Discrete Comput Geom 20:499–514 (1998)

Discrete & Computational
Geometry

Pushing Disks Together—The Continuous-Motion Case*

M. Bern[1] and A. Sahai[2]

[1]Xerox Palo Alto Research Center, 3333 Coyote Hill Rd.,
Palo Alto, CA 94304, USA
bern@parc.xerox.com

[2]MIT Laboratory for Computer Science, 545 Technology Square,
Cambridge, MA 02139, USA
amits@theory.lcs.mit.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 $|c'_i c'_j| \le |c_i c_j|$ 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 most the volume of $U = \bigcup_{i=1}^n D_i$?

Kneser [9] and Thue Poulsen [10] (see [8] for more background) first asked this question, specifically for the case of unit-radius disks in the plane. An affirmative answer is known for the case that $n \le d + 1$ [6], [2]. An affirmative answer is also known for the case of unit disks in the plane, when disks are assumed to have a *continuous contraction*. That is, there are continuous functions $c_i(t)$ mapping [0, 1] to $\mathbb{R}^2$, such that $c_i(0) = c_i$, $c_i(1) = c'_i$, and $t' \ge t$ implies $|c_i(t')c_j(t')| \le |c_i(t)c_j(t)|$. Hadwiger [6] stated this result

* The work of the second author was performed while at Xerox PARC.

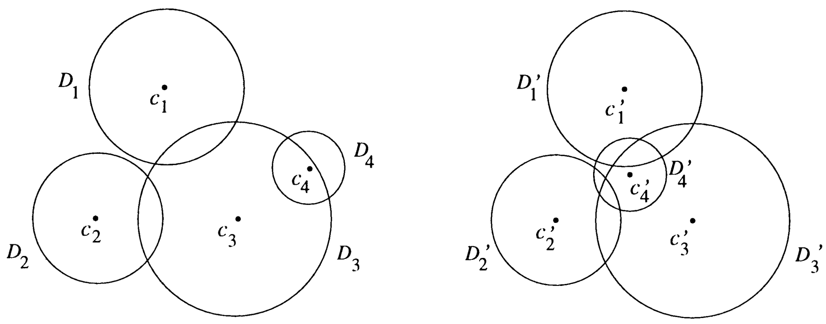

Fig. 1. There is no continuous shrinking motion mapping disks $\mathcal{D}$ to $\mathcal{D}'$, because c_4 must pass through an edge of $c_1c_2c_3$ to get to its final location.

without proof, attributing it to Habicht; Bollobás [1] supplied a proof. The continuous motion assumption does indeed define a special case, as shown in Fig. 1.

In this paper we show that the answer in the latter case remains affirmative if we remove the assumption that each $r_i = 1$. Our argument is quite different from Bollobás's argument, which 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 contraction.

2. The Dual Complex

Edelsbrunner [5] 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 *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 zero; 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 shortest 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 $c_i c_j$ 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.

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 $c_i c_j$ between two disk centers if the restricted power diagram cells of c_i and c_j share a boundary side, and it contains triangle $c_i c_j c_k$ 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 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 Fig. 2.

We 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

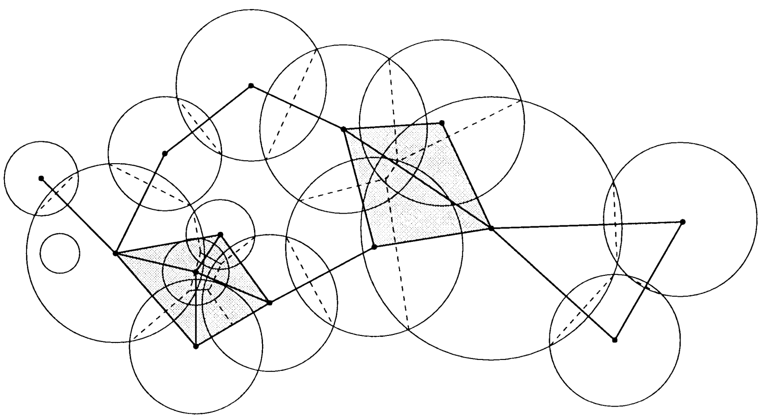

Fig. 2. The restricted power diagram and the dual complex.

