

Analyzing Bounding Boxes for Object Intersection

SUBHASH SURI

Washington University

PHILIP M. HUBBARD

Interval Research Corporation

and

JOHN F. HUGHES

Brown University

Heuristics that exploit bounding boxes are common in algorithms for rendering, modeling, and animation. While experience has shown that bounding boxes improve the performance of these algorithms in practice, the previous theoretical analysis has concluded that bounding boxes perform poorly in the worst case. This paper reconciles this discrepancy by analyzing intersections among n geometric objects in terms of two parameters: α , an upper bound on the *aspect ratio* or elongatedness of each object; and σ , an upper bound on the *scale factor* or size disparity between the largest and smallest objects. Letting K_o and K_b be the number of intersecting object pairs and bounding box pairs, respectively, we analyze a ratio measure of the bounding boxes' efficiency, $\rho = K_b/(n + K_o)$. The analysis proves that $\rho = O(\alpha \sqrt{\sigma \log^2 \sigma})$ and $\rho = \Omega(\alpha \sqrt{\sigma})$.

One important consequence is that if α and σ are small constants (as is often the case in practice), then $K_b = O(K_o) + O(n)$, so an algorithm that uses bounding boxes has time complexity proportional to the number of actual object intersections. This theoretical result validates the efficiency that bounding boxes have demonstrated in practice. Another consequence of our analysis is a proof of the output-sensitivity of an algorithm for reporting all intersecting pairs in a set of n convex polyhedra with constant α and σ . The algorithm takes time $O(n \log^{d-1} n + K_o \log^{d-1} n)$ for dimension $d = 2, 3$. This running time improves on the performance of previous algorithms, which make no assumptions about α and σ .

A preliminary version of this paper was presented at the 9th Symposium on Discrete Algorithms, San Francisco, CA, 1998.

S. Suri's research was supported in part by NSF grant CCR-9501494. P. M. Hubbard's research was conducted while at Washington University, St. Louis. J. F. Hughes' research was supported in part by the NSF Science & Technology Center for Computer Graphics and Scientific Visualization, Sun, Microsoft, and NASA.

Authors' addresses: S. Suri, Computer Science, Washington University, St. Louis, MO 63130; P. M. Hubbard, Interval Research Corporation, Palo Alto, CA 94304; J. F. Hughes, Computer Science, Brown University, Providence, RI 02912.

Permission to make digital/hard copy of part or all of this work for personal or classroom use is granted without fee provided that the copies are not made or distributed for profit or commercial advantage, the copyright notice, the title of the publication, and its date appear, and notice is given that copying is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists, requires prior specific permission and/or a fee.

© 2000 ACM 0730-0301/99/0700-0257 \$5.00

Categories and Subject Descriptors: F.2.2 [Analysis of Algorithms and Problem Complexity]: Nonnumerical Algorithms and Problems—*Geometrical problems and computations*; I.1.2 [Symbolic and Algebraic Manipulation]: Algorithms—*Analysis of algorithms*; I.3.0 [Computer Graphics]: General; I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling; I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism

General Terms: Algorithms, Performance, Theory

Additional Key Words and Phrases: Aspect ratio, bounding boxes, collision detection

1. INTRODUCTION

Many computer graphics algorithms improve their performance by using *bounding boxes*. The bounding box of a geometric object is a simple volume that encloses the object, forming a conservative approximation to the object. The most common form is an *axis-aligned* bounding box, whose extent in each dimension of the space is bounded by the minimum and maximum coordinates of the object in that dimension (see Figure 1(a) for an example).

Bounding boxes are useful in algorithms that should process only objects that intersect. Two objects intersect only if their bounding boxes intersect, and intersection testing is almost always more efficient for objects' bounding boxes than for the objects themselves. Thus, bounding boxes allow an algorithm to quickly perform a "trivial reject" test that prevents more costly processing in unnecessary cases. This heuristic appears in algorithms for rendering, from traditional algorithms for visible-surface determination [Foley et al. 1996] to algorithms that optimize clipping through view-frustum culling [Greene 1994], and recent image-based techniques that reconstruct new images from the reprojected pixels of reference images [McMillan 1997]. Bounding boxes are also common in algorithms for modeling, from techniques that define complex shapes as Boolean combinations of simpler shapes [Hoffmann 1989] to techniques that verify the clearance of parts in an assembly [Garcia-Alonso et al. 1995]. Animation algorithms also exploit bounding boxes, especially collision-detection algorithms for path planning [Latombe 1991] and the simulation of physically-based motion [Cohen et al. 1995; Klosowski et al. 1998; Moore and Wilhelms 1988].

While empirical evidence demonstrates that the bounding box heuristic improves performance in practice, the goal of *proving* that bounding boxes maintain high performance in the worst case has remained elusive. To understand the difficulties in such a proof, consider the use of bounding boxes when detecting pairs of colliding objects from a set \mathcal{S} of n polyhedra. Let K_o be the number of colliding pairs of objects and let K_b be the number of colliding pairs of bounding boxes. Figure 1 (b) shows an example in which $K_b = \Omega(n^2)$ while $K_o = O(1)$, meaning that the bounding box heuristic adds only unnecessary overhead, and a collision-detection algorithm that uses the heuristic is slower than one that naively tests every pair of objects for collision.




Fig. 1. (a) A polygonal object and its axis-aligned bounding box. (b) An example with $K_b = \Omega(n^2)$ and $K_o = O(1)$.

Intuitively, the poor performance in this example is due to the pathological shapes of the objects in \mathcal{S} . In this paper we identify two natural measures of the degree to which object shapes are pathological and analyze the bounding box heuristic in terms of these measures. We show that if the *aspect ratio* α and *scale factor* σ are bounded by small constants (as is generally the case in practice), then the bounding box heuristic avoids poor performance in the worst case.

The *aspect ratio* measures the *elongatedness* of an object. In classical geometry, the aspect ratio of a rectangle is defined as the ratio of its length to its width. This definition can be extended in a variety of ways to general objects and dimensions greater than two. It is often defined as the ratio between the volumes of the smallest ball enclosing the object and the largest ball contained in the object. We find it convenient to use the volumes of L_∞ -norm balls in the d -space.¹ Given a solid object P in d -space, let $b(P)$ denote the smallest L_∞ ball containing P and let $c(P)$ denote the largest L_∞ ball contained in P . The aspect ratio of P is defined as

$$\alpha(P) = \frac{\text{vol}(b(P))}{\text{vol}(c(P))},$$

where $\text{vol}(P)$ denotes the d -dimensional volume of P . We call $b(P)$ the *enclosing box*, and $c(P)$ the *core* of P . Thus, the aspect ratio measures the volume of the enclosing box relative to the core. For a set of objects

¹In two dimensions, for instance, the L_∞ ball of radius r and center o is the axis-aligned square of side length $2r$, with center o . The choice of the norm affects only the dimension-dependent constant factors, so our results also apply to L_1 or L_2 balls, with appropriate changes in the multiplicative constant.

