

Union of Random Minkowski Sums and Network Vulnerability Analysis*

Pankaj K. Agarwal[†] Sariel Har-Peled[‡] Haim Kaplan[§] Micha Sharir[¶]

June 29, 2014

Abstract

Let $\mathcal{C} = \{C_1, \dots, C_n\}$ be a set of n pairwise-disjoint convex sets of constant description complexity, and let π be a probability density function (density for short) over the non-negative reals. For each i , let K_i be the Minkowski sum of C_i with a disk of radius r_i , where each r_i is a random non-negative number drawn independently from the distribution determined by π . We show that the expected complexity of the union of K_1, \dots, K_n is $O(n^{1+\varepsilon})$ for any $\varepsilon > 0$; here the constant of proportionality depends on ε and the description complexity of the sets in \mathcal{C} , but not on π . If each C_i is a convex polygon with at most s vertices, then we show that the expected complexity of the union is $O(s^2 n \log n)$.

Our bounds hold in a more general model in which we are given an arbitrary multi-set $\Theta = \{\theta_1, \dots, \theta_n\}$ of expansion radii, each a non-negative real number. We assign them to the members of \mathcal{C} by a random permutation, where all permutations are equally likely to be chosen; the expectations are now with respect to these permutations.

We also present an application of our results to a problem that arises in analyzing the vulnerability of a network to a physical attack.

1 Introduction

Union of random Minkowski sums. Let $\mathcal{C} = \{C_1, \dots, C_n\}$ be a set of n pairwise-disjoint convex sets of constant description complexity, i.e., the boundary of each C_i is defined by a constant number

*A preliminary version of this paper appeared in *Proc. 29th Annual Symposium of Computational Geometry, 2013*, pp. 177–186. Work by Pankaj Agarwal and Micha Sharir has been supported by Grant 2012/229 from the U.S.-Israel Binational Science Foundation. Work by Pankaj Agarwal has also been supported by NSF under grants CCF-09-40671, CCF-10-12254, and CCF-11-61359, by ARO grant W911NF-08-1-0452, and by an ERDC contract W9132V-11-C-0003. Work by Sariel Har-Peled has been supported by NSF under grants CCF-0915984 and CCF-1217462. Work by Haim Kaplan has been supported by grant 822/10 from the Israel Science Foundation, grant 1161/2011 from the German-Israeli Science Foundation, and by the Israeli Centers for Research Excellence (I-CORE) program (center no. 4/11). Work by Micha Sharir has also been supported by NSF Grant CCF-08-30272, by Grants 338/09 and 892/13 from the Israel Science Foundation, by the Israeli Centers for Research Excellence (I-CORE) program (center no. 4/11), and by the Hermann Minkowski–MINERVA Center for Geometry at Tel Aviv University.

[†]Department of Computer Science, Box 90129, Duke University, Durham, NC 27708-0129, USA; pankaj@cs.duke.edu

[‡]Department of Computer Science; University of Illinois; 201 N. Goodwin Avenue; Urbana, IL, 61801, USA; sariel@illinois.edu

[§]School of Computer Science, Tel Aviv University, Tel Aviv 69978, Israel; haimk@post.tau.ac.il

[¶]School of Computer Science, Tel Aviv University, Tel Aviv 69978, Israel; michas@post.tau.ac.il

of algebraic arcs of constant maximum degree. Let $D(r)$ denote the disk of radius r centered at the origin. We consider the setup where we are given a sequence $\mathbf{r} = \langle r_1, \dots, r_n \rangle$ of non-negative numbers, called *expansion distances* (or *radii*). We set $K_i = C_i \oplus D(r_i)$, the Minkowski sum of C_i with $D(r_i)$. The boundary of K_i , denoted by ∂K_i , consists of $O(1)$ algebraic arcs of bounded degree. We refer to the endpoints of the arcs of ∂K_i as the *vertices* of K_i . If C_i is a convex polygon with s vertices, then ∂K_i is an alternating concatenation of line segments and circular arcs, where each segment is a parallel shift, by distance r_i , of an edge of C_i , and each circular arc is of radius r_i and is centered at a vertex of C_i ; see Figure 1. Let $\mathcal{K} = \{K_1, \dots, K_n\}$, and let $\mathcal{U} = \mathcal{U}(\mathcal{K}) = \bigcup_{i=1}^n K_i$. The *combinatorial complexity* of \mathcal{U} , denoted by $\psi(\mathcal{C}, \mathbf{r})$, is defined to be the number of vertices of \mathcal{U} , each of which is either a vertex of some K_i or an intersection point of the boundaries of a pair of K_i 's, lying on $\partial \mathcal{U}$. We do not make any assumptions on the shape and location of the sets in \mathcal{C} , except for requiring them to be pairwise disjoint.

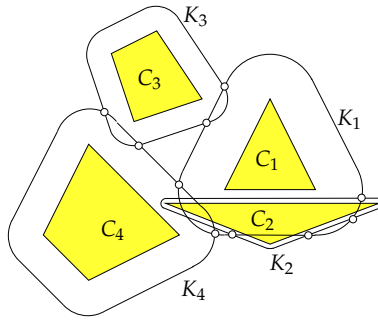


Figure 1: Pairwise-disjoint convex polygons and their Minkowski sums with disks of different radii. The vertices of the union of these sums are highlighted.

Our goal is to obtain an upper bound on the *expected* combinatorial complexity of \mathcal{U} , under a suitable probabilistic model for choosing the expansion radii \mathbf{r} of the members of \mathcal{C} —see below for the precise models that we will use.

Network vulnerability analysis. Our motivation for studying the above problem comes from the problem of analyzing the vulnerability of a network to a physical attack (e.g., electromagnetic pulse (EMP) attacks, military bombing, or natural disasters [13]), as studied in [2]. Specifically, let $\mathcal{G} = (V, \mathcal{E})$ be a planar graph embedded in the plane, where V is a set of points in the plane and $\mathcal{E} = \{e_1, \dots, e_n\}$ is a set of n segments (often called *links*) with pairwise-disjoint relative interiors, whose endpoints are points of V . For a point $q \in \mathbb{R}^2$ and an edge e , let $d(q, e) = \min_{p \in e} \|q - p\|$ denote the (minimum) distance between q and e . Let $\varphi : \mathbb{R}_{\geq 0} \rightarrow [0, 1]$ denote the *edge failure probability function*, so that the probability of an edge e to be damaged by a physical attack at a location q is $\varphi(d(q, e))$. In this model, the failure probability only depends on the distance of the point of attack from e . We assume that $1 - \varphi$ is a cumulative distribution function (cdf), or, equivalently, that $\varphi(0) = 1$, $\varphi(\infty) = 0$, and φ is monotonically decreasing. A typical example is $\varphi(x) = \max\{1 - x, 0\}$, where the cdf is the uniform distribution on $[0, 1]$.

For each $e_i \in \mathcal{E}$, let $f_i(q) = \varphi(d(q, e_i))$. The function $\Phi(q, \mathcal{E}) = \sum_{i=1}^n f_i(q)$ gives the expected number of links of \mathcal{E} damaged by a physical attack at a location q ; see Figures 2 and 3. Set

$$\Phi(\mathcal{E}) = \max_{q \in \mathbb{R}^2} \Phi(q, \mathcal{E}).$$

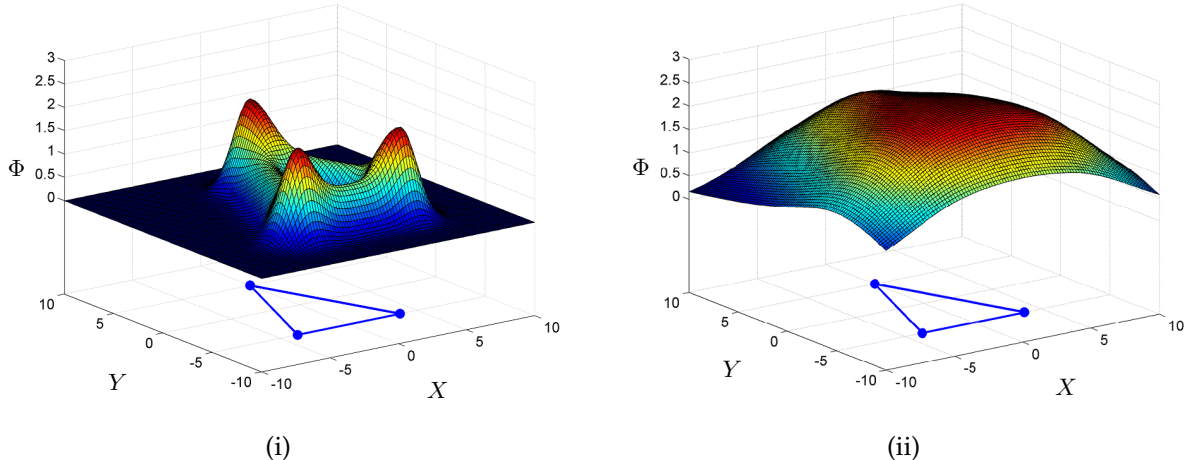


Figure 2: Expected damage for a triangle network and Gaussian probability distribution function with (i) small variance, (ii) large variance.

Our ideal goal is to compute $\Phi(\mathcal{E})$ and a location q^* such that $\Phi(q^*, \mathcal{E}) = \Phi(\mathcal{E})$. We refer to such a point q^* as a *most vulnerable* location for \mathcal{G} . As evident from Figure 3, the function Φ can be quite complex, and it is generally hard to compute $\Phi(\mathcal{E})$ exactly, so we focus on computing it approximately. More precisely, given an error parameter $\delta > 0$, we seek a point $\tilde{q} \in \mathbb{R}^2$ for which $\Phi(\tilde{q}, \mathcal{E}) \geq (1 - \delta)\Phi(\mathcal{E})$ (a so-called *approximately most vulnerable* location). Agarwal *et al.* [2] proposed a Monte Carlo algorithm for this task. As it turns out, the problem can be reduced to the problem of estimating the maximum depth in an arrangement of random Minkowski sums of the form considered above, under the density model (one of the models for the random choice of expansion distances that we use, which is defined later in the introduction), and its performance then depends on the expected complexity of $\mathcal{U}(\mathcal{K})$. Here \mathcal{K} is a collection of Minkowski sums of the form $e_i \oplus D(r_i)$, for a sample of edges $e_i \in \mathcal{E}$ and for suitable random choices of the r_i 's, from the distribution $1 - \varphi$. We adapt and simplify the algorithm in [2] and prove a better bound on its performance by using the sharp (near-linear) bound on the complexity of $\mathcal{U}(\mathcal{K})$ that we derive in this paper; see below and Section 5 for details.

Related work. (i) *Union of geometric objects.* There is extensive work on bounding the complexity of the union of a set of geometric objects, especially in \mathbb{R}^2 and \mathbb{R}^3 , and optimal or near-optimal bounds have been obtained for many interesting cases. We refer the reader to the survey paper by Agarwal *et al.* [5] for a comprehensive summary of most of the known results on this topic. For a set of n planar objects, each of *constant description complexity*, the complexity of their union can be $\Theta(n^2)$ in the worst case, but many linear or near-linear bounds are known for special restricted cases. For example, a classical result of Kedem *et al.* [18] asserts that the union of a set of pseudo-disks in \mathbb{R}^2 has linear complexity. It is also shown in [18] that the Minkowski sums of a set of pairwise-disjoint planar convex objects with a fixed common convex set is a family of pseudo-disks. Hence, in our setting, if all the r_i 's were equal, the result of [18] would then imply that the complexity of $\mathcal{U}(\mathcal{K})$ is $O(n)$. On the other hand, an adversarial choice of the r_i 's may result in a union \mathcal{U} with $\Theta(n^2)$ complexity; see Figure 4. Another observation is that if all the distances r_i are sufficiently large, the expanded sets will all be *fat* (see, e.g., [5] for details), in which case $\mathcal{U}(\mathcal{K})$ has near-linear complexity.

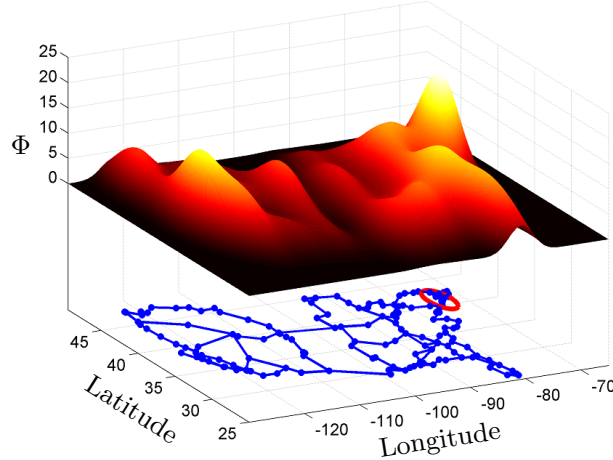


Figure 3: Expected damage for a complex fiber network. This figure is taken from [2].




Figure 4: A bad choice of expansion distances may cause \mathcal{U} to have quadratic complexity.

(ii) Network vulnerability analysis. Most of the early work on network vulnerability analysis considered a small number of isolated, independent failures; see, e.g., [9, 23] and the references therein. Since physical networks rely heavily on their physical infrastructure, they are vulnerable to physical attacks such as electromagnetic pulse (EMP) attacks as well as natural disasters [13, 27], not to mention military bombing and other similar kinds of attack. This has led to recent work on analyzing the vulnerability of a network under geographically correlated failures due to a physical attack at a single location [1, 2, 21, 22, 27]. Most papers on this topic have studied a deterministic model for the damage caused by such an attack, which assumes that a physical attack at a location x causes the failure of all links that intersect some simple geometric region (e.g., a vertical segment of unit length, a unit square, or a unit disk) centered at x . The impact of an attack is measured in terms of its effect on the connectivity of the network, (e.g., how many links fail, how many pairs of nodes get disconnected, etc.), and the goal is to find the location of attack that causes the maximum damage to the network. In the simpler model studied in [2] and in the present paper, the damage is measured by the number of failed links. This kind of problem is something that both attackers and planners of such networks would like to solve. The former for obvious reasons, and the latter for identifying the most vulnerable portions of the network, in order to protect them better.

