Binary Space Partitions for Axis-Parallel Line Segments: Size-Height Tradeoffs

Sunil Arya
Department of Computer Science
The Hong Kong University of Science and Technology
Clear Water Bay, Kowloon, Hong Kong.
arya@cs.ust.hk

Abstract

We present worst-case lower bounds on the minimum size of a binary space partition (BSP) tree as a function of its height, for a set $\mathcal{S}$ of n axis-parallel line segments in the plane. We assume that the BSP uses only axis-parallel cutting lines. These lower bounds imply that, in the worst case, a BSP tree of height $O(\log n)$ must have size $\Omega(n \log n)$ and a BSP tree of size $O(n)$ must have height $\Omega(n^\delta)$, where δ is a suitable constant.
Key words: Computational geometry, binary space partitions, line segments.

1 Introduction

Given a set $\mathcal{S}$ of n interior-disjoint objects in $\mathbb{R}^d$, the *binary space partition (BSP)* for $\mathcal{S}$ is a subdivision obtained by recursively cutting space into two parts by a hyperplane, terminating when each cell is intersected by at most a constant number of the given objects. BSPs can be represented in a natural way by binary trees, and they have numerous applications including hidden surface removal [5], ray tracing [7], visibility problems [11], solid geometry [12], and spatial databases [10]. The efficiency of BSP-based algorithms depends on the size and the height of the tree. For example, the time taken to answer point location and ray tracing queries depends on the height of the tree, while the space requirement depends on the size of the tree.

BSPs have been extensively studied for objects in two, three, and higher dimensions. Here we restrict our attention to BSPs for objects in the plane. Paterson and Yao [8] showed that any set of n line segments admits a BSP of $O(n \log n)$ size and $O(\log n)$ height. They conjectured that it is possible to achieve a size bound of $O(n)$; however, recently, Tóth [13] disproved this conjecture by showing that, in the worst case, $\Omega(n \log n / \log \log n)$ size is necessary.

When the line segments are axis-parallel, Paterson and Yao [9] showed that there exists a BSP of $3n$ size. This bound was improved to $2n$ by d'Amore and Franciosa [2]; a nearly matching lower bound was presented by Dumitrescu et al. [4]. Besides axis-parallel line segments, linear size BSPs are known for three other special cases: *fat* objects, *homothetic* objects, and line segments where the ratio between the lengths of the longest and shortest segment is bounded by a constant [3].

In this paper we consider the following natural question: Is it possible to construct a BSP of $O(n)$ size and $O(\log n)$ height for any set $\mathcal{S}$ of n axis-parallel segments? Not only is this question interesting for its theoretical value, this subproblem also arises naturally when constructing balanced BSPs for axis-parallel rectangles in three dimensions.[1] We give a negative answer to this question, under the assumption that the BSP is only allowed to use axis-parallel cutting lines. Henceforth, we call such a BSP an *orthogonal* BSP. While the restriction to orthogonal BSPs is a limitation of our analysis, we note that all the BSP constructions in the literature for axis-parallel segments satisfy this model and, thus, our analysis applies to them.

More precisely, we present a worst-case lower bound of $\Omega(n \log n / \log h)$ on the minimum size of an orthogonal BSP that has height at most h. Thus, in the worst case, a linear size orthogonal BSP must be highly unbalanced (that is, have height $\Omega(n^{\delta})$, where δ is a suitable constant). On the other hand, an orthogonal BSP of height $O(\log n)$ must have size at least $\Omega(n \log n / \log \log n)$. We also present a stronger lower bound (detailed in Theorem 1(ii)) that enables us to improve the size bound in the last claim to $\Omega(n \log n)$.