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 a deformation retract of* U.

The next lemma states that the area of U can be computed by a depth-3 inclusion–exclusion formula.

Lemma 2 (Edelsbrunner). *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 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* $\mathcal{E}$.

Lemma 2 shows that, to compute the area of a configuration of disks, it suffices to know only the lengths of edges appearing in the dual complex. The last 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 $\mathcal{E}$; there is one such vertex in Fig. 2.

Lemma 3. *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. We assume that 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 "area of" and ∂ to denote "boundary of."

Suppose that the lengths of c_1c_2 and c_2c_3 are fixed, while $|c_1c_3|$ is decreasing smoothly (differentiably) 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 4. *The following statements hold*:

(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

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

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 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. However, $\mu(D_1 \cap D_3)$ 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 Fig. 3.

Then, viewing D_2 and D_3 as fixed, the shaded area $D_2 \backslash (D_1 \cup D_3)$ 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 $D_2 \backslash (D_1 \cup D_3)$ 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

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

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

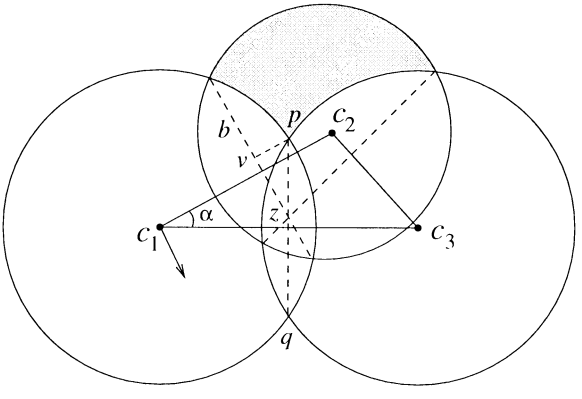

Fig. 3. The shaded area is increasing at $|vp|$ times the speed of D_1, which is $1/\sin\alpha$ times faster than c_1c_3 shrinks.

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 5. *Suppose that the disks move so that c_1c_3 shrinks at a constant rate of -1, but the lengths of the quadrilateral edges are held constant. Then $\mu(U)$ decreases with rate equal to the length of the shared edge of the restricted power diagram cells of c_1 and c_3.*

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 let w be the power center of D_1, D_3, and D_4. Note that z and w are colinear with p and q. 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, $D_2^- = D_2 \backslash (D_1 \cup D_3)$ and $D_4^- = D_4 \backslash (D_1 \cup D_3)$. 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 4. Hence, $d\mu(D_2^-)/dt$ must be zero if z is outside the segment pq, or $|pz|$ if z is on pq. Similarly, $d\mu(D_4^-)/dt$ must be zero if w is outside the segment pq, or $|wq|$ if not. Also, $-d\mu(D_1 \cup D_3)/dt$ equals $|pq|$ and $\mu(U)$ equals $\mu(D_1 \cup D_2) + \mu(D_2^-) + \mu(D_4^-)$. So $-d\mu(U)/dt$ must equal the length of the shortest of the segments pq, zq, pw, or zw, which is exactly the shared edge of the restricted power diagram cells of c_1 and c_3. □

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 [7].

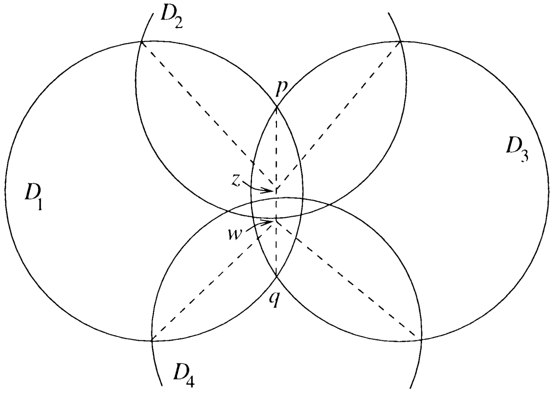

Fig. 4. A four-sided quadruple intersection.

