

Finding the Maximal Empty Disk Containing a Query Point*

Haim Kaplan[†] Micha Sharir[‡]

January 1, 2012

Abstract

Let P be a set of n points in the plane. We present an efficient algorithm for preprocessing P , so that, for a given query point q , we can quickly report the largest disk that contains q but its interior is disjoint from P . The storage required by the data structure is $O(n \log n)$, the preprocessing cost is $O(n \log^2 n)$, and a query takes $O(\log^2 n)$ time.

*Work by Haim Kaplan was partially supported by Grant 2006/204 from the U.S.–Israel Binational Science Foundation, and by Grant 822/10 from the Israel Science Fund. Work by Micha Sharir was partially supported by NSF Grant CCR-08-30272, by Grant 2006/194 from the U.S.–Israel Binational Science Foundation, by Grant 338/09 from the Israel Science Fund, and by the Hermann Minkowski–MINERVA Center for Geometry at Tel Aviv University.

[†]School of Computer Science, Tel Aviv University, Tel Aviv 69978, Israel. E-mail: haimk@post.tau.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. E-mail: michas@post.tau.ac.il

1 Introduction

We consider the problem of preprocessing a given set P of n points in the plane into a data structure such that for any given query point q we can efficiently find the largest disk that contains q but its interior does not contain any point of P . This problem models scenarios where one wants to efficiently estimate the size of the “free space” (free of points of P , that is) surrounding a query point and locate the center of this free space. Concrete examples of such a scenario are: (i) We want to transmit a message over the largest possible area (which is a disk if propagation proceeds uniformly in all directions) such that we are sure that q receives the message but none of the eavesdroppers at the points of P hears it. (ii) We want to position a sprinkler to water the largest possible disk that includes our favorite plant at q but does not include any of the locations in P , which are facilities that we do not want to get wet.

Note that it suffices to consider query points q that lie inside the convex hull $CH(P)$ of P because otherwise there are arbitrarily large disks containing q and disjoint from P .

Our result. We present an efficient algorithm that requires $O(n \log n)$ storage, $O(n \log^2 n)$ preprocessing time, and can answer a query in $O(\log^2 n)$ time.

Background. To the best of our knowledge this problem was first introduced by Augustine *et al.* [1], who design a data structure that requires $O(n^2)$ space, $O(n^2 \log n)$ preprocessing time, and can answer a query in $O(\log n)$ time.

A much easier problem, which has been studied much earlier, is that of finding the largest empty disk for a set P of n points in the plane (where, to make sense of the problem, the center of the disk has to lie, say, inside the convex hull of P —recall the remark made earlier concerning points outside the hull). This problem arises in several areas such as data mining, database management, visualization, and VLSI design, and it has a very simple solution. Indeed, the largest empty disk has to be centered at a vertex of the Voronoi diagram of P , or at a crossing of an edge of the diagram with the convex hull boundary, and therefore can be found in $O(n \log n)$ time.

We note that Augustine *et al.* [1] also consider a similar problem, where the goal is to report the largest-area axis-parallel rectangle that contains the query point, is contained in some fixed bounding box B , and its interior is disjoint from P . Their solution for this variant also requires nearly quadratic storage and preprocessing; in a companion paper [4] (see also [3]), we present a more efficient solution, which requires only nearly linear storage and preprocessing, for that latter problem.

2 Preliminaries and an easy quadratic solution

Let P be a set of n points in the plane. We call a disk whose interior is disjoint from P a *P -empty disk*. We wish to preprocess P into a data structure so that, given a query point q , we can efficiently find the largest P -empty disk containing q . The following easy lemma

completes the observation made earlier, that it suffices to consider only query points q that lie inside the convex hull $CH(P)$ of P .

Lemma 2.1. *There are arbitrarily large P -empty disks containing q if and only if q does not lie in the interior of the convex hull $CH(P)$ of P .*

Proof. If q does not lie in the interior of $CH(P)$ then there is a line ℓ through q that defines an halfplane whose interior is P -empty; clearly there are arbitrarily large P -empty disks in this halfplane tangent to ℓ at q .

