



ELSEVIER

Computational Geometry 16 (2000) 103–127

Computational  
Geometry

Theory and Applications

www.elsevier.nl/locate/comgeo

## Cylindrical static and kinetic binary space partitions

Pankaj K. Agarwal<sup>a,\*</sup>, Leonidas J. Guibas<sup>b,2</sup>, T.M. Murali<sup>c,3</sup>, Jeffrey Scott Vitter<sup>a,4</sup>

<sup>a</sup> Center for Geometric Computing, Department of Computer Science, Duke University, Box 90129,  
Durham, NC 27708-0129, USA

<sup>b</sup> Computer Science Department, Gates 374, Stanford University, Stanford, CA 94305, USA

<sup>c</sup> Compaq Computer Corporation, Cambridge Research Lab, One Kendall Square, Bldg. 700, Cambridge, MA 02139, USA

Communicated by J.-R. Sack; submitted 3 April 1999; accepted 31 January 2000

---

### Abstract

We describe the first known algorithm for efficiently maintaining a Binary Space Partition (BSP) for  $n$  continuously moving segments in the plane, whose interiors remain disjoint throughout the motion. Under reasonable assumptions on the motion, we show that the total number of times this BSP changes is  $O(n^2)$ , and that we can update the BSP in  $O(\log n)$  expected time per change. Throughout the motion, the expected size of the BSP is  $O(n \log n)$ .

We also consider the problem of constructing a BSP for  $n$  static triangles with pairwise-disjoint interiors in  $\mathbb{R}^3$ . We present a randomized algorithm that constructs a BSP of size  $O(n^2)$  in  $O(n^2 \log^2 n)$  expected time. We also describe a deterministic algorithm that constructs a BSP of size  $O((n+k) \log^2 n)$  and height  $O(\log n)$  in  $O((n+k) \log^3 n)$  time, where  $k$  is the number of intersection points between the edges of the projections of the triangles onto the  $xy$ -plane. This is the first known algorithm that constructs a BSP of  $O(\log n)$  height for disjoint triangles in  $\mathbb{R}^3$ . © 2000 Elsevier Science B.V. All rights reserved.

**Keywords:** Binary space partitions; Hierarchical data structures; Kinetic data structures; Computational geometry; Computer graphics

---

\* Corresponding author.

*E-mail addresses:* pankaj@cs.duke.edu (P.K. Agarwal), guibas@cs.stanford.edu (L.J. Guibas), murali@crl.dec.com (T.M. Murali), jsv@cs.duke.edu (J.S. Vitter).

<sup>1</sup> Support was provided by National Science Foundation research grant CCR-93-01259, by Army Research Office MURI grant DAAH04-96-1-0013, by a Sloan fellowship, by a National Science Foundation NYI award and matching funds from Xerox Corp, and by a grant from the U.S.–Israeli Binational Science Foundation.

<sup>2</sup> Support was provided in part by National Science Foundation grant CCR-9623851 and by US Army MURI grant DAAH04-96-1-007.

<sup>3</sup> This research was done when this author was affiliated with Brown University and was a visiting scholar at Duke University. Support was provided in part by National Science Foundation research grant CCR-9522047 and by Army Research Office MURI grant DAAH04-96-1-0013.

<sup>4</sup> Support was provided in part by National Science Foundation research grant CCR-9522047, by Army Research Office grant DAAH04-93-G-0076, and by Army Research Office MURI grant DAAH04-96-1-0013.

## 1. Introduction

The Binary Space Partition (BSP, also known as BSP tree), originally proposed by Schumacker et al. [32] and further refined by Fuchs et al. [19], is a hierarchical partitioning of space widely used in several areas, including computer graphics (visibility determination [4,34], global illumination [9], shadow generation [12,13], and ray tracing [25]), solid modeling [26,28,35], geometric data repair [24], robotics [5], network design [22], and surface simplification [3]. Key to the BSP's success is that it serves both as a model for an object (or a set of objects) and as a data structure for querying the object.

Informally, a BSP  $\mathcal{B}$  for a set of objects is a binary tree, where each node  $v$  is associated with a convex region  $\Delta_v$ . The regions associated with the children of  $v$  are obtained by splitting  $\Delta_v$  with a hyperplane. If  $v$  is a leaf of  $\mathcal{B}$ , then the interior of  $\Delta_v$  does not intersect any object.<sup>5</sup> The regions associated with the leaves of the tree form a convex decomposition of space. The faces of the decomposition induced by the leaves intersect the objects and divide them into fragments; these fragments are stored at appropriate nodes of the BSP. The efficiency of BSP-based algorithms depends on the number of nodes in the tree and on the height of the tree. As a result, several algorithms for constructing BSPs of small size and/or small height have been proposed [4,10,19,29,30,34,35].

In this paper, we study *cylindrical* BSPs in which all the cuts that do not contain any input object are made by hyperplanes parallel to the same fixed direction. We address two problems. The first problem can be formulated as follows. Let  $S$  be a set of  $n$  interior-disjoint segments in the plane, each moving along a continuous path. We want to maintain the BSP for  $S$  as the segments in  $S$  move. We assume that the segments move in such a way that they never intersect, except possibly at their endpoints. Most of the work to date deals with constructing a BSP for a set of static segments that do not move. Paterson and Yao propose a randomized algorithm that constructs a BSP of  $O(n \log n)$  size in  $\Theta(n \log n)$  time for a set of  $n$  segments in the plane [29]. They also propose a deterministic algorithm, based on a divide-and-conquer approach, that constructs a BSP of size  $O(n \log n)$  in  $\Theta(n \log n)$  time [29]. Both of these algorithms are not “robust”, in the sense that a small motion of one of the segments may cause many changes in the tree, or may cause non-local changes. Therefore, they are ill-suited for maintaining a BSP for a set of moving segments.

There have been a few attempts to update BSPs when the objects defining them move. Naylor describes a method to implement dynamic changes in a BSP, where the static objects are represented by a balanced BSP (computed in a preprocessing stage), and then the moving objects are inserted at each time step into the static tree [27]. Using the same assumption that moving objects are known a priori, Torres proposes the augmentation of BSPs with additional separating planes, which may localize the updates needed for deletion and re-insertion of moving objects in a BSP [36]. This approach tries to exploit the spatial coherence of the dynamic changes in the tree by introducing additional cutting planes. Chrysanthou suggests a more general approach, which does not make any distinction between static and moving objects [14]. By keeping additional information about topological adjacencies in the tree, the algorithm performs insertions and deletions of a node in a more localized way. But all these prior efforts boil down to deleting moving objects from their earlier positions and re-inserting them in their current positions after some time interval has elapsed. Such approaches suffer from the fundamental problem that it is very difficult to know how to choose the correct time interval size: if the interval is too small, then the BSP does not in fact change combinatorially, and the deletion/re-insertion is just wasted computation; if it is

---

<sup>5</sup> We assume that the objects are  $(d - 1)$ -dimensional polytopes in  $\mathbb{R}^d$ .

too big, then important intermediate changes to the BSP can be missed, which may affect applications that use the tree.

Our algorithm, instead, treats the BSP as a *kinetic data structure*, a paradigm introduced by Basch et al. [6]; see also the survey by Guibas [20]. We view the equations of the cuts made at the nodes of the BSP and the edges and faces of the subdivision induced by the BSP as functions of time. The cuts and the edges and faces of the subdivision change continuously with time. However, “combinatorial” changes in the BSP and in the subdivision (we precisely define this notion later) occur only at certain times. We explicitly take advantage of the continuity of the motion of the objects involved so as to generate updates to the BSP only when actual events cause the BSP to change combinatorially.

In Section 3, we describe a randomized kinetic algorithm for maintaining a BSP for moving segments in the plane. We assume that the segment motions are *oblivious* to the random bits used by the algorithm; our algorithm chooses a random permutation of the segments at the beginning of time, and we assume that no agent determining the motion of the segments has access to any information about this random permutation. Following Basch et al. [6], we assume that each moving segment has a posted flight plan that gives full or partial information about the segment’s current motion. Whenever a flight plan changes (possibly due to an external agent), our algorithm is notified and it updates a global event queue to reflect the change. We first derive a randomized algorithm for computing a BSP for a set of static segments, which combines ideas from Paterson and Yao’s randomized and deterministic algorithms, but is also robust, in the sense described earlier. The “combinatorial structure” of the BSP constructed by this algorithm changes only when the  $x$ -coordinates of a pair of segment endpoints become equal. We prove that at any given instant, we need to consider only  $O(n)$  such endpoint pairs. Furthermore, the set of pairs we need to consider changes only when the combinatorial structure of the BSP changes. We show that under our assumption on the segment motions, the BSP can be updated in  $O(\log n)$  expected time at each event. We also show that if  $k$  of the segments of  $S$  move along “pseudo-algebraic” paths, and the remaining segments of  $S$  are stationary, then the expected number of changes in the BSP is  $O(kn \log n)$ . As far as we know, this is the first nontrivial algorithm for maintaining a BSP for moving segments in the plane.

Next, we study the problem of computing a BSP for a set  $S$  of  $n$  interior-disjoint triangles in  $\mathbb{R}^3$ . Paterson and Yao [29] describe a randomized incremental algorithm that constructs a BSP of size  $O(n^2)$  in expected time  $O(n^3)$ . They also show that their algorithm can be made deterministic without affecting its asymptotic running time. It has been an open problem whether a BSP for  $n$  triangles in  $\mathbb{R}^3$  can be constructed in near-quadratic time. Sub-quadratic bounds are known for special cases: Paterson and Yao’s algorithm for orthogonal rectangles [30], de Berg’s result for fat polyhedra [17], and the technique of Agarwal et al. [2] for fat orthogonal rectangles. However, none of these approaches leads to a near-quadratic-time algorithm for triangles in  $\mathbb{R}^3$ . The bottleneck in analyzing the expected running time of the Paterson–Yao algorithm is that no nontrivial bound is known on the number of vertices in the convex subdivision of  $\mathbb{R}^3$  induced by the BSP constructed by the algorithm. Known techniques for analyzing randomized algorithms, such as the Clarkson–Shor framework [16] or backwards analysis [33], cannot be used to obtain a near-quadratic bound on this quantity, since the BSP constructed by the algorithm is not canonical; it strongly depends on the order in which triangles are processed.

In Section 4, we present a randomized algorithm that constructs a BSP for  $S$  of size  $O(n^2)$  in  $O(n^2 \log^2 n)$  expected time.<sup>6</sup> Our algorithm is a variant of the randomized Paterson–Yao algorithm. We

---

<sup>6</sup> Our algorithm constructs a BSP of expected size  $O(n^2)$ . We can make the size bound deterministic by repeatedly running the algorithm until it constructs a BSP of size  $O(n^2)$ . This process affects only the constant factor in the running time.

construct the BSP for  $S$  in such a way that there is a close relationship between the BSP and the planar arrangement of the lines supporting the edges of the  $xy$ -projections of the triangles in  $S$ . We use results on random sampling [16] and on arrangements of lines [18] to bound the expected number of vertices in the convex subdivision of  $\mathbb{R}^3$  induced by the BSP and the expected running time of the algorithm.

Finally, we present a deterministic algorithm in Section 5 for constructing a BSP for a set  $S$  of  $n$  triangles in  $\mathbb{R}^3$ . If  $k$  is the number of intersection points of the  $xy$ -projections of the edges of triangles in  $S$ , then the algorithm constructs a BSP of size  $O((n+k)\log^2 n)$  in time  $O((n+k)\log^3 n)$ ; if  $k \ll n^2$ , the deterministic algorithm constructs a much smaller BSP than do Paterson and Yao's and our randomized algorithms. Another nice property of our deterministic algorithm is that the height of the BSP it constructs is  $O(\log n)$ , which is useful for ray-shooting queries, for example. It was an open problem whether BSPs of near-quadratic size and  $O(\log n)$  height could be constructed for  $n$  triangles in  $\mathbb{R}^3$ . The height of the BSP constructed by the randomized algorithms (both ours and Paterson and Yao's) can be  $\Omega(n)$ , e.g., when  $S$  is the set of faces of a convex polytope.

Before proceeding further, we give a formal definition of a BSP. A *binary space partition*  $\mathcal{B}$  for a set  $S$  of convex  $(d-1)$ -polytopes in  $\mathbb{R}^d$  with pairwise-disjoint interiors is a binary tree defined as follows: Each node  $v$  in  $\mathcal{B}$  is associated with a convex  $d$ -polytope  $\Delta_v$  and a set of  $(d-1)$ -polytopes  $S_v = \{s \cap \Delta_v \mid s \in S\}$ . The polytope associated with the root of  $\mathcal{B}$  is  $\mathbb{R}^d$  itself. If  $S_v$  is empty, then node  $v$  is a leaf of  $\mathcal{B}$ . Otherwise, we partition  $\Delta_v$  into two convex polytopes by a *cutting hyperplane*  $H_v$ . We refer to the polytope  $H_v \cap \Delta_v$  as the *cut* made at  $v$ . At  $v$ , we store the equation of  $H_v$  and the set  $\{s \mid s \subseteq H_v, s \in S_v\}$  of polytopes in  $S_v$  that lie in  $H_v$ . If we let  $H_v^+$  be the positive halfspace and  $H_v^-$  be the negative halfspace bounded by  $H_v$ , the polytopes associated with the left and right children of  $v$  are  $\Delta_v \cap H_v^-$  and  $\Delta_v \cap H_v^+$ , respectively. The left subtree of  $v$  is a BSP for  $S_v^- = \{s \cap H_v^- \mid s \in S_v\}$  and the right subtree of  $v$  is a BSP for  $S_v^+ = \{s \cap H_v^+ \mid s \in S_v\}$ . The size of  $\mathcal{B}$  is the sum of the number of nodes in  $\mathcal{B}$  and the total number of polytopes stored at all the nodes in  $\mathcal{B}$ .

