

Kinetic convex hulls, Delaunay triangulations and connectivity structures in the black-box model

Citation for published version (APA):

Berg, de, M. T., Roeloffzen, M. J. M., & Speckmann, B. (2012). Kinetic convex hulls, Delaunay triangulations and connectivity structures in the black-box model. *Journal of Computational Geometry*, 3(1), 222-249.

Document status and date:

Published: 01/01/2012

Document Version:

Publisher's PDF, also known as Version of Record (includes final page, issue and volume numbers)

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.tue.nl/taverne

Take down policy

If you believe that this document breaches copyright please contact us at:

openaccess@tue.nl

providing details and we will investigate your claim.

KINETIC CONVEX HULLS, DELAUNAY TRIANGULATIONS AND CONNECTIVITY STRUCTURES IN THE BLACK-BOX MODEL*

Mark de Berg,[†] Marcel Roeloffzen,[†] and Bettina Speckmann[†]

ABSTRACT. Over the past decade, the kinetic-data-structures framework has become the standard in computational geometry for dealing with moving objects. A fundamental assumption underlying the framework is that the motions of the objects are known in advance. This assumption severely limits the applicability of KDSs. We study KDSs in the *black-box model*, which is a hybrid of the KDS model and the traditional time-slicing approach. In this more practical model we receive the position of each object at regular time steps and we have an upper bound on d_{\max} , the maximum displacement of any point in one time step.

We study the maintenance of the convex hull and the Delaunay triangulation of a planar point set P in the black-box model, under the following assumption on d_{\max} : there is some constant k such that for any point $p \in P$ the disk of radius d_{\max} contains at most k points. We analyze our algorithms in terms of Δ_k , the so-called k -spread of P . We show how to update the convex hull at each time step in $O(\min(n, k\Delta_k \log n) \log n)$ amortized time. For the Delaunay triangulation our main contribution is an analysis of the standard edge-flipping approach; we show that the number of flips is $O(k^2 \Delta_k^2)$ at each time step.

1 Introduction

Motivation. Algorithms dealing with objects in motion traditionally discretize time and recompute the structure of interest at every time step from scratch. This can be wasteful, especially if the time steps are small: then the objects will have moved only slightly, and the structure may not have changed at all. Ideally an object gets attention if and only if its new location triggers an actual change in the structure. *Kinetic data structures (KDSs)*, introduced by Basch *et al.* [4], try to do exactly that: they maintain not only the structure itself, but also additional information that helps to find out when and where the structure will undergo a “real” (combinatorial) change. Instead of sampling the object locations at regular time intervals, KDSs follow an event-driven approach. They maintain a collection of simple geometric tests—these are called *certificates*—with the property that as long as these certificates remain valid, the structure of interest does not change combinatorially. A KDS computes for each certificate the nearest time in the future when it will fail and puts all these failure times into an event queue. Whenever there is an event—that is, a certificate failure—the KDS is updated. Note that the fact that we know which certificate has failed

*Marcel Roeloffzen was supported by the Netherlands’ Organisation for Scientific Research (NWO) under project no. 600.065.120. Bettina Speckmann was supported by the Netherlands’ Organisation for Scientific Research (NWO) under project no. 639.022.707.

[†]Technical University Eindhoven, {mberg, mroeloff, speckman}@win.tue.nl

when we handle an event gives us valuable information to update the attribute efficiently. See one of the surveys by Guibas [14, 15, 16] for more information and results on KDSs.

A basic assumption in the KDS framework is that the object trajectories are known. This is necessary to compute the failure times of the certificates, which is essential for the event-driven approach taken in the KDS framework. This assumption severely limits the applicability of the framework. When tracking moving objects, for instance, one gets the object locations only at (probably regular) time steps in an online manner—no detailed knowledge of future trajectories is available. The same is true for physical simulations, where successive locations are computed by a numerical integrator. Our goal is to study the kinetic maintenance of two fundamental geometric structures—convex hulls and Delaunay triangulations—in a less restrictive setting: instead of assuming knowledge of the trajectories, we assume only that we know upper bounds on the speeds of the objects and that we get their positions at regular time steps.

Related work. We are not the first to observe that the basic assumption in the KDS framework is not always valid. The need for a hybrid model, which combines ideas from KDSs with a traditional time-slicing approach, was already noted in the survey by Agarwal *et al.* [2]. Since then there have been several papers in this direction, as discussed next.

One way to cope with this problem is to verify all certificates at each time step and treat the certificate failures accordingly. Agarwal *et al.* [1] observe that treating certificate failures in the wrong order can cause the resulting structure to be incorrect. For kinetic tournaments and kinetic range trees they provide KDSs that can deal with out-of-order certificates and guarantee the structure is correct most of the time, and if it is not correct, it is close to being correct. However, their approach does not easily generalize to other problems and treating all certificate failures may still be slower than recomputing, since a failed certificate can produce new certificates that have also failed. An example of this is sorting, where the certificates indicate that two adjacent numbers are in the right order. If a certificate fails we simply swap the number, remove one certificate and add two new ones. Repairing the list using swaps takes $\Omega(n^2)$ time in the worst case, whereas sorting can be done in $O(n \log n)$ time.

Gao *et al.* [13] study spanners for sets of n moving points in a model where one does not know the trajectories in advance but receives only the positions at each time step. They call this the *blackbox replacement model*—we simply call it the *black-box model*—and show how to update the spanner at each time step in $O(n + k \log \alpha)$ time. Here α is the *spread* of the point set, and k is the number of changes to the hierarchical structure defining their spanner.

Mount *et al.* [21] also study the maintenance of geometric structures in a setting where the trajectories are unknown. They separate the concerns of tracking the points and updating the geometric structure into two modules: the motion processor (MP) is responsible for tracking the points, and the incremental motion algorithm (IM) is responsible for maintaining the geometric structure. The MP monitors the points to see whether they move “as expected”, and notifies the IM when this is no longer the case or some other important event happens. The IM then recomputes the structure, possibly querying the MP for the

location of certain points, and gives the MP new motion estimates. Mount *et al.* describe a protocol trying to minimize the interaction between the modules, and they prove that under certain conditions their protocol has good competitive ratio. Their approach goes back to the work of Kahan [18] on certain kinetic 1-dimensional problems. See also the more recent works by Cho *et al.* [9] and by Yi and Zhang [24].

The following papers show how to repair a triangulation after the vertices have moved. Shewchuk [23] considers d -dimensional Delaunay triangulations. He introduces star splaying, which estimates the neighborhood of each vertex and then resolves all inconsistencies between neighborhoods until the new Delaunay triangulation is found. The worst-case expected running time is $O(n^{\lceil d/2 \rceil + 1} + n^2 \log n)$, but the algorithm runs in linear time when the degree of each vertex is $O(1)$. Agarwal *et al.* [3] repair an (arbitrary) planar triangulation by finding “inverted” triangles and then finding regions that can be locally re-triangulated. After $O(n)$ time to find all inverted triangles, they use $O(k^2 \log k)$ time to find and re-triangulate the local regions, where k is the total complexity of these regions. (The analysis by Agarwal *et al.* is more refined and depends on additional parameters that indicate how entangled the triangulation is.)

Experimental work has also been done on kinetic Delaunay triangulations. De Castro *et al.* [6] describe how to easily determine a tolerance region for each point, such that as long as the point remains within its tolerance region we do not have to check its certificates. They then give experimental results showing that for fairly stable Delaunay triangulations this filtering method is faster than traditional KDSs. Russel argues in his thesis [22] that in practice a naive traditional KDS for a Delaunay triangulation is never faster than rebuilding and presents a filtering approach that is faster than rebuilding when the number of certificate failures is small.

The theoretical results discussed above typically express the running time in terms of the number of changes to the structure at hand, without further analyzing this number. This is not surprising, since without assumptions on the maximum displacements of the points one cannot say much about the number of changes. This “abstract” analysis is nice since it makes the results general, but on the other hand it becomes hard to decide whether it is better to use these kinetic algorithms or to simply recompute the structure from scratch at each time step. This is the goal of our paper: to develop KDSs in the black-box model that are—under certain assumptions on the trajectories—provably more efficient than recomputing the structure from scratch.

Our results. We study black-box KDSs for the convex hull and the Delaunay triangulation of a set P of n points moving in the plane. We also show how the Delaunay triangulation can be used to efficiently maintain the connectivity structure of a set of disks in the black-box model. As mentioned, we make assumptions on the point movements and time steps to obtain provably efficient solutions. In particular, the time steps should be small enough so that there is some coherence between the positions of the points in consecutive time steps—otherwise we cannot do better than recomputing the structure from scratch. Furthermore, we will assume in most of our results that P is fairly evenly distributed at each time step. Next we discuss our assumptions in detail, and state our results.

For a point $p \in P$, let $\text{NN}_k(p, P)$ denote the k -th nearest neighbor of p in $P \setminus \{p\}$.

Let $\text{dist}(p, q)$ denote the Euclidean distance between two points p and q , and define

$$\text{mindist}_k(P) := \min_{p \in P} \text{dist}(p, \text{NN}_k(p, P)).$$

Our basic assumption is that d_{\max} , the maximum displacement of any point during one time step, satisfies the *Displacement Assumption* that $d_{\max} \leq \text{mindist}_k(P)$, for some small k . Note that $\text{mindist}_k(P)$ may change as the points move. Thus a more precise statement of the Displacement Assumption is that d_{\max} is bounded by the minimum value of $\text{mindist}_k(P)$ over all time steps—see Section 2. We believe that in many practical applications, the sampling rate will be such that the Displacement Assumption is satisfied.

To describe the distribution of P we use the concept of *k-spread*, as introduced by Erickson [11] and defined as follows. Let $\text{diam}(P)$ denote the diameter of P . Then the *k-spread* of P , denoted by $\Delta_k(P)$, is defined as

$$\Delta_k(P) := \text{diam}(P) / \text{mindist}_k(P).$$

Note that the 1-spread of P is simply the standard spread. The 1-spread is not very suitable for moving points, however, as it blows up as soon as two points get very close to each other. The k -spread is more robust since it allows up to k objects to get very close to each other without causing a blow-up in the k -spread. Our analyses will be in terms of Δ_k , the maximum k -spread over all time steps, where k is such that the motions satisfy the Displacement Assumption.

For the convex-hull problem, we present an algorithm that updates the convex hull at each time step in $O(\min(n, k\Delta_k \log n) \log n)$ amortized time. We also present a variant of the algorithm whose running time does not depend on the k -spread: for any set of moving points satisfying the Displacement Assumption, it updates the convex hull in $O(n \log k)$ time. Lastly we show how to generalise our convex hull algorithm to higher dimensions.

For the Delaunay triangulation we consider two straightforward algorithms. The first one moves the points one at a time from their old to their new locations, meanwhile updating the Delaunay triangulation using edge flips. The second algorithm deletes each point from the triangulation and re-inserts it at its new location; the triangle into which the new location lies is found by walking in the triangulation. Our main contribution lies in the analysis of these simple approaches under the Displacement Assumption and in terms of Δ_k . For example, we show that the simple flipping algorithm performs only $O(k^2 \Delta_k^2)$ flips. We also show how to use the Delaunay triangulation to maintain the connected components of the intersection graph of a set of unit disks. Updating this connectivity structure takes $O(k^2 \Delta_k^2)$ time per time step.

2 Preliminaries

In this section we introduce some notation, and we discuss a few basic issues regarding the black-box model and the concept of k -spread. Although some of our results extend to higher dimensions, we will focus here on the case where P is a set of points moving in the plane.