For the converse statement, let B be some fixed disk containing $P \cup \{q\}$. Clearly if there are arbitrarily large P -empty disks containing q then there are arbitrarily large P -empty disks containing q on their boundary. For any such disk D , $D \cap B$ contains the intersection of B with a wedge W whose apex is q and whose bounding rays go through the points in $\partial B \cap \partial D$. The wedge W is P -empty ($W \cap B$ is P -empty since it is contained in D and $W \setminus B$ is empty since there are no points of P outside B), and its opening angle tends to π as D grows. In the limit, W becomes a halfplane containing q and openly disjoint from P . Hence q does not lie in the interior of $CH(P)$. \square

Let $\text{Vor}(P)$ denote the Voronoi diagram of P .

For a point q in the interior of $CH(P)$, let $D_{\max}(q)$ denote an arbitrary largest P -empty disk containing q , and let $c_{\max}(q)$ denote its center. (In general this disk might not be unique. For example we may have two congruent Delaunay disks, that is, maximal P -empty disks centered at vertices of $\text{Vor}(P)$, both containing q , and larger than any other P -empty disk containing q .)

The analysis proceeds through the following stages.

Lemma 2.2. *Let q be in the interior of $CH(P)$. Then $c_{\max}(q)$ lies on a Voronoi edge or at a Voronoi vertex.*

Proof. If D is a P -empty disk that contains q and is centered at some point c not on the Voronoi diagram then D touches at most one point of P . If it does not touch any point of P , we can get a larger P -empty disk containing q by expanding D about c until it touches some (unique) point $a \in P$. If D already touches a point a of P then we can get a larger P -empty disk containing q by moving the center of D along the ray from a through c , away from a , keeping it touching a , until it touches another point b of P . This has to occur, for otherwise we get arbitrarily large P -empty disks containing q , contrary to Lemma 2.1. It follows that a disk which is not centered at a Voronoi edge or vertex is not a maximal P -empty disk containing q . \square

Suppose that $c_{\max}(q)$ lies in the relative interior of a Voronoi edge e_{ab} defined by two sites $a, b \in P$. There is at least one direction along e_{ab} so that we can increase the size of the disk by moving its center in that direction away from $c_{\max}(q)$ while keeping it touching a and b and P -empty. Since $D_{\max}(q)$ is the largest P -empty disk containing q , we conclude that in this case q must lie on the boundary of $D_{\max}(q)$. Moreover, this boundary is split into two




Figure 1: The case where $c_{\max}(q)$ lies in the relative interior a Voronoi edge.

arcs by a and b , and q has to lie on the *smaller* arc; see Figure 1. Note that in this case we have $\angle aqb > \pi/2$. (In particular, ab cannot be a diameter of $D_{\max}(q)$.)

Continuing the analysis of this scenario, we note that in this case $c_{\max}(q)$ is equally nearest to a , b , and q . That is, $c_{\max}(q)$ is a vertex of the Voronoi cell $V(q)$ of q in the augmented Voronoi diagram $\text{Vor}(P \cup \{q\})$. Then a and b are two consecutive neighbors of q ; that is, the Voronoi edges e_{aq} , e_{bq} are consecutive along $\partial V(q)$. The angle between these edges, within $V(q)$, is $\pi - \angle aqb$, and is therefore smaller than $\pi/2$. See Figure 2. Since the sum of the exterior angles of a convex polygon is 2π , it follows that $V(q)$ has at most three vertices that can be the location of $c_{\max}(q)$. We summarize the observations made so far in the following lemma.

Lemma 2.3. *If $c_{\max}(q)$ lies in the relative interior of a Voronoi edge e_{ab} then q lies on the smaller arc connecting a and b along $\partial D_{\max}(q)$. In this case $c_{\max}(q)$ is a vertex of the Voronoi cell $V(q)$ of q in the Voronoi diagram $\text{Vor}(P \cup \{q\})$. Furthermore, $c_{\max}(q)$ is one of the at most three vertices of $V(q)$ with an acute interior angle.*

The preceding analysis therefore suggests the following high-level approach: Identify, without constructing $V(q)$ explicitly (whose complexity can be as high as $\Theta(n)$), the at most three vertices of this cell which might be the location of $c_{\max}(q)$, compute, in $O(1)$ time, the radii of the corresponding largest P -empty disks centered at these vertices, and output the center with the largest radius. This is however only one part of the solution, because we also have to consider Voronoi vertices of $\text{Vor}(P)$ as possible locations for $c_{\max}(q)$.

