Spheres, Molecules, and Hidden Surface Removal*

Dan Halperin† Mark H. Overmars‡

Abstract

We devise techniques to manipulate a collection of loosely interpenetrating spheres in three-dimensional space. Our study is motivated by the representation and manipulation of molecular configurations, modeled by a collection of spheres. We analyze the sphere model and point to its favorable properties that make it more easy to manipulate than an arbitrary collection of spheres. For this special sphere model we present efficient algorithms for computing its union boundary and for hidden surface removal. The efficiency and practicality of our approach are demonstrated by experiments on actual protein data.

1 Introduction

In this paper we devise techniques to represent and manipulate a collection of overlapping spheres in three-dimensional space. Often, manipulating three-dimensional geometric objects that intersect is time consuming. However, by imposing several constraints on the spheres and their interaction, we obtain a setting where efficient and practical techniques are quite easy to obtain. Our study is motivated by the representation and manipulation of molecular configurations, modeled by a collection of spheres.

A common approach to representing the three-dimensional geometric structure of a molecule, is to represent each of its atoms by a "hard" sphere. It is also common to assume that the *nuclear arrangement*, i.e., the relative displacement of the spheres, is fixed (it is often the so-called equilibrium nuclear configuration). There are recommended values for the radius of each atom sphere and for the distance between the centers of every pair of spheres. In this model, the spheres are allowed to interpenetrate one another, therefore it is sometimes referred to as the "fused spheres" model (see, e.g., Figure 1). The envelope surface of the fused spheres may be regarded as a formal *molecular surface*. It is evident, that various properties of molecules are disregarded in this simple model. However, in spite of its approximate nature, it has proven useful in many practical applications. For more background material and references, see, e.g., the survey paper by Mezey [16].

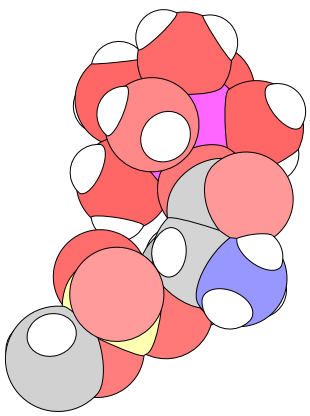

Figure 1: Molecule 1 with 42 atoms.

We study the hard sphere model from a *computational geometry* point of view, that is, we study the combinatorial and algorithmic behavior of a collection of n (possible intersecting) spheres in 3-space, having some special properties. We make several simple observation, showing that, because of these special properties, the spheres in this model can be effi-

*The first author acknowledges support by a Rothschild Postdoctoral Fellowship, by a grant from the Stanford Integrated Manufacturing Association (SIMA), and by NSF/ARPA Research Grant IRI-9306544. The second author was partially supported by the ESPRIT III Basic Research Action No. 7141 (project ALCOM II) and by the Netherlands Organization for Scientific Research (N.W.O.).

†Robotics Laboratory, Department of Computer Science, Stanford University, Stanford, CA 94305, USA. Email: halperin@cs.stanford.edu.

‡Department of Computer Science, Utrecht University, P.O.Box 80.089, 3508 TB, Utrecht, the Netherlands. Email: markov@cs.ruu.nl.

ciently manipulated. For example, we show that the maximum combinatorial complexity of the boundary of the union of the fused spheres is $O(n)$, whereas for an arbitrary collection of spheres (even unit spheres) this may be $\Theta(n^2)$ (see, e.g., [13]). These results are described in Section 2.

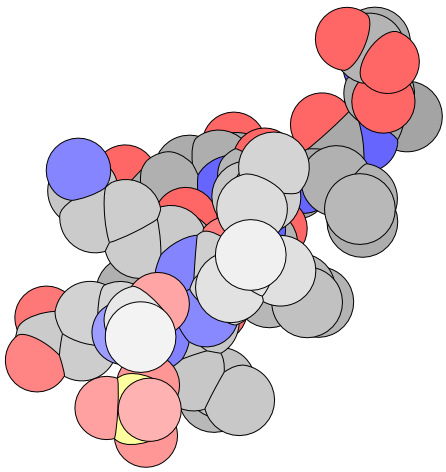

Figure 2: Molecule 2 with 100 atoms.

In Section 3 we take advantage of the favorable behavior of the model, to devise a data structure for answering *intersection queries* of the form: Given a hard sphere model M of a molecule and a query atom sphere Q (both given in a fixed placement in 3-space), report which spheres of M are intersected by Q. In Section 4 we use this data structure to efficiently compute the boundary of the union of the spheres. We also show that this algorithm can be used to compute a related type of molecular surfaces, so-called *solvent accessible surfaces*. Finally, in Section 5, we present efficient algorithms for hidden surface removal, either by computing a depth order or by computing the visibility map of a molecule.

Our work is closely related to the study of *fatness*. It has been shown that certain problems in computational geometry can be solved much more efficiently when the objects involved have no long and thin parts (see, e.g., [15, 17, 20]). Clearly, spheres are fat. In 3-space though fatness is not enough to guarantee efficient algorithms when the objects are allowed to intersect. Fortunately, the extra properties in the hard sphere model provide enough additional constraints on the objects to allow for efficient algorithms. It exemplifies, yet another time, that in practical situations a collection of geometric objects having certain additional properties (obtained from observing real-life situations), behaves favorably.

We support our claim for efficiency and practicality of the data structure and the algorithms that we present by reporting results of experiments that were carried out on actual protein data retrieved from the "Protein Data Bank" at Brookhaven National Laboratory [1, 2]. In Figures 1 till 4 you find pictures of some of the molecules that we used. All of these pictures were created using the algorithms described in Section 5.