For our purposes,  $S$  is either a set of  $n$  segments in the plane or a set of  $n$  triangles in  $\mathbb{R}^3$ . A unifying feature of all the BSPs constructed by our algorithms is that the region  $\Delta_v$  associated with each node  $v$  is a *cylindrical* cell in the sense that  $\Delta_v$  may contain top and bottom faces that are contained in objects belonging to  $S$ , but all other faces are vertical. In the plane,  $\Delta_v$  is a trapezoid; in  $\mathbb{R}^3$ ,  $\Delta_v$  may have large complexity, as it can contain many vertical faces.

## 2. Static algorithm for segments

Let  $S$  be a set of  $n$  interior-disjoint segments in the plane. In this section, we describe a randomized algorithm for computing a BSP  $\mathcal{B}$  for  $S$  when the segments in  $S$  are stationary. In the next section, we explain how to maintain  $\mathcal{B}$  as each segment in  $S$  moves along a continuous path.

Our algorithm makes two types of cuts: a *point* cut is a vertical cut through an endpoint of a segment and an *edge* cut is a cut along a segment. Edge cuts are always contained totally within input segments; therefore, they do not cross any other input segment. For each node  $v \in \mathcal{B}$ , the corresponding polygon  $\Delta_v$  is a trapezoid; the left and right boundaries of the trapezoid are bounded by point cuts, and the top and bottom boundaries are bounded by edge cuts.

We now describe our static algorithm. We start by choosing a random permutation  $\langle s_1, s_2, \dots, s_n \rangle$  of  $S$ . We say that  $s_i$  has a *higher priority* than  $s_j$  if  $i < j$ . We add the segments in decreasing order of priority and maintain a BSP for the segments added so far. Let  $S^i = \{s_1, s_2, \dots, s_i\}$  be the set of the first

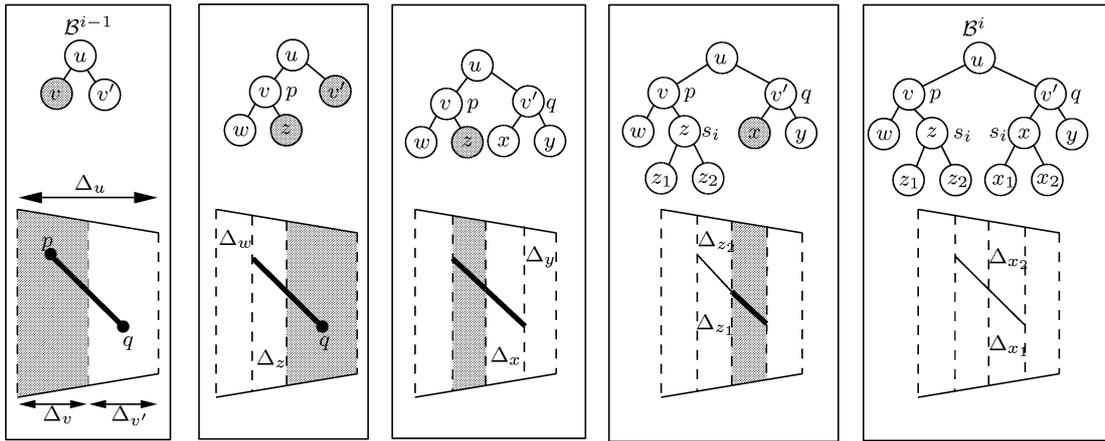


Fig. 1. The BSP  $\mathcal{B}^{i-1}$ , the sequence of cuts made in the  $i$ th stage, and the BSP  $\mathcal{B}^i$ . At each step, the shaded trapezoid is split. Portions of  $s_i$  that lie in the interior of a trapezoid that corresponds to a current leaf node are drawn using thick lines. The label next to a node signifies the cut made at that node.

$i$  segments in the permutation. Our algorithm works in  $n$  stages. At the beginning of the  $i$ th stage, where  $i > 0$ , we have a BSP  $\mathcal{B}^{i-1}$  for  $S^{i-1}$ ;  $\mathcal{B}^0$  consists of a single node  $v$ , where  $\Delta_v$  is the entire plane. In the  $i$ th stage, we add  $s_i$  and compute a BSP  $\mathcal{B}^i$  for  $S^i$  as follows:

1. Suppose  $p$  and  $q$  are the left and right endpoints of  $s_i$ , respectively. Let  $v$  be the leaf of  $\mathcal{B}^{i-1}$  such that  $\Delta_v$  contains  $p$ . We partition  $\Delta_v$  into two trapezoids  $\Delta_v^-$  and  $\Delta_v^+$  using a point cut defined by  $p$ , where  $\Delta_v^-$  lies to the left of the cut. We create two children  $w$  and  $z$  of  $v$ , with  $w$  being the left child of  $v$ . We set  $\Delta_w = \Delta_v^-$  and  $\Delta_z = \Delta_v^+$  and store  $p$  at  $v$ . We then perform a similar step for  $q$ .
2. For each trapezoid  $\Delta_x$  that intersects  $s_i$ , we store  $s_i$  at  $x$ , and split  $\Delta_x$  into two trapezoids by making an edge cut along  $s_i$ . We again create two children  $w$  and  $z$  of  $x$ , with  $w$  being the left child. We set  $\Delta_w$  to be the sub-trapezoid of  $\Delta_x$  lying below the cut and  $\Delta_z$  to be the sub-trapezoid of  $\Delta_x$  lying above the cut.

The resulting tree is the BSP  $\mathcal{B}^i$  for  $S^i$ . See Fig. 1 for an example of constructing  $\mathcal{B}^i$  from  $\mathcal{B}^{i-1}$ .

This completes the description of our algorithm. Note that once we fix the permutation, the algorithm is deterministic and constructs a unique BSP. In order to analyze the algorithm, we need a few definitions. We refer to the vertical segment drawn upwards (respectively, downwards) from an endpoint  $p$  as the *upper* (respectively, *lower*) *thread* of  $p$ . We call the segment containing the other endpoint of a thread the *stopper* of that thread. Note that the priority of the stopper of a thread of  $p$  is higher than that of the segment containing  $p$ . We can prove the following lemma about each thread.

**Lemma 2.1.** *Let  $p$  be an endpoint of a segment  $s \in S$ . The expected number of segments crossed by each of  $p$ 's threads is  $O(\log n)$ .*

**Proof.** Let  $\sigma_1, \sigma_2, \dots$  be the sequence of segments in  $S$  that intersect the top thread  $\rho$  of  $p$ , sorted in increasing order of the  $y$ -coordinates of their intersection with  $\rho$ ; clearly, there are at most  $n$  segments in this sequence. The segment  $\sigma_i$  is crossed by  $\rho$  if and only if  $s$  is inserted before any of the segments  $\sigma_1, \sigma_2, \dots, \sigma_{i-1}, \sigma_i$ . Since  $\mathcal{B}$  is constructed by inserting the segments of  $S$  in random order, the probability that  $\rho$  crosses  $\sigma_i$  is  $1/(i + 1)$ . Therefore the expected number of segments crossing  $\rho$  is at most

$\sum_{i=1}^n 1/(i+1) = O(\log n)$ . We can similarly show that the expected number of segments crossing  $p$ 's lower thread is  $O(\log n)$ .  $\square$

We can use this lemma to bound the size and height of  $\mathcal{B}$ .

**Theorem 2.2.** *The expected size of the BSP constructed by the above algorithm is  $O(n \log n)$  and the height of the BSP is  $O(\log n)$ , where the second bound holds with high probability.*

**Proof.** In order to bound the size of  $\mathcal{B}$ , the BSP constructed by the algorithm, it is enough to count the total number of cuts made in  $\mathcal{B}$ , since a cut is made at each interior node of  $\mathcal{B}$ . Clearly, there are at most  $2n$  point cuts made in  $\mathcal{B}$ . If an edge cut  $e$  is made at a node  $v$ , we charge  $e$  to the right endpoint of  $e$ . Suppose  $s$  is the segment in  $S$  containing  $e$ . The right endpoint of  $e$  is either the right endpoint of  $s$  or the intersection point of  $s$  with a thread of a segment whose priority is higher than  $s$ . In this way, we charge each endpoint and the intersection point of a segment and a thread at most once. As a result, Lemma 2.1 implies that the expected total number of edge cuts is  $O(n \log n)$ , which proves that the expected size of  $\mathcal{B}$  is  $O(n \log n)$ .

To bound  $\mathcal{B}$ 's height, we first bound the *depth* of an arbitrary point  $p$  in the plane, i.e., the number of nodes in the path from the root of  $\mathcal{B}$  to the leaf  $v \in \mathcal{B}$  such that  $\Delta_v$  contains  $p$ . We bound the number of nodes on this path that are split by edge cuts and point cuts separately.

Let  $\sigma_1, \sigma_2, \dots$  be the ordered sequence of segments in  $S$  intersected by a vertical ray starting at  $p$  and pointing in the  $(+y)$ -direction. An ancestor of  $v$  is split by an edge cut through  $\sigma_i$  if and only if  $\sigma_i$  has higher priority than  $\sigma_1, \sigma_2, \dots, \sigma_{i-1}$ . This event happens with probability  $1/i$ . Hence, the expected value of  $X$ , the number of ancestors of  $v$  that are split by edge cuts, is  $H_n = O(\log n)$ . We can actually prove that this bound on  $X$  holds with high probability. Since  $X$  is the sum of independent 0–1 random variables, using Chernoff's bound [23, p. 68], we have that for any constant  $\alpha \geq 1$ ,

$$\Pr[X > \alpha H_n] \leq \left( \frac{e^{\alpha-1}}{\alpha^\alpha} \right)^{H_n} = O(n^{-\alpha \ln \alpha + \alpha - 1}).$$

In particular, for any constant  $c$  we can choose  $\alpha$  so that  $\Pr[X > \alpha H_n] < 1/n^c$ , which shows that the value of  $X$  is  $O(\log n)$  with high probability.

We now consider the ancestors of  $v$  that are split by point cuts. Let  $\pi_1, \pi_2, \dots$  be the left segment endpoints that lie to the left of  $p$ . An ancestor of  $v$  is split by a point cut through  $\pi_i$  only if the segment with  $\pi_i$  as endpoint has higher priority than the segments with  $\pi_1, \pi_2, \dots, \pi_{i-1}$  as endpoints. A similar analysis to the one above proves that the number of ancestors of  $v$  that are split by points cuts is  $O(\log n)$  with high probability. Thus, the depth of any point  $p$  in the plane is  $O(\log n)$  with high probability.

The segments in  $S$  and the vertical lines passing through every segment endpoint decompose the plane into  $O(n^2)$  trapezoids. Any two points in one of these trapezoids will be contained in the same leaf of any BSP that our algorithm constructs, independent of the permutation we choose at the beginning of the algorithm. Hence, the height of BSP is the maximum depth of  $O(n^2)$  points, one in each such trapezoid. Since the depth of each point is  $O(\log n)$  with probability  $1 - 1/n^c$ , the height of  $\mathcal{B}$  is also  $O(\log n)$  with probability  $1 - 1/n^{c-2}$ , if we choose  $c \geq 3$ . This argument completes the proof of the lemma.  $\square$

### 3. Kinetic algorithm for segments

We now describe how to maintain the static BSP as the segments in  $S$  move continuously, under the assumption that their interiors remain pairwise disjoint throughout the motion. We parameterize the motion of the segments by time and use  $t$  to denote time. For a given time instant  $t$ , we will use  $t^-$  and  $t^+$  to denote the time instants  $t - \varepsilon$  and  $t + \varepsilon$ , respectively, where  $\varepsilon > 0$  is a sufficiently small constant.

Let  $s_i \in S$  be a segment with endpoints  $p$  and  $q$ . We assume that the position of  $p$  at time  $t$  is  $p(t) = (x_p(t), y_p(t))$ , where  $x_p(t)$  and  $y_p(t)$  are continuous functions of time;  $q(t)$  is specified similarly. The position of  $s_i$  at time  $t$  is  $s_i(t) = (p(t), q(t))$ ; if  $s_i$  is moving rigidly, then the equations for its endpoints are not independent. Our algorithm and the analysis work even if the endpoints of  $s_i$  move independently. Let  $S(t)$  denote the set  $S$  at time  $t$ . We assume that we choose a random permutation  $\pi$  of  $S$  once in the very beginning (at  $t = 0$ ), and that  $\pi$  does not change with time. Let  $\mathcal{B}(t)$  denote the BSP constructed by the static algorithm when applied on  $S(t)$ , using  $\pi$  as the permutation to decide the priority of the segments. We describe an algorithm that updates the BSP under the following assumption.

(★) *There is no correlation between the motion of the segments in  $S$  and their priorities. Therefore, the chosen permutation  $\pi$  always behaves like a random permutation, and Lemma 2.1 and Theorem 2.2 hold at all times.*

We first give an important definition. The *combinatorial structure* of  $\mathcal{B}$  is a binary tree, each of whose internal nodes  $v$  is associated with the set of segments  $S_v$  and with the segment endpoint (respectively, segment) defining the point (respectively, edge) cut made at  $v$ . We will use the combinatorial structure of the BSP crucially in our algorithm.