In practice, though, it is hard to be certain in advance whether a link will fail by a nearby physical attack. To address this situation, Agarwal *et al.* [2] introduced the simple probabilistic framework for modeling the vulnerability of a network under a physical attack, as described above.

One of the problems that they studied is to compute the largest expected number of links damaged by a physical attack. They described an approximation algorithm for this problem whose expected running time is quadratic in the worst case. A major motivation for the present study is to improve the efficiency of this algorithm and to somewhat simplify it at the same time.

Finally, we note that the study in this paper has potential applications in other contexts, where one wishes to analyze the combinatorial and topological structure of the Minkowski sums (or rather convolutions) of a set of geometric objects (or a function over the ambient space) with some *kernel function* (most notably a Gaussian kernel), or to perform certain computations on the resulting configuration. Problems of this kind arise in many applications, including statistical learning, computer vision, and robotics; see, e.g., [12, 19] and references therein.

Our models. We consider two probabilistic models for choosing the sequence $\mathbf{r} = \langle r_1, \dots, r_n \rangle$ of expansion distances:

The density model. We are given an arbitrary density (or a probability mass function) π over the non-negative reals; for each $1 \leq i \leq n$, we take r_i to be a random value drawn independently from the distribution determined by π .

The permutation model. We are given a multi-set $\Theta = \{\theta_1, \dots, \theta_n\}$ of n arbitrary non-negative real numbers. We draw a random permutation σ on $[1 : n]$, where all permutations are equally likely to be chosen, and assign $r_i := \theta_{\sigma(i)}$ to C_i for each $i = 1, \dots, n$.

Our goal is to prove sharp bounds on the expected complexity of the union $\mathcal{U}(\mathcal{K})$ under these two models. More precisely, for the density model, let $\psi(\mathcal{C}, \pi)$ denote the expected value of $\psi(\mathcal{C}, \mathbf{r})$ (recall that $\psi(\mathcal{C}, \mathbf{r})$ is the complexity of $\mathcal{U}(\mathcal{K})$ under a specific assignment sequence \mathbf{r} of expansion distances), where the expectation is taken over the random choices of \mathbf{r} , made from π in the manner specified above. Set $\psi(\mathcal{C}) = \max \psi(\mathcal{C}, \pi)$, where the maximum is taken over all probability density (mass) functions. For the permutation model, in an analogous manner, we let $\psi(\mathcal{C}, \Theta)$ denote the expected value of $\psi(\mathcal{C}, \mathbf{r})$, where the expectation is taken over the choices of \mathbf{r} , obtained by randomly shuffling the members of Θ . Then, with a slight overloading of the notation, we define $\psi(\mathcal{C}) = \max \psi(\mathcal{C}, \Theta)$, where the maximum is over all possible choices of the multi-set Θ . We wish to obtain an upper bound on $\psi(\mathcal{C})$ under both models.

We note that the permutation model is more general than the density model, in the sense that an upper bound on $\psi(\mathcal{C})$ under the permutation model immediately implies the same bound on $\psi(\mathcal{C})$ under the density model. Indeed, consider some given density π , out of whose distribution the distances r_i are to be sampled (in the density model). Interpret such a random sample as a 2-stage process, where we first sample from the distribution of π a multi-set Θ of n such distances, in the standard manner of independent repeated draws, and then assign the elements of Θ to the sets C_i using a random permutation (it is easily checked that this reinterpretation does not change the probability space). Let \mathbf{r} be the resulting sequence of expansion radii for the members of \mathcal{C} . Using the new interpretation, the expectation of $\psi(\mathcal{C}, \mathbf{r})$ (under the density model), conditioned on the fixed Θ , is at most $\psi(\mathcal{C})$ under the permutation model. Since this bound holds for every Θ , the unconditional expected value of $\psi(\mathcal{C}, \mathbf{r})$ (under the density model) is also at most $\psi(\mathcal{C})$ under the permutation model. Since this holds for every *density*, the claim follows.

We do not know whether the opposite inequality also holds. A natural reduction from the permutation model to the density model would be to take the input set Θ of the n expansion distances and regard it as a discrete mass distribution (where each of its members can be picked

with probability $1/n$). But then, since the draws made in the density model are independent, most of the draws will not be permutations of Θ , so this approach will not turn $\psi(\mathcal{C})$ under the density model into an upper bound for $\psi(\mathcal{C})$ under the permutation model.

Our results. The main results of this paper are near-linear upper bounds on $\psi(\mathcal{C})$ under the two models discussed above. Since the permutation model is more general, in the sense made above, we state, and prove, our results in this model only. We obtain (in Section 2) the following bound for the general case.

Theorem 1.1. *Let $\mathcal{C} = \{C_1, \dots, C_n\}$ be a set of n pairwise-disjoint convex sets of constant description complexity in \mathbb{R}^2 . Then the value of $\psi(\mathcal{C})$ under the permutation model, for any multi-set Θ of expansion radii, is $O(n^{1+\varepsilon})$, for any $\varepsilon > 0$; the constant of proportionality depends on ε and the description complexity of the members of \mathcal{C} , but not on Θ .*

If \mathcal{C} is a collection of convex polygons, we obtain a slightly improved bound by using a different, more geometric argument.

Theorem 1.2. *Let $\mathcal{C} = \{C_1, \dots, C_n\}$ be a set of n pairwise-disjoint convex polygons in \mathbb{R}^2 , where each C_i has at most s vertices. Then the maximum value of $\psi(\mathcal{C})$ under the permutation model, for any multi-set Θ of expansion radii, is $O(s^2 n \log n)$, with an absolute constant of proportionality that does not depend on Θ .*

For simplicity, we first prove Theorem 1.2 in Section 3 for the special case where \mathcal{C} is a set of n segments with pairwise-disjoint relative interiors. Then we extend the proof, in a reasonably straightforward manner, to polygons in Section 4. The version involving segments admits a somewhat cleaner proof, and is sufficient for the application to network vulnerability analysis.

Using the Clarkson-Shor argument [11], we also obtain the following corollary, which will be needed for the analysis in Section 5.

Corollary 1.3. *Let $\mathcal{C} = \{C_1, \dots, C_n\}$ be a set of n pairwise-disjoint convex sets of constant description complexity. Let r_1, \dots, r_n be the random expansion distances, obtained under the permutation model, for any multi-set Θ of expansion radii, that are assigned to C_1, \dots, C_n , respectively, and set $\mathcal{K} = \{C_i \oplus D(r_i) \mid 1 \leq i \leq n\}$. Then, for any $1 \leq k \leq n$, the expected number of vertices in the arrangement $\mathcal{A}(\mathcal{K})$ whose depth is at most k is $O(n^{1+\varepsilon} k^{1-\varepsilon})$, for any $\varepsilon > 0$; the constant of proportionality depends on ε and the description complexity of the members of \mathcal{C} but not on Θ . If each C_i is a convex polygon with at most s vertices, then the bound improves to $O(s^2 n k \log(n/k))$, with an absolute, Θ -independent constant of proportionality.*

We note that the variants of Theorem 1.1, Theorem 1.2, and Corollary 1.3 under the density model have identical statements, except that now the constants of proportionality are independent of the density function π .

Using Theorem 1.2 and Corollary 1.3, we present (in Section 5) an efficient Monte-Carlo δ -approximation algorithm for computing an approximately most vulnerable location for a network, as defined earlier. Our algorithm is a somewhat simpler, and considerably more efficient, variant of the algorithm proposed by Agarwal et al. [2], and the general approach is similar to the approximation algorithms presented in [3, 4, 6] for computing the depth in an arrangement of a set of objects. Specifically, we establish the following result.

Theorem 1.4. *Given a set \mathcal{E} of n segments in \mathbb{R}^2 with pairwise-disjoint relative interiors, an edge-failure-probability function φ such that $1 - \varphi$ is a cdf, and a constant $0 < \delta < 1$, one can compute, in $O(\delta^{-4} n \log^3 n)$ time, a location $\tilde{q} \in \mathbb{R}^2$, such that $\Phi(\tilde{q}, \mathcal{E}) \geq (1 - \delta)\Phi(\mathcal{E})$ with probability at least $1 - 1/n^c$, for arbitrarily large c ; the constant of proportionality in the running-time bound depends on c .*

2 The Case of Convex Sets

In this section we prove Theorem 1.1. We have a collection $\mathcal{C} = \{C_1, \dots, C_n\}$ of n pairwise-disjoint compact convex sets in the plane, each of constant description complexity. Let Θ be a multi-set of n non-negative real numbers $0 \leq \theta_1 \leq \theta_2 \leq \dots \leq \theta_n$. We choose a random permutation σ of $[1 : n]$, where all permutations are equally likely to be chosen, put $r_i = \theta_{\sigma(i)}$ for $i = 1, \dots, n$, and form the Minkowski sums $K_i = C_i \oplus D(r_i)$, for $i = 1, \dots, n$. We put $\mathcal{K} = \{K_1, \dots, K_n\}$. We prove a near-linear upper bound on the expected complexity of $\mathcal{U}(\mathcal{K})$, as follows.

Fix a parameter t , whose value will be determined later, and put $\rho = \theta_{n-t}$, the $(t+1)$ -st largest distance in Θ . Put $\mathcal{C}^+ = \{C_i \in \mathcal{C} \mid \sigma(i) > n-t\}$ and $\mathcal{C}^- = \{C_i \in \mathcal{C} \mid \sigma(i) \leq n-t\}$. That is, \mathcal{C}^+ is the set of the t members of \mathcal{C} that were assigned the t largest distances in Θ , and \mathcal{C}^- is the complementary subset.

By construction, \mathcal{C}^+ is a random subset of \mathcal{C} of size t (where all t -element subsets of \mathcal{C} are equally likely to arise as \mathcal{C}^+). Moreover, conditioned on the choice of \mathcal{C}^+ , the set \mathcal{C}^- is fixed, and the subset Θ^- of the $n-t$ distances in Θ that are assigned to them is also fixed. Furthermore, the permutation that assigns the elements of Θ^- to the sets in \mathcal{C}^- is a random permutation.

For each $C_i \in \mathcal{C}^+$, put $K_i^* = C_i \oplus D(\rho)$. Put $(\mathcal{K}^*)^+ = \{K_i^* \mid C_i \in \mathcal{C}^+\}$, and let \mathcal{U}^* denote the union of $(\mathcal{K}^*)^+$. Note that $\mathcal{U}^* \subseteq \mathcal{U}$, because $K_i^* \subseteq K_i$ for each $C_i \in \mathcal{C}^+$. Since the C_i 's are pairwise-disjoint and we now add the same disk to each of them, $(\mathcal{K}^*)^+$ is a collection of pseudo-disks [18], and therefore \mathcal{U}^* has $O(t)$ complexity. More precisely, $\partial\mathcal{U}^*$ has at most $6t - 12$ vertices that are intersection points of boundaries of pairs of sets of $(\mathcal{K}^*)^+$ [18].

Let \mathcal{V} denote the *vertical decomposition* of the complement of \mathcal{U}^* ; it consists of $O(t)$ *pseudo-trapezoids*, each defined by at most four elements of $(\mathcal{K}^*)^+$ (that is, of \mathcal{C}^+). See [25] for more details concerning vertical decompositions. For short, we refer to these pseudo-trapezoids as *trapezoids*.

In a similar manner, for each $C_i \in \mathcal{C}^-$, define $K_i^* = C_i \oplus D(\rho)$; note that $K_i \subseteq K_i^*$ for each such i . Put $(\mathcal{K}^*)^- = \{K_i^* \mid C_i \in \mathcal{C}^-\}$. Since \mathcal{C}^+ is a random sample of \mathcal{C} of size t , the following lemma follows from a standard random-sampling argument; see [20, Section 4.6] for a proof.

Lemma 2.1. *With probability at least $1 - O\left(\frac{1}{n^{c-4}}\right)$, every (open) trapezoid τ of \mathcal{V} intersects at most $k := \frac{cn}{t} \ln n$ of the sets of $(\mathcal{K}^*)^-$, for sufficiently large $c > 4$.*

For each trapezoid τ of \mathcal{V} , let \mathcal{C}_τ^- denote the collection of the sets $C_i \in \mathcal{C}^-$ for which K_i^* intersects τ . We form the union \mathcal{U}_τ^- of the “real” (and smaller) corresponding sets K_i , for $C_i \in \mathcal{C}_\tau^-$, and clip it to within τ , i.e., keep only its portion with τ (clearly, no other set $C_i \in \mathcal{C}^-$ can have its real expansion K_i meet τ). Finally, we take all the “larger” sets $C_i \in \mathcal{C}^+$, and form the union of \mathcal{U}_τ^- with the corresponding “real” K_i 's, again clipping it to within τ . The overall union \mathcal{U} is the union of \mathcal{U}^* and of all the modified unions \mathcal{U}_τ^- , for $\tau \in \mathcal{V}$.

This divide-and-conquer process leads to the following recursive estimation of the expected complexity of \mathcal{U} . For $m \leq n$, let \mathcal{C}' be any subset of m sets of \mathcal{C} , and let Θ' be any subset of m elements of Θ , which we enumerate, with some abuse of notation, as $\theta_1, \dots, \theta_m$. Let $T(\mathcal{C}', \Theta')$ denote the expected complexity of the union of the expanded regions $C_i \oplus D(\theta_{\sigma'(i)})$, for $C_i \in \mathcal{C}'$, where the expectation is over the random shuffling permutation σ' (on $[1 : m]$). Let $T(m)$ denote the maximum value of $T(\mathcal{C}', \Theta')$, over all subsets \mathcal{C}' and Θ' of size m each, as just defined.