2.1 A quick-and-dirty interim solution

Before continuing towards an efficient solution, we can already exploit the analysis given so far, to obtain the following quick-and-dirty solution. Consider a query point q , the cell $V(q)$ of q in $\text{Vor}(P \cup \{q\})$, and its at most three special vertices having acute interior angles, which




Figure 2: The cell $V(q)$ in $\text{Vor}(P \cup \{q\})$, superimposed on $\text{Vor}(P)$. The angle between e_{aq} and e_{bq} is smaller than $\pi/2$.

are the candidate placements for $c_{\max}(q)$ (other than Voronoi vertices of $\text{Vor}(P)$). If we move q , $V(q)$ will vary continuously, and its combinatorial structure (namely, the sequence of its edges and the identity of the three acute vertices) will remain fixed until q reaches some critical placement in which one of these attributes changes. It is easy to verify that the sequence of edges of $V(q)$ can change only when one of its vertices is a vertex of the non-augmented diagram $\text{Vor}(P)$; that is, when q lies on the boundary of the Delaunay disk which is centered at that vertex (and circumscribes its dual Delaunay triangle). The “acuteness” of a vertex of $V(q)$, lying on an original Voronoi edge e_{ab} , can change only when q lies on the boundary of the diametral disk determined by a and b , but only if this disk is P -empty. (Note that the diametral disk is P -empty if and only if ab is an edge of the *Gabriel graph* of P .)

So a quick-and-dirty solution is to draw the $O(n)$ disks of the above two types (Delaunay disks and diametral disks determined by edges of the Gabriel graph), form their arrangement \mathcal{A} , and store with each face f of \mathcal{A} the list $L(f)$ of the at most three acute vertices of $V(q)$, for any $q \in f$ (by the preceding argument, the list is the same for all $q \in f$). We also store with each f the largest Delaunay disk, $D(f)$, centered at a vertex of $\text{Vor}(P)$, which contains f . We can compute the lists $L(f)$ and the disks $D(f)$ incrementally, observing that it is straightforward to update $L(f)$ in $O(1)$ time as we cross from one face f to a neighboring face f' . Updating $D(f)$ can be done in $O(\log n)$ time, by dynamically maintaining the list of all Delaunay disks containing the present face f , sorted in the decreasing order of their radii. We then preprocess \mathcal{A} for efficient point location, and obtain the desired data structure. It requires $O(n^2)$ storage, $O(n^2 \log n)$ preprocessing, and a query takes $O(\log n)$ time: Locate the query point q in \mathcal{A} , retrieve the list $L(f)$ of the face f containing q , and search each of its (at most three) elements for a possible location of $c_{\max}(q)$. We return the largest among the maximal P -empty disks centered at the vertices defined by $L(f)$ and the disk $D(f)$. Note

that the elements of $L(f)$ are only stored symbolically, because the actual vertices of $V(q)$ vary continuously as q moves inside f ; still, the combinatorial identification of each of these elements allows us to test it in $O(1)$ time for any concrete query point q .

This straightforward approach matches the performance of the algorithm in [1].

3 A more efficient data structure

3.1 Efficient processing of Delaunay disks

It is relatively easy to improve the algorithm that finds the largest Delaunay disk containing q , using the approach given in [1]. That is, we sort the Voronoi vertices v of $\text{Vor}(P)$ by the radius r_v of the Delaunay disk centered at v , and store them, in this order, at the leaves of a balanced binary tree T . For each node w of T , let Q_w denote the subset of vertices stored at the leaves of the subtree rooted at w , and let R_w denote the set of the corresponding radii r_v . Construct the union U_w of all the Delaunay disks centered at the points of Q_w , and store it at w . U_w has $O(|Q_w|)$ edges and vertices [5] and can be constructed in $O(|Q_w| \log |Q_w|)$ time (e.g., by lifting the disks to planes in three dimensions and by constructing the upper envelope of these planes), for a total of $O(n \log n)$ storage and $O(n \log^2 n)$ preprocessing time. We process each U_w for fast (logarithmic time) point location.

Now, given a query point q , we search through T . At each node w that we visit, we take its right child w_r and check whether $q \in U_{w_r}$. If so, then q is contained in at least one Delaunay disk whose radius belongs to R_{w_r} , and we continue the search at w_r . If $q \notin U_{w_r}$, we continue the search at the left child w_l of w . When the search terminates at some leaf, the disk stored at that leaf is the largest Delaunay disk containing q . The search takes a total of $O(\log^2 n)$ time.