$\mathcal{S} = \{P_1, P_2, \dots, P_n\}$, the aspect ratio is the smallest α such that $\alpha \geq \alpha(P_i)$, for $i = 1, 2, \dots, n$.

The *scale factor* for a set of objects measures the *disparity* between the largest and smallest objects. For a set $\mathcal{S} = \{P_1, P_2, \dots, P_n\}$ of objects in d -space, we say that \mathcal{S} has *scale factor* σ if for all $1 \leq i, j \leq n$,

$$\frac{\text{vol}(b(P_i))}{\text{vol}(b(P_j))} \leq \sigma.$$

The analysis in this paper focuses on the ratio

$$\rho = \frac{K_b}{n + K_o},$$

where K_o is the number of object pairs in \mathcal{S} with nonempty intersection, and K_b is the number of object pairs whose enclosing boxes intersect.² This ratio can be seen as a relative performance measure of the bounding box heuristic, because the denominator represents the best-case that an algorithm using the heuristic can achieve in practice. The denominator's $O(n)$ term reflects the overhead that an algorithm must incur if it does anything to each object, and also makes ρ more meaningful if $K_o = 0$. This ratio would, ideally, be a small constant. Unfortunately, the pathological case of Figure 1(b) shows that without any assumptions on α and σ , we can have $\rho = \Omega(n)$. However, if we include aspect ratio and scale factors in the analysis, we can prove the following theorem, which is the main result of our paper.

THEOREM 1.1 *Let \mathcal{S} be a set of n objects in d dimensions, with aspect bound α and scale factor σ , where d is a constant. Then, $\rho = O(\alpha\sqrt{\sigma \log^2 \sigma})$. Asymptotically, this bound is almost tight, as we can show a family \mathcal{S} achieving $\rho = \Omega(\alpha\sqrt{\sigma})$.*

There are two main implications of this theorem. First, it provides a theoretical justification for the efficiency that the bounding box heuristic shows in practice. In most applications, α and σ are small constants, so ρ is also constant. The theorem then indicates that $K_b = O(K_o) + O(n)$. An algorithm that uses the bounding box heuristic is thus nearly optimal in the asymptotic sense: it does not waste time processing bounding box intersections because their number grows no faster than the number of actual object intersections (plus the practically unavoidable $O(n)$ factor, which matches the overhead the algorithm must incur if it does anything to each object). Poor performance requires uncommon situations in which

²Notice that the L_∞ ball is a more conservative estimate than the axis-aligned bounding box, so K_b is an upper bound on the number of bounding box intersections.

$\alpha\sqrt{\sigma} = \Omega(n)$, as in Figure 1(b). The theorem also shows that performance is affected more by the aspect ratio than the scale factor, so it may be worthwhile to decompose irregularly-shaped objects into more regular pieces to reduce the aspect ratio.

The second implication of the theorem is an output-sensitive algorithm for reporting all pairs of intersecting objects in a set of n *convex polyhedra* in two or three dimensions. By using the bounding box heuristic as described in Section 9, the algorithm can report the K_o pairs of intersecting polyhedra in $O(n\log^{d-1}n + \alpha\sqrt{\sigma}K_o \log^2\sigma \log^{d-1}m)$ time, for $d = 2, 3$, where m is the maximum number of vertices in a polyhedron. (We assume that each polyhedron was preprocessed in linear time for efficient pairwise intersection detection [Dobkin and Kirkpatrick 1990].) Without the aspect and scale bounds, we are not aware of any output-sensitive algorithm for this problem in three dimensions. Even in two dimensions, the best algorithm for finding all intersecting pairs in a set of n convex polygons takes $O(n^{4/3} + K_o)$ time [Gupta et al. 1996]. If α and σ are constants, as is common in practice, then the algorithm runs in time $O(n\log^{d-1}n + K_o \log^{d-1}m)$, for $d = 2, 3$, which is nearly optimal.

2. RELATED WORK

The use of the bounding box heuristic in collision-detection algorithms is representative of its use in other algorithms. Thus, our analysis focuses on collision detection, but we believe that our results extend to other applications.

Most collision-detection algorithms that use bounding boxes can be considered as having two phases, which we call the *broad phase* and *narrow phase*. The basic structure of the algorithms is as follows:

- *Broad phase*: find all pairs of intersecting bounding boxes.
- *Narrow phase*: for each intersecting pair found by the broad phase, perform a detailed intersection test on the corresponding objects.

The broad and narrow phases have distinct characteristics, and often have been treated as independent problems for research.

Efficient algorithms for the broad phase must avoid looking at all $O(n^2)$ pairs of bounding boxes, and they do so by exploiting the specialized structure of bounding boxes. Edelsbrunner [1983] and Mehlhorn [1984] describe provably efficient algorithms for axis-aligned bounding boxes in d -space, algorithms that find the k intersecting pairs in $O(n\log^{d-1}n + k)$ time and $O(n\log^{d-2}n)$ space. A variety of heuristic methods are used in practice [Cohen, et al. 1995; Held et al. 1996; Hubbard 1995], and empirical evidence suggest that these algorithms perform well; the “sweep-and-prune” algorithm implemented in the I-COLLIDE package of Cohen et al. [1995] currently appears to be the method of choice. It might seem desirable to use a broad phase that replaces axis-aligned bounding boxes

with objects' convex hulls, which provide a tighter form of bound. Unfortunately, no provably efficient algorithm is known for finding the intersections between n convex polyhedra in three dimensions. In two dimensions, though, a recent algorithm of Gupta et al. [1996] can report the intersecting pairs of convex polygons in time $O(n^{4/3} + k)$.

The narrow phase solves the problem of determining the contact or interpenetration between two objects. Thus, the performance of a narrow phase algorithm does not depend on n , the number of objects in the set, but rather on the complexity of each object. If the objects are convex polyhedra, then a method due to Dobkin and Kirkpatrick [1990] can decide whether two objects intersect in $O(\log^{d-1}m)$ time, where m is the total number of edges in the two polyhedra, and $d \leq 3$ is the dimension. This algorithm preprocesses the polyhedra in a separate phase that runs in linear time. Using this preprocessing, we can also compute an explicit representation of the intersection of two convex polyhedra in time $O(m)$, as shown by Chazelle [1992]. If only one of the objects in the pair is convex, then intersection detection can be performed in time $O(m \log m)$ [Dobkin et al. 1993]. The problem is more difficult if both polyhedra are nonconvex, and only recently has a subquadratic time algorithm been discovered for deciding if two nonconvex polyhedra intersect [Schömer and Thiel 1995]. This algorithm takes $O(m^{8/5+\epsilon})$ time to determine the first collision between two polyhedra, one of which is stationary and the other translating. While the provable running times of these algorithms are important results, they are primarily of theoretical interest because the algorithms are too complicated to be practical. As an alternative, a variety of heuristic methods have been developed that tend to work well in practice [Gottschalk et al. 1996; Klosowski et al. 1998]. These methods use hierarchies of bounding volumes and tree-descent schemes to determine intersections.