The black-box model. We denote the position of a point p at time t by $p(t)$, and we let $P(t) := \{p(t) : p \in P\}$ denote the point set at time t . In the black-box model, we assume that we receive the positions at regular time steps t_0, t_1, \dots and the goal is to update the structure of interest—the convex hull or the Delaunay triangulation in our case—at each time step. The algorithm need not ask for all new positions at each time step; it may ignore some points if the new locations of these points cannot change the structure. Thus a sublinear update time is potentially feasible—indeed, we will show how to obtain sublinear update time for the convex-hull maintenance, under certain conditions. As stated in the introduction, we assume the sampling rate is such that the points in P do not move too much in one time step, as compared to their inter-distances. More precisely, we assume the sampling rate satisfies the following assumption.

Displacement Assumption: There is a maximum displacement d_{\max} such that

- $d_{\max} \leq \min_{t_i} \text{mindist}_k(P(t_i))$, and
- $\text{dist}(p(t_i), p(t_{i+1})) \leq d_{\max}$ for each point $p \in P$ and any time step t_i .

The k -spread of a point set. Recall that $\Delta_k(P)$, the k -spread of P , is defined as

$$\Delta_k(P) := \text{diam}(P) / \text{mindist}_k(P).$$

The k -spread of a point set can be used to bound the number of points within a region if the diameter of the region is not too large.

Lemma 1. *Let P be a set of points in \mathbb{R}^2 , and let R be a region in \mathbb{R}^2 such that $\text{diam}(R) < \text{mindist}_k(P)$. Then R contains at most k points from P .*

Proof. Assume for a contradiction that R contains $k+1$ points. Let p be an arbitrary point in $R \cap P$, then $\text{dist}(p, q) < \text{mindist}_k(P)$ for any point $q \in R \cap P$. This would imply that $\text{dist}(p, \text{NN}_k(p, P)) < \text{mindist}_k$, contradicting the definition of mindist_k . \square

Corollary 1. *Let B be a minimum bounding square of a point set P in \mathbb{R}^2 , and consider a partitioning of B into a regular grid with $\Delta_k(P) \times \Delta_k(P)$ cells. Then each grid cell contains $O(k)$ points.*

Corollary 1 implies that for any point set P in \mathbb{R}^2 we have $\Delta_k(P) = \Omega(\sqrt{n/k})$.

Remark. A statement similar to the converse of Corollary 1 also holds: if a partitioning of the minimum bounding square B into a regular grid with $\Delta \times \Delta$ cells has at most k points per cell, then $\Delta_{k'}(P) = \Delta$ for some $k' = O(k)$. The best case is when a $\sqrt{n} \times \sqrt{n}$ grid has at most k points per cell, so that $\Delta_k(P) = O(\sqrt{n})$. One may think that then the problems become easy, but this is in fact not the case; for example, we show in Theorem 1 that maintaining the points from P in x -order still takes $\Omega(n \log n)$ time in the black-box model, even for point sets with $\Delta_k(P) = O(\sqrt{n})$ and constant k .

Theorem 1. *Maintaining in x -order a set P of points in the black-box model, even when there is a constant k such that $\Delta_k(P) = O(\sqrt{n})$ requires $\Omega(n \log n)$ time in the worst case per time step.*

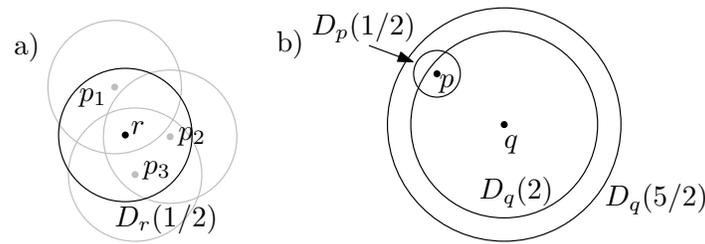


Figure 1: a) r can only be contained in up to k disks $D_p(1/2)$ for $p \in P$. b) $D_q(5/2)$ contains all disks $D_p(1/2)$ for which $p \in D_q(2)$.

Proof. Consider the case where all points are aligned on a $\sqrt{n} \times \sqrt{n}$ grid. (We can avoid degeneracy by placing each point a small distance $\epsilon \ll \text{mindist}_k(P)$ away from its grid coordinate.) For $k = 4$ we have $\Delta_k(P) = \sqrt{2}\sqrt{n} = O(\sqrt{n})$. The points within one column of the grid can have an arbitrary x -order in the next time step. Each column has $\sqrt{n}! \geq (\sqrt{n}/2)^{\sqrt{n}/2}$ different orderings. Since there are \sqrt{n} columns this gives $(\sqrt{n}/2)^{n/2}$ different orderings for P . In the decision tree model this gives a lower bound of $\Omega(\log(\sqrt{n}/2)^{n/2}) = \Omega(n \log n)$ for sorting P on x -coordinate. \square

Remark. For small k , the k -spread can be arbitrarily large. For $k = n - 1$, on the other hand, we have $\Delta_k(P) \leq 2$. Obviously, the k -spread decreases monotonically (though not necessarily strictly monotonically) as k increases. A natural question is whether anything more precise can be said about how $\Delta_k(P)$ changes as k varies. It is easy to see that we cannot say much about the change in k -spread when k is decreased: $\Delta_{k-1}(P)$ cannot be bounded in terms of $\Delta_k(P)$, since $\text{mindist}_{k-1}(P)$ can be arbitrarily much smaller than $\text{mindist}_k(P)$. When k is increased, on the other hand, then at some point the k -spread will go down.

Lemma 2. For any point set P it holds that $\Delta_{k'}(P) \leq \Delta_k(P)/2$ for $k' = 25k$, where $\Delta_k(P)$ denotes the k -spread of P .

Proof. Let $D_r(\alpha)$ denote the open disk centered at $r \in \mathbb{R}^2$ with radius $\alpha \cdot \text{mindist}_k(P)$ and let $P_r(\alpha) = P \cap D_r(\alpha)$. From Lemma 1 it follows that a disk $D_r(1/2)$ contains no more than k points of P . We claim that there is no point $r \in \mathbb{R}^2$ such that $r \in D_p(1/2)$ for more than k points $p \in P$. Assume that a point r exists that is in the disk $D_p(1/2)$ for $k + 1$ different points $p \in P$, then $D_r(1/2)$ contains $k + 1$ points, which contradicts the spread assumption—see Figure 1a. It follows that the intersection depth of the disks $D_p(1/2)$ for all points $p \in P$ is at most k .

Let $q \in P$, then the disks $D_p(1/2)$ for $p \in P_q(2)$ are contained in $D_q(5/2)$ as illustrated in Figure 1b. Because the stabbing depth of the disks $D_p(1/2)$ is at most k the sum of the areas of these disks is at most $k |D_q(5/2)|$, where $|D_q(5/2)|$ is the area of $D_q(5/2)$. This gives:

$$|P_q(2)| \leq \frac{k |D_q(5/2)|}{|D_p(1/2)|} = 25k$$

It follows that $\text{mindist}_{k'}(P) \geq 2\text{mindist}_k(P)$ and $\Delta_{k'}(P) \leq \Delta_k(P)/2$ for $k' = 25$. \square

3 Maintaining the convex hull

Let $\mathcal{CH}(P)$ denote the convex hull of a point set P , and let $\partial\mathcal{CH}(P)$ denote the boundary of $\mathcal{CH}(P)$. In this section we give algorithms to maintain $\mathcal{CH}(P(t))$. From now on, we will use $\mathcal{CH}(t)$ as a shorthand for $\mathcal{CH}(P(t))$. We introduce three algorithms to update the convex hull. The first two work best when the spread Δ_k is small; the second algorithm is slightly faster than the first one, but it needs the floor function as a constant-time operation and it assumes $\text{mindist}_k(P(t))$ is known. The third algorithm works well when $\Delta_k(P)$ is large and (like the first one) only uses the Displacement Assumption.

3.1 The basic algorithm

Our algorithm relies on the following observation, which follows from the fact that the distance between p and $\partial\mathcal{CH}(P)$ can change by only $2d_{\max}$ in a single time step.

Lemma 3. *Consider a point $p \in P$, and let $d_p(t) := \text{dist}(p(t), \partial\mathcal{CH}(t))$. Then p cannot become a vertex of $\mathcal{CH}(P)$ until at least $\frac{d_p(t)}{2d_{\max}}$ time steps have passed.*

Lemma 3 suggests the following simple scheme to maintain $\mathcal{CH}(P)$. Compute the initial convex hull $\mathcal{CH}(t_0)$, and compute for each point $p \in P$ its distance to $\partial\mathcal{CH}(t_0)$. Using this distance and Lemma 3, compute a *time stamp* $t(p)$ for p , which is the number of time steps until p can be a vertex of the convex hull. Thus p can be ignored until the time stamp expires after $t(p)$ time steps. In a generic time step t_i , we now determine the set $Q(t_i)$ of all points whose time stamps expire, compute their convex hull—here we may use $\mathcal{CH}(t_{i-1})$ —and compute new time stamps for the points in $Q(t_i)$.

To implement this algorithm we use an array A , where $A[t_i]$ stores a pointer to a list that contains the points whose time stamps expire at time t_i . We actually work with an array $A[0..n-1]$ with n entries, and let time advance through the array in a cyclic manner (using without loss of generality that $t_i = i$). We bound the time stamps to be at most n steps, and we use an approximation d_p^* of $\text{dist}(p(t), \partial\mathcal{CH}(t))$ to speed up the computations. Our approach is made explicit in Algorithm 1. Note that the algorithm needs to know only d_{\max} to work correctly, it does not need to know bounds on the k -spread.

Algorithm 1: UPDATECH

- 1 $Q(t) \leftarrow$ set of points stored in $A[t]$
 - 2 Compute $\mathcal{CH}(Q(t))$ and set $\mathcal{CH}(t) \leftarrow \mathcal{CH}(Q(t))$.
 - 3 **foreach** $p \in Q(t)$ **do**
 - 4 Compute a lower bound d_p^* on $\text{dist}(p(t), \partial\mathcal{CH}(t))$.
 - 5 $t(p) \leftarrow \min(1 + \lfloor \frac{d_p^*}{2d_{\max}} \rfloor, n)$
 - 6 Add p to $A[(t + t(p)) \bmod n]$.
 - 7 $t \leftarrow (t + 1) \bmod n$
-

It remains to describe how to compute $\mathcal{CH}(Q(t))$ in Step 2 and how to compute the value d_p^* in Step 4. Computing $\mathcal{CH}(Q(t))$ can be done by an optimal convex-hull algorithm

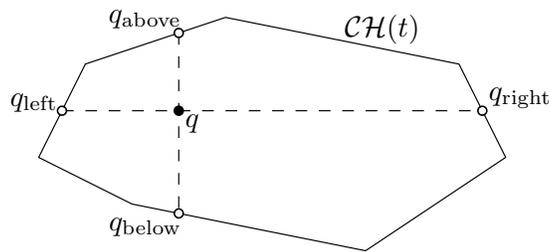


Figure 2: Points straight above, below, left and right of q are used to approximate the minimum distance from q to $\partial\mathcal{CH}(t)$.

in $O(|Q(t)| \log |Q(t)|)$ time [5]. To compute d_p^* we proceed as follows. Let q_{above} be the point on $\partial\mathcal{CH}(t)$ directly above p , and define q_{below} , q_{left} , and q_{right} similarly—see Figure 2. These points can be found in logarithmic time using binary search. Let $d_{\min}(p)$ denote the minimum distance between p and any of the points q_{above} , q_{below} , q_{left} , and q_{right} . Then we set $d_p^* = d_{\min}(p)/\sqrt{2}$. Note that $d_p^* \leq \text{dist}(p, \partial\mathcal{CH}(t)) \leq \sqrt{2} d_p^*$ and that this inequality together with Lemma 3 implies the correctness of Algorithm 1. We get the following result.

