Pushing Disks Together—The Continuous-Motion Case

Marshall Bern
Xerox Palo Alto Research Center
bern@parc.xerox.com

Amit Sahai *
University of California - Berkeley
amits@hkn.eecs.berkeley.edu

Abstract

If disks are moved so that each center-center distance does not increase, must the area of their union also be nonincreasing? We show that the answer is yes, assuming that there is a continuous motion such that each center-center distance is a nonincreasing function of time. This generalizes a previous result on unit disks. Our proof relies on a recent construction of Edelsbrunner and on new isoperimetric inequalities of independent interest. We go on to show analogous results for the intersection and for holes between disks.

1 Introduction

Let $\mathcal{D} = \{D_1, D_2, \ldots, D_n\}$ be a set of overlapping balls in $\mathbb{R}^d$ with centers $c_1, c_2, \ldots, c_n$ and radii $r_1, r_2, \ldots, r_n$. Let $c'_1, c'_2, \ldots, c'_n$ be points such that $|c'_i c'_j| \leq |c_i c_j|$ for each choice of i and j, where $|pq|$ denotes the Euclidean distance between p and q. Let D'_i represent the ball with center c'_i and radius r_i, and let $\mathcal{D}' = \{D'_1, D'_2, \ldots, D'_n\}$. Must the d-dimensional volume of the union $U = \bigcup_{i=1}^n D_i$ be at least the volume of $U' = \bigcup_{i=1}^n D'_i$?

This question was first posed by Kneser [8] and Thue Poulsen [10] (see [7]); it also arises in molecular physics [9]. An affirmative answer is known for the case that $n \leq d+1$. This result was stated by Hadwiger [5] in 1956; Capoyleas and Pach [2] give a proof. An affirmative answer is also known for the case that $d = 2$, each $r_i = 1$, and disks are assumed to have a *continuous shrinking motion*. That is, there are functions $c_i(t)$ mapping $[0,1]$ to $\mathbb{R}^2$, such that $c_i(0) = c_i$, $c_i(1) = c'_i$, and $t' \geq t$ implies $|c_i(t')c_j(t')| \leq |c_i(t)c_j(t)|$. This result was also stated by Hadwiger [5]; a proof is given by Bollobás [1]. The continuous motion assumption does indeed define a special case, as shown in Figure 1.

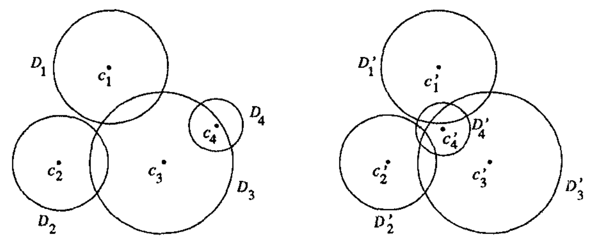

Figure 1. There is no continuous shrinking motion mapping disks $\mathcal{D}$ to $\mathcal{D}'$, because c_4 must pass through an edge of $c_1c_2c_3$ to get to its final location.

In this paper we show that the answer in the latter case remains affirmative if we remove the assumption that each $r_i = 1$. This extension requires a quite different proof, as the argument for the unit-radius case first proves that the perimeter of the union is nonincreasing, a fact that does not hold for disks of various radii. After writing up our results, we learned that Csikós [3] had obtained our main result earlier than ourselves, using a proof much closer to Bollobás's proof.

The remainder of this paper is organized as follows. Section 2 explains the dual complex, a simplicial complex induced by the disks $\mathcal{D}$. Section 3 establishes a new "optimality" property of the dual complex: shrinking the diagonal of a triangulated quadrilateral decreases the area covered by the disks. Section 4 establishes a similar property for interstices between disks. Section 5 shows that it is possible to shrink one edge at a time in the dual complex, while maintaining a meaningful configuration of disks. Section 6 combines all these ingredients and proves our main result. Finally, Section 7 gives two related results: the area of a bounded component of the exterior cannot increase, and the area of the intersection cannot decrease, under continuous shrinking motion.

2 The Dual Complex

Edelsbrunner [4] showed how to associate a simplicial complex called the *dual complex* with a set of overlapping balls, and used this construction in a formula for the volume. We explain the dual complex for the case $d = 2$. As above, let $\mathcal{D} = \{D_1, D_2, \ldots, D_n\}$ be a set of overlapping disks with centers c_1, $c_2, \ldots, c_n$ and radii $r_1, r_2, \ldots, r_n$. Define the

*Work performed while at Xerox PARC.

STOC'96, Philadelphia PA,USA
© 1996 ACM 0-89791-785-5/96/05..$3.50

power distance between arbitrary point p and disk center c_i to be $|pc_i|^2 - r_i^2$. For a point on the boundary of D_i the power distance is 0; for an exterior point it is the square of the tangential distance to D_i, and for an interior point p it is negative the square of half the length of the chord through p.

The *power diagram* of $\mathcal{D}$ is a subdivision of the plane into vertices, relatively open edges, and open cells, defined by the condition that all points in c_i's cell have smaller power distance to c_i than to any other disk center. It is not hard to show that the power distances to c_i and c_j are equal along a straight line, which, if D_i and D_j overlap, contains their mutual chord. Assuming general position, no point of the plane has equal power distance to four disk centers, so each vertex of the power diagram has degree 3. If all r_i's are equal, the power diagram is the same as the well-known Voronoi diagram.