#### 3.1. Critical events

As the segments in  $S$  move continuously, the equations of the cuts associated with the nodes of  $\mathcal{B}$  also change. At the same time, the edges and vertices of the trapezoids in the subdivision of the plane induced by  $\mathcal{B}$  also move. However, the combinatorial structure of  $\mathcal{B}$  changes only when the set  $S_v$  changes for some node  $v \in \mathcal{B}$  or when the segment endpoint or segment defining the cut made at  $v$  changes. Since the segments in  $S$  are interior-disjoint and they move continuously, the set  $S_v$  changes only when the endpoint of a segment in  $S_v$  lies on the left or right edge of  $\Delta_v$ . See Fig. 2 for an example of such an event. If  $S_v$  does not change, then the cut made at  $v$  changes only if the segment defining the cut becomes vertical. We formalize this idea in the following lemma, which is not difficult to prove.

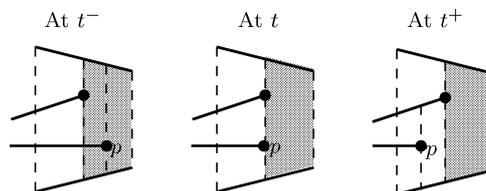


Fig. 2. Endpoint  $p$  lies on the left edge of  $\Delta_v$  (the shaded trapezoid) at  $t$ . The set  $S_v$  changes at time instant  $t$ .

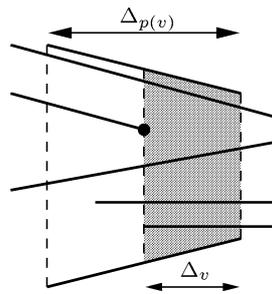


Fig. 3. The shaded trapezoid  $\Delta_v$  is transient.

**Lemma 3.1.** For any time instant  $t$ ,  $\mathcal{B}(t^-)$  and  $\mathcal{B}(t^+)$  have different combinatorial structures if and only if there exists a  $j > 0$  such that either  $s_j$  rotates through a vertical line at time  $t$  or there is a leaf  $v \in \mathcal{B}^{j-1}(t^-)$  such that an endpoint of  $s_j$  lies on the left or right edge of  $\Delta_v$  at time  $t$ .

This lemma implies that the combinatorial structure of  $\mathcal{B}(t)$  changes if and only if for a node  $v \in \mathcal{B}(t)$ ,  $\Delta_v$  shrinks to a vertical segment; we refer to these instants of time as *critical events*. This observation motivates us to call a node  $v$  in  $\mathcal{B}(t)$  *transient* if  $\Delta_v$  does not contain any endpoint in its interior and a point cut is made at the parent  $p(v)$  of  $v$ ; we call  $\Delta_v$  a *transient trapezoid*. See Fig. 3. Note that only edge cuts are made at  $v$  and its descendants. Thus, transient trapezoids are maximal among all those not containing point cuts. The following corollary to Lemma 3.1 is immediate.

**Lemma 3.2.** For any time instant  $t$ ,  $\mathcal{B}(t^-)$  and  $\mathcal{B}(t^+)$  have different combinatorial structures if and only if there exists a transient node  $v$  in  $\mathcal{B}(t^-)$  so that  $\Delta_v$  becomes a vertical segment at time  $t$ .

Transient nodes have some useful properties that are summarized in the following lemma.

**Lemma 3.3.** At any instant  $t$ , all transient nodes in  $\mathcal{B}(t)$  have the following properties. Let  $v$  be a transient node in  $\mathcal{B}(t)$ .

- (i) No proper ancestor of  $v$  is transient.
- (ii) Only edge cuts are made at the descendants of  $v$  (including  $v$  itself). The left (respectively, right) edge of the trapezoid associated with each descendant of  $v$  is a portion of the left (respectively, right) edge of  $\Delta_v$ .
- (iii) The expected number of descendants of  $v$  is  $O(\log n)$ .
- (iv) The number of transient nodes in  $\mathcal{B}(t)$  is at most  $4n$ .

**Proof.** Let  $q$  be the endpoint of a segment in  $S$  through which the point cut at  $p(v)$  is made.

- (i) No proper ancestor  $w$  of  $v$  can be transient since  $\Delta_w$  contains  $q$ .
- (ii) Since  $\Delta_v$  does not contain any endpoints, only edge cuts are made at all the descendants of  $v$ . Each segment that intersects  $\Delta_v$  crosses the left and right boundaries of  $v$ . Hence, the left (respectively, right) edge of the trapezoids associated with each descendant of  $v$  is a portion of the left (respectively, right) edge of  $\Delta_v$ .
- (iii) Each segment that induces an edge cut made at a descendant of  $v$  intersects one of  $q$ 's threads. Hence, by Lemma 2.1, the expected number of descendants of  $v$  is  $O(\log n)$ .

(iv) A point cut is made at the parent of each transient node. There are  $2n$  nodes in  $\mathcal{B}$  that are split by point cuts; each such node has two children.  $\square$

Intuitively, transient nodes are the highest nodes in  $\mathcal{B}(t)$  that can shrink to a vertical segment, thus causing a change in the combinatorial structure of  $\mathcal{B}(t)$ . If a trapezoid contains an endpoint in its interior, it cannot be the next trapezoid to shrink to a segment; and if an edge cut is made at the parent  $p(v)$  of a node  $v$  and  $\Delta_{p(v)}$  does not contain an endpoint, then  $\Delta_{p(v)}$  also shrinks to a segment whenever  $\Delta_v$  shrinks to a segment. Hence, it suffices to keep track of transient nodes to determine all the instants when the combinatorial structure of  $\mathcal{B}(t)$  changes. In the rest of the section, we present our kinetic algorithm motivated by this observation.

### 3.2. Updating the BSP

For a node  $v$  in  $\mathcal{B}$ , let  $\lambda_v$  (respectively,  $\rho_v$ ) denote the endpoint of a segment in  $S$  that induces the point cut containing the left (respectively, right) edge of  $\Delta_v$ . To detect critical events, we maintain the set

$$\Gamma(t) = \{(\lambda_v, \rho_v) \mid v \text{ is a transient node at time } t\}$$

of endpoint pairs inducing the point cuts that bound the left and right edges of each transient node; Lemma 3.3 implies that  $|\Gamma(t)| = O(n)$ . The elements of  $\Gamma(t)$  are certificates that prove that the combinatorial structure of  $\mathcal{B}(t)$  is valid. For each pair  $(\lambda_v, \rho_v)$  in  $\Gamma(t)$ , we use the known flight paths of  $\lambda_v$  and  $\rho_v$  to compute the time at which the  $x$ -coordinates of  $\lambda_v$  and  $\rho_v$  coincide; we store these time values in a global priority queue. In order to expedite the updating of  $\mathcal{B}$  at each critical event, we also store some additional information with the nodes in  $\mathcal{B}$  and the segments in  $S$ :

1. At each node  $v$  of  $\mathcal{B}$ , we store the number  $c_v$  of segment endpoints lying in the interior of  $\Delta_v$  ( $c_v$  helps us to determine the new transient trapezoids at an event).
2. For each endpoint  $p$  of a segment in  $S$ , we maintain the list  $T_p$  (respectively,  $B_p$ ) of segments that the upper (respectively, lower) thread of  $p$  crosses, sorted in the  $(+y)$ -direction (respectively,  $(-y)$ -direction). As the segments move, we will use these lists to update the stoppers of the threads issuing from the segment endpoints.

We first construct  $\mathcal{B}(0)$  using the static algorithm presented in Section 2. Next, we compute the set  $\Gamma(0)$  and insert the corresponding critical events in the priority queue. Then we repeatedly remove the next event from the priority queue and update  $\mathcal{B}$ ,  $\Gamma$ , and the priority queue as required. In the rest of the section, we will prove that if the combinatorial structure of  $\mathcal{B}$  changes at time  $t$ , then we can obtain  $\mathcal{B}(t^+)$  from  $\mathcal{B}(t^-)$  in  $O(\log n)$  expected time. We will also show that at each event, the expected time to update the global event queue is  $O(\log n)$ .

We now describe the procedure for updating the tree at each critical event. Recall that at each such instant  $t$ , there is a segment  $s_j \in S$  such that (i) either  $s_j$  becomes vertical or (ii) there is a leaf  $w \in \mathcal{B}^{j-1}(t^-)$  such that an endpoint  $p$  of  $s_j$  lies on the left or right edge of  $\Delta_w$ . We consider each case separately. Let  $\mathcal{B}^- = \mathcal{B}(t^-)$  and  $\mathcal{B}^+ = \mathcal{B}(t^+)$ . For a node  $z \in \mathcal{B}^-$ , let  $\mathcal{B}_z^-$  denote the subtree of  $\mathcal{B}^-$  rooted at  $z$ ; define  $\mathcal{B}_z^+$  similarly.

*Case (i).* The segment  $s_j$  is vertical. In this case,  $v$  is a transient node in  $\mathcal{B}^-$  with the property that  $\lambda_v$  and  $\rho_v$  are both endpoints of  $s_j$ . See Fig. 4. Let  $\rho_v = p$  and  $\lambda_v = q$ . Let  $u$  be  $v$ 's grandparent in  $\mathcal{B}^-$ ; the trapezoid  $\Delta_u$  contains  $s_j$ . Since  $q$  is to the left of  $p$  at time  $t^-$ ,  $v$  is the left child of the right child of  $u$

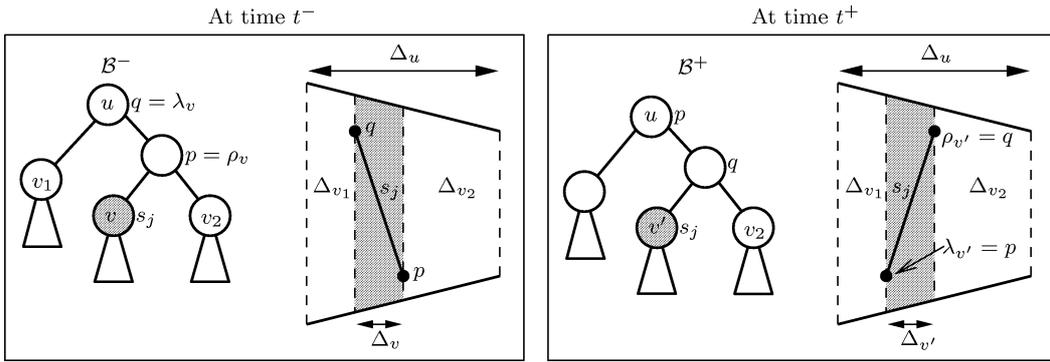


Fig. 4. The case when  $\lambda_v$  and  $\rho_v$  belong to the same segment.

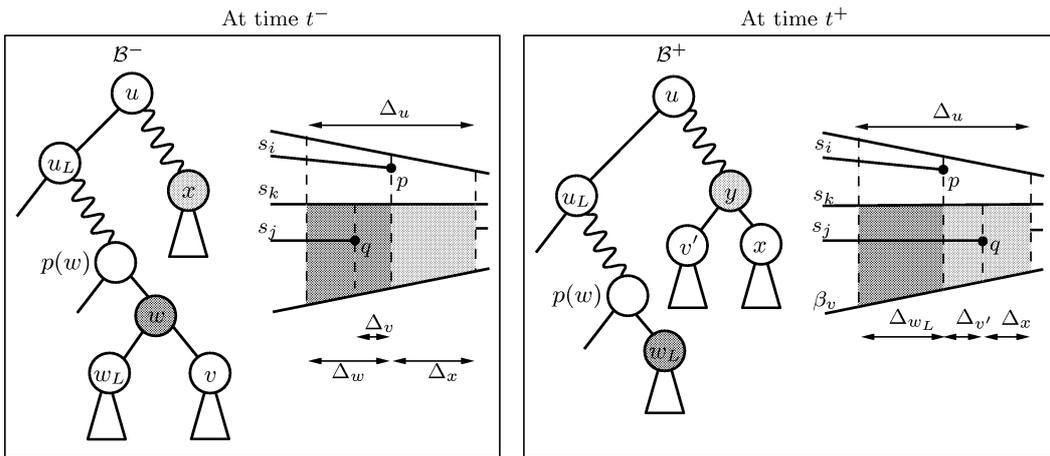


Fig. 5. The case when  $\lambda_v$  and  $\rho_v$  are endpoints of different segments. Arrows mark the horizontal extents of the trapezoids.

in  $\mathcal{B}^-$ . At time  $t^+$ ,  $p$  is to the left of  $q$ . The static algorithm when applied to the segments in  $S$  at time  $t^+$  will make a point cut through  $p$  in  $\Delta_u$ . Thus, we obtain  $\mathcal{B}(t^+)$  by storing  $p$  with  $u$  and  $q$  with the right child of  $u$ .