Lemma 4. *At each time step t , UPDATECH updates the convex hull in $O(|Q(t)| \log n)$ time.*

Next, we analyze the number of time stamps expiring in each time step.

Analysis of the number of expiring time stamps. We perform our analysis in terms of Δ_k , which is an upper bound on the k -spread of P at any time. We first prove a bound on the number of convex-hull vertices.

Lemma 5. *The convex hull $\mathcal{CH}(P)$ of a point set P has $O(\min(n, k\Delta_k))$ vertices.*

Proof. The length of $\partial\mathcal{CH}(P)$ is $\Theta(\text{diam}(P))$, so we can cut $\partial\mathcal{CH}(P)$ into

$$\Theta(\text{diam}(P)/\text{mindist}_k(P)) = O(\Delta_k) \text{ pieces}$$

with a length less than $\text{mindist}_k(P)$. From Lemma 1 we know that each such piece contains at most k points. We also know that the number of vertices on the convex hull can never exceed n . It follows that $\partial\mathcal{CH}(P)$ contains $O(\min(n, k\Delta_k))$ vertices. \square

In the worst case all time stamps expire in a single time step. However, using an amortization argument we show that on average only $O(k\Delta_k \log n)$ points expire in each time step.

Lemma 6. *In each time step an amortized number of $O(\min(n, k\Delta_k \log n))$ time stamps expire.*

Proof. We prove the lemma using the accounting method: each point has an account into which we put a certain amount of money at each time step, and whenever the time stamp of a point expires it has to pay 1 euro from its account. To prove the lemma we need to devise a scheme such that (i) a point always has at least 1 euro in its account when its

time stamp expires, and (ii) the total amount of money handed out at each time step is $O(k\Delta_k \log n)$. Our scheme is that at time step t_i each point p receives

$$\min \left(1, \max \left(\frac{1}{n}, \frac{8\sqrt{2} \cdot d_{\max}}{d_p(t_i)} \right) \right) \text{ euros,}$$

where $d_p(t_i) = \text{dist}(p(t_i), \partial\mathcal{CH}(t_i))$.

To prove (i), consider a point p whose time stamp expires at time t_i , and let $t_j < t_i$ be the previous time step when p 's time stamp expired. If there is no such time step, we can take $j = 0$. Now define $t(p) := t_i - t_j = i - j$ to be the number of time steps from t_j up to t_{i-1} . If $t(p) = n$ then p certainly has enough money in its account at time t_i , so assume this is not the case. Then

$$t(p) = 1 + \left\lfloor \frac{d_p^*(t_j)}{2d_{\max}} \right\rfloor \geq 1 + \left\lfloor \frac{d_p(t_j)}{2\sqrt{2} \cdot d_{\max}} \right\rfloor \geq \frac{d_p(t_j)}{2\sqrt{2} \cdot d_{\max}}.$$

The amount of money received by p from t_j up to t_{i-1} is thus at least

$$t(p) \cdot \frac{8\sqrt{2} \cdot d_{\max}}{\max_{t_j \leq t \leq t_{i-1}} d_p(t)} \geq \frac{4d_p(t_j)}{\max_{t_j \leq t \leq t_{i-1}} d_p(t)}.$$

The distance between p and $\mathcal{CH}(P)$ increases (or decreases) by at most $2d_{\max}$ at each time step, so at any time $t_j \leq t \leq t_{i-1}$ the distance from p to $\partial\mathcal{CH}(P)$ is at most

$$\begin{aligned} d_p(t_j) + t(p) \cdot 2d_{\max} &\leq d_p(t_j) + \left(1 + \frac{d_p(t_j)}{2d_{\max}} \right) \cdot 2d_{\max} \\ &= 2 \cdot d_p(t_j) + 2d_{\max} \leq 4 \cdot d_p(t_j), \end{aligned}$$

where in the last step we assumed that $d_p(t_j) \geq d_{\max}$ (since otherwise p already receives at least 1 euro at time t_j). Hence the total amount of money received by p is at least

$$\frac{4d_p(t_j)}{\max_{t_j \leq t \leq t_{i-1}} d_p(t)} \geq \frac{4d_p(t_j)}{4 \cdot d_p(t_j)} = 1.$$

To prove (ii), we consider the points p such that $d_p(t_i) \leq 8\sqrt{2}n \cdot d_{\max}$; the remaining points get $1/n$ euros each, so in total at most 1 euro. We divide these points into groups G_1, \dots, G_ℓ . Each group G_j contains the points $p \in P$ such that $(j - 1) \cdot d_{\max} \leq d_p(t_i) \leq j \cdot d_{\max}$, where $\ell = 8\sqrt{2}n$. All points from G_j lie in an annulus of diameter $O(\Delta_k \cdot \text{mindist}_k(P(t_i)))$ where the distance between the inner and outer boundary of the annulus is $\Theta(d_{\max})$. Using a simple packing argument and Lemma 1 it follows that such an annulus contains $O(k\Delta_k)$ points. Hence, the total amount we have to pay to all points in a single group G_j is

$$O(k\Delta_k) \cdot \min \left(1, \frac{8\sqrt{2} \cdot d_{\max}}{(j - 1) \cdot d_{\max}} \right) = O \left(\frac{k\Delta_k}{j} \right) \text{ euros}^1.$$

¹Note that when $j = 1$ we get a division by zero, but we assume that $\min(1, \frac{8\sqrt{2} \cdot d_{\max}}{(1-1) \cdot d_{\max}}) = 1$.

Summing this over all groups we see that the amount we pay at each time step is

$$\sum_{j=1}^{8\sqrt{2}n} O\left(\frac{k\Delta_k}{j}\right) = O(k\Delta_k \log n).$$

Every point has only one time stamp, so the number of time stamps is also bounded by n . This results in an $O(\min(n, k\Delta_k \log n))$ bound on the amortized number of time stamps expiring in a time step. \square

Lemma 4 and 6 imply the following theorem.

Theorem 2. *Under the Displacement Assumption, the convex hull of a set P of n points moving in the plane can be maintained in the black-box model in $O(\min(n, k\Delta_k \log n) \log n)$ time amortized per time step, where Δ_k is the maximum k -spread of P at any time.*

3.2 A faster algorithm using the floor function

For small k -spread—more precisely when $\Delta_k < n$ —we can optimize the previous algorithm even further. This optimization requires the algorithm to know $\text{mindist}_k(P(t))$ at each timestep t and assumes a constant-time floor function is available. Our new algorithm still follows the approach described in Algorithm 1; it differs from the implementation described in the previous section in how $\mathcal{CH}(t)$ is computed in Step 2 and how the value d_p^* is computed in Step 4.

Computing $\mathcal{CH}(t)$. First we define a grid with cells of size $\text{mindist}_k(P(t)) \times \text{mindist}_k(P(t))$, which contains all points of $P(t)$. This grid has $O(\Delta_k(P(t))^2)$ cells due to the spread assumption. Now, given the set $Q(t)$ of expired points we need to determine for each point which grid cell contains it. We find the grid cell containing a point in constant time by using the floor function on its coordinates. Although there may be more than n cells in the grid, at most $Q(t)$ of them can actually contain points. We can use hashing techniques to determine, in $O(|Q(t)|)$ time in total, for each non-empty cell which points from $Q(t)$ are contained in it.

The points have only limited movement, so the vertices of $\mathcal{CH}(t)$ will be close to the boundary of $\mathcal{CH}(t-1)$. More precisely, every vertex of $\mathcal{CH}(t)$ has a point on the boundary of $\mathcal{CH}(t-1)$ within distance d_{\max} of it. This implies that only those points of $P(t)$ that lie in a grid cell intersected by $\partial\mathcal{CH}(t-1)$ or in a cell adjacent to such a cell can become a vertex of $\mathcal{CH}(t)$. There are $O(\Delta_k)$ such cells and each of them contains $O(k)$ points. To compute the convex hull efficiently we can divide these cells into four groups: top, right, bottom, and left. Let p_{nw} be the leftmost point of the upper boundary of $\mathcal{CH}(t-1)$ such that a line l_1 with a slope of 1 is tangent to $\mathcal{CH}(t-1)$ in p_{nw} and let p_{ne} be the rightmost point of the upper boundary of $\mathcal{CH}(t-1)$ such that a line l_2 with a slope of -1 is tangent to $\mathcal{CH}(t-1)$ in p_{ne} (see Figure 3). Then the top group of cells are those that are intersected by the part of the upper boundary of $\mathcal{CH}(t-1)$ between p_{nw} and p_{ne} , and cells adjacent to these. Due to the slope of the edges between p_{nw} and p_{ne} only two cells in each column

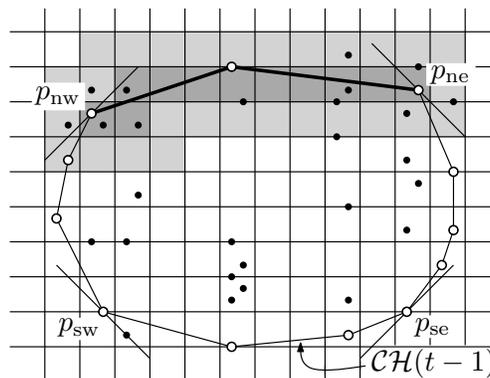


Figure 3: The top group of grid cells is defined as the cells intersected by the part of the boundary of $\mathcal{CH}(t-1)$ between p_{nw} and p_{ne} and those adjacent to these cells.

can be intersected by this top part of $\mathcal{CH}(t-1)$. This means that the top group contains only a constant number of cells in each column. We can then sort the points of the top group by x -coordinate by first sorting the points in each column. Each column of the top group has $O(k)$ points, hence sorting such a column takes $O(k \log k)$. Sorting all points in the top group then takes $O(\min(n, k\Delta_k) \log k)$. After sorting the points we compute the convex hull of points in cells of the top group in $O(\min(n, k\Delta_k))$ time using *Graham's scan* algorithm [5].

For the other groups we use a symmetric approach. We can then merge the convex hulls of the four groups to get $\mathcal{CH}(t)$.

Lemma 7. *The time needed to compute $\mathcal{CH}(t)$, after we have identified the points of $Q(t)$ and determined which cells they are in, is $O(\min(n, k\Delta_k) \log k)$.*

Proof. Finding the points which define where the different groups start and end can simply be done by walking across the boundary of $\mathcal{CH}(t-1)$ and hence takes only $O(\min(n, k\Delta_k))$ time. Identifying the right cells to consider for each part can be done in $O(\Delta_k)$ time because there are only $O(\Delta_k)$ cells to consider and for each group all cells are connected, so we use adjacencies to find them all. Sorting each group either by x or y coordinate can then be done in $O(\min(n, k\Delta_k) \log k)$ time as described above. After sorting, we can compute the convex hull of each group and merge the four convex hulls into one in $O(\min(n, k\Delta_k))$ time. \square

Computing expiration times. The next step is to compute expiration times. Recall that for this we need $d_{\min}(q)$ —the smallest horizontal or vertical distance to $\mathcal{CH}(t)$, as defined in Section 3.1—as an estimate for the distance from a point $q \in P(t)$ to $\partial\mathcal{CH}(t)$. However we do not want to spend $O(\log n)$ time to determine $d_{\min}(q)$ for each point $q \in Q(t)$. Hence, we proceed as follows. We focus on the computation of the vertical distance of the points in $Q(t)$ to the top boundary of $\mathcal{CH}(t)$; computing the vertical distance to the bottom boundary, and computing the horizontal distances to the left and right boundaries, can be done in a similar manner. In other words, we focus on computing for each point $q \in Q(t)$ the point q_{above} (see Figure 2).

