de Berg, Mark; van Wordragen, Geert

Improved Bounds for Discrete Voronoi Games

Published in:
Algorithms and Data Structures - 18th International Symposium, WADS 2023, Proceedings

DOI:
10.1007/978-3-031-38906-1_20

Published: 01/01/2023

Document Version
Peer-reviewed accepted author manuscript, also known as Final accepted manuscript or Post-print

Please cite the original version:
de Berg, M., & van Wordragen, G. (2023). Improved Bounds for Discrete Voronoi Games. In P. Morin, & S. Suri (Eds.), *Algorithms and Data Structures - 18th International Symposium, WADS 2023, Proceedings* (pp. 291-308). (Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics); Vol. 14079 LNCS). Springer. https://doi.org/10.1007/978-3-031-38906-1_20

Improved Bounds for Discrete Voronoi Games*

Mark de Berg† Geert van Wordragen‡

December 14, 2023

Abstract

In the planar one-round discrete Voronoi game, two players $\mathcal{P}$ and $\mathcal{Q}$ compete over a set V of n voters represented by points in $\mathbb{R}^2$. First, $\mathcal{P}$ places a set P of k points, then $\mathcal{Q}$ places a set Q of ℓ points, and then each voter $v \in V$ is won by the player who has placed a point closest to v. It is well known that if $k = \ell = 1$, then $\mathcal{P}$ can always win $n/3$ voters and that this is worst-case optimal. We study the setting where $k > 1$ and $\ell = 1$. We present lower bounds on the number of voters that $\mathcal{P}$ can always win, which improve the existing bounds for all $k \geqslant 4$. As a by-product, we obtain improved bounds on small ε-nets for convex ranges for even numbers of points in general position.

1 Introduction

In the *discrete Voronoi game*, two players compete over a set V of n voters in $\mathbb{R}^d$. First, player $\mathcal{P}$ places a set P of k points, then player $\mathcal{Q}$ places a set Q of ℓ points disjoint from the points in P, and then each voter $v \in V$ is won by the player who has placed a point closest to v. In other words, each player wins the voters located in its Voronoi cells in the Voronoi diagram $\text{Vor}(P \cup Q)$. In case of ties, that is, when a voter v lies on the boundary between a Voronoi cell owned by $\mathcal{P}$ and a Voronoi cell owned by $\mathcal{Q}$, then v is won by player $\mathcal{P}$. Note that $\mathcal{P}$ first places all their k points and then $\mathcal{Q}$ places their ℓ points—hence, this is a *one-round Voronoi game*—and that k and ℓ need not be equal. The one-round discrete Voronoi game was introduced by Banik *et al.* [4].

There is also a version of the Voronoi game where the players compete over a continuous region [1, 9, 12]. For this version a multiple-round variant, where $k = \ell$ and the players place points alternatingly, has been studied as well. We will confine our discussion to the discrete one-round game.

*MdB is supported by the Dutch Research Council (NWO) through Gravitation grant NETWORKS-024.002.003.

†Department of Mathematics and Computer Science, TU Eindhoven, Eindhoven, the Netherlands, `M.T.d.Berg@tue.nl`

‡Department of Computer Science, Aalto University, Espoo, Finland, `Geert.vanWordragen@aalto.fi`

The discrete one-round Voronoi game for $k = \ell = 1$ is closely related to the concept of plurality points in spatial voting theory [14]. In this theory, there is a d-dimensional policy space, and voters are modelled as points indicating their preferred policies. A *plurality point* is then a proposed policy that would win at least $\lceil n/2 \rceil$ voters against any competing policy. Phrased in terms of Voronoi games, this means that $\mathcal{P}$ can place a single point that wins at least $\lceil n/2 \rceil$ voters against any single point placed by $\mathcal{Q}$. The discrete Voronoi game with $k > 1$ and $\ell = 1$ can be thought of as an election where a coalition of k parties is colluding against a single other party.

Another way to interpret Voronoi games is as a *competitive facility-location problem*, where two companies want to place facilities so as to attract as many customers as possible, where each customer will visit the nearest facility. Competitive facility location has not only been studied in a (discrete and continuous) spatial setting, but also in a graph-theoretic setting; see e.g. [3, 13, 16].

Previous work. The one-round discrete Voronoi game leads to interesting algorithmic as well as combinatorial problems.

The algorithmic problem is to compute an optimal set of locations for the players. More precisely, for player $\mathcal{P}$ the goal is to compute, given a set V of n voters, a set P of k points that wins a maximum number of voters under the assumption that player $\mathcal{Q}$ responds optimally. For player $\mathcal{Q}$ the goal is to compute, given a voter set V and a set P of points placed by $\mathcal{P}$, a set Q of ℓ points that wins as many voters from V as possible. These problems were studied in $\mathbb{R}^1$ by Banik *et al.* [4] for the case $k = \ell$. They showed that an optimal set for $\mathcal{P}$ can be computed in $O(n^{k-\lambda_k})$ time, for some $0 < \lambda_k < 1$, and that an optimal set for $\mathcal{Q}$ can be computed in $O(n)$ if the voters are given in sorted order. The former result was improved by De Berg *et al.* [7], who presented an algorithm with $O(k^4 n)$ running time. They also showed that in $\mathbb{R}^2$ the problem for $\mathcal{P}$ is Σ_2^P-hard. The problem for $\mathcal{P}$ in the special case $k = \ell = 1$, is equivalent to finding the so-called Tukey median of V. This can be done in $O(n^{d-1} + n \log n)$ time, as shown by Chan [10].

The combinatorial problem is to prove worst-case bounds on the number of voters that player $\mathcal{P}$ can win, assuming player $\mathcal{Q}$ responds optimally. Tight bounds are only known for $k = \ell = 1$, where Chawla *et al.* [11] showed the following: for any set V of n voters in $\mathbb{R}^d$, player $\mathcal{P}$ can win at least $\lceil n/(d+1) \rceil$ voters and at most $\lceil n/2 \rceil$ voters, and these bounds are tight. Situations where $\mathcal{P}$ can win $\lceil n/2 \rceil$ voters are particularly interesting, as these correspond to the existence of a plurality point in voting theory. The bounds just mentioned imply that a plurality point does not always exist. In fact, a plurality point only exists for certain very symmetric point sets, as shown by Wu *et al.* [17]. De Berg *et al.* [6] showed how to test in $O(n \log n)$ time if a voter set admits a plurality point.

The combinatorial problem for $k > 1$ and $\ell = 1$ was studied by Banik *et al.* [5]. Here player $\mathcal{P}$ will never be able to win more than $\left(1 - \frac{1}{2k}\right) n$ voters, because

$k = 1$	$k = 2$	$k = 3$	$k = 4$	$k = 5$	arbitrary k	reference
1/3	3/7	7/15	15/31	21/41	$1 - \frac{42}{k}$	Banik *et al.* [5]
			1/2	11/21	$1 - \frac{20\frac{5}{8}}{k}$	this paper

Table 1: Lower bounds on the fraction of voters that $\mathcal{P}$ can win on any voter set in $\mathbb{R}^2$, when $\mathcal{P}$ has k points and $\mathcal{Q}$ has a single point. The stated fraction of our method for arbitrary k is for $n \to \infty$; the precise bound is $\left(1 - \frac{20\frac{5}{8}}{k}\right) n - 6$.

player $\mathcal{Q}$ can always win at least half of the voters of the most crowded Voronoi cell in $\text{Vor}(P)$. Banik *et al.* [5] present two methods to derive lower bounds on the number of voters that $\mathcal{P}$ can always win. Below we discuss their results in $\mathbb{R}^2$, but we note that they generalize their methods to $\mathbb{R}^3$.