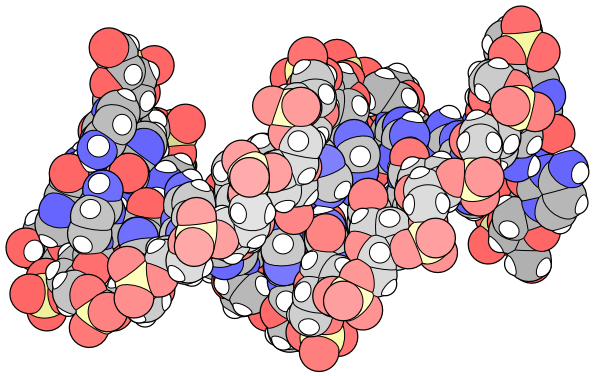

Figure 3: Molecule 3 with 757 atoms.

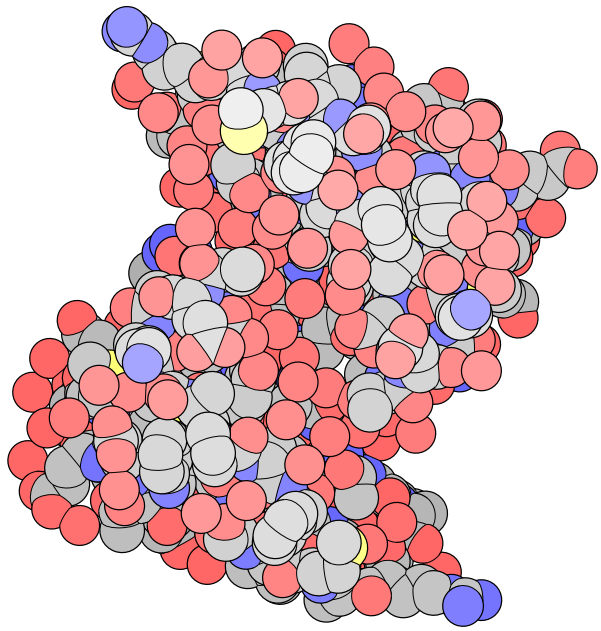

Figure 4: Molecule 4 with 971 atoms.

2 The Hard Sphere Model

Overlapping spheres in three-dimensional space may in general be rather unwieldy objects. The *arrangement* (i.e., the subdivision of space) defined by n spheres in 3-space may have complexity $\Theta(n^3)$ in the worst case and their union boundary may have complexity $\Theta(n^2)$. Moreover, efficient algorithms for manipulating a collection of overlapping spheres are often complex and rather time consuming. However, the atom spheres in the model of a molecule have properties that we can exploit to obtain efficient and simple algorithms for manipulating them. One such property is that, although two spheres may interpenetrate, their centers cannot get too close to one another. The other useful property is that their radii range in a fairly restricted range; see, e.g., [11, p. J-3] for a list of van der Waals radii, which are one type of acceptable radii. Actually, each molecular modelling package seems to use its own slightly different set of radii. The following table shows the list the radii (in Ångstrøm) of the spheres we use to represent the main types of atoms. (Taken from an actual molecular modelling package.)

C	Cal	H	N	O	P	S
1.70	2.07	0.90	1.55	1.52	1.80	1.80

In Theorem 2.1 below, we state the conditions that make the sphere model of a molecule favorable. From this point on, we will distinguish between a *ball* and a *sphere*. Each atom with radius r_i and center at c_i induces a ball $B_i = \{p | d(p, c_i) \leq r_i\}$ and a sphere $S_i = \{p | d(p, c_i) = r_i\}$, where $d(p_1, p_2)$ is the Euclidean distance between the points p_1 and p_2 in 3-space.

Theorem 2.1 *Let $M = \{B_1, \ldots, B_n\}$ be a collection of n balls in 3-space with radii $r_1, \ldots, r_n$ and centers at $c_1, \ldots, c_n$. Let $r_{\min} = \min_i r_i$ and let $r_{\max} = \max_i r_i$. Also let $S = \{S_1, \ldots, S_n\}$ be the collection of spheres such that S_i is the boundary surface of B_i. If there are positive constants k, ρ such that $\frac{r_{\max}}{r_{\min}} < k$ and for each B_i the ball with radius $\rho \cdot r_i$ and concentric with B_i does not contain the center of any other ball in M (besides c_i), then:*

(i) *For each $B_i \in M$, the maximum number of balls in M that intersect it is bounded by a constant.*

(ii) *The maximum combinatorial complexity of the boundary of the union of the balls in M is $O(n)$.*

Proof. Part (i): Consider a ball B_i. Let B be the ball of radius $r_i + 2r_{\max}$, centered at c_i. Clearly, any ball in M that intersects B_i must lie completely inside B. For each B_j that lies completely inside B, let β_j denote the ball of radius $\frac{\rho}{2} \cdot r_{\min}$ centered at c_j. We claim that the β_j's are pairwise interior-disjoint. Suppose the opposite, namely, that there are two such balls β_{j_1} and β_{j_2} whose interiors intersect. This implies that c_{j_1} and c_{j_2} are less than $\rho \cdot r_{\min}$ apart. This, in turn, means that the ball with radius $\rho \cdot r_{j_1}$ and center at c_{j_1} contains the center of the ball B_{j_2}, contradicting the assumptions of the theorem.

