

Algorithms for Bichromatic Line-Segment Problems and Polyhedral Terrains¹

Bernard Chazelle,² Herbert Edelsbrunner,³ Leonidas J. Guibas,⁴
and Micha Sharir⁵

Abstract. We consider a variety of problems on the interaction between two sets of line segments in two and three dimensions. These problems range from counting the number of intersecting pairs between m blue segments and n red segments in the plane (assuming that two line segments are disjoint if they have the same color) to finding the smallest vertical distance between two nonintersecting polyhedral terrains in three-dimensional space. We solve these problems efficiently by using a variant of the segment tree. For the three-dimensional problems we also apply a variety of recent combinatorial and algorithmic techniques involving arrangements of lines in three-dimensional space, as developed in a companion paper.

Key Words. Computational geometry, Line-segment intersection, Segment trees, Lines in space, Polyhedral terrains, Deterministic and randomized algorithms.

1. Introduction. In this paper we study a variety of problems on the interaction between two sets of line segments in two and three dimensions. These problems can all be solved efficiently by using a variant of the well-known *segment tree* (see, e.g., [PS]). We begin by listing the problems and their solutions. More background information and references are given later in the sections that address the individual problems.

PROBLEM 1. Let B be a set of m pairwise disjoint “blue” line segments in the plane, let R be a set of n pairwise disjoint “red” line segments, and define $N = m + n$. Report all bichromatic pairs of intersecting line segments. We solve

¹ Work on this paper by the first author has been supported in part by the National Science Foundation under Grant CCR-9002352. Work by the second author was supported in part by the National Science Foundation under Grant CCR-8714565. The fourth author has been supported in part by the Office of Naval Research under Grant N0014-87-K-0129, by the National Science Foundation under Grant NSF-DCR-83-20085, by grants from the Digital Equipment Corporation and the IBM Corporation, and by a grant from the US–Israeli Binational Science Foundation.

² Department of Computer Science, Princeton University, Princeton, NJ 08544, USA.

³ Department of Computer Science, University of Illinois at Urbana–Champaign, Urbana, IL 61801, USA.

⁴ DEC Systems Research Center, Palo Alto, CA 94301, USA, and Computer Science Department, Stanford University, CA 94305-4055, USA.

⁵ Courant Institute of Mathematical Sciences, New York University, New York, NY 10012, USA, and School of Mathematical Sciences, Tel Aviv University, Tel Aviv, Israel.

this problem in time $O(N \log N + k)$ and linear storage, where k is the number of blue-red intersecting pairs (see Sections 2 and 3.1).⁶

PROBLEM 2. Given the same data as in Problem 1, count how many bichromatic pairs of intersecting line segments there are. The solution to this problem takes time $O(N \log N)$ and linear storage (Sections 2 and 3.1).

PROBLEM 3. Given the same data again and an integer k , find the k th bichromatic intersection point from the left. This point can be found in time $O(N \log N)$ and linear storage (Section 3.2).

PROBLEM 4. Given two convex polytopes, \mathcal{A} and \mathcal{B} , in three dimensions with m and n edges, $N = m + n$, compute the number of faces bounding their *Minkowski sum*, that is, the polytope $\{p + q \mid p \in \mathcal{A}, q \in \mathcal{B}\}$. This takes time $O(N \log N)$ and linear storage (Section 3.3).

PROBLEM 5. A *polyhedral terrain* is the graph of a real bivariate continuous function, or, equivalently, a piecewise linear continuous surface in three dimensions that meets every vertical line in exactly one point. Given two nonintersecting polyhedral terrains with m and n edges, $N = m + n$, find the smallest vertical distance between them. We give an algorithm that solves this problem in $O(N^{4/3+\epsilon})$ time and storage, for any $\epsilon > 0$ (Section 5).

PROBLEM 6. Given the same data as in Problem 5 but allowing the two terrains to intersect, compute their *upper envelope*, that is, pointwise maximum. For this problem we have an algorithm that takes time $O(N^{3/2+\epsilon} + k \log^2 N)$ and worst-case storage $O(N)$, where ϵ is an arbitrary positive constant and k is the number of edges of the resulting envelope (Section 6).

While Problems 1–4 can be solved using a variant of the segment tree plus some standard data structuring techniques, Problems 5 and 6 require additional techniques (developed in [CEG⁺]; see also below) which employ (derandomized) *random sampling* of the data.

The common theme of the solutions to the above problems is that they all exploit the same data structure, the *hereditary segment tree*, which is introduced in this paper. It is very similar to the standard segment tree, except that it stores two sets of line segments in a manner that makes it easy to keep track of their interactions. The data structure is described in detail in Section 2, where we demonstrate its applications to Problems 1 and 2. Using certain modifications and extension of the basic technique of Section 2 we improve the solutions to Problems 1 and 2 in Section 3 and we also give solutions to Problems 3 and 4.

Section 4 discusses another extension of the basic hereditary segment tree used in Sections 5 and 6 to solve Problems 5 and 6. Both problems ask questions about

⁶ Unless the base of a logarithm is explicitly noted it is assumed to be 2.

the interaction between two polyhedral terrains in three-dimensional space. We apply the hereditary segment tree to their xy -projections to obtain a compact representation of the way in which the terrain edges cross above or below one another. This enables us to break the problem into smaller subproblems, with the advantage that in each subproblem we can extend the edges to lines without changing the solution to the subproblem. These subproblems, among several related ones, have been studied in a companion paper [CEG⁺] that deals with combinatorial and algorithmic problems for arrangements of lines in three dimensions. Combining the techniques developed in [CEG⁺] with our hereditary segment tree, we obtain reasonably efficient solutions for Problems 5 and 6.

2. The Basic Hereditary Segment Tree. In this section we introduce our basic data structure, and apply it to Problems 1 and 2. Recall that in these problems we have two collections, $B = \{b_1, b_2, \dots, b_m\}$ and $R = \{r_1, r_2, \dots, r_n\}$, of m "blue" and n "red" line segments in the plane, so that any two line segments of the same color are disjoint. Our goal is to compute (or just count) all bichromatic pairs of intersecting line segments. Define $N = m + n$. The straightforward application of our basic techniques leads to an algorithm that reports all intersecting bichromatic pairs in time $O(N \log^2 N + k)$, where k is the number of such pairs, and uses $O(N \log N)$ storage, and to another algorithm that counts these pairs in time $O(N \log^2 N)$ and uses $O(N \log N)$ storage. Using additional methods we show in Section 3.1 how to remove a factor $\log N$ from each bound.