*Case (ii).* The points  $\lambda_v$  and  $\rho_v$  are endpoints of different segments. Assume that  $\rho_v = p$  (respectively,  $\lambda_v = q$ ) is the right endpoint of the segment  $s_i$  (respectively,  $s_j$ ), that  $s_i$  lies above  $s_j$ , and that the priority of  $s_i$  is higher than that of  $s_j$  (i.e.,  $i < j$ ). The  $x$ -coordinate of  $q$  is less than the  $x$ -coordinate of  $p$  at  $t^-$ . See Fig. 5. We now describe how we update  $\mathcal{B}(t)$  for this case. We will show later how to relax these assumptions. Let  $u$  and  $w$  be the leaves of  $\mathcal{B}^{i-1}(t^-)$  and  $\mathcal{B}^{j-1}(t^-)$ , respectively, at which the point cuts through  $p$  and  $q$ , respectively, are made. At time  $t^-$ , a point cut made through  $q$  divides  $\Delta_w$  into two trapezoids. One of these trapezoids is  $\Delta_v$ , which is transient at time  $t^-$ . By our assumptions about the event,  $v$  is the right child of  $w$ , and  $w$  lies in the left subtree of  $u$ . Let  $u_L$  be the left child of  $u$ , and let  $w_L$  be the left child of  $w$ . Let  $x$  be the leaf of  $\mathcal{B}^{j-1}(t^+)$  that contains  $q$  at time  $t^+$ . Since the combinatorial structures of  $\mathcal{B}^{j-1}(t^-)$  and  $\mathcal{B}^{j-1}(t^+)$  are identical,  $x$  is a leaf of  $\mathcal{B}^{j-1}(t^-)$  too and lies in the right subtree

of  $u$  in  $\mathcal{B}^{j-1}(t^-)$ . Let  $s_k \in S$  be the segment containing the top edge of  $\Delta_x$  in  $\mathcal{B}^{j-1}(t^-)$ . At time  $t^+$ , as  $q$  leaves the trapezoid  $\Delta_w$  and enters  $\Delta_x$ ,  $\Delta_{w_L}$  expands to  $\Delta_w$ ,  $\Delta_v$  disappears, and  $\Delta_x$  is split by a point cut through  $q$  into two trapezoids: a new trapezoid  $\Delta_{v'}$  and the portion of  $\Delta_x$  lying to the right of the cut through  $q$ . At time  $t^-$ ,  $\Delta_w$  is split by a point cut through  $q$  and  $\Delta_{w_L}$  is split by an edge cut along  $s_j$ , while at time  $t^+$ ,  $\Delta_w$  is split by an edge cut along  $s_j$ . Therefore  $\mathcal{B}_w^+$  is the same as  $\mathcal{B}_{w_L}^-$ . To obtain  $\mathcal{B}^+$ , we execute the following steps:

1. We search in the right subtree of  $u$  to locate the leaf  $x$  of  $\mathcal{B}^{j-1}(t^-)$  such that  $\Delta_x$  contains  $q$  at time  $t^+$ .
2. We delete the node  $w$  from  $\mathcal{B}^-$ , and if  $w$  was a left (respectively, right) child of its parent  $p(w)$ , we make  $w_L$  the new left (respectively, right) child of  $p(w)$ .
3. We construct the subtree  $\mathcal{B}_v^+$  by determining the set  $C$  of segments that intersect  $\Delta_{v'}$  (at time  $t^+$ ) and by making edge cuts through the segments in  $C$  in decreasing order of priority. There are two cases to consider:
  - (a) The segment  $s_k$  contains the top edge of  $\Delta_v$ . See Fig. 6. The set  $C$  consists of  $s_j$  and the set of segments intersecting  $\Delta_v$  (at time  $t^-$ ). We find these segments by traversing all the nodes of  $\mathcal{B}_v^-$ .
  - (b) The segment  $s_i$  contains the top edge of  $v$ . See Fig. 7. We set  $s_k$  to be the stopper of the upper thread of  $q$  at time  $t^+$ . As in the previous case, we include  $s_j$  and the segments inducing the edge cuts made in  $\mathcal{B}_v^-$  in  $C$ . In addition,  $C$  contains all segments that appear before  $s_k$  in the upper

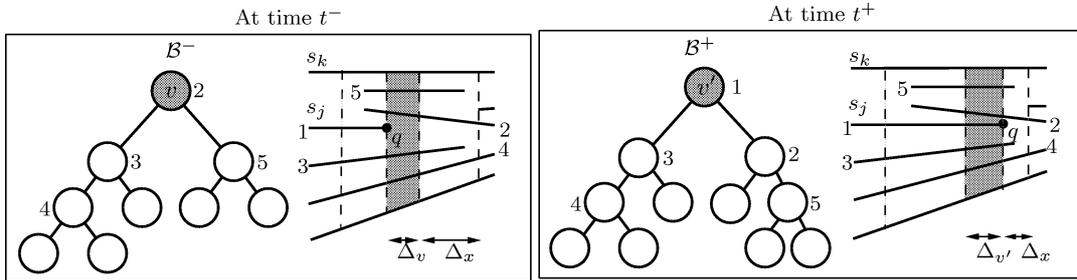


Fig. 6. The edge cuts made in  $\mathcal{B}_v^-$  and  $\mathcal{B}_v^+$ . Segments crossing  $\Delta_v$  and  $\Delta_{v'}$  are labeled with their priorities. The label next to a node is the priority of the segment containing the edge cut made at that node.

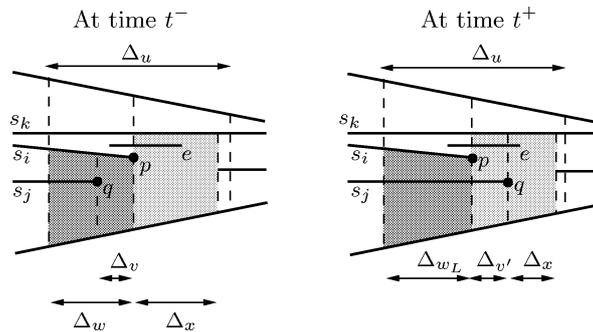


Fig. 7. The case when the top edges of  $\Delta_w$  and  $\Delta_x$  are contained in different segments. The segment  $e$  is intersected by the top thread of  $q$  at  $t^+$  but not at  $t^-$ .

thread of  $p$ . We determine these segments by traversing  $T_p$ , the list of segments that cross the upper thread of  $p$ . Note that these segments also cross the upper thread of  $q$  at  $t^+$ .

Finally, we insert  $s_j$  into  $B_p$ , the list of segments in  $S$  crossed by the lower thread of  $p$ , and update  $T_q$ .

4. We attach  $B_v^+$  to a descendant of  $p(x)$ , the parent of  $x$  in  $B^-$ , as follows: We create a node  $y$  and associate the point cut through  $q$  with it. The left and right subtrees of  $y$  in  $B^+$  are  $B_v^+$  and  $B_x^-$ , respectively. If  $x$  is the left (respectively, right) child of  $p(x)$ , then we add  $y$  as the new left (respectively, right) child of  $p(x)$ .
5. We update the set  $\Gamma(t^+)$ . For a node  $z$ , the number  $c_z$  of endpoints lying in the interior of  $\Delta_z$  changes only if  $z$  lies along the paths in  $B^+$  from  $u$  to the nodes  $p(w)$  and  $y$ . For such a node  $z$ , if  $c_z = 0$  at  $t^+$  and if  $\Delta_{p(z)}$  is split by a point cut, we add  $(\lambda_z, \rho_z)$  to the list  $\Gamma(t^+)$ . On the other hand, if  $c_z \neq 0$  at  $t^+$  but  $z$  is transient at  $t^-$  ( $z$  must be an ancestor of  $x$  in  $B^-$ ), we delete  $(\lambda_v, \rho_z)$  from  $\Gamma(t^+)$ . We also update the priority queue to reflect the changes to  $\Gamma(t^+)$ .

*Other cases.* We now show how we relax the assumptions we made earlier about the relative positions of  $s_i$  and  $s_j$  and their priorities.

1. If  $q$  is the left endpoint of  $s_j$ , the update procedure is the same, except that at time  $t^+$  we do not make an edge cut through  $s_j$  in  $B_v^+$ . See Fig. 8(a).
2. If the  $x$ -coordinate of  $q$  is greater than the  $x$ -coordinate of  $p$  at time  $t^-$ , we adapt a similar procedure as the one described above. See Fig. 8(b). The node  $x$  is again the leaf of  $B^{j-1}(t^+)$  such that  $\Delta_x$  contains  $q$  at  $t^+$ . We can reconstruct  $B_v^+$  as before: if the same segment contains the top edge of  $\Delta_v$  and  $\Delta_x$ , the same set of segments (apart from  $s_j$ ) intersects  $\Delta_v$  and  $\Delta_{v'}$ ; otherwise, among the segments that intersect  $\Delta_v$ , only segments below  $s_i$  intersect  $\Delta_{v'}$ . In both cases,  $s_j$  intersects either  $\Delta_v$  or  $\Delta_{v'}$  depending on whether  $q$  is the left or right endpoint of  $s_j$ . In the second case, we also update  $T_q$  accordingly. The other changes to  $B$  are similar to the cases we have handled; the details are not difficult to work out.
3. If  $p$  is the left endpoint of  $s_i$ , we reflect  $S$  about the  $y$ -axis and reduce the problem to one of the earlier cases.
4. If  $s_i$  lies below  $s_j$ , we reflect  $S$  about the  $x$ -axis, reducing the problem to one of the earlier cases.
5. If the priority of  $s_i$  is less than the priority of  $s_j$ , we swap the roles of  $s_i$  and  $s_j$  and reduce the problem to one of the earlier cases.

This completes the description of our procedure for processing critical events. We now analyze the running time of the update procedure. Assumption  $(\star)$  implies that Lemma 2.1 and Theorem 2.2 hold at times  $t^-$ ,  $t$  and  $t^+$ . We spend  $O(\log n)$  time in Step 1, since we traverse a path in  $B$  to find the node  $x$ .

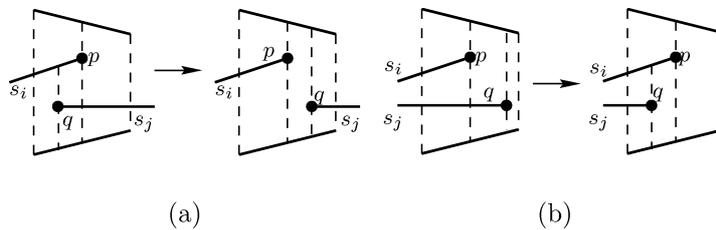


Fig. 8. Some other cases that arise when different segments interact in a critical event.

It is clear that Step 2 takes  $O(1)$  time. In Step 3, we find the segments crossing  $\Delta_{v'}$  and construct  $\mathcal{B}_{v'}^+$  in  $O(\log n)$  expected time, since the expected size of  $\mathcal{B}_{v'}^-$  is  $O(\log n)$  (by Lemma 3.3) and the expected number of segments in  $T_p$  is  $O(\log n)$  (by Lemma 2.1). It is clear that Step 4 takes  $O(1)$  time. Finally, in Step 5, we process  $O(\log n)$  nodes lying in two paths in the tree. By Lemma 3.3, each of the two paths contains at most one transient node. Hence, we insert or delete at most two events from the priority queue, which implies that Step 5 takes  $O(\log n)$  time. We thus obtain the main result of this section.

**Theorem 3.4.** *At each critical event, we can update  $\mathcal{B}(t)$  in  $O(\log n)$  expected time.*

Note that this theorem makes our BSP a kinetic data structure that is efficient, local, and compact, in the sense defined by Basch et al. [6]. However, our BSP is not responsive, since some events may take  $\Omega(n)$  time to process.

We say that the trajectories followed by a set of segments are *pseudo-algebraic* if the segments move so that each pair of endpoints exchanges  $y$ -order only  $O(1)$  times. A special case of pseudo-algebraic trajectories is when the trajectories of all the endpoints are constant-degree polynomials. If the trajectories of  $k$  of the segments in  $S$  are pseudo-algebraic and the remaining segments are stationary, then the total number of critical events is  $O(kn)$ . We spend  $O(\log n)$  expected time to maintain  $\mathcal{B}(t)$  at each event point. Hence, we obtain the following corollary to Theorem 3.4.

**Corollary 3.5.** *Let  $S$  be a set of  $n$  segments in the plane, and let  $G \subseteq S$  be a set of  $k$  segments. Suppose each segment of  $G$  moves along a pseudo-algebraic trajectory and the remaining segments of  $S$  are stationary, the total expected time spent in maintaining  $\mathcal{B}$  is  $O(kn \log n)$ .*

**Remark.** Our update algorithm works correctly even if two or more events occur at the same instant. If the events involve trapezoids in different parts of  $\mathcal{B}$ , it does not matter in which order we process the events. Otherwise, assume that at time  $t$  there are two transient trapezoids  $\Delta_v$  and  $\Delta_w$  with  $\lambda_v = p$ ,  $\rho_v = \lambda_w = q$  and  $\rho_w = r$ , such that the  $x$ -coordinates of  $p$ ,  $q$  and  $r$  are the same. The definition of transient trapezoids implies that only the events involving the pairs  $(p, q)$  and  $(q, r)$  are in the priority queue at time  $t^-$ . At time  $t$ , we process one of these events. Assume without loss of generality that this event involves  $(p, q)$ . When we process this event, we insert an event involving the pair  $(p, r)$  into the priority queue. Our algorithm updates  $\mathcal{B}$  correctly irrespective of whether the event involving  $(q, r)$  or the one involving  $(p, r)$  is processed next.

#### 4. BSPs for triangles: a randomized algorithm

In this section we describe a randomized algorithm for constructing a BSP  $\mathcal{B}$  of expected size  $O(n^2)$  for a set  $S$  of  $n$  triangles with pairwise-disjoint interiors in  $\mathbb{R}^3$ . The expected running time of the algorithm is  $O(n^2 \log^2 n)$ . We describe the algorithm in Section 4.1 and analyze its performance in Section 4.2.