The first method uses a (weak) ε-net for convex ranges on the voter set V, that is, a point set N such that any convex range R containing at least εn voters, will also contain a point from N. Now, if $|Q| = 1$ then the voters won by $\mathcal{Q}$ lie in a single Voronoi cell in $\text{Vor}(P \cup Q)$. Since Voronoi cells are convex, this means that if we set $P := N$ then $\mathcal{P}$ wins at least $(1 - \varepsilon)n$ voters. Banik *et al.* use the ε-net construction for convex ranges by Mustafa and Ray [15]. There is no closed-form expression for the size of their ε-net, but the method can give a $(4/7)$-net of size 2, for instance, and an $(8/15)$-net of size 3. The smallest size for which they obtain an ε-net for some $\varepsilon \leqslant 1/2$, which corresponds to $\mathcal{P}$ winning at least half the voters, is $k = 5$. Banik *et al.* show that the ε-net of Mustafa and Ray can be constructed in $O(kn \log^4 n)$ time. The second method of Banik *et al.* uses an ε-net for disks, instead of convex sets. This is possible because one can show that a point $q \in Q$ that wins α voters, must have a disk around it that covers at least $\lfloor \alpha/6 \rfloor$ voters without containing a point from P. Banik *et al.* then present a $(7/k)$-net for disks of size k, which can be constructed in $O(n^2)$ time. This gives a method that ensures $\mathcal{P}$ wins at least $\left(1 - \frac{42}{k}\right) n$ voters, which is better than the first method when $k \geqslant 137$.

Our results. We study the combinatorial question—how many voters can player $\mathcal{P}$ win from any voter set V of size n, under optimal play from $\mathcal{Q}$—in the planar setting, for $k > 1$ and $\ell = 1$. We obtain the following results, where we assume that V is in general position—no three voters are collinear—and that n is even.

In Section 2 we present an improvement[1] over the ε-net bounds by Mustafa and Ray [15] for convex ranges. This improves the results of Banik *et al.* [5] on the fraction of voters that $\mathcal{P}$ can win when $k \geqslant 4$ and k is relatively small. We do not have a closed-form expression for the size of our ε-net as function of ε. Theorem 2 gives a recurrence on these sizes, and Table 1 shows how our bounds compare to those of Banik *et al.* for $k = 4, 5$ (which follow from the bounds of

[1]Our definition of ε-net is slightly weaker than usual, since a range missing the ε-net may contain up to $\lceil \varepsilon n \rceil$ points, instead of $\lfloor \varepsilon n \rfloor$ points.

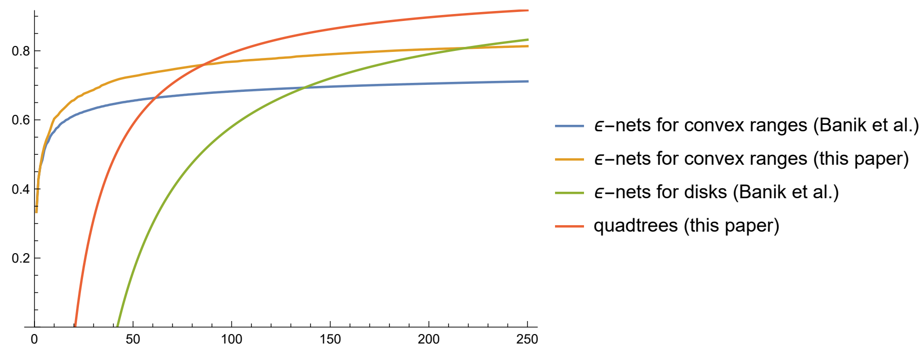

Figure 1: Lower bounds on the fraction of voters that $\mathcal{P}$ can win as a function of k (the number of points of $\mathcal{P}$) when $\mathcal{Q}$ has a single point, for the L_2-metric. The red and green graphs do not intersect, so for large k the quadtree method gives the best solution.

Mustafa and Ray [15]). It is particularly interesting that our bounds improve the smallest k for which $\mathcal{P}$ can win at least half the voters, from $k = 5$ to $k = 4$.

In Section 3 we present a new strategy for player $\mathcal{P}$. Unlike the strategies by Banik *et al.*, it is not based on ε-nets. Instead, it uses a quadtree-based approach. By combining this approach with several other ideas and using our ε-net method as a subroutine, we are able to show that there is a set P of k points that guarantees that $\mathcal{P}$ wins at least $\left(1 - \frac{20\frac{5}{8}}{k}\right) n - 6$ voters, which significantly improves the $\left(1 - \frac{42}{k}\right) n$ bound of Banik *et al.* Fig. 1 show the bounds obtained by the various methods in a graphical way.

We also study the discrete one-round Voronoi game in the L_1-metric, for $k > 1$ and $\ell = 1$. When $k = 1$, player $\mathcal{P}$ can win at least half the voters by placing a point on a multi-dimensional median, that is, a point whose x- and y-coordinate are medians among the x- and y-coordinates of the voter set V [6]. The case $k > 1$ and $\ell = 1$ has, as far as we know, not been studied so far. We first observe that for the L_1-metric, an ε-nets for axis-parallel rectangles can be used to obtain a good set of points for player $\mathcal{P}$. Using known results [2] this implies the results for $2 \leqslant k \leqslant 5$ in Table 1.

2 Better ε-nets for convex ranges

Below we present a new method to construct an ε-net for convex ranges in the plane, which improves the results of Mustafa and Ray [15]. As mentioned in the introduction, this implies improved bounds on the number of voters $\mathcal{P}$ can win with k points when $\mathcal{Q}$ has a single point, for relatively small values of k.

Let L be a set of three concurrent lines and consider the six wedges defined by the lines. Bukh [8] proved that for any continuous measure there is a choice of

L where each of the wedges has equal measure. Instead of a measure, we have V, a point set in general position in the plane, thus we need a generalisation where the wedges contain some specified number of points. In our generalisation, the *weight* of a wedge is given by the number of points from V assigned to it. If a point $v \in V$ lies in the interior of a wedge then we assign v to that wedge, and if v lies on the boundary of two or more wedges we assign v to one of them. (This assignment is not arbitrary, but we will do it is such a way as to obtain the desired number of points in each wedge.) We call this a *wedge assignment*.

Theorem 1. *Let V be a set of n points in general position in the plane, where $n \geqslant 8$ is even. For any given $\alpha, \beta, \gamma \in \mathbb{N}$ such that $2\alpha+2\beta+2\gamma = n$, we can find a set of three concurrent lines that partitions the plane into six wedges such that there is a wedge assignment resulting in wedges whose weights are $\alpha, \beta, \gamma, \alpha, \beta, \gamma$ in counterclockwise order.*

Proof. (Sketch.) Let $\ell(\theta)$ be the directed line making an angle θ with the positive x-axis and that has exactly weight $n/2$ on either side of it, for a suitable assignment of points to the half-planes on either side of $\ell(\theta)$. Consider the line $\ell(\theta)$ for $\theta = 0$. For some point $z = (x, 0) \in \ell(\theta)$, consider the rays $\rho_1, \ldots, \rho_4$ emanating from z, such that the six wedges defined by these rays and $\ell(0)$ have the desired number of voters; see Fig. 2. By varying θ and the point z, we can ensure that the rays $\rho_1, \ldots, \rho_4$ line up in such a way that, together with $\ell(\theta)$, they form three concurrent lines. □ □

We also need the following easy-to-prove observation.

Observation 1. Let L be a set of three lines intersecting in a common point p^*, and consider the six closed wedges defined by L. Any convex set S not containing p^* intersects at most four wedges, and the wedges intersected by S are consecutive in the clockwise order.

Notice that Observation 1 considers all *closed* wedges intersecting S. Thus, when S touches the boundary between two wedges, then both of them are taken into account. Hence, it will not be a problem that Theorem 1 could assign points on the boundary between two wedges to either one of them.

