Binary Space Partitions for Axis-Parallel Segments, Rectangles, and Hyperrectangles*

Adrian Dumitrescu† Joseph S.B. Mitchell‡ Micha Sharir§

June 1, 2001

Abstract

We provide a variety of new results, including upper and lower bounds, as well as simpler proof techniques for the efficient construction of binary space partitions (BSP's) of axis-parallel segments, rectangles, and hyperrectangles. (a) A consequence of the analysis in [1] is that any set of n axis-parallel and pairwise-disjoint line segments in the plane admits a binary space partition of size at most $2n-1$. We establish a worst-case lower bound of $2n-o(n)$ for the size of such a BSP, thus showing that this bound is almost tight in the worst case. (b) We give an improved worst-case lower bound of $\frac{7}{3}n-o(n)$ on the size of a BSP for isothetic pairwise disjoint rectangles. (c) We present simple methods, with equally simple analysis, for constructing BSP's for axis-parallel segments in higher dimensions, simplifying the technique of [10] and improving the constants. (d) We obtain an alternative construction (to that in [10]) of BSP's for collections of axis-parallel rectangles in 3-space. (e) We present a construction of BSP's of size $O(n^{5/3})$ for n axis-parallel pairwise disjoint 2-rectangles in $\mathbb{R}^4$, and give a matching worst-case lower bound of $\Omega(n^{5/3})$ for the size of such a BSP. (f) We extend the results of [10] to axis-parallel k-dimensional rectangles in $\mathbb{R}^d$, for $k < d/2$, and obtain a worst-case tight bound of $\Theta(n^{d/(d-k)})$ for the size of a BSP of n rectangles. Both upper and lower bounds also hold for $d/2 \leq k \leq d-1$ if we allow the rectangles to intersect.

---

*A preliminary version of this paper appears in the *Proc. 17th Annual ACM Symposium on Computational Geometry*, June 3-5, 2001, Medford, MA. Work by Joe Mitchell is partially supported by HRL Laboratories, National Science Foundation (CCR-9732221), NASA Ames Research Center, Northrop-Grumman Corporation, Sandia National Labs, Seagull Technology, and Sun Microsystems. Work by Micha Sharir has been supported by NSF Grant CCR-97-32101, by a grant from the U.S.-Israeli Binational Science Foundation, by a grant from the Israeli Academy of Sciences for a Center of Excellence in Geometric Computing at Tel Aviv University, by the ESPRIT IV LTR project No. 21957 (CGAL), and by the Hermann Minkowski–MINERVA Center for Geometry at Tel Aviv University. Part of the work on this paper was done when Joe Mitchell visited Tel Aviv University during July–August 2000.

†Department of Electrical Engineering and Computer Science, University of Wisconsin at Milwaukee, Milwaukee, WI 53201-0784, USA; `ad@cs.uwm.edu`.

‡Department of Applied Mathematics and Statistics, SUNY at Stony Brook, Stony Brook, NY 11794, USA; `jsbm@ams.sunysb.edu`.

§School of Computer Science, Tel Aviv University, Tel Aviv 69978, Israel and Courant Institute of Mathematical Sciences, New York University, New York, NY 10012, USA; `sharir@math.tau.ac.il`.

1 Introduction

Let S be a collection of n pairwise-disjoint objects in $\mathbb{R}^d$. A *binary space partition* (or *BSP* for short) for S is a recursively defined convex subdivision of space obtained by cutting space into two open regions C, C' by a hyperplane, and by constructing recursively a BSP for $\{s \cap C \mid s \in S\}$ within C and a BSP for $\{s \cap C' \mid s \in S\}$ within C'; the process terminates when each (open) cell of the BSP is intersected by at most one object of S. The input objects are usually assumed to be pairwise disjoint; however, the definition of a BSP we have given is applicable to any set of polyhedral objects in $\mathbb{R}^d$ for which no pair of objects intersect in a full-dimensional set. In particular, the definition applies to arbitrary sets of polyhedral objects S that are not full-dimensional. The definition also applies to sets of curved objects; however, the number of cells of the BSP will depend on the complexity of polyhedral separators between pairs of objects, and a BSP may not exist if the curved objects are allowed to intersect.

Binary space partitions were introduced by the computer graphics community [6, 7, 8] and have numerous applications for rendering, ray shooting and tracing, solid modeling, rectangle tiling, etc. (see [5, 13]).

Ideally, a BSP for S should not split any object of S into pieces, and wind up with each object lying fully in a separate cell or fully on a cutting hyperplane (see [3]). However, in most cases this is impossible; cutting space (and objects) with a hyperplane may create *fragments* of objects that either lie in one of the two open halfspaces bounded by the hyperplane or are contained in the hyperplane. The *size* of a BSP for a set of k-dimensional objects ($k \leq d$) in $\mathbb{R}^d$ is the number of k-dimensional fragments of objects of S that it produces (in the ideal case, the size is n [3]). We remark that one often defines the size of a BSP as the number of convex regions of $\mathbb{R}^d$ in the decomposition. Here, we opt for the definition in terms of the number of object fragments; this choice makes some of the analysis cleaner, particularly in the case of BSPs of low-dimensional hyperrectangles within higher-dimensional spaces.

One of the major research directions in this area is to construct BSP's of small size. Paterson and Yao [9, 10] proved that for a set of n line segments in the plane there exists a BSP of size $O(n \log n)$, and when the segments are *rectilinear* (i.e., parallel to the coordinate axes), there exists a BSP of size at most $3n$. They have also conjectured that any set of n line segments admits a BSP of size $O(n)$. Only very recently, a lower bound of $\Omega(n \frac{\log n}{\log \log n})$, in the worst case, has been obtained by C. Tóth [11]. For both the general case and the axis-parallel case, BSP's with the above size bounds can be computed in $O(n \log n)$ time. Linear size bounds for BSP's have also been obtained (see [2]) for sets of *fat* objects, for sets of line segments having a bounded ($O(1)$) ratio between the lengths of the longest and the shortest segments, and for sets of *homothetic* objects. A linear bound for segments in the plane with a fixed number of orientations was recently obtained by Tóth [12].

The bound $3n$ of [10] on the size of a BSP for n axis-parallel line segments has been (implicitly) improved by d'Amore and Franciosa in [1] to $2n$. What they have explicitly obtained is an upper bound of $4n$ for axis-parallel rectangles in 2D. By specializing their analysis to the case of segments, one obtains the bound $2n$ (or rather $2n - 1$, with an obvious optimization). Once more, in [4], the same approach, accompanied by an improved charging scheme, has reduced the upper bound for the case of axis-parallel rectangles to $3n$.

In this paper we consider the case in which S is a collection of axis-parallel[1] hyperrectangles of various dimensions in $\mathbb{R}^d$. While we do consider some cases in which the hyperrectangles are allowed to intersect, unless otherwise specified, we assume from now on that they are pairwise-disjoint. The simplest such case is that of line segments (or rectangles) in the plane, mentioned above. In higher dimensions, Paterson and Yao [10] show that if S is a set of n (axis-parallel) line segments in $\mathbb{R}^d$, then S admits a BSP of size $O(n^{d/(d-1)})$, and that this bound is tight in the worst case. They also consider the case of rectangles in $\mathbb{R}^3$ and show that the same bound, $O(n^{3/2})$, can also be obtained for rectangles (and also for boxes). They leave as an open problem to obtain sharp bounds for the size of BSP's for higher-dimensional hyperrectangles in dimensions ≥ 4.

In this paper we improve and simplify the analysis leading to the bounds obtained in [10] and derive new bounds for the cases left open in that paper. Specifically, we first consider the case of axis-parallel line segments in the plane. We show, in Section 2, that there exist collections of n (axis-parallel) line segments in the plane for which any rectilinear BSP has size at least $2n - o(n)$, thus showing that the upper bound of [1] is worst-case tight, apart for lower-order terms. This answers the open problem posed by S. Smorodinsky at EuroCG'2000. We also give an improved lower bound for the case of rectangles in the plane, showing that in the worst case a BSP must have size at least $\frac{7}{3}n - o(n)$.

We then consider the case of axis-parallel line segments in $\mathbb{R}^d$, and obtain, in Section 3, a very simple alternative construction of BSP's of size $O(n^{d/(d-1)})$. In addition to being simpler, the constants of proportionality that our method yields appear to be considerably smaller than those produced by the technique of [10]. We also consider the case of rectangles in $\mathbb{R}^3$, and again present an alternative simple construction of a BSP of size $O(n^{3/2})$.