3.2 Efficient solution for disks centered at Voronoi edges

We next consider the other situation, where $c_{\max}(q)$ lies in the relative interior of some Voronoi edge. Let e_{ab} be a Voronoi edge, determined by two points $a, b \in P$ and delimited by the Voronoi vertices v_{abc}, v_{abd} , for two additional points $c, d \in P$, so that v_{abc} (resp., v_{abd}) is equally nearest to a, b, c (resp., to a, b, d). Let D_{ab} denote the diametral disk on ab , and let D_{abc} (resp., D_{abd}) denote the Delaunay disk centered at v_{abc} (resp., v_{abd}), so it circumscribes Δabc (resp., Δabd).

Lemma 3.1. *Any point q for which $c_{\max}(q)$ lies in the relative interior of e_{ab} must lie in*

$$K_{ab} = \text{int}(D_{ab} \setminus (D_{abc} \cap D_{abd})) .$$

Proof. Indeed, we have already argued that such a point q must satisfy $\angle aqb > \pi/2$, which is equivalent to q belonging to the interior of D_{ab} . In addition, the center of the disk circumscribing Δaqb lies in the relative interior of e_{ab} if and only if q lies in the interior of $(D_{abc} \cup D_{abd}) \setminus (D_{abc} \cap D_{abd})$, as is easily verified. The conjunction of these two conditions is $q \in K_{ab}$. \square

The case where ab is an edge of the *Gabriel graph* is shown in Figure 3. In this case ab crosses e_{ab} and D_{ab} is P -empty. We refer to such an edge in short as a *Gabriel edge*. In this case K_{ab} is the union of the interiors of the two shaded lunes.

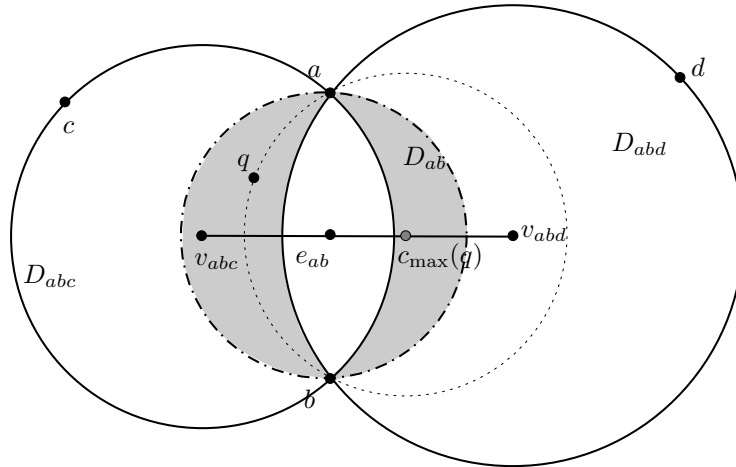


Figure 3: q has to lie in the interior of one of the shaded lunes if $c_{\max}(q)$ lies in the relative interior of the Voronoi edge e_{ab} ; the case where ab is an edge of the Gabriel graph. The corresponding disk $D_{\max}(q)$ is shown dotted.

The case where the Delaunay edge ab does not belong to the Gabriel graph (a *non-Gabriel edge*) is shown in Figure 4. Here K_{ab} is the single shaded lune.

Note that for each lune, in both cases of Gabriel edges and non-Gabriel edges, the area enclosed by the “outer” (longer) arc of the lune and the edge ab is P -empty.

We construct the regions K_{ab} , for all Voronoi edges e_{ab} of $\text{Vor}(P)$, and denote the collection of all these $O(n)$ lunes and double lunes as \mathcal{K} . A crucial property is that no point q can be contained in more than three regions of \mathcal{K} . Indeed, a point q belongs to K_{ab} if and only if the disk circumscribing Δaqb is P -empty, its center lie in the relative interior of the Voronoi edge e_{ab} , and $\angle aqb > \pi/2$. The claim then follows from Lemma 2.3.

Ideally, we could construct the arrangement $\mathcal{A}(\mathcal{K})$, preprocess it for fast point location, and store with each face f of $\mathcal{A}(\mathcal{K})$ the at most three Delaunay edges ab whose regions K_{ab} (fully) contain f . Then, given a query point q , we could locate q in the arrangement, retrieve the at most three corresponding Delaunay edges, and test, in constant time, each of the respective Voronoi edges for possible location of $c_{\max}(q)$. The problem with this approach is that the complexity of $\mathcal{A}(\mathcal{K})$ might be quadratic, as depicted in Figure 5. (Note that the arrangement \mathcal{A} constructed in Section 2.1 is in fact a refinement of $\mathcal{A}(\mathcal{K})$.)