Let us first condition the analysis on a fixed choice of \mathcal{C}^+ . That is, we condition the analysis on the subset of those permutations σ that assign the t largest distances in Θ to the sets of \mathcal{C}^+ . This determines \mathcal{U}^* and \mathcal{V} uniquely. Hence we have a fixed set of trapezoids, and for each trapezoid τ

we have a fixed set \mathcal{C}_τ^- of $k_\tau = |\mathcal{C}_\tau^-|$ sets, whose expansions by ρ meet τ . Moreover, choosing $c = 6$ in Lemma 2.1, with probability $1 - O\left(\frac{1}{n^2}\right)$ our \mathcal{C}^+ is such that $k_\tau \leq \frac{cn}{t} \ln n$ for each $\tau \in \mathcal{V}$. Assume for now that \mathcal{C}^+ does indeed have this property. The set Θ_τ^- of distances assigned to these sets is not fixed, but it is a *random subset* of $\{\theta_1, \dots, \theta_{n-t}\}$ of size $k_\tau \leq \frac{cn}{t} \ln n$, where k_τ depends only on τ . Moreover, the assignment (under the original random permutation σ) of these distances to the sets in \mathcal{C}_τ^- is a *random permutation*. Hence, conditioning further on the choice of Θ_τ^- , the expected complexity of \mathcal{U}_τ^- , before its modification by the expansions of the larger sets of \mathcal{C}^+ , and ignoring its clipping to within τ , is

$$T(\mathcal{C}_\tau^-, \Theta_\tau^-) \leq T(k_\tau) \leq T\left(\frac{cn}{t} \ln n\right).$$

Hence the last expression also bounds the unconditional expected complexity of the unmodified and unclipped \mathcal{U}_τ^- (albeit still conditioned on the choice of \mathcal{C}^+). Summing this over all $O(t)$ trapezoids τ of \mathcal{V} , we get a bound of at most

$$atT\left(\frac{cn}{t} \ln n\right),$$

for a suitable absolute constant a . To this we need to add the number of intersections between the boundaries of the unions \mathcal{U}_τ^- with the respective trapezoid boundaries, and the number of intersections between the boundaries of the t larger expansions and the boundaries of all the expansions, that appear on the boundary of \mathcal{U} . The last two quantities are clearly both at most $O(nt)$ (the constant in this latter expression depends on the description complexity of the sets in \mathcal{C}). Altogether, we obtain that, with probability $1 - O\left(\frac{1}{n^2}\right)$, the complexity of $\mathcal{U}(\mathcal{K})$ is at most (reusing the constant a for simplicity)

$$atT\left(\frac{cn}{t} \ln n\right) + ant,$$

which holds when n is sufficiently large. With probability $O\left(\frac{1}{n^2}\right)$, \mathcal{C}^+ may fail to satisfy the uniform bound $k_\tau \leq \frac{6n}{t} \ln n$, but, since the complexity of $\mathcal{U}(\mathcal{K})$ is always $O(n^2)$, the contribution of these “bad” choices of \mathcal{C}^+ to the overall expectation is $O(1)$. In other words, when n is at least some sufficiently large constant n_0 , we get the recurrence

$$T(n) \leq atT\left(\frac{cn}{t} \ln n\right) + ant + a, \tag{1}$$

and when n is smaller than n_0 , we use the trivial bound $T(n) = O(n^2)$.

The solution of this recurrence is $T(n) \leq An^{1+\varepsilon}$, for any $\varepsilon > 0$, with an appropriate choice of parameters, where A depends on ε and on the other constants appearing in the recurrence. Indeed, fix the value of ε , and set $t = n^{\varepsilon/2}$. Substituting $T(n') \leq A(n')^{1+\varepsilon}$ in the right-hand side of (1), for $n' = \frac{cn}{t} \ln n < n$, an assumption that holds for $n \geq n_0$ and n_0 sufficiently large, we get that

$$\begin{aligned} T(n) &\leq an^{\varepsilon/2} A \left(cn^{1-\varepsilon/2} \ln n \right)^{1+\varepsilon} + an^{1+\varepsilon/2} + a \\ &= A \left(ac^{1+\varepsilon} \right) \cdot n^{1+\varepsilon-\varepsilon^2/2} (\ln n)^{1+\varepsilon} + an^{1+\varepsilon/2} + a, \end{aligned}$$

which is smaller than $An^{1+\varepsilon}$ for sufficiently large choices of A and n_0 (which depend on ε). Specifically, n_0 should be such that $n^{\varepsilon^2/2} > 2ac^{1+\varepsilon} (\ln n)^{1+\varepsilon}$ for all $n \geq n_0$. For this choice, a sufficiently large choice of A ensures that the right-hand side is indeed at most $An^{1+\varepsilon}$. Finally, by choosing A sufficiently large we can also ensure that the bound $O(n^2)$, for any $n \leq n_0$, is at most $A_n^{1+\varepsilon}$.

Remark. It might be possible to improve the bound on the expected complexity of $\mathcal{U}(\mathcal{K})$ to $O(n \text{ polylog}(n))$ by replacing the above divide-and-conquer approach with a randomized incremental construction, as in [15].

3 The Case of Segments

Let $\mathcal{E} = \{e_1, \dots, e_n\}$ be a collection of n line segments in the plane with pairwise-disjoint relative interiors, and as in Section 2, let Θ be a multi-set of n non-negative real numbers $0 \leq \theta_1 \leq \theta_2 \leq \dots \leq \theta_n$. For simplicity, we assume that the segments in \mathcal{E} are in general position, i.e., no segment is vertical, no two of them share an endpoint, and no two are parallel. Moreover, we assume that the expansion distances in Θ are positive and in “general position” with respect to \mathcal{E} , so as to ensure that, no matter which permutation we draw, the expanded sets of \mathcal{K} are also in general position—no pair of them are tangent and no three have a common boundary point.¹ Using the standard symbolic perturbation techniques (see, e.g., [25, Chapter 7]), the proof can be extended when \mathcal{E} or Θ is not in general position or when some of the expansion distances are 0; we omit here the routine details.

For each $1 \leq i \leq n$, let a_i, b_i be the left and right endpoints, respectively, of e_i (as mentioned, we assume, with no loss of generality, that no segment in \mathcal{E} is vertical). We draw a random permutation σ of $\{1, \dots, n\}$, and, for each $1 \leq i \leq n$, we put $r_i = \theta_{\sigma(i)}$. We then form the Minkowski sums $K_i = e_i \oplus D(r_i)$, for $i = 1, \dots, n$. We refer to such a K_i as a *racetrack*. Its boundary consists of two semicircles γ_i^- and γ_i^+ , centered at the respective endpoints a_i and b_i of e_i , and of two parallel copies, e_i^- and e_i^+ , of e_i ; we use e_i^- (resp., e_i^+) to denote the straight edge of K_i lying below (resp., above) e_i . Let a_i^-, a_i^+ (resp., b_i^-, b_i^+) denote the left (resp., right) endpoints of e_i^-, e_i^+ , respectively. We regard K_i as the union of two disks D_i^-, D_i^+ of radius r_i centered at the respective endpoints a_i, b_i of e_i , and a rectangle R_i of width $2r_i$ having e_i as a midline. The left endpoint a_i splits the edge $a_i^- a_i^+$ of R_i into two segments of equal length, and similarly b_i splits the edge $b_i^- b_i^+$ of R_i into two segments of equal lengths. We refer to these four segments $a_i a_i^-, a_i a_i^+, b_i b_i^-, b_i b_i^+$ as the *portals* of R_i ; see Figure 5(i). Set $\mathcal{K} = \{K_i \mid 1 \leq i \leq n\}$, $\mathcal{D} = \{D_i^+, D_i^- \mid 1 \leq i \leq n\}$, and $\mathcal{R} = \{R_i \mid 1 \leq i \leq n\}$.

As above, let $\mathcal{U} = \mathcal{U}(\mathcal{K})$ denote the union of \mathcal{K} . We show that the expected number of vertices on $\partial\mathcal{U}$ is $O(n \log n)$, where the expectation is over the choice of the random permutation σ that produces the distances r_1, \dots, r_n .

We classify the vertices of $\partial\mathcal{U}$ into three types (see Figure 5(ii)):

- (i) *CC-vertices*, which lie on two semicircular arcs of the respective pair of racetrack boundaries;
- (ii) *RR-vertices*, which lie on two straight-line edges; and
- (iii) *CR-vertices*, which lie on a semicircular arc and on a straight-line edge.

Bounding the number of CC-vertices is trivial because they are also vertices of $\mathcal{U}(\mathcal{D})$, the union of the $2n$ disks D_i^-, D_i^+ , so their number is $O(n)$ [5, 18]. We therefore focus on bounding the expected number of RR- and CR-vertices of $\partial\mathcal{U}$.

¹When we carry over the analysis to the density model, the latter assumption will hold with probability 1 when π is Lebesgue-continuous, but may fail for a discrete probability mass distribution. In the latter situation, we can use symbolic perturbations to turn π into a density in general position, without affecting the asymptotic bound that we are after.




Figure 5: (i) Segment e_i , racetrack K_i , and its constituents rectangle R_i and disks D_i^- , D_i^+ . (ii) Union of two racetracks. α, β are RR-vertices, ζ is a CR-vertex, and ξ is a CC-vertex; α is a non-terminal vertex and β is a terminal vertex (because of the edge of R_j it lies on, which ends inside K_i).

3.1 RR-vertices

Let v be an RR-vertex of \mathcal{U} , lying on ∂R_i and ∂R_j , the rectangles of two respective segments e_i and e_j . Denote the edges of R_i and R_j containing v as $e_i^\pm \in \{e_i^-, e_i^+\}$ and $e_j^\pm \in \{e_j^-, e_j^+\}$, respectively. A vertex v is *terminal* if either a subsegment of e_i^\pm connecting v to one of the endpoints of e_i^\pm is fully contained in K_j , or a subsegment of e_j^\pm connecting v to one of the endpoints of e_j^\pm is fully contained in K_i ; otherwise v is a *non-terminal* vertex. For example, in Figure 5(ii), β is a terminal vertex, and α is a non-terminal vertex. There are at most $4n$ terminal vertices on $\partial\mathcal{U}$, so it suffices to bound the expected number of non-terminal vertices.

Our strategy is first to describe a scheme that charges each non-terminal RR-vertex v to one of the portals of one of the two rectangles of \mathcal{R} on whose boundary v lies, and then to prove that the expected number of vertices charged to each portal is $O(\log n)$. The bound $O(n \log n)$ on the expected number of (non-terminal) RR-vertices then follows.

Let v be a non-terminal RR-vertex lying on $\partial R_i \cap \partial R_j$, for two respective input segments e_i and e_j . To simplify the notation, we rename e_i and e_j as e_1 and e_2 . Note that v is an intersection of e_1^+ or e_1^- with e_2^+ or e_2^- . Since the analysis is the same for each of these four choices we will say, using the notations already introduced, that v is the intersection of e_1^\pm with e_2^\pm where e_1^\pm (resp., e_2^\pm) is either e_1^+ (resp., e_2^+) or e_1^- (resp., e_2^-).

For $i = 1, 2$, let g_i^\pm denote the (unique) portion of e_i^\pm between v and an endpoint w_i so that, locally near v , g_i^\pm is contained in the second racetrack K_{3-i} (that is, in R_{3-i}). Now take Q_i to be the rectangle with $g_i^\pm = vw_i$ as one of its sides, and with the orthogonal projection of g_i^\pm onto e_i as the opposite parallel side. We denote this edge of Q_i , which is part of e_i , by E_i . We let A_i be the side perpendicular to E_i and incident to w_i . Note that A_i is one of the four portals of R_i . We denote by A_i^* the side of Q_i that is parallel to A_i (and incident to v). See Figure 6.

Each rectangle Q_i , $i = 1, 2$, has a complementary rectangle Q_i' on the same side of e_i , which is openly disjoint from Q_i , so that $Q_i \cup Q_i'$ is the half of the rectangular portion R_i , between e_i and e_i^\pm , of the full racetrack K_i . Again, see Figure 6.

Since $E_i \subset e_i$ for $i = 1, 2$, it follows that E_1 and E_2 are disjoint. Since v is a non-terminal RR-vertex, w_1 lies outside K_2 and w_2 lies outside K_1 . So, as we walk along the edge vw_i of Q_i from v to w_i , we enter, locally near v , into the other racetrack K_{3-i} , but then we have to exit it before we get to w_i . Note that either of these walks, say from v to w_1 , may enter Q_2 or keep away from Q_2 (and enter instead the complementary rectangle Q_2'); see Figures 6 and 8(i) for the former situation, and Figure 8(ii) for the latter one.




Figure 6: The rectangles Q_1 and Q_2 defined for a non-terminal RR-vertex v .

Another important property of the rectangles Q_1 and Q_2 , which follows from their definition, is that they are “oppositely oriented”, when viewed from v , in the following sense. When we view from v one of Q_1 and Q_2 , say Q_1 , and turn counterclockwise, we first see E_1 and then A_1 , and when we view Q_2 and turn counterclockwise, we first see A_2 and then E_2 .

So far our choice of which among the segments defining v is denoted by e_1 and which is denoted by e_2 was arbitrary. But in the rest of our analysis we will use e_1 to denote the segment such that when we view Q_1 from v and turn counterclockwise, we first see E_1 and then A_1 . The other segment is denoted by e_2 . (This is indeed the situation depicted in Figure 6.)

The following lemma provides the key ingredient for our charging scheme.