We then (Sections 4,5) consider higher-dimensional cases. We provide the first nontrivial bounds on the worst-case size of BSP's for hyperrectangles of dimension greater than 1 in higher dimensions, showing that, for $k < d/2$ there exist BSP's of size $O(n^{d/(d-k)})$ for a set of n k-rectangles in d dimensions, and that this bound is tight in the worst case. This bound subsumes the bound $\Theta(n^{d/(d-1)})$ for segments, mentioned above. In fact, both upper and lower bounds hold for any $k \leq d-1$ if our rectangles are allowed to intersect.

The next simplest case that is not covered by the results reported so far is the case $k = 2$ and $d = 4$. We show that a BSP of size $O(n^{5/3})$ exists for a set of n disjoint (axis-parallel) 2-rectangles in $\mathbb{R}^4$, improving the bound $O(n^2)$ that follows from the general bounds (for possibly intersecting rectangles) stated in the preceding paragraph. We also have a matching lower bound of $\Omega(n^{5/3}))$, showing that our upper bound is tight.

Our results are summarized in Table 1.

We make one final remark regarding the definition of BSP and how it applies to sets of hyperrectangles that are not in general position[2] in higher dimensions. One can modify the definition of BSP in order to require that lower-dimensional subconfigurations of objects that are contained in a cutting hyperplane are decomposed according to a partitioning of the hyperplane, and then recursively down dimensions. This modification is useful in some

[1]From now on, we freely drop the adjective "axis-parallel"; in all cases, the segments, rectangles, or hyperrectangles we consider are axis-parallel.

[2]Here, a set of hyperrectangles is said to be in *general position* if for each $i = 1, \dots, d$ the x_i coordinates that define the extents of the hyperrectangles are all distinct.

d	k	upper bound	lower bound
2	1	$2n-1$ (†)	$2n - o(n)$
2	2	$3n$ (†)	$7n/3 - o(n)$
d	1	$O(n^{d/(d-1)})$ ($*$)	$\Omega(n^{d/(d-1)})$ ($*$)
3	2	$O(n^{3/2})$ ($*$)	$\Omega(n^{3/2})$ ($*$)
d	$k < d/2$	$O(n^{d/(d-k)})$	$\Omega(n^{d/(d-k)})$
d	$k \leq d-1$ intersecting	$O(n^{d/(d-k)})$	$\Omega(n^{d/(d-k)})$
4	2	$O(n^{5/3})$	$\Omega(n^{5/3})$

Table 1: Summary of our bounds. A bound tagged by (†) indicates a known bound, given for reference. A bound tagged by ($*$) indicates a known bound, rederived here with a simpler proof.

applications; e.g., see Vaněček [14]. It does lead to different complexity bounds on the size of the resulting decomposition, since, for instance, a set of n line segments that lie in a common hyperplane h in $\mathbb{R}^4$ require only a single cutting hyperplane (h) for a BSP by our definition, resulting in two cells and size n (since there are n object fragments). However, the n segments may form a configuration in the three-dimensional space of the hyperplane h such that the decomposition of h by a three-dimensional BSP of the segments requires size $\Omega(n^{3/2})$. (See the discussion in Section 6 of [10].) In this paper, we use the strict definition of BSP defined earlier (and introduced in [8]), not requiring that lower dimensional subconfigurations be recursively decomposed. If the input objects are in general position, there is no modification necessary to our stated bounds on the sizes of BSPs if one were to consider the modified definition of BSP; in case the input is not in general position, though, our complexity bounds would require appropriate modification under the modified definition.

2 Segments and Rectangles in Two Dimensions

We start with the case in which S is a set of n axis-parallel pairwise-disjoint line segments in the plane. By applying the method of d'Amore and Franciosa [1], which is designed for the case of *rectangles* (based on a minor variant of Paterson and Yao's method [10]), we readily obtain the following result for line segments:

Theorem 2.1 *A set of n axis-parallel and pairwise-disjoint line segments in the plane admits a binary space auto-partition of size at most $2n-1$. This BSP can be computed using $O(n \log n)$ time and space and has the additional property that no input segment is cut more than once.*

We turn our attention now to establishing the tightness of the upper bound. The best construction known prior to this work is a *cycle of thickness* $n/4$, shown in Figure 1(a) for $n = 20$, which requires a BSP of size at least $5n/4$.

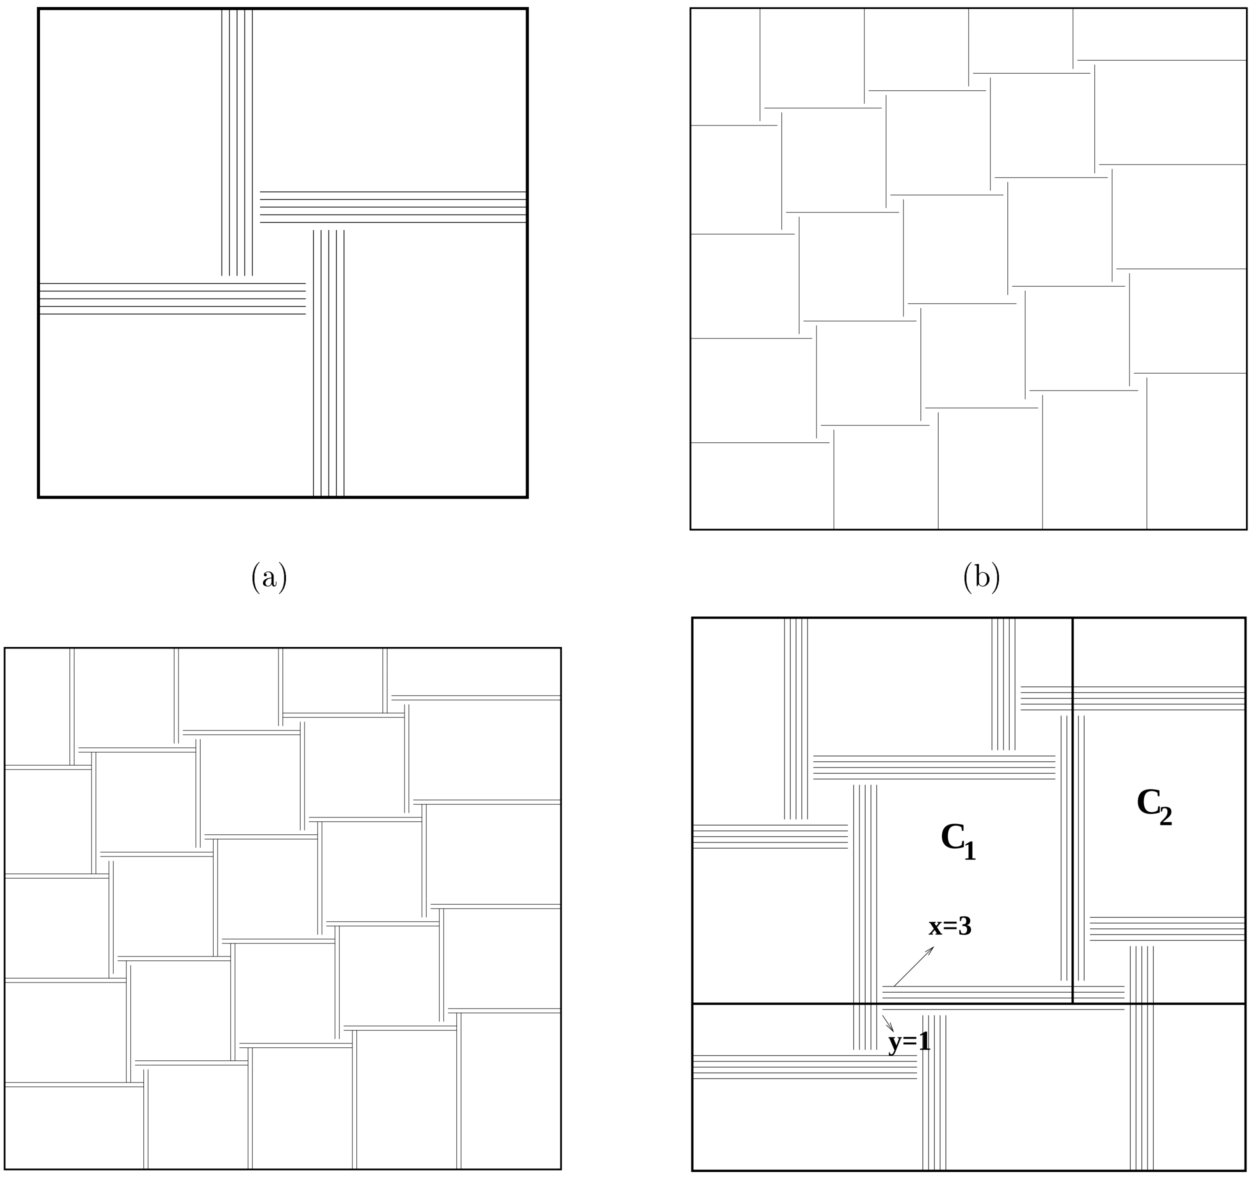

Figure 1: (a). A cycle configuration of thickness $w = 5$. (b). A 4×4 1-grid. (c). A 4×4 2-grid (double grid). (d). Charging scheme in a 5-grid.

