

Medial Spheres for Shape Approximation

Svetlana Stolpner, Paul Kry, Kaleem Siddiqi

Abstract—We study the problem of approximating a 3D solid with a union of overlapping spheres. In comparison with a state-of-the-art approach, our method offers more than an order of magnitude speed-up and achieves a tighter approximation in terms of volume difference with the original solid, while using fewer spheres. The spheres generated by our method are internal and tangent to the solid’s boundary, which permits an exact error analysis, fast updates under local feature size preserving deformation, and conservative dilation. We show that our dilated spheres offer superior time and error performance in approximate separation distance tests than the state-of-the-art method for sphere set approximation for the class of (σ, θ) -fat solids. We envision that our sphere-based approximation will also prove useful for a range of other applications, including shape matching and shape segmentation.

Index Terms—medial axis, shape approximation, sphere-based representations

I. INTRODUCTION

The choice of representation of 3D volumetric data is an important question in motion planning, solid modeling, computer vision, computer graphics, medical imaging and computer aided design. Whereas representations that approximate object boundaries with triangles are popular, representations of solids as unions of spheres are a valuable alternative. As argued in [1], such representations are robust to noise and resolution changes. Further, sphere-sphere distance tests are significantly faster than distance tests between other volumetric primitives.

In recent years, a number of applications have made use of sphere-based representations. For example, hierarchies of spheres are used for collision detection in [2] because sphere-sphere intersection tests are fast and simple. Sphere representations are used in the application of soft shadow generation [3], where a low-frequency object representation is sufficient. Additionally, sphere-based representations are used for efficient level-of-detail rendering [4], shape matching [5], and shape deformation [6].

In the above applications, it is desirable to be able to quickly generate an approximation of a solid that provides a tight fit using a small number of spheres. Finding the minimum number of spheres that covers a set of points on the object boundary is NP-hard by reduction from Set Cover [7]. For this reason, heuristics are used to find a small set of approximating spheres in practice.

A popular strategy for approximating a solid with a small number of well-fitting spheres is to start by computing the Voronoi spheres of a set of points on the solid’s boundary ([7], [8], [2]). The centres of these spheres, the Voronoi vertices, are inherently clustered near rounded corners because there

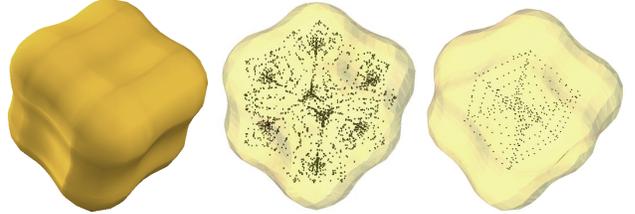


Fig. 1. A polyhedron (Left) and two distributions of sphere centres for a sphere-based approximation this object. (Centre) The subset of the Voronoi vertices of points sampled on the polyhedron boundary that lies near its medial surface. (Right) Sphere centres computed with our method.

are many Voronoi vertices equidistant from 4 boundary points at these locations (e.g. Figure 1(Centre)). Thus, each Voronoi sphere does not necessarily capture a unique salient feature of the polyhedron. Subsequently, to generate a small number of spheres offering a tight fit to a solid, such methods require an optimization step to remove and redistribute spheres, making them computationally expensive.

We have recently proposed a method to approximate the salient portions of the medial surface of a polyhedron using a well-spaced collection of points [9]. The spheres produced are internal and tangent to the boundary of the solid and there is at most one sphere centre per cubic region of space (cf. Figure 1(Right)). In this article we investigate the application of this method to the problem of quickly generating a small set of well-fitting spheres to a polyhedron. Compared to the state-of-the-art method for approximating solids with spheres, we show that our method is significantly faster and provides a tighter fit in terms of volumetric error. When a model undergoes local feature size preserving deformation, we show how our sphere approximation can be quickly updated and show how the volumetric error of the new sphere sets can be evaluated. We then use our sphere-based representation to compute approximate separation distance. To allow this application, we propose a method to improve the coverage of the solid’s boundary and describe how an efficient bounding volume hierarchy of the sphere sets can be built to accelerate distance tests. We show experimental results comparing the performance of our method with the leading method in sphere-based approximation for the computation of approximate separation distance.

II. BACKGROUND AND PREVIOUS WORK

We begin by providing some necessary definitions. Given a 3D solid Ω with boundary \mathcal{B} , the *medial surface* \mathcal{MS} of Ω is the locus of centres of maximal inscribed balls in Ω , called *medial points*. A maximal inscribed ball in Ω is called a *medial ball* of Ω , and its boundary is called a *medial sphere*. A very important property of medial representation is

The authors are with the School of Computer Science and the Centre for Intelligent Machines at McGill University, 3480 University Street, Montréal, PQ H3A 2A7, Canada. E-mail: {sveta,siddiqi}@cim.mcgill.ca,kry@cs.mcgill.ca.

that the union of all medial balls of Ω is Ω . Recently, many promising methods for the computation of medial balls have been proposed, examples of which include [10], [11] and [12]. For an overview of algorithms for the computation of medial representations as well as their properties please refer to [13].

When approximating a solid with spheres, a popular strategy is to start with a subset of the medial spheres of the solid. Voronoi spheres, defined below, are often used to approximate medial spheres.

Given a set of point sites $P = \{p_1, p_2, \dots, p_n\}$, $p_i \in \mathbb{R}^3$, the *Voronoi diagram* of P is a partition of \mathbb{R}^3 into Voronoi cells $V(p_i)$ with respect to the Euclidean distance d_E s.t.

$$V(p_i) = \{x \in \mathbb{R}^3 \mid d_E^2(x, p_i) \leq d_E^2(x, p_j), \forall j \neq i\}. \quad (1)$$

A vertex of $V(p_i)$ is called a *Voronoi vertex* and is the centre of a *Voronoi sphere* that touches 4 or more points of P , but does not contain any point of P in its interior.

When the points P are sampled on the boundary \mathcal{B} of a solid Ω , Amenta et al. [11] show that a subset of the Voronoi vertices of P , the *poles*, converges to the medial surface of Ω as the sampling density of P approaches infinity. As the union of medial balls of a solid reconstructs the solid, Voronoi spheres are proposed for shape approximation in [8] and [2].

