

On the Union Complexity of Diametral Disks*

Boris Aronov[†] Muriel Dulieu[‡] Rom Pinchasi[§] Micha Sharir[¶]

Submitted: Sep 2, 2012; Accepted: Jun 2, 2013; Published: Jun 7, 2013

Abstract

Let S be a set of n points in the plane, and let $\mathcal{D}$ be an arbitrary set of disks, each having a pair of points of S as a diameter. We show that the combinatorial complexity of the union of the disks in $\mathcal{D}$ is $O(n^{3/2})$, and provide a lower bound construction with $\Omega(n^{4/3})$ complexity. If the points of S are in convex position, the upper bound drops to $O(n \log n)$.

The problem. Let S be a set of n points in the plane, and let E be a collection of unordered pairs of points of S . For each pair $\{a, b\} \in E$ let D_{ab} denote the disk with diameter ab ; we refer to D_{ab} as the *diametral disk* determined by $\{a, b\}$. Let $\mathcal{D}$ denote the resulting collection of disks, and let U denote their union. Let the *complexity* of U be measured by the number of its *vertices*, namely intersection points of two disk boundaries that lie on the boundary ∂U of U . In this paper we establish the following upper and lower bounds on the complexity of U .

*Work on this paper by Boris Aronov, Muriel Dulieu, and Micha Sharir has been supported by Grant 2006/194 from the U.S.-Israel Binational Science Foundation. Work on this paper by Boris Aronov has also been supported by NSA MSP Grant H98230-10-1-0210. Work on this paper by Boris Aronov and Muriel Dulieu has also been supported by NSF Grants CCF 08-30691 and CCF 11-17336. Work on this paper by Rom Pinchasi has been supported by ISF grant (grant No. 1357/12) and by BSF grant (grant No. 2008290) from the U.S.-Israel Binational Science Foundation. Work by Micha Sharir has also been supported by NSF Grant CCF-08-30272, by Grant 338/09 from the Israel Science Fund, by the Israeli Centers of Research Excellence (I-CORE) program (Center No. 4/11), and by the Hermann Minkowski-MINERVA Center for Geometry at Tel Aviv University. Most of work on this paper was done at the Centre Interfacultaire Bernoulli, EPFL, Lausanne, Switzerland, in September 2010, and was partially supported by the Swiss National Science Foundation. A preliminary version of this work was presented at the *2011 Fall Workshop on Computational Geometry*.

[†]Department of Computer Science and Engineering, Polytechnic Institute of NYU, Brooklyn, NY 11201-3840, USA; aronov@poly.edu.

[‡]Department of Computer Science and Engineering, Polytechnic Institute of NYU, Brooklyn, NY 11201-3840, USA; mdulieu@gmail.com.

[§]Faculty of Mathematics, The Technion, Haifa 32000, Israel; room@math.technion.ac.il.

[¶]School of Computer Science, Tel Aviv University, Tel Aviv 69978 Israel and Courant Institute of Mathematical Sciences, New York University, New York, NY 10012, USA; michas@post.tau.ac.il.

Theorem 1. *The maximum possible complexity of the union of any number of diametral disks, determined by any set of pairs of points from an n -point set in the plane, is $\Omega(n^{4/3})$ and $O(n^{3/2})$.*

If the points of S are in convex position, the upper bound on the union complexity improves considerably:

Theorem 2. *The complexity of the union of any number of diametral disks, determined by any set of pairs of points from an n -point set in convex position in the plane, is $O(n \log n)$.*

Preliminary properties. As is well known, the complexity of U , being the union of $|E|$ disks, is $O(|E|)$ [4], so the problem becomes more challenging as $|E|$ grows. We call a disk $D_{ab} \in \mathcal{D}$ *essential* if it appears on the boundary of the union U at some point which does not belong to any other disk; that is, it contributes a nonempty arc, or several arcs, to ∂U . We eliminate all non-essential disks, noting that every vertex of U is also a vertex of the union of the essential disks. (Notice that this latter union can have more features on its boundary, because the removal of the non-essential disks might expose new vertices which now lie on the union boundary.) On the other hand, as each essential disk contributes to the boundary, their number provides a natural lower bound, and, in view of the linear bound in [4], (up to a constant multiplicative factor) also an upper bound, on the union complexity. In other words, it suffices to establish the upper bounds in Theorems 1 and 2 on the number of essential disks in $\mathcal{D}$ and this holds trivially also for the lower bound in Theorem 1.