Lemma 3.1. *Let v be a non-terminal RR-vertex. Then, in the terminology defined above, one of the edges, say E_1 , has to intersect either the portal A_2 of the other rectangle Q_2 , or the portal A_2' of the complementary rectangle Q_2' .*

Proof. For $i = 1, 2$, we associate with Q_i a *viewing arc* Γ_i , consisting of all orientations of the rays that emanate from v and intersect Q_i . Each Γ_i is a quarter-circular arc (of angular span 90°), which is partitioned into two subarcs Γ_i^A, Γ_i^E , at the orientation at which v sees the opposite vertex of the corresponding rectangle Q_i ; Γ_i^A (resp., Γ_i^E) is the subarc in which we view A_i (resp., E_i). See Figure 7.

Moreover, the opposite orientations of Q_1 and Q_2 mean that, as we trace these arcs in counterclockwise direction, Γ_1^E precedes Γ_1^A , whereas Γ_2^E succeeds Γ_2^A . That is, the clockwise endpoint of Γ_1 is adjacent to Γ_1^E and we call it the *E-endpoint* of Γ_1 , and the counterclockwise endpoint of Γ_1 , called the *A-endpoint*, is adjacent to Γ_1^A . Symmetrically, the clockwise endpoint of Γ_2 is its A-endpoint and is adjacent to Γ_2^A , and its counterclockwise endpoint is the E-endpoint, adjacent to Γ_2^E . See Figures 7 and 8.

Finally, the overlapping of Q_1 and Q_2 near v mean that the arcs Γ_1 and Γ_2 overlap too. Let $\Gamma_0 := \Gamma_1 \cap \Gamma_2$.

The viewing arcs Γ_1 and Γ_2 can overlap in one of the following two ways. (Since each arc is of angular span 90° , no other kind of overlap, such as nesting of one arc inside the other, is possible.)

AA-overlap: The clockwise and the counterclockwise endpoints of Γ_0 are the A-endpoints of Γ_2 and Γ_1 , respectively. See Figure 8(i).




Figure 7: The two viewing arcs Γ_1, Γ_2 , their partitions into Γ_1^E, Γ_1^A and Γ_2^A, Γ_2^E , mnemonically denoted by the respective symbols A, E , and their overlap.




Figure 8: (i) One possible interaction between Q_1 and Q_2 . The overlap is of type AA , the intersection Γ_0 of the viewing arcs is delimited by the orientations of \vec{vw}_1 and \vec{vw}_2 . (ii) Another possible interaction between Q_1 and Q_2 . The overlap is of type EE , the intersection Γ_0 of the viewing arcs does not contain the orientations of \vec{vw}_1 and \vec{vw}_2 .

EE-overlap: The clockwise and the counterclockwise endpoints of Γ_0 are the E -endpoints of Γ_1 and Γ_2 , respectively. See Figure 8(ii).

We now assume that none of the four intersections (between one of the segments and a suitable portal of the other rectangle), mentioned in the statement of the lemma, occur. We reach a contradiction by showing that under this assumption neither type of overlap can happen.

AA-overlap. For $i = 1, 2$, let $\rho_i(\theta)$, for $\theta \in \Gamma_i$, denote the length of the intersection of Q_i with the ray emanating from v in direction θ . Note that ρ_i is *bimodal*: it increases to its maximum, which occurs at the direction to the vertex of Q_i opposite to v , and then decreases back (each of the two pieces is a simple trigonometric function of θ).

Write $\Gamma_0 = [\alpha, \beta]$. Since the overlap of Γ_1 and Γ_2 is an AA -overlap, α is the orientation of \vec{vw}_2 , and we have $\rho_2(\alpha) > \rho_1(\alpha)$ (we have to exit Q_1 before we reach w_2). Symmetrically, we have $\rho_2(\beta) < \rho_1(\beta)$. See Figure 8(i). Hence, by continuity, there must exist $\theta \in \Gamma_0$ where $\rho_1(\theta) = \rho_2(\theta)$.

We claim however that this equality is impossible. Indeed, by taking into account the partitions $\Gamma_1 = \Gamma_1^E \cup \Gamma_1^A, \Gamma_2 = \Gamma_2^A \cup \Gamma_2^E$, and by overlaying them within Γ_0 , we see that Γ_0 is partitioned into




Figure 9: Illustrating the argument that ρ_1 and ρ_2 cannot intersect in an AA-overlap.

at most three subarcs, each being the intersection (within Γ_0) of one of Γ_1^E, Γ_1^A with one of Γ_2^A, Γ_2^E . See Figure 9 (and also Figure 7, where the overlap has only two subarcs, of $\Gamma_1^E \cap \Gamma_2^E$ and $\Gamma_1^A \cap \Gamma_2^E$). Since E_1 and E_2 are disjoint, and since, by assumption, no E -edge of any rectangle intersects the A -edge of the other rectangle, the intersection $\rho_1(\theta) = \rho_2(\theta)$ can only occur within $\Gamma_1^A \cap \Gamma_2^A$. As we trace Γ_0 from α to β , we start with $\rho_2 > \rho_1$, so this still holds as we reach $\Gamma_1^A \cap \Gamma_2^A$. However, the bimodality of ρ_1, ρ_2 and the different orientations of Q_1, Q_2 mean that ρ_1 is *decreasing* on Γ_1^A , whereas ρ_2 is *increasing* on Γ_2^A , so no intersection of these functions can occur within $\Gamma_1^A \cap \Gamma_2^A$, a contradiction that shows that an AA-overlap is impossible.

EE-overlap. We follow the same notations as in the analysis of AA-overlaps, but use different arguments, which bring to bear the complementary rectangles Q'_1, Q'_2 .

Consider the clockwise endpoint α of Γ_0 , which, by construction, is the E-endpoint of Γ_1 , incident to Γ_1^E . Consider first the subcase where $\rho_1(\alpha) > \rho_2(\alpha)$. That is, the edge A_1^* of Q_1 connecting v to E_1 crosses and exits Q_2 before reaching E_1 ; it may exit Q_2 either at E_2 (as depicted in Figure 8(ii)) or at A_2 (as depicted in Figure 10).




Figure 10: (i) Another instance of an EE-overlap. (ii) The intersection of the viewing arcs; it consists of three subarcs (in counterclockwise order): $\Gamma_1^E \cap \Gamma_2^A$, $\Gamma_1^A \cap \Gamma_2^A$, and $\Gamma_1^A \cap \Gamma_2^E$. Here w_2 lies in the diametral disk (not drawn) spanned by the edge A_1^* of Q_1 .

If A_1^* exits Q_2 at E_2 then we follow E_2 into the complementary rectangle Q'_1 . By our assumption (that no intersection as stated in the lemma occurs) E_2 cannot exit Q'_1 through its anchor side A'_1 (as depicted in Figure 11(i)). So E_2 must end inside Q'_1 , at an endpoint q_2 (see Figure 11(ii)). But then the right angle $q_2 w_2 v$ must either cross the anchor A'_1 twice, or be fully contained in Q'_1 . In the latter case w_2 lies in $Q'_1 \subset R_1$, contrary to the assumption that v is non-terminal, and in the former case w_2 lies in the disk with A'_1 as a diameter, which is also contained in K_1 , and again we have a contradiction.




Figure 11: Another instance of an EE-overlap. The intersection of the viewing arcs consists of just $\Gamma_1^E \cap \Gamma_2^E$. (i) E_2 crosses the anchor A'_1 of the complementary rectangle Q'_1 . (ii) E_2 ends inside Q'_1 .

If A_1^* exits Q_2 at A_2 , the argument is simpler, because then w_2 is contained in the disk with A_1^* as a diameter, which is contained in K_1 , again contrary to the assumption that v is non-terminal (see Figure 10).

A fully symmetric argument leads to a contradiction in the case where $\rho_1(\beta) < \rho_2(\beta)$. It therefore remains to consider the case where $\rho_1(\alpha) < \rho_2(\alpha)$ and $\rho_1(\beta) > \rho_2(\beta)$. Here we argue exactly as in the case of AA-overlaps, using the bimodality of ρ_1 and ρ_2 , that this case cannot happen. (Figure 9 depicts the situation in this case too.) Specifically, there has to exist an intersection point of ρ_1 and ρ_2 within Γ_0 , and it can only occur at $\Gamma_1^A \cap \Gamma_2^A$. But over this subarc ρ_1 is decreasing and ρ_2 is increasing, and we enter this subarc with $\rho_1 < \rho_2$, so these functions cannot intersect within this arc. This completes the argument showing that our assumption implies that an EE-overlap is not possible.

We conclude that one of the intersections stated in the lemma must exist. \square

The charging scheme. We charge v to a portal (A_2 or A'_2) of R_2 that intersects E_1 or to a portal (A_1 or A'_1) of R_1 that intersects E_2 . At least one such intersection must exist by Lemma 3.1. A useful property of this charging, which will be needed in the next part of the analysis, is given by the following lemma.

Lemma 3.2. *Let v be a non-terminal RR-vertex, lying on $\partial R_i \cap \partial R_j$, which is charged to a portal h_j of R_j . Then e_i , traced from its intersection with h_j into R_j , gets further away from e_j (that is, the distance from the traced point to e_j keeps increasing).*

Proof. Suppose to the contrary that e_i approaches e_j and assume, without loss of generality, that e_j is horizontal, that v lies on e_j^- , and that h_j is the left-lower portal of R_j . In this case e_i has positive slope. See Figure 12.




Figure 12: Illustrating the proof of Lemma 3.2. Only the lower portions of R_i and R_j are shown.

Let q denote the endpoint of h_j incident to e_j and let q' denote the lower endpoint of h_j . (Note that q' is w_j if v is charged to the portal A_j of R_j which is also a portal of Q_j and q' is the endpoint of A'_j , the portal of the complementary rectangle Q'_j , otherwise.) Since v is a non-terminal RR-vertex, the segment $\vec{vq'}$, as we trace it from v , enters R_i (this follows as e_i has positive slope) and then exits it before reaching q' . The exit point lies on a portal h_i of R_i . Since e_i intersects h_j , it follows that h_i must also cross h_j . But then q' must lie inside the diametral disk spanned by h_i , and thus it lies inside K_i , a contradiction (to v being non-terminal) that completes the proof. \square

The expected number of vertices charged to a portal. Fix a segment of \mathcal{E} , denote it as e_0 , and rename the other segments as e_1, \dots, e_{n-1} . Assume, for simplicity, that e_0 is horizontal. We bound the expected number of vertices charged to the lower-left portal, denoted by g , of the rectangle R_0 (which is incident to the left endpoint, a_0 , of e_0); symmetric arguments will apply to the other three portals of R_0 . Given a specific permutation (r_0, \dots, r_{n-1}) of the input set of distances Θ , let $\chi_{\text{RR}}(g; r_0, \dots, r_{n-1})$ denote the number of vertices charged to g if e_i is expanded by r_i , for $i = 0, \dots, n-1$. We wish to bound $\bar{\chi}_{\text{RR}}(g)$, the expected value of $\chi_{\text{RR}}(g; r_0, \dots, r_{n-1})$ with respect to the random choice of the r_i 's, as effected by randomly shuffling them (by a random permutation acting on Θ).

We first fix a value r (one of the values $\theta_i \in \Theta$) of r_0 and bound $\bar{\chi}_{\text{RR}}(g \mid r)$, the expected number of vertices charged to g conditioned on the choice $r_0 = r$; the expectation is taken over those permutations that fix $r_0 = r$; they can be regarded as *random* permutations of the remaining elements of Θ . Then we bound $\bar{\chi}_{\text{RR}}(g)$ by taking the expectation of the resulting bound over the choice of r_0 .

So fix $r_0 = r$. Set $K_0 = e_0 \oplus D(r)$, and let ℓ_0^- denote the line supporting e_0^- . We have $g = a_0 a_0^-$, and observe that all these objects depend only on r_0 , so they are now fixed. By our charging scheme, if a vertex $v \in \partial R_0 \cap \partial R_j$ is charged to the portal g , then $v \in e_0^-$, and e_j intersects g . Furthermore, by Lemma 3.2, the slope of e_j is negative. Let $\mathcal{E}_g \subseteq \mathcal{E} \setminus \{e_0\}$ be the set of segments that intersect g and have negative slopes; the set \mathcal{E}_g depends on the choice of $r_0 = r$ but not on (the shuffled distances) r_1, \dots, r_{n-1} .

For a fixed permutation (r_1, \dots, r_{n-1}) , set $K_l := e_l \oplus D(r_l)$, for $e_l \in \mathcal{E}_g$, $\mathcal{K}_g = \{K_l \mid e_l \in \mathcal{E}_g\}$, and $\mathcal{U}_g = \mathcal{U}(\mathcal{K}_g) \cap e_0^-$. We call a vertex of \mathcal{U}_g an *R-vertex* if it lies on ∂R_i for some $e_i \in \mathcal{E}_g$ (as opposed to lying on some semicircular arc). If a non-terminal RR-vertex v is charged to the portal g , then v is an R-vertex of \mathcal{U}_g (for the specific choice $r_0 = r$). It thus suffices to bound the expected number of R-vertices on \mathcal{U}_g , where the expectation is taken over the random shuffles of r_1, \dots, r_{n-1} .

Consider a segment $e_i \in \mathcal{E}_g$. If $\ell_0^- \cap e_i \neq \emptyset$ then we put $q_i = \ell_0^- \cap e_i$. If $\ell_0^- \cap e_i = \emptyset$, then let λ_i denote the line perpendicular to e_i through b_i (the right endpoint of e_i), and define q_i to be the




Figure 13: Segments in \mathcal{E}_g and the union \mathcal{U}_g that they induce on e_0^- .

intersection of λ_i with ℓ_0^- . (We may assume that q_i lies to the right of a_0^- , for otherwise no expansion of e_i will be such that R_i intersects the edge e_0^- .) Define $r_i^* = 0$ if $\ell_0^- \cap e_i \neq \emptyset$ and $r_i^* = |b_i q_i|$ otherwise. For simplicity, write \mathcal{E}_g as $\langle e_1, \dots, e_m \rangle$, for some $m < n$, ordered so that q_1, \dots, q_m appear on ℓ_0^- from right to left in this order; see Figure 14. We remark that q_1, \dots, q_m are independent of the values of r_1, \dots, r_m , and that the order e_1, \dots, e_m may be different from the order of the intercepts of these segments along g (e.g., see the segments e_1 and e_2 in Figure 14).