Our analysis of the bounding box heuristic is related to the idea of "realistic input models," which has become a topic of recent interest in computational geometry. In a recent paper, de Berg et al. [1997] suggested classifying various models of realistic input models into four main classes: fatness, density, clutter, and cover complexity. Briefly, an object is fat if it does not have long and skinny parts; a scene has low clutter if any cube not containing a vertex of an object intersects at most a constant number of objects; a scene has low density if a ball of radius r intersects only a constant number of objects whose minimum enclosing ball has radius at least r ; the cover complexity is a measure of the relative sparseness of an object's neighborhood. One of the first nontrivial results in this direction is by Matoušek et al. [1994], who showed that the union of n fat triangles has complexity $O(n \log \log n)$, as opposed to $\Theta(n^2)$ for arbitrary triangles; a triangle is fat if its minimum angle exceeds δ , for a constant $\delta > 0$. Efrat and Sharir [1997] generalize this result to show that the union of n convex objects has complexity $O(n^{1+\epsilon})$, provided that each object is fat and each pair of objects intersects only in a constant number of points. Additional

results on fat or uncluttered objects can be found in de Berg [1995]; Halperin and Overmars [1994]; and van der Stappen et al. [1993].

3. ANALYSIS OVERVIEW

Our proof for the upper bound on ρ consists of three steps. We first consider the case of arbitrary α but fixed σ (Section 4). Next, in Section 5, we allow both α and σ to be arbitrary, but assume that there are only two kinds of objects: one with box sizes α , core sizes 1, and the other with box sizes $\alpha\sigma$, core size σ (the two extreme ends of the scale factor). Finally, in Section 6, we handle the general case, where objects can have any box size in the range $[\alpha, \alpha\sigma]$. We first detail our proof for two dimensions, and then sketch how to extend it to arbitrary dimensions in Section 7.

4. ARBITRARY ASPECT RATIO BUT FIXED SCALE

We start by assuming that the set \mathcal{S} has scale factor one, that is, $\sigma = 1$; the aspect ratio bound α can be arbitrary. (Any constant bound for σ will work for our proof; we assume one for convenience. The most straightforward way to enforce this scale bound is to make every object's enclosing box to be the same size.) We show that in this case $\rho(\mathcal{S}) = O(\alpha)$. We describe our proof in two dimensions; the extension to higher dimensions is quite straightforward, and is sketched in Section 7.

Without loss of generality, let us assume that each object P in \mathcal{S} has $\text{vol}(c(P)) \geq 1$, and $\text{vol}(b(P)) \leq \alpha$. Recall that a L_∞ box of volume α in two dimensions is a square of side length $\sqrt{\alpha}$. We call this a *size α box*. Consider a tiling of the plane by size α boxes that covers the portion of the plane occupied by the bounding boxes of the objects, namely, $\cup b(P_i)$; see Figure 2. We consider each box semiopen, so that the boundary shared by two boxes belongs to the one on the left, or above. Thus, each point of the plane belongs to at most one box.

We assume an underlying unit lattice in the plane and assign each object P to the (unique) *lexicographically smallest lattice point* contained in P . (Such a point exists because the core is closed and has volume at least one.) Let $m(q)$ be the number of objects assigned to a lattice point q , and let M_i denote the total number of objects assigned to the lattice points contained in a box B_i . That is,

$$M_i = \sum_{q \in B_i} m(q),$$

where $q \in B_i$ means that the lattice point q lies in the box B_i . Since the boxes in the tiling are disjoint, we have the equality $\sum_i M_i = n$. We derive the bounds on K_b and K_o in terms of M_i .

LEMMA 4.1 *Given a set of objects \mathcal{S} with aspect bound α and scale bound $\sigma = 1$, let B_1, B_2, \dots, B_p denote a tiling by size α boxes, as defined above,*




Fig. 2. Tiling of the plane by boxes of size α . The unit size core for the object in B_1 is also shown.

and let M_i denote the total number of objects assigned to lattice points in B_i , for $i = 1, 2, \dots, p$. Then

$$K_b \leq 25 \sum_{i=1}^p M_i^2.$$

PROOF. Consider an object P assigned to B_i , and let P_j be another object whose box intersects $b(P)$. Suppose P_j is assigned to the box B_j . Since $b(P) \cap b(P_j) \neq \emptyset$, the L_∞ norm distance between the boxes B_i and B_j is at most $2\sqrt{\alpha}$. This means that B_j is among the 25 boxes that lie within $2\sqrt{\alpha}$ wide corridor around B_i (including B_i). Suppose that the boxes are labeled B_1, B_2, \dots, B_p in the row-major order—top to bottom, left to right in each row. Assume that the number of columns in the box tiling is k . Then, the preceding discussion shows that if the boxes of objects P_i and P_j intersect, and these objects are assigned to boxes B_i and B_j , then we must have

$$j = i + ck + d,$$

where $c, d \in \{0, \pm 1, \pm 2\}$. (The box B_j can be at most two rows and two columns away from B_i . For instance, the box preceding two rows and two columns from B_i is B_{i-2k-2} ; see Figure 3. The number of box pair intersections contributed by B_i and B_j is clearly no more than $M_i M_j$. Thus, the total number of such intersections is bounded by




Fig. 3. A box, shown in dark at the center, and its 24 neighbors.

$$\sum_{i=1}^p \sum_{j=i+ck+d} M_i M_j,$$

where $c, d \in \{0, \pm 1, \pm 2\}$. Recalling that $x_1 x_2 \leq (x_1^2 + x_2^2)/2$, for reals x_1, x_2 , we can bound the intersection count by

$$\sum_{i=1}^p \sum_{j=i+ck+d} \frac{1}{2}(M_i^2 + M_j^2).$$

There are 5 possible values for c and d each, and so altogether 25 values for j for each i . Since each index can appear once as the i and once as the j , we get that the maximum number of intersections is at most

$$25 \sum_{i=1}^p M_i^2.$$

This completes the proof of the lemma. \square

Next, we establish a lower bound on the number of intersecting object pairs. We need the following elementary fact.