Lemma 3. *Fix a point $a \in S$ and consider the set E_a of all the neighbors of a in E . Let U_a denote the union of the disks D_{ab} , for $b \in E_a$, and let C_a denote the convex hull of $E_a \cup \{a\}$. Then (i) $C_a \subset U_a$ and (ii) the disks D_{ab} that appear on ∂U_a (in the above sense, of contributing at least one nonempty arc to ∂U_a) are precisely those for which b is a vertex of C_a .*

Proof. Let D_{ab} and D_{ac} be two disks in $\mathcal{D}$ with the common diametral endpoint a . The boundaries of these disks meet at a and at another point q which has to lie on the line passing through b and c ; q can lie either between b and c (as in Figure 1(a)) or past b or c (as in Figure 1(b)). In either case the entire triangle abc is covered by the union $D_{ab} \cup D_{ac}$.

To prove (i), we note that C_a is the union of all triangles ab_1b_2 , for $b_1, b_2 \in E_a$ (in fact, it suffices to take only those triangles for which b_1, b_2 are consecutive hull vertices), and the preceding argument implies that each of these triangles is contained in U_a .

To prove (ii), consider a point $b \in E_a$ which is not a vertex of C_a . There exist two points $b_1, b_2 \in E_a$ that are consecutive vertices of C_a such that b lies in the triangle ab_1b_2 ; see Figure 2. Let b' be the intersection point of the line ab with the line b_1b_2 ($b' = b$ when b lies on the segment b_1b_2). Note that D_{ab} is fully contained in $D_{ab'}$, which is fully contained in $D_{ab_1} \cup D_{ab_2}$. For the latter property, note that $\partial D_{ab'}$ passes through a and through the second intersection point q of ∂D_{ab_1} and ∂D_{ab_2} (because, by construction, $\angle aqb' = \pi/2$). Hence $D_{ab'}$ belongs to the pencil of disks determined by a and q . Since the center of $D_{ab'}$ lies between those of D_{ab_1} and D_{ab_2} , it is contained in the union of these two disks. Note




Figure 1: Two diametral disks with a common diametral endpoint.




Figure 2: Illustrating the proof that only hull neighbors of a can generate essential disks with a . In both figures, D_{ab} is the darkly shaded disk, and $D_{ab'}$ is the lightly-shaded disk.

that $\partial D_{ab'}$ meets $\partial(D_{ab_1} \cup D_{ab_2})$ only at a and q , and thus does not contribute any *arc* to ∂U_a .

To see that the converse also holds, consider a vertex b of C_a . Draw a supporting line ℓ to C_a through b , so that $\ell \cap C_a = \{b\}$, and mark on it the foot q of the perpendicular to ℓ from a . Clearly, q belongs to ∂D_{ab} , but it cannot belong to any other disk $D_{ab'}$ because, by construction, the angle $\angle aqb'$ is acute. This completes the proof of property (ii). $\square$

In other words, every essential disk in $\{D_{ab} \mid b \in E_a\}$ is such that b is a vertex of C_a (note that not all those disks necessarily contribute to the overall union U). We may therefore remove all other disks from $\mathcal{D}$. Hence, in what follows we will assume that the set E_a of all the neighbors in E of any $a \in S$ lie in convex position, and that either a lies in their hull or $E_a \cup \{a\}$ is also in convex position.

Proof of the upper bound in Theorem 1. We turn the set E into a *directed* geometric graph G . The vertices of G are the points in S . We orient an edge (drawn as a straight line segment) from a to b if $\{a, b\} \in E$, and some portion of the semicircle of ∂D_{ab} which lies counterclockwise from a to b appears on ∂U ; we refer to this semicircle as the *CCW-semicircle* of D_{ab} . Since we only deal with essential disks, each pair in E becomes either a single edge or a pair of anti-parallel edges of G . See Figure 3.




Figure 3: The directed graph G .

Without loss of generality we may assume that a constant fraction (in fact an eighth) of the directed edges of G form an angle of at most $\pi/8$ with the y -axis and point downwards. We delete all other edges from the graph G and bound the number of the remaining edges. This way we may lose a constant factor of at most eight in the upper bound. Let G' denote the new directed geometric graph.

We use the well-known fact that, for any directed graph $\tilde{G}(V, E)$ on n vertices, there exists a partition of V into sets A and B so that at least $|E|/4$ edges of $\tilde{G}$ emanate from vertices of A and terminate at vertices of B .¹

Applying this partitioning to G' , we obtain a new directed bipartite geometric graph H , without losing more than a factor of four in the number of edges; overall, we have $|E(G)| \leq 32|E(H)|$.

Let a *directed* $K_{2,2}$ be (a graph isomorphic to) the directed graph on four vertices a, b, u, v with edges (a, u) , (a, v) , (b, u) , (b, v) .

Lemma 4. H does not contain a copy of a directed $K_{2,2}$ as a subgraph.