The *regular triangulation* of $\mathcal{D}$ is an embedded planar graph. It contains each disk center c_i with a nonempty cell in the power diagram, and it contains the edge c_ic_j if and only if the power diagram cells of c_i and c_j share a boundary side. By a well-known transformation, if each c_i has coordinates (x_i, y_i), the regular triangulation is the projection onto the xy-plane of the lower convex hull of the points $\hat{c}_i = (x_i, y_i, x_i^2 + y_i^2 - r_i^2)$. If all r_i's are equal, the regular triangulation is the same as the well-known Delaunay triangulation.

We now focus attention on the region of the plane covered by the union of the disks $U = \bigcup_{i=1}^{n} D_i$. The *restricted power diagram* is the power diagram restricted to U. Notice that each point of U has nonpositive power distance to its closest disk center.

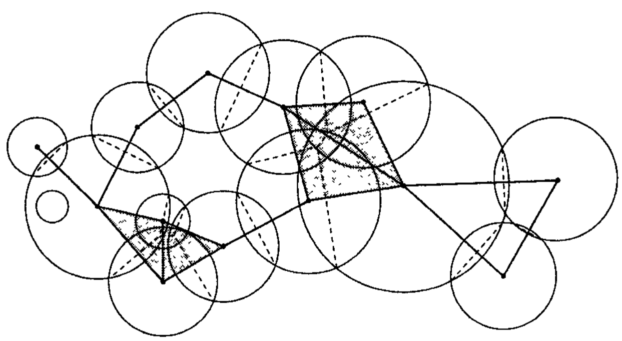

Figure 2. The restricted power diagram and the dual complex.

We denote Edelsbrunner's dual complex by $\mathcal{E}$ (or by $\mathcal{E}(t)$ when we want to make its dependence upon time explicit); it is formed as follows. The dual complex contains each disk center c_i with a nonempty cell in the restricted power diagram. It contains the edge c_ic_j between two disk centers if the restricted power diagram cells of c_i and c_j share a boundary side, and it contains triangle $c_ic_jc_k$ if the cells of c_i, c_j and c_k share a vertex. In order that the dual complex contains no faces more complex than triangles, we shall assume that the power diagram does not contain vertices of degree greater than 3; this can be ensured by slightly perturbing the disk centers. Notice that $\mathcal{E}$ may contain the three sides of a triangle without containing the triangle itself; in this way, it differs from an embedded planar graph such as the regular triangulation. See Figure 2.

We shall make use of both topological and geometric properties of the dual complex. The basic topological result, Lemma 1 below, is a consequence of the "nerve theorem" of algebraic topology. Intuitively speaking, Lemma 1 means that $\mathcal{E}$ and U have corresponding connected components with corresponding holes.

Lemma 1 (Edelsbrunner). *The region covered by $\mathcal{E}$ (that is, the union of its vertices, edges, and faces) is homotopy equivalent to U.*

Edelsbrunner showed that the area of U can be computed by a depth-3 inclusion-exclusion formula: sum the areas of all disks corresponding to vertices of $\mathcal{E}$, then—to correct for double counting—subtract off the areas of pairwise intersections corresponding to edges of $\mathcal{E}$, and finally add back in the areas of triple intersections corresponding to triangles of $\mathcal{E}$.

The next lemma gives another topological fact about $\mathcal{E}$. We call a vertex of $\mathcal{E}$ an *interior* vertex if it lies interior to the region covered by $\mathcal{E}$.

Lemma 2. *If c_i is an interior vertex of $\mathcal{E}$, then the perimeter of D_i is contained within the union of all the other disks.*

Proof: If c_i is an interior vertex, then the perimeter of D_i must be covered by the restricted power diagram cells of other disk centers, and hence by the corresponding disks. ∎

3 Three and Four Disks

We start by considering just three disks, D_1, D_2, and D_3, with centers c_1, c_2, and c_3 and union U. The disks are moving with time t; we sometimes make this dependence explicit by writing expressions such as $c_1(t)$ and $U(t)$. We use $\mu()$ to denote the area function and ∂ to denote the boundary operator.

Suppose that the lengths of c_1c_2 and c_2c_3 are fixed, while $|c_1c_3|$ is decreasing with time. Normalize this decrease so that $d|c_1c_3|/dt = -1$. Let z denote the *power center* of the three disks, that is, the point with equal power distance to c_1, c_2, and c_3.

Lemma 3. *The following holds:*

1. *If $D_1 \cap D_3$ is empty or is contained in D_2, then $\mu(U)$ is unchanging.*
2. *Otherwise, if $D_1 \cap D_2 \cap D_3$ is empty or either D_1 or D_3 contains the intersection of the other two disks, then $-d\mu(U)/dt$ equals the length of the mutual chord of D_1 and D_3.*
3. *If $D_1 \cap D_2 \cap D_3$ is nonempty and no disk contains the intersection of the other two, let p and q be the points of $\partial D_1 \cap \partial D_3$ inside and outside D_2, respectively. Then $-d\mu(U)/dt = |zq|$.*

Proof: Consider the inclusion-exclusion formula

$$\begin{aligned} \mu(U) \;=\; & \mu(D_1) + \mu(D_2) + \mu(D_3) - \\ & \mu(D_1 \cap D_2) - \mu(D_1 \cap D_3) - \mu(D_2 \cap D_3) + \\ & \mu(D_1 \cap D_2 \cap D_3). \end{aligned}$$