LEMMA 4.2 Consider nonnegative numbers a_1, a_2, \dots, a_n , and b_1, b_2, \dots, b_n . Then,

$$\frac{a_1 + a_2 + \dots + a_n}{b_1 + b_2 + \dots + b_n} \leq \max_{1 \leq i \leq n} \frac{a_i}{b_i}.$$

PROOF. Let m denote the index for which the ratio a_i/b_i is maximized. Since $b_i(a_m/b_m) \geq a_i$, summing it over all i , we get

$$\frac{a_m}{b_m} \sum_{i=1}^n b_i \geq \sum_{i=1}^n a_i.$$

Dividing both sides by $\sum_{i=1}^n b_i$ completes the proof of the lemma. \square

Let us now focus on objects assigned to a box B_i in our tiling. If L_i is the number of intersecting pairs among objects assigned to B_i , then we have the following:

$$\begin{aligned}
 \rho(\mathcal{F}) &= \frac{K_b}{n + K_o} \\
 &\leq \frac{25 \sum_{i=1}^p M_i^2}{n + \sum_{i=1}^p L_i} \\
 &\leq \frac{25(M_1^2 + M_2^2 + \cdots + M_p^2)}{(M_1 + L_1) + (M_2 + L_2) + \cdots + (M_p + L_p)} \\
 &\leq \max_{1 \leq i \leq p} \frac{25M_i^2}{(M_i + L_i)},
 \end{aligned}$$

where the second to last inequality follows from the fact that $\sum_i M_i = n$, and the last inequality follows from the preceding lemma. We establish an upper bound on the right-hand side of this inequality by proving a lower bound on the denominator term.

Fix a box B_i in the following discussion, where $1 \leq i \leq p$. Consider a lattice point q in it and the $m(q)$ objects assigned to q . These objects all have q in common, and therefore we get at least $\binom{m(q)}{2}$ object pair intersections. (Observe that each object is assigned to a unique lattice point, and so we count each intersection at most once.) Thus, the total number of pairwise intersections L_i among objects assigned to B_i is at least

$$\sum_{q \in B_i} \binom{m(q)}{2}.$$

We show that the ratio $25M_i^2/(M_i + L_i)$ never exceeds $c\alpha$, where c is an absolute constant. Considering M_i fixed, this ratio is maximized when L_i is minimized.

LEMMA 4.3 *Let x_1, x_2, \dots, x_n be nonnegative numbers that sum to z , with $n \leq N$, where z and N are fixed. Then, the minimum value of $\sum_{i=1}^n \binom{x_i}{2}$ is $z(z - n)/2n$, which is achieved when $n = N$ and $x_i = z/n$, for $i = 1, 2, \dots, n$.*

PROOF. We observe the following equalities:

$$\sum_{i=1}^n \binom{x_i}{2} = \sum_{i=1}^n \frac{x_i(x_i - 1)}{2} = \frac{1}{2}(\sum_{i=1}^n x_i^2 - z).$$

Thus, $\sum_{i=1}^n \binom{x_i}{2}$ is minimized when $\sum_{i=1}^n x_i^2/2$ is minimized. First, considering n fixed, Cauchy's inequality [Hardy et al. 1988] implies that the sum is minimized when $x_i = z/n$. Next, since $n(z/n)^2 < (n + 1)(z/(n + 1))^2$, the sum decreases for increasing n , and so is minimized for $n = N$. The lemma follows. \square

Since no square box of size α can have more than $2\lceil \alpha \rceil$ lattice points in it, we get a lower bound on L_i by setting $m(q) = M_i/2\lceil \alpha \rceil$, for all q . Thus,

$$L_i \geq \frac{1}{2} M_i \left(\frac{M_i}{2\lceil \alpha \rceil} - 1 \right).$$

LEMMA 4.4 $\rho(\mathcal{F}) = O(\alpha)$.

PROOF. Using the bound for L_i above, we have

$$\rho(\mathcal{F}) \leq \frac{25M_i^2}{\frac{1}{2} M_i \left(\frac{M_i}{2\lceil \alpha \rceil} - 1 \right) + M_i} \leq 100\lceil \alpha \rceil.$$

This completes the proof. \square

THEOREM 4.5 *Let \mathcal{F} be a set of n objects in the plane, with aspect bound α and scale bound $\sigma = 1$. Then, $\rho(\mathcal{F}) = O(\alpha)$.*

5. OBJECTS OF TWO FIXED SIZES

In this section we generalize the result of the previous section to the case where objects come from the two extreme ends of the scale: their box size is either α or $\alpha\sigma$. To simplify our analysis, we assume that $\alpha = 4^a$ and $\sigma = 4^b$ for some integers $a, b > 0$. (Otherwise, just use the next nearest powers of 4 as upper bounds for α, σ . In d dimensions, α and σ are assumed to be integral powers of 2^d .)

Let us call an object *large* if its enclosing box has size $\alpha\sigma$, and *small* otherwise. Clearly, there are only three kinds of intersections: large-large, small-small, and large-small. Let K_b^l, K_b^s and K_b^{sl} , respectively, count these intersections for the enclosing boxes. So, for example, K_b^{sl} is the number of pairs consisting of one large and one small object whose boxes intersect. Similarly, define the terms K_o^l, K_o^s and K_o^{sl} for object pair intersections. The ratio bound can now be restated as

$$\begin{aligned} \rho(\mathcal{F}) &= \frac{K_b^l + K_b^s + K_b^{sl}}{K_o^l + K_o^s + K_o^{sl} + n} \\ &\leq 3 \max \left\{ \frac{K_b^l}{K'}, \frac{K_b^s}{K'}, \frac{K_b^{sl}}{K'} \right\} \end{aligned} \quad (1)$$

where $K' = K_o^l + K_o^s + K_o^{sl} + n$. We know from the result of the previous section that $K_b^l/K', K_b^s/K' \leq c\alpha$, for some constant c . So, we only need to establish a bound on the third ratio, $K_b^{sl}/(K_o^l + K_o^s + K_o^{sl} + n)$, which we do as follows.

Let us again tile the plane with boxes of volume $\alpha\sigma$. Call these boxes B_1, B_2, \dots, B_p . Underlying this tiling are two grids: a *level σ grid*, which divides the boxes into cells of size σ , and a *level 1 grid*, which divides the boxes into cells of size 1. The level σ grid has vertices at coordinates $(i\sqrt{\sigma}, j\sqrt{\sigma})$, while the finer grid has vertices at coordinates (i, j) , for integers i, j . The level σ grid is used to reason about large objects, while the level 1 grid is used for small objects. We mimic the proof of the previous section and assign objects of each class to an appropriate box. In order to do this, we need to define subboxes of size α within each size $\alpha\sigma$ original box.

Consider a large box B_i . The level σ grid partitions B_i into α boxes of volume σ each. Next, we also partition B_i into σ subboxes, each of volume α . Since $\alpha = 4^a$ and $\sigma = 4^b$, for integers $a, b > 0$, these subboxes are perfectly aligned with both the level 1 and level σ grids. (Along a side of B_i , the σ grid has vertices at distance multiples of $\sqrt{\sigma} = 2^b$, while the vertices of the subboxes lie at distance multiples of $\sqrt{\alpha} = 2^a$.) We label the σ subboxes within B_i as $B_{i1}, B_{i2}, \dots, B_{i\sigma}$, in row major order. Figure 4 illustrates these definitions by showing two boxes side by side.