Proof. Suppose to the contrary that H does contain such a subgraph. That is, there exist four points a, b, u, v in S so that all ordered pairs in $\{a, b\} \times \{u, v\}$ are in H .

¹To see this, assign independently each vertex $v \in V$ to A or to B with equal probability. An edge $(u, v) \in E$ will have $u \in A$ and $v \in B$ with probability $1/4$, so the expected number of such edges is $|E|/4$, so at least one assignment has at least $|E|/4$ edges directed from A to B .

Suppose first that a, b, u, v are in convex position. Since all edges of H point downwards, a and b are both higher than u and v , so ab and uv are edges of the hull. Hence, the edges of H consist of the two remaining hull edges and the two diagonals of the hull. That is, in this case au, av, bu, bv form a self-intersecting quadrangle (see Figure 4).

Otherwise, a, b, u, v are not in convex position. In this case, as is easily verified, au, av, bu , and bv form a non-convex simple quadrangle (see Figure 5).

We consider separately these two cases. We start with the more difficult case in which au, av, bu , and bv form a self-intersecting quadrangle, where some pair of edges, say av and bu , cross. This implies that, up to relabeling, the situation is as depicted in Figure 4. That is, a, b, v , and u are in convex position, appearing in this clockwise order along the boundary of their convex hull Q , and the CCW-semicircles of all four disks D_{au} , D_{av} , D_{bu} , and D_{bv} appear on the boundary of U .

For any pair of points p, q we denote by ℓ_{pq} the directed line passing through p and q and oriented from p to q . In the hypothetical configuration, the right halfplane of ℓ_{bv} , that is, the halfplane delimited by ℓ_{bv} and lying on its right side, contains a , u , and the CCW-semicircle of D_{bv} ; see Figure 4. By assumption, there exists a point q on this




Figure 4: Illustrating the impossibility of a self-intersecting copy of a directed $K_{2,2}$ in H .

semicircle which appears on the boundary of the union of all the disks and thus also on that of the four disks D_{au} , D_{av} , D_{bu} , and D_{bv} . The preceding arguments imply that q cannot lie inside the convex hull Q .

If q lies to the left of the line ℓ_{uv} (in the halfplane containing Q ; consider the point q in Figure 4—in the figure q also lies to the left of ℓ_{ab} , but the proof does not exploit this property), then we must have $\angle uqb = \angle uqv + \angle vqb > \angle vqb = \frac{\pi}{2}$; since the disk D_{bu} has bu as a diameter, it follows that q must lie inside D_{bu} , contrary to our assumption. A symmetric argument applies when q lies to the right of ℓ_{ab} (consider the point q' in Figure 4). It therefore remains to consider the case where q lies on the opposite sides of both lines, which can happen only when these lines meet at a point w to the right of ℓ_{bv} , and when q lies in the wedge W formed at w by the intersection of these two opposite halfplanes. Refer to q'' in Figure 4; to conform with the notation in the figure, we refer to this point as q'' in the remainder of the argument.

We claim that in this case the CCW-semicircle of D_{au} cannot appear on the union boundary. Indeed, suppose to the contrary that there exists a point p in that semicircle that lies on the union boundary. By assumption, p has to lie to the right of ℓ_{au} , and outside the triangle $q''bv$. Hence, p cannot lie both to the right of ℓ_{ab} and to the left of ℓ_{uv} .

If p lies to the left of ℓ_{ab} , then $\angle upb > \angle upa = \frac{\pi}{2}$, so p lies inside D_{bu} , a contradiction. The symmetric case where p lies to the right of ℓ_{uv} is handled similarly. These contradictions imply that au , av , bu , and bv cannot form a self-intersecting quadrangle.

Next we analyze the remaining case in which au , av , bu , and bv form a non-convex simple quadrangle.




Figure 5: The case where au , av , bu , and bv form a non-convex simple quadrilateral.

There are two cases to consider here.

Case 1. One of $\{a, b\}$, say b , is not a vertex of the convex hull of $\{a, b, u, v\}$. We assume the situation is as depicted in Figure 5(a). We claim that no point of the CCW-semicircle of D_{bu} can appear on the boundary of U , contradicting the fact that bu is a directed edge in G . To see this, let p be any point on the CCW-semicircle of D_{bu} . By definition of H , we have $\angle aub \leq \pi/4$ and $\angle uab < \angle uav \leq \pi/4$. It thus follows that $\angle uba = \pi - \angle aub - \angle uab > \pi/2$. Since $\angle uab \leq \pi/4$ it follows that a lies outside D_{bu} . Now if p and b lie on opposite sides of au then $\angle upa > \angle upb = \pi/2$ and then p lies inside D_{au} . If p and b lie on the same side of au then $\angle upa > \angle uba \geq \pi/2$, again implying that p lies inside D_{au} . In either case p cannot appear on ∂U , as claimed, and the resulting contradiction shows that Case 1 is impossible.