An early method for the approximation of an object with spheres bounds each mesh triangle with a sphere but does not necessarily offer a tight fit [14]. Also, a number of methods have been proposed recently for the simplification of sphere representations ([7], [15]), but they do not explicitly aim to ensure tightness of fit. We now review existing methods that approximate a 3D solid with spheres with the goal of providing a tight approximation to the solid using a small number of spheres.

There are two methods which approximate an object Ω with tight-fitting spheres, starting with a set of Voronoi spheres. Hubbard [8] greedily selects adjacent Voronoi spheres for merging when their bounding sphere has the best tightness of fit. Bradshaw and O’Sullivan [2] improve the tightness of fit of Hubbard’s algorithm by using an adaptive greedy strategy that adds sample points to the boundary of Ω to generate new Voronoi spheres as needed. Tightness of fit is evaluated as the maximum distance between each sphere and Ω . As this quantity is difficult to compute exactly, an approximation is used.

Wang *et al.*[3] propose a variational optimization method that improves on the performance of [2] and is feasible for approximations having up to several hundred spheres. The error measure used is an approximation to total sphere volume outside Ω . Approximate volumetric error is also used in [16] to compare the performance of the methods of [8], [2], and an octree-based method for approximating deforming objects with spheres. In the present work, we also use a volume-based error measure.

III. COMPUTATION OF MEDIAL SPHERES

We now describe a modification of the method first introduced in [9] to approximate a solid Ω with a union of medial balls, such that the density of the centers of the medial balls is user-prescribed.

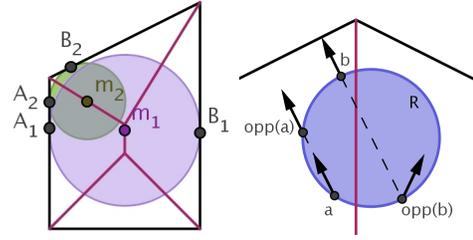


Fig. 2. In this 2D example, the boundary of an object is shown in black and its medial axis in red. Left, the object angle of the small circle with centre m_1 is greater than that of big circle with centre m_2 since $\angle A_1 m_1 B_1 > \angle A_2 m_2 B_2$. Right, arrows show the directions to nearest locations on the boundary of the object to points on R . In this example, the medial surface intersects the line segment $(b, opp(b))$ because $\mathcal{B}(b) \neq \mathcal{B}(opp(b))$.

Suppose that a point $m \in \mathcal{MS}$ is equidistant from two points $A, B \in \mathcal{B}$. Angle $\angle AmB$ (see Figure 2 left) is the *object angle*. The complete medial surface of a polyhedron can be a complicated structure and typically only a subset of the medial surface is sought. The object angle is a valuable simplification criterion for the medial surface ([10], [12]). Our goal is to locate a small number of medial balls such that the volume of their union approximates the volume of Ω well. As shown in [12], removal of medial balls having a small object angle has a small impact on the volume of the reconstructed object (refer to Figure 2, left). In order to reduce the number of spheres generated and to achieve a tight volumetric approximation, we will look for those medial points having a large object angle.

The method we use is based on the analysis of the nearest boundary points to a set of query points inside Ω . Let $\mathcal{B}(p)$ be the nearest point on \mathcal{B} to a point p . As shown in [17], using arguments similar to those in [18], the medial surface intersects a line segment (p, q) , where $q = p + \gamma(p - \mathcal{B}(p))$, iff $\mathcal{B}(p) \neq \mathcal{B}(q)$, for any scalar value of γ . Refer to Figure 2, right, for an illustration. This property is the basis for the algorithm we use for detecting medial points in convex regions.

We consider a regular partition of space into voxels with side length σ . For each voxel interior to or intersected by Ω , we circumscribe v with a sphere S and sample points on S . We then consider pairs of points (p, q) such that $q = p + \gamma(p - \mathcal{B}(p))$, and p and q both lie on S . For those pairs of points (p, q) that have different nearest boundary points, we perform binary search on the segment (p, q) to determine a location within a user-chosen tolerance ϵ of the medial surface on (p, q) . We also estimate the object angle of this approximate medial point, as described in [9], and discard points with low object angle estimates. We output at most one point within ϵ of the medial surface inside voxel v and inside Ω . This point is the centre of a sphere whose radius is the distance from this point to its nearest point on \mathcal{B} . When exact point-to-mesh distance computations are used, the computed spheres are internal and tangent to the boundary \mathcal{B} . To summarize, we produce a set of spheres that are interior and tangent to the solid’s boundary and whose centres are distributed such that at most one sphere centre lies in one voxel. In contrast with the method presented in [9], this method allows locating sphere centres in voxels intersected by the boundary \mathcal{B} in addition to

those completely inside Ω . Thus, the current method offers an improved fit. Using spheres that are interior and tangent to \mathcal{B} is essential for the error analysis in Section IV-A, the updates under deformations in Section V-A, and the sphere dilation process in Section VI-A.

IV. VOLUMETRIC ERROR FOR UNIONS OF BALLS

In this section, we evaluate the global quality of our sphere-based approximation to a polyhedron using a volumetric measure, and present comparative results against a leading method that approximates polyhedra with spheres.

A. Volumetric Error: Exact and Lower Bound

We evaluate tightness of fit of an approximation of Ω with a union of balls U as the volume of U outside Ω plus the volume of Ω outside U . Let $\text{vol}(\cdot)$ denote volume and \bar{A} be the compliment of a set A . Then

$$\text{Err}_\Omega(U) = \text{vol}(U \cap \bar{\Omega}) + \text{vol}(\Omega \cap \bar{U}) \quad (2)$$

is the error of the approximation of Ω with U . We state the following lemma.

Lemma 1: $|\text{vol}(U) - \text{vol}(\Omega)| \leq \text{Err}_\Omega(U)$ and when $\text{vol}(U \cap \bar{\Omega}) = 0$, $\text{vol}(\Omega) - \text{vol}(U) = \text{Err}_\Omega(U)$.