Figure 14: Segments in \mathcal{E}_g and the points that they induce on ℓ_0^- .

For $i = 1, \dots, m$, let r_i be, as above, the (random) expansion distance chosen for e_i , and set $J_i = R_i \cap \ell_0^-$. If $r_i \leq r_i^*$ then $J_i = \emptyset$, and if $r_i > r_i^*$ then J_i is an interval containing q_i . Let \mathcal{U}_0 be the union of the intervals J_i , and let $\mu(r; r_1, \dots, r_m)$ be the number of connected components of \mathcal{U}_0 . Clearly, each R-vertex of \mathcal{U}_g is an endpoint of a component of \mathcal{U}_0 , which implies that $\chi_{RR}(g; r, r_1, \dots, r_{n-1}) \leq 2\mu(r; r_1, \dots, r_m)$. It therefore suffices to bound $\bar{\mu}(r)$, the expected value of $\mu(r; r_1, \dots, r_m)$ over the random shuffles of r_1, \dots, r_m .

For each $e_i \in \mathcal{E}_g$, let β_i be the length of the segment connecting a_0^- to its orthogonal projection on e_i . As is easily checked, we have $\beta_i < r$. It is also clear that if $r_i \geq \beta_i$ then the entire segment $q_i a_0^-$ is contained in K_i .

Lemma 3.3. *In the preceding notations, the expected value of $\bar{\mu}(r)$ is $O(\log n)$.*

Proof. Assume that $r = \theta_{n-k+1}$, for some $k \in \{1, \dots, n\}$. We claim that in this case $\bar{\mu}(r) \leq n/k$.

For $i = 1, \dots, m$, if $r_i > r$ then $r_i > \beta_i$ and therefore $a_0^- \in J_i$. Hence, if i is the smallest index for which $r_i > r$ (assuming that such an index exists), then \mathcal{U}_0 has at most i connected components: the one containing J_i and at most $i - 1$ intervals to its right.

Recall that we condition the analysis on the choice of $r_0 = r$, and that we are currently assuming that r_0 is the k th largest value of Θ . For this fixed value of r_0 , the set \mathcal{E}_g is fixed.

Order the segments in $\mathcal{E}_0 := \mathcal{E} \setminus e_0$ by placing first the m segments of \mathcal{E}_g in their order as defined above, and then place the remaining $n - m - 1$ segments in an arbitrary order. Clearly this reshuffling of the segments does not affect the property that the expansion distances in $\Theta_0 := \Theta \setminus \{r\}$ that are assigned to them form a random permutation of Θ_0 .

In this context, $\bar{\mu}(r)$ is upper bounded by the expected value of the index j of the first segment e_j in \mathcal{E}_0 that gets one of the $k - 1$ distances larger than r . (In general, the two quantities are not equal, because we set $\mu(r; r_1, \dots, r_m) = m$ when j is greater than m , that is, in case no segment of \mathcal{E}_g gets a larger distance.)

As is well known, the expected value of the index j of the first appearance of one of $k - 1$ distinguished elements in a random permutation of $n - 1$ elements is n/k , from which our claim follows. This follows since the expected size of a set in an ordered partition of n elements into k sets is n/k , e.g., as in [14, p. 175, Problem 2]. Here we partition $n - 1 - (k - 1)$ elements into k sets, by placing the distinguished elements between them, and we are interested in the expected size of the first group, plus the distinguished element following it, that is $(n - 1 - (k - 1))/k + 1 = n/k$. Note that the case $k = 1$ is special, because no index can get a larger value, but the resulting expectation, namely n , serves as an upper bound for $\bar{\mu}(r)$.

Since $r = \theta_{n-k+1}$ with probability $1/n$, for every k , we have

$$\mathbb{E}[\bar{\mu}(r)] = \sum_{k=1}^n \frac{1}{n} \cdot \bar{\mu}(\theta_{n-k+1}) \leq \frac{1}{n} \sum_{k=1}^n \frac{n}{k} = \sum_{k=1}^n \frac{1}{k} = O(\log n).$$

□

Putting it all together. Lemma 3.3 proves that the expected number of non-terminal vertices of \mathcal{U} charged to a fixed portal of some rectangle in \mathcal{R} is $O(\log n)$. By Lemma 3.1, each non-terminal RR-vertex of \mathcal{U} is charged to one of the $4n$ portals of the rectangles in \mathcal{R} . Repeating this analysis for all these $4n$ portals, the expected number of non-terminal RR-vertices in \mathcal{U} is $O(n \log n)$. Adding the linear bound on the number of terminal RR-vertices, we obtain the following result.

Lemma 3.4. *The expected number of RR-vertices of $\mathcal{U}(\mathcal{K})$ is $O(n \log n)$.*

3.2 CR-vertices

Next, we bound the expected number of CR-vertices of \mathcal{U} . Using a standard notation, we call a vertex $v \in \mathcal{U}$ lying on $\partial K_i \cap \partial K_j$ *regular* if ∂K_i and ∂K_j intersect at two points (one of which is v); otherwise v is called *irregular*. By a result of Pach and Sharir [24], the number of regular vertices on $\partial \mathcal{U}$ is proportional to n plus the number of irregular vertices on $\partial \mathcal{U}$. Since the expected number of RR- and CC-vertices on $\partial \mathcal{U}$ is $O(n \log n)$, the number of regular CR-vertices on $\partial \mathcal{U}$ is $O(n \log n + \kappa)$, where κ is the number of irregular CR-vertices on $\partial \mathcal{U}$. It thus suffices to prove that the expected value of κ is $O(n \log n)$.

Geometric properties of CR-vertices. We begin by establishing a few simple geometric lemmas.

Lemma 3.5. *Let D and D' be two disks of respective radii r, r' and centers o, o' . Assume that $r' \geq r$ and that $o' \in D$. Then $D' \cap \partial D$ is an arc of angular extent at least $2\pi/3$, whose midpoint lies at the radius vector of D from o through o' .*

Proof. We may assume that D is not fully contained in D' , for otherwise the claim is trivial. Consider then the triangle $oo'p$, where p is one of the intersection points of ∂D and $\partial D'$. Put $|oo'| = d \leq r$, and let $\angle o'op = \theta$; see Figure 15. Then

$$\cos \theta = \frac{r^2 + d^2 - r'^2}{2dr} \leq \frac{d^2}{2dr} = \frac{d}{2r} \leq \frac{1}{2}.$$




Figure 15: Illustration of the proof of Lemma 3.5.

Hence $\theta \geq \pi/3$. Since the angular extent of $D' \cap \partial D$ is 2θ , the claim follows. The property concerning the center of the arc $D' \cap \partial D$ is also obvious. \square

Corollary 3.6. *Let D and D' be two disks of radii r and r' and centers o and o' , respectively, such that $o' \in D$. Let D_1 be some sector of D of angle $\pi/3$, and let γ_1 denote the circular portion of ∂D_1 .*

(a) *If $o' \in D_1$ and $r < r'$, then γ_1 is fully contained in D' .*

(b) *If $o' \in D \setminus D_1$, then either D' is disjoint from γ_1 or $D' \cap \gamma_1$ consists of one or two arcs, each containing an endpoint of γ_1 . See Figure 16(ii) and (iii).*

(In (b), no relationship between r and r' is assumed.)

Proof. See Figure 16(i). Claim (a) follows from the preceding lemma, since $D' \cap \partial D$ is an arc of angular extent at least $2\pi/3$ whose midpoint lies at a point on γ_1 .




Figure 16: Illustration of Corollary 3.6: (i) The setup for (a). (ii) One possible scenario in the proof of (b), where $\delta \cap \gamma_1$ consists of a single arc. (iii) Another scenario in the proof of (b), where $\delta \cap \gamma_1$ consists of two arcs.

For (b), $D' \cap \partial D$ is a connected arc δ , whose midpoint lies in direction $\vec{o}\vec{o}'$ and thus outside γ_1 (recall that $o' \notin D_1$), and $D' \cap \gamma_1 = \delta \cap \gamma_1$. The intersection of two arcs of the same circle consists of zero, one, or two connected subarcs. In the first case the claim is obvious. In the third case, each of the arcs δ, γ_1 must contain both endpoints of the other arc, so (b) follows. In the second case, the only situation that we need to rule out is when $\delta \cap \gamma_1$ is contained in the relative interior of γ_1 , in which case δ , and its midpoint, are contained in γ_1 , contrary to assumption. Hence (b) holds in this case too. \square

Fix a segment of \mathcal{E} , call it e_0 , and rename the other segments to be e_1, \dots, e_{n-1} . ∂K_0 has two semicircular arcs, each corresponding to a different endpoint of e_0 . We fix one of the semicircular arcs of K_0 and denote it by γ_0 . Let r_0 be the random distance assigned to e_0 , let D_0 be the disk of radius r_0 containing γ_0 on its boundary, and let $H_0 \subset D_0$ be the half-disk spanned by γ_0 .

Partition H_0 into three sectors of angular extent $\pi/3$ each, denoted as H_{01}, H_{02}, H_{03} . Let $\gamma_{0i} \subset \gamma_0$ denote the arc bounding H_{0i} , for $i = 1, 2, 3$. Here we call a vertex $v \in \partial \mathcal{U}$ formed by $\gamma_{0i} \cap \partial K_j$, for some j , a *terminal* vertex if the portion of γ_{0i} between v and one of its endpoints is contained in K_j , and a *non-terminal* vertex otherwise. There are at most $O(1)$ terminal vertices on γ_0 , for an overall bound of $O(n)$ on the number of such vertices, so it suffices to bound the expected number of non-terminal irregular CR-vertices on each subarc γ_{0i} , for $i = 1, 2, 3$.

Let $\mathcal{E}(r_0)$ denote the set of all segments $e_j \neq e_0$ that intersect the disk D_0 , and, for $i = 1, 2, 3$, let $\mathcal{E}_i(r_0) \subseteq \mathcal{E}(r_0)$ denote the set of all segments $e_j \neq e_0$ that intersect the sector H_{0i} . Set $m_i := m_i(r_0) = |\mathcal{E}_i(r_0)|$. Segments in $\mathcal{E}(r_0) \setminus \mathcal{E}_i(r_0)$ intersect D_0 but are disjoint from H_{0i} . (The parameter r_0 is to remind us that all these sets depend (only) on the choice of r_0 .)

Lemma 3.7. *Let $e_j \in \mathcal{E} \setminus \mathcal{E}(r_0)$. If \mathcal{U} has a CR-vertex $v \in \gamma_{0i} \cap \partial K_j$, for some $i = 1, 2, 3$, then v is either a regular vertex or a terminal vertex.*

Proof. Let c denote the center of D_0 (it is an endpoint of e_0), and consider the interaction between K_j and D_0 . We split the analysis into the following two cases.

Case 1, $r_j \leq r_0$: Regard D_0 as $D_0^* \oplus D(r_j)$, where D_0^* is the disk of radius $r_0 - r_j$ centered at c . By assumption, D_0^* and e_j are disjoint, implying that D_0 and K_j are pseudo-disks (cf. [18]), that is, their boundaries intersect in two points, one of which is v ; denote the other point as v' . See Figure 17(a).

If only v lies on γ_0 , then v must be a terminal vertex, so assume that both v and v' lie on γ_0 . We claim that ∂K_j and ∂K_0 can intersect only at v and v' , implying that v is regular. Indeed, v and v' partition ∂K_j into two connected pieces. One piece is inside D_0 , locally near v and v' , and cannot intersect ∂K_i in a point other than v and v' without intersecting D_0 in a third point (other than v and v'), contradicting the property that D_0 and K_j are pseudo-disks. The other connected piece of ∂K_j between v and v' is separated from $\partial K_i \setminus \gamma_0$ by the line through v and v' and therefore cannot contain intersections other than v and v' between ∂K_j and ∂K_i .

Case 2, $r_j > r_0$: Let $K_j^* = e_j \oplus D(r_j - r_0)$. K_j can now be regarded as $K_j^* \oplus D(r_0)$. If $c \notin K_j^*$, then by the result of [18], $D_0 = c \oplus D(r_0)$ and K_j are pseudo-disks; see Figure 17(b). Therefore, the argument given above for the case where $r_0 \geq r_j$ implies the lemma in this case as well. Finally, $c \in K_j^*$ implies that K_j contains D_0 , so this case cannot occur (it contradicts the existence of v). See Figure 17(c). \square

Using Lemmas 3.5 and 3.7, we obtain the following property.

Lemma 3.8. *Let $v \in \gamma_{0i} \cap \partial K_j$ be a non-terminal, irregular CR-vertex of \mathcal{U} . Then (i) $e_j \in \mathcal{E}_i(r_0)$, and (ii) for all $e_l \in \mathcal{E}_i(r_0)$, $r_l < r_0$.*

Proof. (i) Lemma 3.7 implies that $e_j \in \mathcal{E}(r_0)$. Suppose to the contrary that $e_j \in \mathcal{E}(r_0) \setminus \mathcal{E}_i(r_0)$. Let v' be the orthogonal projection of v onto e_j (v' exists since v is a CR-vertex). Assume first that v' is inside D_0 (as in Figure 18(i)), and consider the disk D' of radius r_j centered at v' . Clearly, $D' \subseteq K_j$. Part (b) of Corollary 3.6 implies that D' intersects γ_{0i} at an arc or a pair of arcs, each containing an endpoint of γ_{0i} , i.e., v is a terminal vertex, contrary to assumption.