Case 2. One of $\{u, v\}$, say v , is not a vertex of the convex hull of $\{a, b, u, v\}$. We assume the situation is as depicted in Figure 5(b). We claim that no point of the CCW-semicircle of D_{av} can appear on the boundary of U , contradicting the fact that av is a directed edge in G . To see this, let p be any point on the CCW-semicircle of D_{av} .

By definition of H , we have $\angle uav \leq \pi/4$ and $\angle auv < \angle aub \leq \pi/4$. It thus follows that $\angle uva = \pi - \angle uav - \angle auv > \pi/2$. Since $\angle auv \leq \pi/4$ it follows that u lies outside D_{av} . Now if p and v lie on opposite sides of au then $\angle upa > \angle vpa = \pi/2$ and then p lies inside D_{au} . If p and v lie on the same side of au then $\angle upa > \angle uva \geq \pi/2$, again implying that p lies inside D_{au} . In either case p cannot appear on ∂U , as claimed, and the resulting contradiction shows that Case 2 is impossible too. $\square$

We conclude that the directed graph H does not contain a copy of a directed $K_{2,2}$. This is well known to imply (see the theorem of Kővari-Sós-Turán [5]) that H (and therefore G) can have at most $O(n^{3/2})$ edges. It follows that $\mathcal{D}$ contains at most $O(n^{3/2})$ essential disks, and thus its union has complexity $O(n^{3/2})$, completing the proof of the upper bound in Theorem 1.

Lower bound. We next show that the complexity of U can be $\Omega(n^{4/3})$ in the worst case. Informally, the construction consists of the following steps. (i) We take a construction of $n/2$ points and $n/2$ lines in the plane that have $\Omega(n^{4/3})$ incidences between them (see, e.g., [2]). (ii) We apply a transformation in which each of these lines is bent into a very flat parabola, and so that all incidences are preserved. (iii) We take a far-away point on each parabola, and our set is the collection of original points and the new far points on the parabolas. (iv) For each parabola, we take all the diametral disks, each formed by an original point lying on the parabola and by the far point representing the parabola. All the relevant points lie in convex position (on the parabola), and we show that all the corresponding disks are essential within their own union. (v) Finally, we argue that if the far points are indeed sufficiently far and if the parabolas are sufficiently flat, all the constructed disks are essential within the overall union. Since we have $\Omega(n^{4/3})$ disks, the lower bound follows.

In more detail, we construct the following set S of $\Theta(n)$ points. One half of the set, S_1 , consists of the vertices of the $\sqrt{n} \times \sqrt{n}$ integer grid, centered at the origin (for this, assume, without loss of generality, that $\sqrt{n}$ is an odd integer). As shown by Erdős and others (see, e.g., [2]), there exists a set L of $m = \Theta(n)$ distinct lines, each containing $\Theta(n^{1/3})$ points of S_1 . By applying an appropriate projective transformation, which only slightly moves the points of S_1 , we may assume that no two lines in L are parallel. We then shrink and rotate the plane, so that S_1 is contained in the unit disk around the origin, and so that no line of L is vertical. All these transformations preserve the $\Theta(n^{4/3})$ incidences between the points of S_1 and the lines of L .

We next take a circle C centered at the origin with a sufficiently large radius r , and form the set S_2 of the intersection points of C with the lines of L . Each line in L intersects C at two points, but we put in S_2 only the point with positive x -coordinate (when r is sufficiently large there will be exactly one such point).

Lemma 5. *For all sufficiently large values of r the following condition holds: For any pair q_1, q_2 of distinct points of S_2 and for every point p in the unit disk around the origin, the disk D_{pq_1} does not intersect the disk of radius $5/r$ around q_2 .*

Proof. As no two lines in L are parallel, for sufficiently large r the distance between every pair of points in S_2 can be assumed to be larger than 10.

We need to lower bound the distance from q_2 to the center c of D_{pq_1} . Write $p = (2a, 2b)$, with $a^2 + b^2 < \frac{1}{4}$. Assume, without loss of generality, that $q_1 = (r, 0)$, and write $q_2 = (r \cos \theta, r \sin \theta)$, with $\theta > 0$. We then have $c = (\frac{r}{2} + a, b)$.