The only terms that change with time are $\mu(D_1 \cap D_3)$ and $\mu(D_1 \cap D_2 \cap D_3)$. If $D_1 \cap D_3$ is empty, then both these terms

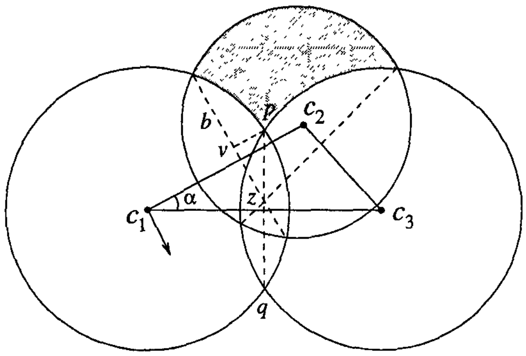

Figure 3. The shaded area is increasing at $|vp|$ times the speed of D_1, which is $1/\sin\alpha$ times faster than c_1c_3 shrinks.

are zero. If $D_1 \cap D_3 \subset D_2$, then the changes in these two terms cancel each other out. Hence statement (1) is true.

If $D_1 \cap D_2 \cap D_3$ is empty or either D_1 or D_3 contains the intersection of the other two disks, then $\mu(D_1 \cap D_2 \cap D_3)$ is unchanging. But $\mu(D_1 \cap D_3)$ is increasing with derivative equal to the length of the mutual chord of D_1 and D_3. So statement (2) is true.

Statement (3) is the most difficult. Let α be the measure of $\angle c_2c_1c_3$, let b be the mutual chord of D_1 and D_2, and let v be the point where the line through p perpendicular to b intersects b. See Figure 3.

Then, viewing D_2 and D_3 as fixed, the shaded area $D_2 \setminus (D_1 \cup D_3)$ grows as D_1 rotates about c_2, moving further into D_3. Instantaneously, D_1 is moving perpendicularly to c_1c_2, hence parallel to b. So the instantaneous gain in the area of $D_2 \setminus (D_1 \cup D_3)$ is $|vp|$ times the speed of D_1. And D_1 is moving exactly $1/\sin\alpha$ times faster than c_1c_3 shrinks. This same correction factor holds whether or not $\angle c_2c_1c_3$ is acute. This means that

$$\frac{d\mu(D_2 \setminus (D_1 \cup D_3))}{dt} = \frac{|vp|}{\sin\alpha}.$$

Note that $\angle pzv$ also measures α. Hence $|zp| = |vp|/\sin\alpha = d\mu(D_2 \setminus (D_1 \cup D_3))/dt$. Finally, since $|pq| = -d\mu(D_1 \cup D_3)/dt$ and $\mu(U) = \mu(D_2 \setminus (D_1 \cup D_3)) + \mu(D_1 \cup D_3)$, $|zq| = -d\mu(U)/dt$, as claimed. ∎

Now add one more disk. Let D_1, D_2, D_3, and D_4 be four moving disks with centers c_1, c_2, c_3, and c_4, and union U. Suppose that the dual complex of these four disks includes all the edges of the quadrilateral $c_1c_2c_3c_4$ along with c_1c_3.

Lemma 4. *Suppose that the disks move so that c_1c_3 shrinks but the lengths of the quadrilateral edges are held constant. Then the area of U is decreasing.*

Proof: Let p and q be the two points where ∂D_1 intersects ∂D_3, with p closer than q to c_2 by power distance. Let z be the power center of D_1, D_2, and D_3, and w be the power center of D_1, D_3, and D_4. Assume $d|c_1c_3|/dt = -1$.

Let D_2^- be D_2 minus the other three disks, and D_4^- be D_4 minus the other three disks. Notice that because c_1c_3 occurs in the dual complex, $D_2^- = D_2 \setminus (D_1 \cup D_3)$, and $D_4^- - = D_4 \setminus (D_1 \cup D_3)$. Consider the three disks D_1, D_2, and D_3. By the assumptions of the lemma, these three disks must fall into case (2) or (3) of Lemma 3. Hence, $d\mu(D_2^-)/dt$ must be either 0 or $|pz|$, and similarly, $d\mu(D_4^-)/dt$ must be either 0 or $|wq|$. Now $-d\mu(D_1 \cup D_3)/dt$ equals $|pq|$, and z and w must lie on pq. Since c_1c_3 is an edge in the dual complex, it must also be the case that the order of points on pq is p, z, w, q. Hence, $-d\mu(U)/dt \geq |zw|$ and $\mu(U)$ decreases as c_1c_3 shrinks. In the special position in which $z = w$, $d\mu(U)/dt$ is initially zero but becomes negative as zw grows. ∎

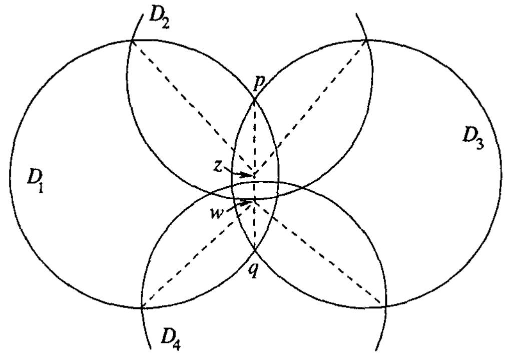

Figure 4. A four-sided quadruple intersection.