Consider the points of $Q(t)$ that lie horizontally between the points p_{nw} and p_{ne} . These points can be grouped into subsets of points lying in the same column of grid cells. The width of such a column is $\text{mindist}_k(P(t))$. Now consider the top part of $\partial\mathcal{CH}(t)$, that is, the part of the boundary from p_{nw} to p_{ne} . Because the slope of the edges in the top part lies between -1 and $+1$, the length of the top part inside a single column is at most $\sqrt{2} \cdot \text{mindist}_k(P(t))$. From this and Lemma 1 it follows that the top part has $O(k)$ vertices inside each column. Thus we proceed as follows. First, we walk along $\partial\mathcal{CH}(t)$ to determine for each column in between p_{nw} and p_{ne} the vertices of the top part of $\partial\mathcal{CH}(t)$ lying inside that column. This can be done in $O(\min(n, k\Delta_k))$ time. Next, for each point of $Q(t)$ in between p_{nw} and p_{ne} we determine in $O(1)$ time the column it lies in (using the floor function) and we perform a binary search on the top part of $\partial\mathcal{CH}(t)$ inside the column to determine q_{above} . Thus, the total time needed to compute q_{above} for the points of $Q(t)$ in between p_{nw} and p_{ne} is $O(\min(n, k\Delta_k) + Q(t) \log k)$.

If q is to the left of p_{nw} or to the right of p_{ne} we cannot determine q_{above} so easily. The reason is that the part of $\partial\mathcal{CH}(t)$ inside the corresponding column can have many vertices, so the binary search will not work in $O(\log k)$ time. The next lemma states that this is not a problem, because q_{above} is only needed to determine $d_{\min}(q)$ for points in between p_{nw} and p_{ne} .

Lemma 8. *If a point $q \in Q(t)$ is to the left of p_{nw} or to the right of p_{ne} then q_{above} does not define $d_{\min}(q)$.*

Proof. Since both cases are symmetric, we assume that q is to the left of p_{nw} . Because of the definition of p_{nw} the edge of $\mathcal{CH}(t)$ that is straight above q will have a slope of at least 1. It follows that $\text{dist}(q, q_{\text{left}}) < \text{dist}(q, q_{\text{above}})$. Hence, q_{above} does not define $d_{\min}(q)$. \square

Thus all the distances needed to compute $d_{\min}(q)$ for all $q \in Q(t)$ can be computed in linear time in total.

Lemma 9. *After having computed $\mathcal{CH}(t)$ we can compute the new expiration times for the points in $Q(t)$ in $O(\min(n, k\Delta_k) + Q(t) \log k)$ time.*

Combining Lemma's 6, 7 and 9 we get the following theorem. Note that the difference in running time, compared to Theorem 2, is that one factor $O(\log n)$ is replaced by $O(\log k)$.

Theorem 3. *Under the Displacement Assumption and assuming we know $\text{mindist}_k(P(t))$ at any time step t , the convex hull of a set P of n points moving in the plane can be maintained in the black-box model in $O(\min(n, k\Delta_k \log n) \log k)$ amortized time, where Δ_k is the maximum k -spread of P at any time.*

Remark. As an example, consider the best possible case, namely that $\Delta_k = O(\sqrt{n})$ for $k = O(1)$. In this case the amortized update time is $O(\sqrt{n} \log n)$. Note that in this case we do not need hashing techniques to determine the non-empty cells, since there are only $O(n)$ cells in the grid.

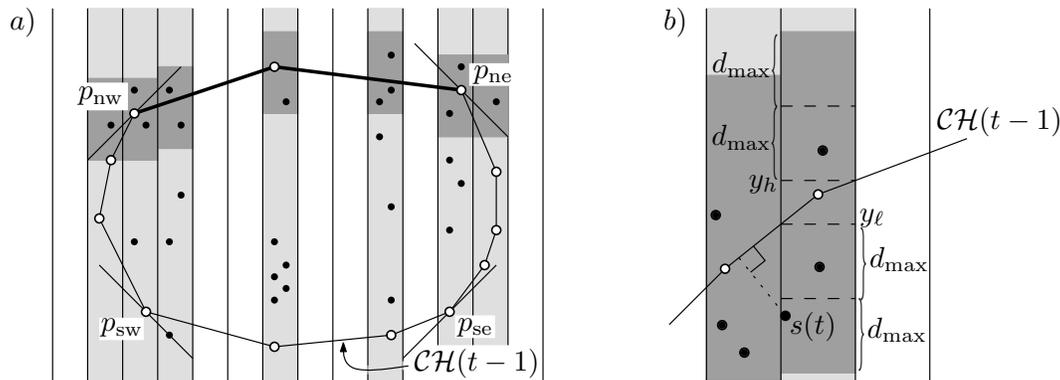


Figure 4: a) Potential vertices of the northern section of $\mathcal{CH}(t)$ must be in the dark gray regions. b) Potential vertices are between $y_l(C) - 2d_{max}$ and $y_h(C) + 2d_{max}$.

3.3 A near-linear-time algorithm without spread assumption

Next we show how the convex hull of a point set with a high k -spread can still be maintained in near-linear time per step, under the Displacement Assumption. Our new algorithm does not use the approach with time stamps; since we allow linear time this is not needed. Our goal is to ensure the algorithm uses $O(n \log k)$ time, rather than the $O(n \log n)$ time that would be needed if we were to reconstruct the convex hull from scratch.

Consider a partitioning of the plane into vertical columns of width d_{max} . We maintain a left-to-right ordered list \mathcal{L}_{col} of the columns that contain at least one point from P . For each column C we also maintain the set $P(C)$ of points inside that column. By the Displacement Assumption each point can either stay in its column or move to a neighboring column in a single time step. Hence we can update \mathcal{L}_{col} and the sets $P(C)$ in $O(1)$ time per point at each time step. We also maintain a bottom-to-top sorted list \mathcal{L}_{row} of the rows with height d_{max} that contain at least one point from P .

After we have updated \mathcal{L}_{col} and the sets $P(C)$ at time t each set $P(C)$ contains those points that are in C at time t . Now suppose we want to compute $\mathcal{CH}(t)$ from $\mathcal{CH}(t-1)$. We divide $\mathcal{CH}(t-1)$ into four sections using points $p_{nw}, p_{ne}, p_{sw}, p_{se}$. The points p_{nw} and p_{ne} are vertices on the upper boundary of $\mathcal{CH}(t-1)$ that have tangent lines with a slope of 1 and -1 respectively. The points p_{sw} and p_{se} are on the lower boundary of $\mathcal{CH}(t-1)$ and they have tangent lines with a slope of -1 and 1 respectively (see Figure 4a). We focus on the northern section $\mathcal{CH}_n(t-1)$ of the convex hull boundary between p_{nw} and p_{ne} .

We want to find all points of $P(t)$ that are within distance d_{max} of the northern section of the convex hull. We inspect the columns of \mathcal{L}_{col} from left to right. For each column C we see if it is intersected by $\mathcal{CH}_n(t-1)$. If this is the case we find the lowest and highest y -coordinate of this intersection, denoted by $y_l(C)$ and $y_h(C)$ respectively. Now for each point $p(t) \in P(C)$ we test if its y -coordinate $y(p(t))$ is greater than $y_l(C) - 2d_{max}$. We define $P^*(C) \subseteq P(C)$ as the set of points which satisfy this criterion. The slope of edges of $\mathcal{CH}_n(t-1)$ is between 1 and -1 , which guarantees that all points of $P(C)$ that are within distance d_{max} of $\mathcal{CH}_n(t-1)$ are in $P^*(C)$. Note that using $y_l(C) - d_{max}$ as bound on the

y -coordinate is not sufficient, since a point $s(t) \in C$ may have a shorter distance to the part of $\mathcal{CH}_n(t-1)$ in a neighboring column (see Figure 4b). Additionally we inspect the columns C_{nw} and C_{ne} which are to the left of the column containing p_{nw} and right of the column containing p_{ne} as shown in Figure 4a. For C_{nw} we find the set $P^*(C_{nw})$ of points which have a y -coordinate greater than $y(p_{nw}) - 2d_{\max}$.

Every point $p(t) \in P(t)$ for which $\text{dist}(p(t), \mathcal{CH}_n(t-1)) \leq d_{\max}$ is in the set $P^*(C)$ for some column C . The points in $P^*(C)$ all have a y -coordinate of at least $y_\ell(C) - 2d_{\max}$. Also no point in $P(C)$ can have a y -coordinate of more than $y_h + 2d_{\max} \leq y_\ell(C) + 3d_{\max}$ since a point $p(t)$ can be no more than d_{\max} away from $\mathcal{CH}(t-1)$. For every column C the points of $P^*(C)$ are contained in a $d_{\max} \times 5d_{\max}$ rectangle and thus $P^*(C)$ contains $O(k)$ points. It follows that we can sort them in $O(k \log k)$ time per column and sort the points in $\bigcup_{C \in \mathcal{L}_{\text{col}}} P^*(C)$ in $O(n \log k)$ time. It then takes $O(n)$ time to compute the convex hull of $\bigcup_{C \in \mathcal{L}_{\text{col}}} P^*(C)$.

In a similar fashion we compute the convex hulls for the points that are within distance d_{\max} of the eastern, southern or western part of $\partial\mathcal{CH}(t-1)$. We then merge the four convex hulls in $O(n)$ time to obtain $\mathcal{CH}(t)$.

Theorem 4. *Under the Displacement Assumption, the convex hull of a set P of n points moving in the plane can be maintained in the black-box model in $O(n \log k)$ time per time step.*

3.4 Maintaining the convex hull in \mathbb{R}^3

The algorithm we described in Section 3.1 generalizes to higher dimensions. We follow the steps of Algorithm UPDATECH to update the 3-dimensional convex hull. The array A again holds pointers to lists of points that expire at a certain time step. To compute new time stamps for each point $p(t) \in Q(t)$ we again find a lower bound $d_p^*(t)$ on the distance to the boundary of the convex hull. We set $d_p^*(t) = d_{\min}(p)/\sqrt{3}$, where $d_{\min}(p)$ is the minimum distance in $+x$ -, $-x$ -, $+y$ -, $-y$ -, $+z$ -, or $-z$ -direction from $p(t)$ to $\partial\mathcal{CH}(t)$. We find $d_{\min}(p)$ by shooting axis-aligned rays in all six possible directions. The real distance $d_p(t)$ between $p(t)$ and $\partial\mathcal{CH}(t-1)$ is then bounded as follows:

$$d_p^*(t) \leq d_p(t) \leq d_p^*(t)\sqrt{3}.$$

We show that the convex hull in \mathbb{R}^3 has complexity $O(\min(n, k\Delta_k^2))$ and that we can update the convex hull in sublinear time (for small spread).