We now have all the tools to prove our new bounds on ε-nets for convex ranges. The guarantee they give is slightly weaker than usual: where ordinarily

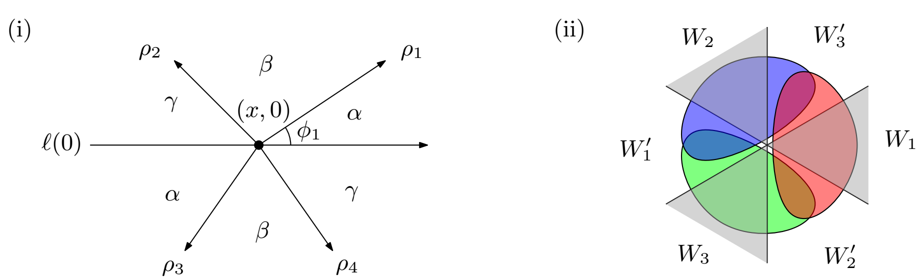

Figure 2: Illustrations for the proofs of Theorem 1 and Theorem 2.

placing an ε-net for n points means a range not intersecting the ε-net can contain at most εn (and thus at most $\lfloor \varepsilon n \rfloor$) points, our *ceiling-based* ε-nets only guarantee that such a range contains at most $\lceil \varepsilon n \rceil$ points.

Theorem 2. *Let ε_k be the smallest value such that any finite point set in $\mathbb{R}^2$ admits a weak ε_k-net of size k for convex ranges. Then for any set V of $n \geqslant 8$ points in general position, with n even, and any $r_1, r_2, s \in \mathbb{N}_0$, we can make a ceiling-based ε-net for V with*

$$\varepsilon = \frac{1}{2}\left(\frac{1}{\varepsilon_{r_1}} + \frac{2}{\varepsilon_{r_2}}\right)^{-1} + \frac{1}{2}\varepsilon_s.$$

Proof. (Sketch.) Let $\mu := \frac{1}{2}\left(\frac{1}{\varepsilon_{r_1}} + \frac{2}{\varepsilon_{r_2}}\right)^{-1}$. We apply Theorem 1 with $\beta = \gamma = \left\lceil \frac{\mu}{\varepsilon_{r_2}} n \right\rceil$ and $\alpha = \frac{n}{2} - 2\beta$, which means $\alpha \leqslant \left\lfloor \frac{\mu}{\varepsilon_{r_1}} n \right\rfloor$, giving us a set L of three concurrent lines. To show that there exists a weak ceiling-based $(\varepsilon_{r_1+2r_2+3s+1})$-net N for V, number the wedges defined by L as $W_1, W_3', W_2, W_1', W_3, W_2'$ in clockwise order, as in Fig. 2(ii). Let $V_i \subset V$ and $V_i' \subset V$ be the subsets of points assigned to W_i and W_i', respectively. We can assume without loss of generality that $|V_1| = |V_1'| = \alpha$, and $|V_2| = |V_2'| = \beta$, and $|V_3| = |V_3'| = \gamma$. We add the following points to our net N: (i) the common intersection of the lines in L, denoted by p^*; (ii) an ε_{r_1}-net for V_1, an ε_{r_2}-net for V_2, and an ε_{r_2}-net for V_3; (iii) for each of the three collections of three consecutive wedges—these collections are indicated in red, green, and blue in Fig. 2(ii)—an ε_s-net. By construction, the size of our net N is $1 + r_1 + 2r_2 + 3s$. In the full paper we show that N is a ceiling-based $(\mu + \frac{1}{2}\varepsilon_s)$-net. □ □

Note that $\varepsilon_0 = 1$, since if the net is empty, a range can contain all n points from V. Moreover, $\varepsilon_1 = 2/3$, and $\varepsilon_2 = 4/7$, and $\varepsilon_3 \leqslant 8/15$ by the results of Mustafa and Ray [15]. Using Theorem 2 we can then set up a recursion to obtain ceiling-based ε-nets with $k \geqslant 4$ points, by finding the best choice of r_1, r_2, s such that $k = r_1 + 2r_2 + 3s + 1$. This gives $\varepsilon_4 \leqslant \frac{1}{2}$, by setting $r_1, r_2 = 0$ and $s = 1$. Hence, for even n, player $\mathcal{P}$ can always place four points to win at least as many voters as player $\mathcal{Q}$, as opposed to the five that were proven in earlier work. Note that this also holds for $n \leqslant 8$, since then player $\mathcal{P}$ can simply pick four points coinciding with four of the at most eight voters. A similar statement holds for larger k when $n \leqslant 8$.

3 A quadtree-based strategy for player $\mathcal{P}$

The algorithm. First, we construct a *compressed quadtree* $\mathcal{T}$ on the voter set V. This gives a tree structure where each node ν is associated with a square or a donut. We will refer to the square or donut associated to a quadtree node ν as the *cell* of that node, and denote it by $\sigma(\nu)$. We assume that no voter in V lies on the boundary of a cell $\sigma(\nu)$, which can be ensured by picking the square corresponding to root($\mathcal{T}$) suitably. Donut cells in a compressed quadtree do not contain voters, and their corresponding nodes are leaves in the compressed quadtree. We denote the set of children of a node ν by $C(\nu)$. For a square quadtree cell σ, we denote its four quadrants by $\textsc{ne}(\sigma)$, $\textsc{se}(\sigma)$, $\textsc{sw}(\sigma)$, and $\textsc{nw}(\sigma)$.

We define the *size* of a square σ, denoted by $\mathit{size}(\sigma)$, to be its edge length. Let $\operatorname{dist}(\sigma_1, \sigma_2)$ denote the distance between the boundaries of two squares σ_1, σ_2. The distance between two quadtree cells satisfies the following property. Note that the property also holds when the cells are nested.

Observation 2. Let σ_1 and σ_2 be square cells corresponding to two nodes in $\mathcal{T}$ such that σ_1 and σ_2 are intersected by a common horizontal (or vertical) line. If $\operatorname{dist}(\sigma_1, \sigma_2) > 0$ then $\operatorname{dist}(\sigma_1, \sigma_2) \geqslant \min(\mathit{size}(\sigma_1), \mathit{size}(\sigma_2))$.

The idea of our algorithm to generate the k points played by player $\mathcal{P}$ is as follows. We pick a parameter m, which depends on k, and then we recursively traverse the tree $\mathcal{T}$ to generate a set $\mathcal{R}$ of regions, each containing between $m+1$ and $4m$ points. Each region $R(\nu) \in \mathcal{R}$ will be a quadtree cell $\sigma(\nu)$ minus the quadtree cells $\sigma(\mu)$ of certain nodes μ in the subtree rooted at ν. For each region $R \in \mathcal{R}$, we then generate a set of points that we put into P. The exact procedure to generate the set $\mathcal{R}$ of regions is described by Algorithm 1, which is called with $\nu = \operatorname{root}(\mathcal{T})$.

Algorithm 1 *MakeRegions*(ν, m)

Input: A node ν in $\mathcal{T}$ and a parameter m
Output: A pair $(\mathcal{R}, V_{\text{free}})$, where $\mathcal{R}$ is a set of regions containing at least $m+1$ and at most $4m$ voters, and V_{free} contains the voters in the subtree rooted at ν that are not yet covered by a region in $\mathcal{R}$.

```
if ν is a leaf node then
    Return (∅, {v}) if ν contains a voter v, and return (∅, ∅) otherwise
else
    Recursively call MakeRegions(μ, m) for all children μ ∈ C(ν). Let R
    be the union of the returned sets of regions, and let V_free be the union
    of the sets of returned free voters.
    if |V_free| ⩽ m then
        Return (R, V_free)
    else
        R(ν) ← σ(ν) \ ⋃_{R∈R} R                ▷ Note that V_free = R(ν) ∩ V.
        Return (R ∪ {R(ν)}, ∅)
```

We use the regions in $\mathcal{R}$ to place the points for player $\mathcal{P}$, as follows. For a region $R := R(\nu)$ in $\mathcal{R}$, define $\sigma(R) := \sigma(\nu)$ to be the cell of the node ν for which R was generated. For each $R \in \mathcal{R}$, player $\mathcal{P}$ will place a grid of 3×3 points inside $\sigma(R)$, plus four points outside $\sigma(R)$, as shown in Fig. 3(i). (Some points placed for R may coincide with points placed for some $R' \neq R$, but this

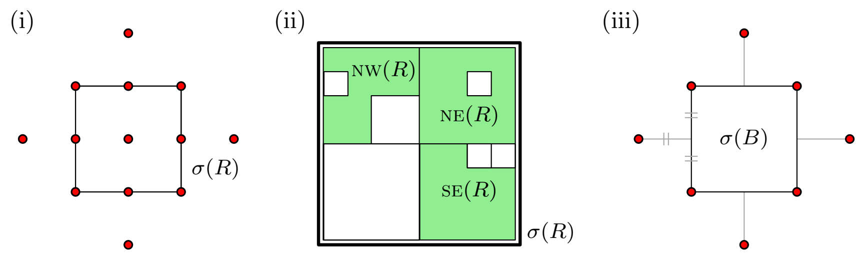

Figure 3: (i) The 13 points (in red) placed in P for a region $R \in \mathcal{R}$. A block B such that $\sigma(B)$ is one of the quadrants of $\sigma(R)$ is called a type-I block. (ii) A region R (shown in green) and its blocks (that is, its child regions). The white area is covered by regions that have been created earlier. Since $\textsc{sw}(\sigma(R))$ has already been fully covered, $\textsc{sw}(R)$ does not exist. (iii) The eight points placed in P for a type-II block $B \in \mathcal{B}$.

will only help to reduce the number of points placed by $\mathcal{P}$.)