2 Preliminaries

Let $\mathcal{S}$ be a set of n interior-disjoint line segments in the plane. A BSP for $\mathcal{S}$ is a tree defined as follows. The root of the tree is associated with $\mathbb{R}^2$. Inductively, suppose that a node v is associated with a convex polygon P_v. If the set $\mathcal{S}_v = \{s \cap P_v \,:\, s \in \mathcal{S}\}$ is empty, then v is a leaf. Otherwise we partition P_v into two polygons by a *cutting line* ℓ_v. We store with v the segments in $\mathcal{S}_v$ (if any) contained within ℓ_v. The left and right child of v are associated with the polygons $P_v \cap \ell_v^+$ and $P_v \cap \ell_v^-$, respectively, where ℓ_v^+ and ℓ_v^- denote the positive and negative open halfplanes bounded by ℓ_v. The *size* of the BSP is the total number of segments stored at all the internal nodes, and its *height* is the length of the longest path from the root to a leaf.

3 Lower Bounds

Our lower bound example consists of a set $\mathcal{S}$ of $2n$ segments arranged as shown in Fig. 1. We refer to this configuration as a *stair-case* of size n. Let $S_h(n)$ denote the minimum size of an orthogonal BSP for a stair-case of size n, subject to the condition that its height is at most h.

The following lemma is useful in deriving a recurrence inequality for $S_h(n)$.

Lemma 1 *Let $\mathcal{S}$ be a set of n interior-disjoint axis-parallel line segments, which is partitioned into two sets $\mathcal{S}_1$ and $\mathcal{S}_2$. Let h be a positive integer. Let s, s_1, and s_2 be the minimum size of an orthogonal BSP of height at most h, for $\mathcal{S}, \mathcal{S}_1$, and $\mathcal{S}_2$, respectively. Then $s \geq s_1 + s_2$. (If it is not possible to construct an orthogonal BSP of height at most h for $\mathcal{S}$, then $s = \infty$ and we regard this inequality as true.)*

Proof Let T denote a minimum-size orthogonal BSP of height at most h for $\mathcal{S}$. Let s'_1 and s'_2 be the number of segments stored at the internal nodes of T that are generated from $\mathcal{S}_1$ and $\mathcal{S}_2$, respectively. Clearly $s = s'_1 + s'_2$. Further, T is obviously an orthogonal BSP of height at most h

[1] Specifically, if a BSP of $O(n)$ size and $O(\log n)$ height could be constructed for axis-parallel segments in the plane, then it would follow from known results [4, 9] that we could construct a BSP of optimal size (that is, $O(n^{3/2})$ size) and $O(\log n)$ height for axis-parallel rectangles in three dimensions.

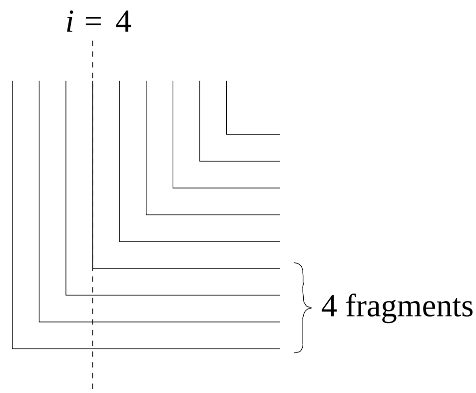

Figure 1: A stair-case of size 9.

for $\mathcal{S}_1$, after eliminating the segments stored at the internal nodes of T that are generated from $\mathcal{S}_2$. It follows that $s_1 \leq s_1'$. Similarly $s_2 \leq s_2'$. Thus $s \geq s_1 + s_2$. □

Let ℓ denote the cutting line associated with the root of a minimum-size orthogonal BSP of height at most h for $\mathcal{S}$. Without loss of generality, assume that ℓ is vertical. Assign labels $1, 2, \ldots, n$ to the vertical segments of the stair-case, from left to right. There are two possible cases: (i) ℓ passes through the ith vertical segment, $1 \leq i \leq n$, and (ii) ℓ lies between the ith and $(i+1)$th vertical segment, $1 \leq i \leq n-1$.