Our first step is to construct a segment tree, \mathcal{T} , on the interval decomposition of the x -axis induced by the x -coordinates of the endpoints of the given line segments. More specifically, \mathcal{T} is defined as follows. Let us assume that the x -coordinates of the $2N$ endpoints are all distinct; if they are not we can get this property by simulating an arbitrarily small perturbation of the x -coordinates. As a result, the $2N$ x -coordinates decompose the x -axis into $2N + 1$ (atomic) intervals. We define these intervals as closed sets. \mathcal{T} is a minimum height ordered binary tree whose i th leaf corresponds to the i th atomic interval from the left. Each interior node μ represents an interval I_μ , which is the union of the intervals associated with the leaves of the subtree rooted at μ . Alternatively, we can think of μ as representing the vertical slab $\sigma_\mu = \{(x, y) \mid x \in I_\mu, y \in \mathcal{R}\}$. In the standard segment tree, each node μ (internal and external) has an associated list L_μ , consisting of all line segments s , with the property that the vertical projection of s contains I_μ but does not contain I_κ , where κ is the parent of μ . In our version we maintain four lists with each node μ . First we split L_μ into two *standard lists*, B_μ and R_μ , consisting of the blue and red line segments in L_μ . Furthermore, we associate with μ two additional *hereditary lists* B_μ^* and R_μ^* . The list B_μ^* contains the blue line segments stored in L_ν for all proper descendants ν of μ ; R_μ^* is the red analog of B_μ^* . In other words, whenever we store a line segment s in some B_ν or R_ν , we also store the line segment in the hereditary list of each proper ancestor μ of ν (see Figure 2.1).

Even though a line segment s can be stored in the standard lists of more than one descendant of a node μ , we still keep only one copy of s in the appropriate




Fig. 2.1. The segment s is stored in the standard lists of all marked nodes and in the hereditary lists of all ancestors of the marked nodes. These ancestors lie on two paths in the tree.

hereditary list of μ . For the applications that follow, it is helpful to regard each element of a standard or hereditary list of a node μ as representing the subsegment $s \cap \sigma_\mu$ of the corresponding line segment s . As suggested by this interpretation, we refer to line segments stored in the standard lists as *long segments* and to those stored in the hereditary lists as *short segments*.

This is our basic hereditary segment tree, which is borrowed from [C1]. In some of the applications given below we modify it slightly to improve its performance. Modifications different from the ones to be described have been also used in [VW] for point enclosure problems.

In a standard segment tree, the total size of all the lists L_μ is $O(N \log N)$. Our first observation is that this is also the total size of all hereditary lists. Indeed, each line segment s is stored in $O(\log N)$ standard lists. The interior nodes that store s in their hereditary lists form two paths both starting at the root of \mathcal{T} ; thus s is stored in at most $O(\log N)$ hereditary lists and therefore $O(\log N)$ times altogether (see also [VW]).

We next show how to use the hereditary segment tree to solve Problems 1 and 2. We first construct the tree \mathcal{T} for the given collections B and R in time $O(N \log N)$, using standard techniques. At each node μ of \mathcal{T} we compute (or, in Problem 2, just count) the intersections between

- (i) a line segment in B_μ and one in R_μ ,
- (ii) a line segment in B_μ and one in R_μ^* , and
- (iii) a line segment in B_μ^* and one in R_μ .

We refer to type (i) intersections as “long–long intersections,” and to the two other types as “long–short intersections.” According to our interpretation, in all three cases we consider only the portions of the line segments within σ_μ .

† Note that by repeating this step for all nodes μ of \mathcal{T} we detect all bichromatic intersecting pairs; moreover, each pair is detected exactly once. Indeed, suppose $b_i \cap r_j \neq \emptyset$. Then there is a unique leaf λ of \mathcal{T} such that the point $p = b_i \cap r_j$ lies in σ_λ (assuming that the x -coordinate of p is different from all endpoints of atomic intervals). Moreover, b_i is stored in the standard list of some unique node along the path from λ to the root, and so is r_j . Let these nodes be κ_i and κ_j , respectively.

If $\kappa_i = \kappa_j$, then p will be detected as a long-long intersection at κ_i ; otherwise, it will be detected as a long-short intersection at either κ_i or κ_j , whichever is higher in \mathcal{F} . This observation is significant because when we deal with problems that involve counting intersections, rather than reporting them, we can simply sum up the numbers computed at the nodes.

COMPUTING LONG-LONG INTERSECTIONS. Let $m_\mu = |B_\mu|$ and $n_\mu = |R_\mu|$. Sort each of these lists in increasing y -order; since the line segments in each list are pairwise disjoint, each of these orders is fixed throughout the slab σ_μ . Let $(b_1, b_2, \dots, b_{m_\mu})$ and $(r_1, r_2, \dots, r_{n_\mu})$ be the two sorted lists. Next we merge the two lists twice, once according to their order along the left boundary of σ_μ , and once according to their order along the right boundary. This allows us to determine for each blue line segment $b \in B_\mu$ two indices, i_b and j_b , such that b lies between r_{i_b} and r_{i_b+1} along the left boundary and between r_{j_b} and r_{j_b+1} along the right boundary. However, then the red line segments intersecting b within σ_μ are precisely $r_{i_b+1}, \dots, r_{j_b}$ if $i_b < j_b$, and $r_{j_b+1}, \dots, r_{i_b}$ otherwise. We can therefore enumerate all long-long intersections within σ_μ in time $O(m_\mu \log m_\mu + n_\mu \log n_\mu + k_\mu)$, where k_μ is the number of such intersecting pairs. Alternatively, we can calculate the number k_μ in time $O(m_\mu \log m_\mu + n_\mu \log n_\mu)$.