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 4. If I has four sides at the degree-4 configuration, then $\mu(I)$ is at a local maximum.*

Proof. Lemma 5 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 4, 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 $d\mu(I)/dt = -|zw|$, with z and w as in Fig. 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 contraction.

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, $c_1c_2c_3 \backslash (D_1 \cup D_2 \cup D_3)$. 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 6. *Assume D_1, D_2, and D_3 are pairwise disjoint and D_2 does not intersect c_1c_3. If z and c_2 lie on the same side of line c_1c_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 Fig. 5. As in the figure, let v be the intersection of ∂D_1 with c_1c_2, let q be the intersection of ∂D_1 with c_1c_3, let p be the projection of q onto c_1c_2, and let 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:

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

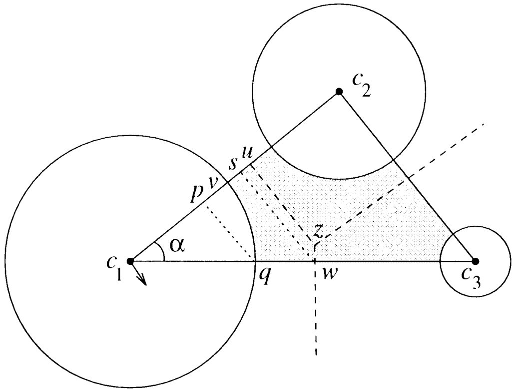

Fig. 5. The shaded area decreases at a rate proportional to $|wz|$.

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 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 Fig. 6, B is the filled-in area. The next lemma generalizes Lemma 6.

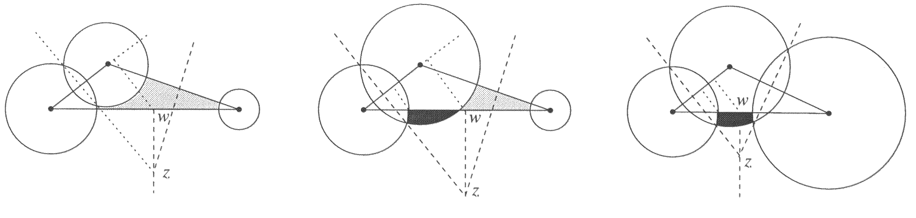

Fig. 6. In each case the shaded area minus the filled-in area increases at a rate proportional to $|wz|$.

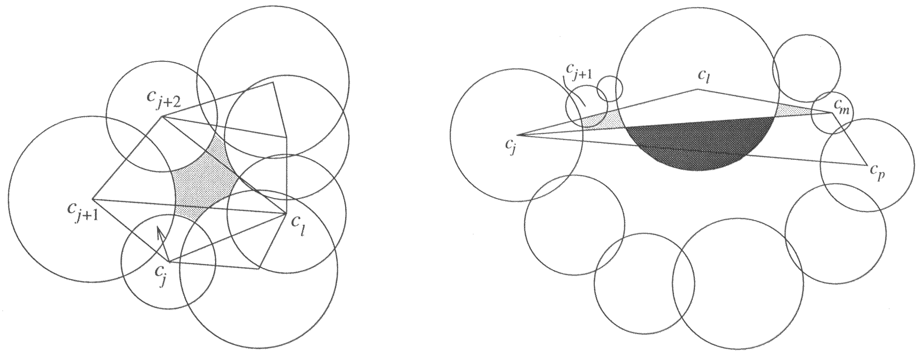

Fig. 7. Shrinking an (a) exterior or (b) interior edge of $\mathcal{R}(t)$.

Lemma 7. *If z and c_2 lie on the same side of line c_1c_3, then*

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

otherwise it is $|wz|$.

Proof. The argument is similar to Lemma 6. 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 Fig. 6. In each case, 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 7 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 Fig. 7(b).) 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 $\mathcal{D} = \{D_1, D_2, \dots, D_n\}$. 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 contraction taking $\mathcal{D}$ to $\mathcal{D}'$.