Proof: Observe that $\text{vol}(U) = \text{vol}(U \cap \Omega) + \text{vol}(U \cap \bar{\Omega})$ and $\text{vol}(\Omega) = \text{vol}(\Omega \cap U) + \text{vol}(\Omega \cap \bar{U})$. It follows that $\text{vol}(U) - \text{vol}(\Omega) = \text{vol}(U \cap \bar{\Omega}) - \text{vol}(\Omega \cap \bar{U}) \leq \text{Err}_\Omega(U)$. Likewise, $\text{vol}(\Omega) - \text{vol}(U) = \text{vol}(\Omega \cap \bar{U}) - \text{vol}(U \cap \bar{\Omega}) \leq \text{Err}_\Omega(U)$. If $\text{vol}(U \cap \bar{\Omega}) = 0$, $\text{vol}(\Omega) - \text{vol}(U) = \text{Err}_\Omega(U)$. ■

The next section explains how to compute $\text{vol}(U)$ exactly.

B. Unions of Balls: Tools

In this section we describe tools for the analysis of unions of balls, proposed in [19], [20], that we will use in this work. First, we define a special space filling diagram for a set of balls, called the *power diagram* [19].

Given a set of balls, $B = \{(c_1, r_1), \dots, (c_n, r_n)\}$ with centres $c_i \in \mathbb{R}^3$ and radii $r_i \in \mathbb{R}$, the *power diagram* of B , denoted $PD(B)$, is a partition of \mathbb{R}^3 into convex *power cells* $P(b_i)$, s.t.

$$P(b_i) = \{x \in \mathbb{R}^3 \mid d_E^2(x, c_i) - r_i^2 \leq d_E^2(x, c_j) - r_j^2, \forall j \neq i\}. \quad (3)$$

Note that power cells (Eq. 3) are identical to Voronoi cells (Eq. 1) when all the ball radii are the same. Figure 3 shows a 2D example of the power diagram of a set of disks.

We decompose a union of balls into convex cells by intersecting each ball with its power cell. Consider the dual complex $DC(B)$ of the decomposition. $DC(B)$ contains a vertex i for each sphere s_i , an edge (i, j) whenever balls b_i and b_j share a face of $PD(B)$, a triangle (i, j, k) whenever balls b_i, b_j, b_k share an edge of $PD(B)$, and tetrahedron (i, j, k, l) whenever balls b_i, b_j, b_k, b_l share a vertex of $PD(B)$. We will use the relationships captured by this complex to improve the surface coverage of a set of medial balls in Section VI-A.

As shown in [20], to find the total volume of a union of balls, $\text{vol}(\cup_i b_i)$, one need only consider the balls corresponding to vertices, edges, triangles and tetrahedra of $DC(B)$:

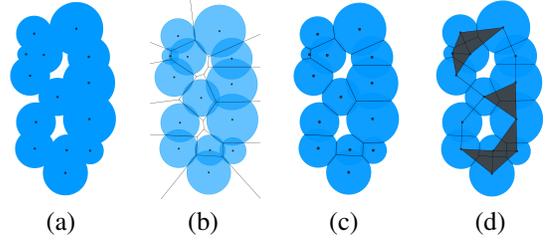


Fig. 3. (a) A union of a set of disks; (b) its power diagram overlaid; (c) the union of the set of disks decomposed using the power diagram; (d) the dual of this decomposition.

$\text{vol}(\cup_i b_i) = \sum_{i \in DC(B)} \text{vol}(b_i) - \sum_{(i,j) \in DC(B)} \text{vol}(b_i \cap b_j) + \sum_{(i,j,k) \in DC(B)} \text{vol}(b_i \cap b_j \cap b_k) - \sum_{(i,j,k,l) \in DC(B)} \text{vol}(b_i \cap b_j \cap b_k \cap b_l)$. This simple formula makes the computation of the volume of a union of balls efficient.

C. Experimental Results and Discussion

In this section, we compare the sphere-based approximations computed with our method and those computed using a leading method in terms of the volumetric error of the approximation. The method of Bradshaw and O’Sullivan [2] for approximating solids with spheres, which we will refer to as AMAA, is the state-of-the-art method for approximating solids tightly using several hundred spheres or more. The sphere approximations generated in [3] provide a tighter fit than those computed by AMAA, but the method is only feasible for generating small sphere sets (approximately 128 spheres). As our method is able to generate a large number of spheres quickly, we compare its performance in generating sphere sets of cardinality greater than 128 to that of AMAA.

Let S_D be the approximation to Ω computed with our distance-based method and let S_V be the approximation to Ω computed with the Voronoi-based AMAA method. Let B_D and B_V be the closed balls corresponding to the spheres S_D and S_V , respectively. Let U_D and U_V be the corresponding unions of balls. As explained in Section III, U_D is completely contained in Ω , i.e., $\text{vol}(U_D \cap \bar{\Omega}) = 0$. By Lemma 1, $\text{Err}(U_D) = \text{vol}(\Omega) - \text{vol}(U_D)$, while $\text{Err}(U_V) \geq |\text{vol}(U_V) - \text{vol}(\Omega)|$. Since in our experiments we approximate a variety of objects of different sizes, we will define the normalized error of the approximation offered by sphere set S to be $\text{nerr}(S) = \frac{\text{Err}_\Omega(U)}{\text{vol}(\Omega)}$, where U is the union of the balls corresponding to S .

We tabulate the exact error of S_D given by $\text{nerr}(S_D)$, and a lower bound on the error of S_V , $\text{nerr}(S_V)$, for 12 models of varying geometric complexity in Table I. Approximate volumetric error has been used in the literature ([3], [16]) to evaluate the quality of sphere-based approximations and provides a global measure of fit. AMAA construction proceeds top-down by building a hierarchy with a fixed branching factor and depth. In Table I, S_V are the leaves of an 8-ary hierarchy of depth 4 (i.e., maximum number of spheres generated is 512). In finding the appropriate set of spheres S_D , we found the largest voxel resolution such that $|S_D| < |S_V|$. Fast construction of a tight binary hierarchy for the spheres S_D is discussed in Section VI-C. Figure 4 shows error and

