Faster Algorithms for Growing Prioritized Disks and Rectangles*

Hee-Kap Ahn[1], Sang Won Bae[2], Jongmin Choi[1], Matias Korman[3], Wolfgang Mulzer[4], Eunjin Oh[1], Ji-won Park[5], André van Renssen[6], and Antoine Vigneron[7]

[1]Department of Computer Science and Engineering, POSTECH, Republic of Korea, {heekap, icothos, jin9082}@postech.ac.kr
[2]Division of Computer Science and Engineering, Kyonggi University, Republic of Korea, swbae@kgu.ac.kr
[3]Department of Computer Science, Tufts University, United States, matias.korman@tufts.edu
[4]Institut für Informatik, Freie Universität Berlin, Germany, mulzer@inf.fu-berlin.de
[5]School of Computing, KAIST, Republic of Korea, wldnjs1727@kaist.ac.kr
[6]School of Computer Science, University of Sydney, Sydney, Australia, andre.vanrenssen@sydney.edu.au
[7]School of Electrical and Computer Engineering, UNIST, Republic of Korea, antoine@unist.ac.kr

Abstract

Motivated by map labeling, Funke, Krumpe, and Storandt [IWOCA 2016] introduced the following problem: we are given a sequence of n disks in the plane. Initially, all disks have radius 0, and they grow at constant, but possibly different, speeds. Whenever two disks touch, the one with the higher index disappears. The goal is to determine the elimination order, i.e., the order in which the disks disappear. We provide the first general subquadratic algorithm for this problem. Our solution extends to other shapes (e.g., rectangles), and it works in any fixed dimension.

We also describe an alternative algorithm that is based on quadtrees. Its running time is $O\big(n\big(\log n + \min\{\log \Delta, \log \Phi\}\big)\big)$, where Δ is the ratio of the fastest and the slowest growth rate and Φ is the ratio of the largest and the smallest distance between two disk centers. This improves the running times of previous algorithms by Funke, Krumpe, and Storandt [IWOCA 2016], Bahrdt *et al.* [ALENEX 2017], and Funke and Storandt [EuroCG 2017].

Finally, we give an $\Omega(n \log n)$ lower bound, showing that our quadtree algorithms are almost tight.

1 Introduction

Suppose we have a digital map in which certain locations are marked with textual labels. As we zoom out, the visible area increases, while the individual features become smaller. To keep readability, the labels must grow during the zooming process. At some point, the labels will collide, and some of them have to be eliminated to avoid clutter. For an efficient implementation, we would like to determine when and in which order the labels disappear.

This can be formalized as follows: we are given a sequence $D_1, \ldots, D_n$ of n *growing disks*. Each disk D_i starts out as a point $p_i \in \mathbb{R}^2$, and it grows with a fixed *growth rate* $v_i > 0$. Thus, at any time $t \geq 0$, the disk D_i is centered at p_i and has radius $v_i t$. The index i corresponds to the *priority* (a smaller

*A preliminary version appeared as H.-K.-Ahn, S. W. Bae, J. Choi, M. Korman, W. Mulzer, E. Oh, J.-W. Park, A. v. Renssen, A. Vigneron. *Faster Algorithms for Growing Prioritized Disks and Rectangles.* Proc. 28th ISAAC, pp. 3:1–3:13. The work by H.-K. Ahn, J. Choi, E. Oh was supported by the MSIT (Ministry of Science and ICT), Korea, under the SW Starlab support program (IITP–2017–0–00905) supervised by the IITP (Institute for Information & communications Technology Promotion.). S.W. Bae was supported by Basic Science Research Program through the National Research Foundation of Korea (NRF) funded by the Ministry of Education (2015R1D1A1A01057220 and 2018R1D1A1B07042755). M. Korman was supported in part by NSF awards CCF-1422311 and KAKENHI No. 17K12635, Japan. W. Mulzer was supported in part by DFG Grants MU 3501/1 and MU 3501/2 and by ERC StG 757609. J.-W. Park was supported by the NRF Grant 2011-0030044 (SRC-GAIA) funded by the Korea government (MSIP). A. v. Renssen was supported by JST ERATO Grant Number JPMJER1201, Japan. A. Vigneron was supported by the 2016 Research Fund (1.160054.01) of UNIST (Ulsan National Institute of Science and Technology).

Shape	Time	Space	Method	Sec.
Balls/Boxes, $\mathbb{R}^d$	$O(dn^2)$	$O(n)$	Priority sort	§2
Disks, $\mathbb{R}^2$	expected $O\big(n^{\frac{5}{3}+\varepsilon}\big)$	$O\big(n^{\frac{5}{3}+\varepsilon}\big)$	Bucketing	§3
Rectangles, $\mathbb{R}^2$	$O(n^{\frac{11}{6}+\varepsilon})$	$O\big(n^{\frac{11}{6}+\varepsilon}\big)$		
$\mathcal{SA}_k$, $k \geq 4$	$O\big(n^{2-\frac{1}{2k-2}+\varepsilon}\big)$	$O\big(n^{2-\frac{1}{2k-2}+\varepsilon}\big)$		
Cubes, $\mathbb{R}^d$	$O\big(n \log^{d+2} n\big)$	$O\big(n \log^{d+1} n\big)$	Orthogonality	§4
Disks, $\mathbb{R}^2$	$O(n \log \Phi \min\{\log \Delta, \log \Phi\})$	$O(n \log \Phi)$	Quadtree	§5.1
Disks, $\mathbb{R}^2$	$O(n(\log n + \min\{\log \Delta, \log \Phi\}))$	$O(n)$	Compressed quadtree	§5.2

Table 1: Summary of our results. The $O(dn^2)$-time algorithm in the first row works for growing objects of any shape in $\mathbb{R}^d$ such that the touching time of any two of them can be computed in $O(d)$ steps. $\mathcal{SA}_k$ stands for any semialgebraic shape that is described with k parameters. Φ denotes the *spread* of the disk centers and $\Delta = \max_i v_i / \min_j v_j$ is the maximum ratio between two growth rates.

index means a higher priority). When two disks meet, we eliminate the one with lower priority from the arrangement. More precisely, for $1 \leq i < j \leq n$, let $t(i,j) = |p_i p_j|/(v_i + v_j)$ be the time when D_i and D_j touch. Then, if neither D_i nor D_j has been removed before time $t(i,j)$, we eliminate D_j at this time, while D_i remains. Our goal is to determine the *elimination order*, that is, the instants of time and the order in which the disks are removed. In this version, we chose to represent the labels as disks, but many other shapes are possible, e.g., to model rectangular labels or map icons with more complex boundary shapes. Thus, it is desirable to have a solution that is widely applicable and adapts easily to small variations in the problem statement.

This problem was introduced by Funke, Krumpe, and Storandt [10]. They observed that a straightforward simulation of the growth process with a priority queue runs in time $O(n^2 \log n)$. They also gave an algorithm that takes expected time $O\big(n(\log^6 n + \Delta^2 \log^2 n + \Delta^4 \log n)\big)$, where $\Delta = \max_i v_i / \min_j v_j$ is the maximum ratio between two growth rates. Subsequently, Bahrdt *et al.* [2] improved this to an algorithm that runs in worst-case time $O\big(\Delta^2 n(\log n + \Delta^2)\big)$. This generalizes to growing balls in arbitrary fixed dimension d, with running time $O\big(\Delta^d n(\log n + \Delta^d)\big)$. Recently, Funke and Storandt [11] presented two further parameterized algorithms for the problem. The first algorithm runs in time $O\big(n \log \Delta(\log n + \Delta^{d-1})\big)$, for arbitrary dimension d, while the second algorithm is restricted to the plane and runs in time $O\big(Cn \log^{O(1)} n\big)$, where C denotes the number of distinct growth rates. If we are interested in finding only the first pair of touching disks, our problem becomes the *weighted closest pair* of the disk centers. Formann showed how to compute it in $O(n \log n)$ time, which is optimal [9].

Our results. We present a simple algorithm that runs in time $O(dn^2)$ in any fixed dimension d (Section 2). For a faster running time, this method can be combined with bucketing and an advanced query data structure for lower envelopes of algebraic surfaces [1, 15] (Section 3). This yields an algorithm with $O(n^{5/3+\varepsilon})$ expected time for disks and $O(n^{11/6+\varepsilon})$ expected time for rectangles in two dimensions. These are the first subquadratic-time algorithms for the problem. More generally, we show that the elimination sequence of a set of n growing objects of any semi-algebraic shape described with k parameters can be computed in subquadratic time, for any fixed $k \geq 4$. In Section 4, we consider the case of growing squares. These objects are much simpler, and we can use ray shooting techniques to get a near-linear running time of $O(n \log^{d+2} n)$.

We also consider a completely different approach based on quadtrees (Section 5). The running time of these algorithms also depends on the *spread* Φ of the disk centers (the ratio of the maximum and the minimum distance between two disk centers) and the ratio Δ between the fastest and slowest growth rate. Table 1 summarizes our results. Finally, we give an $\Omega(n \log n)$ lower bound with a simple reduction from sorting. Our method using compressed quadtrees is thus nearly optimal.

Note. Parallel to our work, Castermans *et al.* [5] considered a related problem for growing squares in the plane. Whenever two squares meet, they are replaced by a new one located at their weighted center. They are also interested in the elimination/replacement sequence. Even though our algorithms are slightly faster (by polylogarithmic factors) and more general (their algorithm can only handle squares), they are not comparable, since our techniques do not apply in the setting where replacements are possible.

Notation. For any $1 \leq i \leq n$, we denote by t_i the time when disk D_i is eliminated. Since D_1 remains throughout, we set $t_1 = \infty$. We denote by $t(i,j) = |p_i p_j|/(v_i + v_j)$ the time at which the disks D_i and D_j would touch, supposing that no other disk has interfered. We assume general position in the sense that all times $t(i,j)$, for $i \neq j$, are pairwise distinct.

2 A simple quadratic algorithm

We present a simple iterative way to determine the elimination times t_i. This method will be useful in our bucketing scheme. As noted above, we have $t_1 = \infty$. We need a way to determine t_i if $t_1, \ldots, t_{i-1}$ are known, for $i \geq 2$. The next lemma shows how to do this: we just need to find the disk of higher priority that first touches D_i and is still alive at that point.

Lemma 2.1. *Let $i \in \{2, \ldots, n\}$, and let*

$$j^* = \operatorname{argmin}_{j=1,\ldots,i-1}\{t(i,j) \mid t(i,j) \leq t_j\}. \tag{1}$$

Then, the disk D_i is eliminated by the disk D_{j^*}, and $t_i = t(i,j^*)$.

Proof. On the one hand, we have $t_i \leq t(i,j^*)$, because at time $t(i,j^*)$, the disk D_i would meet the disk D_{j^*} that has higher priority and that has not been eliminated yet (by (1), we have $t(i,j^*) \leq t_{j^*}$). On the other hand, we have $t_i \geq t(i,j^*)$, because every disk that D_i could meet before time $t(i,j^*)$ either has lower priority or has been eliminated before the encounter. □

The condition (1) can be implemented with a straightforward **for**-loop. This leads to Algorithm 1.