Figure 4 shows a four-sided four-way intersection. In this case, $d\mu(D_2^-)/dt = |pz|$, $d\mu(D_4^-)/dt = |wq|$, and hence $-d\mu(U)/dt = |zw|$.

The following theorem is our new "isoperimetric" result. The special case in which all disks have the same radius is equivalent to a well-known isoperimetric inequality: among all quadrilaterals with given side lengths, the area is maximized when the vertices are cocircular [6].

Theorem 1. *Let D_1, D_2, D_3, and D_4 be disks with centers c_1, c_2, c_3, and c_4. Let U denote the union and I the intersection of D_1, D_2, D_3, and D_4. Among all quadrilaterals $c_1c_2c_3c_4$ with fixed side lengths, $\mu(U)$ is maximized when the power diagram has a vertex of degree four. If I has four sides at the degree-4 configuration, then $\mu(I)$ is at a local maximum.*

Proof: Lemma 4 implies that at the degree-4 configuration, decreasing the length of either diagonal decreases the area of the union. Thus the degree-4 configuration is at least a local maximum. However, decreasing the length of a diagonal in the dual complex cannot remove the diagonal from the dual complex, so in fact the degree-4 configuration must also be a global maximum for $\mu(U)$.

If I is bounded by all four disks, then

$$\mu(I) = \mu(D_1 \cap D_2 \cap D_3) + \mu(D_1 \cap D_3 \cap D_4) - \mu(D_1 \cap D_3).$$

We apply Lemma 3, case (3), to each of the triple intersections, and case (2) to the pairwise intersection, to conclude that if $d|c_1c_3|/dt = -1$ then $d\mu(U)/dt = -|zw|$, with z and w as in Figure 4. Hence decreasing the length of either diagonal at the degree-4 configuration decreases the area of the intersection. ∎

4 Gaps between Disks

In this section, our aim is to prove that the area of a region surrounded by disks cannot increase when the disks move under a continuous shrinking motion.

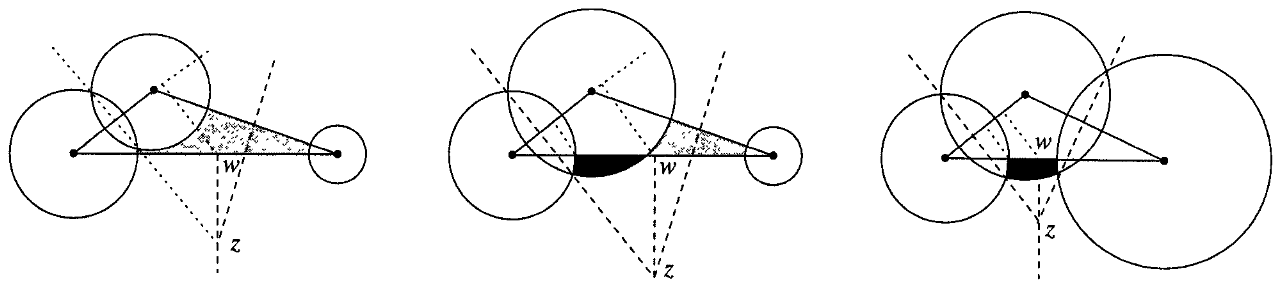

Figure 6. In each case the light-shaded area minus the dark-shaded area increases at a rate proportional to $|wz|$.

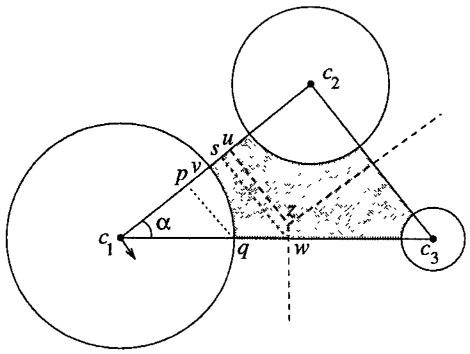

Figure 5. The shaded area decreases at a rate proportional to $|wz|$.

As above, we start with three disks, D_1, D_2, and D_3, moving so that $|c_1c_2|$ and $|c_2c_3|$ are fixed, while $d|c_1c_3|/dt = -1$. Let A denote the triangle minus the disks, $c_1c_2c_3 \setminus (D_1 \cup D_2 \cup D_3)$. Let z be the power center of the three disks. Let w be the point along the line c_1c_3 where the power distance from c_1 equals the power distance from c_3. Assuming c_1c_3 to be horizontal, points z and w lie on the same vertical line.

Lemma 5. *Assume D_1, D_2, and D_3 are pairwise disjoint and D_2 does not intersect c_1c_3. If z lies in triangle $c_1c_2c_3$, then $d\mu(A)/dt = -|wz|$, otherwise $d\mu(A)/dt = |wz|$.*

Proof: Region A is bounded by three arcs and three line segments, as shown in Figure 5. As in the figure, let v be the intersection of ∂D_1 with c_1c_2, q be the intersection of ∂D_1 with c_1c_3, p be the projection of q onto c_1c_2, and u be the point on c_1c_2 with equal power distance to c_1 and c_2.