We need a lower bound for $\|cq_2\| - \|cq_1\|$. For this we write

$$\|cq_2\| - \|cq_1\| = \frac{\|cq_2\|^2 - \|cq_1\|^2}{\|cq_2\| + \|cq_1\|} \geq \frac{\|cq_2\|^2 - \|cq_1\|^2}{4r}.$$

Substituting the above expressions for c , q_1 , and q_2 , we obtain

$$\begin{aligned} \|cq_2\| - \|cq_1\| &\geq \frac{1}{4r} ((r(\cos \theta - 1/2) - a)^2 + (r \sin \theta - b)^2 - (r/2 - a)^2 - b^2) \\ &= \frac{1}{4r} (r^2((\cos \theta - 1/2)^2 + \sin^2 \theta - 1/4) - 2r(a(\cos \theta - 1/2) + b \sin \theta - a/2)) \\ &= \frac{r}{4}(1 - \cos \theta) + \frac{1}{2}(a(1 - \cos \theta) - b \sin \theta). \end{aligned}$$

Clearly, when r is sufficiently large, the term $\frac{1}{2}a(1 - \cos \theta)$ is negligible compared with the leading term, so we can write, say,

$$\|cq_2\| - \|cq_1\| \geq \frac{2r \sin^2(\theta/2)}{5} - \frac{1}{4} \sin \theta.$$

We have assumed that $\|q_1 q_2\| \geq 10$, so $\sin(\theta/2) = \|q_1 q_2\|/(2r) \geq 5/r$. In this case the second term in the above expression is smaller than half the first term, so we have

$$\|cq_2\| - \|cq_1\| \geq \frac{r \sin^2(\theta/2)}{5} \geq \frac{5}{r},$$

completing the proof. □

We now put $S := S_1 \cup S_2$; we have $|S| = \Theta(n)$. See Figure 6.

The construction so far is not the end of the story. We next distort $S = S_1 \cup S_2$ by mapping each point $p = (x, y)$ to the point $p^* = (x, y + \varepsilon x^2)$, where $\varepsilon > 0$ is a sufficiently small parameter. This turns each line $\ell = \{y = a_\ell x + b_\ell\} \in L$ into the parabola $\ell^* = \{y = a_\ell x + b_\ell + \varepsilon x^2\}$, which is still incident to $\Theta(n^{1/3})$ points of the perturbed copy S_1^* of S_1 . The final step in the construction uses the set $S^* = S_1^* \cup S_2^*$ and defines the following set $\mathcal{D}$ of diametral disks on S^* . For each line $\ell \in L$ we add to $\mathcal{D}$ all diametral disks of the form $D_{p^* q^*}$, where q^* is the point $S_2^* \cap \ell^*$ and where $p^* \in S_1^* \cap \ell^*$. Clearly, $|\mathcal{D}| = \Theta(n^{4/3})$. We claim that all disks in $\mathcal{D}$ are essential, which readily implies the asserted lower bound.

To show that all disks in $\mathcal{D}$ are essential, consider a line $\ell \in L$ and its associated parabola ℓ^* . Let $p_1^*, \dots, p_k^*$ denote the points of S_1^* on ℓ^* , appearing there in this left-to-right order, and let q^* denote the single point of S_2^* on ℓ^* ; see Figure 7. Since all these points are in convex position, it follows from Lemma 3 that all the disks $D_{p_i^* q^*}$, for




Figure 6: The (still unperturbed) lower bound construction.

$i = 1, \dots, k$, are essential within their own union U_{q^*} . Moreover, if ε is sufficiently small, the boundaries of these disks intersect only within the disk $\tilde{D}_q$ of radius $1/r$ around q^* (when $\varepsilon = 0$, all these disks are tangent to each other at $q^* = q$). Hence, each of these disks, except for the largest one, can appear along ∂U_{q^*} (and therefore also along ∂U) only within $\tilde{D}_q$. By Lemma 5, these appearances are disjoint from all the disks constructed on the other points of S_2^* , and therefore the corresponding disks are all essential in U .

This establishes the lower bound, and thereby completes the proof of Theorem 1. $\square$



Figure 7: The convex chain $p_1^*, \dots, p_k^*, q^*$ lies on a common parabola.

Proof of Theorem 2. Consider next the case where the points of S are in convex position, and denote by $\text{CH}(S)$ their convex hull. As in the general case, we are also given a set E of unordered pairs of points of S , and a corresponding set $\mathcal{D}$ of diametral disks D_{ab} , for $\{a, b\} \in E$. As before, we let U denote the union of these disks.

The proof is a variation of a classical result of Füredi [3] on the number of repeated distances between n points in convex position, which reduces the problem to that of

bounding the number of 1s in a 0 – 1 matrix which does not contain a copy of a certain 2×3 forbidden submatrix. We apply a different reduction to the very same problem, and obtain the same upper bound, namely $O(n \log n)$, as in [3].

To do so, we construct, as above, a directed geometric graph G on S so that, for each essential disk $D_{au} \in \mathcal{D}$ for which some portion of its CCW-semicircle from a to u appears on ∂U , we add the directed edge (a, u) to G .