Algorithm 1 A quadratic time algorithm

```
1:  function EliminationOrder(p_1, ..., p_n, v_1, ..., v_n)
2:      t_1 ← ∞
3:      for i ← 2, ..., n do
4:          t_i ← t(i, 1)
5:          for j ← 2, i − 1 do
6:              if t_j ≥ t(i, j) and t_i ≥ t(i, j) then
7:                  t_i ← t(i, j)
8:      S ← ⟨D_1, ..., D_n⟩
9:      Sort S using key t_i for each disk D_i
10:     return S
```

Theorem 2.2. *Algorithm 1 computes the elimination order of a set of prioritized disks in $O(n^2)$ time and $O(n)$ space. It generalizes to growing objects of any shape in $\mathbb{R}^d$ such that the touching time of any pair of them can be computed in $O(d)$ steps, with running time $O(dn^2)$.*

Proof. The correctness follows directly from Lemma 2.1. The running time and space analysis is straightforward. Lemma 2.1 is purely combinatorial and requires only that the times $t(i,j)$ are well defined. Thus, Algorithm 1 can be generalized by using an appropriate $O(d)$-time subroutine for computing $t(i,j)$. Then, the claim is immediate. □

In particular, Theorem 2.2 shows that the elimination order for d-dimensional balls or rectangles can be computed in $O(dn^2)$ time, for any $d \geq 1$.

3 A subquadratic algorithm using bucketing

We now improve Algorithm 1 with bucketing and an appropriate data structure. The main idea is as follows: in Algorithm 1, we go through the disks by decreasing order of priority and determine for each one which disk eliminates it. This is done by examining each disk of higher priority individually, leading to a quadratic running time. To avoid this, we partition the disks into *buckets* of size m, where m will be fixed later. For each bucket B, we build a data structure that can find in sublinear time the disk in B that eliminates the current disk. Thus, we process the m disks in each bucket with a single query, at the expense of an additional overhead for building the data structure. Furthermore, we must compute the elimination events within each bucket, using Algorithm 1.

Elimination queries. We now describe the query data structure that is used in a single bucket. Let $B \subseteq \{1, \ldots, n\}$ be a contiguous set of m indices, and suppose we know the elimination time t_j of every disks D_j with $j \in B$. In an *elimination query*, we are given a query index $q > \max B$, and we ask for the disk D_{j^*} with $j^* \in B$, that eliminates the query disk D_q (i.e., the first disk in B that meets D_q and that has not been eliminated yet by *any* other disk of higher priority). The argument from Lemma 2.1 shows that we can find j^* by the following slight adaptation of (1):

$$j^* = \operatorname{argmin}_{j \in B} \big\{ t(q, j) \mid t(q, j) \leq t_j \big\}. \tag{2}$$

This leads to a natural interpretation of elimination queries as vertical ray shooting with four-dimensional lower envelopes: a growing disk D corresponds to a point $(x, y, v) \in \mathbb{R}^3$, where (x, y) is the center of D and v is the growth rate. For each $j \in B$, consider the function $f_j : \mathbb{R}^3 \to \mathbb{R}$ defined by

$$f_j(x, y, v) = \begin{cases} t\big((x, y, v), j\big), & \text{if } t\big((x, y, v), j\big) < t_j, \\ \infty, & \text{otherwise,} \end{cases}$$

where $t\big((x, y, v), j\big)$ denotes the time when D_j and the growing disk given by (x, y, v) touch. For $q > \max B$, let $(x_q, y_q, v_q) \in \mathbb{R}^3$ be the point that represents D_q. Then, (2) tells us that the elimination query q corresponds to finding the value $\min_{j \in B} f_j(x_q, y_q, v_q)$ vertically above (x_q, y_q, v_q) and the index $\operatorname{argmin}_{j \in B} f_j(x_q, y_q, v_q)$ of the function that attains it. The pointwise minimum $E(w) = \min_{j \in B} f_j(w)$ is called the *lower envelope* of the functions f_j, and the problem of determining the function that achieves this minimum for a given $w \in \mathbb{R}^3$ is called a *vertical ray shooting query*. Vertical ray shooting in lower envelopes is a well-studied problem in computational geometry [12]. In our case, we can apply the following result of Agarwal *et al.* [1]:

Theorem 3.1 (Theorem 3.3 in Agarwal *et al.* [1])**.** *Let $\mathcal{F}$ be a given collection of m trivariate, possibly partially defined functions, all algebraic of constant maximum degree, and whose domains of definition (if they are partially defined) are each defined by a constant number of algebraic equalities and inequalities of constant maximum degree. Then, for any $\varepsilon > 0$, the lower envelope $E_{\mathcal{F}}$ of $\mathcal{F}$ can be computed in randomized expected time $O(m^{3+\varepsilon})$, and stored in a data structure of size $O(m^{3+\varepsilon})$, so that, given any query point $w \in \mathbb{R}^3$, we can compute $E_{\mathcal{F}}(w)$, as well as the function(s) attaining $E_{\mathcal{F}}$ at w in $O(\log^2 m)$ time.*

The data structure of Agarwal *et al.* [1] works by first computing the 0-, 1-, and 2-dimensional faces of the *minimization diagram* for $\mathcal{F}$, i.e., the maximally connected regions in $\mathbb{R}^3$ in which the lower envelope of $\mathcal{F}$ is achieved by four, three, or two functions from $\mathcal{F}$. For this, we go over all pairs f, g of distinct functions in $\mathcal{F}$, and we determine the part of the two-dimensional surface $\Gamma = \{x \in \mathbb{R}^3 \mid f(x) = g(x)\}$ where the lower envelope is achieved by f and g. This can be done in expected time $O(m^{1+\varepsilon})$ with a randomized incremental construction, by inserting the functions from $\mathcal{F} \setminus \{f, g\}$ in a random order and keeping track of (the trapezoidal decomposition of) the part of Γ that is not cut off by any function inserted so far [1, Theorem 2.3]. Once the faces of the minimization diagram are available, it can be converted into a suitably well-behaved subdivision of $\mathbb{R}^3$ [1, Lemma 3.1] to which the point location structure of Preparata and Tamassia can be applied [18]. The structure of Preparata and Tamassia

requires $O(m^{3+\varepsilon})$ space and preprocessing time, while a query needs $O(\log^2 m)$ steps. The result from Theorem 3.1 follows. This directly translates into the following lemma on elimination queries:

Lemma 3.2. *Let $B \subseteq \{1, \ldots, n\}$ with $|B| = m$. Then, for any fixed $\varepsilon > 0$, elimination queries for B can be answered in $O(\log^2 m)$ time, using space and expected preprocessing time $O(m^{3+\varepsilon})$.*

Bucketing. With all the tools in place, we can now describe our subquadratic algorithm. We group the disks into $\lceil n/m \rceil$ *buckets* $B_1, \ldots, B_{\lceil n/m \rceil}$, such that the kth bucket B_k contains the disks $D_{(k-1)m+1}$, $\ldots$, D_{km} (the last bucket might not be full). There are $O(n/m)$ buckets, each with at most m disks. As in Algorithm 1, we compute the elimination times $t_1, \ldots, t_n$, in this order. As soon as the elimination times of all the disks in a bucket B_k have been determined, we construct the elimination query data structure for B_k. By Lemma 3.2, for each bucket, this takes $O(m^{3+\varepsilon})$ expected time and space, for a total of $O(nm^{2+\varepsilon})$ expected time and space.

Now, in order to determine the elimination time t_i of a disk D_i, we must check all the previous buckets (as well as the bucket containing D_i). We first perform elimination queries for all previous buckets, that is, the buckets B_k with $1 \leq k \leq \lfloor (i-1)/m \rfloor$. There are $O(n/m)$ such queries, each taking $O(\log^2 m)$ time. Then, we handle the disks that are in the same bucket as D_i by inspecting all of them, which takes $O(m)$ time. We return the smallest of all the resulting elimination times. The total time, for $i = 1, \ldots, n$, is thus $O((n^2/m) \log^2 m + nm)$. Overall, we obtain an expected running time of $O(nm^{2+\varepsilon} + (n^2/m) \log^2 m)$. We balance the terms by setting $m = \lfloor n^{1/3} \rfloor$, to get an algorithm that takes $O(n^{5/3+\varepsilon})$ space and expected time.

Theorem 3.3. *The elimination sequence of a set of n growing disks can be computed in $O(n^{5/3+\varepsilon})$ space and expected time, for any fixed $\varepsilon > 0$.*

Generalizations. The subquadratic algorithm generalizes to other shapes. For example, consider the problem of growing rectangles in $\mathbb{R}^2$. Each rectangle is given by the x- and y-coordinates of two opposing corners at time $t = 1$ (this lets us deduce the center and the speed of the rectangle). These are four parameters, so elimination queries now can be handled by vertical ray shooting in lower envelopes in $\mathbb{R}^5$. In this setting, we employ a general data structure for point location in arrangements of high-dimensional surfaces or surface patches. Chazelle *et al.* [8] presented such a structure, based on geometric divide and conquer. Their result requires a decomposition of (subsets of) this arrangement into *elementary cells* of constant complexity and depends on the complexity of such a decomposition. The best relevant bounds are due to Koltun [15]. By plugging his combinatorial bounds into the scheme of Chazelle *et al.* [8], he obtains the following theorem:

Theorem 3.4 (Theorem 5.1 in [15])**.** *A collection $\mathcal{F}$ of m fixed-degree algebraic surfaces or surface patches in $\mathbb{R}^d$, for $d \geq 4$, can be preprocessed in time $O(m^{2d-4+\varepsilon})$ into a data structure of size $O(m^{2d-4+\varepsilon})$ such that the cell in the arrangement of $\mathcal{F}$ that contains a query point can be located in time $O(\log m)$.*

A few remarks are in order: the bound in Theorem 3.4 holds for any fixed $\varepsilon > 0$. The preprocessing algorithm is deterministic (for this, we need to replace the random sampling step in the scheme of Chazelle *et al.* [8] with the deterministic ε-net construction by Chazelle and Matoušek [6, Theorem 4.6], which can be done in a black box fashion). Since the point location actually takes place in the vertical decomposition of the arrangement of $\mathcal{F}$, Theorem 3.4 can also be used for vertical ray shooting. Thus, for our case of m growing rectangles, Theorem 3.4 means that a data structure for elimination queries with query time $O(\log m)$ can be constructed in $O(m^{6+\varepsilon})$ space and time, for any fixed $\varepsilon > 0$.

We now apply the same approach as for growing disks. With bucket size m, we then obtain an algorithm that determines the elimination order in $O(nm^{5+\varepsilon} + (n^2/m) \log m)$ time. Setting $m = \lfloor n^{1/6} \rfloor$ to balance the terms, we get the following result:

Theorem 3.5. *The elimination sequence of a set of n growing rectangles can be computed in $O(n^{11/6+\varepsilon})$ deterministic time and space, for any $\varepsilon > 0$.*

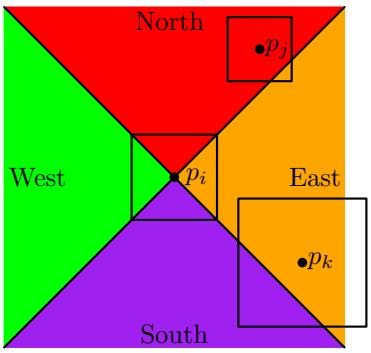

Figure 1: The four quadrants for a growing square D_i with center p_i. If the center p_j of D_j lies in the north quadrant, then the possible elimination time between D_j and D_i is determined by the y-coordinates (and similarly for D_k).