As before, we think of D_2 and D_3 as fixed and D_1 as rotating about c_2. Hence A is unchanging along c_2c_3 and along its shared boundaries with D_2 and D_3. Along its shared boundary with D_1, A is shrinking by $|pv|$ times the speed of D_1. Along c_1c_2, A shrinks as its boundary segment sweeps out a trapezoidal region. We write an integral for the loss of area per unit of D_1 motion:

$$\frac{1}{|c_1c_2|}\int_{r_2}^{|c_1c_2|-r_1} x\,dx \;=\; -r_1 + \frac{1}{2|c_1c_2|}\left(|c_1c_2|^2 + r_1^2 - r_2^2\right).$$

This quantity has a neat geometric interpretation: some easy algebra shows that it is the length of vu. Finally, along c_1c_3, A grows by $\cos\alpha$ times

$$\frac{1}{|c_1c_3|}\int_{r_3}^{|c_1c_3|-r_1} x\,dx \;=\; -r_1 + \frac{1}{2|c_1c_3|}\left(|c_1c_3|^2 + r_1^2 - r_3^2\right).$$

The factor of $\cos\alpha$ accounts account for the direction of D_1's motion. The second factor has a geometric interpretation; it is $|qw|$.

We can combine these three quantities geometrically. We project qw onto c_1c_2 to multiply its length by $\cos\alpha$. If the projection of qw is ps, the net gain in area per unit motion of D_1 is $|ps| - |pv| - |vu|$, which is $|us|$ if u lies on the c_1 side of s, and $-|us|$ otherwise. Now since D_1 is moving $1/\sin\alpha$ times faster than c_1c_3 shrinks, we must divide this quantity by $\sin\alpha$. This can be done geometrically by projecting us onto wz, yielding the lemma. ∎

We now consider other topologies of disks and triangle. Let B denote the region of the plane below line c_1c_3 that is covered by D_2 but not by D_1 nor D_3. In Figure 6, B is the darker shaded area. The next lemma generalizes Lemma 5.

Lemma 6. *If z lies in triangle $c_1c_2c_3$, then*

$$d(\mu(A) - \mu(B))/dt = -|wz|,$$

otherwise it is $|wz|$.

Proof: The argument is similar to Lemma 5. The segments representing changing boundaries of A and B are projected onto the c_1c_2 line and then onto the wz line, as shown in Figure 6. In each case, regardless of the topology of the triangle and disks, w and z mark the endpoints of the projection. ∎

Now imagine adding a fourth disk D_4 below the c_1c_3 line, so that the cells of c_1 and c_3 share an edge in the power diagram of c_1, c_2, c_3, and c_4. Let Q^- denote the quadrilateral minus the disks, $c_1c_2c_3c_4 \setminus \bigcup_{i=1}^{4} D_i$. Applying Lemma 6 to each of D_1, D_2, D_3 and to D_1, D_3, D_4 shows that $d\mu(Q^-)/dt$ is negative. In fact, $d\mu(Q^-)/dt$ is at most negative the length of the c_1-c_3 power diagram edge. (It would be exactly negative the length, except that area B may actually protrude from the quadrilateral as in Figure 8.) Thus $\mu(Q^-)$ is maximized when the power diagram of the disk centers has a vertex of degree four; this adds yet another isoperimetric result.

We now turn to the case of n disks $\mathcal{D} = \{D_1, D_2, \ldots, D_n\}$. Assume that no disk is covered by the union of the other disks, but D_i is a disk whose perimeter is covered by the union of all the other disks. Let D_i^- denote $D_i \setminus \bigcup_{j \neq i} D_j$. Assume there is a continuous shrinking motion taking $\mathcal{D}$ to $\mathcal{D}'$.

Lemma 7. *Throughout the continuous shrinking motion, $d\mu(D_i^-)/dt \leq 0$. In other words, the area covered by D_i alone is nonincreasing.*

Proof: Assume without loss of generality that D_i^- has a single connected component, as multiple components can be handled separately. Renumber so that $D_1, D_2, \ldots, D_k$ cover the perimeter of D_i, $i > k$, with $D_i^- = D_i \setminus \bigcup_{j=1}^{k} D_j$. We may assume that the boundary of D_i^- contains, in order, exactly one arc from each of $D_1, D_2, \ldots, D_k$, for if some D_j contributed more than one arc, then a disk contributing an arc between the two D_j arcs must be contained in $D_i \cup D_j$. Thus $c_1 c_2 \ldots c_k$ is a simple polygon, which we shall assume is in general position. Let $\mathcal{R}$ $(= \mathcal{R}(t))$ denote the *constrained regular triangulation* of $c_1 c_2 \ldots c_k$, meaning the unique triangulation of $c_1 c_2 \ldots c_k$ in which each convex quadrilateral is triangulated as it would be by the regular triangulation of its vertices.

We shall simulate a small continuous shrinking motion of $\mathcal{D}$ by shrinking one edge of $\mathcal{R}$ at a time. At all times, $\mathcal{R}(t)$ will be a triangulated simple polygon, hence each edge length can be adjusted independently. After we shrink an edge, we may have to change $\mathcal{R}(t)$ combinatorially; for example, shrinking edge $c_j c_{j+1}$ in Figure 7 may "flip" the diagonal of quadrilateral $c_j c_{j+1} c_{j+2} c_l$ from $c_{j+1} c_l$ to $c_j c_{j+2}$. Since triangles are rigid, when all edges of $\mathcal{R}(t)$ have reached their final lengths, the configuration of disks must be at the desired endpoint. A technicality: as the disks move, polygon $c_1 c_2 \ldots c_k$ may overlap itself. In this case, we think of the polygon as embedded on a plane with Riemann sheets.