Figure 17: (a): The case when $r_0 \geq r_j$ and $\partial K_j \cap \gamma_0$ contains the two intersection points of ∂D_0 and ∂K_j . (b) The case when $r_j > r_0$ and $c \notin K_j^*$. (c) The case when $r_j > r_0$ and $c \in K_j^*$.

Assume then that $v' \notin D_0$. Since e_j intersects D_0 , there exists a point $u \in e_j \cap \partial D_0$ such that uv' is openly disjoint from D_0 . By our assumption that $e_j \in \mathcal{E}(r_0) \setminus \mathcal{E}_i(r_0)$, $u \notin \gamma_{0i}$. Let δ denote the arc of ∂D_0 that connects u and v and lies, locally near u , on the side of e_j that contains v . See Figure 18(ii). It can be verified that δ is the smaller of the two arcs of ∂D_0 connecting u and v . Clearly, δ contains an endpoint of γ_{0i} , and the desired contradiction will be reached by showing that δ is contained in K_j .

The arc δ lies inside a wedge W , with apex u , formed by the ray uv and the line supporting e_j , and also inside the strip Σ bounded by e_j and the parallel line through v . Let b_j be the endpoint of e_j that lies on the boundary of W , γ_j^+ the circular arc of K_j centered at b_j , and C_j^+ the circle (centered at b_j) containing the arc γ_j^+ . Since $\delta \subset W \cap \Sigma$, δ can intersect ∂K_j only at a point on the arc γ_j^+ . Therefore, if $\delta \not\subset K_j$, then δ escapes K_j through a point $w \in \gamma_j^+$. Note that the endpoint v of δ lies outside C_j^+ , so the portion of δ between v and w intersects C_j^+ at a point other than w . Finally, the endpoint u of δ lies inside K_j , so the portion of δ between w and u , which lies outside D_j^+ locally near w , also intersects γ_j^+ at a point different from w ; δ enters the disk D_j^+ bounded by C_j^+ at this point. This, however, implies that δ intersect C_j^+ at three points, a contradiction. Hence $\delta \subset K_j$.

To establish (ii), notice that Part (a) of Corollary 3.6 implies that $r_l < r_0$ for all $e_l \in \mathcal{E}_i(r_0)$, because otherwise we would have $\gamma_{0i} \subset K_l$ and γ_{0i} would not contain any vertex of $\partial \mathcal{U}$. \square

The expected number of non-terminal vertices on γ_0 . We are now ready to bound the expected number of non-terminal irregular CR-vertices of \mathcal{U} that lie on the semi-circular arc γ_0 of K_0 . Note that γ_0 is not fixed, as it depends on the value of r_0 . Let $\chi_{\text{CR}}(\gamma_0; r_0, r_1, \dots, r_{n-1})$ denote the number

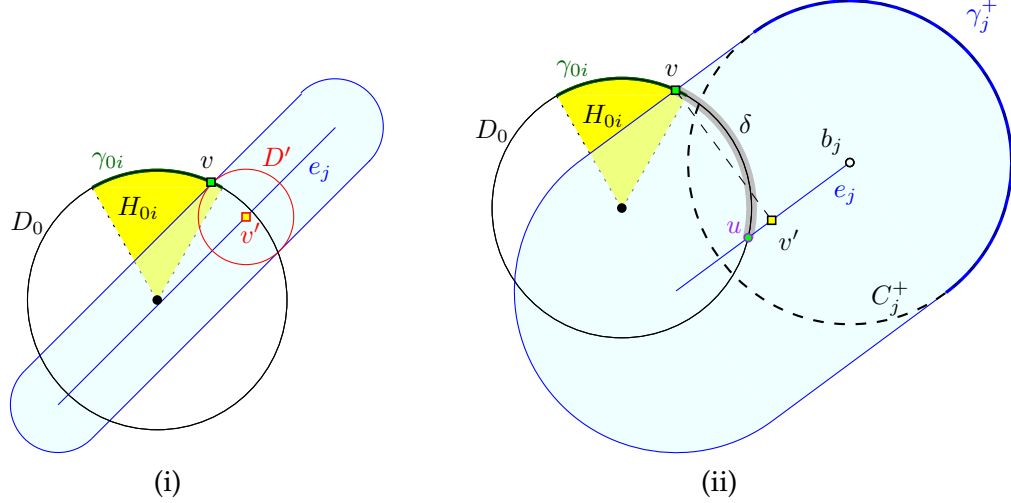


Figure 18: Illustration of the proof of Lemma 3.8: (i) The case where $v' \in D_0$. (ii) The case where $v' \notin D_0$.

of non-terminal irregular vertices on γ_0 , assuming that r_i is the expansion distance of K_i , for $i = 0, \dots, n-1$. Our goal is to bound

$$\bar{\chi}_{\text{CR}}(\gamma_0) = \mathbb{E}[\chi_{\text{CR}}(\gamma_0; r_0, \dots, r_{n-1})]$$

where the expectation is over all the random permutations assigning these distances to the segments of \mathcal{E} . As for RR-vertices, we first fix the value of r_0 to, say, r , and bound $\chi_{\text{CR}}(\gamma_0 \mid r)$, the expected value of $\chi_{\text{CR}}(\gamma_0, r, r_1, \dots, r_{n-1})$, where the expectation is taken over the random shuffles of r_1, \dots, r_{n-1} , and then bound $\bar{\chi}_{\text{CR}}(\gamma_0)$ by taking the expectation of the bound over the choice of r_0 .

Lemma 3.9. *Using the notation above, $\bar{\chi}_{\text{CR}}(\gamma_0) = O(\log n)$.*

Proof. Following the above scheme, suppose that the value r_0 is indeed fixed to r , so γ_0 and γ_{0i} , $1 \leq i \leq 3$, are fixed. As above, set $m_i = |\mathcal{E}_i(r)|$, for $i = 1, 2, 3$; the sets $\mathcal{E}_i(r)$ and their sizes m_i are also fixed. We bound the expected number of non-terminal irregular vertices on γ_{0i} , for a fixed $i \in \{1, 2, 3\}$. By Lemma 3.8, any such vertex lies on the boundary of \mathcal{D}_0 , the intersection of γ_{0i} with the union of $\{K_l \mid e_l \in \mathcal{E}_i(r)\}$. Equivalently, it suffices to bound the expected number of connected components of \mathcal{D}_0 that lie in the interior of γ_{0i} . By Lemma 3.8, if $r_l \geq r$ for any $e_l \in \mathcal{E}_i(r)$, then there are no such components.

Assume that $r = \theta_{n-k+1}$ for some $k \in \{1, \dots, n\}$. To bound $\chi_{\text{CR}}(\gamma_0 \mid r)$, we first bound the probability p that all the m_i radii that are assigned to the segments of $\mathcal{E}_i(r)$ are smaller than r . We

have

$$\begin{aligned}
p &= \frac{\binom{n-k}{m_i}}{\binom{n-1}{m_i}} = \frac{(n-k)(n-k-1)\cdots(n-k-m_i+1)}{(n-1)(n-2)\cdots(n-m_i)} \\
&= \left(1 - \frac{k-1}{n-1}\right) \left(1 - \frac{k-1}{n-2}\right) \cdots \left(1 - \frac{k-1}{n-m_i}\right) \\
&< \left(1 - \frac{k-1}{n-1}\right)^{m_i} < e^{-(k-1)m_i/(n-1)}.
\end{aligned}$$

By Lemma 3.8 (ii), with probability $1 - p$, there are no connected components. In the complementary case, when all the m_i radii under consideration are smaller than r , we pessimistically bound the number of connected components by $2m_i$ — each segment of $\mathcal{E}_i(r)$ can generate at most two connected components. In other words, when $k \geq 2$, the expected number of connected components of \mathcal{J}_0 is at most

$$\begin{aligned}
2m_i p &< 2m_i e^{-(k-1)m_i/(n-1)} \\
&= \frac{2(n-1)}{k-1} \cdot \left(((k-1)m_i/(n-1)) e^{-(k-1)m_i/(n-1)} \right) \\
&< \frac{2(n-1)}{e(k-1)},
\end{aligned}$$

because the maximum value of the expression xe^{-x} is e^{-1} . For $k = 1$, the value of p is q , so the expected number of connected components in this case is $2m_i \leq 2(n-1)$.

Since $r = \theta_{n-k+1}$ with probability $1/n$ for every k , we have

$$\begin{aligned}
\mathbb{E}[\bar{\chi}_{\text{CR}}(\gamma_0)] &= \mathbb{E}[\chi_{\text{CR}}(\gamma_0 \mid r)] = \sum_{k=1}^n \frac{1}{n} \cdot \mathbb{E}[\chi_{\text{CR}}(\gamma_0 \mid \theta_{n-k+1})] \\
&\leq \frac{1}{n} \left[2(n-1) + \sum_{k=2}^n \frac{2(n-1)}{e(k-1)} \right] \\
&= O\left(\sum_{k=1}^n \frac{1}{k} \right) = O(\log n).
\end{aligned}$$

□

Summing this bound over all three subarcs of γ_0 and adding the constant bound on the number of terminal (irregular) vertices, we obtain that the expected number of irregular CR-vertices of \mathcal{U} on γ_0 is $O(\log n)$. Summing these expectations over the $2n$ semicircular arcs of the racetracks in \mathcal{K} , and adding the bounds on the number of regular CR-vertices we obtain the following lemma.

Lemma 3.10. *The expected number of CR-vertices on $\mathcal{U}(\mathcal{K})$ is $O(n \log n)$.*

Combining Lemma 3.4, Lemma 3.10, and the linear bound on the number of CC-vertices, completes the proof of Theorem 1.2 for the case of segments.

4 The Case of Polygons

In this section we consider the case where the objects of \mathcal{C} are n convex polygons, each with at most s vertices. For simplicity, we prove Theorem 1.2 when each C_i is a convex s -gon—if C_i has fewer than s vertices, we can split some of its edges into multiple edges so that it has exactly s vertices. We reduce this case to the case of segments treated above. A straightforward reduction that just takes the individual edges of the s -gons as our set of segments does not work, since edges of the same polygon are all expanded by the same distance. Nevertheless, we can overcome this difficulty as follows.

Let $\mathcal{C} = \{C_1, \dots, C_n\}$ be the set of the n given polygons, and consider a fixed assignment of expansion distances r_i to the polygons C_i . For each i , enumerate the edges of C_i as $e_{i1}, e_{i2}, \dots, e_{is}$; the order of enumeration is not important. Let v be a vertex of $\partial\mathcal{U}$, lying on the boundaries of K_i and K_j , for some $1 \leq i < j \leq n$. Then there exist an edge e_{ip} of C_i and an edge e_{jq} of C_j such that v lies on $\partial(e_{ip} \oplus D(r_i))$ and on $\partial(e_{jq} \oplus D(r_j))$; the choice of e_{ip} is unique if the portion of ∂K_i containing v is a straight edge, and, when that portion is a circular arc, any of the two edges incident to the center of the corresponding disk can be taken to be e_{ip} . A similar property holds for e_{jq} .

The following stronger property holds too. For each $1 \leq p \leq s$, let \mathcal{C}_p be the set of edges $\{e_{1p}, e_{2p}, \dots, e_{np}\}$, and let $\mathcal{K}_p = \{e_{1p} \oplus D(r_1), \dots, e_{np} \oplus D(r_n)\}$. Then, as is easily verified, our vertex v is a vertex of the union $\mathcal{U}(\mathcal{K}_p \cup \mathcal{K}_q)$. Moreover, for each p , the expansion distances r_i of the edges e_{ip} of \mathcal{C}_p are all the elements of Θ , each appearing once, and their assignment to the segments of \mathcal{C}_p is a *random permutation*. Fix a pair of indices $1 \leq p < q \leq s$, and note that each expansion distance r_i is assigned to exactly two segments of $\mathcal{C}_p \cup \mathcal{C}_q$, namely, to e_{ip} and e_{iq} .

We now repeat the analysis given in the preceding section for the collection $\mathcal{C}_p \cup \mathcal{C}_q$, and make the following observations. First, the analysis of CC-vertices remains the same, since the complexity of the union of any family of disks is linear.

Second, in the analysis of RR- and CR-vertices, the exploitation of the random nature of the distances r_i comes into play only after we have fixed one segment (that we call e_0) and its expansion distance r_0 , and consider the expected number of RR-vertices and CR-vertices on the boundary of $K_0 = e_0 \oplus D(r_0)$, conditioned on the fixed choice of r_0 . Suppose, without loss of generality, that e_0 belongs to \mathcal{C}_p . We first ignore its sibling e'_0 in \mathcal{C}_q (from the same polygon), which receives the same expansion distance r_0 ; e'_0 can form only $O(1)$ vertices of \mathcal{U} with e_0 .² The interaction of e_0 with the other segments of \mathcal{C}_p behaves exactly as in Section 3, and yields an expected number of $O(\log n)$ RR-vertices of $\mathcal{U}(\mathcal{K}_p)$ charged to the portals of R_0 and an expected number of $O(\log n)$ CR-vertices charged to circular arcs of K_0 . Similarly, The interaction of e_0 with the other segments of \mathcal{C}_q (excluding e'_0) is also identical to that in Section 3, and yields an additional expected number of $O(\log n)$ vertices of $\mathcal{U}(\{e_0\} \cup \mathcal{K}_q)$ charged to an portals and circular arcs of K_0 . Since any vertex of $\mathcal{U}(\mathcal{K}_p \cup \mathcal{K}_q)$ involving e_0 must be one of these two kinds of vertices, we obtain a bound of $O(\log n)$ on the expected number of such vertices, and summing this bound over all segments e_0 of $\mathcal{C}_p \cup \mathcal{C}_q$, we conclude that the expected complexity of $\mathcal{U}(\mathcal{K}_p \cup \mathcal{K}_q)$ is $O(n \log n)$. (Note also that the analysis just given manages to finesse the issue of segments sharing endpoints.)