More generally, we can use Theorem 3.4 to handle elimination queries for regions defined by any semi-algebraic shape of constant complexity. If the shape of the object is described with $k \geq 4$ parameters, elimination queries translate to vertical ray shooting in the lower envelope of m surfaces or surface patches in $\mathbb{R}^{k+1}$. The bucketing approach and Theorem 3.4 then yield an algorithm that needs $O(nm^{2k-3+\varepsilon} + (n^2/m)\log n)$ time and space. We balance the terms with $m = \lfloor n^{1/(2k-2)} \rfloor$. This gives an overall running time of $O\left(n^{\frac{4k-5}{2k-2}+\varepsilon}\right)$, which is subquadratic for any fixed $k \geq 4$.

Theorem 3.6. *The elimination sequence of a set of n growing objects of any semi-algebraic shape, each described with $k \geq 4$ parameters can be computed in $O\left(n^{2-\frac{1}{2k-2}+\varepsilon}\right)$ deterministic time and space, for any $\varepsilon > 0$.*

4 Growing cubes

Axis-aligned cubes in $\mathbb{R}^d$ are given by $d+1$ parameters. Thus, the general approach from Section 3 applies. However, for axis-aligned cubes, much better data structures for orthogonal range searching and for planar ray-shooting can be leveraged for elimination queries. In this section, we combine bucketing with orthogonal range searching and ray shooting techniques to achieve an almost linear bound.

To simplify the presentation, we first focus on the case $d = 2$. A sequence of n growing squares is given by the centers $p_1, \ldots, p_n$ and the growth rates $v_1, \ldots, v_n$. At time $t \geq 0$, each square D_i has edge length $2v_i t$. We consider the four *quadrants* around each center $p_i = (x_i, y_i)$. The *north*, *east*, *south*, and *west* quadrants are, respectively, $\{(x,y) \in \mathbb{R}^2 \mid y - y_i \geq |x - x_i|\}$, $\{(x,y) \in \mathbb{R}^2 \mid x - x_i \geq |y - y_i|\}$, $\{(x,y) \in \mathbb{R}^2 \mid -(y - y_i) \geq |x - x_i|\}$, and $\{(x,y) \in \mathbb{R}^2 \mid -(x - x_i) \geq |y - y_i|\}$, see Figure 1.

The quadrants determine which sides of the growing square D_i are relevant for the elimination time. For example, suppose that the center p_j of the disk D_j is in the north quadrant of p_i. Then, D_i and D_j will meet when the top side of D_i touches the bottom side of D_j, and the possible elimination time of D_i and D_j is $t(i,j) = (y_j - y_i)/(v_i + v_j)$. This can be used for more efficient elimination queries as follows: suppose we have a set $B \subset \{1, \ldots, n\}$ of m growing cubes, and let $q > \max B$ such that all centers p_j with $j \in B$ lie in the north quadrant of p_q. Then, an elimination query for q in B just depends on the y-coordinates of the disk centers and the growth rates, the x-coordinates become irrelevant. Thus, we can reduce these elimination queries to two-dimensional ray-shooting.

Lemma 4.1. *Let $B \subseteq \{1, \ldots, n\}$, $|B| = m$. We can preprocess B in $O(m \log m)$ time into a data structure of $O(m)$ space, so that elimination queries can be answered with $O(\log m)$ time, given that the centers of the squares in B lie in the north quadrant of the query square D_q.*

Proof. We equip the plane with a coordinate system in which the horizontal direction is labeled t and the vertical direction is labeled y. For each $j \in B$, consider the line segment $f_j : t \mapsto y_j - v_j t$, defined for

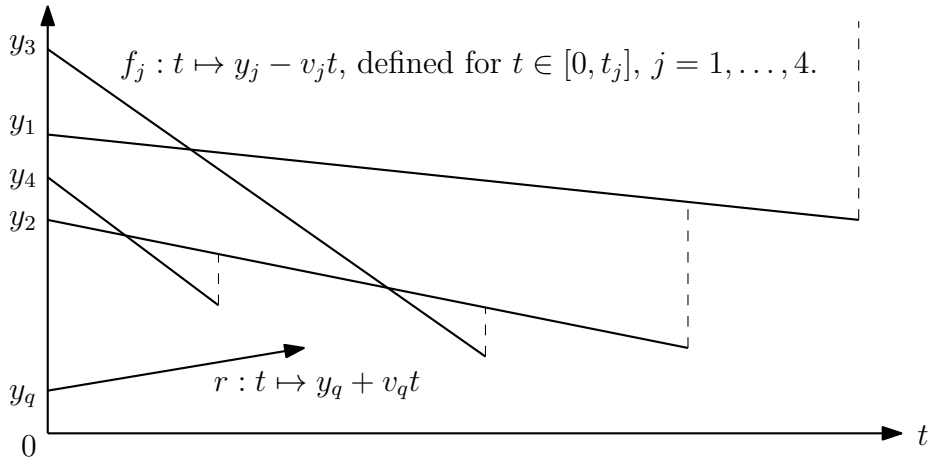

Figure 2: The lower envelope of four line segments. An elimination query for a square D_q with center (x_q, y_q) and growth rate v_q consists of shooting a ray $t \mapsto y_q + v_q t$ from below.

$t \in [0, t_j]$; see Figure 2. Let $E(t) = \min_{j \in B} f_j(t)$ be the lower envelope of the line segments. An *edge* of E is a maximal contiguous interval where the minimum is achieved by a single function f_j. All the line segments f_j begin on the y-axis. This means that for any pair f_i, f_j of line segments, f_i and f_j can alternate at most twice on E (i.e., there can be an edge from f_j, later an edge from f_i, and again later an edge from f_j, but after that f_i cannot appear again). Combinatorially, this corresponds to a *Davenport-Schinzel sequence* of order 2 with alphabet size m, i.e., a sequence σ of symbols from a finite set Σ with m elements such that for any two symbols $a, b \in \Sigma$, the pattern $a \ldots b \ldots a \ldots b$ does not appear in σ [19]. It is well known that such a sequence σ has length at most $2m - 1 = O(m)$ [16, Chapter 7.1].

An elimination query for a square D_q with center (x_q, y_q) and growth rate v_q translates to shooting a ray r from the point $(0, y_q)$ with slope v_q. Since r and all line segments f_j originate from the y-axis, and since r starts below the lower envelope E, the point where r meets the first line segment f_j must lie on E. Thus, we first compute E in $O(m \log m)$ time [14], and then we build a ray-shooting data structure for E. The latter can be done with $O(m)$ preprocessing time and space, and a query time of $O(\log m)$ [7]. The result follows. □

We now show how to handle more general elimination queries where we do not require B to be in the north quadrant of D_q. This is done using methods from orthogonal range searching.

Lemma 4.2. *Let $B \subseteq \{1, \ldots, n\}$, $|B| = m$. We can preprocess B in time $O(m \log^3 m)$ into a data structure of $O(m \log^2 m)$ space, so that elimination queries can be answered in $O(\log^3 m)$ time.*

Proof. For any square D_q, we must be able to perform elimination queries on B for all four quadrants that are defined by the center p_q of D_q.

This is done as follows: let P be the centers of the squares in B. We build a two-dimensional range tree for B [3, Chapter 5.3], where the coordinate axes have been rotated by an angle of $\pi/4$. More precisely, let the rotated coordinates be z_1 and z_2. We sort P according to the z_1 coordinate, and we build a perfect binary tree T on P for this order, such that the leaves of T correspond to the points in P. The height of T is $O(\log m)$, and each node ν of T corresponds to a subset P_ν of P, namely the leaves in the subtree under ν. Each point of P appears in $O(\log m)$ subsets. Next, we determine the z_2-order in each subset P_ν, and we build a binary tree T_ν on each P_ν for that order, as before. Again, for every node ν of T, each node μ of T_ν corresponds to a subset $P_{\nu\mu}$ of P_ν. We call the subsets $P_{\nu\mu}$ the *canonical sets*, see Figure 3(left) for an example. Each point in P_ν appears in $O(\log m)$ subsets $P_{\nu\mu}$, so the total size of all canonical sets is $O(m \log^2 m)$. For each canonical set $P_{\nu\mu}$, we build four elimination query structures as in Lemma 4.1, one for each quadrant. The total preprocessing time is $O(m \log^3 m)$, and the total space requirement is $O(m \log^2 m)$.

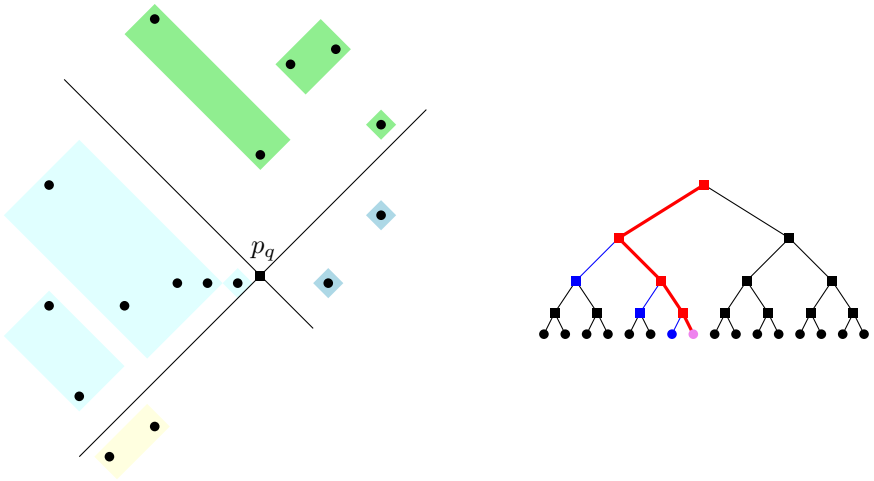

Figure 3: (left) A point set B and the canonical sets for a query point p_q; (right) the bucketing scheme: each node in the binary tree corresponds to a bucket. The current square is D_8 (the eighth leaf from the left). To process D_8, we query all buckets that are left children of the nodes along the path from the root to the leaf for D_8.

Now, to process a query q, we determine the four quadrants of p_q, and for each quadrant Q, we find the canonical sets that constitute a partition of $P \cap Q$. For this, we locate the z_1-coordinate of the vertical boundary of Q in T, and we take all the nodes that are left or right children of the nodes on this search path, depending on whether Q lies to the left or to the right of p_q in the rotated coordinate system. Then, we perform an analogous query in the second level tree for each such node, using the z_2-coordinate of the horizontal boundary of Q. This gives $O(\log^2 m)$ canonical sets that constitute a partition of $P \cap Q$. We do an elimination query for each canonical set, and we return the element that gives the minimum elimination time for D_q. We repeat this for all four quadrants, and we find the element with the minimum overall elimination time. Since we query $O(\log^2 m)$ canonical sets in total, this takes $O(\log^3 m)$ time. □

As in Section 3, we now apply Lemma 4.2 together with the bucketing technique. This time, however, the bucketing is done in a slightly different way: we construct a perfect binary tree T whose leaves represent the squares $D_1, \ldots, D_n$, in that order. A node $\nu \in T$ represents the subset B_ν of disks that consists of the leaves in the subtree rooted in ν.

As soon as the elimination times of all the disks B_ν associated with a node ν of T have been determined, we compute the elimination query structure from Lemma 4.2 for B_ν. Thus, when processing a disk D_i, the disks D_j, $j < i$, can be partitioned into $O(\log n)$ buckets for which elimination query structures have been constructed (at most one node per level in the tree), see Figure 3(right). Hence, we can find t_i in $O(\log^4 n)$ time by querying all these structures. Since each disk apperas in $O(\log n)$ structures, the total preprocessing time is $O(n \log^4 n)$ and the total space requirement is $O(n \log^3 n)$. In higher dimensions, these bounds increase by a factor $O(\log n)$ per dimension, as we need one more level in the range tree. The following theorem summarizes our result.