Now each member of the large object set (resp. small object set) contains at least one grid point of the large (resp. small) grid. Just as in the previous section, we assign each object to a unique grid point (say, the one with lexicographically smallest coordinates). Let X_i denote the number of large objects assigned to all the grid points in B_i . Let y_{ij} , for $j = 1, 2, \dots, \sigma$, denote the number of small objects assigned to the subbox B_{ij} . Define also $Y_i = \sum_{j=1}^{\sigma} y_{ij}$ to be the total number of small objects assigned to level one grid points in B_i .

We estimate an upper bound on K_b^{sl} and a lower bound on K_o^{sl} , in terms of X_i and Y_i . Fix a box B_i . The enclosing box of a large object P_i , assigned to B_i , can intersect the box of a small object P_j , assigned to B_j , only if B_j is one of the 25 neighbors of B_i (including itself) that form the two layers of boxes around B_i (see Figure 3 again). Let B_i^m be the box with a maximum number of small objects among the 25 neighbors of B_i , and let Y_i^m be the count of the small objects in B_i^m . That is, $Y_i^m = \max_j \{Y_j \mid B_j \text{ is one of 25 neighbors}$




Fig. 4. The box on the left shows large grid and the one on the right shows small grid as well as the subboxes. In this figure, $\alpha = 4$ and $\sigma = 16$.

of B_i , and B_i^m is the box corresponding to Y_i^m . Then, we have the following upper bound:

$$K_b^{sl} \leq 25 \sum_{i=1}^p X_i Y_i^m.$$

Next, we estimate lower bounds on the number of object pair intersections. Let L_i denote the number of object pair intersections among the large objects assigned to B_i , and let S_i denote the object pair intersections among the small objects assigned to B_i . Since there are only α grid points for the large objects in B_i , by Lemma 4.3, we have

$$L_i \geq \frac{1}{2} X_i \left(\frac{X_i}{\alpha} - 1 \right). \tag{2}$$

Similarly, each of the subboxes B_{ij} , for $j = 1, 2, \dots, \sigma$, has α grid points of the level 1 grid. Thus, we also have

$$S_i \geq \sum_{j=1}^{\sigma} \sigma \frac{y_{ij}}{2} \left(\frac{y_{ij}}{\alpha} - 1 \right). \tag{3}$$

In deriving our bound, we use the conservative estimate of $\sum_{i=1}^p (L_i + S_i)$ for K_o ; that is, only count the intersections between two large or two small objects. We also use the notation S_i^m for the number of object-pair intersections among the small objects assigned to B_i^m . We have the following inequalities:

$$\begin{aligned}
 \frac{K_b^{sl}}{n + K_o} &\leq \frac{25 \sum_{i=1}^p X_i Y_i^m}{n + \sum_{i=1}^p (L_i + S_i)} \\
 &= \frac{25^2 \sum_{i=1}^p X_i Y_i^m}{25 \sum_{i=1}^p (X_i + L_i + Y_i + S_i)} \\
 &\leq \frac{25^2 \sum_{i=1}^p X_i Y_i^m}{\sum_{i=1}^p (X_i + L_i + Y_i^m + S_i^m)} \\
 &\leq \max_{1 \leq i \leq p} \frac{25^2 X_i Y_i^m}{X_i + L_i + Y_i^m + S_i^m},
 \end{aligned}$$

where the second inequality follows from the fact that $\sum_{i=1}^p (X_i + Y_i) = n$; the third follows from the fact that a particular box B_i^m can contribute the Y_i^m term to at most its 25 neighbors; and the final inequality follows from Lemma 4.2. The remaining step of the proof now is to show that the above inequality is $O(\alpha \sqrt{\sigma})$. First, by summing up the terms in Eqs. (2) and (3), we observe the following:

$$X_i + L_i + Y_i^m + S_i^m \geq \frac{X_i^2 + \sum_{j=1}^{\sigma} (y_{ij}^m)^2}{2\alpha},$$

where we recall that $\sum_{j=1}^{\sigma} y_{ij}^m = Y_i^m$. Thus, we have

$$\begin{aligned}
 \frac{X_i Y_i^m}{X_i + Y_i^m + L_i + S_i^m} &\leq \frac{2\alpha X_i Y_i^m}{X_i^2 + \sum_{j=1}^{\sigma} (y_{ij}^m)^2} \\
 &\leq \frac{2\alpha X_i Y_i^m}{X_i^2 + \sigma \left(\frac{Y_i^m}{\sigma}\right)^2} \\
 &\leq \frac{2\alpha \sigma X_i Y_i^m}{\sigma X_i^2 + (Y_i^m)^2},
 \end{aligned}$$

where once again Cauchy's inequality is invoked to show that $\sum_{j=1}^{\sigma} (y_{ij}^m)^2 \geq \sigma(Y_i^m/\sigma)^2$. It can be easily shown that this ratio is at most $2\alpha\sqrt{\sigma}$, as follows. If $Y_i^m \leq \sqrt{\sigma}X_i$, then we have

$$\frac{2\alpha\sigma X_i Y_i^m}{\sigma X_i^2 + (Y_i^m)^2} \leq \frac{2\alpha\sigma X_i \sqrt{\sigma} X_i}{\sigma X_i^2} \leq 2\alpha\sqrt{\sigma}.$$

Otherwise, $Y_i^m > \sqrt{\sigma}X_i$, and we have

$$\frac{2\alpha\sigma X_i Y_i^m}{\sigma X_i^2 + (Y_i^m)^2} \leq \frac{2\alpha\sigma \frac{Y_i^m}{\sqrt{\sigma}} Y_i^m}{(Y_i^m)^2} \leq 2\alpha\sqrt{\sigma}.$$

This shows that $K_b^{sl}/(n + K_o) = O(\alpha\sqrt{\sigma})$. Combining this with Inequality (1), we get the desired result, which is stated in the following theorem.