Hence, by volume consideration, the total number of balls that are completely contained in B cannot exceed

$$\frac{(r_i + 2r_{\max})^3}{(\frac{\rho}{2} \cdot r_{\min})^3} \leq \frac{(3r_{\max})^3}{(\frac{\rho}{2} \cdot r_{\min})^3} \leq 216 \cdot (\frac{k}{\rho})^3.$$

Part (ii) follows from Part (i) because Part (i) implies that the number of features involving B_i on the union boundary is bounded by a constant. Indeed, if the balls are in general position then a feature involving a ball B_i is either a face on the boundary of B_i, or an intersection arc of B_i with another ball, or an intersection point of B_i with two other balls. All these features belong to the two-dimensional subdivision (or *arrangement*) formed on the boundary of B_i by a constant number of circles, each being the intersection of B_i with another ball. The complexity of this arrangement is evidently bounded by a constant. Thus, the overall complexity of the union boundary is $O(n)$. □

The following table gives the values of k, ρ and the maximal and average number of balls intersecting a single ball for our four example molecules (Figures 1 to 4).

mol.	k	ρ	max	aver.
1	2.30	0.64	14	5.0
2	1.18	0.71	12	6.4
3	2.00	0.63	15	7.4
4	1.18	0.71	12	6.3

As can be seen k is small and ρ is large as required, resulting in a small number of intersections per ball (being much smaller than the worst-case bound of $216(k/\rho)^3$).

We now turn to discuss the subdivision of three-dimensional space induced by a collection of spheres; the *arrangement* of the spheres. The study of arrangements plays a fundamental role in geometric computing. We denote the arrangement of the spheres in M by $\mathcal{A}(M)$, that is, $\mathcal{A}(M)$ is the subdivision of 3-space into cells of dimensions $0, 1, 2$ and 3, induced by the spheres in M (see, e.g., [4, pp. 144–150]). The combinatorial complexity of the arrangement $\mathcal{A}(M)$ is defined to be the number of cells of various dimensions in this subdivision. For arbitrary

sets of spheres the complexity can be $\Theta(n^3)$. The following is an immediate corollary of Theorem 2.1 (i):

Corollary 2.2 *The complexity of the arrangement $\mathcal{A}(M)$ of the spheres in M as defined in Theorem 2.1 is $\Theta(n)$.*

For most algorithmic uses, however, a raw arrangement is an unwieldy structure. The difficulty is that many cells in an arrangement can have very complex topologies, and thus navigating around them is difficult. What we often want is a further refinement of the cells of an arrangement into pieces that are each homeomorphic to a ball and have constant description complexity.

A prevailing and general technique for decomposing arrangements is the *vertical decomposition* (see, e.g., [3]), defined as follows. For every sphere S_i, we call the curve of intersection of S_i and the horizontal plane through the center of S_i the *equator* of the sphere. Let E be the collection of curves on the spheres in M consisting of intersection curves between any pair of spheres and the equators of the spheres. Let γ be a curve in E. We extend a vertical segment upward and downwards from every point of γ until it hits a sphere in M or extends to infinity. We repeat this process for every curve in E. As a result we get a collection of vertical walls that together with the spheres in M subdivide 3-space into xy-monotone 3D cells. We then project each cell onto the xy-plane, and extend vertical segments (with respect to the y-direction) from every vertex of the projection and from every x-extreme point in the projection of a cell, such that the segments are maximal and contained inside the projection of the cell. We then extend each of these segments into a vertical wall contained inside the original 3D cell. For details and illustrations of vertical decompositions for arrangements of spheres in 3-space, see [4].

Clarkson et al. [4] show that the complexity of the vertical decomposition of an arrangement of spheres is dominated by the complexity of the vertical walls erected from curves in E. They show a slightly supercubic upper bound on the complexity of the vertical decomposition of an arbitrary collection of n spheres in 3-space. In the next theorem we show another favorable property of the collection of spheres that we study:

Theorem 2.3 *The complexity of the vertical decomposition of the arrangement $\mathcal{A}(M)$ for a collection of spheres M as defined in Theorem 2.1 is $O(n^2)$.*

Proof. We distinguish between two portions of 3-space: inside the union of the balls (i.e., the balls corresponding to the spheres in M), and outside the union of the balls. The complexity of the vertical decomposition of $\mathcal{A}(M)$ inside the union of the balls is evidently $\Theta(n)$, by Theorem 2.1 (i). Hence, from this point on we bound the complexity of the vertical walls that lie completely outside the union of the balls. These walls are erected from curves that lie on the union boundary. Recall that the equators of the spheres are also considered curves in the arrangement $\mathcal{A}(M)$.

Fix an edge γ of $\mathcal{A}(M)$ that lies on the union boundary, and let H_γ be the vertical surface which is the union of vertical lines through points on γ. Let F be the collection of faces of $\mathcal{A}(M)$ that lie on the union boundary, where a face is a maximal portion of a sphere in M that does not meet any other sphere or the equator of that sphere, and is not contained in any other sphere. By Theorem 2.1 (parts (i) and (ii)) F consists of $O(n)$ faces (or surface patches), each bounded by a small number (bounded by a constant) of low-degree algebraic curves.