Theorem 4.3. *The elimination sequence of a set of n axis-aligned cubes in fixed dimension $d = O(1)$ can be computed in $O(n \log^{d+2} n)$ time, using $O(n \log^{d+1} n)$ space.*

5 Quadtree-based approach

Algorithm 1 can be improved further by noticing that in order to find out when a disk is eliminated, it suffices to check only disks that are *nearby*. Thus, we need a suitable data structure to maintain the proximity relations between the disks. For this, we use a *quadtree*. The nodes in the quadtree allow us

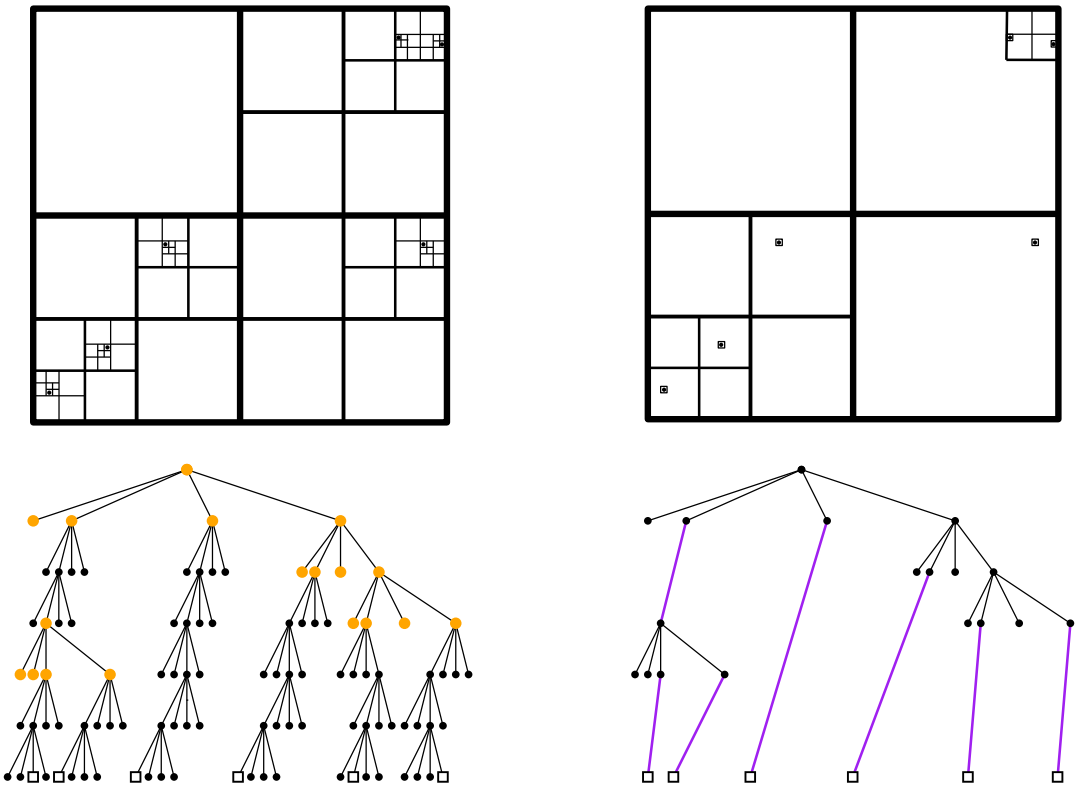

Figure 4: A quadtree and its compressed version: (left) a quadtree for 6 disk centers, where the subdivision process stops once a cell contains at most one disk center and the diameter of the cell becomes smaller than a quarter of the smallest distance between disk centers. The nodes that appear also in the compressed quadtree are marked; (right) the compressed quadtree. The compressed edges are shown in bold purple.

to approximate the disks by square cells and to efficiently maintain subsets of nearby disks during the growth process. The main drawback is that the complexity of the algorithm will depend on the structure of the point set and not just on the number of points. By compressing the quadtree, we can reduce this dependency, but not completely eliminate it. More precisely, let Φ denote the *spread* of the disk centers and Δ denote the ratio of the growth rates, i.e.,

$$\Phi = \frac{\max_{1 \le i < j \le n} |p_i p_j|}{\min_{1 \le i < j \le n} |p_i p_j|}$$

and

$$\Delta = \frac{\max_{i \in \{1,\ldots,n\}} v_i}{\min_{j \in \{1,\ldots,n\}} v_j}.$$

We provide two algorithms: the first algorithm uses a quadtree and runs in $O(n \log \Phi \min\{\log \Phi, \log \Delta\})$ time and $O(n \log \Phi)$ space. The second algorithm uses a compressed quadtree, and it needs $O(n(\log n + \min\{\log \Phi, \log \Delta\}))$ time and $O(n)$ space. To simplify the notation, we set $\alpha = \min\{\log \Phi, \log \Delta\}$.

5.1 A simple quadtree-based algorithm

We give a simple quadtree-based algorithm to compute the elimination sequence for a set of growing disks in the plane. For this, we first construct a modified quadtree $\mathcal{Q}$ for the disk centers: suppose that the maximum distance between two disk centers is 1 and the minimum distance is $1/\Phi$. Then, $\mathcal{Q}$ will have depth $O(\log \Phi)$, and each disk center will be in a leaf cell of diameter $O(1/\Phi)$. The growth process of a disk D_i can be approximated by tracing the quadtree cells from the leaf cell that contains p_i to the

root. It turns out that if two disks collide, then their respective cell approximations must be "close" in the quadtree and of approximately the same size (roughly up to a factor of Δ). This means that for each approximate disk (i.e., quadtree cell), there are $O(\alpha)$ other approximate disks that could potentially eliminate it. Thus, we reduce the number of disk pairs that need to be checked in the main loop of Algorithm 1 from $O(n^2)$ to $O(n\alpha \log \Phi)$. Details follow.

Without loss of generality, we assume that all disk centers lie in the unit square $[0,1]^2$, and that the diameter of this point set is 1. We construct a *quadtree* $\mathcal{Q}$ for the disk centers [17], with a slight modification. The quadtree is a rooted tree in which every internal node has four children. Each node ν of $\mathcal{Q}$ has an associated square *cell* $b(\nu)$. To obtain $\mathcal{Q}$, we recursively split the unit square. In each step, the current node ν is partitioned into four congruent quadrants (cells) if its corresponding cell $b(\nu)$ contains one or more disk centers. We stop when each cell at the bottom level contains at most one disk center and the diameter of the cell becomes smaller than a quarter of the smallest distance between two disk centers. Note that in our quadtree, all leaf cells have the same size and are small in relation to the smallest distance between two disk centers; see Figure 4. The quadtree $\mathcal{Q}$ has depth $O(\log \Phi)$, and it can be constructed in $O(n \log \Phi)$ time and space.

We introduce some notation. For a node $\nu \in \mathcal{Q}$, we let $p(\nu)$ be the parent node of ν. We denote by $|\nu|$ the diameter of the cell $b(\nu)$. For two nodes $\nu, \mu \in \mathcal{Q}$, we write $d(\nu, \mu)$ for the smallest distance between a point in $b(\nu)$ and a point in $b(\mu)$. For a point q and a node $\nu \in \mathcal{Q}$, we write $d(q, \nu)$ for the smallest distance between q and a point in $b(\nu)$. The next definitions show what it means to *approximate* a disk by a quadtree cell. For $t \geq 0$, we let D_i^t be the disk D_i at time t. We say that D_i^t *occupies* a node ν if

(i) the disk center p_i lies in the cell for ν, i.e., $p_i \in b(\nu)$;

(ii) ν is a leaf or D_i^t covers the whole cell for ν, i.e., $b(\nu) \subseteq D_i^t$; and

(iii) D_i^t has not been eliminated before time t.

We denote by $\nu(i,t)$ the node of the largest cell of $\mathcal{Q}$ that is occupied by D_i^t. We may think of $\nu(i,t)$ as an approximate representation in $\mathcal{Q}$ of disk D_i at time t. The next lemma shows that if two disks meet, then their respective approximate representations must be close and of approximately the same size.

Lemma 5.1. *Let $i \in \{2, \ldots, n\}$, and let D_j, $j \in \{1, \ldots, i-1\}$) be the disk that eliminates D_i, i.e., $t_i = t(i,j)$. Then,*

$$d\left(\nu(i,t_i), \nu(j,t_i)\right) \leq 2\left(|\nu(i,t_i)| + |\nu(j,t_i)|\right),$$

and

$$\frac{1}{4\Delta} \leq \frac{|\nu(i,t_i)|}{|\nu(j,t_i)|} \leq 4\Delta.$$

Proof. We state three simple facts from the construction of $\mathcal{Q}$ and from the definition of $\nu(\cdot,\cdot)$:

(i) all non-empty leaf cells have the same diameter (this is how we construct $\mathcal{Q}$);

(ii) for any $k \in \{1, \ldots, n\}$ and $t > 0$, if $\nu(k,t)$ is not a leaf, then $|\nu(k,t)| \leq 2v_k t$ (D_k^t has radius $v_k t$ and covers $\nu(k,t)$); and

(iii) for any $k \in \{1, \ldots, n\}$ and $t \geq 0$, we have $|\nu(k,t)| \geq v_k t/2$ (D_k^t has radius $v_k t$ and does not cover the parent of $\nu(k,t)$, which has diameter $2\,|\nu(k,t)|$).

For the first claim, on the distance between the cells, let $q = \partial D_i^{t_i} \cap \partial D_j^{t_i}$ be the point where D_i and D_j touch; see Figure 5. By fact (iii), we have $v_i t_i \leq 2\,|\nu(i,t_i)|$ and $v_j t_i \leq 2\,|\nu(j,t_i)|$. Hence, it follows that

$$d\left(\nu(i,t_i), \nu(j,t_i)\right) \leq d\left(q, \nu(i,t_i)\right) + d\left(q, \nu(j,t_i)\right) \leq v_i t_i + v_j t_i \leq 2\,|\nu(i,t_i)| + 2\,|\nu(j,t_i)|,$$

as claimed. Now we prove the second claim, on the ratio of the cell sizes; see Table 2 for an overview of the case distinction. Suppose first that $v_i \geq v_j$. If $\nu(j,t_i)$ is a leaf, we have $|\nu(i,t_i)| \,/\, |\nu(j,t_i)| \geq 1$, by fact (i). If $\nu(j,t_i)$ is not a leaf, it follows from facts (ii) and (iii) that

$$\frac{|\nu(i,t_i)|}{|\nu(j,t_i)|} \geq \frac{v_i t_i/2}{2 v_j t_i} \geq \frac{1}{4} \geq \frac{1}{4\Delta}.$$

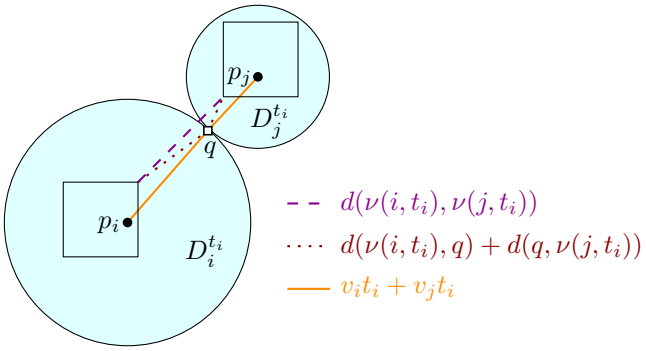

Figure 5: When two disks meet, the two cells that represent them are close.