Lemma 10. *The convex hull of a point set P in \mathbb{R}^3 has complexity $O(\min(n, k\Delta_k^2))$.*

Proof. Consider overlaying P with a grid with cells of size $\text{mindist}_k/2 \times \text{mindist}_k/2 \times \text{mindist}_k/2$. By definition of Δ_k the entire point set P can be covered by $2\Delta_k \times 2\Delta_k \times 2\Delta_k$ such cells. The convex hull of P is a convex polytope within this grid and, hence, its boundary can intersect only $O(\Delta_k^2)$ cells. Each cell has a diameter less than mindist_k . Following Lemma 1, which is easily generalized to higher dimensions, each cell contains at most k points of P . It follows that there are at most $O(\min(n, k\Delta_k^2))$ points of P on the convex hull boundary. \square

Lemma 11. *In \mathbb{R}^3 , algorithm UPDATECH takes time*

$$O\left(\left(k\Delta_k^2 + \Delta_k\sqrt{k|Q(t)|} + |Q(t)|\right)\log^{O(1)}n\right).$$

Proof. Computing the convex hull in three dimensions takes $O(m \log f)$ time using Chan's output-sensitive algorithm [7], where m is the number of points and f is the complexity of the convex hull. The complexity of the convex hull is $O(\min(n, k\Delta_k^2))$ under the spread assumption as shown in Lemma 10. Hence, computing the convex hull can be done in $O(|Q(t)| \log n)$ time. The ray-shooting queries we need in order to compute $d_p^*(t)$ can be done using Chan's algorithm for ray shooting [7], which takes

$$O(f \log q + \sqrt{fq} \log^{O(1)} f + q \log f) \text{ time,}$$

where f is again the complexity of the convex hull and q is the number of queries performed. We perform $q = 6|Q(t)|$ queries on a convex hull with $f = O(\min(n, k\Delta_k^2))$ complexity. Because $|Q(t)| \leq n$ and $f = O(\min(n, k\Delta_k^2))$, all logarithmic factors are upper bounded by $O(\log n)$. Hence, the final bound is as claimed. \square

For the number of expired points in a single time step we use a similar amortization scheme as in the proof of Lemma 6 where each point $p(t_i)$ gets

$$\min\left(1, \max\left(\frac{1}{n}, \frac{8\sqrt{3} d_{\max}}{d_p^*(t_i)}\right)\right) \text{ euros.}$$

Using a similar analysis as in Lemma 6 gives us that $|Q(t)| = O(\min(n, k\Delta_k^2 \log n))$ amortized. The main difference in the analysis is that in 2D the groups G_1, \dots, G_ℓ were formed by annuli that could be covered by $O(\Delta_k)$ disks with diameter mindist_k , whereas in 3D the groups are formed by convex shells that can be covered by $O(\Delta_k^2)$ balls with diameter mindist_k . Combining this with Lemma 11 we get the following result.

Theorem 5. *Under the Displacement Assumption, the convex hull of a set P of n points moving in \mathbb{R}^3 can be maintained in the black-box model in $O(\min(n, k\Delta_k^2 \log n) \log^{O(1)} n)$ time amortized per step, where Δ_k is the maximum k -spread of P at any time.*

When the spread is optimal— $\Delta_k = O(n^{1/3})$ and $k = O(1)$ —this is slightly faster than rebuilding from scratch; it even runs in sublinear time. The method also extends to higher dimensions which gives a running time of

$$O((\min(n, k\Delta_k^{d-1}))^{\lfloor d/2 \rfloor} \log^{O(1)} n),$$

under the assumption that d is a constant and Δ_k^{d-1} is polynomial in n . The output sensitive convex hull algorithm by Chan runs in

$$O(n \log f + (nf)^{1-1/(\lfloor d/2 \rfloor + 1)} \log^{O(1)} n) \text{ time.}$$

Our algorithm is asymptotically faster than Chans algorithm depending on the value of k and Δ_k , but it is never slower. The reason for this is that our bound is the same as Chans except that in our case n is replaced by $|Q(t)|$ which is at most n .

4 Maintaining the Delaunay triangulation

We denote the Delaunay triangulation of a point set P by $\mathcal{DT}(P)$, and we use $\mathcal{DT}(t)$ as a shorthand for $\mathcal{DT}(P(t))$. We describe two algorithms to recompute $\mathcal{DT}(t)$ given $\mathcal{DT}(t-1)$. Both algorithms use the same global approach: they update the position of each point in turn and change the Delaunay triangulation accordingly. Let $P = \{p_1, \dots, p_n\}$ and define $Q_i = \{p_1(t), \dots, p_i(t), p_{i+1}(t-1), \dots, p_n(t-1)\}$. Thus $Q_0 = P(t-1)$ and $Q_n = P(t)$. Updating the position of p_i now means computing $\mathcal{DT}(Q_i)$ from $\mathcal{DT}(Q_{i-1})$. This can be done in two ways.

- **MOVEANDFLIP:** The point $p_i(t-1)$ is moved along a straight line to its new position $p_i(t)$ while maintaining the Delaunay triangulation. For each point we follow the traditional KDS approach. We compute all in-circle certificates that involve p_i and sort them by their failure times—the failure time of a certificate is the position along the line $p_i(t-1)p_i(t)$ where the certificate becomes false. Each time a certificate fails we flip an edge of the Delaunay triangulation and update the collection of certificates.
- **INSERTANDDELETE:** We first walk in $\mathcal{DT}(Q_{i-1})$ from $p_i(t-1)$ to $p_i(t)$ to determine the triangle τ of $\mathcal{DT}(Q_{i-1})$ containing $p_i(t)$. Then we insert $p_i(t)$ into $\mathcal{DT}(Q_{i-1})$ in the usual way, namely by adding edges from $p_i(t)$ to the vertices of τ and then performing edge flips until we have the Delaunay triangulation of $Q_{i-1} \cup \{p_i(t)\}$ [5]. Finally, we delete $p_i(t-1)$ using the algorithm by Devillers [10] or we delete $p_i(t-1)$ with all its edges and then use the algorithm by Wang and Chin [8] to repair the Delaunay triangulation.

Let $D(p, r)$ denote the disk of radius r centered at point p . Note that $D(p_i(t-1), d_{\max})$ is exactly the region reachable from $p_i(t-1)$ in one time step. Now consider two points $p_j, p_\ell \in P$. We say that $p_j p_\ell$ is a *potential edge* at time t if there is a placement of each point $p_i \in P$ somewhere in its disk $D(p_i(t-1), d_{\max})$ such that $p_j p_\ell$ is an edge in the Delaunay triangulation of the resulting point set. We then call p_j and p_ℓ *potential neighbors*. Theorem 8 below states the number of potential edges is $O(k^2 \Delta_k^2)$. With this result in hand, we can analyze our update algorithms.

Theorem 6. *Computing $\mathcal{DT}(t)$ from $\mathcal{DT}(t-1)$ using MOVEANDFLIP in the black-box model under the Displacement Assumption requires $O(k^2 \Delta_k^2)$ flips in total and takes $O(k^2 \Delta_k^2 \log n)$ time, where Δ_k is the k -spread at time $t-1$.*

Proof. Consider the movement of p_i from $p_i(t-1)$ to $p_i(t)$. A flip occurs when p_i becomes co-circular with three other points, say a, b, c , and $\text{circ}(a, b, c)$, the circle defined by a, b, c , is empty. Next we observe that each point a, b, c must form a potential edge with p_i : at the time of the flip p_i is on $\text{circ}(a, b, c)$, which contains no other points, hence ap_i, bp_i and cp_i are edges of the Delaunay triangulation at that time. Let $E(p_i)$ denote the set of potential edges with p_i as an endpoint, and let $N(p_i) = \{q \in Q_i \mid qp_i \in E(p_i)\}$ be the set of potential neighbors of p_i . The number of empty circles defined by $N(p_i)$ is equal to the number of Delaunay triangles in $\mathcal{DT}(N(p_i))$, which is $O(|N(p_i)|) = O(|E(p_i)|)$. Using Theorem 8 we

conclude that the total number of flips for moving all points $p_i \in P$ is

$$\sum_{p_i \in P} O(|E(p_i)|) = O(k^2 \Delta_k^2).$$

At each flip we update the Delaunay triangulation in $O(1)$ time. We must also update the event queue storing the certificate failure times, which takes $O(\log n)$ time since we have to delete and insert only $O(1)$ certificates into the queue. The running time thus becomes $O(k^2 \Delta_k^2 \log n)$. \square

The $O(\log n)$ factor in the running time for the MOVEANDFLIP approach stems from the event queue on the certificates. The INSERTANDDELETE approach avoids this factor.

Theorem 7. *Under the Displacement Assumption computing $\mathcal{DT}(t)$ from $\mathcal{DT}(t-1)$ using INSERTANDDELETE takes $O(k^2 \Delta_k^2)$ time in the black-box model, where Δ_k is the k -spread at time $t-1$.*

Proof. Consider the update of the position of p_i . Note that whenever we cross an edge $p_j p_\ell$ during the walk with p_i , then $p_j p_\ell$ would be part of a flip in the MOVEANDFLIP strategy. This implies that the number of edges crossed is at most linear in the number of potential edges. Moreover, inserting $p_i(t)$ takes time linear in the degree of $p_i(t)$ in $\mathcal{DT}(Q_{i-1} \cup \{p_i(t)\})$ [5] and deletion of $p_i(t-1)$ takes linear time in the degree of $p_i(t-1)$ [8]. Overall, updating the position of p_i takes linear time in the number of potential edges involving p_i , so the total time is $O(k^2 \Delta_k^2)$ by Theorem 8. The algorithm by Wang and Chin [8] is somewhat complicated and might be slow when deleting points with a small degree. In that case it may be more efficient to use the asymptotically slower, but simpler algorithm by Devillers [10]. \square

Analysis of the number of potential edges. Potential edges are defined on the point set $P(t-1)$ at a single time step, so we drop the time parameter and use $P = \{p_1, \dots, p_n\}$ to denote the set of points we are dealing with. We also define $\text{mindist}_k := \text{mindist}_k(P)$ and $\Delta_k := \Delta_k(P)$. (The latter is a slight abuse of notation as in fact Δ_k was defined as an upper bound on the k -spread at any time.) We first give a necessary condition for two points p_j and p_ℓ to form a potential edge. We call a disk *empty* if its interior does not contain any point from P . The following lemma follows from the empty-disk property of Delaunay triangulations and the fact that points move by at most d_{\max} in a single time step.

Lemma 12. *If $p_j p_\ell$ is a potential edge then there is an empty disk D^- such that p_j and p_ℓ are within distance $2d_{\max}$ of D^- .*

Proof. Assume that $p_j p_\ell$ is a potential edge. Then there is a set of points $P' = \{p'_1, \dots, p'_n\}$, where $p'_i \in D(p_i, d_{\max})$ for all i , such that $p'_j p'_\ell$ is an edge in $\mathcal{DT}(P')$. Let $D = D(c, r)$ denote an empty disk with p'_j and p'_ℓ on its boundary. For now assume $r > d_{\max}$, and let $D^- = D(c, r - d_{\max})$; see Figure 5. For any point p'_i , we know that $|cp'_i| \geq r$ and

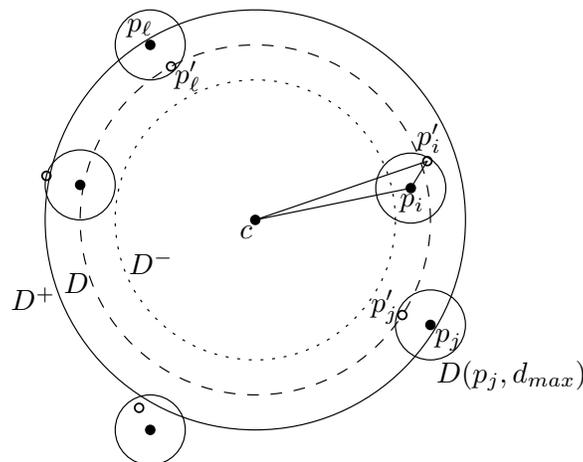


Figure 5: $p_j p_l$ is a potential edge only if a disk D^- exists that has p_j and p_l within distance $2d_{\max}$ and does not contain any other points of P .