For every face f in F, let f' be $f \cap H_\gamma$. f' is a collection of a small number of low degree algebraic curves on H_γ. By standard arguments, the complexity of the vertical wall extended from γ and that lies outside the union, is determined by the complexity of the lower envelope (or upper envelope, or both) defined by the collection $F' = \{f' | f \in F\}$ with regard to the curve γ. Since the curves F' are either pairwise disjoint, or a pair of curves in F' meet at an endpoint of at least one of the curves, and they do not intersect otherwise, the complexity of the envelope is $O(n)$. We repeat this argument for every edge γ on the union boundary. There are $O(n)$ such edges, and the upper bound follows. □

The bound is sharp. One can easily construct a set of (even non-intersecting) unit spheres that achieve $\Theta(n^2)$ complexity. It is though possible to construct an $O(n)$ decomposition of the arrangement. Consider a three-dimensional grid whose unit size (along each of the coordinate axes) is $2 \times r_{\max}$. The grid induces a partitioning of 3-space into axis parallel unit cubes (those cubes whose vertices are grid points and that do not contain any grid point in their interior). By applying the same type of arguments as in the proof of Theorem 2.1, one can show that every cube in this partitioning intersects at most some fixed number of spheres. Moreover, each sphere intersects at most 8 cubes. Now take the collection of $O(n)$ non-empty cubes. It is easy to see that the vertical decomposition of this set has linear complexity. For each of the cubes consider the arrangement of spheres that intersect it, restricted to the cube and construct its vertical decomposition (again restricted to the cube). This will give a decomposition of the cube in a bounded

number of simple pieces. As there are $O(n)$ cubes to decompose the total complexity will be linear. It is easy to see that the number of adjacencies between pieces is linear as well, which is an important property in certain applications (like motion planning).

Theorem 2.4 *For a collection of spheres M as defined in Theorem 2.1 there exists a decomposition of the arrangement $\mathcal{A}(M)$ into simple pieces of total complexity $O(n)$.*

3 Intersection Queries

Given a molecule it is required that we are able to answer certain queries efficiently. In particular intersection queries play an important role. For example, in computer aided drug design, one wishes to manipulate a molecule in the presence of another molecule to see whether they fit together. During such manipulation the molecules should not intersect. Hence, with every change in the manipulated molecule an intersection query with the other molecule must be performed. Also, it is often required to know whether particular positions in 3-space lie inside or outside the molecule, i.e., whether a point intersects the molecule or not.

In this section we will devise a data structure that can be used to answer intersection queries with either a point or with a ball whose radius is bounded by $r_{\max}$. The query should report all atom balls that contain the point or intersect the query ball. We will use this structure in the next section to compute the molecule boundary. Clearly a point is a degenerated ball. Hence, a data structure that can answer queries for balls can also answer point queries. Below we only describe a structure for query balls. When implementing the structure though a distinction is made because point queries can be answered faster (i.e., with smaller constants in the time bounds).

Our first approach to solving this problem transforms it into an orthogonal range query problem. Let M be the set of n balls (as defined in Theorem 2.1) to store and let C be the set of centers of the balls. We store the set C in a three-dimensional range tree that can answer the following type of query: Given a query axis-parallel box, report the points of C contained in the box. To report all the balls of M intersecting a query ball Q with radius r_Q we construct a ball Q' with radius $r_Q + r_{\max}$ and concentric with Q. Clearly, the center of a ball of M intersecting Q must lie inside Q'. Next, let $\bar{Q}'$ denote the axis-parallel bounding box of Q'. We query the range search structure on C with the box $\bar{Q}'$. Using the same technique as in the proof of Theorem 2.1(i), it can be shown that the number of answers is bounded by a constant. For each of these answers we check whether the corresponding ball actually intersects Q. This takes $O(1)$ time. Using a standard range tree (see, e.g., [18, Section 2.3.4]) we obtain a solution that requires $O(n \log^2 n)$ preprocessing time and space and answers queries in time $O(\log^2 n)$.

However, due to the special properties of the collection of spheres that we study, they admit a simpler and more efficient data structure. Like in the previous section, we subdivide space into cubes of size $2 \times r_{\max}$. For each ball in M we compute the grid cubes that it intersects. Let $\mathcal{C}$ be the set of non-empty grid cubes. The size of $\mathcal{C}$ is bounded by $O(n)$. We arrange the cubes of $\mathcal{C}$ in a balanced binary search tree, ordered by the lexicographic order of the (x, y, z) coordinates of the say, bottom-left-front vertex of the cube. With each non-empty cube we store the list of (at most a constant number of) balls of M that intersect it.

Given a query ball Q, we compute all the (at most 8) grid cubes it intersects, and search for each of these cubes in the binary tree, to see whether it is non-empty. If it exists we check the balls stored for intersection with Q. We might find some balls more than once but duplicates can easily be removed. The total number of balls tested will be $O(1)$. This leads to the following result:

Theorem 3.1 *Given a collection of n balls M as in Theorem 2.1, one can construct a data structure using $O(n)$ space and $O(n \log n)$ preprocessing time, to answer intersection queries for balls whose radii are not greater than $r_{\max}$, in time $O(\log n)$.*

The binary search tree in the above approach is only used to locate the different cubes. Hence, we can easily replace it by a hashing structure. Using the perfect hashing results in [7] this leads to:

Theorem 3.2 *Given a collection M of n balls as defined in Theorem 2.1, one can construct a data structure using $O(n)$ space, to answer intersection queries for balls whose radii are not greater than $r_{\max}$, in $O(1)$ time. The randomized preprocessing time of the structure is $O(n)$.*

Recently, Fjällström and Petersson [8] have carried out a comparative study of the behavior in practice of various three-dimensional range search structures, motivated by simulation of deformation processes. An interesting result of their study is that range trees do not perform well in practice (for the special test data that they experiment with), as compared to other methods that they examined. In particular the best performing method in their study is a grid-like method, similar to the one mentioned above.

We implemented this data structure although we used double hashing rather than perfect hashing.