	Lower Bound		Upper Bound
$v_i \geq v_j$	$\nu(j,t_i)$ is leaf	Fact (i)	Facts (ii), (iii), as $\nu(i,t_i)$ is never a leaf
	$\nu(j,t_i)$ is not leaf	Facts (ii), (iii)	
$v_i < v_j$	$\nu(i,t_i)$ is leaf	Fact (i)	Facts (ii), (iii), as $\nu(j,t_i)$ is never a leaf
	$\nu(i,t_i)$ is not leaf	Facts (ii), (iii)	

Table 2: Schematic overview for bounding the ratio of the cell sizes.

By construction, the leaf cell that contains p_i has diameter smaller than a quarter of the smallest distance between disk centers. Hence, the node $\nu(i,t_i)$ is not a leaf. Thus, by facts (ii) and (iii),

$$\frac{|\nu(i,t_i)|}{|\nu(j,t_i)|} \leq \frac{2v_it_i}{v_jt_i/2} \leq 4 \cdot \frac{\max_i v_i}{\min_j v_j} = 4\Delta.$$

The argument for $v_j > v_i$ is analogous: if $\nu(i,t_i)$ is a leaf, then $|\nu(j,t_i)| \,/\, |\nu(i,t_i)| \geq 1$, by fact (i). If not,

$$\frac{|\nu(j,t_i)|}{|\nu(i,t_i)|} \geq \frac{v_jt_i/2}{2v_it_i} > \frac{1}{4} \geq \frac{1}{4\Delta},$$

by facts (ii) and (iii). Now, the node $\nu(j,t_i)$ cannot be a leaf, so by facts (ii) and (iii)

$$\frac{|\nu(j,t_i)|}{|\nu(i,t_i)|} \leq \frac{2v_jt_i}{v_it_i/2} \leq 4 \cdot \frac{\max_j v_j}{\min_i v_i} = 4\Delta.$$

The lemma follows. □

Lemma 5.1 gives a necessary condition for the event that two quadtree cells could lead to an elimination event. This allows us to focus on a limited number of candidate pairs. More formally, we say that two unrelated[1] nodes $\nu, \mu \in \mathcal{Q}$ form a *candidate pair* if

$$d(\nu,\mu) \leq 2\,(|\nu| + |\mu|) \tag{*}$$

and

$$\frac{|\nu|}{4\Delta} \leq |\mu| \leq 4\Delta|\nu|. \tag{**}$$

[1]That is, no node is an ancestor or descendant of the other node.

We say that ν *forms the candidate pair* (ν, μ) with μ. We denote by $\mathsf{CNP}(\nu)$ the set of candidate pairs formed by ν. By Lemma 5.1, to determine the elimination order, we need to consider only candidate pairs. The following lemma uses a simple volume argument to bound their number.

Lemma 5.2. *Let $\nu \in \mathcal{Q}$. Then, $\mathsf{CNP}(\nu)$ contains $O(\alpha)$ candidate pairs (ν, μ) with $|\nu| \leq |\mu|$. The total number of candidate pairs is $O(n\alpha \log \Phi)$.*

Proof. Fix a node $\nu \in \mathcal{Q}$. We claim that in each level of $\mathcal{Q}$, there are at most $O(1)$ nodes μ with $|\nu| \leq |\mu|$ such that (ν, μ) is a candidate pair. This follows from a simple volume argument: fix a level i of $\mathcal{Q}$, and let X be the set of all nodes μ at level i with $|\nu| \leq |\mu|$ and $(\nu, \mu) \in \mathsf{CNP}(\nu)$. By (*) and our assumption $|\nu| \leq |\mu|$, we have for all $\mu \in X$ that

$$d(\nu, \mu) \leq 2(|\nu| + |\mu|) \leq 4|\mu|.$$

Hence, all cells $\mu \in X$ lie in a region of diameter at most $9|\mu|$. Since these cells have pairwise disjoint interiors and the same size $|\mu|$, it follows that $|X| = O(1)$, as claimed.

Furthermore, by the definition of Φ and by (**), we have $|\mu| = O(\min\{\Phi, \Delta\})|\nu|$. This implies that the levels of ν and μ in $\mathcal{Q}$ differ by $O(\alpha)$. Hence, $\mathsf{CNP}(\nu)$ contains $O(\alpha)$ candidate pairs (ν, μ) with $|\nu| \leq |\mu|$. Since $\mathcal{Q}$ has $O(n \log \Phi)$ nodes, and since the symmetry of (*) and (**) shows that $(\nu, \mu) \in \mathsf{CNP}(\nu)$ if and only if $(\mu, \nu) \in \mathsf{CNP}(\mu)$, it follows that there are $O(n\alpha \log \Phi)$ candidate pairs overall. □

Now we can describe our algorithm. As in Algorithm 1, we handle the disks in decreasing order of priority. Recall that by (1), for each disk D_i, we need to find the disk D_j, $1 \leq j < i$, that first meets D_i and is still alive. By Lemma 5.1, it suffices to focus on disks that correspond to candidate pairs. Throughout the algorithm, we compute for each node μ in $\mathcal{Q}$ the index $D(\mu)$ of the (at most one) disk that occupies it at some point in time. When processing a disk D_i, we start from the leaf node for D_i, and we simulate the growth of D_i by tracing a path to the root. As we follow the leaf-root path for D_i, we check all the candidate pairs $(\nu, \mu) \in \mathsf{CNP}(\nu)$ for the current node ν, to see if a node μ is occupied by a disk that could eliminate D_i (i.e., a disk of higher priority that is still alive when it meets D_i). These disks have been computed in previous iterations of the algorithm. We continue along the leaf-root path for D_i until it is clear that D_i has been eliminated before it can occupy the current node ν, setting the $D(\nu)$ variables accordingly. The pseudocode in Algorithm 2 provides the details. Initially, we set all $D(\nu) = \perp$. We use $\tau(\nu, i)$ to denote the first time at which $b(\nu)$ is covered by the disk D_i. We will describe below how the sets $\mathsf{CNP}(\nu)$ can be computed efficiently.

Algorithm 2 A simple quadtree-based algorithm

```
1:  function EliminationOrder(p_1, ..., p_n, v_1, ..., v_n)
2:      Q ← ConstructQuadTree(p_1, ..., p_n)
3:      D(ν) ←⊥ for every node ν of Q
4:      for i ← 1, ..., n do
5:          ν ← getLeaf(p_i)
6:          t ← ∞
7:          while ν ≠⊥ and t ≥ τ(ν, i) do
8:              D(ν) ← i
9:              for (ν, μ) in CNP(ν) do
10:                 if D(μ) ≠⊥ and t_{D(μ)} ≥ t(i, D(μ)) then
11:                     t ← min(t, t(i, D(μ)))
12:             ν ← p(ν)
13:         t_i ← t
14:     S ← (D_1, ..., D_n)
15:     Sort S using key t_i for each disk D_i
16:     return S
```

Theorem 5.3. *The elimination sequence of n growing disks can be computed in $O(n\alpha \log \Phi)$ time and $O(n \log \Phi)$ space, where $\alpha = \min\{\log \Phi, \log \Delta\}$.*

Proof. The outer **for**-loop iterates over the input disks by decreasing order of priority. In the **while**-loop, the algorithm traverses each node $\nu \in \mathcal{Q}$ from the leaf-node with p_i to the root. It updates $D(\nu)$, until it encounters a node ν with $t < \tau(\nu, i)$. The inner **for**-loop iterates over every candidate pair (ν, μ) in $\mathsf{CNP}(\nu)$. It checks if disk D_i and $D(\mu)$ might touch by computing the time $t(i, D(\mu))$; if so, it updates the tentative elimination time for D_i. To show correctness, we prove that the algorithm maintains the following invariant: after i iterations of the **for**-loop, the algorithm has correctly computed the elimination times $t_1, \ldots, t_i$ for $D_1, \ldots, D_i$. Furthermore, for each node ν of $\mathcal{Q}$, we have $D(\nu) \in \{1, \ldots, i, \perp\}$, and if there is a point in time when ν is occupied by the disk D_j, $1 \leq j \leq i$, then $D(\nu) = j$.

The invariant holds after the first iteration. This is because t is set to ∞ in Line 6, and all $D(\mu)$'s are initalized to $\perp$, so that the **while**-loop will proceed all the way to the root and set all $D(\nu)$'s along the leaf-root path for p_1 to 1. Now suppose that $i \geq 2$ and that D_i is eliminated by the disk D_j, with $1 \leq j < i$. By Lemma 5.1, the pair $(\nu, \mu) = (\nu(i, t_i), \nu(j, t_i))$ is a candidate pair in $\mathsf{CNP}(\nu)$. Furthermore, ν lies on the leaf-root path for p_i and by the inductive hypothesis, we have $D(\mu) = j$. Again by the inductive hypothesis and by the test $t_{D(\mu)} \geq t(i, D(\mu))$ in Line 10, we have $t \geq t_i$ throughout the **while**-loop. Thus, the **while**-loop will visit ν and the candidate pair (ν, μ) will be considered. At this point, Algorithm 2 will detect the elimination event and set t to t_i. After that, the test $t \geq \tau(\nu, i)$ in Line 7 ensures that the remaining variables $D(\nu)$ are set correctly. In particular, the algorithm does not overwrite any other such values from previous iterations. Thus, the invariant is maintained.

We now turn to the running time. We can compute $\mathcal{Q}$ in Line 2 in $O(n \log \Phi)$ time and space. Furthermore, by Lemma 5.2, there are $O(n\alpha \log \Phi)$ candidate pairs overall, so that the total time for the **for**-loop, excluding the time for finding the candidate pairs, is $O(n\alpha \log \Phi)$. We find the candidate pairs without additional asymptotic overhead as follows: to determine the leaf for p_i, we walk down from the root, going from one node that contains p_i to the next. Suppose that we are currently at level ℓ, and that the node ν contains p_i in $b(\nu)$. We find all nodes μ at level ℓ with $d(\nu, \mu) \leq 2(|\nu| + |\mu|)$, and we store them in a list for level ℓ. For this, we can use appropriate pointers in $\mathcal{Q}$, or we perform an appropriate root-leaf traversal of $\mathcal{Q}$ that keeps track of all nodes close to the nodes containing p_i. By a simple volume argument, there are $O(1)$ such nodes at each level, for a total of $O(\log \Phi)$. To find $\mathsf{CNP}(\nu)$ for a node ν, we consider the list for the parent of ν that is $O(\alpha)$ levels above ν (including the parent), and we perform a depth-first search in $\mathcal{Q}$ of these nodes and their descendants for all candidate pairs. This takes $O(\alpha + |\mathsf{CNP}(\nu)|)$ time, and it can be implemented so that the space requirement does not exceed $O(n \log \Phi)$, since we can process the candidate pairs immediately as we discover them. Finally, since $\Phi = \Omega(\sqrt{n})$, the sorting step in Line 16 does not increase the asymptotic running time or space.[2] □

5.2 Using a compressed quadtree

We now speed up Algorithm 2 with the help of a *compressed quadtree* $\mathcal{Q}_C$. In $\mathcal{Q}_C$, the number of nodes is reduced to $O(n)$ by replacing certain long paths in the uncompressed quadtree $\mathcal{Q}$ with single *compressed* edges. Now, the definition of the candidate pairs becomes more tricky, because the cells identified in Lemma 5.1 might no longer be present after the compression. Thus, we need a way to map candidate pairs in $\mathcal{Q}$ to candidate pairs in $\mathcal{Q}_C$. It is important that no elimination event is missed and that the number of candidate pairs is small. To achieve this we project a candidate pair (ν, μ) in $\mathcal{Q}$ to the pair given by the lowest ancestors of ν and of μ that appear in $\mathcal{Q}_C$. With these *compressed candidate pairs*, we can essentially run Algorithm 2 on $\mathcal{Q}_C$, with minor modifications. To complete our result, we must show that the compressed candidate pairs are few and can be found efficiently, and that no elimination event is missed. Details follow.