Note that each $R \in \mathcal{R}$ contains more than m voters and the regions in $\mathcal{R}$ are disjoint. Hence, $|\mathcal{R}| < n/m$ and $|P| < 13n/m$. A compressed quadtree can be constructed in $\mathcal{O}(n \log n)$ time, and the rest of the construction takes $\mathcal{O}(n)$ time. The following lemma summarizes the construction.

Lemma 1. *The quadtree-based strategy described above places fewer than $13n/m$ points for player $\mathcal{P}$ and runs in $\mathcal{O}(n \log n)$ time.*

An analysis of the number of voters player $\mathcal{Q}$ can win. To analyze the number of voters that $\mathcal{Q}$ can win, it will be convenient to look at the "child regions"of the regions in $\mathcal{R}$, as defined next. Recall that for a region $R := R(\nu)$ in $\mathcal{R}$, we defined $\sigma(R) := \sigma(\nu)$. Let $\textsc{ne}(R) := R \cap \textsc{ne}(\sigma(R))$ be the part of R in the $\textsc{ne}$-quadrant of $\sigma(R)$. We call $\textsc{ne}(R)$ a *child region* of R. The child regions $\textsc{se}(R)$, $\textsc{sw}(R)$, and $\textsc{nw}(R)$ are defined similarly; see Fig 3(ii) for an example.

Let $\mathcal{B}$ be the set of non-empty child regions of the regions in $\mathcal{R}$. From now on, we will refer to the child regions in $\mathcal{B}$ as *blocks*. Blocks are not necessarily rectangles, and they can contains holes and even be disconnected. For a block $B \in \mathcal{B}$, we denote its parent region in $\mathcal{R}$ by $\mathrm{pa}(B)$, and we let $\sigma(B)$ denote the quadtree cell corresponding to B. For instance, if $B = \textsc{ne}(\mathrm{pa}(B))$ then $\sigma(B) = \textsc{ne}(\sigma(\mathrm{pa}(B)))$.

Note that at the end of Algorithm 1, the set V_{free} need not be empty. Thus the blocks in $\mathcal{B}$ may not cover all voters. Hence, we add a special *root block* B_0 to $\mathcal{B}$, with $\sigma(B_0) := \sigma(\mathrm{root}(\mathcal{T}))$ and which consists of the part of $\sigma(\mathrm{root}(\mathcal{T}))$ not covered by other blocks. Note that we do not add any points to P for B_0.

Because we will later refine our strategy for player $\mathcal{P}$, it will be convenient to analyze the number of voters that $\mathcal{Q}$ can win in an abstract setting. Our analysis requires the collection $\mathcal{B}$ of blocks and the set P of points played by $\mathcal{P}$ to have the following properties.

($\mathcal{B}$.1) The blocks in $\mathcal{B}$ are generated in a bottom-up manner using the compressed quadtree $\mathcal{T}$. More precisely, there is a collection $N(\mathcal{B})$ of nodes in $\mathcal{T}$ that is in one-to-one correspondence with the blocks in $\mathcal{B}$ such that the following holds:

> Let $B(\nu)$ be the block corresponding to a node $\nu \in N(\mathcal{B})$. Then $B(\nu) = \sigma(\nu) \setminus \bigcup_\mu B(\mu)$, where the union is taken over all nodes $\mu \in N(\mathcal{B})$ that are a descendant of ν.

We also require that the blocks in $\mathcal{B}$ together cover all voters.

($\mathcal{B}$.2) For each block $B \in \mathcal{B}$, except possibly the root block B_0, the point set P includes the 13 points shown in Fig. 3(i) for the cell that is the parent of $\sigma(B)$, or it includes the eight points shown in Fig. 3(iii). In the former case we call B a *type-I block*, in the latter case we call B a *type-II block*. Note that in both cases P includes the four corners of $\sigma(B)$.

Observe that ($\mathcal{B}$.1) implies that the blocks $B \in \mathcal{B}$ are disjoint. Moreover, property ($\mathcal{B}$.2) implies the following. For a square σ, define *plus*(σ) to be the plus-shaped region consisting of five equal-sized squares whose central square is σ.

Observation 3. Let q be a point played by player $\mathcal{Q}$ and let $B \in \mathcal{B}$ be a block. If q wins a voter v that lies in $\sigma(B)$, then $q \in$ *plus*$(\sigma(B))$. Furthermore, if $q \in \sigma(B)$ then q can only win voters in *plus*$(\sigma(B))$.

It is easy to see that the sets $\mathcal{B}$ and P generated by the construction described above have properties ($\mathcal{B}$.1) and ($\mathcal{B}$.2). We proceed to analyze the number of blocks from which a point q played by $\mathcal{Q}$ can win voters, assuming the set $\mathcal{B}$ of blocks has the properties stated above.

We will need the following observation. It follows from ($\mathcal{B}$.1), which states that a block B completely covers the part of $\sigma(B)$ not covered by blocks that have been created earlier in the bottom-up process.

Observation 4. If $\sigma(B) \subset \sigma(B')$ for two blocks $B, B' \in \mathcal{B}$ then $B' \cap \sigma(B) = \emptyset$.

The following lemma states that the set P of points played by player $\mathcal{P}$ includes all vertices of the blocks in $\mathcal{B}$, except possibly the corners of the root block B_0.

Lemma 2. *Let p be a vertex of a block $B \in \mathcal{B}$. Then $p \in P$, except possibly when p is a corner of $\sigma(B_0)$.*

Proof. Property ($\mathcal{B}$.1) states that the blocks in $\mathcal{B}$ are created in a bottom-up order. We will prove the lemma by induction on this (partial) order.

Consider a block $B \in \mathcal{B}$ and let p be a vertex of B. Let s be a sufficiently small square centered at p and let s_1, s_2, s_3, s_4 be its quadrants. There are two cases; see Fig. 4.

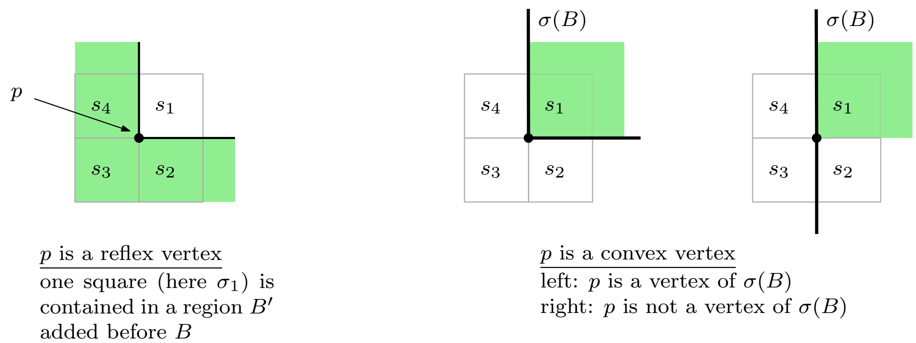

Figure 4: Illustration for the proof of Lemma 2.