Summing this bound over all $O(s^2)$ choices of p and q , we obtain the bound asserted in Theorem 1.2. The constant of proportionality in the bound that this analysis yields is $O(s^2)$.

²As a matter of fact, e_0 and e'_0 do not generate any vertex of the full union $\mathcal{U}(\mathcal{C})$, but they might generate vertices of the partial union $\mathcal{U}(\mathcal{C}_p \cup \mathcal{C}_q)$.




Figure 19: Discretizing the function f_e for an edge e .

5 Network Vulnerability Analysis

Let $\mathcal{E} = \{e_1, \dots, e_n\}$ be a set of n segments in the plane with pairwise-disjoint relative interiors, and let $\varphi : \mathbb{R}_{\geq 0} \rightarrow [0, 1]$ be an edge failure probability function such that $1 - \varphi$ is a cdf. For each segment e_i , define the function $f_i : \mathbb{R}^2 \rightarrow [0, 1]$ by $f_i(q) = \varphi(d(q, e_i))$, for $q \in \mathbb{R}^2$, where $d(q, e_i)$ is the distance from q to e_i , and set

$$\Phi(q, \mathcal{E}) = \sum_{i=1}^n f_i(q).$$

In this section we present a Monte-Carlo algorithm, which is an adaptation and a simplification of the algorithm described in [2], for computing a location \tilde{q} such that $\Phi(\tilde{q}, \mathcal{E}) \geq (1 - \delta)\Phi(\mathcal{E})$, where $0 < \delta < 1$ is some prespecified error parameter, and where

$$\Phi(\mathcal{E}) = \max_{q \in \mathbb{R}^2} \Phi(q, \mathcal{E}).$$

The expected running time of the algorithm is a considerable improvement over the algorithm in [2]; this improvement is a consequence of the bounds obtained in the preceding sections.

To obtain the algorithm we first discretize each f_i by choosing a finite family \mathcal{K}_i of *super-level sets* of f_i (each of the form $\{q \in \mathbb{R}^2 \mid f_i(q) \geq t\}$), and reduce the problem of computing $\Phi(\mathcal{E})$ to that of computing the maximum depth in the arrangement $\mathcal{A}(\mathcal{K})$ of $\mathcal{K} = \bigcup_i \mathcal{K}_i$. Our algorithm then uses a sampling-based method for estimating the maximum depth in $\mathcal{A}(\mathcal{K})$, and thereby avoids the need to construct $\mathcal{A}(\mathcal{K})$ explicitly.

In more detail, set $m = \lceil 2n/\delta \rceil$. For each $1 \leq j < m$, let $r_j = \varphi^{-1}(1 - j/m)$, and let, for $i = 1, \dots, n$, $K_{ij} = e_i \oplus D(r_j)$ be the racetrack formed by the Minkowski sum of e_i with the disk of radius r_j centered at the origin. Note that r_j increases with j . Set $\tilde{\varphi} = \{r_j \mid 1 \leq j < m\}$, $\mathcal{K}_i = \{K_{ij} \mid 1 \leq j < m\}$, and $\mathcal{K} = \bigcup_{1 \leq i \leq n} \mathcal{K}_i$. See Figure 19. Note that we cannot afford to, and indeed do not, compute \mathcal{K} explicitly, as its cardinality (which is quadratic in n) is too large.

For a point $q \in \mathbb{R}^2$ and for a subset $X \subseteq \mathcal{K}$, let $\Delta(q, X)$, the *depth* of q with respect to X , be the number of racetracks of X that contain q in their interior, and let

$$\Delta(X) = \max_{q \in \mathbb{R}^2} \Delta(q, X).$$

The following lemma (whose proof, which is straightforward, can be found in [2]) shows that the maximum depth of \mathcal{K} approximates $\Phi(\varepsilon)$.

Lemma 5.1 (Agarwal et al. [2]). (i) $\Phi(q, \varepsilon) \geq \frac{\Delta(q, \mathcal{K})}{m}$, for each point $q \in \mathbb{R}^2$.

(ii) $\Phi(\varepsilon) \geq \frac{\Delta(\mathcal{K})}{m} \geq (1 - \frac{1}{2}\delta)\Phi(\varepsilon)$.

By Lemma 5.1, it suffices to compute a point \tilde{q} of depth at least $(1 - \frac{1}{2}\delta)\Delta(\mathcal{K})$ in $\mathcal{A}(\mathcal{K})$; by (i) and (ii) we will then have

$$\Phi(\tilde{q}, \varepsilon) \geq \frac{\Delta(\tilde{q}, \mathcal{K})}{m} \geq (1 - \frac{1}{2}\delta) \frac{\Delta(\mathcal{K})}{m} \geq (1 - \frac{1}{2}\delta)^2 \Phi(\varepsilon) > (1 - \delta)\Phi(\varepsilon).$$

We describe a Monte-Carlo algorithm for computing such a point \tilde{q} , which is a simpler variant of the algorithm described in [6] (see also [4]), but we first need the following definitions.

For a point $q \in \mathbb{R}^2$ and for a subset $X \subseteq \mathcal{K}$, let $\omega(q, X) = \frac{\Delta(q, X)}{|X|}$ be the *fractional depth* of q with respect to X , and let

$$\omega(X) = \max_{q \in \mathbb{R}^2} \omega(q, X) = \frac{\Delta(X)}{|X|}.$$

We observe that $\Delta(\mathcal{K}) \geq m - 1$ because the depth near each e_i is at least $m - 1$. Hence,

$$\omega(\mathcal{K}) \geq \frac{m - 1}{|\mathcal{K}|} = \frac{m - 1}{(m - 1)n} = \frac{1}{n}. \quad (2)$$

Our algorithm estimates fractional depths of samples of \mathcal{K} and computes a point \tilde{q} such that $\omega(\tilde{q}, \mathcal{K}) \geq (1 - \frac{1}{2}\delta)\omega(\mathcal{K})$. By definition, this is equivalent to $\Delta(\tilde{q}, \mathcal{K}) \geq (1 - \frac{1}{2}\delta)\Delta(\mathcal{K})$, which is what we need.

We also need the following concept from the theory of random sampling.

For two parameters $0 < \rho, \varepsilon < 1$, we call a subset $A \subseteq \mathcal{K}$ a (ρ, ε) -*approximation* if the following holds for all $q \in \mathbb{R}^2$:

$$|\omega(q, \mathcal{K}) - \omega(q, A)| \leq \begin{cases} \varepsilon\omega(q, \mathcal{K}) & \text{if } \omega(q, \mathcal{K}) \geq \rho, \\ \varepsilon\rho & \text{if } \omega(q, \mathcal{K}) < \rho. \end{cases} \quad (3)$$

This notion of (ρ, ε) -approximation is a special case of the notion of *relative (ρ, ε) -approximation* defined in [16] for general range spaces with finite VC-dimension. The special case at hand applies to the so-called dual range space $(\mathcal{K}, \mathbb{R}^2)$, where the ground set \mathcal{K} is our collection of racetracks, and where each point $q \in \mathbb{R}^2$ defines a range equal to the set of racetracks containing q ; here $\Delta(q, \mathcal{K})$ is the size of the range defined by q , and $\omega(q, \mathcal{K})$ is its *relative size*. Since $(\mathcal{K}, \mathbb{R}^2)$ has finite VC-dimension (see, e.g., [25]), it follows from a result in [16] that, for any integer b , a random subset of size

$$v(\rho, \varepsilon) := \frac{cb}{\varepsilon^2\rho} \ln n \quad (4)$$

is a (ρ, ε) -approximation of \mathcal{K} with probability at least $1 - 1/n^b$, where c is a sufficiently large constant (proportional to the VC-dimension of our range space). In what follows we fix b to be a sufficiently large integer, so as to guarantee (via the probability union bound) that, with high

probability, all the samples that we construct in the algorithm will have the desired approximation property.

The algorithm works in two phases. The first phase finds a value $\rho \geq 1/n$ such that $\omega(\mathcal{K}) \in [\rho, 2\rho]$. The second phase exploits this “localization” of $\omega(\mathcal{K})$ to compute the desired point \tilde{q} .

The first phase performs a decreasing exponential search: For $i \geq 1$, the i th step of the search tests whether $\omega(\mathcal{K}) \leq 1/2^i$. If the answer is YES, the algorithm moves to the $(i + 1)$ -st step; otherwise it switches to the second phase. Since we always have $\omega(\mathcal{K}) \geq 1/n$ (see (2)), the first phase consists of at most $\lceil \log_2 n \rceil$ steps.

At the i th step of the first phase, we fix the parameters $\rho_i = 1/2^i$ and $\varepsilon = 1/8$, and construct a $(2\rho_i, \varepsilon)$ -approximation of \mathcal{K} by choosing a random subset $\mathcal{R}_i \subset \mathcal{K}$ of size $v_i = v(2\rho_i, \varepsilon) = O(2^i \log n)$. We construct $\mathcal{A}(\mathcal{R}_i)$, e.g., using the randomized incremental algorithm described in [25, Chapter 4] and compute $\omega(\mathcal{R}_i)$ by traversing the arrangement $\mathcal{A}(\mathcal{R}_i)$. Then, if

$$\omega(\mathcal{R}_i) \leq (1 - 2\varepsilon)\rho_i = \frac{3}{4}\rho_i,$$

we continue to step $i + 1$ of the first phase. Otherwise, we switch to the second phase of the algorithm (which is described below). The following lemma establishes the important properties of the first phase.

Lemma 5.2. *When the algorithm reaches step i of the first phase, we have $\omega(\mathcal{K}) \leq \rho_{i-1}$ and $\omega(\mathcal{K})$ is within the interval $[\omega(\mathcal{R}_i) - \frac{1}{4}\rho_i, \omega(\mathcal{R}_i) + \frac{1}{4}\rho_i]$.*

Proof. The proof is by induction on the steps of the algorithm. Assume that the algorithm is in step i . Then by induction $\omega(\mathcal{K}) \leq \rho_{i-1} = 2\rho_i$ (for $i = 1$ this is trivial since $\rho_{i-1} = 1$). This, together with \mathcal{R}_i being a $(2\rho_i, \varepsilon)$ -approximation of \mathcal{K} , implies by (3) that

$$\omega(\mathcal{K}) = \omega(q^*, \mathcal{K}) \leq \omega(q^*, \mathcal{R}_i) + 2\varepsilon\rho_i \leq \omega(\mathcal{R}_i) + 2\varepsilon\rho_i = \omega(\mathcal{R}_i) + \frac{1}{4}\rho_i, \quad (5)$$

where q^* is a point satisfying $\omega(q^*, \mathcal{K}) = \omega(\mathcal{K})$. Furthermore, using (3) again (in the opposite direction), we conclude that

$$\omega(\mathcal{K}) \geq \omega(q_i, \mathcal{K}) \geq \omega(q_i, \mathcal{R}_i) - 2\varepsilon\rho_i = \omega(\mathcal{R}_i) - \frac{1}{4}\rho_i,$$

where q_i is a point satisfying $\omega(q_i, \mathcal{R}_i) = \omega(\mathcal{R}_i)$. So we conclude that $\omega(\mathcal{K})$ is in the interval specified by the lemma.

The algorithm continues to step $i + 1$ if $\omega(\mathcal{R}_i) \leq \frac{3}{4}\rho_i$. But then by (5) we get that $\omega(\mathcal{K}) \leq \frac{3}{4}\rho_i + \frac{1}{4}\rho_i = \rho_i$ as required. \square

Suppose that the algorithm decides to terminate the first phase and continue to the second phase, at step i . Then, by Lemma 4.2, we have that $\omega(\mathcal{K}) \in [\omega(\mathcal{R}_i) - \frac{1}{4}\rho_i, \omega(\mathcal{R}_i) + \frac{1}{4}\rho_i]$. Since, by construction, $\omega(\mathcal{R}_i) > \frac{3}{4}\rho_i$, the ratio between the endpoints of this interval is at most 2, as is easily checked, so if we set $\rho = \omega(\mathcal{R}_i) - \frac{1}{4}\rho_i$ then $\omega(\mathcal{K}) \in [\rho, 2\rho]$ as required upon entering the second phase.

In the second phase, we set $\rho = \omega(\mathcal{R}_i) - \frac{1}{4}\rho_i$ and $\varepsilon = \delta/4$, and construct a (ρ, ε) -approximation of \mathcal{K} by choosing, as above, a random subset \mathcal{R} of size $v = v(\rho, \varepsilon) = O(2^i \log n)$. We compute $\mathcal{A}(\mathcal{R})$, using the randomized incremental algorithm in [25], and return a point $\tilde{q} \in \mathbb{R}^2$ of maximum depth in $\mathcal{A}(\mathcal{R})$.

This completes the description of the algorithm.