THEOREM 5.1 *Suppose \mathcal{S} is a set of n objects in the plane such that each object has aspect ratio at most α , and the enclosing box of each object has size either α or $\alpha\sigma$. Then, $\rho(\mathcal{S}) = O(\alpha\sqrt{\sigma})$.*

6. THE GENERAL CASE

We now are in a position to prove our main theorem. Suppose \mathcal{S} is a set of n objects, with aspect ratio bound α and scale factor σ . Recall that for simplicity we assume that both α and σ are powers of four. We partition the set \mathcal{S} into $O(\log \sigma)$ classes, $\mathcal{C}_0, \mathcal{C}_1, \dots, \mathcal{C}_k$, for $k = \log \sigma$, such that an object P belongs to class \mathcal{C}_i if $2^i \leq \text{vol}(c(P)) < 2^{i+1}$. (Equivalently, the enclosing boxes of objects in class \mathcal{C}_i have volumes between $\alpha 2^i$ and $\alpha 2^{i+1}$.) Each class behaves like a fixed size family (the case considered in Section 4), and so we have $\rho(\mathcal{C}_i) = O(\alpha)$, for $i = 0, 1, \dots, \log \sigma$. Any pair of classes behaves like the case considered in Section 5, implying that $\rho(\mathcal{C}_i \cup \mathcal{C}_j) = O(\alpha\sqrt{\sigma})$, for $0 \leq i, j \leq \log \sigma$. We can now formalize this argument to show that $\rho(\mathcal{S}) = O(\alpha\sqrt{\sigma} \log^2 \sigma)$.

Let K_b^{ij} , for $0 \leq i, j \leq \log \sigma$, denote the number of object pairs (P, P') whose enclosing boxes intersect such that $P \in \mathcal{C}_i$ and $P' \in \mathcal{C}_j$. Similarly, define K_o^{ij} . Then, we have the following:

$$\begin{aligned} \rho(\mathcal{S}) &= \frac{\sum_i \sum_j K_b^{ij}}{\sum_i \sum_j K_o^{ij} + n} \\ &\leq \left(\frac{\max_{i,j} K_b^{ij}}{\sum_i \sum_j K_o^{ij} + n} \right) \log^2 \sigma \\ &\leq O(\alpha\sqrt{\sigma} \log^2 \sigma) \end{aligned}$$

where the second inequality follows from the fact that i, j are each bounded by $\log \sigma$, and the last inequality follows directly from Theorem 5.1. This proves our main result, which we restate in the following theorem.

THEOREM 6.1 *Let \mathcal{S} be a set of n objects in the plane, with aspect ratio bound α and scale factor bound σ . Then, $\rho(\mathcal{S}) = O(\alpha \sqrt{\sigma} \log^2 \sigma)$.*

7. EXTENSION TO HIGHER DIMENSIONS

The 2-dimensional result might lead one to suspect that the bound in d dimensions, for $d \geq 2$, will be $O(\alpha \sigma^{1/d})$. In fact, the asymptotic bound in d dimension turns out to be the same as in two dimensions—only the constant factors are different. A closer examination shows that the exponent on σ in Theorem 6.1 arises not from the dimension, but rather from Cauchy’s inequality.

Our proof of Theorem 6.1 extends easily to d dimensions, for $d \geq 3$. The structure of the proof remains exactly the same. We tile the d -dimensional space with boxes (L_∞ balls). The main difference arises in the number of neighboring boxes for a given box B_i . While in the plane, a box has at most 5^2 neighboring boxes in the two surrounding layers, this number increases to 5^d in d dimensions. Since our arguments have been volume based, they hold in d dimensions as well. Our main theorem in d dimensions can be stated as follows.

THEOREM 7.1 *Let \mathcal{S} be a set of n objects in d -space, with aspect ratio bound α and scale factor bound σ . Then, $\rho(\mathcal{S}) = O(\alpha \sqrt{\sigma} \log^2 \sigma)$, where the constant is about 5^d .*

8. LOWER BOUND CONSTRUCTIONS

We first describe a construction of a family \mathcal{S} with $\sigma = 1$, which shows $\rho(\mathcal{S}) = \Omega(\alpha)$. The construction works in any dimension d , but for ease of exposition, we describe it in two dimensions. See Figure 5 for illustration.

Consider a square box B of size α in the standard position, namely, $B = [0, \sqrt{\alpha}] \times [0, \sqrt{\alpha}]$. We can pack roughly α unit boxes in B , in a regular grid pattern; the number is $\lfloor \sqrt{\alpha} \rfloor^2$ to be exact. We convert each of these unit boxes into a polyhedral object of aspect ratio α , by attaching two “wire” extensions at the two endpoints of its main diagonal. Specifically, consider one such unit box u , the endpoints of whose main diagonal have coordinates (a_1, a_2) and (b_1, b_2) . The b endpoint of u is connected to the point $(\sqrt{\alpha}, \sqrt{\alpha})$ with a Manhattan path, whose i th edge is parallel to the *positive* i -coordinate axes and has length $\sqrt{\alpha} - b_i$. Similarly, the a endpoint of u is connected to the origin with a Manhattan path, whose i th edge is parallel to the *negative* i -coordinate axes and has length a_i . It is easy to see that each unit box, together with the two wire extensions forms a polyhedral




Fig. 5. The lower bound construction showing $\rho(\mathcal{F}) = \Omega(\alpha)$.

object with aspect ratio α . By a small perturbation, we can ensure that no two objects intersect. The bounding boxes of each object pair intersect, however, and so we have at least $\alpha - 2$ bounding box intersections in B .

We can group our n objects into $\lfloor n/\alpha \rfloor$ groups, each group corresponding to a α -size box as above. This gives us

$$K_b \geq \left\lfloor \frac{n}{\alpha} \right\rfloor \times \binom{\alpha}{2} = \Omega(n\alpha).$$

On the other hand, $K_o = 0$, and thus, $\rho(\mathcal{F}) = \Omega(n\alpha/n) = \Omega(\alpha)$.

We next generalize this construction to establish a lower bound of $\Omega(\alpha\sqrt{\sigma})$, assuming that $\alpha\sigma \leq n$. See Figure 6.

We take a square box B' of volume $4\alpha\sigma$. We divide the lower right quadrant of B' into α subboxes of size σ . We take a copy of the construction of Figure 5, scale it up by a factor of σ , and put it in place of the lower right quadrant of B' . We extend the wires attached to each object to the corners c, d of B' . Thus, the smallest enclosing box of each object is now exactly B' , and aspect ratio is 4α . These are the big objects. Next, we take the upper-left quadrant, divide it into σ subboxes of size α each. At each α -size subbox, we place a copy of the construction in Figure 5. These are the small objects.