Figure 8: Left: (a, u) is an acute chord. Right: (a, u) is an obtuse chord.

We distinguish between three types of edges (a, u) of G ; we refer to them collectively as *chords* (of $\text{CH}(S)$):

- (i) **An acute chord:** Every pair of lines ℓ_a, ℓ_u that support $\text{CH}(S)$ at a and u , respectively, either do not meet at all in the right halfplane of ℓ_{au} , or meet there at an acute angle (the angle at which the intersection point of ℓ_a and ℓ_u sees the segment au). See Figure 8(left).
- (ii) **An obtuse chord:** Every pair of supporting lines ℓ_a, ℓ_u , as above, meet in the right halfplane of ℓ_{au} at an obtuse angle (at which their intersection point sees au). See Figure 8(right).
- (iii) **A right-angle chord:** There exist a pair of supporting lines ℓ_a, ℓ_u , as above, that meet in the right halfplane of ℓ_{au} at a right angle. As is easily verified, a chord which is neither acute nor obtuse must be a right-angle chord.

The number of right-angle chords is $O(n)$, as easily follows by an argument based on rotating calipers [7]. (As we rotate a right-angle wedge so that it contains $\text{CH}(S)$ and touches it at two points, the right-angle chords are exactly those which connect the appropriately directed pairs of contact points, and these pairs change only $O(n)$ times.)

To obtain an upper bound on the number of acute and obtuse chords we employ an argument similar to that used by Füredi [3], which involves forbidden submatrices in a 0 – 1 matrix. Specifically, the argument is based on the following lemma.

Lemma 6 (Füredi [3]). *Let A be an $n \times n$ 0 – 1 matrix which does not contain the 2×3 submatrix*

$$A_0 = \begin{pmatrix} 1 & 1 & * \\ 1 & * & 1 \end{pmatrix},$$

where $$ denotes any value. That is, there do not exist a pair of rows $i < j$ and a triple of columns $k < \ell < m$, such that $A_{ik} = A_{i\ell} = A_{jk} = A_{jm} = 1$. Then the number of 1s in A is $O(n \log n)$.*

Corollary 7. *Let G be a directed graph on a set S of n points in convex position in the plane. If G does not contain five vertices a_1, a_2, u_3, u_2, u_1 , appearing in this counterclockwise order along the convex hull of S , such that (a_1, u_1) , (a_2, u_1) , (a_1, u_2) , and (a_2, u_3) are all directed edges in G , then G has $O(n \log n)$ edges.*

Proof. Enumerate the points of S in counterclockwise order along the hull as $p_1, \dots, p_n$, starting at an arbitrary point, and construct two 0 – 1 matrices $A^{(1)}$, $A^{(2)}$, which are variants of the adjacency matrix of G , where $A_{i,j}^{(1)} = 1$ iff (p_i, p_j) is an edge of G , and $i < j$, and $A_{n+1-i,j}^{(2)} = 1$ iff (p_i, p_j) is an edge of G , and $i > j$. Hence, $A^{(1)}$ is an upper triangular matrix, and the nonzero entries of the matrix $A^{(2)}$ lie above the reflected main diagonal. The condition imposed on G is easily seen to imply the nonexistence of A_0 in $A^{(1)}$, and symmetric considerations imply that $A^{(2)}$ also does not contain A_0 as a submatrix. By Lemma 6, each of $A^{(1)}$ and $A^{(2)}$ has $O(n \log n)$ 1s. Since these matrices record all edges of G , the corollary follows. $\square$

We are now ready to bound the number of acute chords. We consider the subgraph G_a of G consisting only of the acute chords, and claim that G_a does not contain as a subgraph the directed graph in the statement of Corollary 7. That is, there do not exist five vertices a_1, a_2, u_3, u_2, u_1 of $\text{CH}(S)$ in this counterclockwise order along $\partial \text{CH}(S)$, so that (a_1, u_1) , (a_2, u_1) , (a_1, u_2) , and (a_2, u_3) are all directed edges in G_a . See Figure 9(left) for an illustration of this forbidden subgraph.

To prove the claim, assume to the contrary that G_a does contain such a subgraph. Consider the lines $\ell_{a_1 a_2}$ and $\ell_{u_1 u_2}$, and assume that they meet at a point w that lies to the right of $\ell_{a_1 u_1}$; the analysis is much simpler if they do not meet in the right halfplane of $\ell_{a_1 u_1}$ and we handle this case separately below. Arguing as in the proof of the upper bound in Theorem 1, the existence of the three chords (a_1, u_1) , (a_1, u_2) , and (a_2, u_1) in G_a implies that any point q on the CCW-semicircle of $D_{a_1 u_1}$ which appears on the union boundary must lie in the wedge W which is the intersection of the right halfplane of $\ell_{a_1 a_2}$ and the left halfplane of $\ell_{u_1 u_2}$ (and therefore has w as an apex; see Figure 9(left)). We claim that the angle at which w sees $a_1 u_1$ is acute. This in turn will imply that q must also see $a_1 u_1$ at an acute angle and thus cannot lie on $\partial D_{a_1 u_1}$, contrary to assumption.