Lemma 8. *The area covered by D_i alone cannot increase with the continuous contraction.*

Proof. Assume without loss of generality that D_i^- has a single connected component, as multiple components can be handled separately. Renumber so that $D_1, D_2, \dots, D_k$ 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 $D_1, D_2, \dots, D_k$, 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 $c_1 c_2 \cdots c_k$ is a simple polygon, which we assume is in general position. Let $\mathcal{R}$ $(= \mathcal{R}(t))$ denote the *constrained regular triangulation* of $c_1 c_2 \cdots c_k$, meaning the unique triangulation of $c_1 c_2 \cdots c_k$ in which each convex quadrilateral is triangulated as it would be by the regular triangulation of its vertices.

We simulate a small continuous contraction of $\mathcal{D}$ by smoothly 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 Fig. 7(a) 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 $c_1 c_2 \cdots c_k$ may overlap itself. In this case, we think of the polygon as isometrically immersed in the plane, or alternatively embedded on a plane with Riemann sheets.

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 Fig. 7(a). 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 7 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 Fig. 7(b)). This accounts for both shrinking and growing boundaries of D_i^-. Notice that a disk such as D_{j+1} in Fig. 7(b) moves rigidly with a side of $c_j c_l c_m c_p$, 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 Global Embeddings

In order to prove that the area of the union of the disks cannot increase, we shrink edges of the dual complex $\mathcal{E}$ one at a time, while keeping constant all other dual complex edge lengths. However, 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 Fig. 8. After this shrinking, the disks will be realizable on a cone but not on the plane. Notice, however, that as 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.

We define a *disk complex* to be a finite set of planar disks and intersections of these

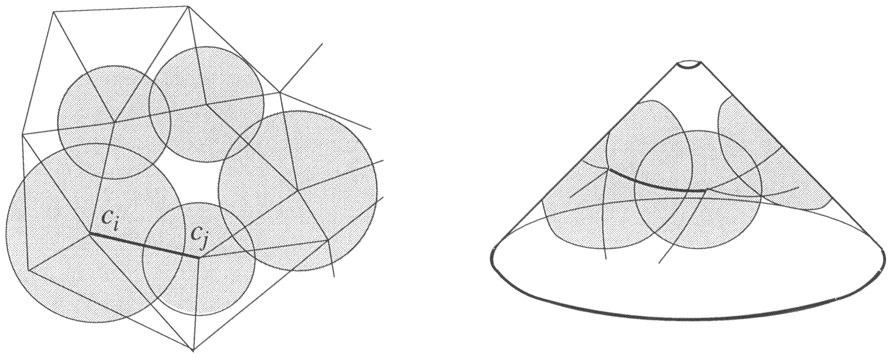

Fig. 8. Shrinking edge $c_i c_j$ wraps the plane into a cone.

disks, such that any set of disks with nonempty common intersection can be isometrically embedded in the plane in a way that is consistent with their intersection pattern. We can think of the disks as disjoint disks on the plane and the intersections as arcs *drawn* on the disks. The embedding requirement ensures that drawings are mutually consistent. In particular, all drawings of any given intersection must be congruent.

Another way to view a disk complex is as a two-dimensional Riemannian manifold with boundary, with a finite number of coordinate neighborhoods, each mapping to a disk.[1] Such a manifold has constant curvature zero.

Area still makes sense for a disk complex. Simply measure the area of each intersection of disks involving D_i within the drawing of D_i, and apply inclusion–exclusion (or alternatively weight a k-way intersection by $1/k$).

The restricted power diagram and the dual complex also carry over to disk complexes. Simply define the power distance from a point p in D_i to the center c_i as the power distance in the planar embedding of D_i. Since each cell in the restricted power diagram is a subset of a single disk, we do not need to define the power distance from a point p to the center of a disk that does not contain p. The dual complex $\mathcal{E}$ is defined exactly as before: it contains each disk center with a nonempty restricted power diagram cell, each edge between centers with cells sharing a side, and each triangle between centers with cells sharing a vertex.