Theorem 2.2 *There exists a set S of n disjoint axis-parallel line segments with the property that any binary space auto-partition of S has size at least $2n - o(n)$.*

Proof: If $s, t \in S$, we say that *s cuts t* if by extending the supporting line of s (within the cell of the current BSP containing s and t), the segment t is the first that is cut in two parts. Any segment can cut at most two other segments (on either side of s). This defines a directed simple graph $G = (S, E)$ with vertex set S, and edge set E given by the (asymmetric) cut relation. A cycle is *minimal* if no proper subset of it is a cycle.

In general, the size of the BSP is $n + c_n$, where c_n denotes the number of cuts that are made during the partitioning process. Since $c_n \geq \frac{n}{4}$ in the example of Figure 1(a), it yields the lower bound $n + \frac{n}{4} = \frac{5n}{4}$.

Our construction for the lower bound is based on a grid-like configuration of segments, as illustrated in Figure 1(b,c). We parameterize it by the number of anchored horizontal segments on the left side of R, and the number of anchored vertical segments on the bottom side of R: if these numbers are k and l respectively, we have a $k \times l$ *simple grid* (or a $k \times l$ *1-grid*); see Figure 1(b). We restrict our attention to the set $\mathcal{C}$ of minimal cycles of length 4 associated with the (large and small) cells in the grid. It is easy to see that a BSP is obtained only when all cycles in the set $\mathcal{C}$ have been cut.

First, consider a $k \times k$ 1-grid, to obtain a lower bound on c_n. The number of cells is $e = k^2 + (k-1)^2$, k^2 of which are small cells, and the total number of segments is $n = 2k(k+1)$. Each small cell is a cycle and requires at least one cut by itself. We have $c_n \geq k^2 = \frac{n}{2} - o(n)$, from which we get a lower bound of $1.5n - o(n)$ on the size of any BSP for this configuration.

In order to improve this lower bound, we consider a $k \times k$ *m-grid*, in which each single segment of the simple $k \times k$ 1-grid is replaced by a set of m parallel segments of equal length. A 2-grid is shown in Figure 1(c), and a portion of a 5-grid is shown in Figure 1(d). Here there are $n = 2mk(k+1)$ segments and $e = k^2 + (k-1)^2$ cells.

We claim that $c_n \geq (m - \frac{1}{2})e$; selecting $m = k$ implies a lower bound of $2n - o(n)$ on the size of any auto-partition BSP for this configuration, since

$$\frac{c_n}{n} = \frac{(2k^2 - 2k + 1)(2k - 1)}{4k^2(k+1)} = 1 - o(1).$$

(To be more precise, the lower-order term is $O(n^{2/3})$.)

For a given cell C, denote by $A(C)$ the set of at most $4m$ segments *associated with* C; namely, $A(C)$ consists of m segments on each of the four sides of C. We say that cell C is *touched* by a cutting line l if either (i) l intersects the interior of the cell, viewed as a convex region (square), or (ii) l is one of the supporting lines of the $4m$ segments associated with C. Case (i) defines a *middle cut*, case (ii) defines a *boundary cut*.

The *(current) thickness* of C, denoted $w(C)$, is the minimum number of segments that one needs to cut such that there are no more cycles determined by C. We observe an equivalent characterization of the (current) thickness: the minimum number of segments that one needs to cut, such that $w(C)$ is strictly decreased. In the example of Figure 1(d), the thickness of the cell C_1 after the first horizontal cut is 3, and one needs to cut at least 3 segments to reduce it. After the second vertical cut, the thickness becomes 2. The thickness of each cell is decreased during the cutting process, from m to 0, through one or more cuts.

We employ a charging scheme that charges each cell of the grid with (at least) $m - \frac{1}{2}$. It maintains the following invariant: If a cell has not been touched, its current charge is 0; if a cell has been touched, and its current thickness is w ($0 \leq w \leq m-1$), its charge is (at least) $m - w - \frac{1}{2}$. Let us examine the first cut that touches a fixed cell C_1. The number of locally cut segments in $A(C_1)$ is $c = m$. The charge ($= m$) is distributed to the cell and its neighbor as illustrated in Figure 1(d): $m - x - \frac{1}{2} = m - w_1 - \frac{1}{2}$ to cell C_1 and $m - y - \frac{1}{2}$ to its neighbor cell. The total charge is $(m - x - \frac{1}{2}) + (m - y - \frac{1}{2}) = m$, since $x + y + 1 = m$. Thus, the invariant holds after the first cut that touches C_1 is made.

For any of the subsequent cuts, we distinguish two cases:

Case 1: The cut does not reduce the thickness of any adjacent cell. If the cut does not reduce the current thickness of C_1, we can just ignore the excess charge. If it does reduce w_1, let $w_1' < w_1$ be its reduced thickness. We only have to ensure that the number c of locally cut segments in $A(C_1)$ satisfies $c \geq w_1 - w_1'$ to maintain the invariant. But this is clear, since $c \geq w_1$ from our second (equivalent) characterization of thickness.

Case 2: The cut reduces the thickness w_2 of the adjacent cell C_2 (as well as the thickness w_1 of C_1, otherwise the excess charge is ignored). Denote by $\Delta_1 = w_1 - w_1'$, $\Delta_2 = w_2 - w_2'$, the two reductions. We only have to ensure that the number c of locally cut segments in $A(C_1)$ satisfies $c \geq \Delta_1 + \Delta_2$ to maintain the invariant. Since $w_1' + w_2' = m - 1$ (by the assumption of reduction), this is equivalent to $c + m - 1 \geq w_1 + w_2$ which follows from the inequalities $c \geq w_2$ and $m - 1 \geq w_1$. The first is implied by our equivalent characterization of thickness (more generally, $c \geq \max(w_1, w_2)$), and the second is true since this is not the first cut for C_1.

After the partitioning process ends, each cell has thickness 0, and its charge is (at least) $m - \frac{1}{2}$, as desired. □

We turn now to the case in which S is a set of n disjoint axis-parallel rectangles. The best known upper bound on the size of the BSP is $3n$, given in the recent paper of Berman et al. [4], improving the prior bound of $4n$ [1]. Our construction for line segments gives immediately a lower bound of $2n - o(n)$ (a similar bound was also obtained independently in [4]; however, their bound applies only for rectangles, not for segments). We are able to show an even better lower bound for rectangles:

Theorem 2.3 *There exists a set S of n disjoint axis-parallel rectangles (in fact, of n unit squares, as in Figure 2) with the property that any binary space auto-partition of S has size at least $\frac{7}{3}n - o(n)$.*

Proof: Consider a $k \times k$ square configuration (with $n = k^2$ unit squares) with the pattern shown in Figure 2. We have $s = (k-1)^2$ "junctions," corresponding to the small cells in the grid configuration (of size $\epsilon > 0$), at each of which 4 squares are "meeting". Consider any orthogonal cut in the BSP tree, having (physical) length l. We distinguish 3 types of parts of our cut: (i) the *border* part, if any, of length $l_b \geq 0$ lies on a side or in between the 2 sides of 2 adjacent squares (adjacency in the grid refers to the N,S,E,W squares only); (ii) the *aligned* part, if any, of length $l_a \geq 0$ measures the at most 2 parts of length ≤ 1 adjacent to the border part; (iii) the *middle* part, if any, of length $l_m \geq 0$ encompasses the rest of the length.

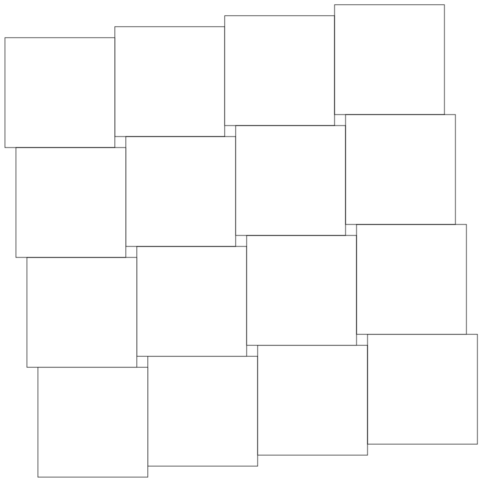

Figure 2: A set of disjoint unit squares used in the lower bound of Theorem 2.3.