The following table shows, for each of the four test molecules: their size, the building time, the time to answer 1000 point intersection tests and the time to answer 1000 intersection tests for a ball with radius $r_{\max}$. All running times are in seconds, on an Indigo R3000 workstation (25.4 specmarks).

mol.	n	build	point query	ball query
1	42	0.01	0.14	0.20
2	100	0.01	0.13	0.22
3	757	0.08	0.13	0.19
4	971	0.12	0.14	0.22

As can be seen the query time seems independent of the size of the molecule. Also, the query time bounds are very low allowing, e.g., for testing whether a molecule with 100 atoms intersects a molecule with 1000 atoms in 0.02 seconds, i.e., fast enough for interactive use while manipulating the molecule.

4 Computing the Boundary

We now use the data structure of the previous section to efficiently compute the outer cell of the boundary of the union of the spheres. We assume that the union consists of one connected component, which is often the case in our application setting. If the union is composed of several pairwise-disjoint connected components, then additional measures must be taken; see below for more details.

The boundary of the outer cell can be viewed as a formal molecular surface, which is of interest in various applications in molecular biology. Moreover, we will show that the same algorithm, operating on a closely related set of balls, solves another problem in computational biology, sometimes referred to as computing the *approximate solvent accessible surface* (see, e.g., [14, 5, 16]). This involves the interaction of solute and solvent molecules in a solution. In a simplified model, the question may be formulated as follows: Which parts of a molecular surface of a solute molecule are accessible to the solvent molecules, where the latter are modeled by spheres (a single sphere each). We will assume that the solute molecules are modeled by the hard sphere model.

An approach to solving this problem is proposed in [14]: Roll a sphere representing the solvent molecule on a reference surface, to obtain a new surface described by the center of the rolling sphere. The reference surface, in our case, is the boundary of the union of the balls in the hard sphere model. In most cases, the solvent "sphere" is assumed to be fairly small. We denote the solvent molecule sphere by R.

We can rephrase the above problem in terms of *motion planning in robotics*: Let R be a spherical robot moving in 3-space among spherical obstacles (the balls in M). Describe the obstacles in a configuration space where every point represents a possible placement of R by the position of the center of R. In this formulation, the solvent accessible surface becomes the boundary of the outer cell of the free portion of the configuration space. Unlike prevailing formulations of motion-planning problems, the "obstacles" in our case may intersect. However, as expressed in Theorem 2.1, these obstacles have other, favorable, properties.

To compute the boundary of the outer cell of the free configuration space, we follow a common practice in motion planning, and compute the Minkowski (vector) difference[1] of each obstacle and the robot, to obtain the configuration obstacles. Now, the Minkowski difference of two spheres is again a (larger) sphere. Hence, our goal is actually to compute the outer component of the union boundary for a collection of balls in 3-space.

The solute molecule is modeled by balls as in Theorem 2.1. The radius of each ball in M is increased by r'—the radius of the solvent sphere. We assume r' to be of the same order of magnitude as $r_{\max}$. Note that, the expanded balls obey the conditions of Theorem 2.1 for different constants $r'_{\min}$, $r'_{\max}$, k' and ρ'. Therefore, an algorithm for computing the union boundary for a collection M of balls as in Theorem 2.1, is also applicable to computing solvent accessible surfaces. Thus, we present an algorithm for computing the union boundary of a collection of balls M as defined in Theorem 2.1.

The scheme of the algorithm is to compute, for each $B_i \in M$, the contribution of its boundary to the union boundary and then to combine all this information to give the final output of the algorithm. The algorithm consists of four steps:

1. For each ball identify the other balls intersecting it.
2. For each ball compute its (potentially null) contribution to the union boundary.
3. Transform the local information into global structures describing the connected components of the union boundary.
4. Filter out portions of the union boundary that do not belong to the desired connected component of the "outer cell".

If all we need is a list of the faces of the union boundary (as is the case in the preparatory stage of

[1] The Minkowski difference of two spatial sets A and B, $A - B$, is the set $\{p - q | p \in A, q \in B\}$.

the hidden surface removal algorithm of Section 5) then we only need Steps 1 and 2.

For Step 1, we use the data structure described in Theorem 3.2. The cost of the query for each ball is $O(1)$, hence the total cost of this step is $O(n)$ randomized time (or $O(n \log n)$ deterministic).

As for Step 2, consider a ball B_i and the family of (a constant number of) balls intersecting it. Let B_j be a ball intersecting B_i. If B_j fully contains B_i, then we stop the process for B_i, as it cannot contribute in any way to the union boundary. If B_j is fully contained in B_i, then we ignore B_j, for obvious reasons. Otherwise, we compute the intersection between the spheres S_i and S_j, which is a circle C_{ij} on S_i. The circle C_{ij} partitions S_i into two parts: One part may appear on the union boundary, and we will refer to it as the *free* part of S_i with respect to S_j, and the other is completely contained in S_j and therefore cannot appear on the union boundary. We repeat the process for each ball intersecting B_i, to get a collection of circles on S_i. These circles form a 2D arrangement $\mathcal{A}_i$ on S_i. A face of $\mathcal{A}_i$ belongs to the union boundary if and only if it belongs to the free portion defined by each circle C_{ij}. Since the number of circles on S_i defined in that way is constant, the arrangement can be computed by a brute force method in constant time, together with the attribute whether a face is free or not (a free face is guaranteed to appear on the union boundary). The complexity of the arrangement $\mathcal{A}_i$ is evidently constant. For each ball in M, the above procedure will take constant time.