If p is a reflex vertex of B, then B covers three of the four squares s_1, s_2, s_3, s_4. The remaining square must already have been covered by a region B' created before B, by Observation 4. By induction, we may conclude that $p \in P$.

If p is a convex vertex, then exactly one of the four squares s_1, s_2, s_3, s_4, say s_1, is contained in B. If p is a corner of $\sigma(B)$, then $p \in P$ by property ($\mathcal{B}$.2). Otherwise, at least one square $\sigma_i \neq \sigma_1$, say σ_2, is contained in $\sigma(B)$. We can now use the same argument as before: p is a vertex of a region B' created before B, and so $p \in P$ by induction. Note that this not only holds when p lies on an edge of $\sigma(B)$, as in Fig. 4, but also when p lies in the interior of $\sigma(B)$. □ □

Now consider a point q played by player $\mathcal{Q}$, and assume without loss of generality that $q \in \sigma(B_0)$. We first show that q can win voters from at most five blocks $B \in \mathcal{B}$; later we will improve this to at most three blocks. We may assume that the horizontal and vertical lines through q do not pass through a vertex of any block $B \in \mathcal{B}$. This is without loss of generality, because an infinitesimal perturbation of q ensures this property, while such a perturbation does not change which voters are won by q. (The latter is true because voters at equal distance from q and P are won by player $\mathcal{P}$.)

Define $B(q) \in \mathcal{B}$ to be the block containing q. We start by looking more closely at which voters q might win from a block $B \neq B(q)$. Define $V(B) := V \cap B$ to be the voters lying in B. Let ρ_{left} be the axis-aligned ray emanating from q and going the left, and define $\rho_{\text{up}}, \rho_{\text{right}}, \rho_{\text{down}}$ similarly. Let e be the first edge of B that is hit by ρ_{right} and define

$$V_{\text{right}}(B) := \{v \in V(B) : \ v \text{ lies in the horizontal half-strip whose left edge is } e\}.$$

Define the sets $V_{\text{up}}(B)$, $V_{\text{down}}(B)$, and $V_{\text{left}}(B)$ similarly. See Fig. 5(i), where the voters from $V_{\text{right}}(B)$ are shown in dark green, the voters from $V_{\text{up}}(B)$ and $V_{\text{down}}(B)$ are shown in orange and blue, respectively, and $V_{\text{left}}(B) = \emptyset$. Because P contains all vertices of B by Lemma 2, the only voters from $V(B)$ that q can possibly win are the voters in $V_{\text{left}}(B) \cup V_{\text{up}}(B) \cup V_{\text{right}}(B) \cup V_{\text{down}}(B)$. (In fact, we could restrict these four sets even a bit more, but this is not needed for our arguments.)

Let $B_{\text{right}} \neq B(q)$ be the first block in $\mathcal{B}$ hit by ρ_{right}, and define $B_{\text{left}}, B_{\text{up}}, B_{\text{down}}$ similarly for the rays $\rho_{\text{left}}, \rho_{\text{up}}, \rho_{\text{down}}$. The next lemma states that there is only one block $B \neq B(q)$ for which q might be able to win voters in $V_{\text{right}}(B)$, namely B_{right}. Similarly, q can only when voters from $V_{\text{left}}(B)$ for $B = B_{\text{left}}$, and so on.

Lemma 3. *If q wins voters from $V_{\text{right}}(B)$, where $B \neq B(q)$, then $B = B_{\text{right}}$.*

Proof. Suppose for a contradiction that q wins voters from $V_{\text{right}}(B)$ for some block $B \notin \{B(q), B_{\text{right}}\}$. We distinguish two cases.

Case I: $q \notin \sigma(B_{\text{right}})$.

Since the corners of $\sigma(B_{\text{right}})$ are in P by ($\mathcal{B}$.2), the point q cannot win voters to the right of $\sigma(B_{\text{right}})$. Hence, if q wins voters from $V_{\text{right}}(B)$, then B must lie at least partially inside $\sigma(B_{\text{right}})$. Now consider $\sigma(B)$. We cannot have $\sigma(B_{\text{right}}) \subset \sigma(B)$ by Observation 4. Hence, $\sigma(B) \subset \sigma(B_{\text{right}})$ and so $B_{\text{right}} \cap \sigma(B) = \emptyset$. We now have two subcases, illustrated in Fig. 5(ii).

- If the left edge of $\sigma(B)$ is contained in the left edge of $\sigma(B_{\text{right}})$, then ρ_{right} would hit B before B_{right}, contradicting the definition of B_{right}.
- On the other hand, if the left edge of $\sigma(B)$ is not contained in the left edge of $\sigma(B_{\text{right}})$, then $\text{dist}(\sigma(B_{\text{right}}), \sigma(B)) \geqslant \textit{size}(\sigma(B))$ by Observation 2. Since P contains the four corners of $\sigma(B)$, this contradicts that q wins voters from $V_{\text{right}}(B)$.

Case II: $q \in \sigma(B_{\text{right}})$.

We cannot have $\sigma(B_{\text{right}}) \subset \sigma(B(q))$, otherwise $B(q) \cap \sigma(B_{\text{right}}) = \emptyset$ by Observation 4, which contradicts $q \in B(q)$. Hence, $\sigma(B(q)) \subset \sigma(B_{\text{right}})$ and $B_{\text{right}} \cap \sigma(B(q)) = \emptyset$.

Consider the square σ with the same size of $\sigma(B(q))$ and immediately to the right of $\sigma(B(q)))$; see the grey square in Fig. 5(ii). We must have $\sigma \subset \sigma(B_{\text{right}})$, otherwise the right edge of $\sigma(B(q))$ would be contained in the right edge of $\sigma(B_{\text{right}})$ and so ρ_{right} would exit $\sigma(B_{\text{right}})$ before it can hit B_{right}. By Observation 3, point q cannot win voters to the right of σ. Hence, $B \cap \sigma \neq \emptyset$. Now consider the relative position of $\sigma(B)$ and σ. There are two subcases.

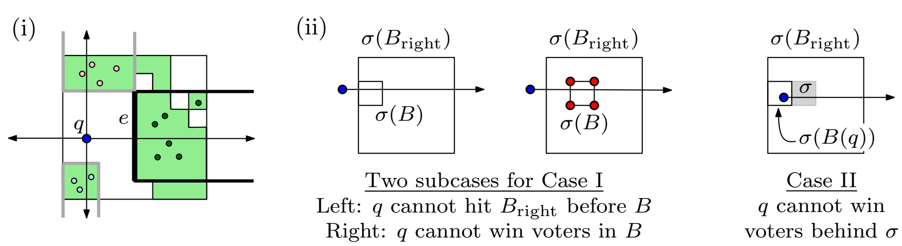

Figure 5: (i) The sets of voters that q might be able to win in the green block B. (ii) Illustration for the proof of Lemma 3.

- If $\sigma(B) \subset \sigma$, then either the distance from q to B is at least $size(\sigma(B))$ by Observation 2, contradicting that q wins voters from $V_{\text{right}}(B)$; or $\sigma(B)$ lies immediately to the right of $\sigma(B(q))$, in which case ρ_{right} cannot hit B_{right} before B.
- Otherwise, $\sigma \subset \sigma(B)$. If $\sigma(B) \subset \sigma(B_{\text{right}})$, then $B_{\text{right}} \cap \sigma(B) = \emptyset$ by Observation 4, contradicting (since $\sigma \subset \sigma(B)$) that ρ_{right} hits B_{right} before B. Hence, $\sigma(B_{\text{right}}) \subset \sigma(B)$. But then $B \cap \sigma(B_{\text{right}}) = \emptyset$, which contradicts that $B \cap \sigma \neq \emptyset$.

□ □

Lemma 3 implies that q can only win voters from the five blocks $B(q)$, B_{left}, B_{right}, B_{up}, and B_{down}. The next lemma shows that q cannot win voters from all these blocks simultaneously.