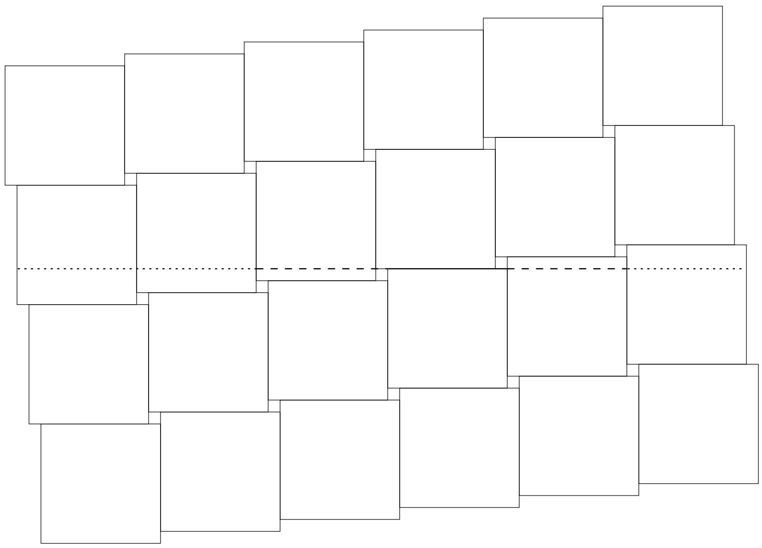

Figure 3: A 4×6 square configuration (an instance for $T(3,5)$); the three portions of the cut are: border (solid), aligned (dashed) and middle (dotted)

We have $l = l_b + l_a + l_m, \quad l_b \leq 1 + \epsilon, \quad l_a \leq 2$. As before, let c_n denote the number of cuts (additional number of object parts in the BSP tree); then, the size of the BSP tree is $n + c_n$. It is easy to see that each of the $(k-1)^2$ junctions will create at least one cut. This gives us a lower bound of $2n - o(n)$ on the size of the BSP tree. To account for more, we will prove a lower bound on the number of additional parts generated by middle portions of the cuts: unless a rectangular subcell R obtained during the BSP tree construction has one "short" side, any cut of R will generate additional parts due to the middle part component, which are unaccounted for by the junctions inside the rectangle. For $i, j \geq 0$, denote by $T(i,j)$ the minimum number of additional "middle" parts in the BSP tree attributed to the middle portions of the cut, obtained when a rectangle containing a complete $i \times j$ array of interior junctions is cut by a line. It is easy to see that $T(3,3) \geq 1$ (a 3×3 array of interior junctions is illustrated in Figure 2). A lower bound on the size of a BSP tree (with $n = k^2$) for a $k \times k$ square configuration is

$$n + c_n = k^2 + (k-1)^2 + T(k-1, k-1).$$

We proceed to prove by induction on $i + j$ the following

Claim 2.4 *For $i, j \geq 3$, $T(i,j) \geq \frac{(i-2)(j-2)}{3}$.*

Proof: The basis $i = j = 3$ holds by the above observation. Without loss of generality assume the orthogonal cut splits the rectangle containing the $i \times j$ array of junctions into 2 subrectangles containing $i \times j_1$ and $i \times j_2$ arrays of junctions, with $j - 1 \leq j_1 + j_2 \leq j, \ j_1 \leq j_2$. Then

$$T(i,j) \geq T(i,j_1) + T(i,j_2) + i - 2.$$

(For example, the bounding rectangle of the square configuration in Figure 3 would contribute $T(3,5)$ "middle" parts, while the 2 resulting subrectangles after the horizontal cut is made would each contribute $T(1,5)$ "middle" parts.) We distinguish 3 cases:

Case 1. $j \leq 5$. Then,

$$T(i,j) \geq i - 2 \geq \frac{(i-2)(j-2)}{3}.$$

The last inequality is satisfied by the choice of j.

Case 2. $j \geq 6, \ j_1 \leq 2$. Since $j_2 \geq j - 3 \geq 3$, we can use the inductive bound on $T(i, j_2)$.

$$T(i,j) \geq T(i,j_2) + i - 2 \geq T(i, j-3) + i - 2$$

$$\geq \frac{(i-2)(j-5)}{3} + i - 2 = \frac{(i-2)(j-2)}{3}.$$

Case 3. $j \geq 6, \ j_1, j_2 \geq 3$. Using the inductive bounds on both terms yields

$$T(i,j) \geq T(i,j_1) + T(i,j_2) + i - 2$$

$$\geq \frac{(i-2)(j_1-2)}{3} + \frac{(i-2)(j_2-2)}{3} + i - 2$$

$$= \frac{(i-2)(j_1+j_2-1)}{3} \geq \frac{(i-2)(j-2)}{3}.$$

□

Our lower bound on the size of the BSP tree becomes

$$k^2 + (k-1)^2 + \frac{(k-3)^2}{3} = \frac{7}{3}n - o(n).$$

We note that, most likely, the constant in the lower bound offered by our square configuration cannot be improved substantially, if at all; certainly, a BSP of size smaller than $2.5n$ can be constructed for this configuration: using a $3 \times k$ strip cutting yields a BSP of size $\approx \frac{22}{9}k^2 \approx 2.444n$, using a $4 \times k$ strip cutting yields a BSP of size $\approx \frac{39}{16}k^2 \approx 2.437n$, etc. □

3 Segments in Higher Dimensions and Rectangles in $\mathbb{R}^3$

Segments in Three Dimensions. Let $E = X \cup Y \cup Z$ be a set of n axis-parallel segments in 3-space, where X (resp. Y, Z) is the subset of segments of E that are parallel to the x-axis (resp. to the y-axis, z-axis). Put $x = |X|$, $y = |Y|$, $z = |Z|$, so that $x + y + z = n$. For simplicity of presentation, suppose that the segments of E are in *general position*, meaning that no two endpoints of different segments have the same x, y or z-coordinate.

We construct a binary space partition of E in the following simple manner.

(i) If one of x, y, z is zero, say $z = 0$, then we can obtain a BSP of size $O(n)$ by a sequence of horizontal cuts.

(ii) Suppose next that each of x, y, z is at least 1 and that $z \leq x, y$. Then we have $z \leq n/3$ and $x+y \geq 2n/3$. Put $t = \lceil 2\sqrt{\frac{xy}{z}} \rceil \geq 2\sqrt{\max\{x, y\}} \geq 2\sqrt{\frac{n}{3}}$. We partition space into a stack of t horizontal slabs $\sigma_1, \ldots, \sigma_t$ by a sequence of horizontal cuts, so that, if x_i, y_i, z_i denote, respectively, the numbers of segments in X, Y, Z that intersect (the interior of) σ_i, then we require that $x_i \leq x/t$, for each i. We clearly also have $\sum_i y_i \leq y$.

For each slab σ_i, project all the segments of E that intersect σ_i onto the xy-plane. We obtain x_i horizontal segments, y_i vertical segments and z_i points. We partition the segments into subsegments at their intersection points. The number of such points is $k_i \leq x_i y_i$ and the total number of subsegments is $x_i + y_i + 2k_i \leq x_i + y_i + 2x_i y_i$.

We apply the planar binary space partitioning scheme of Theorem 2.1, and note that none of the $z_i \leq z$ singleton points will be split. We lift this planar partitioning scheme into three dimensions, lifting each cut by a line (segment) in the xy-plane to a cut by the vertical plane (strip) containing the line (segment). It follows that the size of the partition within σ_i is at most

$$2(x_i + y_i + 2x_i y_i) + z_i \leq 2\frac{x}{t} + 2y_i + \frac{4xy_i}{t} + z.$$

Hence the overall size of the BSP is at most $2x + 2y + \frac{4xy}{t} + zt \leq 2(x+y) + 4\sqrt{xyz} + z$. We have thus shown:

Theorem 3.1 *Let E be a collection of n segments in 3-space, consisting of x segments parallel to the x-axis, y segments parallel to the y-axis and z segments parallel to the z-axis. Then E admits a BSP of size $4\sqrt{xyz} + 2n - z$, for $z \leq x, y$.*

Remark: The maximum value of this bound is easily seen to be at most $\frac{4}{3\sqrt{3}}n^{3/2}+\frac{5}{3}n$. This improves significantly the constant in the bound given in [10]. The lower bound construction given in [10] yields a BSP of size at least $\frac{1}{3\sqrt{3}}n^{3/2}+n$. This leaves the open problem of tightening the gap of the factor 4 in the constant of proportionality between our upper bound and this lower bound.