COMPUTING LONG-SHORT INTERSECTIONS. Without loss of generality, consider only intersections between B_μ^* and R_μ . Let $m_\mu^* = |B_\mu^*|$ and $n_\mu = |R_\mu|$. As above, we sort the list R_μ in ascending y -order and write $(r_1, r_2, \dots, r_{n_\mu})$ for the resulting sequence. Then, for each $b \in B_\mu^*$, we locate the two endpoints of b (within σ_μ) in the list R_μ using binary search twice. The line segments of R_μ that lie between these two locations are precisely those that intersect b within σ_μ . Thus, in time $O((m_\mu^* + n_\mu) \log n_\mu + k_\mu^*)$ we can find all k_μ^* intersecting pairs between B_μ^* and R_μ . To count these pairs takes time $O((m_\mu^* + n_\mu) \log n_\mu)$. Intersections between R_μ^* and B_μ are handled symmetrically.

Summing up these bounds over all nodes μ of \mathcal{F} , taking also into account the cost of constructing \mathcal{F} , and observing that the total size of all lists is $O(N \log N)$, we obtain the bounds on reporting and counting bichromatic pairs of intersecting line segments mentioned at the beginning of this section.

3. Modifications, Extensions, and Applications. Having introduced our basic data structure, we continue our presentation with a number of modifications, extensions, and further applications, proceeding down the list of problems given in the introduction. Section 3.1 presents the additional techniques needed to improve the performance of the two intersection algorithms of Section 2, and Sections 3.2 and 3.3 discuss Problems 3 and 4.

3.1. Improving the Bichromatic Intersection Algorithms. In this section we show how to improve both the storage and the time bound of the two intersection algorithms in Section 2 using standard data structuring techniques.

Let us first address the issue of reducing the storage requirements. This is done by noting that for most of the listed problems (including Problems 1 and 2) we do not need to maintain the entire segment tree, but only process it node by node. We therefore traverse the tree in preorder so that at any time we maintain only a single path of \mathcal{T} . Assuming we have the standard lists, B_κ and R_κ , as well as the hereditary lists, B_κ^* and R_κ^* , at a node κ , we can construct the lists of a child μ simply by scanning B_κ^* and R_κ^* . To initialize those lists at the root ρ we note that $B_\rho = R_\rho = \emptyset$ and $B_\rho^* = B$ and $R_\rho^* = R$. Notice that a line segment in B_κ and R_κ can be neither in a standard nor in a hereditary list of μ . Let us now argue that the total size of all lists along a path is at most proportional to N . This is easy to see for the standard lists because a line segment can be in at most one standard list along a single path. The argument for the hereditary lists is based on the well-balancedness of the segment tree (and on the assumption that the x -coordinates of the endpoints are pairwise different). If a line segment is in B_κ^* or in R_κ^* , then at least one of its endpoints lies in the vertical slab σ_κ of κ . However, then the number of endpoints in σ_κ is one less than the number of leaves in the subtree rooted at κ . Since the number of leaves in the subtrees of nodes on a path decreases geometrically, the total count and thus the total size of the hereditary lists of these nodes is $O(N)$.

To cut down the time from $O(N \log^2 N)$ to $O(N \log N)$ we need to achieve two goals. One is to avoid having to sort the standard lists at each node of \mathcal{T} , and the other is to speed up the search of elements of hereditary lists in the standard lists.

To address the first issue, we begin by finding a linear extension of the following relation on the blue line segments: $b_i < b_j$ if the x -projections of b_i and b_j overlap and b_j lies above b_i along a vertical line that intersects both. The relation can be computed in time $O(m \log m)$ by a simple left to right sweep, and a linear extension can then be found in time $O(m)$. We sort R in a similar manner. Then we do the preorder traversal and the construction of the lists as described above. Since the standard lists and the hereditary lists of a node μ are derived by scanning B_κ^* and R_κ^* , where κ is the parent of μ , they are in sorted order already. Doing it this way, we maintain the sortedness of all lists throughout the traversal. The total time consumed by this procedure is proportional to the total size of all lists which is $O(N \log N)$ as argued before.

The second step in speeding up the binary search can be achieved by "fractionally cascading" the standard lists (see [CG] for a complete description of the data structuring technique called *fractional cascading*). Here is what we mean by this. As we go down the current path we maintain each standard list in a padded-up form, so that it also contains some elements of the standard lists of ancestor nodes along the path. As we go down from a node κ to one of its children μ , we take every fourth element of the (padded) list B_κ , pass down these elements to μ and merge them with B_μ . Similar action is performed on R_κ and R_μ . These operations neither increase the storage nor the time of the algorithm by more than a constant factor. Note that the fractional cascading technique is performed on-the-fly, which is somewhat different from the standard static version of the technique (as described in [CG]). Also it proceeds downward, which is the reverse

of the more standard upward direction generally used when cascading on trees.

The benefits of the cascading technique are reaped when we back up from μ to κ . We can assume inductively that each endpoint of any (short or long) red line segment $r \in R_\mu^* \cup R_\mu$ has been located among the line segments in the standard blue list B_μ . Since this list also contains a portion of the list B_κ , it is easy, by maintaining appropriate pointers between these lists and by investing an additional constant time per endpoint, to locate all these endpoints among the line segments in B_κ as well. Similar speed-up is achieved in a completely symmetric manner for locating the blue endpoints in the list R_κ . Notice that we thus construct the lists in preorder and evaluate them in an order which is similar to postorder. More specifically, the evaluation at a node κ is done in two steps, once when we back up from its left child and then again when the recursion returns from its right child. The preceding considerations imply our first main result.

THEOREM 3.1. *Given a set of m blue and a set of n red line segments in the plane, $N = m + n$, such that no two line segments with the same color intersect.*

- (i) *The k bichromatic pairs of intersecting line segments can be reported in time $O(N \log N + k)$ and storage $O(N)$.*
- (ii) *The bichromatic pairs of intersecting line segments can be counted in time $O(N \log N)$ and storage $O(N)$.*

The first of these two results is not new; indeed, the algorithm of Mairson and Stolfi [MaS] achieves the same asymptotic bounds, and the time-complexity is also matched by the more general algorithm of Chazelle and Edelsbrunner [CE]. The second result is new. We note that the best-known algorithms for counting intersections between line segments are much inferior to the results of the theorem: they run in time close to $O(n^{4/3})$ [A1], [GOS], [C2], [C3].