												
$ S_V =$	512	506	474	497	512	496	510	444	482	506	499	495
$ S_D =$	495	475	449	483	472	481	487	422	469	500	487	480
$\text{nerr}(S_V) \geq$	0.035	0.248	0.017	0.054	0.141	0.266	0.055	0.091	0.015	0.101	0.480	0.050
$\text{nerr}(S_D) =$	0.034	0.191	0.011	0.033	0.073	0.156	0.038	0.062	0.011	0.060	0.157	0.029
$\text{Time}(S_V) =$	2,342	3,212	4,049	2,672	4,936	1,930	22,845	2,469	7,428	2,870	3,739	9,479
$\text{Time}(S_D) =$	103	138	284	99	171	94	429	278	356	114	159	243
$\# \text{ triangles} =$	6,144	5,104	6,924	5,672	11,116	5,288	46,202	5,110	18,432	5,114	13,320	19,276

TABLE I

SPHERE SETS S_D ARE GENERATED WITH OUR METHOD; SPHERE SETS S_V ARE GENERATED WITH THE METHOD OF [2]. TIMINGS ARE IN SECONDS.

computation time as a function of the number of spheres generated for AMAA and for multiple invocations of our method until the appropriate sphere set was found. Timings are shown on a 3.6 GHz Pentium IV processor with 3 GB of RAM. As can be seen from Table I and Figure 4, our method generates a set of tighter fitting spheres significantly faster than AMAA. In generating our sphere sets, we use a threshold on object angle of 0.6 radians.

Our distribution of sphere centres, where at most one sphere centre is produced per voxel, is typically not an optimal distribution for minimizing the volumetric error. An optimal distribution of spheres minimizing volumetric error would include more large radius medial spheres than small radius medial spheres. In the case of the peanut and tooth models (columns 3 and 9 of Table I) our heuristic necessarily generates a suboptimal sphere centre distribution. However, for these models as well, we observe that our method’s performance in terms of volumetric error is superior to that of AMAA.

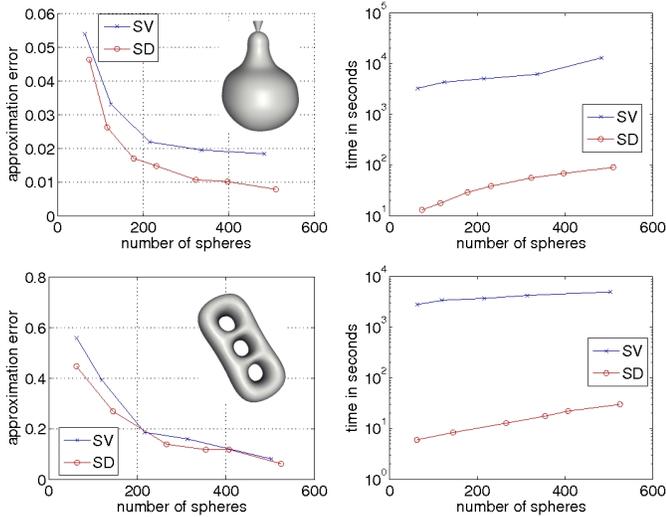


Fig. 4. Error and timing results for the pear (5154 triangles) and triple torus (16000 triangles) models. Note that true SV error may be higher than plotted.

V. FAST UPDATES UNDER DEFORMATION

In this section, we consider the case when Ω is a polyhedron whose triangulated boundary undergoes deformation. We devise an algorithm for quickly updating our spheres sets in this case. We also compute bounds on the volumetric error of the new approximation.

A. Update Method

Let \mathcal{B} be a triangle mesh boundary of Ω . Consider a deformed version of \mathcal{B} , \mathcal{B}' , bounding solid Ω' . We will show how to update the sphere set approximating Ω to build an approximation to Ω' . Consider the set of internal spheres tangent to Ω . Let $F : \mathcal{B} \rightarrow \mathcal{B}'$ be a one-to-one mapping of vertices in \mathcal{B} to those in \mathcal{B}' . For a medial sphere $s = (c, r)$ of Ω , we now explain how the position and radius of the corresponding sphere $s' = (c', r')$ in \mathcal{B}' may be estimated. Let A be a nearest point to c on \mathcal{B} . Let $N_{\mathcal{B}}(A)$ be the inner normal to \mathcal{B} at A . By construction,

$$c = A + rN_{\mathcal{B}}(A). \quad (4)$$

Suppose that A lies on $\Delta v_i v_j v_k$. Let $\lambda_i, \lambda_j, \lambda_k$ be the barycentric coordinates of A in $\Delta v_i v_j v_k$, such that $A = \lambda_i v_i + \lambda_j v_j + \lambda_k v_k$. Then if v'_i, v'_j, v'_k are the vertices of \mathcal{B}' corresponding to v_i, v_j, v_k , let sphere s' be tangent to \mathcal{B}' at $A' = \lambda_i v'_i + \lambda_j v'_j + \lambda_k v'_k$. If s' is a medial sphere, then

$$c' = A' + \gamma N_{\mathcal{B}'}(A'), \quad (5)$$

for some constant γ [21]. Given vertex normals, we estimate $N_{\mathcal{B}'}(A') = \lambda_i N_{\mathcal{B}'}(v'_i) + \lambda_j N_{\mathcal{B}'}(v'_j) + \lambda_k N_{\mathcal{B}'}(v'_k)$. We obtain an initial estimate for c' by letting $\gamma = r$.

Incorrect normal estimates or local shrinking of \mathcal{B}' cause sphere $s' = (c', r)$ to protrude outside \mathcal{B}' . This case is determined by checking the shortest distance from c' to \mathcal{B}' . When this distance is significantly less than γ , the value for γ is reduced until the sphere protrudes a user-chosen tolerable amount outside \mathcal{B}' . Algorithm 1 summarizes the procedure for shrinking spheres.

Algorithm 1 SHRINK($s = (c, r), \mathcal{B}, p, \xi$)

Input: Mesh boundary \mathcal{B} , point p of contact of sphere s with \mathcal{B} , threshold $\xi > 0$.

Output: A sphere that protrudes at most ξ outside \mathcal{B} .

- 1: Let d be the distance from c to \mathcal{B} .
 - 2: **while** $r - d \geq \xi$ **do**
 - 3: Let $r = (r + d)/2$.
 - 4: Let $c = p + r(c - p)/\|c - p\|_2$.
 - 5: Let d be the distance from c to \mathcal{B} .
 - 6: **end while**
 - 7: Return (c, d) .
-

B. Volumetric Error: Upper Bound

After using Algorithm 1 to ensure that no deformed sphere protrudes more than ξ outside the mesh boundary \mathcal{B} , we can find an upper bound for the volumetric error of the new sphere set.