##### 4.1. The randomized algorithm

We start with some definitions. For an object  $s$  in  $\mathbb{R}^3$ , let  $s^*$  denote the  $xy$ -projection of  $s$ . Let  $E$  be the set of edges of the triangles in  $S$ , and let  $E^*$  denote the set  $\{e^* \mid e \in E\}$ . Let  $\mathcal{L}$  be the set of lines in the

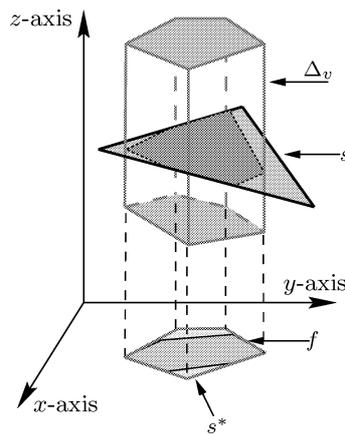


Fig. 9. An active face  $f$  and an active cell  $\Delta_v \in \mathbf{\Delta}(f)$  that is a vertical section of the cylinder erected on  $f$ . The boundary of triangle  $s$  intersects  $\Delta_v$  and the boundary of  $s^*$  intersects  $f$ .

$xy$ -plane supporting the edges in  $E^*$ . We choose a random permutation  $\langle \ell_1, \ell_2, \dots, \ell_{3n} \rangle$  of  $\mathcal{L}$ , and add the lines one by one in this order to compute  $\mathcal{B}$ . Let  $\mathcal{L}^i = \{\ell_1, \ell_2, \dots, \ell_i\}$ .

The algorithm works in  $3n$  stages. In the  $i$ th stage, we add  $\ell_i$  and construct a top subtree<sup>7</sup>  $\mathcal{B}^i$  of  $\mathcal{B}$  by refining the leaves of  $\mathcal{B}^{i-1}$ ;  $\mathcal{B}^0$  consists of one node (corresponding to  $\mathbb{R}^3$ ) and  $\mathcal{B}^{3n}$  is  $\mathcal{B}$ . As usual, we associate a convex polytope  $\Delta_v$  with each node  $v$  of  $\mathcal{B}^{i-1}$ . If  $v$  is a leaf of  $\mathcal{B}^{i-1}$  and no triangle in  $S$  intersects the interior of  $\Delta_v$ , i.e.,  $S_v = \emptyset$ , then  $v$  is a leaf of  $\mathcal{B}$  and we do not refine it further. Otherwise, we partition  $\Delta_v$  into two or more cells; these cells are leaves of  $\mathcal{B}^{i+1}$ .

Before describing the  $i$ th stage of the algorithm in detail, we explain the structure of  $\mathcal{B}^i$ . We need a few definitions first. At a node  $v$  of  $\mathcal{B}$ , the cutting plane  $H_v$  may support a triangle  $s \in S$  such that  $H_v \cap \Delta_v \subseteq s$ , i.e., the portion of  $H_v$  that lies in the interior of  $\Delta_v$  is contained in  $s$ . Such a cutting plane is referred to as a *free cut* and  $s$  is called a *free triangle*. We say that a leaf  $v$  of  $\mathcal{B}^i$  (or the cell  $\Delta_v$ ) is *active* if a triangle in  $S$  intersects the interior of  $\Delta_v$  (i.e.,  $S_v \neq \emptyset$ ); similarly, we say that a face  $f$  in the line arrangement  $\mathcal{A}(\mathcal{L}^i)$  is *active* if a segment in  $E^*$  intersects the interior of  $f$ . For each active leaf  $v$  in  $\mathcal{B}^i$ , the algorithm ensures that  $\Delta_v$  satisfies the following properties:

- (P1) If a triangle  $s \in S$  intersects the interior of  $\Delta_v$ , then the boundary of  $s$  also intersects the interior of  $\Delta_v$ .
- (P2) The cell  $\Delta_v$  is a vertical section of the cylinder  $f \times [-\infty, \infty]$  for exactly one active face  $f$  of  $\mathcal{A}(\mathcal{L}^i)$ ; the vertical section may be truncated by triangles of  $S$  at the top and bottom. See Fig. 9.

In order to execute each stage efficiently, we maintain the following additional information:

- (i) For each active cell  $\Delta \in \mathcal{B}^i$ , we store the subset  $S_\Delta \subseteq S$  of triangles that intersect the interior of  $\Delta$ .
- (ii) We maintain the arrangement  $\mathcal{A}(\mathcal{L}^i)$  as a planar graph [18]. For each active face  $f \in \mathcal{A}(\mathcal{L}^i)$ , we maintain the set  $\mathbf{\Delta}(f)$  of those active cells in  $\mathcal{B}^i$  that lie inside the cylinder  $f \times [-\infty, \infty]$ . Note that by Properties (P1) and (P2), a face  $f \in \mathcal{A}(\mathcal{L}^i)$  is active if and only if  $\mathbf{\Delta}(f) \neq \emptyset$ .

We now describe the  $i$ th stage in detail. Let  $H_i$  be the vertical plane supporting  $\ell_i$ . In the  $i$ th stage, we make a vertical cut inside each active cell that is intersected by  $H_i$ , followed by a number of free

<sup>7</sup> A top subtree of a tree is one that includes the root of the tree.

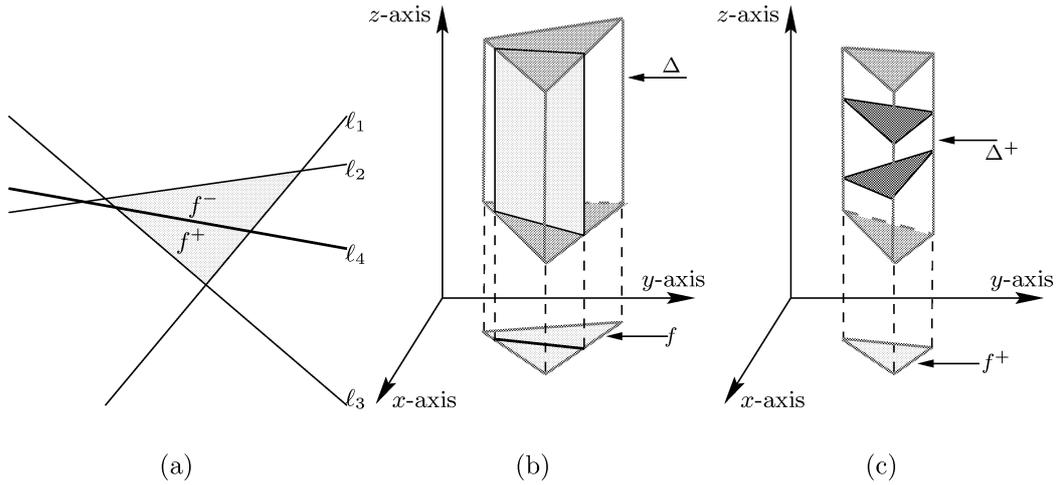


Fig. 10. (a) Tracing  $\ell_4$  (the thick line) through the faces of  $\mathcal{A}(\mathcal{L}^3)$ . The face  $f$  is shaded. (b) Splitting cell  $\Delta \in \mathbf{\Delta}(f)$  by  $H_4$ , the vertical plane containing  $\ell_4$ . (c) The free cuts in  $F_{\Delta^+}$  are ordered by  $z$ -coordinate.

cuts as follows. Let  $H_i^+$  (respectively,  $H_i^-$ ) be the positive (respectively, negative) halfspace supported by  $H_i$ .

1. We trace  $\ell_i$  through the faces of  $\mathcal{A}(\mathcal{L}^{i-1})$ . For each face  $f \in \mathcal{A}(\mathcal{L}^{i-1})$  intersected by  $\ell_i$ , we use  $\ell_i$  to split  $f$  into two faces  $f^+$  and  $f^-$ . See Fig. 10(a). Next, we partition each active cell  $\Delta \in \mathbf{\Delta}(f)$  into two cells  $\Delta^+ = \Delta \cap H_i^+$  and  $\Delta^- = \Delta \cap H_i^-$  (see Fig. 10(b)) and execute the following two steps on  $\Delta$ :
2. We compute the set  $S_{\Delta^+} \subseteq S_{\Delta}$  of triangles that intersect the interior of  $\Delta^+$ . We also compute the set  $F_{\Delta^+} \subseteq S_{\Delta^+}$  of triangles whose boundaries do not cross  $\Delta^+$ . Similarly, we compute the sets  $S_{\Delta^-}$  and  $F_{\Delta^-}$  for  $\Delta^-$ .
3. We split  $\Delta^+$  into a set  $\Psi$  of  $|F_{\Delta^+}| + 1$  cells by making free cuts along each of the triangles in  $F_{\Delta^+}$ . The cells in  $\Psi$  can be ordered by  $z$ -coordinate. Since the triangles in  $S$  are pairwise-disjoint, each triangle  $s \in S_{\Delta^+} \setminus F_{\Delta^+}$  intersects a unique cell  $\Delta' \in \Psi$ . We compute  $\Delta'$  by performing a binary search in  $\Psi$ , and add  $s$  to  $S_{\Delta'}$ . For each cell  $\Delta' \in \Psi$ , we add  $\Delta'$  to the set  $\mathbf{\Delta}(f^+)$  if  $S_{\Delta'} \neq \emptyset$ . Next, we repeat the same procedure for  $\Delta^-$ .

Whenever we split a three-dimensional cell into two, we add two children to the corresponding node in  $\mathcal{B}^{i-1}$  and store the necessary information with the newly created nodes. The resulting tree is  $\mathcal{B}^i$ . The cuts made in Step 3 ensure that  $\mathcal{B}^i$  satisfies property (P1).  $\mathcal{B}^i$  satisfies property (P2) since the cuts made in Step 1 are vertical. Note that a triangle  $s \in S$  does not intersect the interior of any cell after the three lines supporting the edges of  $s^*$  have been processed.

**Remark.** The free cuts made in Step 3 are crucial in keeping the size of the BSP quadratic. Instead, if we simply erect vertical planes as we do in Step 1 of the algorithm, and make cuts along a triangle  $s \in S$  only when all three lines supporting the  $xy$ -projections of the edges of  $s$  have been added, then there are instances of input triangles for which our algorithm will construct a BSP of  $\Omega(n^3)$  size regardless of the initial permutation of the triangles.

#### 4.2. Analysis of the algorithm

We first bound the expected size of  $\mathcal{B}$ . A similar proof is used by Paterson and Yao [29] to analyze their randomized algorithm for constructing BSPs for triangles in  $\mathbb{R}^3$ . The cuts made by the algorithm partition each triangle in  $S$  into a number of sub-polygons; each such sub-polygon is contained in the cutting plane of some node in  $\mathcal{B}$  and is stored at that node. Let  $\nu(S)$  be the total number of polygons stored at the nodes of  $\mathcal{B}$ . The following lemma bounds the size of  $\mathcal{B}$  in terms of  $\nu(S)$ .

**Lemma 4.1.** *The size of  $\mathcal{B}$  is at most  $11\nu(S)$ .*

**Proof.** Recall that the size of  $\mathcal{B}$  is defined to be the sum of the number of nodes in  $\mathcal{B}$  and the total number of triangles stored at all the nodes in  $\mathcal{B}$ . To bound the size of  $\mathcal{B}$ , we count the number of nodes in  $\mathcal{B}$  and then add  $\nu(S)$ . Let  $\nu(E)$  be the number of segments into which the edges of the triangles in  $S$  are partitioned by the cuts in  $\mathcal{B}$ .

We first bound the number of leaves in  $\mathcal{B}$  in terms of  $\nu(S)$  and  $\nu(E)$ . Let  $v$  be the parent of a leaf  $w$  in  $\mathcal{B}$ . Either a free cut or a vertical cut through an edge of a triangle  $s \in S$  is made at  $v$  (since these are the two types of cuts we make in  $\mathcal{B}$ ). We charge  $w$  to the cut made at  $v$ . There are three cases to consider:

1. If we made a face cut at  $v$ , we charge the cut at most twice.
2. If we made a vertical cut at  $v$  and if both children of  $v$  are leaves of  $\mathcal{B}$ , then  $s$  is vertical. In this case, we charge  $s$  twice.
3. If  $w$  is the only child of  $v$  that is a leaf and if a vertical cut is made at  $v$ , we now show that this cut is charged at most once. Suppose the vertical cut at  $v$  passes through an edge  $e$  of  $s$ . It is clear that we have not made a vertical cut passing through  $e$  at any ancestor of  $v$ . Consider the segment  $e' = e \cap \Delta_v$ . Let  $\Phi$  denote the set of segments that  $e'$  is partitioned into by the cuts made at the descendants of  $v$  in  $\mathcal{B}$ . Since we have made a vertical cut through  $e$ , we do not make vertical cuts in  $\mathcal{B}$  through any of the segments in  $\Phi$ . Therefore, we can charge  $w$  to any segment in  $\Phi$ . We charge each such segment at most once in this manner.

This argument implies that the number of leaves in  $\mathcal{B}$  is at most  $2\nu(S) + \nu(E)$ .