3.2. Order Statistics for Bichromatic Intersections. We next turn to Problem 3. Let B and R be two collections of m blue and n red line segments ($N = m + n$) with no intersecting monochromatic pairs. Given an integer k , we wish to find the x -coordinate $\xi(k)$ of the intersection point that has the property that there are exactly $k - 1$ bichromatic intersecting pairs whose intersection points lie to the left of the vertical line $x = \xi(k)$. We assume that the x -coordinates of the intersection points are pairwise distinct as well as distinct from the endpoints of atomic intervals; otherwise, this property can be simulated [EM]. We also assume that k is not larger than the total number of bichromatic intersecting pairs; we already know how to compute the latter number in time $O(N \log N)$ and linear storage.

An efficient solution has been recently obtained for the related problem where each line segment is a line and there is no bichromatic condition. To find $\xi(k)$ in this case takes time $O(N \log N)$ as noted by Cole *et al.* [CSSS]. In contrast, very little is known about the monochromatic case for line segments. The most efficient known way just to compute the total number of intersecting pairs runs in time

$O(N^{4/3}(\log N)^{1/3})$ and linear storage [C2]. However, with the guidance of the hereditary segment tree, we can solve the problem efficiently in the bichromatic case as follows.

Our first step is to locate a leaf, λ , of the tree \mathcal{T} such that σ_λ contains the k th intersection point. The simplest way to find λ is by binary search over the x -coordinates of the endpoints of the given line segments. At each step of the search we are at some x -coordinate ξ , and our goal is to count the number of intersections to the left of the line $x = \xi$. This is done by clipping each line segment to its portion to the left of this line and by applying our solution to Problem 2 to the clipped collection. It follows that we can find λ in time $O(N \log^2 N)$ and linear storage.

Suppose there are $k_0 < k$ intersections to the left of σ_λ , and define $k_1 = k - k_0$. We can now collect all line segments that cross the slab σ_λ into a set Λ and extend these line segments to full lines. Now we apply the algorithm of [CSSS] to find the $(k_1 + k_2)$ th leftmost intersection in the arrangement of these lines, where k_2 is the number of intersections of these lines to the left of σ_λ (which is easily computed in time $O(N \log N)$). This gives us an $O(N \log^2 N)$ algorithm for finding the k th leftmost bichromatic intersection.

Agarwal [A2] has suggested to us a way to cut down the running time to $O(N \log N)$. The bottleneck is the computation of the leaf λ . This can be done by binary search in \mathcal{T} , provided that we can determine quickly how many bichromatic intersections lie in a given slab σ_μ . Suppose that we can do that in time $O(h_\mu + k_\mu d_\mu \log N)$, where

- (i) h_μ is the sum of the sizes of all the lists (standard or hereditary, blue or red) stored at μ or at any of its ancestors,
- (ii) k_μ is the total size of the hereditary lists at μ (i.e., $k_\mu = |B_\mu^*| + |R_\mu^*|$), and
- (iii) d_μ is the number of ancestors of μ .

From our previous discussion, we know that $h_\mu = O(N)$, and therefore the sum $\sum_\mu h_\mu$, taken over all nodes μ on the path π from the root to λ , is $O(N \log N)$. Similarly, k_μ cannot exceed the number of endpoints in the slab σ_μ . Since this number decreases geometrically as we go down the path π , we have

$$\sum_{\mu \in \pi} k_\mu d_\mu \log N \leq (\log N) \sum_{i \geq 0} (i+1)N/2^i = O(N \log N),$$

and therefore we can find λ , and hence the k th leftmost intersection, in $O(N \log N)$ time.

It remains now to show how to compute the number of bichromatic intersections in σ_μ in time $O(h_\mu + k_\mu d_\mu \log N)$. Note that by using the algorithm of Theorem 3.1(ii), we can assume that the intersections in $B_\mu \cup R_\mu \cup B_\mu^* \cup R_\mu^*$ have already been counted. So, it suffices to count the intersections in σ_μ from among $B_\nu \cup R_\nu$, $B_\nu^* \cup R_\nu$, and $R_\nu^* \cup B_\nu$, where ν is an ancestor of μ .

From the ordering of the segments in B_ν and R_ν , we can compute their ordering along the bounding lines of σ_μ , and hence, in $|B_\nu| + |R_\nu|$ time, derive the number of intersections in $B_\nu \cup R_\nu$. Next, consider the intersections in $B_\nu \cup R_\nu^*$ (the case

of $B_v^* \cup R_v$ is similar). The segments of R_v^* fall into two categories: those with an endpoint in σ_μ and those crossing σ_μ all the way through. To handle the latter kind is exactly the same as what we just described with B_v and R_v , and therefore can be done in $|B_v| + |R_v^*|$ time. For the other kind of segments we must perform a binary search in B_v with respect to their endpoints (with the segments clipped to within σ_μ). Since these segments also appear in R_μ^* , we find that the time spent at node v is at most proportional to $|B_v| + |R_v| + |B_v^*| + |R_v^*| + k_\mu \log N$. Summing over all ancestors v of μ , we find that, indeed, the number of bichromatic intersections in σ_μ can be determined in time $O(h_\mu + k_\mu d_\mu \log N)$. Again, the segment tree need not be computed in its entirety. Only linear space is required to apply the algorithm of Theorem 3.1(ii), after which, it suffices to maintain the segment tree lists alongside the path π (as it is discovered top-down). This limits the space requirement to $O(N)$.