Lemma 2: Suppose that for each sphere $(c_i, r_i) \in S$, $d_E(c_i, \mathcal{B}) > r_i - \xi$. Let U be the union of the balls corresponding to S and let U^ξ be the union of the balls corresponding to $S^\xi = \{(c_i, r_i - \xi) | (c_i, r_i) \in S\}$. Then $\text{Err}_\Omega(U) < \text{vol}(\Omega) - \text{vol}(U) + 2(\text{vol}(U) - \text{vol}(U^\xi))$.

Proof: As shown in the proof of Lemma 1, $\text{vol}(\Omega) - \text{vol}(U) = \text{vol}(\Omega \cap \bar{U}) - \text{vol}(U \cap \bar{\Omega})$. Further, recall that $\text{Err}_\Omega(U) = \text{vol}(\Omega \cap \bar{U}) + \text{vol}(U \cap \bar{\Omega})$. It follows that

$$\text{Err}_\Omega(U) = \text{vol}(\Omega) - \text{vol}(U) + 2\text{vol}(U \cap \bar{\Omega}). \quad (6)$$

Let us express $\text{vol}(U) = \text{vol}(U \cap \bar{\Omega}) + \text{vol}(U \cap \Omega)$ and $\text{vol}(U^\xi) = \text{vol}(U^\xi \cap \bar{\Omega}) + \text{vol}(U^\xi \cap \Omega)$. Because $\text{vol}(U^\xi \cap \bar{\Omega}) = 0$, $\text{vol}(U) - \text{vol}(U^\xi) = \text{vol}(U \cap \bar{\Omega}) + \text{vol}(U \cap \Omega) - \text{vol}(U^\xi \cap \Omega)$. Since $\text{vol}(U \cap \Omega) - \text{vol}(U^\xi \cap \Omega) > 0$, $\text{vol}(U \cap \bar{\Omega}) < \text{vol}(U) - \text{vol}(U^\xi)$. Combining with Eq. 6, the lemma follows. ■

Next, we will use this bound to evaluate the quality of our sphere-based approximation to deformed polyhedra.

C. Experimental Results and Discussion

Figure 5 shows times to compute and update a reference set of spheres for two deformation sequences. Note that this fast update under deformations is possible because we use spheres that are interior and tangent to the solid's boundary. Using Lemmas 1 and 2, we compute lower and upper bounds on the volumetric error of our approximations.

Following [11], the *local feature size* of a point A on the boundary of Ω is the distance from A to the medial surface of Ω . Note that our approach will not produce tight sphere-based approximations for deformations that modify the local feature size of a solid, as well as deformations that stretch the boundary. However, we have observed that the method produces tight-fitting sphere-based approximations at a fraction of the time necessary to compute the original set of spheres for those deformations arising from part articulation (such as the horse sequence) and volume preserving elastic deformations (such as the octopus sequence).

VI. APPROXIMATE SEPARATION DISTANCE

For those solids which can be tightly approximated using a small number of spheres relative to the number of surface triangles, separation distance computations can be accelerated by working with the sphere representation instead of the boundary mesh representation. Further, pairwise distance tests are significantly faster between spheres than triangles. In this section, we show how to quickly grow our set of inscribed spheres so as to improve boundary coverage, and how to construct a bounding volume hierarchy of the set of spheres. We introduce a class of solids, called (σ, θ) -fat solid. We then demonstrate experimentally that spheres generated with our method perform faster and with smaller error than the state-of-the-art method AMAA of [2] for approximate separation distance computations for (σ, θ) -fat solids.

A. Improving Boundary Coverage by Conservative Dilation

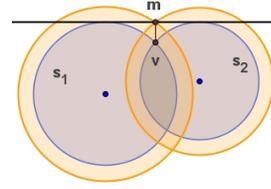


Fig. 6. The envelope of the dark circles s_1 and s_2 is non-differentiable at point v , whose nearest point on the boundary is m . The radii of both internal circles are increased to create circles that contain or touch m .

In Section IV, we aimed to approximate the volume of the solid Ω tightly. For the application of approximate separation distance, we now require the envelope of our set of spheres to provide a tight approximation to the boundary of the solid. By *improving boundary coverage*, we mean to increase the number of boundary points that lie inside some sphere. It is no longer essential that the approximating spheres be interior. We introduce a simple heuristic for conservatively growing the set of internal spheres to improve boundary coverage.

The envelope of the spheres S is a non-differentiable object. Let B be the balls corresponding to S , *i.e.*, $B = \text{cl}(\text{int}(S))$. Recall that the dual complex $DC(B)$ describes the adjacency of balls in the union of balls B . Edges in $DC(B)$ identify pairs of balls intersecting along circular arcs on the envelope of S , while triangles identify triples of spheres intersecting at vertices on the envelope of S .

We can solve for the locations of the vertices of the envelope of the spheres S , called *v-points*, by considering all triangles (b_i, b_j, b_k) in $DC(B)$, finding the 2 intersection points of spheres s_i, s_j and s_k , and ignoring those intersection points that lie inside some spheres of S . We note which triple of spheres contributed to the creation of a *v-point*. For a *v-point* v , we consider the nearest point on the boundary of Ω to v . The radius of each sphere s is increased sufficiently so as to cover the nearest mesh point to each of the *v-points* that is created by s . Figure 6 presents a 2D example. For those singleton edges of $DC(B)$, identifying pairs of balls that touch but do not contribute any *v-points*, we sample 2 random opposite points on the circle of intersection of the two corresponding spheres and proceed to cover their nearest boundary points by growing the spheres.

Let S_D^+ be the sphere-based approximation obtained by growing internal spheres S_D as described. Let S_V be the approximation produced by the AMAA method of [2]. We compare how well the envelope of each sphere set approximates the boundary of a polyhedron by evaluating the signed distance from 1) points on the envelope of the spheres to the polyhedron boundary, and 2) from points on the polyhedron boundary to the envelope of the spheres. Figure 7 shows histograms of signed distances for several models, where positive distance means that the point on the sphere envelope giving the distance measurement is outside the polyhedron, and negative otherwise.