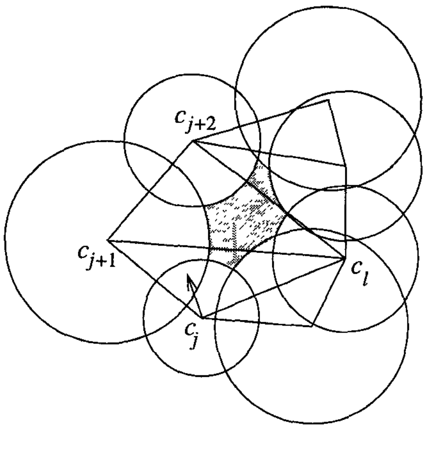

Figure 7. Shrinking an exterior edge of $\mathcal{R}(t)$.

Assume that we shrink an exterior edge $c_j c_{j+1}$ of $\mathcal{R}(t)$ while keeping all other edge lengths fixed, by rotating D_j around the third vertex of the triangle containing $c_j c_{j+1}$, as shown in Figure 7. Since D_i^- is losing area along its boundary with D_j and all other boundaries are fixed, this shrinking cannot increase $\mu(D_i^-)$.

Now assume that we shrink a diagonal $c_j c_m$ of a quadrilateral $c_j c_l c_m c_p$ in $\mathcal{R}(t)$. Since $c_j c_m$ is a diagonal of the regular triangulation, the power cells of c_j and c_m share an edge. Two applications of Lemma 6 reveal that the sum of the A areas minus the sum of the B areas is decreasing with rate proportional to the length of the shared power diagram edge. This quantity is exactly the area of $c_j c_l c_m c_p$ not covered by $D_j \cup D_l \cup D_m \cup D_p$, minus the area by which D_l and D_p protrude outside the far sides of the quadrilateral (the area of D_l below $c_j c_p$ in Figure 8). This accounts for both shrinking and growing boundaries of D_i^-. Notice that a disk such as D_{j+1} in Figure 8 moves rigidly with a side of

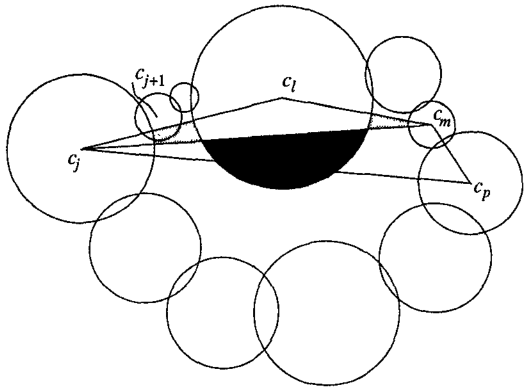

Figure 8. Shrinking an interior edge of $\mathcal{R}(t)$.

$c_j c_l c_m c_p$, so its overlap with $c_j c_l c_m c_p$ remains fixed. Hence shrinking $c_j c_m$ cannot increase $\mu(D_i^-)$. ∎

5 Disks without Embeddings

In order to prove that the area of the union of the disks cannot increase, we shall shrink edges of the dual complex $\mathcal{E}$ one at a time, just as Lemma 7 shrinks edges of the regular triangulation of a simple polygon. But in carrying out this plan, we encounter a difficulty. When we shrink an edge of a complicated $\mathcal{E}$, the configuration of disks will not in general remain planar-realizable.

For example, imagine shrinking an edge of a cycle surrounding a hole in $\mathcal{E}$ as shown in Figure 9. After this shrinking, the disks will be realizable on a cone but not on the plane. Notice, however, that so long as no disk covers the apex of the cone, each disk appears flat, that is, the region covered by that disk, including its overlaps with other disks, is planar embeddable. When we shrink edges of $\mathcal{E}$, we could imagine the disks $\mathcal{D}$ to be embedded in some developable surface such as the cone in Figure 9, but it is easier to think of $\mathcal{D}$ without reference to any embedding: it is simply an abstract collection of overlapping two-dimensional disks.

We define a *disk complex* to be a collection of disks, $\{D_1, D_2, \ldots, D_n\}$, such that for each D_i we know its radius r_i and its overlaps with other disks. For example, each D_i may have its own coordinate system, say with origin at the center c_i, that gives the coordinates of each nonempty lune $D_i \cap D_j$. Of course, D_i and D_j must agree on the size and shape of $D_i \cap D_j$ and its overlaps with other disks. Each D_i must be *intrinsically flat*, meaning that its overlaps with other disks are indistinguishable from overlaps in a configuration of planar disks. If D_i does not completely contain another disk, then D_i is intrinsically flat if and only if the sum of angles, subtended by closest points of lunes, around c_i is 2π.

Area is an integral, defined locally, so it still makes sense for a disk complex. Since each disk is intrinsically flat, the power distance from a point in D_i to c_i does not change, so the restricted power diagram and the dual complex $\mathcal{E}$ remain well-defined.

Lemma 8. *Assume that $\mathcal{D}$ is a complex of intrinsically flat disks, such that each disk center appears in the dual complex $\mathcal{E}$ and none of them is an interior vertex. Shrinking*

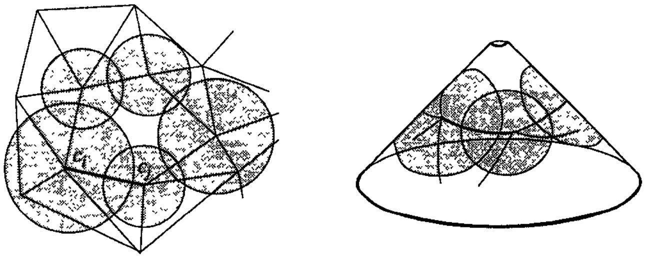

Figure 9. Shrinking edge $c_i c_j$ wraps the plane into a cone.

any single edge of $\mathcal{E}$ while keeping the others fixed maintains a collection of intrinsically flat disks.