To bound this quantity, for each triangle  $s \in S$ , consider the arrangement  $\mathcal{A}_s$  on  $s$  formed by the intersection of  $s$  and the cuts in  $\mathcal{B}$ . Let  $e_s$  be the number of edges in  $\mathcal{A}_s$  that are portions of edges of  $s$  and let  $f_s$  be the number of faces in  $\mathcal{A}_s$ . Since at most three edges on the boundary of a face in  $\mathcal{A}_s$  are also portions of the edges of  $s$ , we have  $e_s \leq 3f_s$ . Summing over all triangles  $s \in S$ , we have  $\nu(E) \leq 3\nu(S)$ . Hence, the number of leaves in  $\mathcal{B}$  is at most  $5\nu(S)$ , which implies that the number of nodes in  $\mathcal{B}$  is at most  $10\nu(S)$ , thus proving the lemma.  $\square$

Thus, it suffices to bound the expectation  $E[\nu(S)]$  to bound the expected size of  $\mathcal{B}$ . To that end, we count the expected number of new sub-polygons created in the  $i$ th stage, and sum the result over all stages. We bound the number  $\nu_s^i$  of new sub-polygons into which a triangle  $s \in S$  is partitioned by the cuts made in the  $i$ th stage, and sum the resulting bound over all triangles in  $S$ . Note that the vertical cuts made in the  $i$ th stage are contained in the vertical plane  $H_i$  containing  $\ell_i$ .

Fix a triangle  $s \in S$ . For  $1 \leq k \leq i$ , let  $\lambda_k = H_k \cap s$  be the segment formed by the intersection of  $H_k$  and  $s$ , and let  $\Lambda_i$  be the set of resulting segments. Note that the endpoints of each  $\lambda_k$  lie on the boundary of  $s$ . To calculate  $E[\nu_s^i]$ , consider the segment arrangement  $\mathcal{A}(\Lambda_i)$  on  $s$ . We call a face of  $\mathcal{A}(\Lambda_i)$  a *boundary* face if it is adjacent to an edge of  $s$ ; otherwise, it is an *interior* face. See Fig. 11. Recall that for

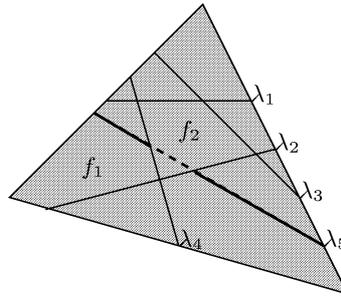


Fig. 11. The arrangement  $\mathcal{A}(\{\lambda_1, \lambda_2, \lambda_3, \lambda_4\})$  on triangle  $s$  (the shaded triangle). The face  $f_1$  is a boundary face and the face  $f_2$  is an interior face. The segment  $\lambda_5$  does not partition  $f_2$ .

a leaf  $v \in \mathcal{B}^{i-1}$ , we partition the cell  $\Delta_v$  only if  $\Delta_v$  is active. Property (P1) implies that the cuts made in the  $i$ th stage do not cross the interior of any interior face of  $\mathcal{A}(\Lambda_{i-1})$ , since such a face cannot intersect the interior of any active cell  $\Delta_v$ . Hence,  $v_s^i$  is the number of boundary faces of  $\mathcal{A}(\Lambda_{i-1})$  that are intersected by  $\lambda_i$ . (If property (P1) did not hold,  $v_s^i$  would be *all* the regions of  $\mathcal{A}(\Lambda_{i-1})$  that are intersected by  $\lambda_i$ .)

For  $1 \leq k \leq i$ , let  $\mu(\Lambda_i, k)$  denote the number of boundary faces in the arrangement  $\mathcal{A}(\Lambda_i \setminus \{\lambda_k\})$  that are intersected by  $\lambda_k$ . Observe that the sum  $\sum_{1 \leq k \leq i} \mu(\Lambda_i, k)$  equals the total number of edges bounding the boundary faces of  $\mathcal{A}(\Lambda_i)$ . By the zone theorem [11,18], the total number of edges of the boundary faces of  $\mathcal{A}(\Lambda_i)$  is  $O(i)$ . Hence,

$$\sum_{1 \leq k \leq i} \mu(\Lambda_i, k) = O(i).$$

Since the lines in  $\mathcal{L}^i$  are chosen randomly from the set  $\mathcal{L}$ ,  $\lambda_i$  can be any of the lines in  $\Lambda_i$  with equal probability. Therefore,

$$E[v_s^i] = \frac{1}{i} \sum_{1 \leq k \leq i} \mu(\Lambda_i, k) = O(1).$$

Hence, the total number of pieces created in the  $i$ th stage is  $O(n)$ . Summing over  $i$ , we find that the total number of sub-polygons into which the triangles in  $S$  are partitioned over the course of the entire algorithm is  $O(n^2)$ . The following lemma is immediate.

**Lemma 4.2.** *The expected size of the BSP constructed by the algorithm is  $O(n^2)$ .*

**Remark.** Since each cell  $\Delta_v \in \mathcal{B}$  is cylindrical, each vertex  $p$  of  $\Delta_v$  is contained in one of the triangles  $s$  that contains the non-vertical faces of  $\Delta_v$ . In fact,  $p$  is a vertex of the arrangement  $\mathcal{A}(\{\lambda_1, \lambda_2, \dots, \lambda_{3n}\})$  on  $s$  defined above. Thus, the preceding argument implies that the expected value of the total number of vertices of the nodes of  $\mathcal{B}$  is also  $O(n^2)$ . However, the height of  $\mathcal{B}$  can be  $\Omega(n)$ , e.g., if the triangles in  $S$  form a convex polytope.

Before analyzing the running time of the algorithm, we establish a relation between the projected edges intersecting an active face  $f \in \mathcal{A}(\mathcal{L}^{i-1})$  and the triangles intersecting the cells in  $\mathbf{\Delta}(f)$ . For such an active face  $f$ , let  $k_f$  be the number of projected edges in  $E^*$  that intersect the interior of  $f$ . By Property (P1), if a triangle  $s \in S$  intersects the interior of a cell  $\Delta \in \mathbf{\Delta}(f)$ , i.e.,  $s \in S_\Delta$ , then the boundary of  $s$  also

intersects the interior of  $\Delta$ . Therefore, an edge of  $s^*$  intersects the interior of  $f$ . Since  $s$  intersects the interior of only one cell in  $\mathbf{\Delta}(f)$ , we obtain

$$\sum_{\Delta \in \mathbf{\Delta}(f)} |S_{\Delta}| \leq k_f. \quad (4.1)$$

We now analyze the expected running time of the algorithm. We count the time spent during the  $i$ th stage in inserting the line  $\ell_i$  and then add this time over all stages of the algorithm. The zone theorem implies that in Step 1 of the algorithm, we spend  $O(i)$  time in tracing  $\ell_i$  through  $\mathcal{A}(\mathcal{L}^{i-1})$ . While processing an active face  $f$  of  $\mathcal{A}(\mathcal{L}^{i-1})$  that intersects  $\ell_i$ , for each cell  $\Delta \in \mathbf{\Delta}(f)$ , we spend  $O(1)$  time in Step 1 and  $O(|S_{\Delta}|)$  time in Step 2. In Step 3, for each triangle  $s \in S_{\Delta^+} \setminus F_{\Delta^+}$ , we spend  $O(\log |F_{\Delta^+}|)$  time in the binary search used to find the cell in the set  $\Psi$  that intersects  $s$ . Hence, the total time spent in Step 3 for the face  $f$  is  $O(|S_{\Delta}| \log |S_{\Delta}|)$ . Thus, (4.1) implies that the total time spent in processing  $f$  is

$$\sum_{\Delta \in \mathbf{\Delta}(f)} O(|S_{\Delta}| \log |S_{\Delta}|) = O(k_f \log k_f).$$

Let  $Z$  be the set of all active faces of  $\mathcal{A}(\mathcal{L}^{i-1})$  that are intersected by  $\ell_i$ . The total time spent in the  $i$ th stage is

$$\sum_{f \in Z} O(k_f \log k_f).$$

We now bound this sum. If we denote the number of vertices on the boundary of a face  $f$  by  $|f|$ , then by the zone theorem, we have  $\sum_{f \in Z} |f| = O(i)$ . Consider the vertical decomposition  $\mathcal{A}^{\parallel}(\mathcal{L}^{i-1})$  of  $\mathcal{A}(\mathcal{L}^{i-1})$ . Each face  $f \in \mathcal{A}(\mathcal{L}^{i-1})$  is decomposed into  $O(|f|)$  trapezoids in  $\mathcal{A}^{\parallel}(\mathcal{L}^{i-1})$ . By standard random-sampling arguments (see Clarkson [15]), the expected number of edges in  $E^*$  that intersect the interior of any such trapezoid is  $O((n \log i)/i)$ . This implies that for a face  $f \in \mathcal{A}(\mathcal{L}^{i-1})$ , the expected value of  $k_f$  is  $O(|f|(n \log i)/i)$ . Hence, the expected time spent in the  $i$ th stage is

$$\sum_{f \in Z} O(k_f \log k_f) = \sum_{f \in Z} O\left(|f| \left(\frac{n \log i}{i}\right) \log n\right) = O(n \log^2 n),$$

which implies the following theorem.

**Theorem 4.3.** *Let  $S$  be a set of  $n$  non-intersecting triangles in  $\mathbb{R}^3$ . We can compute a BSP for  $S$  of expected size  $O(n^2)$  in expected time  $O(n^2 \log^2 n)$ .*

## 5. BSPs for triangles: a deterministic algorithm

In this section, we describe a deterministic algorithm for computing a BSP for a set  $S$  of  $n$  triangles in  $\mathbb{R}^3$ . As in the previous section, let  $E$  denote the set of edges of triangles in  $S$ , and let  $E^* = \{e^* \mid e \in E\}$  be the set of  $xy$ -projections of the edges in  $E$ . Let  $k$  be the number of intersections between the edges in  $E^*$ . Our algorithm constructs a BSP  $\mathcal{B}$  of size  $O((n+k) \log^2 n)$  and height  $O(\log n)$  in  $O((n+k) \log^3 n)$  time. The algorithm is a three-dimensional extension of Paterson and Yao's algorithm for constructing a BSP for segments in the plane [29]. The cuts we make are either free cuts contained in triangles of  $S$  or vertical extensions of the cuts made by the Paterson–Yao algorithm when applied to the  $xy$ -projections of the edges of the triangles in  $S$ . Before presenting our algorithm, we give some definitions.

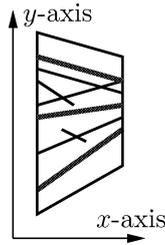


Fig. 12. Anchored segments in  $E_v^*$  (these segments are drawn thick).

As in the previous section, each node  $v$  of  $\mathcal{B}$  is associated with a cylindrical cell  $\Delta_v$ , but the top and bottom faces of  $\Delta_v$  are now trapezoids. Let  $\Delta_v^*$  denote the  $xy$ -projection of the top (or bottom) face of  $\Delta_v$ ; two of the edges of  $\Delta_v^*$  are perpendicular to the  $x$ -axis. We refer to the faces of  $\Delta_v$  passing through these edges as the *left* and *right* faces and to the other two vertical faces of  $\Delta_v$  as the *front* and *back* faces.

Let  $E_v^*$  be the set of  $xy$ -projections of the segments in  $E$  that intersect the interior of  $\Delta_v$  and are clipped within  $\Delta_v^*$ . A segment  $\gamma \in E_v^*$  is called *anchored* if its endpoints lie on the two parallel edges of  $\Delta_v^*$  and  $\gamma$  does not intersect any other segment of  $E^*$ .<sup>8</sup> Fig. 12 shows an example. The anchored segments in  $E_v^*$  can be linearly ordered by  $y$ -coordinate (since they are disjoint). Let  $A_v$  be the set of anchored segments in  $E_v^*$  and let  $P_v$  be the set of intersection points between the segments of  $E_v^*$ . Finally, let  $F_v \subseteq S_v$  be the set of all free triangles in  $S_v$ . Recall that a triangle  $s \in S_v$  is free with respect to  $\Delta_v$  if no edge of  $s$  intersects the interior of  $\Delta_v$ ;  $s$  partitions  $\Delta_v$  into two cylindrical cells. Since  $\Delta_v$  is a cylindrical cell, the triangles in  $F_v$  can be sorted by their  $z$ -coordinates. Before describing the algorithm, we state two useful lemmas that are easy to prove.

**Lemma 5.1.** *Let  $v$  be a node in  $\mathcal{B}$  such that  $F_v \neq \emptyset$ . Then one of the following conditions holds:*

- (i) *There exists a triangle  $t \in F_v$  such that the plane containing  $t$  splits  $\Delta_v$  into two cells  $\Delta_w$  and  $\Delta_z$  and  $|P_w|, |P_z| \leq 2|P_v|/3$ .*
- (ii) *There exist two consecutive triangles  $t_1, t_2 \in F_v$  such that the planes containing  $t_1$  and  $t_2$  split  $\Delta_v$  into three cells  $\Delta_w, \Delta_y$  and  $\Delta_z$  where the top and bottom faces of  $\Delta_y$  are contained in  $t_1$  and  $t_2$ ,  $|P_y| > |P_v|/3$ , and  $F_y = \emptyset$ .*

**Lemma 5.2.** *Let  $v$  be a node in  $\mathcal{B}$  such that  $F_v = \emptyset$  and  $A_v \neq \emptyset$ . Then one of the following conditions holds:*

- (i) *There exists an anchored edge  $e \in A_v$  such that the vertical plane containing  $e$  splits  $\Delta_v$  into two cells  $\Delta_w$  and  $\Delta_z$  and  $|P_w|, |P_z| \leq 2|P_v|/3$ .*
- (ii) *There exist two consecutive anchored edges  $e_1, e_2 \in A_v$  such that the vertical planes containing  $e_1$  and  $e_2$  split  $\Delta_v$  into three cells  $\Delta_w, \Delta_y$  and  $\Delta_z$  where the front and back faces of  $\Delta_y$  are bounded by the planes containing  $e_1$  and  $e_2$ ,  $|P_y| > |P_v|/3$ , and  $A_y = \emptyset$ .*

<sup>8</sup> This definition requires that  $\gamma$  not intersect the  $xy$ -projection of any segment  $\gamma' \in E$  even when  $\gamma'$  does not intersect  $\Delta_v$ . An alternative definition could have required that  $\gamma$  not intersect any other segment in  $E_v^*$ . The reason we use our definition will become clear when we describe our algorithm.