In Step 3 we will represent each connected component of the union boundary by a *quad-edge* structure [10]. To this end we have to augment the arrangements $\mathcal{A}_i$ slightly. If $\mathcal{A}_i$ is the whole sphere S_i we split it into two parts with some circle. Next, if a boundary component of $\mathcal{A}_i$ is a simple circle C we split C into two arcs adding two vertices. (If C lies in both $\mathcal{A}_i$ and $\mathcal{A}_j$ the same vertices should be added in both arrangements.) Finally, if a free face of $\mathcal{A}_i$ contains holes we split it by adding extra arcs. (To make these additions canonical, we can fix a direction d, and add all the extra arcs along great circles that are intersection of the sphere with planes parallel to the direction d.) After this step, which can easily be performed in time $O(n)$, the union boundary will consist of simple faces where each face is bounded by at least two edges and each edge bounds exactly two faces (assuming general position).

Now we proceed as follows: We take some free face f of some $\mathcal{A}_i$, that was not treated before, and determine its boundary. In this way we obtain for each edge e bounding f pointers to the previous and next edge along f, as required for the quad-edge structure. Also for each edge e bounding f we determine the arrangement $\mathcal{A}_j$ containing the face f' on the opposite side of f (j can be i because of the extra arcs we added). f' can be found in time $O(1)$ because only a constant number of spheres intersect S_i. We locate e in $\mathcal{A}_j$, add pointers between the copies of e and recursively treat f' in $\mathcal{A}_j$ if it was not visited before. In this way we continue until we have located the whole connected component of the boundary containing f. If not all free faces have been visited we take one of the non-visited faces and continue from there. It is easy to see that this requires time $O(n)$ in total.

As a result of Steps 2 and 3 we get a quad-edge representation of each connected component of the union boundary of the balls in M. Recall that we are only interested in the boundary of the outer cell, which we assumed to be connected. At this stage we need to identify the outer cell (Step 4). We search for the ball $B_l \in M$ having the point with largest z-coordinate on its boundary. We then determine an edge on the boundary of the face f_l of B_l containing the highest point. This face clearly belongs to the outer cell. From this face we scan the entire connected component containing f_l, using the quad-edge structure. We mark every edge found throughout the scan, as belonging to the outer cell. Finally we delete any edge that has not been marked in the previous step. The remaining structure describes the outer cell. This takes time proportional to the size of the union boundary which is $O(n)$.

As for space requirements, the data structure of Theorem 3.2 requires $O(n)$ space. The additional structures described above are easily verified to require linear space.

Theorem 4.1 *The outer portion of the boundary of the union of a connected collection of balls as defined in Theorem 2.1 can be computed in $O(n)$ randomized time (or $O(n \log n)$ deterministic), using $O(n)$ space.*

If the union of the balls is not connected, additional machinery is required to identify and discard components lying in internal cells. This can be done in time $O(n \log n)$ by computing the decomposition of $\mathcal{A}(M)$ as described in Theorem 2.4 and traversing the outer cell of this arrangement. Details will be given in the full paper.

We only implemented Steps 1 and 2 of the above algorithm because that is all that is required for hidden surface removal. The following table shows the running time for our four molecules.

mol.	Step 1	Step 2
1	0.02	0.2
2	0.05	0.7
3	0.42	6.5
4	0.54	7.5

As can be seen Step 2 dominates the amount of time required, while theoretically both steps take time $O(n)$. It is though easy to verify that the operations involved in Step 2 are much more complicated and, hence, the constants in the bound are much higher.

5 Hidden Surface Removal

One of the tasks of molecular modelling packages is to display molecules. This should preferably be done so fast that the user can interact with the model, e.g., by turning it around to look at it from different directions or by moving different molecules with respect to each other, for example to see whether they fit together in some nice way. Hence, one needs fast algorithms for hidden surface removal among sets of intersecting spheres.

In practice one normally uses the *Z-buffer algorithm* for this. Almost all 3-D graphics workstation available nowadays have such an algorithm implemented in their graphics hardware. Unfortunately, such hardware can only handle polyhedral objects. Hence, as a first step, one has to approximate the spheres by triangular meshes. Different methods exist but all lead to a large number of faces. To get a reasonable sphere one needs at least 100 triangles. So a molecule of 1000 atoms requires the drawing of 100,000 triangles in 3-space. Even special graphics hardware (except maybe the most expensive types) cannot display such molecules fast. It would typically take between one and a few seconds.

A second approach used in computer graphics is the *painter algorithm*. Here one tries to define a depth order on the objects, sorting them from back to front. Next one draws the objects in this order on top of each other (like a painter) where each new object hides the parts of other objects that lie below it. Such an approach does not require special graphics hardware and, hence, can also be used on, e.g., laser printers. The problem though is that one needs a valid depth order on the objects. Such an order does not always exist (there can be cyclic overlap among objects). Also, for intersecting object such an order obviously does not exist.

For sets of non-intersecting spheres an easy depth order exists: simply sort the spheres by z-coordinate of their center. But for intersecting spheres this does not apply. There is though no need to draw the entire spheres. It is good enough to draw the pieces of the spheres that constitute the boundary of the union, as computed in the previous section. Clearly, pieces of the spheres that do not contribute to the boundary cannot be visible. Also the parts of the sphere that point away from the viewing direction cannot be seen. We will show that for the remaining part of the boundary a depth order does exist.

Theorem 5.1 *Let $S = S_1, ..., S_n$ be a collection of spheres sorted by decreasing z-coordinate of their center (i.e., from back to front). For each S_i let H_i be the hemi-sphere facing the viewing direction. Let $H_i^1, ..., H_i^{j_i}$ be the collection of maximal pieces of H_i that are part of the boundary of the union of the spheres. Then*

$$H_1^1, ..., H_1^{j_1}, H_2^1, ..., H_2^{j_2}, ..., H_n^1, ..., H_n^{j_n}$$

is a valid depth order for the pieces on the boundary.