We overcome this problem by first decomposing \mathcal{K} and its lunes into a constant number of subcollections, each consisting of portions of the lunes, so that, within each subproblem that we create, the regions are *pairwise openly disjoint*, and so the planar map that they form has linear complexity. Overall, we obtain a structure that uses only linear storage, and the query time, which is dominated by the point location in these maps, is only $O(\log n)$. (The




Figure 4: q has to lie in the interior of the shaded lune if $c_{\max}(q)$ lies in the relative interior of the Voronoi edge e_{ab} ; the case where ab is not an edge of the Gabriel graph.

overall query time, though, is $O(\log^2 n)$, because of the search for the largest Delaunay disk containing q ; see Section 3.1.)

The intersection pattern of the lunes. The following simple lemma provides the main technical tool for decomposing the lunes.

Lemma 3.2. *Let K_{ab} and K_{cd} be two distinct lunes (the four points a, b, c, d need not all be distinct) with a nonempty intersection. Then, for each point $q \in K_{ab} \cap K_{cd}$, there exists a line ℓ through q that (weakly) separates a, b from c, d . That is, one closed halfplane bounded by ℓ contains a and b , and the other closed halfplane contains c and d .*

Proof. Refer to Figure 6. Let D_{abq} (resp., D_{cdq}) denote the disk circumscribing the triangle




Figure 5: Illustrating a scenario in which the arrangement of the lunes K_{ab} has quadratic complexity.

Δabq (resp., Δcdq). By the properties of the lunes, as analyzed in Lemma 3.1, both disks are P -empty. The circle C_{abq} bounding D_{abq} and the circle C_{cdq} bounding D_{cdq} meet at q and therefore either (i) they also intersect at a second point q' , or (ii) they are tangent to each other at q . Let ℓ be the line connecting q and q' in case (i) or the common tangent line at q in case (ii). As is easily checked, ℓ separates the arcs $\alpha = C_{abq} \setminus D_{cdq}$ and $\gamma = C_{cdq} \setminus D_{abq}$. Moreover, since D_{abq} and D_{cdq} are both P -empty, a and b must lie in the closure of α and c and d must lie in the closure of γ , thereby completing the proof of the lemma. \square

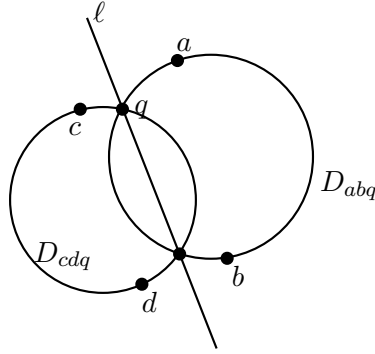


Figure 6: Illustrating the proof of Lemma 3.2.

3.3 Partitioning the lunes

Having these observations at hand, we now proceed to eliminate all intersections between the lunes by an appropriate problem decomposition, as promised above. Rotate the coordinate frame so that no two points have the same x - or y -coordinate. We first partition \mathcal{K} into two subcollections \mathcal{K}^+ and \mathcal{K}^- , where \mathcal{K}^+ (resp., \mathcal{K}^-) consists of all the lunes which lie above (resp., below) the lines containing their “bases” (i.e., the corresponding Delaunay edges). We describe the processing of \mathcal{K}^+ ; processing \mathcal{K}^- is done in a fully symmetric manner. We refer to the lunes in \mathcal{K}^+ as *upper lunes*.

We next distinguish between lunes in \mathcal{K}^+ whose bases have positive slope and those whose bases have negative slope, and treat each subcollection separately. We will only consider the first subcollection, since the other one is handled in a fully symmetric manner, and, for simplicity of notation, we will continue to denote it as \mathcal{K}^+ .

Let K_{ab} be one of these upper lunes, with, say, a lying to the left of b (and with ab having positive slope). We note that some portion of K_{ab} extends to the left of the vertical line through a and that K_{ab} lies fully to the left of the vertical line through b (this latter property holds because K_{ab} is always contained in the diametral disk of ab).