For case (i), we obtain a stair-case of size $i-1$ to the left of ℓ; the size of the corresponding BSP is thus $S_{h-1}(i-1)$. See Fig. 1. To the right of ℓ we obtain i fragments (from the i lowest horizontal segments) and a stair-case of size $n-i$; it follows from Lemma 1 that the size of the corresponding BSP is at least $S_{h-1}(n-i)+i$. Thus, for case (i), we have

$$S_h(n) \geq \min_{1 \leq i \leq n} S_{h-1}(i-1) + S_{h-1}(n-i) + i + 1, \tag{1}$$

where 1 has been added to account for the fragment contained in ℓ.

Similarly, for case (ii), we can write

$$S_h(n) \geq \min_{1 \leq i \leq n-1} S_{h-1}(i) + S_{h-1}(n-i) + i \geq \min_{1 \leq i \leq n-1} S_{h-1}(i-1) + S_{h-1}(n-i) + i + 2. \tag{2}$$

In the last inequality we have used the fact that $S_{h-1}(i) \geq S_{h-1}(i-1)+2$, which follows immediately from Lemma 1. Combining the two cases, it is easy to see that Eq. (1) dominates. Moreover, we can restrict the minimization variable i to the range $1 \leq i \leq \lceil n/2 \rceil$. Including the boundary cases, we have

$$\begin{array}{rcll} S_h(0) & = & 0, & \text{for } h \geq 0, \\ S_0(n) & = & \infty, & \text{for } n \geq 1, \\ S_h(n) & \geq & \min_{1 \leq i \leq \lceil n/2 \rceil} S_{h-1}(i-1) + S_{h-1}(n-i) + i + 1, & \text{for } n \geq 1, h \geq 1. \end{array} \tag{3}$$

We define $\overline{S}_h(n)$ using the same recurrence relation as given here for $S_h(n)$, except with the inequality in the third equation replaced by equality. Clearly a lower bound for $\overline{S}_h(n)$ also applies to $S_h(n)$. To compute a lower bound for $\overline{S}_h(n)$, we find it convenient to expand the recurrence relation in the form of a tree as follows. The root of the tree corresponds to the term $\overline{S}_h(n)$ and is *labelled* n. Suppose that at the top level of the recursion, we assign a value k to the minimization variable i. We represent this by two nodes labelled $k-1$ and $n-k$ corresponding to the terms $\overline{S}_{h-1}(k-1)$ and $\overline{S}_{h-1}(n-k)$, respectively. The node labelled $k-1$ is made the left child of the root, and the node labelled $n-k$ is made the right child (note that the label assigned to the left child is never more than the label assigned to the right child). We continue recursively in this manner to expand the terms generated, leading to an equivalent recursion tree. This tree has height at most h and its leaves correspond to the boundary cases. We can get many different trees depending on the value assigned to the minimization variable for each term generated. Let $\mathcal{T}_h(n)$ denote the set of all the labelled trees that this process can generate, starting with the term $\overline{S}_h(n)$ associated with the root of the tree. (Fig. 2 shows one possible tree for $n = 11$ and $h = 5$.)

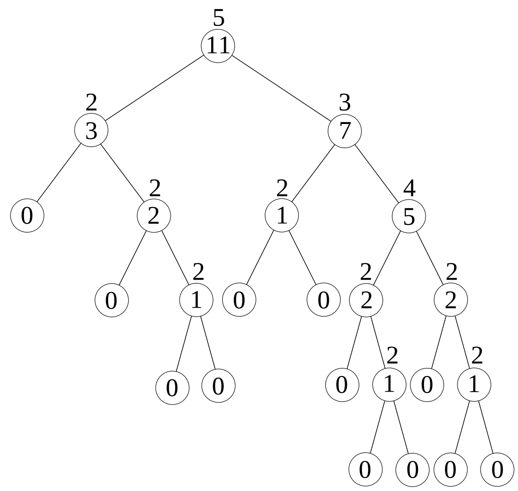

Figure 2: A labelled tree T for $n = 11$ and $h = 5$. The labels of the nodes are shown inside the circles. For internal nodes, their cost is shown above the circle.