Segments in Higher Dimensions. Let $E = X_1 \cup X_2 \cup \cdots \cup X_d$ be a set of n axis-parallel segments in d-space, where X_i is the subset of segments of E that are parallel to the x_i-axis, for $i = 1, \ldots, d$. Put $n_i = |X_i|$, for $i = 1, \ldots, d$, so that $n_1 + \cdots + n_d = n$.

We re-establish the following result of [10] with a simpler proof which also gives better constants of proportionality. As noted in [10], the upper bound is tight in the worst case—see also Section 4 below for an extended lower bound.

Theorem 3.2 *Let E be a collection of n segments in d-space, for $d \geq 3$, consisting of n_i segments parallel to the x_i-axis, for $i = 1, \ldots, d$. Then E admits a BSP of size at most*

$$(2d-2)(n_1 n_2 \cdots n_d)^{1/(d-1)} + 2(n_1 + n_2 + \cdots + n_d).$$

Proof: We proceed by induction on d, where the base case $d = 3$ has already been treated. We assume, for simplicity of presentation, that the segments of E are in *general position*, meaning that no two endpoints of different segments have an equal coordinate.

(i) If one of the n_i's is zero, say $n_d = 0$, then we can obtain a BSP of linear size by a sequence of cuts orthogonal to the x_d-axis.

(ii) Suppose next that each of the n_i's is at least 1 and that $n_d \leq n_{d-1} \leq \cdots \leq n_1$. Put

$$t = \left\lceil \frac{(n_1 n_2 \cdots n_{d-1})^{1/(d-1)}}{n_d^{(d-2)/(d-1)}} \right\rceil \geq n_1^{1/(d-1)} \geq (n/d)^{1/(d-1)}.$$

We partition space into a stack of t slabs $\sigma_1, \ldots, \sigma_t$ by a sequence of cuts orthogonal to the x_d-axis, so that the following property holds. Let $n_i^{(\xi)}$ denote the number of segments in X_i that intersect (the interior of) the slab σ_ξ. We require that $n_1^{(\xi)} \leq n_1/t$, for each ξ. We clearly also have $\sum_\xi n_i^{(\xi)} \leq n_i$, for $i = 2, \ldots, d-1$.

For each slab σ_ξ, project all the segments of E that cross σ_ξ onto the hyperplane $x_d = 0$. We obtain a collection of $n_1^{(\xi)} + \cdots + n_{d-1}^{(\xi)}$ segments which, by our general position assumption, are pairwise disjoint (as long as $d > 3$), and $n_d^{(\xi)}$ points. We apply the partitioning algorithm for $d-1$ dimensions, provided by the induction hypothesis, to the projected set, lifting, along the x_d dimension, each $(d-2)$-dimensional cut performed by this algorithm to a $(d-1)$-dimensional cut (within σ_ξ). Note that the presence of points in the input has little effect on the algorithm and adds only a linear term to the size of the resulting BSP: We simply ignore the points and apply the algorithm only to the segments. When we are done, we take the cells of the resulting BSP that contain the input points, and split any such cell that contains more than one point into subcells, say by a sequence of parallel cuts.

By the induction hypothesis, the size of the resulting BSP is at most

$$\sum_{\xi}\left[(2d-4)\left(n_1^{(\xi)}n_2^{(\xi)}\cdots n_{d-1}^{(\xi)}\right)^{\frac{1}{d-2}}+2\left(n_1^{(\xi)}+n_2^{(\xi)}+\cdots+n_d^{(\xi)}\right)\right]\leq$$

$$(2d-4)\left(\frac{n_1}{t}\right)^{\frac{1}{d-2}}\cdot\sum_{\xi}\left(n_2^{(\xi)}\cdots n_{d-1}^{(\xi)}\right)^{\frac{1}{d-2}}+2(n_1+n_2+\cdots+n_{d-1})+2tn_d.$$

We need the following easy inequality:
Claim: Let m be a positive integer and let $a_1,\ldots,a_m,b_1,\ldots,\ b_m$ be nonnegative. Then

$$(a_1a_2\cdots a_m)^{1/m}+(b_1b_2\cdots b_m)^{1/m}\leq$$

$$\Big[(a_1+b_1)(a_2+b_2)\cdots(a_m+b_m)\Big]^{1/m}.$$

Proof: By induction on m. In the case $m=1$, there is nothing to prove. For $m>1$ we have, using Hölder's inequality,

$$(a_1a_2\cdots a_m)^{1/m}+(b_1b_2\cdots b_m)^{1/m}\leq$$

$$(a_1+b_1)^{1/m}\cdot\left[(a_2\cdots a_m)^{1/(m-1)}+(b_2\cdots b_m)^{1/(m-1)}\right]^{(m-1)/m}.$$

Combining this with the induction hypothesis, the claim follows. □

Hence, applying this claim repeatedly, we conclude that the size of our BSP is at most

$$(2d-4)\left(\frac{n_1}{t}\right)^{1/(d-2)}\cdot(n_2\cdots n_{d-1})^{1/(d-2)}+2(n_1+n_2+\cdots+n_{d-1})+2tn_d=$$

$$(2d-4)\left(\frac{n_1n_2\cdots n_{d-1}}{t}\right)^{1/(d-2)}+2(n_1+n_2+\cdots+n_{d-1})+2tn_d.$$

By the choice of t, this becomes at most

$$(2d-4+2)\,(n_1n_2\cdots n_d)^{1/(d-1)}+2(n_1+n_2+\cdots+n_d).$$

This establishes the induction step and thus completes the proof of the theorem. □

Remark: Theorem 4.1 given below subsumes in general Theorems 3.1 and 3.2. We have considered separately these theorems because they also apply to situations where the sizes of the sets X_i are unbalanced and because their more careful analysis leads to smaller constants of proportionality.

Rectangles in Three Dimensions. Let R be a set of n pairwise-disjoint axis-parallel rectangles in 3-space, and let E denote the set of their edges. Write $E=X\cup Y\cup Z$, as above, and put $x=|X|$, $y=|Y|$, $z=|Z|$, so that $x+y+z=4n$. We establish the following theorem; the upper bound $O(n^{3/2})$ was also obtained in [10].

Theorem 3.3 *Let R be a collection of n axis-parallel rectangles in 3-space, having a total of x edges parallel to the x-axis, y edges parallel to the y-axis and z edges parallel to the z-axis. Then R admits a BSP of size*

$$O(n(\min\{x,y,z\})^{1/2}+n)=O(n^{3/2}).$$

Proof: We construct a binary space partition of R in the following manner.

(i) If one of x, y, z is zero, say $z=0$, then all rectangles are horizontal, and we can obtain a linear-size BSP as above.

(ii) Suppose next that each of x, y, z is at least 1 and that $z \le x, y$. Then we have $z \le 4n/3$ and $x+y \ge 8n/3$. Put

$$t=\left\lceil \frac{x+y}{\sqrt{z}} \right\rceil \ge \frac{(8n/3)}{\sqrt{4n/3}}=\frac{4\sqrt{3}}{3}\sqrt{n}.$$

We partition space into t horizontal slabs $\sigma_1, \ldots, \sigma_t$, as above, so that $x_i+y_i \le (x+y)/t$, for each i, where x_i, y_i, z_i are as defined in Section 3. We then have

$$x_iy_i \le \left(\frac{x_i+y_i}{2}\right)^2 \le \frac{(x+y)^2}{4t^2} \le \frac{z}{4}.$$

Fix a slab $\sigma=\sigma_i$, and consider the set R_σ of rectangles that intersect σ. These rectangles are of two kinds: (a) rectangles that have a horizontal edge in the interior of σ; (b) vertical rectangles whose boundary crosses σ only at two vertical segments, implying that they have no horizontal edge inside σ. Note that rectangles of type (b) contribute only to the z_i-count within σ but not to the x_i and y_i-counts. The rectangles of type (a) are either horizontal rectangles that are fully contained in σ or vertical rectangles that 'start' or 'end' (or both) within σ. We refer to the portions within σ of all these rectangles as *black* rectangles. We refer to the rectangles of type (b) as *red.* Their number is at most $\frac{z_i}{2} \le \frac{z}{2}$.

We project σ onto the xy-plane. The projections of the red rectangles are red segments that are pairwise disjoint and are also disjoint from the projection of any black rectangle. Those black projections can be either segments or rectangles, and they can intersect (or overlap) each other.

Let G be the (nonuniform) grid formed in the xy-plane by the horizontal and vertical lines that contain the edges of the projections of the black rectangles. We refer to the atomic rectangles of G as *pixels.* We classify those pixels into *red pixels*, which are those that are intersected by a red segment, and the remaining *black pixels.* Note that there are a total of $O(x_iy_i)$ pixels. The black pixels can be grouped into *black strips*, which are maximal sets of consecutive black pixels within a single column of G.