Lemma 4. *Point q can win voters from at most three of the blocks $B(q)$, B_{left}, B_{right}, B_{up}, and B_{down}.*

Proof. First suppose that the size of $\sigma(B(q))$ is at most the size of any of the four cells $\sigma(B_{\text{left}}), \ldots, \sigma(B_{\text{down}})$ from which q wins voters. By Observation 4, this implies that all four blocks $B_{\text{left}}, B_{\text{right}}, B_{\text{up}}, B_{\text{down}}$ lie outside $\sigma(B(q))$. Then it is easy to see that q can win voters from at most two of the four blocks $B_{\text{left}}, B_{\text{right}}, B_{\text{up}}, B_{\text{down}}$, because all four corners of $\sigma(B(q))$ are in P by ($\mathcal{B}$.2). For instance, if q lies in the NE-quadrant of $\sigma(B(q))$, then q can only win voters from B_{right} and B_{up}; the other cases are symmetrical.

Now suppose that $\sigma(B(q))$ is larger than $\sigma(B_{\text{right}})$, which we assume without loss of generality to be a smallest cell from which q wins voters among the four cells $\sigma(B_{\text{left}}), \ldots, \sigma(B_{\text{down}})$. We have two cases.

Case I: $q \notin \sigma(\text{pa}(B_{\text{right}}))$.

Note that $\sigma(B_{\text{right}})$ must either be the NW- or SW-quadrant of $\sigma(\text{pa}(B_{\text{right}}))$, because otherwise $q \notin plus(B_{\text{right}})$ and q cannot win voters from B_{right} by Observation 3. Assume without loss of generality that $\sigma(B_{\text{right}}) = \text{NW}(\sigma(\text{pa}(B_{\text{right}})))$. Then q must be located in the square of the same size as $\sigma(B_{\text{right}})$ and immediately to its left. In fact, q must lie in the right half of this square. We now define two blocks, σ and σ' that play a crucial role in the proof. Their definition depends on whether B_{right} is a type-I or a type-II block.

- If B_{right} is a type-I block, then we define $\sigma' := \sigma(B_{\text{right}})$ and we define σ to be the square of the same size as σ' and immediately to its left. Note that $q \in \sigma$, since q wins voters from $\sigma(B_{\text{right}})$. See Fig. 6(i).
- If B_{right} is a type-II block, then we define $\sigma' := \text{NW}(\sigma(B_{\text{right}}))$ or $\sigma' = \text{SW}(\sigma(B_{\text{right}}))$ and we define σ to be the square of the same size as σ' and immediately to its left. Whether we choose $\sigma' := \text{NW}(\sigma(B_{\text{right}}))$ or $\sigma' = \text{SW}(\sigma(B_{\text{right}}))$ depends on the position of q: the choice is made such that the square σ to the left of σ' contains q. See Fig. 6(ii) for an example. Since we will not use $\text{pa}(\text{pa}(\sigma'))$ in the proof, these two choices

are symmetric as far as the proof is concerned—we only need to swap the up- and down-direction.

We now continue with the proof of Case I. All statements referring to σ and σ' will hold for both definitions just given.

Observe that $\sigma(B_{\text{down}}) \neq \sigma$, since otherwise $\sigma(B_{\text{down}}) \subset \sigma(B(q))$, contradicting by Observation 4 that $q \in B(q)$. We will now consider three subcases. In each subcase we argue that either we are done—we will have shown that q wins voters from at most one of the blocks B_{down}, B_{left}, and B_{up}—or q cannot win voters from B_{down}. After discussing the three subcases, we then continue the proof under the assumption that q does not win voters from B_{down}.

- *Subcase (i): $\sigma(B_{\text{down}}) = \textsc{se}(\text{pa}(\sigma))$ and B_{down} is a type-I block.*
In the case all four corners of σ are in P. If $q \in \textsc{ne}(\sigma)$ then q can only win voters from B_{up} and if $q \in \textsc{se}(\sigma)$ then q can only win voters from B_{down} (this is in addition to voters won from $B(q)$ and B_{right}), and so we are done.

- *Subcase (ii): $\sigma(B_{\text{down}}) = \textsc{se}(\text{pa}(\sigma))$ and B_{down} is a type-II block.*
If $q \in \textsc{se}(\sigma)$ then q cannot win voters from B_{up} or B_{left}, and so we are done. Otherwise q cannot win voters from B_{down}, as claimed.

- *Subcase (iii): $\sigma(B_{\text{down}}) \neq \textsc{se}(\text{pa}(\sigma))$.*
If $B_{\text{down}} \cap \text{pa}(\sigma) \neq \emptyset$ then both B_{down} and $B(q)$ intersect $\text{pa}(\sigma)$, and both $\sigma(B_{\text{down}})$ and $\sigma(B(q))$ contain $\text{pa}(\sigma)$. But this is impossible due to Observation 4. Hence, B_{down} must lie below $\text{pa}(\sigma)$. We claim that then q cannot win voters from B_{down}. The closest q can be to B_{down} is when it lies on the bottom line segment of σ. Hence, any voter in B_{down} won by q must be closer to that segment than to p_3, and also than to p_2. (See Fig. 6 for the locations of p_2 and p_3.) But this is clearly impossible. Hence, q cannot win voters from B_{down}.

Thus, in the remainder of the proof for Case I we can assume that q does not win voters in B_{down}. Hence, it suffices to show that q cannot win voters from

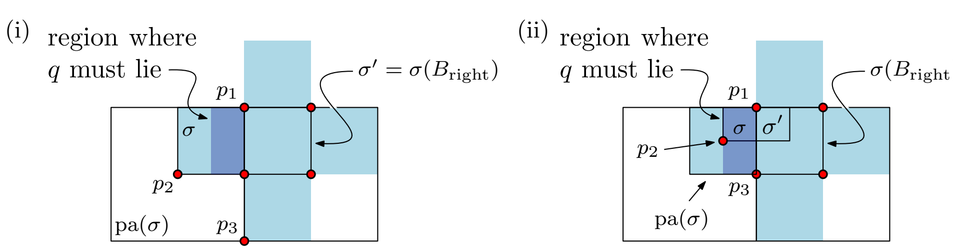

Figure 6: Two cases for the definition of σ' and σ, (i) when B_{right} is a type-I block and (ii) when B_{right} is a type-II block. In the latter case σ' and σ' could also lie in the bottom half of their parent regions, depending on where q lies.

B_{up} and B_{left} simultaneously. To this end, we assume q wins a voter v_{up} from B_{up} and a voter v_{left} from B_{left} and then derive a contradiction.

Let e_{up} be the first edge of B_{up} hit by ρ_{up} and let e_{left} be defined analogously; see Fig. 7. Note that e_{up} and e_{left} must lie outside $\mathrm{pa}(\sigma)$, otherwise we obtain a contradiction with Observation 4.

Let $\ell_{\mathrm{hor}}(q)$ be the horizontal line through q.

Claim. v_{left} must lie above $\ell_{\mathrm{hor}}(q)$.

Proof of Claim. We need to show that the perpendicular bisector of q and p_2 will always intersect the left edge of $\mathrm{pa}(\sigma)$ above $\ell_{\mathrm{hor}}(q)$. For the situation in Fig. 6(i) this is relatively easy to see, since q lies relatively far to the left compared to p_2. For situation in Fig. 6(ii), it follows from the following argument. To win voters in B_{right}, the point q must lie inside the circle C through p_1, p_2, p_3. Now, suppose q actually lies on C and let $\alpha := \angle qzp_2$, where z is the center of C. Thus the bisector of q and p_3 has slope $-\tan(\alpha/2)$. The result then follows from the fact that $2\tan(\alpha/2) - \sin\alpha > 0$ for $0 < \alpha < \pi/2$. $\square$

For v_{up} the situation is slightly different: q can, in fact, win voters to the right of $\ell_{\mathrm{vert}}(q)$, the vertical line through q. In that case, however, it cannot win v_{left}.

Claim. If q wins a voter from B_{up} right of $\ell_{\mathrm{vert}}(q)$, then q cannot win a voter from B_{left}.

Proof of Claim. It follows from Observation 2 that e_{up} must overlap with the top edge of σ. Because the edges e_{up} and e_{left} cannot intersect, one of them must end when or before the two meet.