As we can see from the plots in Figure 7, the heuristic for growing internal spheres S_D generates spheres S_D^+ whose

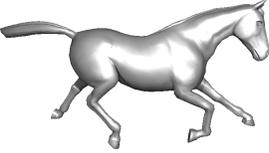
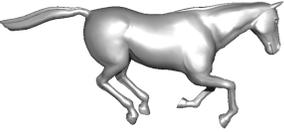
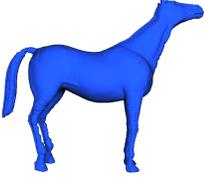
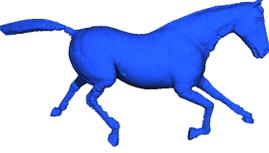
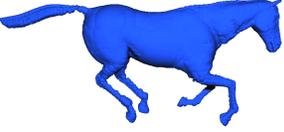
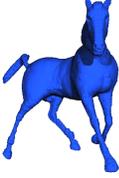
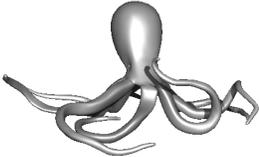
16857 triangles				
4307 spheres				
rel. speed	100%	4.04%	4.10%	4.08%
vol. error	[0.0110, 0.0110]	[0.0331, 0.0473]	[0.0362, 0.0509]	[0.0370, 0.0516]
39632 triangles				
9801 spheres				
rel. speed	100%	19.5%	19.9%	19.6%
vol. error	[0.0211, 0.0211]	[0.0256, 0.0585]	[0.0253, 0.0583]	[0.0252, 0.0579]
	(a)	(b)	(c)	(d)

Fig. 5. (a) Reference mesh (top) approximated using internal and tangent spheres (bottom). (b)-(d) These sphere sets (bottom) approximate deformed versions of the reference mesh (top). The parameter ξ is 0.25% of the maximum dimension of the reference mesh bounding box. The reference set of spheres took 146.8 and 226.6 seconds to compute for the horse and octopus reference poses, respectively. The relative speed is the fraction of the time to compute the sphere set compared to computing the reference sphere set. Also shown is the range of the volumetric error normalized by volume of the polyhedra, where the lower bound for volumetric error is given by Lemma 1 and the upper bound by Lemma 2.

envelope covers a significant portion of the boundary of the polyhedra, without generating spheres that protrude a great amount outside the polyhedra, as is often the case with S_V spheres. The mean absolute error for S_D^+ is smaller than that for S_V , except for the last model, where both means are small. Note that since neither sphere set is bounding, neither the envelope of our sphere sets S_D^+ , nor that of S_V , contain the polyhedra. As can be seen from Figure 7, the sphere sets S_D^+ capture the shape of the polyhedra more faithfully than the S_V spheres, in particular for the duck, bunny and octopus models.

Because the spheres S_D are internal and tangent to the boundary of Ω , points on the envelope of S_D are inside or on the boundary of Ω . Using the method we have just described, boundary coverage is improved in the vicinity of points on the envelope of S_D . This heuristic does not consider the geometry of Ω explicitly. Whenever the (one-sided) Hausdorff distance from the boundary of Ω to the set of spheres is small, this strategy works well to improve boundary coverage. This strategy also offers the advantage of improving the fit of the envelope of spheres without an expensive optimization procedure.

B. (σ, θ) -fat Solids

In this section, we characterize the class of solids whose sphere-based approximations computed using the method we have proposed are connected. Consider the subset of the medial surface containing only points with object angle greater or equal to θ , \mathcal{MS}_θ . Let σ be the side length of the voxels used to generate the approximation to the medial surface. Let S be the subset of medial spheres of Ω , such that each sphere has object angle greater than θ , and such that each voxel with side length σ containing a medial point with object angle greater or equal to θ contributes a medial sphere to S . We say that a solid Ω is (σ, θ) -fat with respect to given values of σ and θ if $\sigma \leq \frac{r_0}{\sqrt{3}}$ for r_0 the smallest medial sphere radius for medial points in \mathcal{MS}_θ . Solids that are (σ, θ) -fat have the following property:

Lemma 3: Suppose that \mathcal{MS}_θ of Ω is connected. If Ω is (σ, θ) -fat, then the union of the spheres in S is connected.

Proof: Consider a pair of adjacent voxels with side length σ , each containing medial points in \mathcal{MS}_θ . The maximum distance between any pair of points in these voxels is $2\sqrt{3}\sigma$. If $r \geq \sqrt{3}\sigma$ for all radii r of medial points in the two voxels considered, then the two spheres touch. Since Ω is (σ, θ) -fat, medial spheres in adjacent voxels always touch. Since the set