Let $T \in \mathcal{T}_h(n)$ be any labelled tree. Based on the value assigned to the minimization variable at the internal nodes, we can deduce the value of the term $\overline{S}_h(n)$ associated with its root. We formalize this by defining the notion of the *cost* of tree T, denoted *cost*(T), as follows. Define the cost of an internal node to be $k+1$, where k is the value of the minimization variable associated with it. If any leaf in T has a non-zero label, then define the cost of T to be ∞; otherwise, define the cost of T to be the total cost of its internal nodes. (For example, the cost of the tree in Fig. 2 is 28.) It is clear from Eq. (3) that the cost of a tree represents the value of the term $\overline{S}_h(n)$ associated with its root. It is now obvious that

$$\overline{S}_h(n) = \min_{T \in \mathcal{T}_h(n)} \mathit{cost}(T). \tag{4}$$

We compute a lower bound on the cost of T. To this end, we assign an *address* to each node in T based on the path taken to reach it, starting from the root of T. For example, if we take a right, right, left, and a right branch to reach a node, its address is *RRLR*. We define the *L-depth* of a node to be the number of L's in its address, and define the *L-path-length* of a tree T, denoted *L-path-length*(T), to be the sum of the L-depth of all the leaves in T.

In Lemma 2, we show that *L-path-length*(T) is a lower bound on *cost*(T) and, in Lemma 3, we establish lower bounds on *L-path-length*(T).[2]

Lemma 2 *Let $T \in \mathcal{T}_h(n)$ be a tree with all leaves labelled 0. Then the following are true:*

(i) $int(T) = n$ and $leaves(T) = n+1$, where $int(T)$ and $leaves(T)$ denote the number of internal nodes and leaves, respectively, in T.

(ii) $cost(T) = L\text{-}path\text{-}length(T) + n$.

Proof The proof is by induction on n. If $n = 0$, then T consists of just one node labelled 0. Thus, $int(T) = 0, leaves(T) = 1, cost(T) = 0, L\text{-}path\text{-}length(T) = 0$, so (i) and (ii) hold.

Next suppose that the claim in the lemma holds for all $n < m$, where m is an integer greater than 0. We will show that the claim holds for $n = m$. Let T_L and T_R denote the left and right subtree, respectively, of the root of T. If k is the value assigned to the minimization variable associated with the root of T, then $T_L \in \mathcal{T}_{h-1}(k-1)$ and $T_R \in \mathcal{T}_{h-1}(m-k)$. Also, the leaves in T_L and T_R are all labelled 0. Thus, we can apply the induction hypothesis to T_L and T_R, which implies that $int(T_L) = k-1, leaves(T_L) = k, int(T_R) = m-k$, and $leaves(T_R) = m-k+1$. Thus, $int(T) = int(T_L) + int(T_R) + 1 = m$ and $leaves(T) = leaves(T_L) + leaves(T_R) = m+1$, which proves (i).

To prove (ii), observe that for a leaf in T_L, its L-depth in T is one more than its L-depth in T_L, and for a leaf in T_R, its L-depth in T is the same as its L-depth in T_R. Thus $L\text{-}path\text{-}length(T) = L\text{-}path\text{-}length(T_L) + L\text{-}path\text{-}length(T_R) + leaves(T_L)$. Also, by definition, we have $cost(T) = cost(T_L) + cost(T_R) + k + 1$ and, by the induction hypothesis, we have $cost(T_L) = L\text{-}path\text{-}length(T_L) + k - 1$ and $cost(T_R) = L\text{-}path\text{-}length(T_R) + m - k$. It is now easy to see that $cost(T) = L\text{-}path\text{-}length(T) + m$, which completes the proof. □

Lemma 3 *Let $T \in \mathcal{T}_h(n)$ be a tree with all leaves labelled 0. Then*

(i) $L\text{-}path\text{-}length(T) = \Omega(n \log n / \log h)$.

(ii) $L\text{-}path\text{-}length(T) = \Omega(nr)$, where r is an integer, $0 \le r \le h$, such that $\sum_{d=0}^{r} \binom{h}{d} \le n/4$.