We form from \mathcal{K}^+ two new collections. We take each upper lune K_{ab} in \mathcal{K}^+ , with a to the left of (and below) b , and split it into two open regions by drawing a vertical line through a . (The union of the regions is K_{ab} without its intersection segment with the vertical line

through a .) One region, denoted K_{ab}^T , lies vertically above ab , and the other, denoted K_{ab}^L , lies fully to the left of a . We denote by \mathcal{K}^T (resp., \mathcal{K}^L) the collection of all the top portions K_{ab}^T (resp., left portions K_{ab}^L) of the lunes of \mathcal{K}^+ . See Figure 7 for an illustration.




Figure 7: Partitioning an upper lune (with a base ab of positive slope) into a top portion and a left portion.

Lemma 3.3. *The regions of \mathcal{K}^T are pairwise disjoint.*

Proof. Suppose to the contrary that there exist two distinct top regions K_{ab}^T and K_{cd}^T in \mathcal{K}^T , with a to the left of b and c to the left of d , such that K_{ab}^T and K_{cd}^T have a nonempty intersection, and let q be a point in both regions. By Lemma 3.2 there exists a line ℓ through q that weakly separates a and b from c and d . This however is impossible because q lies vertically above both segments ab and cd . See Figure 8(a). \square

Lemma 3.4. *The regions of \mathcal{K}^L are pairwise disjoint.*

Proof. Suppose to the contrary that there exist two distinct left regions K_{ab}^L and K_{cd}^L in \mathcal{K}^L , with a to the left of b and c to the left of d , such that K_{ab}^L and K_{cd}^L have a nonempty intersection, and let $q \in K_{ab}^L \cap K_{cd}^L$. By Lemma 3.2 there exists a line ℓ through q that weakly separates a and b from c and d . By construction, the two segments ab and cd lie to the right of the vertical line λ through q ; one of them, say, ab lies in the upper-right quadrant Q^+ formed by ℓ and λ , and the other, cd , lies in the lower-right quadrant Q^- . See Figure 8(b). The angle at q of at least one of these quadrants is at most $\pi/2$; without loss of generality assume that this is the case for Q^- . Since q lies in K_{cd}^L , which is fully contained in the interior of the diametral disk of cd , it follows that $\angle cqd < \pi/2$, contradicting the properties of the lunes as studied at the beginning of this section. This contradiction completes the proof of the lemma. \square




Figure 8: (a) Illustrating the proof of Lemma 3.3. (b) Illustrating the proof of Lemma 3.4.

3.4 The data structure

We thus obtain the following algorithm. Partition \mathcal{K} into the four subcollections of upper lunes with positive-slope bases, upper lunes with negative-slope bases, lower lunes with positive-slope bases, and lower lunes with negative-slope bases. Each of these subcollections is processed in a similar manner, and we describe only the processing of the first subcollection, which, as above, we denote as \mathcal{K}^+ .

We decompose each lune K in \mathcal{K}^+ into its top subregion K^T and its left subregion K^L , by the recipe given above, and let \mathcal{K}^T (resp., \mathcal{K}^L) denote the collection of the resulting top (resp., left) subregions. By Lemmas 3.3 and 3.4, each collection consists of pairwise disjoint regions, and so forms a planar map of linear complexity (so that each subregion in the collection is a single separate face of the map). We process each of the two resulting maps for fast (logarithmic time) point location queries. Altogether we obtain eight such maps.

Now, given a query point q , we locate it in each of these eight maps, and retrieve the at most three lunes that contain q . Each such lune yields a candidate maximal empty disk containing q and centered at the respective Voronoi edge. (If q happens to lie on a vertical edge of a map, which bounds some region K_{ab}^T or K_{ab}^L , then the Voronoi edge e_{ab} is one of the three candidate edges.) This part of the data structure requires $O(n)$ space, take $O(n \log n)$ time to construct and querying it takes $O(\log n)$ time. We then find the largest empty Delaunay disk (centered at a Voronoi vertex) containing q , as described in Section 3.1, and output the largest of the at the four candidate disks. The overall complexity is dominated by the costs of handling Delaunay disks.

We thus obtain the main result of this paper.

Theorem 3.5. *One can preprocess a set P of n points in the plane in $O(n \log^2 n)$ time into a data structure of size $O(n \log n)$, so that, given a query point q (in the interior of the convex hull of P), we can compute, in $O(\log^2 n)$ time, the largest P -empty disk containing q .*

See the appendix for a discussion and some open problems.