We can speak of the lengths of edges within the dual complex: simply the distance between the disk centers in the isometric embedding of the two disks. Similarly, we can speak of angles within triangles of the dual complex, as the three disks defining the triangle must have a nonempty common intersection. We say that a disk D_i in a disk complex has a *covered perimeter* if, in the planar embedding of D_i, each point on the boundary of D_i is contained in some other (closed) disk D_j as well.

---

[1] Observe, however, that our definition does not allow all such manifolds. For example, because we do not allow disks to have intersections with more than one connected component, a pair of disks that wrap into a cylinder is disallowed.

Lemma 9. *Let $\mathcal{D}$ be a disk complex such that no disk has a covered perimeter. Then if the length of any edge in the dual complex is reduced by some positive amount, while maintaining the lengths of all other edges in the dual complex, the configuration of disks will remain a disk complex.*

Proof. Let the changing edge be $c_i c_j$. We must show how to maintain the drawings of disks. We distinguish three cases, depending on the number of triangles of the dual complex $\mathcal{E}$ bounded by $c_i c_j$.

First assume $c_i c_j$ bounds no triangles in $\mathcal{E}$. Since D_i and D_j intersect, $D_i \cup D_j$ has a planar embedding. Then no disk covers a vertex of the lune $D_i \cap D_j$. If disk D_k intersects $D_i \cap D_j$, then D_k intersects one, but not both, of $D_i \backslash D_j$ and $D_j \backslash D_i$. (If it intersected both, then $D_k \subset D_i \cup D_j$.) In the former case we consider D_k to be "attached to" D_i and in the latter case D_j. (D_i is itself attached to D_i.) When $|c_i c_j|$ changes, the arcs bounding D_i-attached disks change in the drawings of D_j-attached disks, and vice versa. Other drawings do not change.

We can think of the changes in the drawings of disks as occurring in a drawing of $D_i \cup D_j$. Each disk moves rigidly with either D_i or D_j, and changes occur only where the two sets overlap. In particular, the drawings of disks that do not intersect $D_i \cap D_j$ do not change, since these disks move rigidly with respect to all the disks they intersect.

Second assume $c_i c_j$ bounds exactly one triangle $c_i c_j c_k$. Then D_k covers one of the vertices of $D_i \cap D_j$, and no disk covers the other vertex. As in the first case, we may think of the changes as occurring in a planar drawing of $D_i \cup D_j \cup D_k$. Such a planar drawing must exist, because D_i, D_j, and D_k must have a nonempty common intersection for triangle $c_i c_j c_k$ to appear in $\mathcal{E}$.

If disk D_l, $l \neq k$, intersects $D_i \cap D_j$, it intersects one, but not both, of $D_i \backslash (D_j \cup D_k)$ and $D_j \backslash (D_i \cup D_k)$. (If it intersected both, then either D_l covers the uncovered vertex of $D_i \cap D_j$ or D_l covers all of $D_i \cap D_j \cap D_k$, a contradiction to $c_i c_j c_k$ being in the dual complex.) In the former case we consider D_l to be attached to D_i and in the latter case to D_j. We update the drawings of disks attached to D_i and D_j as in the first case.

We must also update the drawing of disk D_k. On this disk, the lunes $D_i \cap D_k$ and $D_j \cap D_k$ do not change size, but the angle $\angle c_i c_k c_j$ between them changes. (We can think of c_i as rotating about c_k.) The lunes of disks attached to D_i, and moreover the lunes intersecting those lunes and so forth, must rotate along with $D_i \cap D_k$. See Fig. 9. Since D_k does not have covered perimeter, these lunes do not go all the way around D_k. Since no disk intersects both of $D_i \backslash (D_j \cup D_k)$ and $D_j \backslash (D_i \cup D_k)$, no disk covers all of D_k's perimeter within $D_i \cap D_j$. Hence no lune is attached by a chain of lunes to both D_i and D_j, and we can update the drawing of D_k.