$|p_i p'_i| \leq d_{\max}$. It follows from the triangle inequality that $|cp_i| \geq r - d_{\max}$ and that p_i cannot be inside D^- .

For p'_j it holds that $|cp'_j| = r$ and $|p_j p'_j| \leq d_{\max}$. By the triangle inequality we get that $|cp_j| \leq r + d_{\max}$. The same holds for p_l , which proves that $p_j p_l$ can be a potential edge only if there is a disk D^- that does not contain any other points of P but has p_j and p_l within distance $2d_{\max}$ of its boundary. It remains to consider the case $r \leq d_{\max}$. Then $|p_j p_l| \leq 4d_{\max}$. The overlap region of $D(p_j, 2d_{\max})$ and $D(p_l, 2d_{\max})$ is non-empty. If the overlap has positive area then we can place a (possibly very small) disk D inside $D(p_j, 2d_{\max}) \cap D(p_l, 2d_{\max})$ that does not contain any points from P and, hence, satisfies the conditions of the lemma. If the overlap is a single point q —note that this point could happen to be a point in P —then we can place a small empty disk touching q and satisfying the conditions of the lemma. □

Let E_{pot} denote the set of potential edges defined by P . For each potential edge $p_j p_l$ we pick a disk as in Lemma 12, which we call its *witness disk*. We split the set of witness disks into three subsets based on the size of the disks:

- \mathcal{D}_S : the *small* witness disks, which have radius at most $16 \cdot \text{mindist}_k$,
- \mathcal{D}_M : the *medium-size* witness disks, which have radius between $16 \cdot \text{mindist}_k$ and $\text{diam}(P)$,
- \mathcal{D}_L : the *large* witness disks, which have radius larger than $\text{diam}(P)$.

The number of potential edges contributed by disks in \mathcal{D}_S is easy to bound: if a potential edge pq has a small witness disk, then q must lie within $O(\text{mindist}_k)$ distance of p , and so by Lemma 1 there are only $O(k)$ such points for any point p .

Lemma 13. *The witness disks in \mathcal{D}_S contribute $O(kn)$ potential edges.*

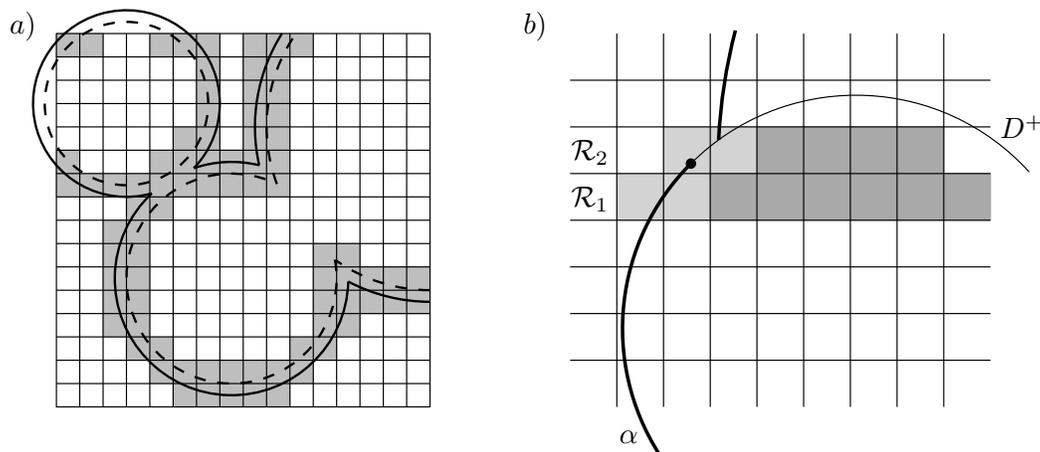


Figure 6: a) Points in $U^+ \setminus U$ must be in cells intersected by ∂U^+ or ∂U . b) To the right of a cluster of up to three cells intersected left arcs of ∂U^+ is a cluster of $O(\text{size}(\sigma)\text{mindist}_k)$ cells in the interior of U^+ .

To prove bounds for \mathcal{D}_L and \mathcal{D}_M we need the following lemma. For a square σ , let $\text{size}(\sigma)$ denote its edge length and let $\text{union}(\mathcal{D})$ denote the union of a set \mathcal{D} of disks.

Lemma 14. *Let σ be a square with edge length $\text{size}(\sigma)$ and let $P_\sigma = P \cap \sigma$. Let \mathcal{D} be a set of disks with radius at least $\text{size}(\sigma)/4$ that do not contain any points of P_σ . Then the number of points in P_σ within distance $2d_{\max}$ of $\text{union}(\mathcal{D})$ is $O(k \cdot \text{size}(\sigma)/\text{mindist}_k)$.*

Proof. Define $\mathcal{D}^+ := \{D(c, r + 2d_{\max}) : D(c, r) \in \mathcal{D}\}$. Set $U := \text{union}(\mathcal{D})$ and $U^+ := \text{union}(\mathcal{D}^+)$. Since the disks in \mathcal{D} are empty, all the points of P_σ that are within distance $2d_{\max}$ of $\text{union}(\mathcal{D})$ lie in $U^+ \setminus U$. We overlay the square σ by a grid whose cells have size $4 \cdot \text{mindist}_k$. Because $d_{\max} \leq \text{mindist}_k$, each grid cell intersecting $U^+ \setminus U$ must intersect ∂U^+ or ∂U (or both); see Figure 6a. Since the grid cells have size $4 \cdot \text{mindist}_k$ it follows from Lemma 1 that they each contain $O(k)$ points. It remains to prove that the number of grid cells intersected by ∂U^+ or ∂U is $O(\text{size}(\sigma)/\text{mindist}_k)$.

We show how to count the cells intersecting ∂U^+ ; the cells intersecting ∂U can be counted similarly. We split the boundary of each disk $D^+ \in \mathcal{D}^+$ into a *left arc*, *right arc*, *top arc*, and *bottom arc* at the points where the tangent lines have slope $+1$ and -1 . The boundary ∂U^+ consists of parts of these arcs. We count the cells intersecting ∂U^+ separately for each type of arc. We argue that in every row only $O(1)$ cells can be intersected by left arcs. Let \mathcal{R}_1 be a row that is intersected by a left arc α from disk D^+ so that arc α crosses the row entirely. The arc α intersects at most two cells (see Figure 6b). To the right of these intersected cells are $\Theta(\text{size}(\sigma)/\text{mindist}_k)$ cells (if they exist) that are contained inside disk D^+ and cannot be intersected by any other left arcs. Let \mathcal{R}_2 be a row that contains an endpoint of α . Again only two cells in \mathcal{R}_2 can be intersected by α but there may be more cells to the right that are not fully contained in D^+ as illustrated in Figure 6b. However since each disk has at least radius 16mindist_k and each cell has a size of 4mindist_k there can only be two such cells. To the right of this cluster of up to three cells that may be

intersected by left arcs there is a cluster of $\Theta(\text{size}(\sigma)/\text{mindist}_k)$ cells that is contained in D^+ and cannot be intersected by left arcs. It follows that in each row only $O(1)$ cells are intersected by left arcs. The total number of grid cells intersecting a left arc on ∂U^+ is therefore proportional to the number of rows, which is $O(\text{size}(\sigma)/\text{mindist}_k)$. For the right, top, and bottom arcs we can use a similar argument. \square

We can now prove that the large witness disks contribute $O(k^2\Delta_k^2)$ potential edges.

Lemma 15. *The witness disks in \mathcal{D}_L contribute $O(k^2\Delta_k^2)$ potential edges.*

Proof. Obviously, P is contained inside a $\text{diam}(P) \times \text{diam}(P)$ square. Because the disks in \mathcal{D}_L have radius at least $\text{diam}(P) = \Delta_k \cdot \text{mindist}_k$ we can apply Lemma 14 to conclude there are only $O(k\Delta_k)$ points within distance $2d_{\max}$ of $\text{union}(\mathcal{D}_L)$. These points define $O(k^2\Delta_k^2)$ potential edges. \square

It remains to bound the number of potential edges contributed by the medium-size disks. For $2 \leq i \leq \log_4 \Delta_k$, define

$$\mathcal{D}_M^i := \{D \in \mathcal{D}_M : 4^i \cdot \text{mindist}_k \leq \text{radius}(D) < 4^{i+1} \cdot \text{mindist}_k\}.$$

We bound the number of potential edges contributed by a subset \mathcal{D}_M^i in terms of the area of $\text{union}(\mathcal{D}_M^i)$.

Lemma 16. *The witness disks in \mathcal{D}_M^i contribute $O(k^2 \frac{A_i}{\text{mindist}_k^2})$ potential edges, where A_i is the area of $\text{union}(\mathcal{D}_M^i)$.*

Proof. We overlay the plane with a grid whose cells have size $4^{i+1} \cdot \text{mindist}_k$. Any two points defining a potential edge with a witness disk in \mathcal{D}_M^i have distance at most

$$4^{i+1} \cdot \text{mindist}_k + 4d_{\max} \leq 4^{i+2} \cdot \text{mindist}_k$$

from each other, and so the points in any grid cell can form potential edges with points in only $O(1)$ other cells. Lemma 14 implies that in any cell only $O(k4^{i+1})$ points are within distance $2d_{\max}$ of $\text{union}(\mathcal{D}_M^i)$. Hence in total the points in any cell C can contribute only $O((k4^{i+1})^2) = O(k^24^{2i})$ potential edges with witness disks from \mathcal{D}_M^i . Now we just have to count the number of cells within distance $2d_{\max}$ of $\text{union}(\mathcal{D}_M^i)$. Let

$$(\mathcal{D}_M^i)^+ := \{D(c, r + 2d_{\max}) : D(c, r) \in \mathcal{D}_M^i\}$$

and set $U_i^+ := \text{union}((\mathcal{D}_M^i)^+)$. Note that the area of U_i^+ is $O(A_i)$. We need to count the number of cells intersecting U_i^+ . Each cell intersecting U_i^+ either contains at least $1/4$ of a disk with radius at least 4^imindist_k or it is adjacent to such a cell. Hence, the total number of intersected cells is bounded by

$$O(\text{area}(U_i^+)/\text{area of one cell}) = O(A_i/(4^{i+1} \cdot \text{mindist}_k)^2).$$

We already showed that the points in any given cell contribute $O(k^24^{2i})$ potential edges in total, so the total number of potential edges is $O(k^2 A_i/\text{mindist}_k^2)$. \square

Recall that $i \leq \log_4 \Delta_k$. Since $\text{diam}(P) = \Delta_k \cdot \text{mindist}_k$, we know that the diameter of union(\mathcal{D}_M^i) is at most

$$\Delta_k \cdot \text{mindist}_k + 4^{i+1} \cdot \text{mindist}_k = O(\Delta_k \cdot \text{mindist}_k).$$

Hence, $A_i = O(\Delta_k^2 \cdot \text{mindist}_k^2)$, and so Lemma 16 implies that \mathcal{D}_M^i contributes $O(k^2 \Delta_k^2)$ potential edges. Since the number of subsets \mathcal{D}_M^i is $O(\log \Delta_k)$, It follows that the total number of potential edges contributed by medium-size disks is $O(k^2 \Delta_k^2 \log \Delta_k)$.

We can get rid of the $O(\log \Delta_k)$ factor by not considering each subset \mathcal{D}_M^i in isolation, but also considering their interaction. From Lemma 16 we know that the set \mathcal{D}_M^i can only contribute many edges if $A_i = \text{area}(U_i)$ is large, where $U_i = \text{union}(\mathcal{D}_M^i)$. In the next lemma we show that A_i cannot be large for all i .