Proof We say that a node in T is an *L-node* if its address ends with an L (that is, it is a left child), or if it is the root of T. If an L-node is a leaf in T, we call it an *L-leaf*.

We construct a new tree T' from T as follows. The nodes of T' correspond in a one-to-one manner to the L-nodes of T. For each L-node in T at L-depth $d, d \ge 0$, the corresponding node in T' is at depth d. Moreover, the tree T' satisfies the following property. Let x be an L-node in T at L-depth d, and let C_x be the set of descendants of x at L-depth $d+1$; then in T' the nodes corresponding to the nodes in C_x are children of the node corresponding to x. (Fig. 3 shows the tree T' for the tree T given in Fig. 2.) Since $|C_x|$ is at most h, the degree of any node in T' is at most h. Also, the height of T' is the same as the maximum L-depth of a node in T, and so it is at most h.

[2]Recently, Theocharis Malamatos has pointed it out to us that this quantity was studied earlier by Callahan and Kosaraju [1], in the context of constructing the well-separated pair decomposition of points. Our analysis in Lemma 3 is similar to theirs.

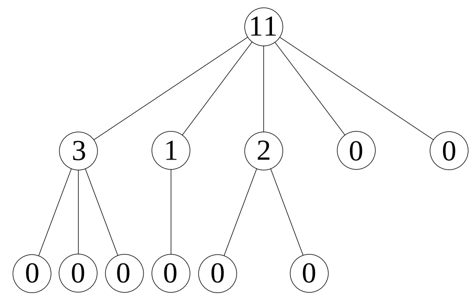

Figure 3: Tree T' for the tree T shown in Fig. 2.

It is clear from the construction of T' that the external path length of T' (that is, the sum of the depth of all its leaves) is the same as the sum of the L-depth of the L-leaves in T, and is, therefore, a lower bound on *L-path-length*(T). Observe that the sibling of a leaf in T that is a right child must be a leaf (since the label of the right child is never less than the label of its sibling) and so, by Lemma 2(i), the number of L-leaves in T is at least $(n+1)/2$. Thus the number of leaves in T' is at least $(n+1)/2$. Since the height of T' and the degree of its nodes are both at most h, it follows from standard results [6, Section 2.3.4.5] that its external path length is minimized when the tree is balanced and this quantity is at least $\Omega(n \log_h n)$. This completes the proof of (i).

In order to obtain a stronger bound on *L-path-length*(T), we note that for any $d, 0 \le d \le h$, the number of L-nodes in T at L-depth d is at most $\binom{h}{d}$ (because the addresses of any two nodes are distinct and the address of any such node can be formed by choosing d out of h positions, assigning an L to these positions, and an R to all remaining positions to the left of the right-most L). Thus, the number of nodes in T' at depth d is at most $\binom{h}{d}$. It follows that the number of leaves in T' that can occur at depth r or less, where r is as given in part (ii) of the lemma, is at most $n/4$. Recalling that the number of leaves in T' is at least $(n+1)/2$, this implies that at least $n/4$ leaves are at depth $r+1$ or more. Thus the external path length of T is at least $n(r+1)/4$, which proves (ii). □

Recall that the cost of a tree $T \in \mathcal{T}_h(n)$ is either ∞ or all its leaves are labelled 0. By Eq. (4), Lemmas 2 and 3, it follows that the lower bounds on *L-path-length*(T) given in Lemma 3 also apply to $S_h(n)$. We have thus shown the following:

Theorem 1 *There exists a set $\mathcal{S}$ of n interior-disjoint axis-parallel line segments such that the size of any orthogonal BSP for $\mathcal{S}$ of height at most h satisfies the following two lower bounds:*

(i) $\Omega(n \log n / \log h)$.

(ii) $\Omega(nr)$, *where r is an integer, $0 \le r \le h$, such that* $\sum_{d=0}^{r} \binom{h}{d} \le n/8$.