We begin with the formal definition of the compressed quadtree. Let $\mathcal{Q}$ be the (uncompressed) quadtree for the n disk centers, as in Section 5.1. We describe how to obtain the compressed quadtree $\mathcal{Q}_C$ from $\mathcal{Q}$. A node ν in $\mathcal{Q}$ is *empty* if $b(\nu)$ does not contain a disk-center, and *non-empty* otherwise. A

[2] By a packing argument, the spread of any d-dimensional n-point set is $\Omega(n^{1/d})$: if any two points have distance at least 1, the point set must cover at least $\Omega(n)$ units of volume and hence must have diameter $\Omega(n^{1/d})$.

singular path σ in $\mathcal{Q}$ is a path $\nu_1, \nu_2, \ldots, \nu_k$ of nodes such that (i) ν_k is a non-empty leaf or has at least two non-empty children; and (ii) for $i = 1, \ldots, k-1$, the node ν_{i+1} is the only non-empty child of ν_i. We call σ *maximal* if it cannot be extended by the parent of ν_1 (either because ν_1 is the root or because $p(\nu_1)$ has two non-empty children). For each maximal singular path $\sigma = \nu_1, \ldots, \nu_k$ in $\mathcal{Q}$, we remove from $\mathcal{Q}$ all proper descendants of ν_1 that are not descendants of ν_k, together with their incident edges. Then, we add a new *compressed edge* between ν_1 and ν_k. The resulting tree $\mathcal{Q}_C$ has $O(n)$ nodes. Each internal node has one or four children.[3] See Figure 4 (right) for an illustration. There are algorithms to compute $\mathcal{Q}_C$ in $O(n \log n)$ time and $O(n)$ space (see, e.g. Har-Peled's book [13, Theorem 2.9] or Buchin *et al.* for a version that does not need the floor function [4, Appendix A]).

We now describe how to map the candidate pairs from $\mathcal{Q}$ to candidate pairs in $\mathcal{Q}_C$. A node ν from $\mathcal{Q}$ may appear as a node in $\mathcal{Q}_C$ or not. We let $\pi(\nu)$ be the lowest ancestor node of ν (including ν) in $\mathcal{Q}$ that appears also in $\mathcal{Q}_C$. We call $\pi(\nu)$ the *upward projection* of ν in $\mathcal{Q}_C$. For a node ν that appears in $\mathcal{Q}_C$, we denote by $\Pi(\nu) = \{\nu' \in \mathcal{Q} \mid \pi(\nu') = \nu\}$ the set of all nodes in $\mathcal{Q}$ that project to ν. If ν is a leaf or has four children in $\mathcal{Q}_C$, then $\Pi(\nu) = \{\nu\}$. Otherwise, if ν has one child in $\mathcal{Q}_C$, then $\Pi(\nu)$ contains the nodes of the maximal singular path that starts in ν, without the last vertex. Furthermore, for $\nu \in \mathcal{Q}_C$, we write $p_C(\nu)$ the parent of ν in $\mathcal{Q}_C$ (we set $p_C(\nu) = \perp$, if ν is the root). We define the set of *compressed candidate pairs* $\mathsf{CNP}_C(\nu)$ for ν in $\mathcal{Q}_C$ as

$$\mathsf{CNP}_C(\nu) = \big\{(\pi(\nu'), \pi(\mu)) \mid \nu' \in \Pi(\nu) \text{ and } (\nu', \mu) \in \mathsf{CNP}(\nu') \text{ and } \pi(\nu') \neq \pi(\mu)\big\}. \qquad (***)$$

In other words, we obtain the compressed candidate pairs for a node ν in $\mathcal{Q}_C$ by taking the upward projections of all candidate pairs in $\mathcal{Q}$ where the first component projects to ν.

We can now describe our modified algorithm, Algorithm 3. Essentially, Algorithm 3 works in the same way as Algorithm 2: we go through the disks by decreasing order of priority, and for each disk D_i, we simulate the growth process of D_i by walking along the leaf-root path for the disk center p_i in the compressed quadtree $\mathcal{Q}_C$, while checking for elimination events with nearby disks of higher priority. There are two differences: first, instead of the candidate pairs $\mathsf{CNP}(\nu)$ of the current node ν, we now use the compressed candidate pairs $\mathsf{CNP}_C(\nu)$ to check for elimination events; second, the termination condition of the **while**-loop (line 7) is modified: instead of comparing t with the time $\tau(\nu, i)$ when D_i first covers $b(\nu)$, we use $\tau_C(\nu, i)$, the time when D_i first covers the box for any node in $\Pi(\nu)$. This ensures that we do not miss an elimination event on a singular path starting at ν. Next, we argue that Algorithm 3 correctly computes the elimination sequence. Then, we will discuss an efficient implementation. The correctness for Algorithm 3 follows from essentially the same argument as for Algorithm 2.

Lemma 5.4. *Algorithm 3 correctly computes the elimination sequence.*

Proof. We prove that Algorithm 3 maintains the following invariant: after i iterations of the **for**-loop, we have the correct elimination times $t_1, \ldots, t_i$ for $D_1, \ldots, D_i$. Furthermore, for each node ν of $\mathcal{Q}_C$, we have $D_C(\nu) \in \{1, \ldots, i, \perp\}$, and if there is a point in time when a node in $\Pi(\nu)$ is covered by the disk D_j, $1 \leq j \leq i$, then $D_C(\nu) = j$ (this is well defined, because the nodes along a singular path can be occupied by at most one disk).

The invariant holds for $i = 1$, because t is set to ∞ in Line 6 and for all the nodes ν in $\mathcal{Q}_C$, the value $D_C(\nu)$ is initialized to $\perp$, so that no elimination event will be detected in Line 10. The values $D_C(\nu)$ on the leaf-root path of p_1 are all set to 1, as desired. Now, suppose that $i \geq 2$ and that D_i is eliminated by the disk D_j, with $1 \leq j < i$. By Lemma 5.1, the pair $(\nu, \mu) = (\nu(i, t_i), \nu(j, t_i))$ is a candidate pair in $\mathsf{CNP}(\nu)$. Furthermore, ν lies on the leaf-root path for p_i in $\mathcal{Q}$. By (***), it follows that $(\pi(\nu), \pi(\mu))$ is a compressed candidate pair in $\mathsf{CNP}(\pi(\nu))$ and that $\pi(\nu)$ lies on the leaf-root path for p_i in $\mathcal{Q}_C$. By the inductive hypothesis, we have $D_C(\pi(\mu)) = j$. Again by the inductive hypothesis and by the test $t_{D_C(\mu)} \geq t(i, D_C(\mu))$ in Line 10, we have $t \geq t_i$ throughout the **while**-loop. Thus, the **while**-loop will visit $\pi(\nu)$ and consider the compressed candidate pair $(\pi(\nu), \pi(\mu))$. At this point, Algorithm 3 will detect the elimination event and set t to t_i. After that, the test $t \geq \tau_C(\nu, i)$ in Line 7 ensures that the remaining variables $D_C(\nu)$ are set correctly. In particular, the algorithm does not overwrite any other such values from previous iterations. Thus, the invariant is maintained and the correctness of Algorithm 3 follows. □

[3] According to our definition, the compressed quadtree may contain empty leaves, namely empty leaves that are children of nodes with at least two non-empty children. This empty leaves do not belong to any singular path.

Algorithm 3 Using the Compressed Quadtree

function EliminationOrder($p_1, \dots, p_n$, $v_1, \dots, v_n$)
  $\mathcal{Q}_C \leftarrow$ ConstructCompressedQuadTree($p_1, \dots, p_n$)
  $D_C(\nu) \leftarrow \perp$ for every node ν of $\mathcal{Q}$
  for $i \leftarrow 1, \dots, n$ **do**
    $\nu \leftarrow$ getCompressedLeaf(p_i)
    $t \leftarrow \infty$
    while $\nu \neq \perp$ and $t \geq \tau_C(\nu, i)$ **do**
      $D_C(\nu) \leftarrow i$
      for (ν, μ) in $\mathsf{CNP}_C(\nu)$ **do**
        if $D_C(\mu) \neq \perp$ and $t_{D_C(\mu)} \geq t(i, D_C(\mu))$ **then**
          $t \leftarrow \min\Big(t, t\big(i, D_C(\mu)\big)\Big)$
      $\nu \leftarrow p_C(\nu)$
    $t_i \leftarrow t$
  $S \leftarrow (D_1, \dots, D_n)$
  Sort S using key t_i for each disk D_i
  return S

Next, we discuss how to implement Algorithm 3 efficiently. First we must bound the number of compressed candidate pairs. For this, we need two closure properties for candidate pairs in $\mathcal{Q}$: if two nodes in $\mathcal{Q}$ form a candidate pair, then so do their parents; and if two nodes (ν, μ) form a candidate pair in which ν lies at a lower level than μ, then all the ancestors of ν up to the level of μ also form candidate pairs with μ. See Figure 6 for an illustration.

Lemma 5.5. *Let ν, μ be nodes of $\mathcal{Q}$ such that $(\nu, \mu) \in \mathsf{CNP}(\nu)$ and $p(\nu) \neq p(\mu)$. Then, we have*

(i) $(p(\nu), p(\mu)) \in \mathsf{CNP}(p(\nu))$, i.e., the parents of ν and μ also form a candidate pair; and

(ii) if $|\nu| \leq |\mu|$, then $(\nu', \mu) \in \mathsf{CNP}(\nu')$ for any ancestor ν' of ν with $|\nu'| \leq |\mu|$, i.e., all the ancestors of ν up to the level of μ also form candidate pairs with μ.

