

On one-round discrete voronoi games

Citation for published version (APA):

de Berg, M. T., Kisfaludi-Bak, S., & Mehr, M. (2019). On one-round discrete voronoi games. *arXiv*, [1902.09234v1].

Document status and date:

Published: 25/02/2019

Document Version:

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

Please check the document version of this publication:

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

[Link to publication](#)

General rights

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

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

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

www.tue.nl/taverne

Take down policy

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

openaccess@tue.nl

providing details and we will investigate your claim.

On One-Round Discrete Voronoi Games*

Mark de Berg[§]

Sándor Kisfaludi-Bak[§]

Mehran Mehr[§]

Abstract

Let V be a multiset of n points in \mathbb{R}^d , which we call voters, and let $k \geq 1$ and $\ell \geq 1$ be two given constants. We consider the following game, where two players \mathcal{P} and \mathcal{Q} compete over the voters in V : First, player \mathcal{P} selects k points in \mathbb{R}^d , and then player \mathcal{Q} selects ℓ points in \mathbb{R}^d . Player \mathcal{P} wins a voter $v \in V$ iff $\text{dist}(v, P) \leq \text{dist}(v, Q)$, where $\text{dist}(v, P) := \min_{p \in P} \text{dist}(v, p)$ and $\text{dist}(v, Q)$ is defined similarly. Player \mathcal{P} wins the game if he wins at least half the voters. The algorithmic problem we study is the following: given V , k , and ℓ , how efficiently can we decide if player \mathcal{P} has a winning strategy, that is, if \mathcal{P} can select his k points such that he wins the game no matter where \mathcal{Q} places her points.

Banik *et al.* devised a singly-exponential algorithm for the game in \mathbb{R}^1 , for the case $k = \ell$. We improve their result by presenting the first polynomial-time algorithm for the game in \mathbb{R}^1 . Our algorithm can handle arbitrary values of k and ℓ . We also show that if $d \geq 2$, deciding if player \mathcal{P} has a winning strategy is Σ_2^P -hard when k and ℓ are part of the input. Finally, we prove that for any dimension d , the problem is contained in the complexity class $\exists\forall\mathbb{R}$, and we give an algorithm that works in polynomial time for fixed k and ℓ .

1 Introduction

Voronoi games, as introduced by Ahn *et al.* [1], can be viewed as competitive facility-location problems in which two players \mathcal{P} and \mathcal{Q} want to place their facilities in order to maximize their market area. The Voronoi game of Ahn *et al.* is played in a bounded region $R \subset \mathbb{R}^2$, and the facilities of the players are modeled as points in this region. Each player gets the same number, k , of facilities, which they have to place alternately. The market area of \mathcal{P} (and similarly of \mathcal{Q}) is now given by the area of the region of all points $q \in R$ whose closest facility was placed by \mathcal{P} , that is, it is the total area of the Voronoi cells of \mathcal{P} 's facilities in the Voronoi diagram of the facilities of \mathcal{P} and \mathcal{Q} . Ahn *et al.* proved that for $k > 1$ and when the region R is a circle or a segment, the second player can win the game by a payoff of $1/2 + \varepsilon$, for some $\varepsilon > 0$, where the first player can ensure ε is arbitrarily small.

The one-round Voronoi game introduced by Cheong *et al.* [8] is similar to the Voronoi game of Ahn *et al.*, except that the first player must first place all his k facilities, after which the second player places all her k facilities. They considered the problem where R is a square, and they showed that when k is large enough the first player can always win a fraction $1/2 + \alpha$ of the area of R for some $\alpha > 0$. Fekete and Meijer [11] considered the problem on a rectangle R of aspect ratio $\rho \leq 1$. They showed that the first player wins more than half the area of R , unless $k \geq 3$ and $\rho > \sqrt{2}/n$, or $k = 2$ and $\rho > \sqrt{3}/2$. They also showed that if R is a polygon with holes, then computing the

*This work was partially supported by the Netherlands Organization for Scientific Research NWO under project no. 024.002.003.

[§]Department of Mathematics and Computer Science, TU Eindhoven, the Netherlands. Email: mdberg@win.tue.nl, s.kisfaludi.bak@tue.nl, mehran.mehr@gmail.com

locations of the facilities for the second player that maximize the area she wins, against a given set of facilities of the first player is NP-hard.

One-round discrete Voronoi games. In this paper we are interested in *discrete (Euclidean) one-round Voronoi games*, where the players do not compete for area but for a discrete set of points. That is, instead of the region R one is given a set V of n points in a geometric space, and a point $v \in V$ is won by the player owning the facility closest to v . (Another discrete variant of Voronoi games is played on graphs [17, 19] but we restrict our attention to the geometric variant.) More formally, the problem we study is defined as follows.