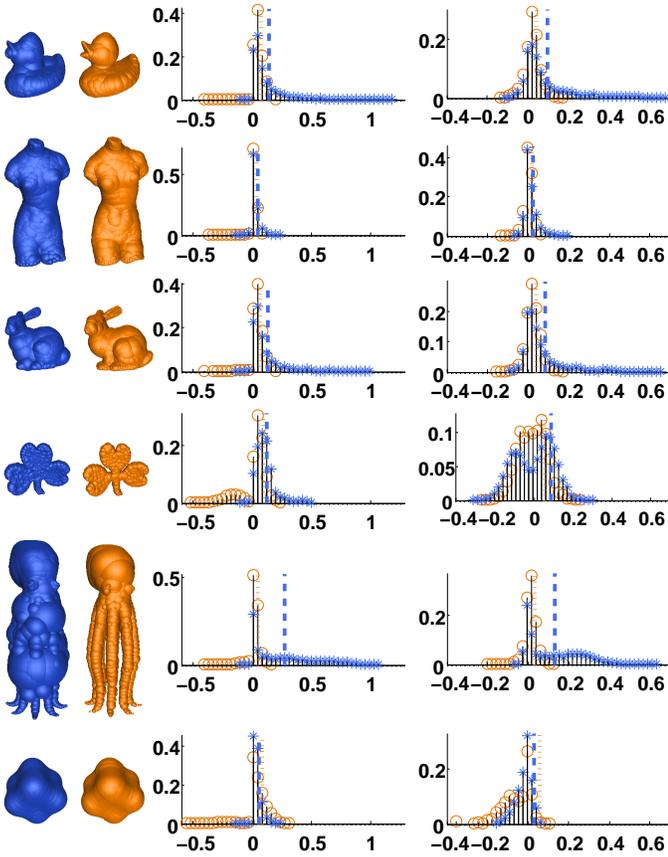


Fig. 7. Comparison of the S_V (col. 1) and S_D^+ (col. 2) approximations for various models. Col. 3: histograms of signed distances from points sampled on the boundary of polyhedra to points on the envelopes of S_D^+ and S_V . Col. 4: histograms of signed distances from points sampled on meshes approximating the envelopes of S_D^+ and S_V to the boundary of polyhedra. Stars denote S_V spheres, while circles denote S_D^+ spheres. Means of the absolute distance are shown as a dashed line for S_V and a dotted line for S_D^+ .

of voxels considered is connected, the union of the spheres in S is connected. ■

Solids that are (σ, θ) -fat do not have sharp corners or narrow parts with respect to the voxel resolution. The boundaries of those solids that are (σ, θ) -fat can be better approximated using our sphere sets than other solids.

C. Hierarchy Construction using Rectangle Swept Spheres

In order to use our set of spheres S_D^+ to perform fast proximity queries, we fit a bounding volume hierarchy to S_D^+ . We observe that the medial surface simplified by object angle is often composed of relatively flat sheets, along which the radius of the medial spheres varies smoothly. A Rectangle Swept Sphere (RSS) is therefore a suitable bounding volume for medial spheres (see Figure 8). Such a bounding volume was introduced for bounding mesh triangles in the exact proximity query package PQP [22]. The hierarchy is constructed top-down, each parent has 2 children, and each RSS is fitted to the set of leaf spheres it bounds. We find the orientation of the RSS rectangle and the partitioning plane by using CGAL’s



Fig. 8. RSS

implementation of linear least square fitting of spheres.¹ The radius is found by using a procedure similar to that in [22].

Alternatively, one may compute a bounding volume hierarchy using spheres as the bounding volume. Given our sphere set approximation, we compute the hierarchy using a similar top-down procedure as when using RSSs as the bounding volume. To compute the minimum bounding sphere of a set of spheres, we use CGAL’s implementation of the algorithm in [23].

D. Experimental Results and Discussion

We evaluate the usefulness of our sphere approximations for measuring approximate separation distance in terms of computation time and accuracy of the computations. In our experiments, we perform a physics simulation by placing pairs of the same object in a box, simulating random gravity, allowing objects to continuously collide inside the box, and measuring separation distance using both the sphere sets and the triangulated boundaries at each frame.

We compare the performance of separation distance tests for the set of spheres computed with our method and that computed with the AMAA method of [2]. When generating spheres for this application, we subdivided those voxels that are intersected by the boundary of Ω into 8 voxels and thus, output at most 8 medial spheres for boundary voxels. The threshold on object angle is 0.6 radians. We grow our initial sphere set S_D using the strategy outlined in Section VI-A to create a new sphere set S_D^+ . This step takes 1-5 seconds for the models in Table II. We compute both an RSS hierarchy and a sphere hierarchy of our set of spheres using the method described in Section VI-C. AMAA constructs a bounding volume hierarchy where bounding elements are spheres. Timings for AMAA in Table I include time for building a hierarchy with branching factor 8 (as AMAA is a top-down approach). In proximity query experiments, we consider hierarchies with branching factor 2. Computation of a binary AMAA hierarchy for the models shown in Table II takes significantly longer: from 3 hours (for the eight model) to 14 hours (for the octopus model). In addition to a binary sphere hierarchy, we also construct an RSS hierarchy of the AMAA spheres. Constructing an RSS hierarchy for the sphere sets considered takes a fraction of a second on average.

We evaluate the accuracy of a separation distance test as the difference between the distance between two polyhedra and the distance between their approximations with spheres. As in each frame of the simulation we consider pairs of non-intersecting solids at a variety of separation distances and relative orientations, this error provides a meaningful measure of the quality of the sphere set approximation for this task. Table II presents error statistics for the different sphere sets. Our sphere sets display a smaller average error and generally smaller maximum error than those of AMAA. The models in Table I that are not represented in Table II are the horse, hippo, dragon and cow models. These models are not (σ, θ) -fat with respect to the voxel resolution σ and the object angle θ considered because they have very narrow

¹Computational Geometry Algorithms Library, www.cgal.org.

regions, such as their ears and tails. The simulation we ran is particularly demanding of the quality of the approximation of such extremal regions as they minimize the separation distance more often than other regions. To approximate these models well using spheres by using our strategy, a very small value of σ needs to be used, resulting in a large number of spheres. We were not able to generate such large AMAA sphere sets for these models to draw comparative results.

We compare average per-frame time for performing approximate separation distance tests using S_D^+ and S_V , where both sphere hierarchies (SH) and RSS hierarchies (RSSH) are considered. These results are summarized in Table II. Timings are shown for a 3.6 GHz Pentium 4 CPU with 3 GB of RAM. In our experiments, we find that building an RSS hierarchy of spheres significantly improves distance query time compared to using a sphere hierarchy for both our and AMAA sphere sets, as the RSS offers greater tightness of fit. For these models, we see that the fastest performance is achieved by using our sphere sets and the RSS hierarchy. We believe that because our spheres are nearly medial, RSS bounding volumes provide them with a particularly tight fit. For the models shown, the fastest performance is achieved by using our sphere sets and the RSS hierarchy, even when $|S_D^+| > |S_V|$.