To prove that the angle at which w sees $a_1 u_1$ is acute we make use of the existence of the fourth acute chord (a_2, u_3) . Take any supporting line ℓ_{a_2} to $\text{CH}(S)$ at a_2 and any supporting line ℓ_{u_3} to $\text{CH}(S)$ at u_3 . Because a_1, a_2, u_3, u_2 , and u_1 are in convex position in this counterclockwise order on $\text{CH}(S)$, it follows that ℓ_{a_2} and ℓ_{u_3} must meet at a point z that lies in the wedge opposite to W . By definition, z sees the segment $a_2 u_3$ at an




Figure 9: Left: An impossible realization of A_0 in the acute case. Right: The obtuse case.

acute angle. This implies that z sees the segment a_1u_1 at an acute angle, which in turn implies that w sees the segment a_1u_1 at an acute angle, thereby reaching the desired contradiction.

To complete the analysis, notice that if $\ell_{a_1a_2}$ and $\ell_{u_1u_2}$ do not meet in the right halfplane of $\ell_{a_1u_1}$, then the wedge W does not exist, leaving no place for the point q , and thereby leading to an immediate contradiction.

Since G_a does not contain the forbidden subgraph of Corollary 7, we conclude that it has at most $O(n \log n)$ edges.

We next consider the case of obtuse chords, and handle it using a “mirror image” of the preceding argument. Specifically, we now consider the directed subgraph G_o of G which consists only of the obtuse chords, and claim that G_o does not contain as a subgraph the “reflection” of the subgraph in the statement of Corollary 7. That is, there do not exist five vertices a_1, a_2, u_3, u_2, u_1 of $\text{CH}(S)$ in this *clockwise* order along $\partial \text{CH}(S)$, so that (a_1, u_1) , (a_2, u_1) , (a_1, u_2) , and (a_2, u_3) are all directed edges in G_o . See Figure 9(right) for an illustration of this forbidden subgraph.

Assume to the contrary that G_o does contain a copy of the above subgraph. Consider the lines $\ell_{a_2a_1}$ and $\ell_{u_2u_1}$, and assume that they meet at a point w in the right halfplane of $\ell_{a_1u_1}$ (as the analysis will imply, they must indeed meet in this manner). By arguments similar to those in the proof of the upper bound in Theorem 1, the existence of the three chords (a_1, u_1) , (a_1, u_2) , and (a_2, u_1) in G_o implies that any point q on the CCW-semicircle of $D_{a_1u_1}$ which appears on the union boundary must lie in the triangle bounded by $\ell_{a_1a_2}$, $\ell_{u_1u_2}$, and $\ell_{a_1u_1}$. We claim that the angle at which w sees the segment a_1u_1 is obtuse, which in turn will imply that q must also see a_1u_1 at an obtuse angle and thus cannot lie on $\partial D_{a_1u_1}$, contrary to assumption.

To prove that the angle at which w sees a_1u_1 is obtuse we make use of the existence of the fourth obtuse chord (a_2, u_3) . Take any supporting line ℓ_{a_2} to $\text{CH}(S)$ at a_2 and any supporting line ℓ_{u_3} to $\text{CH}(S)$ at u_3 . By definition, ℓ_{a_2} and ℓ_{u_3} meet at a point z , and $\angle a_2zu_3$ is obtuse. Since the points a_2, a_1, u_1, u_2 , and u_3 lie in convex position in that counterclockwise order on the boundary of $\text{CH}(S)$, it follows that w lies in the triangle bounded by ℓ_{a_2}, ℓ_{u_3} , and a_2u_3 . We have $\angle a_1wu_1 > \angle a_1wu_3 = \angle a_2wu_3 > \angle a_2zu_3$. Hence $\angle a_1wu_1$ is obtuse, thereby reaching the desired contradiction. (Note that, as mentioned above, the analysis does indeed imply that $\ell_{a_1a_2}$ and $\ell_{u_1u_2}$ meet to the right of $\ell_{a_1u_1}$.)

Again, since G_o does not contain the (reflection of the) forbidden subgraph of Corollary 7, we conclude that it has at most $O(n \log n)$ edges. Hence G has $O(n \log n)$ edges, and this implies the asserted bound on the number of essential disks in $\mathcal{D}$, and thus also on the complexity of the union U .