We now apply the 2-dimensional BSP construction (provided by Theorem 2.1) to the collection of black strips and red segments. We obtain a decomposition of the xy-plane into $O(x_iy_i+z_i)$ rectangular subregions. Moreover, any red segment or black strip is split by the algorithm at most once.

We lift the BSP just constructed in the z-direction, to obtain a partition of the slab σ by vertical planes orthogonal to the x- and the y-axes. Let K be a cell produced by this partitioning. If K projects to a (portion of a) black strip, then it needs further partitioning, which we do as follows. Ignoring black rectangles that overlap the boundary of K (which are not part of the subproblem at K anyway), any other black rectangle that intersects K crosses it from left to right, i.e., neither of its edges that are orthogonal to the x-axis meets K. Project K onto the yz-plane. By the observation just made, the n_K black rectangles that intersect K project to a collection of n_K pairwise-disjoint segments, and we can again apply the 2-dimensional BSP construction within this projection, effectively obtaining a BSP for K that uses only cuts parallel to the x-axis, whose size is $O(n_K)$. We claim that $\sum_K n_K = O(x_i y_i)$. Indeed, a black rectangle that is counted in n_K must have an edge parallel to the x-axis that intersects K. This follows from the fact that any black rectangle that violates this property must be horizontal and its xy-projection must cover that of K completely. However, K is delimited from above and from below (in the y-direction) by red pixels, which no horizontal black rectangle can cross. This contradiction establishes the asserted property. Now an (x-parallel) edge of a black rectangle can cross at most x_i black regions, and since we have only y_i such edges, we conclude that $\sum_K n_K = O(x_i y_i)$.

We have thus constructed a BSP of size $O(x_i y_i + z_i) = O(z)$ for each of the t slabs σ_i, thus obtaining an overall BSP of size $O(zt) = O((x+y)z^{1/2})$. This completes the proof of the theorem. □

4 Arbitrary Hyperrectangles in Higher Dimensions

Let $\mathcal{R}$ be a set of n axis-parallel k-dimensional hyperrectangles (k-rectangles, or just rectangles, in short) in $\mathbb{R}^d$. We assume that $k < d/2$ and, for simplicity, that the k-rectangles are in *general position*, as above. We note that this assumption implies that no pair of rectangles intersect.

Each rectangle $r \in \mathcal{R}$ has k *extent coordinates*, i.e., coordinates x_i for which the projection of r onto the x_i-axis is an interval with nonempty interior, and $d-k$ *fixed coordinates* (those for which this projection is a singleton point).

Let K be an axis-parallel box in $\mathbb{R}^d$. Let r be a rectangle in $\mathcal{R}$ and put $r' = r \cap K$. We say that r is an *x_i-pass-through* in K if the projection of r' on the x_i-axis is equal to the projection of K on the same axis. We denote by $\mathbf{pt}(r, K)$ the tuple of coordinates for which r is a pass-through in K. The main result of this section is

Theorem 4.1 *(a) A set $\mathcal{R}$ of n axis-parallel k-rectangles in d-space, as above, admits a BSP of size $O(n^{d/(d-k)})$. (b) There exist sets $\mathcal{R}$ of n axis-parallel k-rectangles in $\mathbb{R}^d$, as above, for any n, d and $k < d/2$, so that any (rectilinear) binary space auto-partition for $\mathcal{R}$ has size $\Omega(n^{d/(d-k)})$.*

Proof of the upper bound: Let $\mathcal{R}$ be a set of n k-rectangles in d-space, satisfying the above properties. Put $t = cn^{1/(d-k)}$, for some absolute constant $c > 1$. The BSP construction proceeds through d phases, where in the j-th phase we cut each of the cells produced in the preceding phases by hyperplanes orthogonal to the x_j-axis. Each cut that we perform is at some fixed coordinate of some rectangle in $\mathcal{R}$.

In the first phase we slice $\mathbb{R}^d$ by a sequence of $t-1$ hyperplanes orthogonal to the x_1-axis, partitioning space into t slabs, so that each slab σ contains at most n/t rectangles that are either orthogonal to x_1 or have an extent in the x_1-coordinate but are not x_1-pass-throughs in σ.

Suppose we are in the j-th phase. Let σ be a cell (subslab) produced by the previous phases. We assume inductively that, for each subset M of $\{1,\ldots,j-1\}$ of size $|M| \le k$, σ contains at most $n/t^{j-1-|M|}$ rectangles that are pass-throughs in exactly the coordinates in M. (We note that this property holds for $j=2$.)

We cut σ by $O(t)$ cuts orthogonal to the x_j-axis, to ensure that, for each subset M as above, any resulting subslab σ' contains at most $n/t^{j-|M|}$ rectangles that were pass-throughs in σ in exactly the coordinates in M and are not x_j-pass-throughs in σ'. (By the induction hypothesis, σ contains at most $n/t^{j-1-|M|}$ such rectangles, so it is easy to cut this number down by a factor of t for each σ'.) In addition, one also has the property that for each subset M of $\{1,\ldots,j\}$ of size $\le k$ that contains j, σ' contains at most $n/t^{j-|M|}$ rectangles that are pass-throughs in exactly the coordinates in M. (These bounds are simply 'carried over' from the inductively assumed bounds for σ and $j-1$.)

This establishes the inductive property for j, and thus allows us to continue in this manner until all d phases are executed.

Let us analyze the performance of this partitioning scheme. We claim that, for $c>1$, none of the final cells can contain any rectangle in their interior. Indeed, let σ be a final cell. By the above property, for each set M of coordinates of size $|M| \le k$, there are at most

$$\frac{n}{t^{d-|M|}} = \frac{1}{c^{d-|M|} n^{\frac{d-|M|}{d-k}-1}} = \frac{1}{c^{d-|M|} n^{(k-|M|)/(d-k)}} \le \frac{1}{c^{d-k}}$$

(portions of) rectangles contained in σ that are pass-throughs in σ in exactly the coordinates in M. By choosing $c>1$ we are guaranteed that the interior of σ does not meet any rectangle of $\mathcal{R}$.

Hence the resulting subdivision is indeed a BSP for $\mathcal{R}$. The number of cells of this BSP is clearly $O(t^d) = O(n^{d/(d-k)})$, with the constant of proportionality depending (exponentially) on d and k. Further, any one rectangle is cut into at most $t^k = n^{k/(d-k)}$ pieces, implying a bound of $O(n^{d/(d-k)})$ on the number of fragments. This completes the proof of the upper bound.

Proof of the lower bound: Put $I = [0, n^{1/(d-k)}+1]$, and let K be the cube I^d. Put $L = \binom{d}{k}$. For each k-tuple τ of coordinates, we construct n k-rectangles whose extent coordinates are those of τ, as follows. Put $E = \{1,2,\ldots,n^{1/(d-k)}\}^{d-k}$. For each $\mathbf{a} \in E$ construct a rectangle $r = r(\mathbf{a})$ whose i-th fixed coordinate is $\mathbf{a}_i + \varepsilon(\mathbf{a},\tau)$, and whose projection on each of the extent coordinates (i.e., those in τ) is I. Here $\varepsilon(\mathbf{a},\tau)$ is a number in $(0,1)$, so that different pairs $(\mathbf{a},\tau)$ are assigned different numbers. We thus obtain a collection $\mathcal{R}$ of a total of Ln rectangles.

We claim that any axis-parallel BSP for $\mathcal{R}$ must consist of $\Omega(n^{d/(d-k)})$ cells. Consider the integer grid $G = \{1,2,\ldots,n^{1/(d-k)}\}^d$. With each $\mathbf{g} \in G$ associate the grid cell $Q(\mathbf{g}) = \prod_{i=1}^d (\mathbf{g}_i, \mathbf{g}_i+1)$. A grid cell $Q = Q(\mathbf{g})$ is crossed by exactly L rectangles of $\mathcal{R}$: For each k-tuple of coordinates there is exactly one rectangle of $\mathcal{R}$ whose extent coordinates are those in the tuple, that crosses Q—it is the rectangle whose fixed coordinates agree with the corresponding elements of $\mathbf{g}$.