Lemma 17. *There is a region $V_i \subseteq U_i$ such that, for all $j \leq i - 2$, $\text{union}(U_j)$ is disjoint from V_i and this region has area at least $A_i/\gamma = \Theta(A_i)$, for some fixed constant $\gamma > 1$.*

Proof. Consider a disk $D \in \mathcal{D}_M^i$ with center c and radius r . Set

$$r^- := r - 4^{i-1} \cdot \text{mindist}_k - 2d_{\max}$$

and define $D^- = D(c, r^-)$. Then D^- must be disjoint from any disk in \mathcal{D}_M^j for $j \leq i - 2$. Indeed, any disk of \mathcal{D}_M^j has radius at most $4^{i-1} \cdot \text{mindist}_k$, so if it were to intersect D^- it would be completely contained in $D(c, r - 2d_{\max})$. This means that all points are at least distance $2d_{\max}$ away from it, and so it cannot be a witness disk of a potential edge.

Note that the radius r^- of any inner disk D^- is

$$r^- = r - 4^{i-1} \cdot \text{mindist}_k - 2d_{\max} \geq r - (4^{i-1} + 2) \cdot \text{mindist}_k.$$

We overlay U_i with a grid whose cells have size $(4^{i-2}) \cdot \text{mindist}_k$. Then any inner disk D^- contains at least one grid cell and D itself intersects a constant number of cells (see Figure 7). Hence, we can charge any cell intersecting a disk D of \mathcal{D}_M^i to a cell that is completely inside some inner disk D^- in such a way that we charge only a constant number of cells to each cell in an inner disk. It follows that

$$A_i = O(\text{area of the union of the inner disks}),$$

which proves the claim. □

With this result we prove that the total number of potential edges contributed by \mathcal{D}_M is $O(k^2 \Delta_k^2)$.

Lemma 18. *The number of potential edges contributed by the witness disks in \mathcal{D}_M is $O(k^2 \Delta_k^2)$.*

Proof. We prove the bound for the subsets for which i is even; the proof for the subsets with i odd is similar. Consider the subsets \mathcal{D}_M^i , for even i , in order of decreasing i . Let A_i^* be the area that is still available for the disks in \mathcal{D}_M^i after considering all disks from \mathcal{D}_M^j

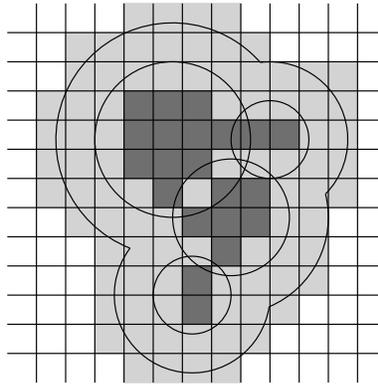


Figure 7: A disk intersects only a constant number of grid cells (light gray) and has at least one grid cell completely contained in its inner area (dark gray).

where $j \geq i + 2$. Thus $A_i \leq A_i^*$. Lemma 17 implies that $A_{i-2}^* \leq A_i^* - (1/\gamma)A_i$ for some constant $\gamma > 1$. In other words, $A_i \leq \gamma(A_i^* - A_{i-2}^*)$. If we define $j_{\max} = \log_4 \Delta_k/2$ then using Lemma 16 we can bound the contribution to the number of potential edges for \mathcal{D}_M^i , with i even, as follows.

$$\begin{aligned} \sum_{j=1}^{j_{\max}} O\left(k^2 \frac{A_{2j}}{\text{mindist}_k^2}\right) &= O\left(\frac{k^2}{\text{mindist}_k^2} \sum_{j=1}^{j_{\max}} \gamma(A_{2j}^* - A_{2j-2}^*)\right) \\ &= O\left(\frac{k^2}{\text{mindist}_k^2} \gamma(A_{2j_{\max}}^* - A_0^*)\right) \\ &= O(k^2 \Delta_k^2) \end{aligned}$$

The last step follows from the fact that

$$A_{2j_{\max}}^* = O(\text{diam}(P)^2) = O(\Delta_k^2 \cdot \text{mindist}_k^2).$$

□

From Lemma's 13, 15 and 18 we conclude the following.

Theorem 8. *Let P be a set of n points. Under the Displacement Assumption, the number of potential edges defined by P is $O(\min(k^2 \Delta_k^2, n^2))$, where Δ_k is k -spread of P .*

We can show this bound is tight in the worst case. One could hope that only a smaller number of edges may actually occur during movement, but even this is not the case: Next we show that $\Omega(k^2 \Delta_k^2)$ edges can occur when moving the points one at a time.

Theorem 9. *For large enough n , $k \geq 1$ and $\Delta_k \geq \sqrt{8n}$, there is a set P of n points with a k -spread of Δ_k at time $t-1$ such that computing $\mathcal{DT}(t)$ from $\mathcal{DT}(t-1)$ under the Displacement Assumption takes $\Omega(k^2 \Delta_k^2)$ time using the MOVEANDFLIP or INSERTANDDELETE approach if the points are moved in a bad order.*

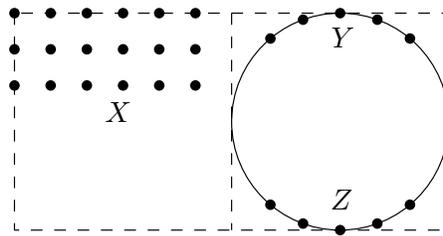


Figure 8: We need $\Omega(k^2\Delta_k^2)$ time to update the Delaunay triangulation when points of Y move in a bad order.

Proof. Consider a set $P = X \cup Y \cup Z$ as specified below. The points in Y and Z will generate the desired $\Omega(k^2\Delta_k^2)$ lowerbound, whereas X contains leftover points which are placed on a grid and do not move to ensure the spread assumption is not violated. Without loss of generality we assume that $\text{mindist}_k = 1$. The points of X are placed in a square of height and width $\Delta_k/(2\sqrt{2})$ and the points of $Y(t-1)$ and $Z(t-1)$ are placed in a similar square (see Figure 8).

We place the points of Y and Z at time $t-1$ along the boundary of a disk $D = D(d, r)$ as follows. We divide D into four sections using lines with a slope of 1 and -1 through the center d of D . Points of $Y(t-1)$ are along the boundary of the top section D_t and points of $Z(t-1)$ along the boundary of the bottom section D_b (see Figure 9). The points of Y will move into D ; whenever a point $y \in Y$ has moved it will have an edge with every point of Z in the Delaunay triangulation. Let $y_1 \dots y_m$ be the points of Y in the order in which we move them and $z_1 \dots z_m$ the points of Z , which do not move. Let P_i be the point set P between time $t-1$ and t just after y_i has moved:

$$P_i := \{y_1(t) \dots y_i(t), y_{i+1}(t-1) \dots y_m(t-1)\} \cup Z \cup X$$

Note that the time parameter for Z and X is omitted as points in these sets do not move.

Each point $y_i(t-1)$ is moved towards d by a distance $2\frac{2^i}{2^{m+1}} = \frac{2^{i+1}}{2^{m+1}} = 2^{-(m-i)}$. Let $C_{i,j}$ be the disk inside D and tangent to D in z_j with radius $r - 2^{-(m-i+1)}$ and let $c_{i,j}$ denote the center of $C_{i,j}$. For every point z_j the point $y_i(t)$ will be contained in $C_{i,j}$.

For $y_i(t)$ and z_j to form an edge in the Delaunay triangulation there has to be a disk that contains $y_i(t)$ and z_j , but no other points of P_i . No point of Z , other than z_j can be inside $C_{i,j}$ as all points of Z are on the boundary of D and $C_{i,j}$ only intersects D in z_j . For each point in $P_i \setminus \{y_i\}$ it holds that the distance to d is at least $r - 2^{-(m-i+1)}$. Now we claim (and will prove later) that every point inside $C_{i,j}$ that is in the top section of D has at most distance $\sqrt{r^2 - r2^{-(m-i)}}$ to d . Since

$$\sqrt{r^2 - r2^{-(m-i)}} \leq r - 2^{-(m-i+1)}$$

it follows that no point of $P_i \setminus \{y_i\}$ is contained in $C_{i,j}$. Therefore there is an edge between $y_i(t)$ and z_j in $\mathcal{DT}(P_i)$. In this manner each point $y_i(t)$ has edges to all points of Z in $\mathcal{DT}(P_i)$. Since both Y and Z contain $\Omega(k\Delta_k)$ points, the total number of edges created is $\Omega(k^2\Delta_k^2)$.

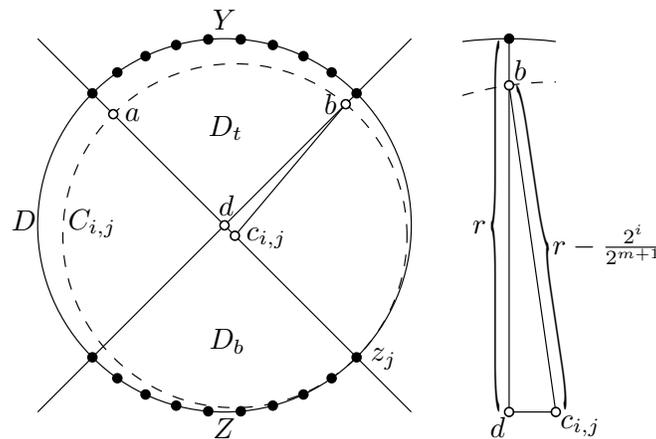


Figure 9: The point of Y and Z are in the top and bottom wedges respectively.

What remains to be proven is that the distance between any point in $C_{i,j} \cap D_t$ and the center d of D is at most $\sqrt{r^2 - r2^{-(m-i)}}$. The shortest distance from the boundary of $C_{i,j}$ to d is achieved by the point a on the line through z_j and d and by construction this point is in D_t . If we go clockwise or counterclockwise along the boundary of $C_{i,j}$ the distance to d only becomes larger until we reach z_j . The worst case situation arises for the point b on one of the bounding edges of D_t (see Figure 9). Here the longest distance from the boundary of $C_{i,j}$ to d is

$$\sqrt{(r - 2^{-(m-i+1)})^2 - (2^{-(m-i+1)})^2} = \sqrt{r^2 - r2^{-(m-i)}}.$$

□

5 Kinetic connectivity of unit disks

The intersection graph of a set \mathcal{R} of regions in the plane is the graph whose vertices are the regions in \mathcal{R} and there is an edge between two regions when the regions intersect each other. Intersection graphs of (unit) disks have been used to model wireless networks; here a disk models the region that is within the transmission radius of some node in the network and the fact that two disks intersect means that the corresponding nodes can communicate directly. In this application, the connected component of a vertex in the intersection graph corresponds to those nodes in the network with which the node can (directly or indirectly) communicate. For mobile networks, it thus becomes interesting to maintain the connected components as the nodes in the network move. This problem has been studied in the standard KDS setting [17]. We study it in the black-box model.

Let $\mathcal{R}(t)$ be the set of disks at time t . We assume that the disks adhere to Displacement Assumption, that is, any disk moves at most d_{\max} per time step. We define Δ_k as the maximum k -spread of the centers of the disks over all time steps.

Unit disks. The simplest case occurs when all the disks in \mathcal{R} have the same radius, which we assume without loss of generality to be 1. In this case we can proceed as follows: we maintain the Delaunay triangulation of the centers, remove (at each time step) all edges that are more than twice as long as the radius of the disks, and compute the connected components of the resulting graph. This is correct because of the following lemma.