If e_{left} ends before meeting (the extension of) e_{up}, then the top endpoint of e_{up}, which is in P by Lemma 2, prevents q from winning voters from B_{left}.

So now assume that e_{up} ends before meeting (the extension of) e_{left}. Then the left endpoint of e_{up}, which we denote by p_4, is in P. Without loss of generality, set $p_1 = (0, 1)$, $p_2 = (-1, 0)$ and $p_4 = (p_x, 1)$. Now, winning voters from B_{up} means q must lie inside the circle C_{up} with center on e_{up} that goes through p_1 and p_4. Thus it has center $c = (\frac{p_x}{2}, 0)$. It must also lie in the circle C through p_1, p_2, p_3 so it can win voters from B_{right}. The line through the circle centers makes an angle $\alpha := \arctan \frac{p_x}{2}$ with the line $x = 0$. The circles intersect at p_1, which lies on the line $x = 0$, so their other intersection point lies on the line ℓ that makes an angle 2α with $x = 0$. If q is to win voters from both B_{right}

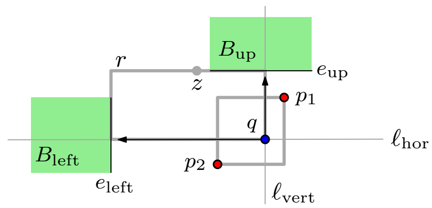

Figure 7: Definition of e_{left} and e_{up}, and r and z.

and B_{up} it must lie between ℓ and $x = 0$. Next we show that this implies that q cannot win voters from B_{left}.

We first show that if $q = \ell \cap C$, then p_4 prevents q from winning voters in B_{left}. Because then p_4 and q both lie on C_{up}, their perpendicular bisector $b(p_4, q)$ is the angular bisector of $\angle qcp_4$. Note that $\angle qcp_4 = 2\alpha$. Indeed, the line through the circle centers makes a right-angled triangle together with $y = 1$ and $x = 0$, so the angle at c must be $\frac{1}{2}\pi - \alpha$. Hence, $\angle qcp_1 = \pi - 2\alpha$, and so $\angle qcp_4 = 2\alpha$. Thus, $b(p_4, q)$ makes an angle α with $y = 1$ and so it intersects the line $x = -2$ at height $y = 1 - (2 + \frac{p_x}{2}) \tan\alpha$ which is $1 - (2 + \tan\alpha)\tan\alpha$. For $0 < 2\alpha < \pi/2$ this is below $\ell_{\mathrm{hor}}(q)$ which lies at $y = \cos 2\alpha$. By the previous Claim, this means that q cannot win voters from B_{left}. Therefore, q cannot win voters from B_{left}.

To finish the proof, we must argue that q cannot win voters from B_{left} either when $q \neq \ell \cap C$. It clear that moving q to the left helps to win voters in B_{left}, so we can assume that $q \in C$. Then it is not hard to see (by following the calculations above) that the best position for q is $\ell \cap C$, for which we just showed that q cannot win voters in B_{left}. This finishes the proof of the claim. □

We can now assume v_{left} lies above $\ell_{\mathrm{vert}}(q)$ and v_{up} lies to the left of $\ell_{\mathrm{vert}}(q)$. We will show that this leads to a contradiction. To this end, consider the rectangle r whose bottom-right corner is q, whose top edge overlaps e_{up} and whose left edge overlaps with e_{left}; see Fig. 7. Then the left edge of r contains the top endpoint of e_{left} and/or the top edge of r contains the left endpoint of e_{up}. By Lemma 2, we thus know that there is a point $p_4 \in P$ lying on the left or top edge of r. Now assume without loss of generality that the top edge of r is at least as long as its left edge, and let $z \in e_{\mathrm{up}}$ be the point such that the $q_x - z_x = z_y - q_y$. Now, if p_4 lies on the left edge of r or to the left of z on the top edge, then p_4 prevents q from winning v_{left}. On the other hand, if p_4 lies to the right of z on the top edge of r, then p_4 prevents q from winning v_{up}. So in both cases we have a contradiction.

Case II: $q \in \sigma(\mathrm{pa}(B_{\mathrm{right}}))$.

Assume without loss of generality that $\sigma(B_{\mathrm{right}})$ is one of the two northern quadrants of $\mathrm{pa}(\sigma(B_{\mathrm{right}}))$. We cannot have $q \in \sigma(B_{\mathrm{right}})$, since together with $\mathit{size}(\sigma(B_{\mathrm{right}})) < \mathit{size}(\sigma(B(q)))$ this contradicts that $q \in B(q)$, by Observation 4. Hence, $q \in \textsc{nw}(\mathrm{pa}(\sigma(B_{\mathrm{right}})))$ and $\sigma(B_{\mathrm{right}}) = \textsc{ne}(\mathrm{pa}(\sigma(B_{\mathrm{right}})))$.

If B_{right} is a type-I block then all corners of $\textsc{nw}(\mathrm{pa}(\sigma(B_{\mathrm{right}})))$ are in P, which (as we saw earlier) implies that q can win voters from at most three blocks. If B_{right} is a type-II block, then we can follow the proof of Case I. (For type-I blocks this is not true. The reason is that in the proof of the first Claim, we use that B_{left} does not lie immediately to the left of σ, which is not true for type-I blocks in Case 2. Note that this still is true for type-II blocks in Case 2.) □ □

By construction, each block contains at most $m < n/|\mathcal{R}|$ voters, where $\mathcal{R}$ is the set of regions created by Algorithm 1. Moreover, $\mathcal{P}$ places 13 points per region in $\mathcal{R}$, and so $k \leqslant 13|\mathcal{R}|$ points in total. Finally, Lemma 4 states that $\mathcal{Q}$ can win voters from at most three blocks. We can conclude the following.

Lemma 5. *Let V be a set of n voters in $\mathbb{R}^2$. For any given k, the quadtree-based strategy described above can guarantee that $\mathcal{P}$ wins at least $\left(1 - \frac{39}{k}\right) n$ voters by placing at most k points, against any single point placed by player $\mathcal{Q}$.*

A more refined strategy for player $\mathcal{P}$. It can be shown that the analysis presented above is tight. Hence, to get a better bound we need a better strategy.

Recall that each region $R \in \mathcal{R}$ contains between $m + 1$ and $4m$ voters. Currently, we use the same 13 points for any R, regardless of the exact number of voters it contains and how they are distributed over the child regions of R. Our refined strategy takes this into account, and also incorporates the ε-nets developed in the previous section, as follows. Let n_R denote the number of voters in a region $R \in \mathcal{R}$. We consider two cases, with several subcases.

- *Case A: $m < n_R \leqslant \frac{16}{11}m$.* We place eight points in total for R, as in Fig. 3(iii). We also add between two and six extra points, depending on the subcase.

  - If $m < n_R \leqslant \frac{7}{6}m$, we add two extra points, forming a $\frac{4}{7}$-net.
  - If $\frac{7}{6}m < n_R \leqslant \frac{5}{4}m$, we add three extra points, forming a $\frac{8}{15}$-net.
  - If $\frac{5}{4}m < n_R \leqslant \frac{4}{3}m$, we add four extra points, forming a $\frac{1}{2}$-net.
  - If $\frac{4}{3}m < n_R \leqslant \frac{7}{5}m$, we add five extra points, forming a $\frac{10}{21}$-net.
  - If $\frac{7}{5}m < n_R \leqslant \frac{16}{11}m$, we add six extra points, forming a $\frac{11}{24}$-net.

  One can show that in each subcase above, player $\mathcal{Q}$ wins at most $\lceil 2m/3 \rceil + 1$ voters from inside R, due to the (ceiling-based) ε-nets. For example, in the first case $\mathcal{Q}$ wins at most $(7m/6) \cdot (4/7) = 2m/3$ voters, in the second subcase $\mathcal{Q}$ wins at most $(5m/4) \cdot (8/15) = 2m/3$ voters, etcetera. Furthermore, one easily verifies that in each subcase we have $\frac{\text{number of voters in } R}{\text{number of points placed}} > m/10$.