Proof. We must check properties (*) and (**) of a candidate pair. For (i), property (*) holds because

$$d(p(\nu), p(\mu)) \leq d(\nu, \mu) \leq 2(|\nu| + |\mu|) \leq 2(|p(\nu)| + p(|\mu|)),$$

where in the first and third inequality we used the fact that $b(\nu) \subset b(p(\nu))$ and $b(\mu) \subset b(p(\mu))$ and in the second inequality we used that (ν, μ) is a candidate pair. Property (**) holds because

$$\frac{|p(\mu)|}{|p(\nu)|} = \frac{|\mu|/2}{|\nu|/2} = \frac{|\mu|}{|\nu|} \in \left[\frac{1}{4\Delta}, 4\Delta\right],$$

since (ν, μ) is a candidate pair. The argument for (ii) is analogous. Since $b(\nu)$ is a subset of $b(\nu')$ and since (ν, μ) is a candidate pair, we have

$$d(\nu', \mu) \leq d(\nu, \mu) \leq 2(|\nu| + |\mu|) \leq 2(|\nu'| + |\mu|),$$

so property (*) holds. Furthermore, since $|\nu| \leq |\nu'| \leq |\mu|$ and since (ν, μ) is a candidate pair, we have

$$\frac{1}{4\Delta} \leq 1 \leq \frac{|\mu|}{|\nu'|} \leq \frac{|\mu|}{|\nu|} \leq 4\Delta,$$

which shows property (**). □

The next lemma provides a way to charge compressed candidate pairs in $\mathcal{Q}_C$ to candidate pairs in $\mathcal{Q}$. More precisely, we show that each compressed candidate pair (ν, μ) in $\mathcal{Q}_C$ can be obtained by taking a candidate pair $(\sigma, \tau) \in \mathsf{CNP}(\nu) \cup \mathsf{CNP}(\mu)$ with $|\sigma| \leq |\tau|$, and by projecting τ upwards; see Figure 7.

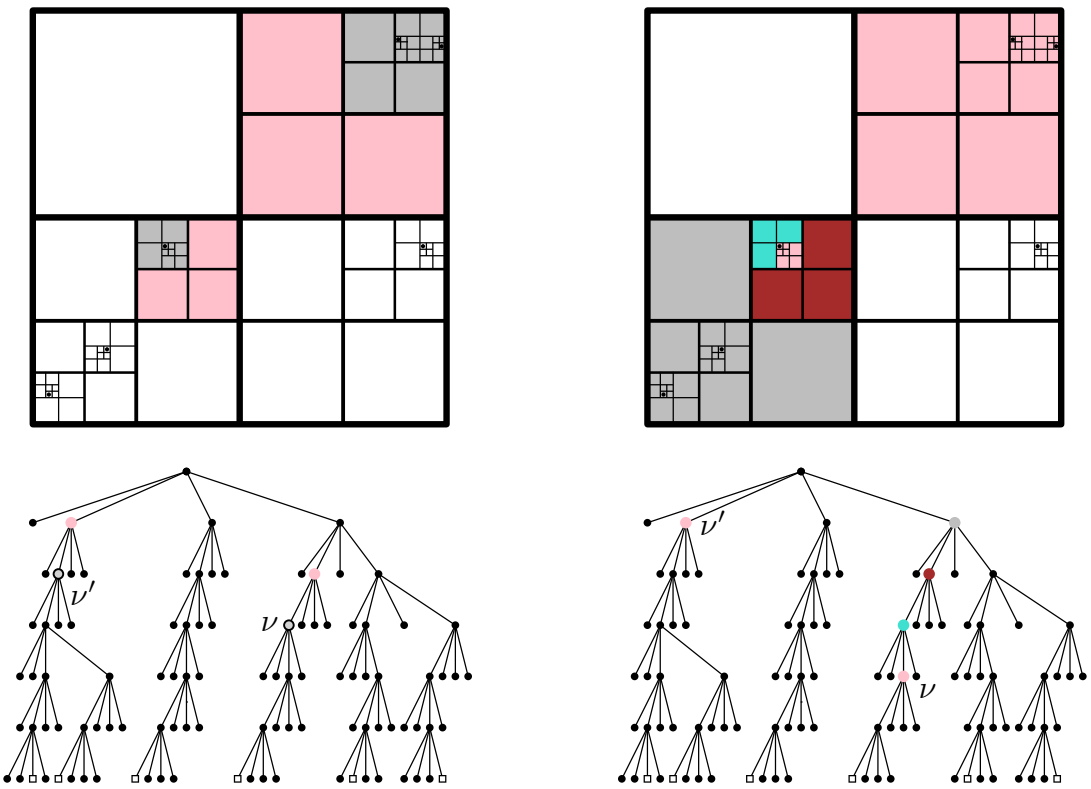

Figure 6: (left) If two nodes ν and ν' form a candidate pair, so do their parents. (right) If $|\nu| \leq |\nu'|$ and ν forms a candidate pair with ν', then so do all ancestors of ν up to the level of ν'.

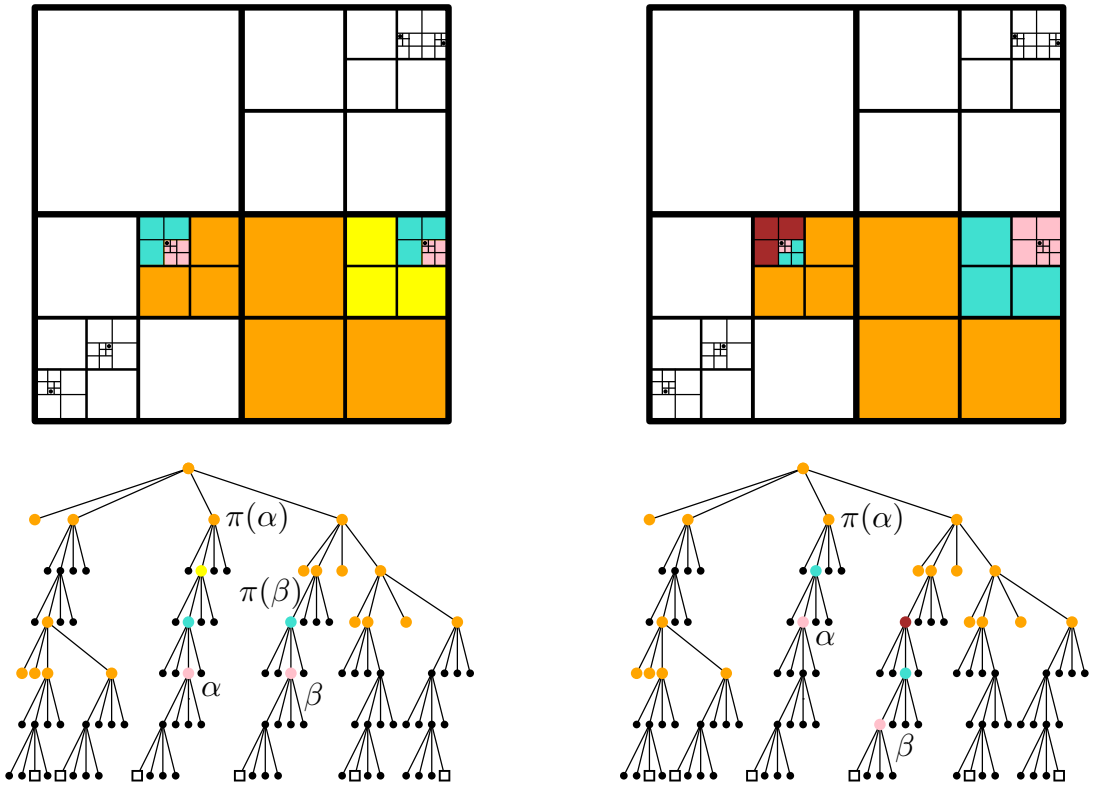

Figure 7: Compressed candidate pairs can be charged to regular candidate pairs in $\mathcal{Q}_C$ where the second node is larger: apply Lemma 5.5(i) until reaching a node in $\mathcal{Q}_C$; (left) If this is the smaller node, the projection of the other node stays the same; (right) if it is the larger node, we apply Lemma 5.5(ii).

Lemma 5.6. *Let ν, μ be two nodes in $\mathcal{Q}_C$ with $(\nu, \mu) \in \mathsf{CNP}_C(\nu)$. Then, there are two nodes σ, τ in $\mathcal{Q}$ such that (i) $(\sigma, \tau) \in \mathsf{CNP}(\sigma)$, i.e., σ and τ form a candidate pair in $\mathcal{Q}$; (ii) $|\sigma| \leq |\tau|$, i.e., σ is not larger than τ; and (iii) $(\nu, \mu) = (\sigma, \pi(\tau))$ or $(\nu, \mu) = (\pi(\tau), \sigma)$, i.e., the compressed candidate pair (ν, μ) is obtained by taking a candidate pair for ν or for μ and by projecting the other component upwards.*

Proof. Since $(\nu, \mu) \in \mathsf{CNP}_C(\nu)$, definition (***) implies that there are two nodes α, β in $\mathcal{Q}$ where (i) α and β form a candidate pair, i.e., $(\alpha, \beta) \in \mathsf{CNP}(\alpha)$; (ii) the notation is such that α is not larger than β, i.e., $|\alpha| \leq |\beta|$; and (iii) the upward projections of α and β (in the right order) give the compressed candidate pair (ν, μ), i.e., $\{\nu, \mu\} = \{\pi(\alpha), \pi(\beta)\}$. We repeatedly apply Lemma 5.5(i) to (α, β), until we meet $\pi(\alpha)$ or $\pi(\beta)$, whichever happens first. By assumption, $\pi(\alpha)$ and $\pi(\beta)$ are distinct, so all the parents along the way are also distinct, and Lemma 5.5(i) is applicable. See Figure 7 for an illustration.

Suppose we meet $\pi(\alpha)$ first. In this case, we set $\sigma = \pi(\alpha)$. Then, it follows that $(\sigma, \tau) \in \mathsf{CNP}(\sigma)$ for some ancestor τ of β in $\mathcal{Q}$ with $|\sigma| \leq |\tau|$. Since $\sigma = \pi(\alpha)$ is encountered first, the upward projections of τ and β in $\mathcal{Q}_C$ are the same, i.e., $\pi(\tau) = \pi(\beta)$. Hence, the pair (σ, τ) has all the desired properties.

Second, suppose we meet $\pi(\beta)$ first. Consider the ancestor α' of α that has the same size as $\pi(\beta)$. If $\pi(\alpha)$ appears on the path from α to α' in $\mathcal{Q}$, we set $\sigma = \pi(\alpha)$ and $\tau = \pi(\beta)$. Otherwise, we set $\sigma = \pi(\beta)$ and $\tau = \alpha'$. In either case, $|\sigma| \leq |\tau|$. Furthermore, by Lemma 5.5(ii), we have $(\sigma, \tau) \in \mathsf{CNP}(\sigma)$. Finally, the upward projections are maintained. Hence, (σ, τ) again has all the desired properties. □

Now, we can bound the number of compressed candidate pairs with a simple charging argument.

Lemma 5.7. *The total number of compressed candidate pairs is $O(n\alpha)$.*

Proof. Let (ν, μ) be a compressed candidate pair. We use Lemma 5.6 to charge (ν, μ) to a pair $(\sigma, \tau) \in \mathsf{CNP}(\nu) \cup \mathsf{CNP}(\mu)$ with $|\sigma| \leq |\tau|$ and $\{\pi(\sigma), \pi(\tau)\} = \{\nu, \mu\}$.

Now, let ν be a node of $\mathcal{Q}_C$. Every candidate pair $(\nu, \mu) \in \mathsf{CNP}(\nu)$ with $|\nu| \leq |\mu|$ is charged at most twice, namely (potentially) by $(\nu, \pi(\mu))$ and by $(\pi(\mu), \nu)$. By Lemma 5.2, there are $O(\alpha)$ candidate pairs $(\nu, \mu) \in \mathsf{CNP}(\nu)$ with $|\nu| \leq |\mu|$. Since $\mathcal{Q}_C$ has $O(n)$ nodes, the claim follows. □

Now, we have enough tools to find all the compressed candidates in Line 9 efficiently.