Since $L > 1$, any BSP for $\mathcal{R}$ must have cuts that cross Q, for otherwise Q will be contained in a single cell of the BSP, which is impossible since that cell is now crossed by more than one rectangle of $\mathcal{R}$. Moreover, we argue that Q must be cut by (portions of) hyperplanes in at least k different orientations (we omit the easy proof). Halt the BSP construction right after Q is cut for the first time by hyperplanes in k different orientations; suppose that they are the first k coordinates. Let r be the unique rectangle of $\mathcal{R}$ that crosses Q and whose extent coordinates are the first k coordinates. For each $i = 1, \ldots, k$ let h_i be a hyperplane orthogonal to x_i that has already crossed Q. Then $r \cap \left(\bigcap_{i=1}^{k} h_i\right)$ is a singleton point v—a vertex of a portion of r that the BSP has just formed. We assign Q to this portion of r, or, more precisely, to the final portion of r that will be formed by the BSP and will have v as a vertex. No such portion can be charged by more than 2^k grid cells, which implies that the number of fragments of rectangles in $\mathcal{R}$ that the BSP has to form is at least proportional to the number of grid cells, so the BSP has size $\Omega(n^{d/(d-k)})$, as asserted. $\square$

Remark: The upper bound of Theorem 4.1 hold also for $d/2 \leq k \leq d-1$, even if the rectangles in $\mathcal{R}$ are allowed to intersect. The lower bound applies also for $d/2 \leq k \leq d-1$, but the construction uses rectangles that intersect. The proofs are essentially identical.

5 Disjoint 2-Rectangles in $\mathbb{R}^4$

Let $\mathcal{R}$ be a set of n axis-parallel *pairwise disjoint* 2-rectangles in $\mathbb{R}^4$. This is the simplest instance not covered by Theorem 4.1.

Theorem 5.1 *(a) A set $\mathcal{R}$ of n axis-parallel pairwise-disjoint 2-rectangles in $\mathbb{R}^4$ admits a BSP of size $O(n^{5/3})$. (b) There exist collections of n (axis-parallel) pairwise-disjoint 2-rectangles in $\mathbb{R}^4$ that only admit (rectilinear) BSP's of size $\Omega(n^{5/3})$.*

Proof of the upper bound: Let K be an axis-parallel box in $\mathbb{R}^4$. Apply the same round-robin construction given in the preceding section, but with $t = c'n^{1/6}$ for an appropriate constant c'. We obtain $O(n^{2/3})$ subcells, so that each subcell σ contains at most $n/t^2 = O(n^{2/3})$ rectangles that are pass-throughs in two coordinates, at most $n/t^3 = O(n^{1/2})$ rectangles that are pass-throughs in exactly one coordinate, and at most $n/t^4 = O(n^{1/3})$ rectangles that are not pass-throughs in any coordinate.

Lemma 5.2 *A subcell σ cannot contain two rectangles r, r' such that r is pass-through in σ in two coordinates and r' is pass-through in σ in the two* complementary *coordinates.*

Proof: Any two such rectangles must intersect, contrary to assumption. $\square$

By Lemma 5.2, it is easily verified that there are only two possible maximal values for the set

$$\mathbf{pt}(\sigma) \equiv \{\mathbf{pt}(r, \sigma) \mid r \text{ is a rectangle that is pass-through in } \sigma \text{ in 2 coordinates}\},$$

up to a permutation of the coordinates; namely:

$$\text{(i)}\ \ \mathbf{pt}(\sigma) \subseteq \{(1,2),(1,3),(2,3)\}$$

$$\text{(ii)}\ \ \mathbf{pt}(\sigma) \subseteq \{(1,4),(2,4),(3,4)\}.$$

Case (i). Consider first case (i). Note that in this case all the rectangles that are pass-throughs in σ in 2 coordinates are orthogonal to the x_4-axis and lie at different heights. We cut σ by $O(n^{1/6})$ cuts orthogonal to the x_4-axis so that each of the $O(n^{5/6})$ subcells contains at most $n^{1/2}$ (portions of) rectangles. In particular, each subcell contains at most $n^{1/2}$ rectangles that are pass-throughs in 2 coordinates, at most $n^{1/3}$ rectangles that are pass-throughs in 1 coordinate, and at most $n^{1/6}$ rectangles that are not pass-throughs in any coordinate. By Lemma 5.2, the extent coordinates of the 2-coordinate-pass-throughs in a subcell may again fall into either case (i) or case (ii) (with a possible new permutation of the coordinates). For a subcase-(i) subcell σ, we cut σ by $O(n^{1/6})$ cuts orthogonal to the x_4 axis, resulting in a total of $O(n)$ subsubcells each containing at most $n^{1/3}$ (portions of) rectangles. For a subcase-(ii) subcell σ, we apply the scheme described below, based on a round-robin cutting of each subcell into $n^{3/6}$ pieces each containing at most $n^{1/6}$ (portions of) rectangles. If we denote by $F(n)$ the maximum size of a BSP that the algorithm constructs for any input set of n pairwise-disjoint axis-parallel 2-rectangles in $\mathbb{R}^4$, then the overall number of final cells produced for cells σ that belong to case (i) is $O(n)F(n^{1/3}) + O(n^{4/3})F(n^{1/6})$.
Case (ii). Next consider case (ii). We execute a round-robin procedure that only makes cuts orthogonal to the x_1, x_2, and x_3-axes. At each stage of this procedure we make $t = O(n^{1/6})$ cuts. This partitions each of the preceding $O(n^{2/3})$ cells into $O(n^{1/2})$ subcells, for a total of $O(n^{7/6})$ subcells. This can be done so that each subcell σ contains

- at most $n^{2/3}/t^2 = n^{1/3}$ rectangles that are pass-throughs in σ in two coordinates, one of which is x_4,
- at most $n^{1/2}/t = n^{1/3}$ rectangles that are pass-throughs in σ in two coordinates, none of which is x_4,
- at most $n^{1/2}/t^2 = n^{1/6}$ rectangles that are pass-throughs in σ in exactly one coordinate,
- and no other rectangles.

Note that, because of Lemma 5.2, the existence of 2-coordinate-pass-throughs of the second category annihilates those pass-throughs of the first category that have complementary extent coordinates. Consequently, the extent coordinates of the 2-coordinate-pass-throughs in σ may again fall into either case (i) or case (ii) (with a possible new permutation of the coordinates).

In subcase (i) we proceed in a manner similar to the one above, cutting σ by $O(n^{1/6})$ cuts orthogonal to the x_4-axis, to obtain $O(n^{1/6})$ subcells, each containing at most $n^{1/6}$ rectangles, for an overall recursive bound of the form $O(n^{4/3}) \cdot F(n^{1/6})$.

In subcase (ii), with, say, $\mathbf{pt}(\sigma) = \{(1,4),(2,4),(3,4)\}$, we again proceed as above, cutting σ in a round-robin fashion by cuts orthogonal to the x_1, x_2, and x_3-axes, making

$t = O(n^{1/6})$ cuts in each round. This can be done so as to eliminate the at most $n^{1/6}$ rectangles that are pass-throughs in one coordinate, as well as all of the at most $n^{1/3}$ 2-coordinate-pass-throughs. Hence this step produces a BSP for σ, whose size is $O(n^{1/2})$, for a total of $O(n^{7/6} \cdot n^{1/2}) = O(n^{5/3})$.

Putting everything together, we obtain the following recurrence for $F(n)$:

$$F(n) = O(n^{5/3}) + O(n) \cdot F(n^{1/3}) + O(n^{4/3}) \cdot F(n^{1/6}),$$

whose solution is easily seen to be

$$F(n) = O(n^{5/3}).$$

Proof of the lower bound: Let G be the $n^{2/3} \times n^{1/3} \times n^{1/3} \times n^{1/3}$ integer grid in 4-space. Let I denote the interval $[0, n^{1/3} + 1]$ and let I' denote the interval $[0, n^{2/3} + 1]$. We construct the following four families of rectangles:

$$\begin{aligned}
\mathcal{R}_1 &= \{\{i + \varepsilon_{i,j}^{(1)}\} \times \{j + \varepsilon_{i,j}^{(1)}\} \times I \times I \mid \\
&\qquad\qquad i = 1, \ldots, n^{2/3},\ j = 1, \ldots, n^{1/3}\} \\
\mathcal{R}_2 &= \{\{i + \varepsilon_{i,j}^{(2)}\} \times I \times \{j + \varepsilon_{i,j}^{(2)}\} \times I \mid \\
&\qquad\qquad i = 1, \ldots, n^{2/3},\ j = 1, \ldots, n^{1/3}\} \\
\mathcal{R}_3 &= \{\{i + \varepsilon_{i,j}^{(3)}\} \times I \times I \times \{j + \varepsilon_{i,j}^{(3)}\} \mid \\
&\qquad\qquad i = 1, \ldots, n^{2/3},\ j = 1, \ldots, n^{1/3}\} \\
\mathcal{R}_4 &= \{I' \times [i + \varepsilon_{i,j,k}^{(4)}, i + \varepsilon_{i,j,k}^{(5)}] \times [j + \varepsilon_{i,j,k}^{(4)}, j + \varepsilon_{i,j,k}^{(5)}] \\
&\qquad \times\ [k + \varepsilon_{i,j,k}^{(4)}, k + \varepsilon_{i,j,k}^{(5)}] \mid i, j, k = 1, \ldots, n^{1/3}\},
\end{aligned}$$