This completes the proof of Theorem 2. $\square$

Remark. Note that the proof, in both acute and obtuse cases, actually goes through if we replace the chord (a_2, u_3) by a chord (a_3, u_3) , where a_3 is any point lying between a_2 and u_3 . That is, we also exclude the existence of the submatrix $\begin{pmatrix} 1 & 1 & * \\ 1 & * & * \\ * & * & 1 \end{pmatrix}$ in both $A^{(1)}$ and $A^{(2)}$. Unfortunately, this in itself does not imply a better upper bound, as already noted by Pettie [6], citing previously known constructions.

Discussion. The main contribution of this paper is to show that the maximum possible complexity of the union of any collection of diametral disks defined over a set of n points in the plane, is neither quadratic nor linear. Still there is a significant gap between our lower and upper bounds, and the main obvious open problem is to close this gap. When the given points are in convex position, the upper bound drops to $O(n \log n)$, but we do not know whether the bound is worst-case tight. Therefore, here too there is a gap to close. A linear lower bound on the number of essential disks is trivial. Is there a superlinear construction?

If, instead of requiring the disks in $\mathcal{D}$ to be diametral, we only insist that every disk in $\mathcal{D}$ pass through a (distinct) pair of points of S , the number of essential disks might be quadratic in the worst case. To see this, take S to be a set of n points in general position in a unit disk, and let γ be a concentric circle with generic radius larger than 2. For each pair a, b of distinct points in S , draw a disk D_{ab} which passes through a and b and is tangent to γ (from the inside). The general position of S and a generic choice of the radius of γ ensure that all the points of tangency between these disks and γ , which clearly lie on the boundary of the union of the disks, are distinct, so all the disks are essential.

Hence, some restriction must be imposed on the disks of $\mathcal{D}$ to obtain subquadratic behavior. Does the union complexity remain subquadratic when (a) in each disk D_{ab} the chord ab has the same central angle, or when (b) these central angles are bounded away from 0 (by an angle independent of n) or sufficiently close to π ?

More ambitiously, does our upper bound continue to hold when one replaces disks by homothetic copies of some fixed convex shape K ? Here one might define, for a given pair of points a, b , the corresponding diametral copy K_{ab} of K to be the smallest homothetic

copy of K that contains both a and b . What is the complexity of the union of any subset of diametral copies of K determined by pairs of points from a set of n points in the plane?

The motivation for studying the problem of the union of diametral disks came from a recent study of a subset of the authors [1] on witness Gabriel graphs. Specifically, a (*negative*) *witness Gabriel graph* for a pair of point sets P and Q in the plane is the geometric graph with vertex set P , where $p, p' \in P$ are connected by an edge if the diametral disk $D_{pp'}$ contains no points of Q . We were considering the problem of testing whether a given geometric graph G , on a given vertex set P and a given witness set Q , is indeed a witness Gabriel graph. For this one needs to verify that the union of the diametral disks corresponding to the edges of G does not contain any point of Q . A natural question that then arises is to obtain bounds on the complexity of such a union, which is what we have done in this paper.

We note that applying the (upper) bounds that we have derived in the context of witness Gabriel graphs requires some care. Two issues arise: (a) The size of G may be large, so unless it is encoded in some compact form, the bound on the union complexity may not lead to a comparably efficient algorithm. (b) One also needs to tackle the complementary problem, of verifying that each disk corresponding to a non-edge of G contains a point of Q .

References

- [1] B. Aronov, M. Dulieu, and F. Hurtado, Witness Gabriel graphs, *Comput. Geom. Theory Appl.*, in press, see also <http://dx.doi.org/10.1016/j.comgeo.2011.06.004>.
- [2] H. Edelsbrunner, *Algorithms in Combinatorial Geometry*, Springer-Verlag, Heidelberg, 1987.
- [3] Z. Füredi, The maximum number of unit distances in a convex n -gon, *J. Combinat. Theory, Ser. A* 55 (1990), 316–320.
- [4] K. Kedem, R. Livné, J. Pach, and M. Sharir, On the union of Jordan regions and collision-free translational motion amidst polygonal obstacles. *Discrete Comput. Geom.* 1(1) (1986), 59–71.
- [5] T. Kővari, V. T. Sós, and P. Turán, On a problem of K. Zarankiewicz, *Colloquium Mathematicum* 3 (1954), 50–57.
- [6] S. Pettie, Degrees of non-linearity in forbidden 0–1 problems, *Discrete Math.* 311 (2011), 2396–2410.
- [7] G.T. Toussaint, Solving geometric problems with the rotating calipers, *Proc. IEEE MELECON'83*, Athens, Greece, (1983), pp. A10.02/1–4.