Correctness. We claim that $\omega(\tilde{q}, \mathcal{K}) \geq (1 - \frac{1}{2}\delta)\omega(\mathcal{K})$. Indeed, let $q^* \in \mathbb{R}^2$ be, as above, a point of maximum depth in $\mathcal{A}(\mathcal{K})$. We apply (3), use the fact that $\omega(q^*, \mathcal{K}) \geq \rho$, and consider two cases. If $\omega(\tilde{q}, \mathcal{K}) \geq \rho$ then

$$\begin{aligned} \omega(\tilde{q}, \mathcal{K}) &\geq \frac{\omega(\tilde{q}, \mathcal{R})}{1 + \frac{1}{4}\delta} \geq \frac{\omega(q^*, \mathcal{R})}{1 + \frac{1}{4}\delta} \geq \frac{1 - \frac{1}{4}\delta}{1 + \frac{1}{4}\delta} \omega(q^*, \mathcal{K}) \\ &\geq (1 - \frac{1}{2}\delta)\omega(q^*, \mathcal{K}) = (1 - \frac{1}{2}\delta)\omega(\mathcal{K}). \end{aligned}$$

On the other hand, if $\omega(\tilde{q}, \mathcal{K}) < \rho$ then

$$\begin{aligned} \omega(\tilde{q}, \mathcal{K}) &\geq \omega(\tilde{q}, \mathcal{R}) - \frac{\delta}{4}\rho \geq \omega(q^*, \mathcal{R}) - \frac{\delta}{4}\rho \\ &\geq (1 - \frac{\delta}{4})\omega(q^*, \mathcal{K}) - \frac{\delta}{4}\rho \\ &\geq (1 - \frac{\delta}{4})\omega(q^*, \mathcal{K}) - \frac{\delta}{4}\omega(q^*, \mathcal{K}) \\ &= (1 - \frac{1}{2}\delta)\omega(q^*, \mathcal{K}). \end{aligned}$$

Hence in both cases the claim holds. As argued earlier, this implies the desired property

$$\Phi(\tilde{q}, \mathcal{E}) \geq (1 - \delta)\Phi(\mathcal{E}).$$

Running time. We now analyze the expected running time of the algorithm. We first note that we do not have to compute the set \mathcal{K} explicitly to obtain a random sample of \mathcal{K} . Indeed, a random racetrack can be chosen by first randomly choosing a segment $e_i \in \mathcal{E}$, and then by choosing (independently) a random racetrack of \mathcal{K}_i . Hence, each sample \mathcal{R}_i can be constructed in $O(v_i)$ time, and the final sample \mathcal{R} in $O(v)$ time.

To analyze the expected time taken by the i th step of the first phase, we bound the expected number of vertices in $\mathcal{A}(\mathcal{R}_i)$.

Lemma 5.3. *The expected number of vertices in the arrangement $\mathcal{A}(\mathcal{R}_i)$ is $O(2^i \log^3 n)$.*

Proof. By Lemma 5.2 if we perform the i th step of the first phase then $\omega(\mathcal{K}) \leq \rho_{i-1} = 2\rho_i$. Therefore, using (3) we have

$$\omega(\mathcal{R}_i) \leq \omega(\mathcal{K}) + 2\varepsilon\rho_i \leq 2\rho_i + 2\varepsilon\rho_i < 3\rho_i.$$

Therefore, $\Delta(\mathcal{R}_i) = \omega(\mathcal{R}_i)|\mathcal{R}_i| \leq 3\rho_i v_i = O(\log n)$. The elements in \mathcal{R}_i are chosen from \mathcal{K} using the 2-stage random sampling mechanism described above, which we can rearrange so that we first choose a random sample \mathcal{E}_i of segments, and then, with this choice fixed, we choose the random expansion distances. This allows us to view \mathcal{R}_i as a set of racetracks over a fixed set \mathcal{E}_i of segments, each of which is the Minkowski sum of a segment of \mathcal{E}_i with a disk of a random radius, where the radii are drawn uniformly at random and independently from the set $\tilde{\varphi}$. There is a minor technical issue: we might choose in \mathcal{E}_i the same segment $e \in \mathcal{E}$ several times, and these copies of e are not pairwise-disjoint. To address this issue, we slightly shift these multiple copies of e so as to make them pairwise-disjoint. Assuming that \mathcal{E} is in general position and that the cdf defining φ is in “general position” with respect to the locations of the segments of \mathcal{E} , as defined in Section 3, this will not affect the asymptotic maximum depth in the arrangement of the sample.

By Corollary 1.3, applied under the density model and conditioned on a fixed choice of \mathcal{E}_i , the expected value of $|\mathcal{A}(\mathcal{R}_i)|$ is

$$\mathbb{E}[|\mathcal{A}(\mathcal{R}_i)|] = O(\Delta(\mathcal{R}_i)v_i \log n) = O(2^i \log^3 n),$$

implying the same bound for the unconditional expectation too. \square

The expected time spent in constructing $\mathcal{A}(\mathcal{R}_i)$ by the randomized incremental algorithm in [25] is $O(v_i \log v_i + |\mathcal{A}(\mathcal{R}_i)|) = O(2^i \log^3 n)$. Hence, the i th step of the first phase takes $O(2^i \log^3 n)$ expected time. Summing this bound over the steps of the first phase, we conclude that the expected time spent in the first phase is $O(n \log^3 n)$.

In the second phase, $|\mathcal{R}| = O(\frac{1}{\delta^2 \rho} \log n) = O(\frac{n}{\delta^2} \log n)$, and the same argument as above, using (3), implies that

$$\omega(\mathcal{R}) \leq \max\{(1 + \frac{\delta}{4})\omega(\mathcal{K}), \omega(\mathcal{K}) + \frac{\delta}{4}\rho\} = O(\rho);$$

where the latter bound follows as $\omega(\mathcal{K}) \in [\rho, 2\rho]$. Hence, $\Delta(\mathcal{R}) = \omega(\mathcal{R}) \cdot |\mathcal{R}| = O(\frac{1}{\delta^2} \log n)$, and the expected size of $\mathcal{A}(\mathcal{R})$ is thus $O(\Delta(\mathcal{R}) \cdot |\mathcal{R}| \log n) = O(\frac{n}{\delta^4} \log^3 n)$. Since this dominates the cost of the other steps in this phase, the second phase takes $O(\frac{n}{\delta^4} \log^3 n)$ expected time.

Putting everything together, we obtain that the expected running time of the procedure is $O(\frac{n}{\delta^4} \log^3 n)$, and it computes, with high probability, a point \tilde{q} such that $\Phi(\tilde{q}, \mathcal{E}) \geq (1 - \delta)\Phi(\mathcal{E})$. This completes the proof of Theorem 1.4.

6 Discussion

We have shown that if we take the Minkowski sums of the members of a family of pairwise-disjoint convex sets, each of constant description complexity, with disks whose radii are randomly chosen according to a suitable probabilistic model, then the expected complexity of the union of the Minkowski sums is near linear. This generalizes the result of Kedem *et al.* [18] and indicates that the complexity of the union of Minkowski sums is quadratic only if the expansion distances are chosen in an adversarial manner. Our model is related to the so-called *realistic input models*, proposed to obtain more refined bounds on the performance of a variety of geometric algorithms [8]. There are also some similarities between our model and the framework of smoothed analysis [26].

A natural collection of open problems is to tighten the bounds in our theorems or prove corresponding lower bounds. In particular, the following questions arise. (i) The $O(n^{1+\epsilon})$ bound of Theorem 1.1 is unlikely to be tight. Is it possible to prove an $O(n \log n)$ upper bound as we did for polygons in Theorem 1.2? (ii) Can the bound in Theorem 1.2 be improved from $O(s^2 n \log n)$ to $O(sn \log n)$? (iii) Is the bound of Theorem 1.2 asymptotically tight, even for segments, or could one prove a tighter $o(n \log n)$ bound? maybe even linear?

It would be useful to prove, or disprove, that the density and permutation models are equivalent, in the sense that the value of $\psi(\mathcal{C})$ is asymptotically the same under both models for any family \mathcal{C} of pairwise-disjoint convex sets. Nevertheless, it is conceivable that there is a large class of density functions for which the density model yields a better upper bound.

Another interesting direction for future research is to explore other problems that can benefit from our model. For example, following the work in this paper, Har-Peled and Raichel [15] showed that the complexity of the multiplicatively weighted Voronoi diagram of n points in the plane is $O(n \text{ polylog } n)$, if the weights are chosen randomly, in either of the two models assumed in this paper. Recall that if the weights are chosen by an adversary, then the complexity of the diagram can be quadratic in the worst case [7]. More recently, Chang *et al.* [10] showed a similar bound for more generalized Voronoi diagrams, where the distance to a site might depend on several parameters that are chosen randomly. We believe that the expected complexity of the multiplicatively-weighted Voronoi diagram of an arbitrary set of points in \mathbb{R}^3 is subquadratic if their weights are chosen using

one of our models and the spread of the weights is at least a constant. Note that, in contrast, the standard Voronoi diagram can have quadratic complexity in the worst case.

Acknowledgments. The authors thank Emo Welzl for useful discussions concerning the two probabilistic models used in the paper, and the anonymous referees for their valuable comments.

References

- [1] P. K. Agarwal, A. Efrat, S. K. Ganjugunte, D. Hay, S. Sankararaman, and G. Zussman, Network vulnerability to single, multiple, and probabilistic physical attacks, in *Proc. 2010 Military Communication Conference*, 2010, 1824–1829.
- [2] P. K. Agarwal, A. Efrat, S. K. Ganjugunte, D. Hay, S. Sankararaman, and G. Zussman, The resilience of WDM networks to probabilistic geographical failures, *IEEE/ACM Trans. Netw.* 21 (2013), 1525–1538. Also in *Proc. 30th IEEE International Conference on Computer Communications*, 2011, 1521–1529.
- [3] P. K. Agarwal, E. Ezra, and M. Sharir, Near-linear approximation algorithms for geometric hitting sets, *Algorithmica* 63 (2012), 1–25.
- [4] P. K. Agarwal, T. Hagerup, R. Ray, M. Sharir, M. Smid, and E. Welzl, Translating a planar object to maximize point containment, *Proc. 10th Annu. European Sympos. Algorithms*, 2002, 42–53.
- [5] P. K. Agarwal, J. Pach and M. Sharir, State of the union (of geometric objects), *Surveys on Discrete and Computational Geometry*, (J. Goodman, J. Pach and R. Pollack, eds.), Amer. Math. Soc., Providence, RI, 2008, 9–48.
- [6] B. Aronov and S. Har-Peled, On approximating the depth and related problems, *SIAM J. Comput.* 38 (2008), 899–921.
- [7] F. Aurenhammer and R. Klein, Voronoi diagrams, in *Handbook of Computational Geometry* (eds. J.-R. Sack and J. Urrutia), Elsevier, Amsterdam, 1999, 201–290.
- [8] M. de Berg, M. J. Katz, A. F. van der Stappen, and J. Vleugels, Realistic input models for geometric algorithms, *Algorithmica* 34 (2002), 81–97.
- [9] R. Bhandari, *Survivable Networks: Algorithms for Diverse Routing*, Kluwer, Norwell, MA, 1998.
- [10] H.-C. Chang, S. Har-Peled, and B. Raichel, From proximity to utility: A Voronoi partition of Pareto optima, *CoRR*, abs/1404.3403, 2014.
- [11] K. L. Clarkson and P. W. Shor, Applications of random sampling in computational geometry II, *Discrete Comput. Geom.* 4 (1989), 387–421.
- [12] H. Edelsbrunner, B. T. Fasy, and G. Rote, Add isotropic Gaussian kernels at own risk: More and more resilient modes in higher dimensions, *Proc. 28th ACM Annu. Sympos. Comput. Geom.*, 2012, 91–100.
- [13] J. S. Foster, E. Gjeldel, W. R. Graham, R. J. Hermann, H. M. Kluepfel, R. L. Lawson, G. K. Soper, L. L. Wood, and J. B. Woodard, *Report of the commission to assess the threat to the United States from electromagnetic pulse (EMP) attack, critical national infrastructures*, Apr. 2008.
- [14] R. L. Graham, D. E. Knuth, and O. Patashnik, *Concrete Mathematics: A Foundation for Computer Science* (2nd ed.), Addison Wesley, Boston, MA, 1994.
- [15] S. Har-Peled and B. Raichel, On the complexity of randomly weighted Voronoi diagrams, in *Proc. 30th Annu. ACM Sympos. Comput. Geom.*, 2014, 232–241.
- [16] S. Har-Peled and M. Sharir, Relative (p, ϵ) -approximations in geometry, *Discrete Comput. Geom.* 45 (2011), 462–496.

- [17] N. L. Johnson, A. W. Kemp, and S. Kotz, *Univariate Discrete Distributions (3rd ed.)*, Wiley & Sons, New York, 2005.
- [18] K. Kedem, R. Livne, J. Pach and M. Sharir, On the union of Jordan regions and collision-free translational motion amidst polygonal obstacles, *Discrete Comput. Geom.* 1 (1986), 59–71.
- [19] L. M. Lifshitz and S. M. Pizer, A multiresolution hierarchical approach to image segmentation based on intensity extrema, *IEEE Trans. Pattern Anal. Mach. Intell.* 12 (1990), 529–540.
- [20] J. Matoušek, *Lectures on Discrete Geometry*, Springer-Verlag, Heidelberg–Berlin, 2002.
- [21] S. Neumayer and E. Modiano, Network reliability with geographically correlated failures, in *Proc. 29th IEEE International Conference on Computer Communications*, 2010, 1658–1666.
- [22] S. Neumayer, G. Zussman, R. Cohen, and E. Modiano, Assessing the vulnerability of the fiber infrastructure to disasters, *IEEE/ACM Trans. Netw.* 19 (2011), 1610–1623.
- [23] C. Ou and B. Mukherjee, *Survivable Optical WDM Networks*, Springer-Verlag, Heidelberg–Berlin, 2005.
- [24] J. Pach and M. Sharir, On the boundary of the union of planar convex sets, *Discrete Comput. Geom.* 21 (1999), 321–328.
- [25] M. Sharir and P. K. Agarwal, *Davenport-Schinzel Sequences and Their Geometric Applications*, Cambridge University Press, Cambridge-New York-Melbourne, 1995.
- [26] D. A. Spielman and S. H. Teng, Smoothed analysis: An attempt to explain the behavior of algorithms in practise, *Comm. ACM* 52 (2009), 76–84.
- [27] W. Wu, B. Moran, J. Manton, and M. Zukerman, Topology design of undersea cables considering survivability under major disasters, in *Proc. International Conference on Advanced Information Networking and Applications Workshops*, 2009, 1154–1159.