Proof. We will prove this by contradiction. So assume there are two pieces H_i^k and H_j^l that are in the wrong order. Clearly $i \neq j$ so let us assume $i < j$. H_i^k and H_j^l are in the wrong order so there must be a ray in the viewing direction that first hits H_i^k and than H_j^l (so H_i^k partially hides H_j^l). Because H_i^k and H_j^l belong to the union boundary, the ray must first hit H_i^k, next the back of H_i, next H_j^l, and finally the back of H_j. But this implies that the back of H_j lies behind the back of H_i and, hence, $j < i$, which contradicts our assumption. □

This leads to the following algorithm. We first compute the boundary of the union. In this way, for each sphere S_i we collect a (constant) number of pieces that it contributes to the boundary. For each of these pieces we cut off the part that does not lie in H_i. Next we sort the spheres by depth and draw the pieces in this order. As shown in the previous section computing the boundary takes deterministic time $O(n \log n)$ and results in $O(n)$ pieces. The sorting of the spheres by depth takes time $O(n \log n)$. Hence, we obtain the following result:

Theorem 5.2 *Given a set of n spheres as defined in Theorem 2.1, a valid depth order consisting of $O(n)$ pieces can be computed in time $O(n \log n)$.*

We implemented this algorithm. Actually, we split it into two phases: In the preprocessing phase we compute the boundary of the union and for each sphere we collect the pieces that are part of the boundary. (See the previous section for time bounds.) In the query phase we are given a particular viewing direction. Now we compute the hemi-spheres (which are dependent on the viewing direction) and cut off the parts that lie outside it. Next we compute the depth order for the spheres and draw the pieces. All pictures in this paper were generated using this approach. The implementation (even though not yet

fully optimized) is rather fast. We did run it on two different machines: An Indigo R3000 machine with entry level graphics and an Indigo R3000 XS24 with hardware Z-buffer. The first machine performs the Z-buffer algorithm in software while the second one has special purpose hardware for this. The following table shows the amount of time (in seconds) required per frame for both the Z-buffer algorithm and our method.

	entry graphics		XS-24	
mol.	Z-buf	new	Z-buf	new
1	0.5	0.2	0.1	0.1
2	1.1	0.5	0.25	0.3
3	5.5	3.6	1.85	2.0
4	7.4	4.6	2.35	2.9

From these figures one might conclude that the new method runs about twice as fast on the entry level system while it is slightly slower on the XS-24. This should though be treated with care. First of all, the pictures produced by the two methods are quite different. The way we implemented it, the Z-buffer algorithm produces a Gouraud shaded image while our new method produces a flat shaded image (because the method draws in 2-d one cannot produce anything else). Our method though draws boundary lines around the spheres and along the intersections (like in Figures 1 till 4) while the Z-buffer method does not (it cannot because it does not know where the intersections are). Also the picture quality differs largely. On 8-bit systems the Gouraud shaded image looks rather displeasing while the images produced by the new method look much nicer. Also there is the issue of precision. For the Z-buffer we used an approximation of the spheres with about 100 triangles. In the new method we used an approximation of the circles and ellipses (projected sphere boundaries and intersections) with 24 edges, which is better. Finally, the difference in time for our new method on the two different platforms is caused because, on the entry level systems, more than half of the time is spent in the actual drawing, not in the computation of the depth order. On the XS-24 this is much less. For more details on the implementation see [9].

So we can conclude that our new approach produces nice images and can be used on simple graphics workstations without 3-D capabilities. Even on machines with 3-D graphics hardware the method behaves comparably. Finally, and probably most important, the new method can be used to produce high quality Postscript pictures which is impossible with the Z-buffer algorithm.

In some applications a depth order is not enough but one wants a combinatorial representation of the visible pieces, the *visibility map*. A lot of work has been done in computational geometry in producing visibility maps. See, e.g., the book of de Berg [6]. For collections of arbitrary intersecting spheres the best known hidden surface removal algorithm runs in randomized expected time $O(n^{2+\epsilon})$ for any $\epsilon > 0$ (using the general algorithm for computing the lower envelope of low-degree algebraic surfaces in 3-space by Sharir[19]). No output-sensitive method, where the time bound depends on the complexity of the resulting visibility map, is known. In the case of a set of non-intersecting spheres the best known method is presented in [12]. It runs in time $O((n+k)\log^2 n)$ where k is the complexity of the visibility map (which can be $\Theta(n^2)$, even for non-intersecting unit spheres). We will now briefly indicate how the method of [12] can be adapted to work for intersecting spheres as well. For more details see the full paper. We assume some familiarity with the method of [12].

For the method of [12] to work one needs to subdivide the objects in a (small) collection of pieces that first of all has a depth order and, secondly, has the property that the union of the projection on the viewing plane of the pieces lying in any depth range has small complexity. One could use the collection of pieces as defined in Theorem 5.1. It is though unclear whether this collection satisfies the second condition. Hence, we define a different set of pieces.