### 5.1. The deterministic algorithm

Let  $I$  be the set of intersection points of  $E^*$ . Suppose  $|I| = k$ . In a pre-processing step, we compute the set  $I$  in  $O((n+k)\log n)$  time using the Bentley–Ottman sweep-line algorithm [7,31]. Our algorithm then constructs  $\mathcal{B}$  in a top-down fashion by maintaining a top subtree of  $\mathcal{B}$ . We say that a leaf  $v$  of the subtree is *active* if  $S_v \neq \emptyset$ . Note that  $v$  is active even if  $S_v$  contains free triangles; in contrast, in Section 4, an active leaf  $v$  has the property that  $S_v$  does not contain any free triangles. We store the set of all active leaves of the current subtree in a list. For each active leaf  $v$ , we maintain the sets  $P_v$ ,  $A_v$ ,  $F_v$  and  $S_v$ . We can easily obtain the set  $E_v^*$  from the sets  $A_v$  and  $S_v \setminus F_v$ .

At each step of the algorithm, we choose an active leaf  $v$ , compute at most two cutting planes, and use these planes to split  $\Delta_v$  into at most three cells. For each child  $w$  of  $v$ , if  $S_w$  is nonempty, we mark  $w$  as being active. Before describing how we compute the cutting planes, we specify how we determine the sets  $P_w$ ,  $F_w$ ,  $S_w$  and  $A_w$  (the procedure is symmetric for the other children of  $v$ ):

$P_w$ : Let  $p \in P_v$  be the intersection point of  $e_1^*$  and  $e_2^*$ , where  $e_1$  and  $e_2$  are triangle edges;  $p \in P_w$  if both  $e_1$  and  $e_2$  intersect  $\Delta_w$  and  $p$  is contained in  $\Delta_w^*$ .

$F_w$ : Let  $s$  be a triangle in  $S_v$ . If  $s$  intersects  $\Delta_w$  but none of the edges of  $s$  intersects the interior of  $\Delta_w$ , then  $s \in F_w$ .

$S_w$ : Since  $S_w$  is the union of  $F_w$  and  $S_v \setminus F_w$ , it is enough to specify how to compute  $S_v \setminus F_w$ . Let  $s$  be a triangle in  $S_v \setminus F_v$ . If an edge of  $s$  intersects the interior of  $\Delta_w$ , then  $s \in S_w \setminus F_w$ .

$A_w$ : Let  $e^* \in E_v^*$ , where  $e$  is an edge of a triangle in  $S_v$ . There are two cases to consider: (i) If  $e^* \in A_v$  and  $e$  intersects  $\Delta_w$ , then  $e^* \in A_w$ . (ii) If  $e^* \notin A_v$ ,  $e$  intersects  $\Delta_w$ , and no point in  $I$  is contained in  $\Delta_w^* \cap e^*$ , then  $e^* \in A_w$ . To detect the second case, for each edge  $e^* \in E_v^*$ , we store the set of points in  $I$  that are contained in  $e^*$  (these points are formed by the intersection of  $e^*$  and other segments in  $E^*$ ).

It is clear that these sets can be computed for all children of  $v$  in  $O(|P_v| + |F_v| + |S_v| + |A_v|)$  time.

We now describe how we compute the cutting planes we use to partition  $v$ . Our algorithm uses three kinds of cuts: a *face* cut is a plane containing a triangle in  $S$  (all face cuts will be free cuts), an *edge* cut is a vertical plane erected on an anchored segment in  $A_v$ , and a *point* cut is a plane perpendicular to the  $x$ -axis passing through a point in  $P_v$ . See Fig. 13. We choose the cutting planes as follows.

1.  $F_v \neq \emptyset$ : We apply Lemma 5.1 to select a set  $\Phi$  of at most two free triangles in  $F_v$  and split  $\Delta_v$  into at most three cells using face cuts contained in the triangles in  $\Phi$ . See Fig. 13(i). Since no triangle in  $\Phi$  intersects any triangle of  $S_v$ , each triangle of  $S_v \setminus \Phi$  belongs to exactly one of the cells we partition  $\Delta_v$  into. We can similarly partition the anchored segments in  $A_v$ .<sup>9</sup>
2.  $F_v = \emptyset$  and  $A_v \neq \emptyset$ : We apply Lemma 5.2 to select a set  $\Upsilon$  of at most two anchored segments in  $A_v$  and split  $\Delta_v$  using edge cuts passing through the segments in  $\Upsilon$ . See Fig. 13(ii). Since no segment in  $A_v \setminus \Upsilon$  intersects the vertical planes erected on the segments in  $\Upsilon$ , we can partition the anchored segments of  $A_v$  between the cells that  $\Delta_v$  is split into.
3.  $F_v = \emptyset$ ,  $A_v = \emptyset$ : We split  $\Delta_v$  into two cells using the point cut through the vertex in  $P_v$  with the median  $x$ -coordinate. See Fig. 13(iii).

<sup>9</sup> If  $\Delta_w$  is one of the cells that we split  $\Delta_v$  into using face cuts, then  $A_w \subseteq A_v$ . This property holds because we defined a segment  $\gamma \in E_v^*$  to be anchored if  $\gamma$  does not intersect any other segment in  $E^*$ , i.e., if no point in  $I$  lies on  $\gamma$ . On the other hand, if we had defined  $\gamma$  to be anchored if  $\gamma$  does not intersect any other segment in  $E_v^*$  (i.e., if no point in  $P_v$  lies on  $\gamma$ ), then it is possible that  $\gamma \in A_w$  but  $\gamma \notin A_v$ . This possibility arises when there is a triangle edge  $\beta$  that intersects  $\Delta_v$  and  $\beta^*$  intersects  $\gamma$  but  $\beta$  does not intersect  $\Delta_w$ .

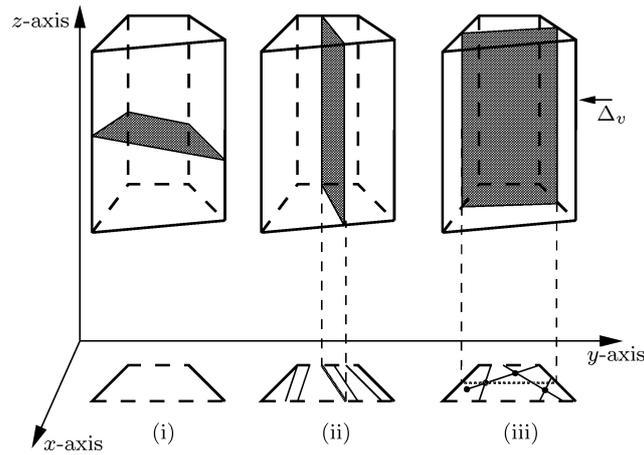


Fig. 13. Cuts made in the deterministic algorithm: (i) free cut, (ii) cut parallel to the  $z$ -axis through an anchored segment, and (iii) cut parallel to the  $yz$ -plane through a vertex of  $P_v$ .

### 5.2. Analysis of the algorithm

We now analyze the performance of the algorithm. We first bound the size of  $\mathcal{B}$ , then the running time of the algorithm, and finally the height of  $\mathcal{B}$ . Let  $v$  be a node in  $\mathcal{B}$  such that  $P_v$  contains  $p$  intersection points,  $A_v$  contains  $a$  anchored segments, and  $F_v$  contains  $f$  free triangles; clearly  $S_v$  contains  $O(p + a + f)$  triangles. Let  $\mathcal{B}_v$  denote the subtree of  $\mathcal{B}$  rooted at  $v$ . Set

$$S(p, a, f) = \max_v |\mathcal{B}_v|,$$

where the maximum is taken over all nodes  $v$  with  $|P_v| = p$ ,  $|A_v| = a$  and  $|F_v| = f$ . We bound  $S(p, a, f)$  by setting up a recurrence for it. Suppose the cutting planes chosen at  $v$  partition  $\Delta_v$  into the cells  $\Delta_w$ ,  $\Delta_y$  and  $\Delta_z$ . We use the convention that  $\Delta_y$  is empty if we chose only one cutting plane at  $v$  and that  $\Delta_y$  lies “between”  $\Delta_w$  and  $\Delta_z$  if we choose two cutting planes at  $v$ . Let  $p_w = |P_w|$ ,  $a_w = |A_w|$  and  $f_w = |F_w|$ ; define  $p_y, a_y, f_y, p_z, a_z$  and  $f_z$  similarly. Note that  $P_w, P_y$  and  $P_z$  are disjoint subsets of  $P_v$ ; therefore,  $p_w + p_y + p_z \leq p$ . We consider three cases:

- (i)  $f \neq 0$ . Since we partition  $\Delta_v$  using free cuts, we have  $a_w + a_y + a_z = a$  and  $f_w + f_y + f_z < f$ . If we use one free cut to partition  $v$  (in this case,  $\Delta_y$  is empty), Lemma 5.1(i) implies that  $p_w, p_z \leq 2p_v/3$ . If we use two (consecutive) free cuts to partition  $v$ , Lemma 5.1(ii) implies that  $p_w + p_z \leq 2p_v/3$ ,  $p_y > p_v/3$  and  $f_y = 0$ .
- (ii)  $f = 0, a \neq 0$ . Since the edge cuts we use to partition  $\Delta_v$  are erected on anchored segments in  $A_v$ , we have  $a_w + a_y + a_z \leq a - 1$ . These cuts may intersect triangles in  $S_v$ , creating free triangles in  $F_w, F_y$  and  $F_z$ . However, there are  $O(p + a)$  triangles in  $S_v$ ; hence, we have that  $f_w + f_y + f_z = O(p + a)$ . If we use one edge cut to partition  $v$  (in this case,  $\Delta_y$  is empty), Lemma 5.2(i) implies that  $p_w, p_z \leq 2p_v/3$ . If we use two (consecutive) edge cuts to partition  $v$ , Lemma 5.2(ii) implies that  $p_w + p_z \leq 2p_v/3, p_y > p_v/3$  and  $a_y = 0$ .
- (iii)  $f = 0, a = 0$ . We split  $\Delta_v$  using a point cut defined by the vertex in  $P_v$  with the median  $x$ -coordinate, which implies that  $p_w, p_z \leq p/2$ . Since the cut may intersect triangles in  $S_v$ , both  $\Delta_w$  and  $\Delta_z$  can

contain anchored segments and free triangles. Since there are  $O(p)$  edges in  $E_v^*$  and  $O(p)$  triangles in  $S_v$ , we have  $a_w, a_z, f_w, f_z = O(p)$ .

If we choose one edge cut or one point cut at  $v$ , then the size of node  $v$  is one. If we pick one face cut at  $v$ , then the size of node  $v$  is two, since we store the free triangle inducing this face cut at  $v$ . If we choose two cuts at  $v$ , assume without loss of generality that  $w$  is a child of  $v$  and that  $y$  and  $z$  are the children of  $x$ , the other child of  $v$ . If we pick two edge cuts at  $v$ , then  $v$  and  $x$  each have size one. If we select two face cuts at  $v$ , then  $v$  and  $x$  each have size two. The preceding discussion implies that we can write the following recurrence for  $S(p, a, f)$ :

$$S(p, a, f) \leq S(p_w, a_w, f_w) + S(p_y, a_y, f_y) + S(p_z, a_z, f_z) + 4, \quad (5.1)$$

where  $p_w + p_y + p_z \leq p$ , and

1.  $a_w + a_z = a$ ,  $f_w + f_z = f - 1$ ,  $p_w, p_z \leq 2p/3$ , and  $p_y = a_y = f_y = 0$ , if  $f \neq 0$  and we apply Lemma 5.1(i).
2.  $a_w + a_y + a_z = a$ , and  $f_w + f_z = f - 2$ ,  $f_y = 0$ ,  $p_w + p_z \leq 2p/3$ , and  $p_y > p/3$ , if  $f \neq 0$  and we apply Lemma 5.1(ii).
3.  $a_w + a_z = a - 1$ ,  $f_w + f_z = O(p + a)$ ,  $p_w, p_z \leq 2p/3$ , and  $p_y = a_y = f_y = 0$ , if  $f = 0$ ,  $a \neq 0$  and we apply Lemma 5.2(i).
4.  $a_w + a_z = a - 2$ ,  $a_y = 0$ ,  $f_w + f_y + f_z = O(p + a)$ ,  $p_w + p_z \leq 2p/3$ , and  $p_y > p/3$ , if  $f = 0$ ,  $a \neq 0$  and we apply Lemma 5.2(ii).
5.  $p_w, p_z \leq p/2$ ,  $a_w, a_z, f_w, f_z = O(p)$ , and  $p_y = a_y = f_y = 0$  if  $f = 0$ ,  $a = 0$ .

Using mathematical induction, we can prove that the solution to this recurrence is

$$S(p, a, f) = O(p \log^2 p + (p + a) \log p + f).$$

Since the root node of  $\mathcal{B}$  has  $n + k$  intersection points, no anchored segments, and no free triangles, and since  $k = O(n^2)$ , we obtain the following lemma.