Let V be a multiset of n points in \mathbb{R}^d , which we call *voters* from now on, and let $k \geq 1$ and $\ell \geq 1$ be two integers. The one-round discrete Voronoi game defined by the triple $\langle V, k, \ell \rangle$ is a single-turn game played between two players \mathcal{P} and \mathcal{Q} . First, player \mathcal{P} places a set P of k points in \mathbb{R}^d , then player \mathcal{Q} places a set Q of ℓ points in \mathbb{R}^d . (These points may coincide with the voters in V .) We call the set P the *strategy of \mathcal{P}* and the set Q the *strategy of \mathcal{Q}* . Player \mathcal{P} wins a voter $v \in V$ if $\text{dist}(v, P) \leq \text{dist}(v, Q)$, where $\text{dist}(v, P)$ and $\text{dist}(v, Q)$ denote the minimum distance between a voter v and the sets P and Q , respectively. Note that this definition favors player \mathcal{P} , since in case of a tie a voter is won by \mathcal{P} . We now define $V[P \succeq Q] := \{v \in V : \text{dist}(v, P) \leq \text{dist}(v, Q)\}$ to be the multiset of voters won by player \mathcal{P} when he uses strategy P and player \mathcal{Q} uses strategy Q . Player \mathcal{P} wins the game $\langle V, k, \ell \rangle$ if he wins at least half the voters in V , that is, when $|V[P \succeq Q]| \geq n/2$; otherwise \mathcal{Q} wins the game. Here $|V[P \succeq Q]|$ denotes the size of the multiset $V[P \succeq Q]$ (counting multiplicities). We now define $\Gamma_{k,\ell}(V)$ as the maximum number of voters that can be won by player \mathcal{P} against an optimal opponent:

$$\Gamma_{k,\ell}(V) := \max_{P \subset \mathbb{R}^d, |P|=k} \min_{Q \subset \mathbb{R}^d, |Q|=\ell} |V[P \succeq Q]|.$$

For a given multiset V of voters, we want to decide if¹ $\Gamma_{k,\ell}(V) \geq n/2$. In other words, we are interested in determining for a given game $\langle V, k, \ell \rangle$ if \mathcal{P} has a *winning strategy*, which is a set of k points such that \mathcal{P} wins the game no matter where \mathcal{Q} places her points.

An important special case, which has already been studied in spatial voting theory for a long time, is when $k = \ell = 1$ [14]. Here the coordinates of a point in V represent the preference of the voter on certain topics, and the point played by \mathcal{Q} represents a certain proposal. If the point played by \mathcal{P} wins against all possible points played by \mathcal{Q} , then the \mathcal{P} 's proposal will win the vote against any other proposal. Note that in the problem definition we gave above, voters at equal distance from P and Q are won by \mathcal{P} , and \mathcal{P} has to win at least half the voters. This is the definition typically used in papers of Voronoi games [3, 4, 5, 6]. In voting theory other variants are studied as well, for instance where points at equal distance to P and Q are not won by either of them, and \mathcal{P} wins the game if he wins more voters than \mathcal{Q} ; see the paper by McKelvey and Wendell [14] who use the term *majority points* for the former variant and the term *plurality points* for the latter variant.

Previous work. Besides algorithmic problems concerning the one-round discrete Voronoi game one can also consider combinatorial problems. In particular, one can ask for bounds on $\Gamma_{k,\ell}(V)$ as a function of n , k , and ℓ . It is known that for any set V in \mathbb{R}^2 and $k = \ell = 1$ we have $\lfloor n/3 \rfloor \leq \Gamma_{1,1}(V) \leq \lceil n/2 \rceil$. This result is based on known bounds for maximum Tukey depth, where the lower bound can be proven using Helly's theorem. It is also known [6] that there is a constant c such that $k = c\ell$ points suffice for \mathcal{P} to win the game, that is, $\Gamma_{c\ell,\ell}(V) \geq n/2$ for any V .

In this paper we focus on the algorithmic problem of computing $\Gamma_{k,\ell}(V)$ for given V , k , and ℓ . The problem of deciding if $\Gamma_{k,\ell}(V) \geq n/2$ was studied for the case $k = \ell = 1$ by Wu *et al.* [20] and

¹One can also require that $\Gamma_{k,\ell}(V) > n/2$; with some small modifications, all the results in this paper can be applied to the case with strict inequality as well.

Lin *et al.* [13], and later by De Berg *et al.* [9] who solve this problem in $O(n \log n)$ time in any fixed dimension d . Their algorithms work when V is a set (not a multiset) and for plurality points instead of majority points. Other algorithmic results are for the setting where the players already placed all but one of their points, and one wants to compute the best locations for the last point of \mathcal{P} and of \mathcal{Q} . Banik *et al.* [5] gave algorithms that find the best location for \mathcal{P} in $O(n^8)$ time and for \mathcal{Q} in $O(n^2)$ time. For the two-round variant of the problem, with $k = \ell = 2$, polynomial algorithms for finding the optimal strategies of both players are also known [4].