Lemma 19 ([12, 20]). *Let $\mathcal{DT}_{\text{short}}$ denote the graph on the disk centers that is equal to the Delaunay graph, but with all edges removed that are more than twice as long as the radius of the disks. Then the connected components of $\mathcal{DT}_{\text{short}}$ are the same as the connected components of the intersection graph of the disks.*

The Delaunay triangulation has linear size, and the connected components can be computed in linear time. Maintaining the Delaunay can be done in $O(k^2\Delta_k^2)$ time, as described in Section 4. Since $k^2\Delta_k^2 = \Omega(n)$ we obtain the following result.

Theorem 10. *Under the Displacement Assumption we can maintain the connected components of the intersection graph of a set of unit disks in the black-box model in $O(k^2\Delta_k^2)$ time, where Δ_k is the maximum k -spread in any time step.*

Arbitrary size disks. When the disks are allowed to have an arbitrarily large radius we cannot use the Delaunay triangulation of the centers. However, as explained by Karavelas [19] we can use the Delaunay graph of the disks instead. More precisely, we can use the dual of the Voronoi diagram where the distance of a point p to a disk D with center q and radius r is $|pq| - r$. Note that this distance is negative when p is inside D .

We define a subgraph G of the Delaunay graph of disks by only using edges from the Delaunay graph between intersecting disks. This graph G has linear size and the same connected components as the intersection graph of the disks [19, Section 6]. However, without any bounds on the radii of the disks we cannot update the Delaunay triangulation for disks efficiently in the same manner as we do for points. Doing so can take $\Omega(n^2)$ time even when the k -spread Δ_k of the centers is $O(\sqrt{n})$.

We can prove this using a construction similar to that of Theorem 9. Let D be a disk with radius \sqrt{n} and define two wedges W_{up} and W_{down} as illustrated in Figure 10. We divide the input disks into two groups \mathcal{R}_{up} and $\mathcal{R}_{\text{down}}$, where the disks in \mathcal{R}_{up} are tangent to D in W_{up} and the disks in $\mathcal{R}_{\text{down}}$ are tangent to D in W_{down} . Each set contains $n/2$ disks. The disks in \mathcal{R}_{up} are further divided into \sqrt{n} rows where the i -th row contains $\sqrt{n}/2$ disks with their centers located on the circular arc defined by the intersection of a circle centered at d —the center of D —and W_{up} . Each disk in the i -th row has a radius of i , so that it is tangent to D . The disks in $\mathcal{R}_{\text{down}}$ are defined similarly.

We can then move each disk of \mathcal{R}_{up} inside D so that it forms Delaunay edges with all disks of $\mathcal{R}_{\text{down}}$. This results in $\Omega(n^2)$ flips and $\Omega(n^2)$ time to update the Delaunay graph. The proof for this is very similar to that of Theorem 9 and therefore omitted here.

Theorem 11. *Maintaining the Delaunay graph (as defined above) of a set of disks with arbitrary radii takes $\Omega(n^2)$ time when using the MOVEANDFLIP or INSERTANDDELETE approach when the k -spread is at least $4\sqrt{n}$.*

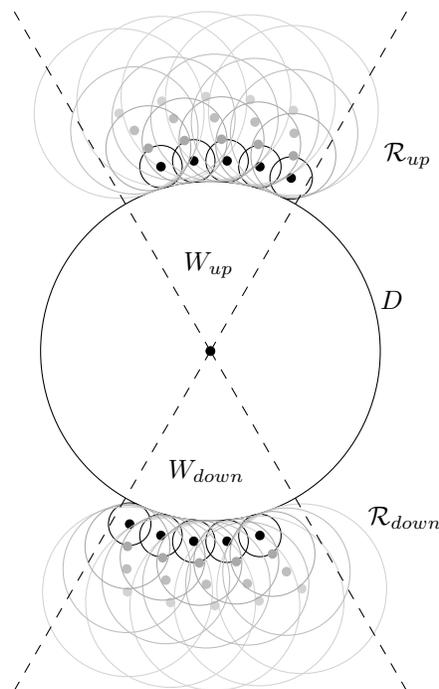


Figure 10: Example where a set of disks where the centers have $\Delta_k = O(\sqrt{n})$ for constant k still requires $\Omega(n^2)$ time to update when moving disks one at a time.

6 Conclusion

We presented algorithms for maintaining the convex hull and the Delaunay triangulation of a planar point set P in the KDS black-box model. The algorithms are simple and do not require knowledge of the k -spread $\Delta_k(P)$ or k : the convex-hull algorithms need to know only d_{\max} , the maximum displacement of any point in one time step, and the Delaunay-triangulation algorithm needs no knowledge at all. Our main contribution lies in the analysis of these algorithms under the Displacement Assumption and in terms of Δ_k .

For the convex-hull maintenance we spend $O(k\Delta_k \log^2 n)$ amortized time. This is optimal up to the logarithmic factors, because the convex hull can undergo $\Omega(k\Delta_k)$ changes in any time step. Moreover, we can show that our bound $O(k\Delta_k \log n)$ on the number of expiring time stamps is tight in the worst case. However, it may be possible to get rid of one logarithmic factor from the time bound by a more clever algorithm. In fact, when we can use the floor function, then we know how to do this. Unfortunately, the algorithm needs to know $\text{mindist}_k(P(t))$, which is perhaps not realistic. It would be interesting to design an algorithm that needs to know only d_{\max} and achieves $O(k\Delta_k \log n)$ update time. Another interesting open problem is whether it is possible to make the time bound worst-case rather than amortized.

For the Delaunay triangulation we have shown that a simple flipping algorithm needs $O(k^2\Delta_k^2)$ flips. The bound is based on an analysis of the number of potential edges—that is, all edges that can possibly arise in one time step. Our bound on the number of potential

edges is tight in the worst case and we show that there is also an $\Omega(k^2\Delta_k^2)$ lower bound on the number of edges that appear in the worst case when moving points one at a time. We also showed how to use the Delaunay triangulation to easily maintain a connectivity structure for units disks and proved that we cannot easily extent this to arbitrary size disks.

It depends on the application how realistic our model is. We expect that most sampling rates are such that the Displacement Assumption is satisfied. A valid question is whether point sets can be expected to have small k -spread. In meshing-type applications, it may be realistic to assume that the k -spread is $O(\sqrt{n})$; our results then imply that the simple flipping approach needs only $O(n)$ flips. In any case, Δ_k seems like a reasonable parameter to measure efficiency. We think it will be interesting to study other structures in the KDS black-box model under the Displacement Assumption and to analyze their performance in terms of the k -spread Δ_k .

References

- [1] M. Abam, P.K. Agarwal, M. de Berg, and H. Yu. Out-of-order event processing in kinetic data structures. In *Proc. 14th Ann. Europ. Sympos. Algorithms*, LNCS 4168, pages 624–635, 2006.
- [2] P.K. Agarwal, L.J. Guibas, H. Edelsbrunner, J. Erickson, M. Isard, S. Har-Peled, J. Hershberger, C. Jensen, L. Kavraki, P. Koehl, M. Lin, D. Manocha, D. Metaxas, B. Mirtich, D. Mount, S. Muthukrishnan, D. Pai, E. Sacks, J. Snoeyink, S. Suri, and O. Wolfson. Algorithmic issues in modelling motion. *ACM Comput. Surv.* 34:550–572 (2002).
- [3] P.K. Agarwal, B. Sadri, and H. Yu. Untangling triangulations through local explorations. In *Proc. 24th ACM Sympos. Comput. Geom.*, pages 288–297, 2008.
- [4] J. Basch, L.J. Guibas, and J. Hershberger. Data structures for mobile data. In *Proc. 8th ACM-SIAM Sympos. Discr. Algorithms*, pages 747–756, 1997.
- [5] M. de Berg, M. van Kreveld, M. Overmars, O. Schwartzkopf. *Computational Geometry: Algorithms and Applications*. Springer-Verlag, Berlin, Germany 3rd edition, 2008.
- [6] P.M. Manhães de Castro, J. Tournois, P. Alliez, and O. Devillers. Filtering relocations on a Delaunay triangulation. In *Proc. Sympos. Geometry Processing*, pages 1465–1474, 2009
- [7] T. Chan. Output-sensitive results on convex hulls, extreme points, and related problems. *Discr. Comput. Geom.* 16: 369–387 (1996).
- [8] C.A. Wang and F. Chin. Finding the constrained Delaunay triangulation and constrained voronoi diagram of a simple polygon in linear-time. In *Proc. 3rd Ann. Europ. Sympos. Algorithms*, LNCS 979, pages 280–294, 1995.
- [9] M. Cho, D.M. Mount, and E. Park. Maintaining nets and net trees under incremental motion. In *Proc. 20th Sympos. Algo. Comput.*, pages 1134–1143, 2009.

- [10] O. Devillers. On deletion in Delaunay triangulations. In *Proc. 15th ACM Sympos. Comput. Geom.*, pages 181–188, 1999.
- [11] J. Erickson. Dense Point Sets Have Sparse Delaunay Triangulations. *Discr. Comput. Geom.* 30:83-115 (2005).
- [12] J. Gao, L.J. Guibas, J. Hershberger, L. Zhang and A. Zhu. Geometric spanner for routing in mobile networks. In *Proc. 2nd ACM Intl Sympos. on Mobile Ad Hoc Networking & Computing.*, pages 45–55, 2001.
- [13] J. Gao, L.J. Guibas, A. Nguyen. Deformable spanners and applications. In *Proc. 20th ACM Sympos. Comput. Geom.*, pages 190–199, 2004.
- [14] L.J. Guibas. Kinetic data structures—a state-of-the-art report. In *Proc. 3rd Workshop Algorithmic Found. Robot.*, pages 191–209, 1998.
- [15] L.J. Guibas. Kinetic data structures. In: D. Mehta and S. Sahni (editors), *Handbook of Data Structures and Applications*, Chapman and Hall/CRC, 2004.
- [16] L.J. Guibas. Motion. In: J. Goodman and J. O’Rourke (eds.), *Handbook of Discrete and Computational Geometry (2nd edition)*, pages 1117–1134. CRC Press, 2004.
- [17] L.J. Guibas, J. Hershberger, S. Suri and L. Zhang. Kinetic Connectivity for Unit Disks. In *Proc. 16th Annu. ACM Sympos. Comput. Geom.*, pages 331–340, 2000
- [18] S. Kahan. A model for data in motion. In *Proc. 23rd ACM Sympos. Theory Comput.*, pages 267–277, 1991.
- [19] M.I. Karavelas. Voronoi Diagrams for Moving Disks and Applications. In *Proceedings of the 7th International Workshop on Algorithms and Data Structures*, pages 62–74, 2001.
- [20] X.-Y. Li, G. Calinescu, and P.-J. Wan. Distributed Construction of Planar Spanner and Routing for Ad Hoc Wireless Networks. In *Proc. 21st IEEE INFOCOM, vol. 3*, pages 1268–1277, 2002.
- [21] D.M. Mount, N.S. Netanyahu, C.D. Piatko, R. Silverman, and A.Y. Wu. A computational framework for incremental motion. In *Proc. 20th ACM Sympos. Comput. Geom.*, pages 200–209, 2004.
- [22] D. Russel (2007). *Kinetic Datastructures in Practise*. Ph.D. thesis. Stanford University: U.S.A.
- [23] R. Shewchuk. Star splaying: an algorithm for repairing Delaunay triangulations and convex hulls. In *Proc. 21st ACM Sympos. Comput. Geom.*, pages 237–246, 2005.
- [24] K. Yi and Q. Zhang. Multi-dimensional online tracking. In *Proc. 20th ACM-SIAM Sympos. Discr. Algo.*, pages 1098–1107, 2009.