Finally assume $c_i c_j$ bounds two triangles $c_i c_j c_k$ and $c_i c_j c_l$. As in the two previous cases, we may think of the changes as occurring in a planar drawing of $D_i \cup D_j \cup D_k \cup D_l$. (Figure 10 below shows an example.) We say that a disk is attached to D_i if it intersects D_i minus the other three disks, and is attached to D_j if it intersects D_j minus the other three disks. Again it is not hard to confirm that no disk is attached to both.

The drawings of disks attached to D_i and D_j (including D_i and D_j themselves) are updated to reflect the new distance $|c_i c_j|$. The drawings of D_i, D_j, D_k, and D_l are updated to reflect the changed exterior angles in quadrilateral $c_i c_k c_j c_l$. Each of these disks is treated like D_k in the second case. □

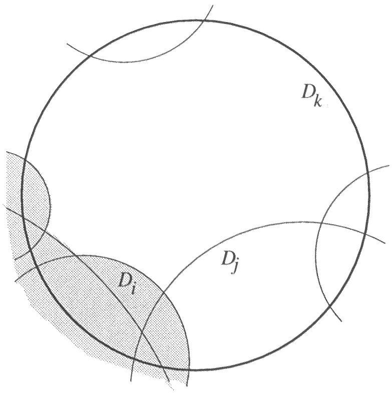

Fig. 9. To update the drawing of D_k, we rotate as a fixed set all lunes (shaded) attached to D_i by a chain of other lunes.

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 contraction. Then the area of the union of the disks cannot increase.*

Proof. We consider a small time interval and show that the area cannot increase; summing over all small intervals then gives the result. Since there are only a finite number of topologies for the dual complex $\mathcal{E}$, we may assume that $\mathcal{E}$ is combinatorially the same at the beginning and end of the small 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 8. Repeat this process, again choosing disks in arbitrary order, until no disk has a covered perimeter. Notice that each removal leaves a new hole, disjoint from previously formed holes. By Lemma 8 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 simulate the continuous contraction of $\mathcal{D}$ by smoothly shrinking edges of $\mathcal{E}^*$ one at a time. (Recall that the area depends on these edges alone.) Lemma 9 shows that shrinking an edge leaves us with a disk complex. 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.

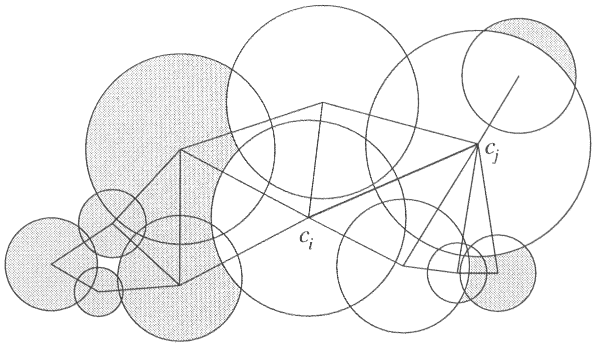

Fig. 10. When $c_i c_j$ shrinks the shaded areas move rigidly.

We assert that a (possibly infinite) sequence of edge shrinkings suffices 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 $\mathcal{E}^*$. This process cannot get caught in an infinite loop, because there is a measure of progress for each diagonal: the area of the union of the four disks in its quadrilateral, which is decreasing by Lemma 5.

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 shaded in Fig. 10) have unchanging intersection patterns, and we may think of them as moving rigidly with the outside edges of the triangles.

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 $\mathcal{E}^*$. Three disks forming a triangle in $\mathcal{E}^*$ must have a nonempty intersection, thus their union can be embedded in the plane. Lemma 4 now guarantees that the area of this union does not increase. Finally assume we shrink an edge $c_i c_j$ that bounds two triangles in $\mathcal{E}^*$, as in Fig. 10. The union of the four disks must be simply connected and hence can be immersed in the plane. Lemma 5 now handles this most difficult case. $\square$

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 contraction. 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 8; 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, as 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 contraction. Then the area of the intersection cannot decrease.*

Proof. Remove all disks from $\mathcal{D}$ that do not contribute a boundary to the intersection $I = D_1 \cap D_2 \cap \cdots \cap D_n$. 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 $\mathcal{F}$ includes all disk centers as exterior vertices.