	Size	Error		Timings		
		Ave.	Max.	SH	RSSH	
	S_D^+	520	0.012	0.074	0.524	0.085
	S_V	498	0.024	0.063	0.268	0.088
	S_D^+	558	0.074	0.228	0.467	0.131
	S_V	397	0.094	0.331	0.627	0.132
	S_D^+	1052	0.026	0.102	0.672	0.134
	S_V	831	0.054	0.144	0.206	0.140
	S_D^+	296	0.009	0.024	0.385	0.066
	S_V	439	0.020	0.075	0.206	0.123
	S_D^+	695	0.116	0.297	0.709	0.143
	S_V	697	0.118	0.599	0.239	0.145
	S_D^+	1389	0.081	0.272	0.373	0.074
	S_V	745	0.109	0.4086	0.130	0.075
	S_D^+	548	0.055	0.173	1.458	0.108
	S_V	579	0.110	0.597	0.250	0.114
	S_D^+	469	0.076	0.325	0.360	0.068
	S_V	502	0.167	0.425	0.123	0.076

TABLE II

TIMING AND ERROR RESULTS FOR SEPARATION DISTANCE TESTS FOR OUR GROWN SPHERES S_D^+ AND SPHERES S_V OF [2]. TIMINGS ARE IN MILLISECONDS. MAXIMUM DIMENSION OF BOUNDING BOXES IS 10.

VII. CONCLUSION

We have described a method to compute a tight-fitting union of spheres approximation to a 3D solid without an expensive sphere redistribution or pruning step. As such, our method is significantly faster than existing methods and can be used to generate sphere set approximations with a larger number of spheres than previously possible. In comparison with the state-of-the-art AMAA method of [2], our method generates fewer spheres, has a smaller volumetric error, and is significantly faster. Because our spheres are internal and tangent to the object, it is possible to compare the volumetric error of approximations produced by our method and another

method, to quickly update our sphere sets when the object deforms, and to quickly dilate the sphere sets to improve boundary coverage. We demonstrate the benefit of using the rectangle swept sphere bounding volume for building a hierarchy of medial spheres. For (σ, θ) -fat solids, we have presented experimental results showing that our sphere sets perform faster and more accurately than those of the AMAA method for approximate separation distance computation. There is a wide variety of additional applications where our sphere approximations can be valuable, that do not necessarily require the boundary of the object to be well approximated, including shape matching, mesh deformation, shape morphing, shape segmentation, approximate Minkowski sums, point location and shadow rendering.

REFERENCES

- [1] V. Ranjan and A. Fournier, "Volume models for volumetric data," *Computer*, vol. 27, no. 7, pp. 28–36, 1994.
- [2] G. Bradshaw and C. O'Sullivan, "Adaptive medial-axis approximation for sphere-tree construction," *ACM Trans. Graph.*, vol. 23, no. 1, 2004.
- [3] R. Wang, K. Zhou, J. Snyder, X. Liu, H. Bao, Q. Peng, and B. Guo, "Variational sphere set approximation for solid objects," *Visual Computer*, vol. 22, no. 9, pp. 612–621, 2006.
- [4] S. Rusinkiewicz and M. Levoy, "QSplat: A multiresolution point rendering system for large meshes," in *SIGGRAPH*, 2000, pp. 343–352.
- [5] A. Shamir, A. Sotzio, and D. Cohen-Or, "Enhanced hierarchical shape matching for shape transformation," *Journ. of Shape Modeling*, vol. 9, no. 2, pp. 203–222, 2003.
- [6] S. Yoshizawa, A. Belyaev, and H.-P. Seidel, "Skeleton-based variational mesh deformations," *Comput. Graph. Forum*, vol. 26, no. 3, pp. 255–264, 2007.
- [7] O. Aichholzer, F. Aurenhammer, B. Kornberger, S. Plantinga, G. Rote, A. Sturm, and G. Vegter, "Recovering structure from r -sampled objects," *Computer Graphics Forum*, vol. 28, no. 5, pp. 1349–1360, 2009.
- [8] P. M. Hubbard, "Approximating polyhedra with spheres for time-critical collision detection," *ACM Trans. Graph.*, vol. 15, no. 3, 1996.
- [9] S. Stolpner, S. Whitesides, and K. Siddiqi, "Sampled medial loci for shape representation," *Computer Vision and Image Understanding*, 2011.
- [10] T. Dey and W. Zhao, "Approximating the medial axis from the Voronoi diagram with a convergence guarantee," *Algorithmica*, vol. 38, no. 1, pp. 179–200, 2003.
- [11] N. Amenta, S. Choi, and R. Kolluri, "The Power Crust, Unions of Balls, and the Medial Axis Transform," *Computational Geometry: Theory and Applications*, vol. 19, no. 2-3, pp. 127–153, 2001.
- [12] M. Foskey, M. C. Lin, and D. Manocha, "Efficient computation of a simplified medial axis," in *SMA*, 2003, pp. 96–107.
- [13] K. Siddiqi and S. Pizer, Eds., *Medial representations: mathematics, algorithms and applications*. Springer, 2008.
- [14] S. Quinlan, "Efficient distance computations between non-convex objects," in *ICRA*, 1994, pp. 3324–3329.
- [15] B. Miklos, J. Giesen, and M. Pauly, "Discrete scale axis representations for 3d geometry," *ACM Trans. Graph.*, vol. 29, July 2010.
- [16] M. Garcia, S. Bayona, P. Toharia, and C. Mendoza, "Comparing sphere-tree generators and hierarchy updates for deformable objects collision detection," in *Intern. Symp. Visual Computing*, 2005, pp. 167–174.
- [17] S. Stolpner, "Medial spheres for shape representation," Ph.D. dissertation, McGill University, 2011.
- [18] S. Stolpner and S. Whitesides, "Medial axis approximation with bounded error," in *Intern. Symp. Voronoi Diagrams*, 2009, pp. 171–180.
- [19] F. Aurenhammer, "Power diagrams: properties, algorithms and applications," *SIAM Journal of Computing*, vol. 16, no. 1, pp. 78–96, 1987.
- [20] H. Edelsbrunner, "The union of balls and its dual shape," in *Symposium on Computational geometry*, 1993, pp. 218–231.
- [21] E. Sherbrooke, N. Patrikalakis, and F. Wolter, "Differential and topological properties of medial axis transforms," *Graphical Models and Image Processing*, vol. 58, no. 6, pp. 574–592, 1996.
- [22] E. Larsen, S. Gottschalk, M. C. Lin, and D. Manocha, "Fast proximity queries with swept sphere volumes," University of North Carolina at Chapel Hill, Tech. Rep., 1999.
- [23] B. Gärtner, "Fast and robust smallest enclosing balls," in *European Symposium on Algorithms*, 1999, pp. 325–338.