References

- [1] J. Augustine, S. Das, A. Maheshwari, S. C. Nandy, S. Roy, and S. Sarvattomananda, Recognizing the largest empty circle and axis-parallel rectangle in a desired location, in [arXiv.org:1004.0558](https://arxiv.org/abs/1004.0558), 2010.
- [2] D. P. Dobkin and D. G. Kirkpatrick, Fast detection of polyhedral intersection, *Theoret. Comput. Sci.* 27 (1983), 241–253.
- [3] H. Kaplan, S. Mozes, Y. Nussbaum and M. Sharir, Submatrix maximum queries in Monge matrices and Monge partial matrices, and their applications, *Proc. 23rd ACM-SIAM Annu. Sympos. Discrete Algorithms*, 2012, to appear.
- [4] H. Kaplan and M. Sharir, Finding the maximal empty rectangle containing a query point, in [arXiv.org:1106.3628](https://arxiv.org/abs/1106.3628), 2011.
- [5] K. Kedem, R. Livne, J. Pach and M. Sharir, On the union of Jordan regions and collision-free translational motion amidst polygonal obstacles, *Discrete Comput. Geom.* 1 (1986), 59–71.

Appendix

In view of the applications discussed in the introduction, a natural interesting generalization of the problem studied in this paper is to handle queries, each consisting of a collection Q of k points, where, say, k is some constant, and the goal is to find the largest disk containing Q (or, rather, its convex hull $CH(Q)$) whose interior is disjoint from P . (Note that when $|Q| > 1$ the problem does not always have a solution.)

To answer such a query with a set Q we observe that the largest P -empty disk D containing Q must contain on its boundary either three points of P , in which case it is a Delaunay disk, or two points of P and one point of Q . Indeed, as is easy to see, in all other cases there is a larger P -empty disk containing Q . Finding the largest Delaunay disk containing a set of two or more points is harder than the case of a single query point since we cannot use the binary search procedure in the arrangements of unions of disks described in Section 3.1.

We can determine whether there exists a P -empty disk containing Q in polylogarithmic time with an appropriate preprocessing of P , as follows. We lift the points of P and the points of Q to the paraboloid $z = x^2 + y^2$ in \mathbb{R}^3 . (For simplicity, we refer to the lifted sets of points also as P and Q , respectively.) In this representation a P -empty disk containing Q corresponds to a hyperplane separating the convex hull $CH(P)$ of P and the convex hull $CH(Q)$ of Q . Such a hyperplane exists if and only if $CH(Q)$ is disjoint from $CH(P)$. It is not difficult to show, using the technique of Dobkin and Kirkpatrick [2] that one can preprocess P in $O(n \log n)$ time, into a linear-size data structure, such that, given a query set Q , we can determine whether $CH(Q)$ and $CH(P)$ are disjoint, in $O(k \log n)$ time.

Briefly and informally, we note that since neither of $CH(Q)$, and $CH(P)$ fully contains the other hull, they intersect if and only if either an edge of $CH(Q)$ intersects a face of

$CH(P)$ or (an edge of) $CH(P)$ intersects a face of $CH(Q)$. Intersection of edges of $CH(Q)$ with $CH(P)$ can be detected in $O(k \log n)$ time using the technique of [2]. Once we have determined that no edge of $CH(Q)$ meets $CH(P)$, we test, for each face f of $CH(Q)$, whether its supporting plane π_f meets $CH(P)$. (This can also be easily done in a total of $O(k \log n)$ time.) If no such intersection is found, f is disjoint from $CH(P)$. Otherwise we obtain a witness point $w \in \pi_f \cap CH(P)$, and we test whether $w \in f$ (and then f intersects $CH(P)$) or not (and then there is no intersection).

We also note that finding the largest P -empty disk whose boundary contains two points of P and one of Q is rather easy to do. We simply query the structure presented at Sections 3.2–3.4 with each point of Q separately, collect the at most $3k$ candidate disks, filter out those that do not fully contain Q , and return the longest among the surviving ones. With an appropriate implementation this takes a total of $O(k \log n + k \log k)$ time.

We leave the problem of finding the largest P -empty Delaunay disk containing Q as an open problem for further research.

Another open problem is to reduce the query time to $O(\log n)$ without increasing significantly the storage and preprocessing costs. As already noted, the bottleneck in the query cost is at the stage that finds the largest Delaunay disk that contains q (see Section 3.1).