Altogether we want $X = n/(1 + \sqrt{\sigma})$ big objects, and $Y = n\sqrt{\sigma}/(1 + \sqrt{\sigma})$ small objects. Since there are a total of α locations for big objects, we superimpose X/α copies of the big object at each location. Similarly, there are $\alpha\sigma$ locations for the small objects, so superimpose $Y/\alpha\sigma$ copies of the small object at each location. (This is where we need the condition $\alpha\sigma \leq n$, since we want to ensure that each location receives at least one object.) Let us now estimate bounds for K_b and K_o . The enclosing box of every big object intersects the enclosing box of every small object, we have

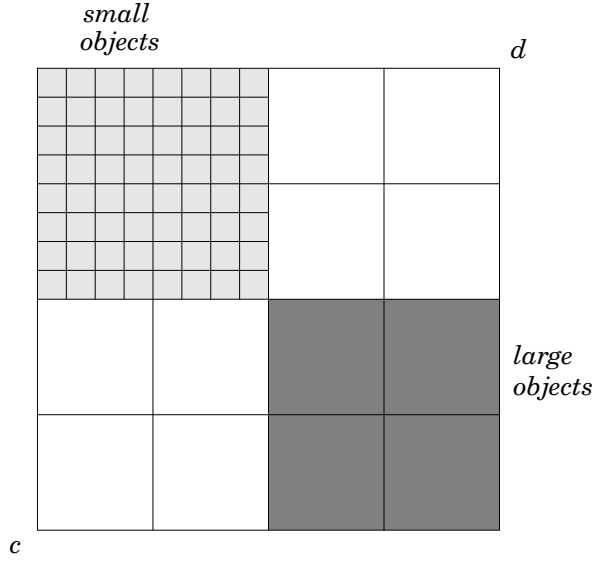


Fig. 6. The lower bound construction, showing $\rho(\mathcal{F}) = \Omega(\alpha\sqrt{\sigma})$.

$$K_b \geq XY \geq \frac{n^2\sqrt{\sigma}}{(1 + \sqrt{\sigma})^2} \tag{4}$$

On the other hand, the only object pair intersections exist between objects assigned to the same location. We therefore have

$$\begin{aligned} K_o &\leq \alpha \binom{X/\alpha}{2} + \alpha\sigma \binom{Y/\alpha\sigma}{2} \\ &\leq \alpha(X/\alpha)^2 + \alpha\sigma(Y/\alpha\sigma)^2 \\ &\leq \frac{\sigma X^2 + Y^2}{\alpha\sigma} \\ &\leq \frac{2n^2}{\alpha(1 + \sqrt{\sigma})^2}. \end{aligned}$$

Thus,

$$\begin{aligned} \rho(\mathcal{F}) &= \frac{K_b}{K_o + n} \\ &\geq \left(\frac{n^2\sqrt{\sigma}}{(1 + \sqrt{\sigma})^2} \right) / \left(\frac{n^2}{2\alpha(1 + \sqrt{\sigma})^2} + n \right) \end{aligned}$$

$$\begin{aligned}
&\geq \frac{\alpha\sqrt{\sigma}}{\frac{1}{2} + \frac{\alpha(1 + \sqrt{\sigma})^2}{n}} \\
&\geq c\alpha\sqrt{\sigma},
\end{aligned}$$

for some constant $c > 0$. (The ratio $\alpha(1 + \sqrt{\sigma})^2/n$ is bounded by a constant, since $\alpha\sigma \leq n$.)

THEOREM 8.1 *There exists a family \mathcal{S} of n objects with aspect ratio bound α and scale factor σ such that $\rho(\mathcal{S}) = \Omega(\alpha\sqrt{\sigma})$, assuming $\alpha\sigma \leq n$.*

9. APPLICATIONS AND CONCLUDING REMARKS

Theorems 6.1 and 7.1 have two interesting consequences. The first is a theoretical validation of the bounding box heuristic mentioned in Section 1. In practice, the object families tend to have bounded aspect ratio and scale factor. Thus, the number of extraneous box intersections is at most a constant factor of the number of actual object-pair intersections. This result needs no assumption about the convexity of the objects.

If the aspect ratio and scale factor grow with n , our theorem indicates their impact on the efficiency of the heuristic. The degradation of the heuristic is smooth, and not abrupt. Furthermore, the result suggests that the dependence on aspect ratio and scale factor is *not symmetric*—the complexity grows linearly with α , but only as a *square root* of σ . It is common in practice to decompose complex objects into simpler parts. Our work suggests that for collision detection purposes, reducing aspect ratio may have higher payoff than reducing scale factor. It would be interesting to verify empirically how this strategy performs in practice.

The second consequence of our theorems is an output-sensitive algorithm for reporting pairwise intersections among polyhedra. We run the two-phase intersection algorithm: the broad phase computes the pairs of objects whose bounding boxes intersect; the narrow phase then checks the object-pairs for actual intersection. If there are K intersecting box-pairs, then the algorithm of Edelsbrunner [1983] finds them in worst-case time $O(n \log^{d-1} n + K)$, where d is the dimension of the ambient space. When the aspect ratio and scale factors are constant, the number of box-pairs K is within a constant factor of the intersecting object-pairs, and therefore the total running time is proportional to output size. The bound is the strongest for *convex* polyhedra in dimensions $d = 2, 3$. We are aware of only one non-trivial result for this problem, which holds in two dimensions. Gupta et al. [Gupta, Janardan, and Smid 1996] give an $O(n^{4/3} + K_o)$ time algorithm for reporting K_o pairs of intersecting convex polygons in the plane. The problem is wide open in three and higher dimensions.

Our theorem leads to a significantly better result in two and three dimensions for small aspect and scale bounds, and nearly optimal result for *convex polyhedra*. Given n polyhedra in two or three dimensions, we can report all pairs whose *bounding boxes* intersect in time $O(n \log^{d-1} n + K_b)$ [Edelsbrunner 1983; Mehlhorn 1984], where K_b is the number of intersecting bounding box pairs. If the polyhedra are convex, then the narrow phase intersection test can be performed in $O(\log^{d-1} m)$ time [Dobkin and Kirkpatrick 1990], assuming that all polyhedra have been preprocessed in linear time; m is the maximum number of vertices in a polyhedron. If the convex polyhedra have aspect ratio at most α and scale factor at most σ , then by Theorem 7.1, the total running time of the algorithm is $O(n \log^{d-1} n + \alpha \sqrt{\sigma} K_o \log^2 \sigma \log^{d-1} m)$, for $d = 2, 3$. If α and σ are constants, then the running time is $O(n \log^{d-1} n + K_o \log^{d-1} m)$, which is nearly optimal.

Finally, Zhou and Suri [1999] have recently improved and extended this work. In particular, they close the gap between the upper and lower bound, showing that $\rho = O(\alpha \sqrt{\sigma})$. Zhou and Suri also consider a natural extension of our model, in which only the *average* aspect ratio is bounded. This may be more practical in situations where a few pathological objects may exist in the scene, but the total aspect ratio of the n objects is bounded by $\alpha_{avg} n$. The main result of Zhou and Suri [1999] for the average aspect ratio is the following theorem:

“Let \mathcal{S} be a set of n objects in d dimensions, with average aspect ratio α_{avg} and scale factor σ , where d is a constant. Then, $\rho(\mathcal{S}) = O(\alpha_{avg}^{2/3} \sigma^{1/3} n^{1/3})$. This bound is tight in the worst case.”

ACKNOWLEDGMENTS

The authors wish to thank Peter Shirley for his valuable comments on earlier versions of the proof.