**Lemma 5.3.** *The size of  $\mathcal{B}$  is  $O((n + k) \log^2 n)$ .*

We now analyze the running time of the algorithm. As we have noted earlier, at each node  $v$ , we can choose the cutting planes and perform the operations to split  $\Delta_v$  in  $O(p + a + f)$  time. If  $T(p, a, f)$  denotes the maximum time taken by our algorithm to construct the subtree of  $\mathcal{B}$  rooted at a node  $v$  with  $|P_v| = p$ ,  $|A_v| = a$  and  $|F_v| = f$  (the maximum is taken over all such nodes  $v$ ), we have

$$T(p, a, f) = T(p_w, a_w, f_w) + T(p_y, a_y, f_y) + T(p_z, a_z, f_z) + O(p + a + f),$$

where  $p_w, a_w, f_w, p_y, a_y, f_y, p_z, a_z$  and  $f_z$  satisfy the same conditions as in (5.1). Using mathematical induction, we can prove that the solution to the above recurrence is

$$T(p, a, f) = O(p \log^3 p + (p + a) \log^2 p + (p + a + f) \log p).$$

Thus, we obtain the following lemma.

**Lemma 5.4.** *The time taken by our algorithm to construct  $\mathcal{B}$  is  $O((n + k) \log^3 n)$ .*

We now prove a lemma that implies that the height of  $\mathcal{B}$  is  $O(\log n)$ . We first need a simple definition: if  $v$  is a node of  $\mathcal{B}$  and  $w$  is a descendant of  $v$  in  $\mathcal{B}$ , then the *distance* between  $v$  and  $w$  is the number of tree edges in the path from  $v$  to  $w$ .

**Lemma 5.5.** *Let  $v$  be a node in  $\mathcal{B}$ . If  $w$  is a descendant of  $v$  and the distance between  $w$  and  $v$  is seven, then  $p_w \leq 2p_v/3$ .*

**Proof.** We first define some notation that will be useful in the proof. For a node  $v \in \mathcal{B}$ , consider the subtree  $\mathcal{T}$  rooted at  $v$  such that if  $w$  is a leaf of  $\mathcal{T}$ , then  $p_w \leq 2p_v/3$  and if  $w$  is an interior node of  $\mathcal{T}$ , then  $p_w > 2p_v/3$ . We use  $d_v$  to denote the height of  $\mathcal{T}$  (the height of a tree is the maximum distance between the root and a leaf of the tree). We claim that for any node  $v \in \mathcal{B}$ ,  $d_v \leq 7$ . Clearly, the lemma is true if we prove this claim.

If we choose one cutting plane at  $v$ , then  $d_v = 1$  since for each child  $w$  of  $v$ , we have  $p_w \leq 2p_v/3$ . Suppose we choose two cutting planes at  $v$ . These cuts split  $\Delta_v$  into three regions  $\Delta_w$ ,  $\Delta_y$  and  $\Delta_z$  such that  $p_y > p_v/3$ ; therefore,  $d_v = d_y + 2$ . If we apply Lemma 5.1(ii) at  $v$ , we have  $F_y = \emptyset$ . Similarly, if we apply Lemma 5.2(ii) at  $v$ , we have  $A_y = \emptyset$ .

We now prove a bound on  $d_y$ . If we choose one cutting plane at  $y$ , we have  $d_y = 1$ , which proves that  $d_v = 3$ . Let us now consider the other possibilities (i.e., we split  $\Delta_y$  using Lemma 5.1(ii) or Lemma 5.2(ii)). Let  $y'$  be the grandchild of  $y$  such that  $p_{y'} > 2p_y/3$ , which implies that  $d_y = d_{y'} + 2$ . Since either  $F_y$  or  $A_y$  is empty, we consider the two possible cases:

1.  $F_y \neq \emptyset$  and  $A_y = \emptyset$ : Since we apply Lemma 5.1(ii) at  $y$ ,  $F_{y'} = \emptyset$ . Further,  $A_{y'} = \emptyset$  since the face cuts that split  $\Delta_y$  do not create any new anchored edges. Therefore, we split  $y'$  using a point cut, which implies that  $d_{y'} = 1$ ; therefore,  $d_y = 3$ .
2.  $F_y = \emptyset$  and  $A_y \neq \emptyset$ : Since we apply Lemma 5.2(ii) at  $y$ ,  $A_{y'} = \emptyset$ . Applying the argument of the previous case to  $y'$ , we have  $d_{y'} = 3$ , which implies that  $d_y = 5$ .

This argument shows that  $d_v \leq 7$  for all nodes  $v \in \mathcal{B}$ .  $\square$

Combining the last three lemmas, we state the main result of this section.

**Theorem 5.6.** *Let  $S$  be a set of  $n$  triangles in  $\mathbb{R}^3$ , and let  $k$  be the number of intersection points of the  $xy$ -projections of the edges of  $S$ . We can compute a BSP of size  $O((n+k)\log^2 n)$  and height  $O(\log n)$  for  $S$  in  $O((n+k)\log^3 n)$  time.*

## 6. Conclusions

In this paper, we first presented an efficient algorithm to maintain a BSP of a set of moving segments in the plane. Currently, we do not know any non-trivial lower bounds for this problem. Agarwal et al. [1] have extended our result and developed an algorithm to maintain BSPs for moving triangles in  $\mathbb{R}^3$ .

We have also presented algorithms for constructing BSPs for triangles in  $\mathbb{R}^3$ . The randomized algorithm constructs a BSP of worst-case optimal size and runs in near-optimal time in the worst case. The deterministic algorithm is near-optimal in the worst-case. However, for inputs such as terrains that actually arise in practice, the number of intersections between the  $xy$ -projections of the triangles is likely to be near-linear. In such cases, our deterministic algorithm constructs BSPs of near-linear size.

There are many interesting open questions regarding BSPs. First of all, our deterministic algorithm is likely to construct BSPs of near-linear size for terrains and urban landscapes, which are common in computer graphics and geographic information systems, but might not be very good for data sets in other application domains (e.g., CAD design). Proving near-linear bounds on BSP size in models that capture

the geometric structure of such inputs will be very useful. Secondly, all our algorithms for triangles in  $\mathbb{R}^3$  construct BSPs of  $\Omega(n^2)$  size even if an  $O(n)$  size BSP exists. This raises the question of constructing a BSP of optimal or near-optimal size for triangles in  $\mathbb{R}^3$ . It is not known whether the problem is *NP*-hard.

## Acknowledgements

We would like to thank the anonymous referees for many useful comments and for pointing out some errors in an earlier version of this paper.

## References

- [1] P.K. Agarwal, J. Erickson, L.J. Guibas, Kinetic binary space partitions for intersecting segments and disjoint triangles, in: Proc. 9th ACM–SIAM Symposium on Discrete Algorithms, 1998, pp. 107–116.
- [2] P.K. Agarwal, E.F. Grove, T.M. Murali, J.S. Vitter, Binary space partitions for fat rectangles, in: Proc. 37th Annu. IEEE Sympos. Found. Comput. Sci., October 1996, pp. 482–491.
- [3] P.K. Agarwal, S. Suri, Surface approximation and geometric partitions, in: Proc. 5th ACM–SIAM Sympos. Discrete Algorithms, 1994, pp. 24–33.
- [4] J.M. Airey, Increasing update rates in the building walkthrough system with automatic model-space subdivision and potentially visible set calculations, Ph.D. Thesis, Department of Computer Science, University of North Carolina, Chapel Hill, 1990.
- [5] C. Ballieux, Motion planning using binary space partitions, Technical Report inf/src/93-25, Utrecht University, 1993.
- [6] J. Basch, L.J. Guibas, J. Hershberger, Data structures for mobile data, in: Proc. 8th ACM–SIAM Sympos. Discrete Algorithms, 1997, pp. 747–756.
- [7] J.L. Bentley, T.A. Ottmann, Algorithms for reporting and counting geometric intersections, *IEEE Trans. Comput.* 28 (1979) 643–647.
- [8] M. Bern, D. Eppstein, P. Plassman, F. Yao, Horizon theorems for lines and polygons, in: J. Goodman, R. Pollack, W. Steiger (Eds.), *Discrete and Computational Geometry: Papers from the DIMACS Special Year*, DIMACS Series in Discrete Mathematics and Theoretical Computer Science, Vol. 6, American Mathematical Society, Providence, RI, 1991, pp. 45–66.
- [9] A.T. Campbell, Modeling global diffuse illumination for image synthesis, Ph.D. Thesis, Department of Computer Sciences, University of Texas, Austin, 1991.
- [10] T. Cassen, K.R. Subramanian, Z. Michalewicz, Near-optimal construction of partitioning trees by evolutionary techniques, in: Proc. of Graphics Interface '95, 1995, pp. 263–271.
- [11] B. Chazelle, L.J. Guibas, D.T. Lee, The power of geometric duality, *BIT* 25 (1985) 76–90.
- [12] N. Chin, S. Feiner, Near real-time shadow generation using BSP trees, in: Proc. SIGGRAPH 89, *Comput. Graph.*, Vol. 23, ACM SIGGRAPH, 1989, pp. 99–106.
- [13] N. Chin, S. Feiner, Fast object-precision shadow generation for areal light sources using BSP trees, in: Proc. 1992 Sympos. Interactive 3D Graphics, 1992, pp. 21–30.
- [14] Y. Chrysanthou, Shadow computation for 3D interaction and animation, Ph.D. Thesis, Queen Mary and Westfield College, University of London, 1996.
- [15] K.L. Clarkson, New applications of random sampling in computational geometry, *Discrete Comput. Geom.* 2 (1987) 195–222.
- [16] K.L. Clarkson, P.W. Shor, Applications of random sampling in computational geometry, II, *Discrete Comput. Geom.* 4 (1989) 387–421.

- [17] M. de Berg, Linear size binary space partitions for fat objects, in: Proc. 3rd Annu. European Sympos. Algorithms, Lecture Notes in Computer Science, Vol. 979, Springer, Berlin, 1995, pp. 252–263.
- [18] H. Edelsbrunner, Algorithms in Combinatorial Geometry, Springer, Heidelberg, 1987.
- [19] H. Fuchs, Z.M. Kedem, B. Naylor, On visible surface generation by a priori tree structures, in: Proc. SIGGRAPH 80, Comput. Graph., Vol. 14, ACM SIGGRAPH, 1980, pp. 124–133.
- [20] L.J. Guibas, Kinetic data structures: A state of the art report, in: P.K. Agarwal, L.E. Kavraki, M.T. Mason (Eds.), Robotics: An Algorithmic Perspective, 1998, pp. 191–209.
- [21] D. Haussler, E. Welzl, Epsilon-nets and simplex range queries, Discrete Comput. Geom. 2 (1987) 127–151.
- [22] C. Mata, J.S.B. Mitchell, Approximation algorithms for geometric tour and network design problems, in: Proc. 11th Annu. ACM Sympos. Comput. Geom., 1995, pp. 360–369.
- [23] R. Motwani, P. Raghavan, Randomized Algorithms, Cambridge University Press, New York, 1995.
- [24] T.M. Murali, T.A. Funkhouser, Consistent solid and boundary representations from arbitrary polygonal data, in: Proc. 1997 Sympos. Interactive 3D Graphics, 1997.
- [25] B. Naylor, W. Thibault, Application of BSP trees to ray-tracing and CSG evaluation, Technical Report GIT-ICS 86/03, Georgia Institute of Tech., School of Information and Computer Science, February 1986.
- [26] B.F. Naylor, SCULPT: an interactive solid modeling tool, in: Proc. Graphics Interface '90, 1990, pp. 138–148.
- [27] B.F. Naylor, Interactive solid geometry via partitioning trees, in: Proc. Graphics Interface '92, 1992, pp. 11–18.
- [28] B.F. Naylor, J. Amanatides, W.C. Thibault, Merging BSP trees yields polyhedral set operations, in: Proc. SIGGRAPH 90, Comput. Graph., Vol. 24, ACM SIGGRAPH, 1990, pp. 115–124.
- [29] M.S. Paterson, F.F. Yao, Efficient binary space partitions for hidden-surface removal and solid modeling, Discrete Comput. Geom. 5 (1990) 485–503.
- [30] M.S. Paterson, F.F. Yao, Optimal binary space partitions for orthogonal objects, J. Algorithms 13 (1992) 99–113.
- [31] F.P. Preparata, M.I. Shamos, Computational Geometry: An Introduction, Springer, New York, 1985.
- [32] R.A. Schumacker, R. Brand, M. Gilliland, W. Sharp, Study for applying computer-generated images to visual simulation, Technical Report AFHRL-TR-69-14, U.S. Air Force Human Resources Laboratory, 1969.
- [33] R. Seidel, Backwards analysis of randomized geometric algorithms, in: J. Pach (Ed.), New Trends in Discrete and Computational Geometry, Springer, Heidelberg, 1993, pp. 37–68.
- [34] S.J. Teller, Visibility computations in densely occluded polyhedral environments, Ph.D. Thesis, Department of Computer Science, University of California, Berkeley, 1992.
- [35] W.C. Thibault, B.F. Naylor, Set operations on polyhedra using binary space partitioning trees, in: Proc. SIGGRAPH 87, Comput. Graph., Vol. 21, ACM SIGGRAPH, 1987, pp. 153–162.
- [36] E. Torres, Optimization of the binary space partition algorithm (BSP) for the visualization of dynamic scenes, in: Eurographics '90, North-Holland, 1990, pp. 507–518.