where the $\varepsilon_{i,j}^{(m)}$'s and $\varepsilon_{i,j,k}^{(m)}$'s are all distinct small positive real numbers (say, at most 0.1). In addition, we require that all the intervals $[\varepsilon_{i,j,k}^{(4)}, \varepsilon_{i,j,k}^{(5)}]$ do not contain any of the numbers $\varepsilon_{i,j}^{(m)}$, for $m = 1, 2, 3$. The elements of $\mathcal{R}_4$ are (long and skinny) 4-dimensional boxes rather than rectangles, but we can replace each of them by its 24 bounding (2-dimensional) rectangles, slightly shifted away from each other, to maintain the general position property. Note that the rectangles in $\mathcal{R} = \bigcup_{i=1}^{4} \mathcal{R}_i$ are pairwise disjoint.

Fix any $1 \le i \le n^{2/3}$, $1 \le j, k, \ell \le n^{1/3}$, and consider the cube

$$\Xi(i, j, k, \ell) = [i, i + 0.1] \times [j, j + 0.1] \times [k, k + 0.1] \times [\ell, \ell + 0.1].$$

We refer to $\Xi(i, j, k, \ell)$ as the *junction* at (i, j, k, ℓ).

Consider any rectilinear BSP for $\mathcal{R}$, namely, one that only uses cuts orthogonal to the coordinate axes. It is clear that each junction $\Xi = \Xi(i, j, k, \ell)$ must be cut by at least one hyperplane of the BSP, or else the BSP will have a final cell that is intersected by more than one rectangle of $\mathcal{R}$.

If the first hyperplane that intersects Ξ is orthogonal to the x_1-axis, then it intersects the unique skinny box of $\mathcal{R}_4$ that crosses Ξ. Such an intersection is a tiny 3-rectangle that lies inside Ξ. Hence the number of fragments of the boxes in $\mathcal{R}_4$ produced by the BSP is at least equal to the number of junctions Ξ with this property.

Consider then a junction $\Xi = \Xi(i, j, k, \ell)$ that is not crossed (for the first time) by any hyperplane orthogonal to the x_1-axis. We make the following claim:

Claim 5.3 *Either the box of $\mathcal{R}_4$ that crosses Ξ is eventually cut as in the preceding paragraph, or else there exists a rectangle in $\mathcal{R}_1 \cup \mathcal{R}_2 \cup \mathcal{R}_3$ that is split by the BSP into subrectangles, so that at least one of them has a vertex inside Ξ.*

Intuitively, the only way to "get rid" of the box of $\mathcal{R}_4$ that cuts Ξ by cuts parallel to x_1 is to cut along each of its 3-D facets; moreover, one cannot get rid of the other three rectangles in $\mathcal{R}_1 \cup \mathcal{R}_2 \cup \mathcal{R}_3$ that cross Ξ without making at least two differently-oriented cuts, each being orthogonal to one of the axes x_2, x_3, x_4.

Proof of the Claim: Suppose, with no loss of generality, that the first hyperplane h of the BSP that cuts Ξ is orthogonal to the x_2-axis. Then h splits each of the two rectangles of $\mathcal{R}_2 \cup \mathcal{R}_3$ that cross Ξ into two subrectangles, and h may fully contain only the rectangle of $\mathcal{R}_1$ that crosses Ξ (if at all). The box of $\mathcal{R}_4$ that crosses Ξ meets at least one of the two subcells into which Ξ is split. It follows that each of the two pieces into which Ξ has been split by h, which is met by all three elements of $\mathcal{R}_2 \cup \mathcal{R}_3 \cup \mathcal{R}_4$ that cross Ξ (and at least one of these two sub-junctions has this property), must be further cut at least once more.

Suppose first that the next cut of such a sub-junction Ξ' is by a hyperplane h' orthogonal to another axis. If h' is orthogonal to the x_1-axis then we can charge Ξ to the cut of the relevant box of $\mathcal{R}_4$, as above. Otherwise, suppose h' is orthogonal to the x_3-axis. Then the rectangle of $\mathcal{R}_3$ that crosses Ξ is cut by h and h' into pieces that have at least one vertex in (the closure of) Ξ'. A similar property holds for the rectangle of $\mathcal{R}_2$ if h' is orthogonal to the x_4-axis.

The only remaining case is when h' is also orthogonal to the x_2-axis. In fact, in general, Ξ may be cut by several such hyperplanes. However, all of these portions are intersected by each of the two rectangles of $\mathcal{R}_2 \cup \mathcal{R}_3$ that meet Ξ, and at least one portion is also crossed by the box of $\mathcal{R}_4$ that meets Ξ.

Let Ξ'' denote such a portion of Ξ, which we assume not to be cut any more by hyperplanes orthogonal to the x_2-axis. But Ξ'' does have to be cut again, because it is still met by more than one rectangle, and any cut in any other direction can be charged uniquely to Ξ, using the arguments in the preceding paragraphs. This completes the proof of the claim., $\Box$

This claimed property implies that the number of fragments of the rectangles in $\mathcal{R}_1 \cup \mathcal{R}_2 \cup \mathcal{R}_3$ that is produced by the BSP is at least equal to the number of junctions Ξ. Since this number is $n^{5/3}$, it follows that this is a lower bound on the size of any (rectilinear) BSP for $\mathcal{R}$. $\Box$

6 Conclusion

In conclusion, we refer again to our summary of results in Table 1. Perhaps the most interesting remaining open problem is to generalize our techniques for (disjoint) 2-rectangles in $\mathbb{R}^4$ to higher dimensions, with the hope of obtaining asymptotically tight bounds for k-rectangles in $\mathbb{R}^d$, for all k and d.

7 Acknowledgments

We thank Shakhar Smorodinsky and Saurabh Sethia for several useful discussions and valuable input on this work.

References

[1] F. d'Amore and P. G. Franciosa, On the optimal binary plane partition for sets of isothetic rectangles, *Inform. Process. Lett.* 44 (1992), 255–259.

[2] M. de Berg, M. de Groot and M. Overmars, New results on binary space partitions in the plane, *Comput. Geom. Theory Appl.* 8 (1997), 317–333.

[3] M. de Berg, M. de Groot and M. Overmars, Perfect binary space partitions, *Comput. Geom. Theory Appl.* 7 (1997), 81–91.

[4] P. Berman, B. DasGupta and S. Muthukrishnan, On the exact size of the binary space partitioning of sets of isothetic rectangles with applications, DIMACS Technical Report 2000-06, July 2000.

[5] P. Berman, B. DasGupta, S. Muthukrishnan and S. Ramaswami, Improved approximation algorithms for rectangle tiling and packing, *Proc. 12th ACM-SIAM Symp. on Disc. Alg.*, (2001), 427–436.

[6] S. Chin and S. Feiner, Near real-time shadow generation using BSP trees, *Comput. Graph.* 23 (1989), 99–106.

[7] S.E. Dorwald, A survey of object-space hidden surface removal, *Intern. J. Comput. Geom. Appl.* 4 (1994), 325–362.

[8] H. Fuchs, Z. M. Kedem, and B. Naylor, On visible surface generation by a priori tree structures, *Comput. Graph.* 14 (1980), 124–133.

[9] M.S. Paterson and F.F. Yao, Efficient binary space partitions for hidden-surface removal and solid modeling, *Discrete Comput. Geom.* 5 (1990), 485–503.

[10] M.S. Paterson and F.F. Yao, Optimal binary space partitions for orthogonal objects, *J. Algorithms* 13 (1992), 99–113.

[11] C. D. Tóth, A note on binary plane partitions, *Proc. 17th ACM Sympos. Comput. Geometry*, 2001.

[12] C. D. Tóth, BSP for line segments with a limited number of directions, Manuscript, ETH Zürich, March, 2001.

[13] W.C. Thibault and B.F. Naylor, Set operations on polyhedra using binary space partitioning trees, *Comput. Graph.* 21(4) (1987), 153–162 (*Proc. SIGGRAPH '87*).

[14] G. Vaněček Jr., Brep-Index: A multidimensional space partitioning tree, *Internat. J. Comput. Geom. Appl.* 1(3) (1991), 243–261.