- *Case B: $\frac{16}{11}m < n_R \leqslant 4m$.* We first place the same set of 13 points as in our original strategy. We add two or four extra points, depending on the subcase.

  - If $\frac{16}{11}m < n_R \leqslant 2m$ we add two extra points, as follows. Consider the four child regions of R. Then we add a centerpoint—in other words, a $\frac{2}{3}$-net of size 1—for the voters in the two child regions with the largest number of voters.
  - If $2m < n_R \leqslant 4m$ we add four extra points, namely a centerpoint for each of the four child regions of R.

Note that in both subcases, $\mathcal{Q}$ wins at most $2m/3$ voters from any child region. For the child regions where we placed a centerpoint, this holds because a child region contains at most m voters by construction. For the two child regions where we did not place a centerpoint in the first subcase, this holds because these child regions contains at most $2m/3$ voters. Furthermore, in both subcases $\frac{\text{number of voters in } R}{\text{number of points placed}} > \frac{16}{165}m$.

Lemma 6. *Let V be a set of n voters in $\mathbb{R}^2$. For any given k, the refined quadtree-based strategy can guarantee that $\mathcal{P}$ wins at least $\left(1 - \frac{20\frac{5}{8}}{k}\right)n - 6$ voters by placing at most k points, against any single point placed by player $\mathcal{Q}$.*

Proof. The proof for the original quadtree-based strategy was based on two facts: First, player $\mathcal{Q}$ can win voters from at most three blocks $B \in \mathcal{B}$; see Lemma 4. Second, any block $B \in \mathcal{B}$ (which was a child region of some $R \in \mathcal{R}$) contains at most m voters.

In the refined strategy, we use a similar argument, but for a set $\mathcal{B}_{\text{new}}$ of blocks defined as follows. For the regions $R \in \mathcal{R}$ that fall into Case A, we put R itself (instead of its child regions) as a type-II block into $\mathcal{B}_{\text{new}}$. For the regions $R \in \mathcal{R}$ that fall into Case B, we put their child regions as type-I blocks into $\mathcal{B}_{\text{new}}$. By Lemma 4, $\mathcal{Q}$ can win voters from at most three blocks in $\mathcal{B}_{\text{new}}$. Moreover, our refined strategy ensures that $\mathcal{Q}$ wins at most $\lceil 2m/3 \rceil + 1$ voters from any $B \in \mathcal{B}_{\text{new}}$. Thus $\mathcal{Q}$ wins at most $2m + 6$ voters in total.

Finally, for each $R \in \mathcal{R}$ we have $\frac{\text{number of voters in } R}{\text{number of points placed}} > \frac{16}{165}m$. Hence, $m < \frac{165}{16k}n$ and so $\mathcal{Q}$ wins at most $\frac{165}{8k}n + 6 = \frac{20\frac{5}{8}}{k}n + 6$ voters. □ □

4 Conclusion

We studied the discrete one-round Voronoi game where player $\mathcal{P}$ can place $k > 1$ points and player $\mathcal{Q}$ can place a single point. We improved the existing bounds on the number of voters player $\mathcal{P}$ can win. For small k this was done by proving new bounds on ε-nets for convex ranges. For large k we used a quadtree-based approach, which uses the ε-nets as a subroutine. The main open problem is: Can player $\mathcal{P}$ always win at least half the voters in the L_2-metric by placing less than four points?

The discrete one-round Voronoi game can also be studied in the L_1-metric, instead of in the L_2-metric as we did here. Our quadtree-based strategy also works well in the version of the problem, as we show in the full version of our paper.

References

[1] Ahn, H., Cheng, S., Cheong, O., Golin, M.J., van Oostrum, R.: Competitive facility location along a highway. In: Proc. 7th Annual International Conference on Computing and Combinatorics (COCOON 2001).

Lecture Notes in Computer Science, vol. 2108, pp. 237–246 (2001). https://doi.org/10.1007/3-540-44679-6_26

[2] Aronov, B., Aurenhammer, F., Hurtado, F., Langerman, S., Rappaport, D., Seara, C., Smorodinsky, S.: Small weak epsilon-nets. Comput. Geom. **42**(5), 455–462 (2009). https://doi.org/10.1016/j.comgeo.2008.02.005

[3] Bandyapadhyay, S., Banik, A., Das, S., Sarkar, H.: Voronoi game on graphs. Theoretical Computer Science **562**, 270–282 (2015). https://doi.org/https://doi.org/10.1016/j.tcs.2014.10.003

[4] Banik, A., Bhattacharya, B.B., Das, S.: Optimal strategies for the one-round discrete Voronoi game on a line. J. Comb. Optim. **26**(4), 655–669 (2013). https://doi.org/10.1007/s10878-011-9447-6

[5] Banik, A., Carufel, J.D., Maheshwari, A., Smid, M.H.M.: Discrete Voronoi games and ε-nets, in two and three dimensions. Comput. Geom. **55**, 41–58 (2016). https://doi.org/10.1016/j.comgeo.2016.02.002

[6] Berg, M.d., Gudmundsson, J., Mehr, M.: Faster algorithms for computing plurality points. ACM Trans. Algorithms **14**(3), 36:1–36:23 (2018). https://doi.org/10.1145/3186990

[7] Berg, M.d., Kisfaludi-Bak, S., Mehr, M.: On one-round discrete Voronoi games. In: Proc. 30th International Symposium on Algorithms and Computation (ISAAC 2019). LIPIcs, vol. 149, pp. 37:1–37:17 (2019). https://doi.org/10.4230/LIPIcs.ISAAC.2019.37

[8] Bukh, B.: A point in many triangles. Electron. J. Comb. **13**(1) (2006), http://www.combinatorics.org/Volume_13/Abstracts/v13i1n10.html

[9] Byrne, T., Fekete, S.P., Kalcsics, J., Kleist, L.: Competitive location problems: Balanced facility location and the one-round manhattan Voronoi game. In: Proc. 15th International Conference on Algorithms and Computation (WALCOM 2021). Lecture Notes in Computer Science, vol. 12635, pp. 103–115 (2021). https://doi.org/10.1007/978-3-030-68211-8_9

[10] Chan, T.M.: An optimal randomized algorithm for maximum tukey depth. In: Munro, J.I. (ed.) Proc. 15th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2004). pp. 430–436 (2004), http://dl.acm.org/citation.cfm?id=982792.982853

[11] Chawla, S., Rajan, U., Ravi, R., Sinha, A.: Min-max payoffs in a two-player location game. Oper. Res. Lett. **34**(5), 499–507 (2006). https://doi.org/10.1016/j.orl.2005.10.002

[12] Cheong, O., Har-Peled, S., Linial, N., Matousek, J.: The one-round Voronoi game. Discret. Comput. Geom. **31**(1), 125–138 (2004). https://doi.org/10.1007/s00454-003-2951-4, https://doi.org/10.1007/s00454-003-2951-4

[13] Dürr, C., Thang, N.K.: Nash equilibria in Voronoi games on graphs. In: Algorithms – ESA 2007. pp. 17–28. Springer Berlin Heidelberg, Berlin, Heidelberg (2007)

[14] McKelvey, R.D., Wendell, R.E.: Voting equilibria in multidimensional choice spaces. Mathematics of Operations Research **1**(2), 144–158 (1976)

[15] Mustafa, N.H., Ray, S.: An optimal extension of the centerpoint theorem. Comput. Geom. **42**, 505–510 (2009). https://doi.org/10.1016/j.comgeo.2007.10.004

[16] Teramoto, S., Demaine, E.D., Uehara, R.: Voronoi game on graphs and its complexity. In: 2006 IEEE Symposium on Computational Intelligence and Games. pp. 265–271 (2006). https://doi.org/10.1109/CIG.2006.311711

[17] Wu, Y., Lin, W., Wang, H., Chao, K.: Computing plurality points and condorcet points in euclidean space. In: Proc. 24th International Symposium on Algorithms and Computation (ISAAC 2013). Lecture Notes in Computer Science, vol. 8283, pp. 688–698. Springer (2013). https://doi.org/10.1007/978-3-642-45030-3_64