Proof: Shrinking an edge $c_i c_j$ of $\mathcal{E}$ changes the angles in the triangles bounded by $c_i c_j$. Assume c_k is a vertex with a changed angle. Because c_k is not an interior vertex of $\mathcal{E}$, there is an angle around c_k that does not lie within a triangle of $\mathcal{E}$. We can adjust this exterior angle in order to maintain the intrinsic flatness of D_k. ∎

6 Putting it Together

We are now ready to state and prove our main result.

Theorem 2. *Assume that disks $\mathcal{D}$ move to $\mathcal{D}'$ via a continuous shrinking motion. Then the area of the union of the disks cannot increase.*

Proof: We shall prove that at all times $d\mu(U)/dt \leq 0$. Consider a small time interval during the continuous shrinking motion. We may assume that the dual complex $\mathcal{E}$ is combinatorially the same at the beginning and end of this time interval. There are only a finite number of different topologies for $\mathcal{E}$, so even if $\mathcal{E}$ changes infinitely often there must be two positions of $\mathcal{D}$ with the same $\mathcal{E}$ that are separated by a nonnegligible time interval.

Remove disks from $\mathcal{D}$ in arbitrary order, without changing U, until no disk is covered by the union of the others. Now if D_i is a disk in $\mathcal{D}$ whose perimeter is covered by the union of the other disks, remove D_i. This removal leaves a hole D_i^- of the form considered in Lemma 7. Repeat this process, again choosing disks in arbitrary order, until no disk has covered perimeter. Notice that each removal leaves a new hole, disjoint from previously formed holes. By Lemma 7 the derivative of the area of these holes is at most zero.

Let $\mathcal{D}^*$ denote the altered set of disks, regarded as a disk complex (although at this stage $\mathcal{D}^*$ is still realizable in the plane). Let $\mathcal{E}^*$ denote the dual complex of $\mathcal{D}^*$. Again it is safe to assume that $\mathcal{E}^*$ is the same at the beginning and end of our time interval.

We shall simulate the continuous shrinking motion of $\mathcal{D}$ by shrinking edges of $\mathcal{E}^*$ one at a time. Lemma 8 shows that shrinking an edge does not destroy the intrinsic flatness of disks. After we shrink an edge we recompute the dual complex. Notice that shrinking edge $c_i c_j$ never removes $c_i c_j$ from the dual complex. In fact, because we have assumed that $\mathcal{E}^*$ is the same at beginning and end, the only changes to $\mathcal{E}^*$ are diagonal flips in quadrilaterals. These diagonals may later flip back when we shrink another edge.

We assert that a finite number of edge shrinkings suffice to move $\mathcal{D}^*$ to its final configuration in the time interval. An exterior edge of $\mathcal{E}^*$ need be shrunk only once, because no subsequent shrinking can increase its length. An interior edge $c_i c_j$ in $\mathcal{E}^*$, however, can grow. For example, shrinking an outer edge of $c_i c_j$'s quadrilateral could flip $c_i c_j$ out of $\mathcal{E}^*$; now shrinking the opposing diagonal grows $c_i c_j$; and finally shrinking another outer edge can flip $c_i c_j$ back into $\mathcal{E}^*$. How do we know that we cannot get caught in an infinite loop? There is a measure of progress for each diagonal: the area of the union of the four disks in its quadrilateral, which is nonincreasing by Lemma 4.

We now assert that when we shrink an edge $c_i c_j$, the area of the union of all disks must decrease. In fact, we need only worry about the change in the area of the union of up to four disks, the disks whose centers are the vertices of triangles of $\mathcal{E}^*$ bounded by $c_i c_j$. Other disks (shown in gray in Figure 10) have unchanging intersection patterns, and we may think of them as moving rigidly with the outside edges of the triangles.

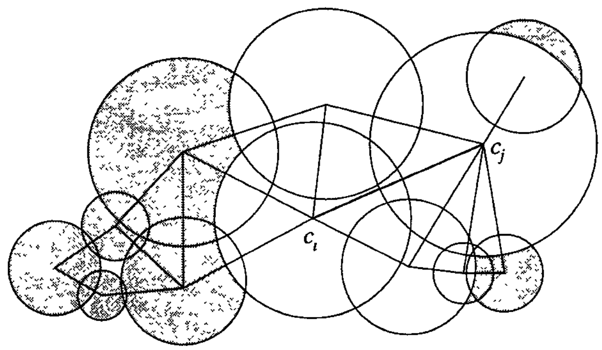

Figure 10. When $c_i c_j$ shrinks the shaded areas move rigidly.

First assume we shrink an edge $c_i c_j$ that bounds no triangles of $\mathcal{E}^*$. Then the area of the union must decrease because $\mu(D_i \cap D_j)$ increases, and all other intersections are unaffected. Next assume we shrink an edge $c_i c_j$ that bounds exactly one triangle in $\mathcal{E}^*$. In this case, Lemma 3 guarantees that the area of the union does not increase. Finally assume we shrink an edge $c_i c_j$ that bounds two triangles in $\mathcal{E}^*$, as in Figure 10. Lemma 4 handles this most difficult case. ∎

7 Related Results

In this section, we give two additional results. For the special case of unit disks, each of these results can be proved by a perimeter argument of the form given by Bollobás [1].