THEOREM 3.2. *The k th leftmost bichromatic intersection for a collection of m blue pairwise disjoint line segments and a collection of n red pairwise disjoint line segments, $N = m + n$, can be computed in time $O(N \log N)$ and linear storage.*

3.3. The Minkowski Sum of Two Convex Polytopes. Let \mathcal{A} and \mathcal{B} be two convex polytopes in three-dimensional space with m and n edges, respectively. The Minkowski (vector) sum $\mathcal{C} = \mathcal{A} + \mathcal{B}$ is defined as $\{p + q \mid p \in \mathcal{A}, q \in \mathcal{B}\}$. As is well known, \mathcal{C} is also a convex polytope, and an alternative way of defining (or constructing) \mathcal{C} is as follows:

Consider the unit sphere S^2 in space, and define on it two spherical maps, $M_{\mathcal{A}}$ and $M_{\mathcal{B}}$, as follows. For each $u \in S^2$ let $f_{\mathcal{A}}(u)$ denote the vertex, edge, or facet of \mathcal{A} that is the intersection of \mathcal{A} with its supporting plane having an outward normal direction u . The map $M_{\mathcal{A}}$ is the corresponding decomposition of S^2 into maximal portions on each of which $f_{\mathcal{A}}$ is invariant. The map $M_{\mathcal{A}}$ is known as the *Gaussian image* of \mathcal{A} . The map $M_{\mathcal{B}}$ is defined in a similar manner. Note that vertices of $M_{\mathcal{A}}$ are the directions of outer normals to the faces of \mathcal{A} , edges of $M_{\mathcal{A}}$ are portions of great circles of directions perpendicular to edges of \mathcal{A} , and faces of $M_{\mathcal{A}}$ correspond to vertices of \mathcal{A} . Similar properties hold for $M_{\mathcal{B}}$. To obtain \mathcal{C} , we overlay the two maps and the resulting decomposition of S^2 is $M_{\mathcal{C}}$. The faces, edges, and vertices of \mathcal{C} and their adjacency structure can then be read off $M_{\mathcal{C}}$ using the correspondence just explained.

Note that \mathcal{C} can have up to $\Theta(mn)$ facets that correspond to intersections between edges of $M_{\mathcal{A}}$ and of $M_{\mathcal{B}}$. As a matter of fact, each such facet is the Minkowski sum of the corresponding edges of \mathcal{A} and \mathcal{B} .

$M_{\mathcal{C}}$ can be computed in time $O(N + k)$, where k is the number of intersecting edge pairs using an algorithm by Guibas and Seidel [GS]. In order to apply their algorithm which overlays two subdivisions of the plane into convex regions, we first have to project $M_{\mathcal{A}}$ and $M_{\mathcal{B}}$ onto a plane. If we carry the projection from the origin of S^2 the circle arcs map to straight line segments. This way only one hemisphere is captured at a time and we have to repeat the algorithm for the other

hemisphere. Each subdivision of the plane can also be viewed as a collection of pairwise disjoint line segments (if they are taken as relatively open sets), so we can apply our solutions to Problem 2 which gives the following result.

THEOREM 3.3. *Given two convex polytopes with m and n edges in three-dimensional space, $N = m + n$, the number of facets bounding their Minkowski sum can be computed in time $O(N \log N)$ and linear storage.*

The Two-Level Hereditary Segment Tree. Next we modify the hereditary segment tree by adding another level of preprocessing. More specifically, we process the lists of the nodes into certain tree structures. This allows us to reduce problems for blue and red line segments to a moderately large number of such problems where each blue line segment intersects each red line segment.

So let B be a set of m pairwise disjoint blue line segments, let R be a set of n pairwise disjoint red line segments, and define $N = m + n$, as usual. Construct the hereditary segment tree, \mathcal{T} , for B and R . Our goal is to reduce the possible types of interactions between the line segments in the standard and the hereditary lists at any node μ of \mathcal{T} as mentioned above. In particular, we would like to divide the set of blue line segments at μ into (not necessarily disjoint) subsets B_i and the set of red line segments at μ into matching subsets R_i so that, for each i ,

- (*) every line segment in B_i intersects every line segment in R_i within the slab σ_μ .

To reduce bichromatic interactions to the kind required, we proceed as follows. For each node μ , we represent the standard list B_μ (and R_μ) by a minimum height ordered binary tree \mathcal{B}_μ (\mathcal{R}_μ). The leaves of \mathcal{B}_μ are associated with the blue line segments in B_μ , in the same order, and for each node τ of \mathcal{B}_μ we define C_τ as the list of blue line segments associated with the leaves below τ ; this list is explicitly stored at τ . Symmetrically, each node τ of \mathcal{R}_μ stores the list D_τ of red line segments associated with the leaves below τ . Consider now, for example, the blue list B_μ . For each red line segment $r \in R_\mu \cup R_\mu^*$ find the sublist of B_μ consisting of line segments that intersect r . Since this is a contiguous sublist, it is the union of $O(\log m)$ lists C_τ stored at nodes of \mathcal{B}_μ , and so we pair r with each of them. In this manner, each node τ acquires a list D_τ of red line segments that cross all of its blue line segments within σ_μ . Thus, (C_τ, D_τ) satisfies condition (*).

Let us do an analysis of the number of pairs (C_τ, D_τ) generated this way, and of the total size of the lists C_τ and D_τ . The number of pairs (C_τ, D_τ) is equal to the number of nodes of the binary trees that represent the standard lists which, in turn, is at most twice the total size of standard lists attached to the nodes of the segment tree. We noted before that the total size of standard lists is $O(N \log N)$ which shows that the number of subproblems of type (*) is $O(N \log N)$. Each line segment, whether blue or red, belongs to at most $2 \log N$ sets C_τ or D_τ of a single tree representing a standard list. It follows that the total size of these sets is at most a factor $2 \log N$ bigger than the total size of standard and hereditary lists. In other words, the total size of the sets C_τ and D_τ is $O(N \log^2 N)$.

5. The Smallest Distance Between Polyhedral Terrains. Let Σ_1 and Σ_2 be two polyhedral terrains with m and n edges so that Σ_2 lies completely above Σ_1 , and define $N = m + n$, as usual. We wish to determine the smallest vertical distance between Σ_1 and Σ_2 , that is, the length of the shortest vertical line segment connecting a point on Σ_1 to a point on Σ_2 . Chazelle and Sharir [CS] give an algorithm for this problem which runs in time $O(N^{2-\delta})$ for some very small $\delta > 0$; in this section we obtain an improved solution that runs in time $O(N^{4/3+\epsilon})$, for any $\epsilon > 0$.

We first observe that the desired smallest vertical distance is attained either between a vertex of Σ_1 and a facet of Σ_2 , between a vertex of Σ_2 and a facet of Σ_1 , or between a point of an edge of Σ_1 and a point of an edge of Σ_2 . In the last case, the vertical projections onto the xy -plane of these two edges intersect. It is this case that consumes the largest amount of time.