This theorem implies that, in the worst case, an orthogonal BSP of $O(n)$ size must be highly unbalanced and, on the other hand, an orthogonal BSP of height $O(\log n)$ must have size $\Omega(n \log n)$. The first claim follows readily from Theorem 1(i), while the proof of the second claim needs more effort and requires the stronger lower bound given in Theorem 1(ii).

Corollary 1.1 *Let c be any constant. There exists a set $\mathcal{S}$ of n interior-disjoint axis-parallel line segments such that any orthogonal BSP for $\mathcal{S}$ that has size at most cn must have height at least $\Omega(n^{\gamma/c})$. Here γ is a constant independent of c.*

Corollary 1.2 *Let c be any constant. There exists a set $\mathcal{S}$ of n interior-disjoint axis-parallel line segments such that any orthogonal BSP for $\mathcal{S}$ that has height at most $c \log n$ must have size at least $\Omega(n \log n / \log(c+1))$.*

Proof Let $h \leq c \log n$ denote the height of an orthogonal BSP for $\mathcal{S}$. We set $r = \lceil \epsilon \log n \rceil$ and show that for sufficiently small ϵ (depending on c), $\sum_{d=0}^{r} \binom{h}{d} \leq n/8$. By Theorem 1(ii), it would follow that the size of the BSP is at least $\Omega(n\epsilon \log n)$.

Using the identity $\binom{m+1}{t} = \binom{m}{t-1} + \binom{m}{t}$ repeatedly, we can easily show that $\sum_{d=0}^{r} \binom{h}{d} \leq \binom{h+r}{r}$. Furthermore

$$\binom{h+r}{r} \leq \frac{(h+r)^r}{r!} \leq \left[\frac{(h+r)e}{r}\right]^r = 2^{r \log\left[\left(1+\frac{h}{r}\right)e\right]},$$

where, in the second step, we have used the fact that $r! > (r/e)^r$ for $r \geq 1$. Since $h \leq c \log n$ and $\epsilon \log n \leq r \leq 2\epsilon \log n$, for sufficiently large n, we can simplify this to obtain

$$\sum_{d=0}^{r} \binom{h}{d} \leq n^{2\epsilon \log\left[\left(1+\frac{c}{\epsilon}\right)e\right]}. \tag{5}$$

It is clear that the exponent of n can be made less than one by choosing a sufficiently small ϵ, depending on c. More precisely, for $\epsilon = 1/(\beta \log(c+1))$, where β is a suitable constant independent of c, the exponent can be made less than one (note that $c \geq 1$, since the BSP must have at least $2n+1$ nodes). This completes the proof. □

Acknowledgements

I would like to thank Mordecai Golin, Chiew-Lan Tai, and an anonymous referee for useful comments.

References

[1] P. B. Callahan and S. R. Kosaraju. A decomposition of multidimensional point sets with applications to k-nearest-neighbors and n-body potential fields. *J. ACM*, 42:67–90, 1995.

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

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

[4] A. Dumitrescu, J. S. B. Mitchell, and M. Sharir. Binary space partitions for axis-parallel segments, rectangles, and hyperrectangles. In *Proc. 17th Annu. ACM Sympos. Comput. Geom.*, pages 141–150, 2001.

[5] H. Fuchs, Z. M. Kedem, and B. Naylor. On visible surface generation by a priori tree structures. *Comput. Graph.*, 14(3):124–133, 1980.

[6] D. E. Knuth. *Fundamental Algorithms*, volume 1 of *The Art of Computer Programming*. Addison-Wesley, Reading, MA, 3rd edition, 1997.

[7] B. Naylor and 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, 1986.

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

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

[10] H. Samet. *Applications of Spatial Data Structures: Computer Graphics, Image Processing, and GIS*. Addison-Wesley, Reading, MA, 1990.

[11] S. J. Teller. *Visibility Computations in Densely Occluded Polyhedral Environments*. PhD thesis, Department of Computer Science, University of California, Berkeley, 1992.

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

[13] C. Tóth. A note on binary plane partitions. In *Proc. 17th Annu. ACM Sympos. Comput. Geom.*, pages 151–156, 2001.