can be done more efficiently using a method of Kirkpatrick and Reisch [16]. Although this gives asymptotically faster results, in practise, due to the complexity of the methods, this will not improve the efficiency. Moreover, these complicated techniques tend to be hard to implement in hardware. Therefore we have chosen to present the methods using the more simple techniques.

References

1. A.M. Andrew, "Another efficient algorithm for convex hulls in two dimensions," *Info Proc. Lett.* **8**, pp. 108-109 (1979).
2. J.L. Bentley and T.A. Ottmann, "Algorithms for reporting and counting geometric intersections," *IEEE Trans. on Computers* **28**, pp. 643-647 (1979).
3. P. van Emde Boas, R. Kaas, and E. Zijlstra, "Design and implementation of an efficient priority queue," *Math. Systems Theory* **10**, pp. 99-127 (1977).
4. P. van Emde Boas, "Preserving order in a forest in less than logarithmic time and linear space," *Info. Proc. Lett.* **6**, pp. 80-82 (1977).
5. B. Chazelle, "Intersecting is easier than sorting," *Proc. 16th Annual ACM Symp. on Theory of Computing*, pp. 125-134 (1984).
6. H. Edelsbrunner and J. van Leeuwen, "Multidimensional data structures and algorithms: a bibliography," *Techn. Rep. F105*, Inst. f. Information Processing, TU Graz (1982).
7. M.L. Fredman, J. Komlos, and E. Szemerédi, "Storing a sparse table with $O(1)$ worst case access time," *J. ACM* **31**, pp. 538-544 (1984).
8. H.N. Gabow, J.L. Bentley, and R.E. Tarjan, "Scaling and related techniques for geometric problems," *Proc. 16th Annual ACM Symposium on Theory of Computing*, pp. 135-143 (1984).
9. R.L. Graham, "An efficient algorithm for determining the convex hull of a finite planar set," *Info. Proc. Lett.* **1**, pp. 132-133 (1972).
10. D.B. Johnson, "A priority queue in which initialization and queue operations take $O(\log \log D)$ time," *Math. Systems Theory* **15**, pp. 295-310 (1982).
11. R.G. Karlsson, "Algorithms in a restricted universe," *Ph.D. thesis, Techn. Rep. CS-84-50*, Dept. of Comp. Science, University of Waterloo (1984).
12. R.G. Karlsson and J.I. Munro, "Proximity on a grid," *Proc. 2nd Symp. on Theoretical Aspects of Comp. Science.*, pp. 187-196, Springer-Verlag Lect. Notes in Comp. Science 182 (1985).
13. R.G. Karlsson and M.H. Overmars, "Scanline algorithms on a grid," *Techn.*

- Rep. RUU-CS-86-18*, Dept. of Comp. Science, University of Utrecht (1986).
14. R.G. Karlsson and M.H. Overmars, "Normalised divide and conquer: A scaling technique for solving multi-dimensional problems,," *Techn. Rep. RUU-CS-86-19 (to appear in Inform. Proc. Lett.)*, Dept. of Comp. Science, University of Utrecht (1986).
 15. J.M. Keil and D.G. Kirkpatrick, "Computational geometry on integer grid," *Proc. 19th Annual Allerton Conference*, pp. 41-50 (1981).
 16. D.G. Kirkpatrick and S. Reisch, "Upperbounds for sorting integers on random access machines," *Theoretical Computer Science* **28**, pp. 263-276 (1984).
 17. D.G. Kirkpatrick and R. Seidel, "The ultimate convex hull algorithm?," *SIAM J. Computing* **15**, pp. 287-299 (1986).
 18. F.W. Myers, "An $O(E \log E + I)$ expected time algorithm for the planar segment intersection problem," *SIAM J. Computing* **14**, pp. 625-637 (1985).
 19. H. Müller, "Rastered point location," *Proc. Workshop on Graphtheoretic Concepts in Computer Science (WG85)*, pp. 281-293, Trauner Verlag (1985).
 20. M.H. Overmars, "Computational geometry on a grid: an overview," *Techn. Rep. RUU-CS-87-4 (to appear in the proceedings of the NATO ASI on theoretical aspects of computer graphics and CAD, Pisa, 1987)*, Dept. of Comp. Science, University of Utrecht (1987).
 21. M.H. Overmars, "Efficient data structures for range searching on a grid," *Techn. Rep. RUU-CS-87-2 (to appear in J. Algorithms)*, Dept. of Comp. Science, University of Utrecht (1987).
 22. M.H. Overmars, "Range searching on a grid," *Proc. Workshop on Graphtheoretic Concepts in Computer Science (WG85)*, pp. 295-305, Trauner Verlag (1985).
 23. F.P. Preparata and S.J. Hong, "Convex hulls of finite sets of points in two and three dimensions," *Comm. ACM* **2**, pp. 87-93 (1977).
 24. F.P. Preparata and M.I. Shamos, "Computational geometry, An introduction," *Springer-Verlag*, New York (1985).
 25. A. Schmitt, "Time and space bounds for hidden line and hidden surface algorithms," *Proc Eurographics 1981*, pp. 43-56, North Holland Publ. Co. (1981).
 26. D.E. Willard, "Log-logarithmic worst-case range queries are possible in space $\Theta(n)$," *Info. Proc. Lett.* **17**, pp. 81-84 (1983).
 27. D.E. Willard, "New trie data structures which support very fast search operations," *J. Comput. Syst. Sci.* **28**, pp. 379-394 (1984).