Theorem 3. *Assume that disks $\mathcal{D}$ move to $\mathcal{D}'$ via a continuous shrinking motion. Then the area of a bounded connected component of the exterior cannot increase (even if it breaks into a number of components).*

Proof: This theorem is a small extension of Lemma 7; the only difference is that the disks surrounding a connected component of the exterior do not necessarily form a simple polygon. This difference does not matter to the proof, so long as we treat the collection of disks as a disk complex rather than a planar-realizable configuration. ∎

Theorem 4. *Assume that disks $\mathcal{D}$ move to $\mathcal{D}'$ via a continuous shrinking motion. Then the area of the intersection cannot decrease.*

Proof: Remove all disks from $\mathcal{D}$ that do not contribute a boundary to the intersection $I = D_1 \cap D_2 \cap \ldots \cap D_n$. In other words, discard each D_i such that $\bigcap_{j \neq i} D_j \subset D_i$.

Now consider the *farthest-point regular triangulation* $\mathcal{F}$ of the remaining disk centers. If center c_i has coordinates (x_i, y_i), the farthest-point regular triangulation is the projection onto the xy-plane of the upper convex hull of the lifted points $\hat{c_i} = (x_i, y_i, x_i^2 + y_i^2 - r_i^2)$. Since each D_i now contributes to the boundary of I, each c_i has a nonempty farthest-point power diagram cell, and hence $\mathcal{F}$ includes all disk centers.

The remainder of the proof shrinks one edge of $\mathcal{F}$ at a time to move from $\mathcal{D}$ to $\mathcal{D}'$. As in the proofs of Lemma 7 and Theorem 2, we can limit attention to three and four disks at a time and think of the other disks as moving rigidly with the outside edges of the changing faces.

Shrinking an exterior edge—one that bounds only one triangle in $\mathcal{F}$—clearly cannot decrease $\mu(I)$. Now consider shrinking the diagonal of a quadrilateral $c_1c_2c_3c_4$ in $\mathcal{F}$. If $c_1c_2c_3c_4$ is not convex, then shrinking either diagonal shrinks both diagonals. In this case, I is growing along its boundaries with each D_i, $1 \leq i \leq 4$, and shrinking nowhere, so $\mu(I)$ cannot decrease. Finally, assume $c_1c_2c_3c_4$ is convex. If the ordinary regular triangulation uses diagonal c_1c_3, then $\mathcal{F}$ uses the opposite diagonal c_2c_4. As shown in the proof of Theorem 1, shrinking c_1c_3 decreases the intersection $D_1 \cap \ldots \cap D_4$; shrinking c_2c_4 grows c_1c_3 and hence increases the intersection. So in this case as well, $\mu(I)$ cannot decrease. ∎

8 Remarks

Observe that our main result, Theorem 2, extends to balls in higher dimensions, so long as all balls are moving parallel to a common (two-dimensional) plane. We conjecture that any continuous shrinking motion of $d + 2$ balls in $\mathbb{R}^d$ can be factored into such planar motions. (Shrinking one edge of a simplex at a time analogously factors a continuous shrinking motion into one-dimensional motions.) Further, the dual complex may be the right tool for solving the case of n balls in $\mathbb{R}^d$ with a continuous shrinking motion. Shrinking a $(d-1)$-simplex seems to be the appropriate generalization of shrinking an edge.

Finally, we admit to having no new insights for the problem without the assumption of continuous shrinking motion. We would be interested in seeing a computer search for a counterexample.

Acknowledgements

We would like to thank Herbert Edelsbrunner for suggesting a possible connection between the disk-pushing problem and the dual complex and David Eppstein for several helpful conversations. We would like to thank Jànos Pach for telling us about Csikós's result, just in time for inclusion in this paper.

References

[1] B. Bollobás. Area of the union of disks. *Elemente der Mathematik* 23 (1968), 60–61.

[2] V. Capoyleas and J. Pach. On the perimeter of a point set in the plane. In *Discrete and Computational Geometry*, DIMACS Series in Disc. Math. and Theoretical Comp. Science, Vol. 6, 1991, 67–76.

[3] B. Csikós. On Poulsen Kneser conjecture. Manuscript, Dept. of Geometry, Eőtvős University, Budapest, 1996.

[4] H. Edelsbrunner. The union of balls and its dual shape. In *Proc. 9th ACM Symp. on Computational Geometry*, 1993, 218–231.

[5] H. Hadwiger. Ungelöste Probleme, Nr. 40. *Elemente der Mathematik* 11 (1956) 60–61.

[6] N. Kazarinoff. *Geometric Inequalities*. Mathematical Association of America, 1961.

[7] V. Klee and S. Wagon. *Old and New Unsolved Problems in Plane Geometry and Number Theory*. Mathematical Association of America, 1991.

[8] M. Kneser. Einige Bemerkungen über das Minkowskische Flächenmaß. *Archiv der Mathematik* 6 (1955) 382–390.

[9] E.H. Lieb. Monotonicity of the molecular electronic energy in the nuclear coordinates. *Journal of Physics B: Atomic and Molecular Physics* 15 (1982) L63–66.

[10] E. Thue Poulsen. Problem 10. *Mathematica Scandinavica* 2 (1954) 346.