We begin by projecting both Σ_1 and Σ_2 vertically onto the xy -plane and obtain two straight-edge plane maps, \mathcal{M}_1 and \mathcal{M}_2 . Then we apply standard point-location techniques (see, e.g., [EGS]) to locate each vertex of \mathcal{M}_1 in \mathcal{M}_2 and vice versa. This allows us to compute the smallest vertical distance between a vertex of one terrain and a facet of the other in total time $O(N \log N)$.

However, we still have to consider the edge-edge intersections, whose number might be quadratic. To this end, regard \mathcal{M}_1 as a collection of pairwise disjoint (relatively open) blue line segments and \mathcal{M}_2 as a similar collection of red line segments, and construct a hereditary segment tree, \mathcal{T} , for the two collections. At each node μ of \mathcal{T} we have the two usual standard lists, B_μ and R_μ , and the two hereditary lists, B_μ^* and R_μ^* . As before, by checking B_μ against R_μ , B_μ^* against R_μ , and R_μ^* against B_μ , for all nodes μ of \mathcal{T} , we encounter every intersection between edges of \mathcal{M}_1 and \mathcal{M}_2 .

Similar to Section 4, our goal is to reduce these interactions to a collection of interactions between subsets of blue line segments and of red line segments, with the property that whenever we test a blue subset B_i against a red subset R_i , condition (*) of Section 4 holds. For every such pair of subsets, the smallest vertical distance between the corresponding edges of Σ_1 and Σ_2 can be obtained by extending all these edges to full lines in three-dimensional space. Doing this, we do not introduce any new bichromatic intersections in the projection. Thus, we can apply the so-called towering algorithm given in [CEG⁺] which, given two sets, B and R , of lines in three-dimensional space, verifies that all lines in B lie above all lines in R . Also, if the verification is successful, it computes the smallest vertical distance between a line in B and a line in R . All this is done using a simple randomized procedure whose expected time and storage requirements are $O(M^{4/3+\epsilon})$, for any $\epsilon > 0$, where M is the total size of the two subsets. The algorithm can be derandomized without asymptotic loss in performance, using the techniques of [C2], [M1], and [M2].

To reduce bichromatic interactions to the kind required above, we proceed as in Section 4, constructing the pairs (C_τ, D_τ) , for τ a node of \mathcal{B}_μ or \mathcal{R}_μ at a node μ of \mathcal{T} , in the manner described there. Because the pairs (C_τ, D_τ) satisfy condition (*) we can apply to it the towering algorithm of [CEG⁺] mentioned earlier. As noted in Section 4, the total size of all these lists is $O(N \log^2 N)$. It follows that we can

determine the smallest vertical distance between any pair of edges of Σ_1 and Σ_2 in time and storage $O(N^{4/3+\epsilon})$, for any $\epsilon > 0$. In summary, we have shown the following result.

THEOREM 5.1. *The smallest vertical distance between two nonintersecting polyhedral terrains with m and n edges, $N = m + n$, can be computed in time and storage $O(N^{4/3+\epsilon})$, for any $\epsilon > 0$.*

6. The Upper Envelope of Polyhedral Terrains. Let Σ_1 and Σ_2 be two polyhedral terrains with m and n edges as in the preceding section but without the disjointness assumption. Our goal is to compute the *upper envelope*, Σ , of Σ_1 and Σ_2 , that is, if we think of Σ_1 and Σ_2 as the graphs of continuous, piecewise linear, bivariate functions, then Σ is the pointwise maximum of them. Special cases and variants of this problem, mainly when one of the terrains is convex, were treated in [MeS] and [S] where solutions that are more efficient than the solution in this section are given. Here we study the general problem. For simplicity we assume non-degenerate position throughout.

Clearly, Σ is also a polyhedral terrain. Each vertex of Σ is either a vertex of Σ_1 , a vertex of Σ_2 , or the intersection between an edge of one terrain and a facet of the other. Similarly, each edge of Σ is either a portion of an edge of Σ_1 or Σ_2 , or it is the intersection between a facet of Σ_1 and a facet of Σ_2 . This implies that k , the number of vertices, edges, and facets of Σ , is $O(mn)$, and, indeed, this upper bound is tight. Our goal is to compute Σ in an output-sensitive manner. We achieve this with an algorithm that runs in time $O(N^{3/2+\epsilon} + k \log^2 N)$, where $N = m + n$.

The above observations make it clear that Σ is easy to construct once we have computed all intersections between edges of Σ_1 and facets of Σ_2 and vice versa. As a matter of fact, it suffices to produce one point on each connected component of the intersection $\Sigma_1 \cap \Sigma_2$. From each such point the component of the intersection can be obtained by a straightforward tracing procedure (see [MeS] or [S] for details). Now, the upper envelope can be constructed by putting things together in the right way. We therefore concentrate on finding all edge-facet intersections between the two terrains. We first describe the data structure used to store Σ_1 and Σ_2 and then show how to use it to find edge-facet intersections.

The Hereditary Segment Tree for Σ_1 and Σ_2 . Project Σ_1 and Σ_2 along the z -axis onto the xy -plane to get two planar subdivisions. For technical reasons we decompose each region of the subdivisions into trapezoids by drawing, from each vertex v , a vertical edge (parallel to the y -axis) upward until it hits the first edge above v and one downward until it hits the first edge below v . With this modification we get two trapezoidal subdivisions, Δ_1 from Σ_1 and Δ_2 from Σ_2 . Each trapezoid has at most two nonvertical edges which we classify as *lower* and *upper* depending on whether the trapezoid lies vertically above or below the edge. Notice, however, that the trapezoids are not necessarily glued edge-to-edge and therefore Δ_1 and Δ_2 might not always be cell complexes. Let B be the set of lower and upper edges of all trapezoids in Δ_1 , and define R as the set of lower




Fig. 6.1. The projection of a terrain onto the xy -plane is a planar subdivision. Each region is cut into trapezoids with unique upper and lower edges. Overlapping edges of neighboring trapezoids are interpreted as nonintersecting.

and upper edges of all trapezoids in Δ_2 . No two edges in B (and in R) cross but there can be overlapping edges. In fact, each edge of the projection of Σ_1 (and of Σ_2) is now represented by two sequences of edges with the union of all edges in either sequence being the old edge. Thus, some of the edges in B (and in R) are bound to overlap. Still, the effect of the overlap is uncritical and we can pretend that two overlapping edges are disjoint, for example, by conceptually shrinking each trapezoid by an infinitesimal amount (see Figure 6.1).