REFERENCES

- CHAZELLE, B. 1992. An optimal algorithm for intersecting three-dimensional convex polyhedra. *SIAM J. Comput.* 21, 4 (Aug. 1992), 671–696.
- COHEN, J. D., LIN, M. C., MANOCHA, D., AND PONAMGI, M. 1995. I-COLLIDE: an interactive and exact collision detection system for large-scale environments. In *Proceedings of the 1995 Symposium on Interactive 3D Graphics* (Monterey, CA, Apr. 9–12, 1995), M. Zyda, Ed. ACM Press, New York, NY, 189ff.
- DE BERG, M. 1995. Linear size binary space partitions for fat objects. In *Proceedings of the Third European Symposium on Algorithms*, 252–263.
- DE BERG, M., KATZ, M., VAN DER STAPPEN, A. F., AND VLEUGELS, J. 1997. Realistic input models for geometric algorithms. In *Proceedings of the 13th Annual Symposium on Computational Geometry* (Nice, France, June 4–6, 1997), J.-D. Boissonnat, Ed. ACM Press, New York, NY, 294–303.
- DOBKIN, D., HERSHBERGER, J., KIRKPATRICK, D., AND SURI, S. 1993. Computing the intersection-depth of polyhedra. *Algorithmica* 9, 518–533.
- DOBKIN, D. P. AND KIRKPATRICK, D. G. 1990. Determining the separation of preprocessed polyhedra—a unified approach. In *Proceedings of the 17th Conference on ICALP*, Lecture Notes in Computer Science, vol. LNCS 443. Springer-Verlag, New York, 400–413.

- EDELSBRUNNER, H. 1983. A new approach to rectangle intersections, Parts I&II. *Int. J. Comput. Math.* 13, 209–229.
- EFRAT, A. AND SHARIR, M. 1997. On the complexity of the union of fat objects in the plane. In *Proceedings of the 13th Annual Symposium on Computational Geometry* (Nice, France, June 4–6, 1997), J.-D. Boissonnat, Ed. ACM Press, New York, NY, 104–112.
- FOLEY, J. D., VAN DAM, A., FEINER, S. K., AND HUGHES, J. F. 1996. *Computer Graphics (in C): Principles and Practice*. 2nd ed. Addison-Wesley systems programming series. Addison-Wesley Longman Publ. Co., Inc., Reading, MA.
- GARCIA-ALONSO, A., SERRANO, N., AND FLAQUER, J. 1994. Solving the collision detection problem. *IEEE Comput. Graph. Appl.* 14, 3 (May 1994), 36–43.
- GOTTSCHALK, S., LIN, M. C., AND MANOCHA, D. 1996. OBBTree: A hierarchical structure for rapid interference detection. In *Proceedings of the 23rd Annual Conference on Computer Graphics (SIGGRAPH '96, New Orleans, LA, Aug. 4–9, 1996)*, J. Fujii, Ed. Annual conference series. ACM Press, New York, NY, 171–180.
- GREENE, N. 1994. Detecting intersection of a rectangular solid and a convex polyhedron. In *Graphics Gems IV*, P. S. Heckbert, Ed. Academic Press Graphics Gems series. Academic Press Prof., Inc., San Diego, CA, 74–82.
- GUPTA, P., JANARDAN, R., AND SMID, M. 1996. Efficient algorithms for counting and reporting pairwise intersection between convex polygons. Tech. Rep. Computer Science, King's College, UK.
- HALPERIN, D. AND OVERMARS, M. H. 1994. Spheres, molecules, and hidden surface removal. In *Proceedings of the 10th Annual Symposium on Computational Geometry* (Stony Brook, NY, June 6–8, 1994), K. Mehlhorn, Ed. ACM Press, New York, NY, 113–122.
- HARDY, H., LITTLEWOOD, J. E., AND PÓLYA, G. 1988. *Inequalities*. Cambridge University Press, New York, NY.
- HELD, M., KLOSOWSKI, J. T., AND MITCHELL, J. S. B. 1996. Collision detection for fly-throughs in virtual environments. In *Proceedings of the 12th Annual Symposium on Computational Geometry* (Philadelphia, PA, May 24–26, 1996), S. Whitesides, Ed. ACM Press, New York, NY, 513–514.
- HOFFMANN, C. M. 1989. *Geometric and Solid Modeling: An Introduction*. Morgan Kaufmann Publishers Inc., San Francisco, CA.
- HUBBARD, P. M. 1995. Collision detection for interactive graphics applications. *IEEE Trans. Visualization Comput. Graph.* 1, 3, 218–230.
- KLOSOWSKI, J. T., HELD, M., MITCHEL, J. S. B., SOWIZRAL, H., AND ZIKAN, K. 1998. Efficient collision detection using bounding volume hierarchies of k-DOPs. *IEEE Trans. Visualization Comput. Graph.* 4, 1, 21–36.
- LATOMBE, J.-C. 1991. *Robot Motion Planning*. Kluwer Academic, Dordrecht, Netherlands.
- MATOUSEK, J., SHARIR, M., SIFRONY, S., AND WELZL, E. 1994. Fat triangles determine linearly many holes. *SIAM J. Comput.* 23, 1 (Feb. 1994), 154–169.
- MCMILLAN, L. 1997. An image-based approach to three-dimensional computer graphics. Ph.D. Dissertation. University of North Carolina at Chapel Hill, Chapel Hill, NC.
- MEHLHORN, K. 1984. *Data Structures and Algorithms 3: Multi-dimensional Searching and Computational Geometry*. EATCS monographs on theoretical computer science. Springer-Verlag, New York, NY.
- MOORE, M. AND WILHELMS, J. 1988. Collision detection and response for computer animation. In *Proceedings of the 15th Annual Conference on Computer Graphics (SIGGRAPH '89, Atlanta, Ga, Aug. 1-5)*, R. J. Beach, Ed. ACM Press, New York, NY, 289–298.
- SCHÖMER, E. AND THIEL, C. 1995. Efficient collision detection for moving polyhedra. In *Proceedings of the 11th Annual Symposium on Computational Geometry* (Vancouver, B.C., Canada, June 5–12, 1995), J. Snoeyink, Ed. ACM Press, New York, NY, 51–60.
- VAN DER STAPPEN, A. F., HALPERIN, D., AND OVERMARS, M. H. 1993. Efficient algorithms for exact motion planning amidst fat obstacles. In *Proceedings of the IEEE International Conference on Robotics and Automation*, IEEE Computer Society, New York, NY, 297–304.
- ZHOU, Y. AND SURI, S. 1999. Analysis of a bounding box heuristic for object intersection. In *Proceedings of the Tenth Annual Symposium on Discrete Algorithms* (Baltimore, MD).

Received: June 1998; revised: July 1999; accepted: July 1999