For each sphere S_i we again consider the hemisphere H_i. We take the pieces of H_i that are contained in the boundary of the union of $\{H_1, ..., H_i\}$, i.e., we ignore the spheres that lie nearer. In a way similar to the proof of Theorem 5.1 one can prove that these pieces constitute a valid depth order. In the full paper we will also prove that for any z-range the union of the projection of the pieces whose sphere center lies in the range, is linear. Applying the results in [12] we obtain:

Theorem 5.3 *Given a set of n spheres as defined in Theorem 2.1, the visibility map can be computed in time $O((n+k)\log^2 n)$, where k is the complexity of the resulting map.*

6 Conclusion

We have considered the hard sphere model of a molecule in a fixed nuclear configuration from a computational geometry point of view. We have shown that the collection of spheres representing the atoms of a molecule in that model behaves favorably in comparison with an arbitrary collection of spheres in 3-space. Using this observation we were able to devise a data structure that efficiently detects which atom spheres of a molecule are intersected by a query

sphere (of bounded radius). Then, we presented an efficient algorithm for computing the boundary (or envelope) of the union of atom spheres. Furthermore, we presented efficient algorithms for hidden surface removal of a molecule. Our results have been supported by experiments carried out on actual protein data, displaying the efficiency and practicality of our approach.

We believe that other (algorithmic) problems in molecular modelling can be solved efficiently using the data structures and properties presented in this paper. The area also provides many other challenging problems, like: Is it possible to preprocess two molecules such that one can efficiently determine whether, in a particular pose, the molecules intersect or not? Also other models of molecules give rise to interesting question. For example, there is the problem of hidden surface removal for the ball-and-stick model where atoms are represented by small balls and the links by thin sticks. We plan to work further on these problems.

Acknowledgements

The authors wish to thank Ed Moret, Raquel Norel and Haim Wolfson for helpful discussions and for supplying various data and references on molecular structures. We would like to thank Geert-Jan Giezeman who did most of the implementation work.

References

[1] E.E. Abola, F.C. Bernstein, S.H. Bryant, T.F. Koetzle and J. Weng, Protein data bank, in *Crystallographic Databases: Information Content, Software Systems, Scientific Applications*, F.H. Allen, G. Bergerhoff and R. Seivers, Eds., Data Comission of the International Union of Crystallography, Bonn/Cambridge/Chester, 1987, pp. 107–132.

[2] F.C. Bernstein, T.F. Koetzle, G.J.B. Williams, E.F. Meyer Jr., M.D. Brice, J.R. Rodgers, O. Kennard, T. Shimanouchi and M. Tasumi, The protein data bank: A computer-based archival file for macromolecular structure, *Journal of Molecular Biology* **112** (1977), pp. 535–542.

[3] B. Chazelle, H. Edelsbrunner, L. Guibas and M. Sharir, A singly-exponential stratification scheme for real semi-algebraic varieties and its applications, *Theoretical Computer Science* **84** (1991), pp. 77–105.

[4] K.L. Clarkson, H. Edelsbrunner, L.J. Guibas, M. Sharir and E. Welzl, Combinatorial complexity bounds for arrangements of curves and spheres, *Discrete and Computational Geometry* **5** (1990), pp. 99–160.

[5] M.L. Connolly, Solvent-accessible surfaces of proteins and nucleic acids, *Science* **221** (1983), pp. 709–713.

[6] M. de Berg, *Ray shooting, depth orders and hidden surface removal*, Springer LNCS 703, Springer-Verlag, 1993.

[7] M. Dietzfelbinger, A. Karlin, K. Mehlhorn, F. Meyer auf der Heide, H. Rohnert and R.E. Tarjan, Dynamic perfect hashing: upper and lower bounds, *Proc. 29th IEEE Symposium on Foundations of Computer Science*, 1988, pp. 524–531.

[8] P.-O. Fjällström and J. Petersson, Evaluation of algorithms for geometrical contact-searching problems, *manuscript*, 1993.

[9] G-J. Giezeman and M.H. Overmars, Fast display of molecular models, in preparation.

[10] L.J. Guibas and J. Stolfi, Primitives for the manipulation of general subdivisions and the computation of Voronoi diagrams, *ACM Transactions on Graphics*, **4** (1985), pp. 74–123.

[11] *Handbook of Biochemistry*, H.A. Sober, Editor, The Chemical Rubber Co., 2nd Edition, 1970.

[12] M.J. Katz, M H. Overmars and M. Sharir, Efficient hidden surface removal for objects with small union size, *Comp. Geom.: Theory and Appl.*, **2** (1992), pp. 223–234.

[13] K. Kedem, R. Livne, J. Pach and M. Sharir, On the union of Jordan regions and collision-free translational motion amidst polygonal obstacles, *Discrete and Computational Geometry* **1** (1986), pp. 59–71.

[14] B. Lee and F.M. Richards, The interpretation of protein structure: Estimation of static accessibility, *J. of Molecular Biology* **55** (1971) pp. 379–400.

[15] J. Matoušek, N. Miller, J. Pach, M. Sharir, S. Sifrony and E. Welzl, Fat triangles determine linearly many holes, *Proc. 32nd IEEE Symposium on Foundations of Computer Science*, 1991, pp. 49–58.

[16] P.G. Mezey, Molecular surfaces, in *Reviews in Computational Chemistry*, Vol. I, K.B. Lipkowitz and D.B. Boyd, Eds., VCH Publishers, 1990.

[17] M. H. Overmars, Point location in fat subdivisions, *Inform. Proc. Lett.*, **44** (1992), pp. 261 – 265.

[18] F.P. Preparata and M.I. Shamos, *Computational Geometry—An Introduction*, Springer Verlag, New York, 1985.

[19] M. Sharir, Almost tight upper bounds for lower envelopes in higher dimensions, *Proc. 34th Annu. IEEE Sympos. Found. Comput. Sci.*, 1993, pp. 498–507.

[20] A.F. van der Stappen, D. Halperin and M. H. Overmars, On the complexity of the free space for motion planning problems, *Comput. Geom.: Theory and Appl.*, 1993, to appear.