So we think of B as a set of $O(m)$ pairwise disjoint blue line segments, and of R as a set of $O(n)$ pairwise disjoint red line segments. The nice effect of cutting regions into trapezoids is that the lower and the upper edges of a single trapezoid are stored in exactly the same standard and hereditary lists when we construct the hereditary segment tree for B and R . Thus, the line segments in the lists come in pairs, and the two line segments of a pair delimit the same trapezoid. As the reader might expect we indeed store B and R in their hereditary segment tree. However, for reasons that will become clear later we do not use a *binary* tree as usual but rather a tree where every node has δ children, where δ is a parameter to be determined later. We use a minimum height tree so its height is $O(\log_\delta N)$ and each edge is stored in $O(\delta \log_\delta N)$ standard and $O(\log_\delta N)$ hereditary lists.

Having constructed the hereditary segment tree in this manner, we next augment it, as in the preceding section, with pairs (C_τ, D_τ) as follows. Let B_μ and R_μ be the standard lists at some node μ of \mathcal{T} . We store B_μ (and R_μ) by a minimum height ordered binary tree whose leaves store the pairs of line segments, each one delimiting a common trapezoid; as in Section 4 we call this tree \mathcal{B}_μ (and \mathcal{R}_μ). If τ is a node of either tree, then C_τ and D_τ are lists of line segments of opposite color so that each line segment in C_τ intersects every line segment in D_τ .

What are the important properties of this data structure? Because the line segments in any list B_μ (or R_μ) come in matching pairs, we can think of the list as a sequence of trapezoids each spanning the vertical strip σ_μ (see Figure 6.2). An endpoint of a line segment $r \in R_\mu \cup R_\mu^*$ is either within such a trapezoid (or on its left or right side) or between two adjacent trapezoids. Similarly, a list C_τ is a sequence of trapezoids spanning σ_μ , but now we have the property that each line segment in D_τ intersects each line segment in C_τ (see Figure 6.2).




Fig. 6.2. To the left we see a sequence of four trapezoids in a vertical strip; each line segment of the opposite color intersects a contiguous sublist of their edges. To the right, the trapezoids are defined by a list C_r , in which case all line segments in D_r intersect all line segments in C_r .

How To Represent the Lists C_r . From now on let us concentrate on edge-facet intersections between red edges and blue facets decomposed into trapezoids. To find all such intersections we store each list C_r of blue line segments (or trapezoids, depending on which interpretation we prefer) in a data structure developed in [CEG⁺]. We recall the relevant properties of this data structure.

Given a set of M lines in three-dimensional space, it is shown in [CEG⁺] how to store them in a data structure of size $O(M^{2+\epsilon})$, for any $\epsilon > 0$, so that given any query line l it can be determined in $O(\log M)$ time whether l lies above all M lines, below all M lines, or neither. We call this data structure the *envelope structure* of the set of M lines. It takes time $O(M^{2+\epsilon})$ to construct it. Given a list C_r , we first go back to the edges of Σ_1 that project onto the edges in C_r , then extend these edges to full lines, and finally store the envelope structure of the resulting set of lines. Because each line segment in D_r intersects each line segment in C_r , its corresponding edge in Σ_2 lies above (or below) all edges of Σ_1 corresponding to line segments in C_r if and only if the same is true for their extensions to lines.

A major drawback of the data structure, as described so far, is that it takes $O(N^{2+\epsilon})$ time and storage. To improve the complexity we construct the envelope structure for a list C_r only if it contains at most $N^{(1-\epsilon)/2}$ line segments. A single list thus takes no more than $O(N)$ time and storage. We now set δ , the number of children of any node in the segment tree, equal to $\lfloor N^{\epsilon/2} \rfloor$. With this choice of δ all lists C_r together take only time $O(N^{3/2+\epsilon})$ for construction because each edge belongs to at most $O(\delta \log_\delta N) = O(N^{\epsilon/2})$ lists C_r , and the cost for each occurrence is $O(N^{(1+\epsilon)/2})$.

Finding Edge-Facet Intersections. Using the envelope structures representing the lists of blue trapezoids, C_r , we can now find intersections between red edges and blue facets. Recall that a red line segment, r (which is the projection of an edge of Σ_2 onto the xy -plane), is stored in the standard and hereditary lists of $O(\delta \log_\delta N) = O(N^{\epsilon/2})$ nodes of the segment tree. Each occurrence of r gives rise to $O(\log N)$ occurrences in lists D_r of the binary trees that represent the standard

lists of blue trapezoids at these nodes. To describe the algorithm let μ be a node of the segment tree with $r \in R_\mu \cup R_\mu^*$, and let τ be a node of \mathcal{B}_μ with $r \in D_\tau$. If r lies above or below all line segments in C_τ (within σ_μ), then we stop right away. Otherwise, we distinguish three cases.

1. If τ is a leaf, then test whether the edge corresponding to r intersects the blue trapezoid in space represented at τ . If it does, then report this intersection.
2. If C_τ is too large to warrant the construction of the envelope structure, then recursively query both children of τ .
3. Here, τ is not a leaf and C_τ is small enough so that τ stores the envelope structure of the corresponding set of lines. In this case test whether the edge corresponding to r lies above all edges corresponding to blue line segments in C_τ , or below all such edges. If it lies above or below all the blue edges, then stop. Otherwise, recursively query both children of τ .

For the analysis of the search algorithm let k_r be the number of intersections between the edge corresponding to r and the facets of Σ_1 . The basic property of hereditary segment trees discussed in Section 2 is easily seen to imply that each such intersection is detected exactly once. The question that we have to answer is how much time each intersection costs and how much overhead we need for the line segment itself.

As noted earlier, r belongs to the standard lists of $O(N^{\epsilon/2})$ nodes μ of the segment tree and to the hereditary lists of only a constant number of nodes μ , where the constant depends on ϵ . For each occurrence we query the tree \mathcal{B}_μ . To get started we run down the tree (Case 2) until we hit the fringe nodes equipped with envelope structures; there are $O(N^{(1+\epsilon)/2})$ of them. Thus, the overhead for each line segment is $O(N^{1/2+\epsilon})$ and the rest of the time will be charged to the k_r intersections caused by r .