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 8 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. Since I is nonempty, the union of the disks is simply connected and hence planar embeddable, so this process is conceptually easier than in the union case.

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 $D_1 \cap \cdots \cap D_4$; shrinking c_2c_4 grows c_1c_3 and hence increases the intersection.

There is one added twist in the case of the intersection: shrinking c_2c_4 can remove c_2c_4 from the dual complex. (In other words, shrinking c_2c_4 moves toward the ambiguous configuration rather than away.) If the exterior edges of $c_1c_2c_3c_4$ have already assumed their final lengths, however, then one of the diagonals must be too long and the other too short, so that c_2c_4 will reach its final length before the ambiguous configuration. It is not hard to show that if all exterior edges of $\mathcal{F}$ have assumed their final lengths, then some quadrilateral in $\mathcal{F}$ must have one diagonal too long (necessarily the one in $\mathcal{F}$, which has not yet changed length) and the other diagonal too short. Shrinking the too-long diagonal moves this quadrilateral to its final configuration, and the rest of the process succeeds by induction. □

8. Remarks

Observe that the proof of our main result, Theorem 2, holds for the more general context of disk complexes. (In this context, the distance between centers is only defined for intersecting disks.) It is actually possible that the question of Kneser and Thue Poulsen has an affirmative answer for globally planar configurations of disks but a negative answer for disk complexes.

Also observe that Theorem 2 extends to balls in higher dimensions, as long as all balls are moving parallel to a common (two-dimensional) plane. We conjecture that any continuous contraction 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 contraction into one-dimensional motions.)

Since our paper first appeared in the ACM Symposium on the Theory of Computing, Csikós [4] has extended the result to higher dimensions $\mathbb{R}^d$. Using the divergence theorem, he generalizes our Lemma 5 to the formula $dV/dt = \sum W_{ij} \cdot d|c_i c_j|/dt$, where V denotes the d-dimensional volume of the union of the balls, the sum is over all dual complex neighbors c_i and c_j, and W_{ij} represents the $(d - 1)$-dimensional volume of the shared boundary of restricted power diagram cells. This formula holds for the case of balls with smooth motion; for the case that the motion is merely continuous, Csikós gives a more technical proof.

Acknowledgments

We would like to thank Herbert Edelsbrunner for suggesting a possible connection between the disk-pushing problem and the dual complex and David Eppstein and David Goldberg for several helpful conversations. We would like to thank János Pach for telling us about Csikós's work, and Balázs Csikós for sending us his papers.

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 Discrete Mathematics and Theoretical Computer Science, Vol. 6, 1991, pp. 67–76.
3. B. Csikós. On the Poulsen–Kneser–Hadwiger conjecture. To appear in *Intuitive Geometry* (Budapest, 1995), Bolyai Society Mathematical Studies 6, Bolyai Society, Budapest, 1997.
4. B. Csikós. On the volume of the union of balls. To appear in *Discrete & Computational Geometry*.
5. H. Edelsbrunner. The union of balls and its dual shape. In *Proc*. 9*th ACM Symp*. *on Computational Geometry*, 1993, pp. 218–231.
6. H. Hadwiger. Ungelöste Probleme, Nr. 40. *Elemente der Mathematik* 11 (1956), 60–61.
7. N. Kazarinoff. *Geometric Inequalities*. Mathematical Association of America, Washington, DC, 1961.
8. V. Klee and S. Wagon. *Old and New Unsolved Problems in Plane Geometry and Number Theory*. Mathematical Association of America, Washington, DC, 1991.
9. M. Kneser. Einige Bemerkungen über das Minkowskische Flächenmaß. *Archiv der Mathematik* 6 (1955), 382–390.
10. E. Thue Poulsen. Problem 10. *Mathematica Scandinavica* 2 (1954), 346.

Received November 6, 1996, *and in revised form June* 16, 1997, *and September* 23, 1997.