Lemma 5.8. *During Algorithm 3, we can enumerate all compressed candidates $\mathsf{CNP}_C(\nu)$ for the nodes ν visited by the algorithm in total time $O(n\alpha)$ and space $O(n)$.*

Proof. During preprocessing, we compute for each node ν in $\mathcal{Q}_C$ *neighbor pointers* to all nodes $\mu \in \mathcal{Q}_C$ where $\Pi(\mu)$ contains a node μ' with $|\nu| = |\mu'|$ and $d(\nu, \mu') \leq 4|\nu|$. For each ν, there are $O(1)$ such pointers (a simple volume argument), and they can be found in $O(n)$ time and space by a top-down traversal.

Now suppose we want to enumerate the compressed candidate pairs for a node ν in $\mathcal{Q}_C$. Let $(\nu, \mu) \in \mathsf{CNP}(\nu)$ be such a candidate. By Lemma 5.6, we have either (i) $(\nu, \mu) = (\nu, \pi(\mu'))$ for a candidate pair $(\nu, \mu') \in \mathsf{CNP}(\nu))$ with $|\nu| \leq |\mu'|$ or (ii) $(\nu, \mu) = (\pi(\nu'), \mu)$, where μ appears in $\mathcal{Q}_C$ and (μ, ν') is a candidate pair in $\mathcal{Q}$ with $|\mu| \leq |\nu'|$.

To find the compressed candidate pairs of type (i), we enumerate all (regular) candidate pairs for ν, using the neighbor pointers for ν and its ancestors in a similar procedure as in Theorem 5.3. By Lemma 5.2, this takes $O(\alpha)$ time and no additional space if we process the compressed candidate pairs immediately without storing them. For compressed candidate pairs of type (ii), we must enumerate all nodes $\mu \in \mathcal{Q}_C$ such that $\Pi(\nu)$ contains a node ν' with $(\mu, \nu') \in \mathsf{CNP}(\mu)$ and $|\mu| \leq |\nu'|$. The crucial observation is that by Lemma 5.5, these nodes form a connected subtree under each neighbor node of ν. Thus, we can find them by following the neighbor pointers for ν and by traversing each such subtree as long as a compressed candidate pair is found. We can check whether a node $\mu \in \mathcal{Q}_C$ forms a compressed candidate pair with ν in $O(1)$ by elementary calculations involving the floor function.[4] Thus, the time is proportional to the number of distinct compressed candidate pairs that are discovered. No additional space is necessary, because the compressed candidate pairs can be processed immediately. Since each node in $\mathcal{Q}_C$ is visited at most once by Algorithm 3, the result now follows from Lemma 5.7. □

[4] We can also do without the floor function if we slightly relax the notion of a compressed candidate pair.

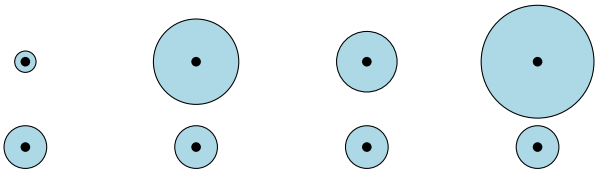

Figure 8: The lower bound reduction: the disk centers are arranged in two rows, where two consecutive disk centers have distance 2. The distance between two rows is 1 (actually, the points are slightly perturbed for general position; this is omitted in the figure). The disks centered at the bottom row grow uniformly with rate 1. The disks in the top row grow at varying speeds in $(1, 2)$. The elimination order of the disks in the top row corresponds to the reverse sorted order of their growth rates.

The following theorem summarizes our result for this section.

Theorem 5.9. *The elimination sequence of n disks can be computed in $O(n \log n + n\alpha)$ time and $O(n)$ space, where $\alpha = \min\{\log \Phi, \log \Delta\}$.*

Proof. By Lemma 5.4, Algorithm 3 correctly computes the elimination sequence. The compressed quadtree $\mathcal{Q}_C$ can be constructed in $O(n \log n)$ time and $O(n)$ space. This is also the time needed for the final sorting step. Since Algorithm 3 visits each node of $\mathcal{Q}_C$ at most once, and since $\mathcal{Q}_C$ has $O(n)$ nodes, the time for the **for**-loop (without the time for computing the compressed candidates) is $O(n\alpha)$, by Lemma 5.7. It uses no additional space. Finally, by Lemma 5.8, the additional time for enumerating the compressed candidate pairs is $O(n\alpha)$, using $O(n)$ space. The result follows. □

6 Lower bound

To complement our results, we provide a lower bound for finding elimination sequences in the algebraic decision tree model. Formann [9] argued that the weighted-closest pair problem has an $\Omega(n \log n)$ lower bound in the algebraic decision tree model, by a reduction from the closest-pair-problem. Since finding the elimination order is a more general problem, this also implies an $\Omega(n \log n)$ lower bound for our problem. Here, we provide a slightly stronger result by showing that the *sorting problem* reduces to finding elimination orders.

Theorem 6.1. *There is a reduction from the sorting problem to the elimination order problem. In particular, it takes at least $\Omega(n \log n)$ time to find the elimination order of a set of n growing disks or squares in the plane under the algebraic decision tree model.*

Proof. We show that the problem of sorting n numbers $v_{n+1}, \ldots, v_{2n} \in (1, 2)$ can be reduced to finding the elimination order of $2n$ disks in the plane. This implies an $\Omega(n \log n)$ lower bound in the algebraic decision tree model.

Our reduction proceeds as follows: Suppose we are given n numbers $v_{n+1}, \ldots, v_{2n} \in (1, 2)$, to be sorted. Set $\varepsilon = 1/10n^3$. We define $2n$ growing disks $D_1, \ldots, D_{2n}$ as follows: for $i = 1, \ldots, n$, we center the disk D_i at $p_i = \big(2i + i^2\varepsilon, 0\big)$ and give it the growth rate $v_i = 1$. For $i = n + 1, \ldots, 2n$, we position the disk D_i at $p_i = (2i + i^2\varepsilon, 1)$ with growth rate v_i as in the input; see Figure 8 for an illustration. Observe that disk D_{n+i} will be eliminated by disk D_i at time $t_{n+i} = t(n + i, i) = 1/(1 + v_{n+i}) < 1/2$, since $t_i > 1/2$ for $1 \le i \le n$. Then, the elimination order of $D_1, \ldots, D_{2n}$ lets us deduce the reverse sorted order of $\{v_{n+1}, \ldots, v_{2n}\}$. An analogous argument also applies to squares. □

We remark that Theorem 6.1 also shows that in general the problem does not become easier if we are interested only in the elimination order and not the exact elimination times.

7 Conclusion

We have presented the first truly subquadratic algorithm for the problem of computing the elimination order and elimination times of a sequence of n growing disks in the plane. Our approach is very general and also applies to other shapes. However, it still falls short of reaching a near-linear time algorithm, except for the special case of growing cubes. Thus, the most pressing question remains: can we compute the elimination order of n growing disks in the plane in $O(n \log n)$ time?

Our algorithm that uses compressed quadtrees comes close, but it depends on additional parameters of the input: if the growth rates vary wildly, or if the points are arranged unevenly, the running times may deteriorate. Perhaps a more careful handling of these inputs could enable us to avoid this dependence. It would also be interesting to see if and how the quadtree approach can be adapted to higher dimensions.

Finally, many further well-motivated variants of the problem are possible. For example, Castermans *et al.* [5] consider the setting where two touching disks are replaced by a new, common, disk, instead of one of them disappearing. It is a promising research direction to explore these variants and to see in how far our techniques are applicable or which new ideas are required.

Acknowledgments. This work was initiated during the 20th Korean Workshop on Computational Geometry. The authors would like to thank the other participants for motivating and insightful discussions. We would also like to thank the anonymous reviewers for their close reading of the paper and for many helpful comments that significantly improved the presentation of the paper.

References

[1] Pankaj K. Agarwal, Boris Aronov, and Micha Sharir. Computing envelopes in four dimensions with applications. *SIAM Journal on Computing*, 26(6):1714–1732, 1997.

[2] Daniel Bahrdt, Michael Becher, Stefan Funke, Filip Krumpe, André Nusser, Martin Seybold, and Sabine Storandt. Growing balls in $\mathbb{R}^d$. In *Proceedings of the 19th Workshop on Algorithm Engineering and Experiments (ALENEX)*, pages 247–258, 2017.

[3] Mark de Berg, Otfried Cheong, Marc van Kreveld, and Mark Overmars. *Computational Geometry: Algorithms and Applications*. Springer-Verlag, 2008.

[4] Kevin Buchin, Maarten Löffler, Pat Morin, and Wolfgang Mulzer. Preprocessing imprecise points for Delaunay triangulation: Simplified and extended. *Algorithmica*, 61(3):674–693, 2011.

[5] Thom Castermans, Bettina Speckmann, Frank Staals, and Kevin Verbeek. Agglomerative clustering of growing squares. In *Proceedings of the 13th Latin American Theoretical Informatics Symposium (LATIN)*, pages 260–274, 2018.

[6] Bernard Chazelle. *The Discrepancy Method. Randomness and complexity*. Cambridge University Press, 2000.

[7] Bernard Chazelle, Herbert Edelsbrunner, Michelangelo Grigni, Leonidas J. Guibas, John Hershberger, Micha Sharir, and Jack Snoeyink. Ray shooting in polygons using geodesic triangulations. *Algorithmica*, 12(1):54–68, 1994.

[8] Bernard Chazelle, Herbert Edelsbrunner, Leonidas J. Guibas, and Micha Sharir. A singly exponential stratification scheme for real semi-algebraic varieties and its applications. *Theoret. Comput. Sci.*, 84(1):77–105, 1991.

[9] Michael Formann. Weighted closest pairs. In *Proceedings of the 10th Symposium on Theoretical Aspects of Computer Science (STACS)*, pages 270–281, 1993.

[10] Stefan Funke, Filip Krumpe, and Sabine Storandt. Crushing disks efficiently. In *Proceedings of the 27th International Workshop on Combinatorial Algorithms (IWOCA)*, pages 43–54, 2016.

[11] Stefan Funke and Sabine Storandt. Parametrized runtimes for ball tournaments. In *Proceedings of the 33rd European Workshop on Computational Geometry (EuroCG)*, pages 221–224, 2017.

[12] Dan Halperin and Micha Sharir. Arrangements. In *Handbook of Discrete and Computational Geometry*, pages 723–762. CRC Press, third edition, 2017.

[13] Sariel Har-Peled. *Geometric Approximation Algorithms*. American Mathematical Society, Boston, MA, USA, 2011.

[14] John Hershberger. Finding the upper envelope of n line segments in $O(n \log n)$ time. *Information Processing Letters*, 33(4):169–174, 1989.

[15] Vladlen Koltun. Almost tight upper bounds for vertical decompositions in four dimensions. *Journal of the ACM*, 51(5):699–730, 2004.

[16] Jiří Matoušek. *Lectures on Discrete Geometry*. Springer-Verlag, 2002.

[17] Joseph S. B. Mitchell and Wolfgang Mulzer. Proximity algorithms. In *Handbook of Discrete and Computational Geometry*, pages 849–874. CRC Press, third edition, 2017.

[18] Franco P. Preparata and Roberto Tamassia. Efficient point location in a convex spatial cell-complex. *SIAM Journal on Computing*, 21(2):267–280, 1992.

[19] Micha Sharir and Pankaj K. Agarwal. *Davenport-Schinzel sequences and their geometric applications*. Cambridge University Press, 1995.