If it turns out that the edge corresponding to r neither lies above all edges corresponding to line segments in C_τ nor below all of them (Case 3), then by continuity of Σ_1 the edge corresponding to r intersects Σ_1 at least once somewhere between the topmost and bottommost line segments in C_τ . Now we need to distinguish two cases depending on whether the trapezoid that intersects r at the point of discussion is in C_τ or it is not (which is possible because there are gaps between the trapezoids of C_τ due to edges that are either too long or too short to be in C_τ). In both cases the intersection causes $O(\log N)$ queries to envelope structures in \mathcal{B}_μ and thus $O(\log^2 N)$ time. This is not all yet because the same intersection point can cause such queries also in other trees \mathcal{B}_μ . We now show that the number of such trees is fortunately only at most some constant, which will lead to the final result of this section.

The $O(N^{\epsilon/2})$ occurrences of r in standard lists amounts to a decomposition of r into the same number of pairwise disjoint subsegments, and each intersection point lies only on one of these subsegments and thus is charged only for the node μ that represents this subsegment of r . The occurrences of r in hereditary lists do not decompose r into disjoint segments, so an intersection point can cause queries for each such occurrence. However, as mentioned above there are only a constant number of occurrences in hereditary lists which shows that the time

spent per intersection point is $O(\log^2 N)$. This implies a total time bound of $O(N^{3/2+\varepsilon} + k \log^2 N)$, where k is the size of the output.

Finally, we note that instead of constructing the hereditary segment tree with envelope structures first and searching it later we can just traverse it (and its subtrees \mathcal{R}_μ and \mathcal{R}_ν) in preorder and construct and search only the structures currently needed (see also Section 3.1). This leaves us with the same time bound but improves the bound on the storage to $O(N)$. We thus summarize the results of this section.

THEOREM 6.1. *Given two polyhedral terrains in three-dimensional space with a total of n edges, we can construct their upper envelope in time $O(n^{3/2+\varepsilon} + k \log^2 n)$ and storage $O(n)$, where ε is an arbitrary but fixed positive constant and k is the number of edges of the upper envelope.*

Acknowledgments. We wish to thank P. K. Agarwal for his comments and helpful suggestions.

References

- [A1] P. K. Agarwal, Intersection and decomposition algorithms for arrangements of curves in the plane, Ph.D. Thesis, New York University, 1989.
- [A2] P. K. Agarwal, Private communication, 1991.
- [C1] B. Chazelle, Reporting and counting segment intersections, *J. Comput. System Sci.* 32(2) (1986), 156–182.
- [C2] B. Chazelle, An optimal convex hull algorithm and new results on cuttings, *Proc. 32nd Ann. IEEE Symp. on Foundations of Computer Science*, 1991, pp. 29–38.
- [C3] B. Chazelle, Cutting hyperplanes for divide-and-conquer, *Discrete Comput. Geom.* 9 (1993), 145–158.
- [CE] B. Chazelle and H. Edelsbrunner, An optimal algorithm for intersecting line segments in the plane, *J. Assoc. Comput. Mach.* 39 (1992), 1–54.
- [CEG⁺] B. Chazelle, H. Edelsbrunner, L. J. Guibas, M. Sharir, and J. Stolfi, Lines in space—combinatorics, algorithms, Tech. Report, Dept. Computer Science, Courant Institute, New York, 1990.
- [CG] B. Chazelle and L. J. Guibas, Fractional cascading: I. A data structuring technique, *Algorithmica* 1 (1986), 133–162.
- [CS] B. Chazelle and M. Sharir, An algorithm for generalized point location and its applications, *J. Symbolic Comput.* 10 (1990), 281–309.
- [CSSS] R. Cole, J. Salowe, W. Steiger, and E. Szemerédi, An optimal-time algorithm for slope selection, *SIAM J. Comput.* 18 (1989), 792–810.
- [EGS] H. Edelsbrunner, L. J. Guibas, and J. Stolfi, Optimal point location in a monotone subdivision, *SIAM J. Comput.* 15 (1986), 317–340.
- [EM] H. Edelsbrunner and E. P. Mücke, Simulation of simplicity: a technique to cope with degenerate cases in geometric algorithms, *ACM Trans. Graphics* 9(1) (1990), 66–104.
- [GOS] L. J. Guibas, M. H. Overmars and M. Sharir, Intersecting line segments, ray shooting, and other applications of geometric partitioning techniques, Tech. Report RUU-CS-88-26, Dept. Computer Science, University of Utrecht, 1988.
- [GS] L. J. Guibas and R. Seidel, Computing convolutions by reciprocal search, *Discrete Comput. Geom.* 2 (1987), 175–193.
- [M1] J. Matoušek, Approximations and optimal geometric divide-and-conquer, *Proc. 23rd Ann. ACM Symp. on Theory of Computing* 1991, pp. 505–511.

- [M2] J. Matoušek, Efficient partition trees, *Proc. 7th Ann. ACM Symp. on Computational Geometry*, 1991, pp. 1–9.
- [MaS] H. Mairson and J. Stolfi, Reporting and counting intersections between two sets of line segments, in *Theoretical Foundations of Computer Graphics and CAD* (R. A. Earnshaw, ed.), NATO ASI Series, Vol. F-40, Springer-Verlag, Berlin, 1988, pp. 307–325.
- [MeS] K. Mehlhorn and K. Simon, Intersecting two polyhedra one of which is convex, in *Proc. Fundamentals of Computer Theory*, Lecture Notes in Computer Science, Vol. 199, Springer-Verlag, Berlin, 1985, pp. 534–542.
- [PS] F. P. Preparata and M. I. Shamos, *Computational Geometry—An Introduction*, Springer-Verlag, New York, 1985.
- [S] M. Sharir, The shortest watchtower and related problems for polyhedral terrains, *Inform. Process. Lett.* **29** (1988), 265–270.
- [VW] V. K. Vaishnavi and D. Wood, Rectilinear line segment intersection, layered segment trees and dynamization, *J. Algorithms* **2** (1982), 160–176.