Our work is inspired by the paper of Banik *et al.* [3] on computing $\Gamma_{k,\ell}(V)$ in \mathbb{R}^1 . They considered the case of arbitrarily large k and ℓ , but where $k = \ell$ (and V is a set instead of a multiset). For this case they showed that depending on the set V either \mathcal{P} or \mathcal{Q} can win the game, and they presented an algorithm to compute $\Gamma_{k,\ell}(V)$ in time $O(n^{k-\lambda_k})$, where $0 < \lambda_k < 1$ is a constant dependent only on k . This raises the question: is the problem NP-hard when k is part of the input?

Our results. We answer the question above negatively, by presenting an algorithm that computes $\Gamma_{k,\ell}(V)$ in \mathbb{R}^1 in polynomial time. Our algorithm works when V is a multiset, and it does not require k and ℓ to be equal. Our algorithm computes $\Gamma_{k,\ell}(V)$ and finds a strategy for \mathcal{P} that wins this many voters in time $O(kn^4)$. The algorithm can be extended to the case when the voters are weighted, requiring only a slight increase in running time.

The algorithm by Banik *et al.* [3] discretizes the problem, by defining a finite set of potential locations for \mathcal{P} to place his points. However, to ensure an optimal strategy for \mathcal{P} , the set of potential locations has exponential size. To overcome this problem we need several new ideas. First of all, we essentially partition the possible strategies into various classes—the concept of thresholds introduced later plays this role—such that for each class we can anticipate the behavior of the optimal strategy for \mathcal{Q} . To compute the best strategy within a certain class we use dynamic programming, in a non-standard (and, unfortunately, rather complicated) way. The subproblems in our dynamic-program are for smaller point sets and smaller values of k and ℓ (actually we will need several other parameters) where the goal of \mathcal{P} will be to push his rightmost point as far to the right as possible to win a certain number of points. One complication in the dynamic program is that it is unclear which small subproblems I' can be used to solve a given subproblem I . The opposite direction—determining for I' which larger subproblems I may use I' in their solution—is easier however, so we use a sweep approach: when the solution to some I' has been determined, we update the solution to larger subproblems I that can use I' .

After establishing that we can compute $\Gamma_{k,\ell}(V)$ in polynomial time in \mathbb{R}^1 , we turn to the higher-dimensional problem. We show that deciding if \mathcal{P} has a winning strategy is Σ_2^P -hard in \mathbb{R}^2 . We also show that for fixed k and ℓ this problem can be solved in polynomial time. Our solution combines algebraic methods [7] with a result of Paterson and Zwick [15] that one can construct a polynomial-size boolean circuits that implements the majority function. The latter result is essential to avoid the appearance of n in the exponent. As a byproduct of the algebraic method, we show that the problem is contained in the complexity class $\exists\forall\mathbb{R}$; see [10] for more information on this complexity class.

2 A polynomial-time algorithm for $d = 1$

In this section we present a polynomial-time algorithm for the 1-dimensional discrete Voronoi game. Our algorithm will employ dynamic programming, and it will be convenient to use n , k , and ℓ as variables in the dynamic program. From now on, we therefore use n^* for the size of the original multiset V , and k^* and ℓ^* for the initial number of points that can be played by \mathcal{P} and \mathcal{Q} , respectively.

2.1 Notation and basic properties

We denote the given multiset of voters by $V := \{v_1, \dots, v_{n^*}\}$, where we assume the voters are numbered from left to right. We also always number the points in the strategies $P := \{p_1, \dots, p_{k^*}\}$ and $Q := \{q_1, \dots, q_{\ell^*}\}$ from left to right. For brevity we make no distinction between a point and its value (that is, its x -coordinate), so that we can for example write $p_1 < q_1$ to indicate that the leftmost point of P is located to the left of the leftmost point of Q .

For a given game $\langle V, k, \ell \rangle$, we say that a strategy P of player \mathcal{P} *realizes* a gain γ if $|V[P \succeq Q]| \geq \gamma$ for any strategy Q of player \mathcal{Q} . Furthermore, we say that a strategy P is *optimal* if it realizes $\Gamma_{k,\ell}(V)$, the maximum possible gain for \mathcal{P} , and we say a strategy Q is *optimal* against a given strategy P if $|V[P \succeq Q]| \leq |V[P \succeq Q']|$ for any strategy Q' .

Trivial, reasonable, and canonical strategies for \mathcal{P} . For $0 \leq n \leq n^*$, define $V_n := \{v_1, \dots, v_n\}$ to be the leftmost n points in V . Suppose we want to compute $\Gamma_{k,\ell}(V_n)$ for some $1 \leq k \leq n$ and $0 \leq \ell \leq n$. The *trivial strategy* of player \mathcal{P} is to place his points at the k points of V_n with the highest multiplicities—here we consider the multiset V_n as a set of distinct points, each with a multiplicity corresponding to the number of times it occurs in V_n —with ties broken arbitrarily. Let $\|V_n\|$ denote the number of distinct points in V_n . Then the trivial strategy is optimal when $k \geq \|V_n\|$ and also when $\ell \geq 2k$: in the former case \mathcal{P} wins all voters with the trivial strategy, and in the latter case \mathcal{Q} can always win all voters not coinciding with a point in P (namely by surrounding each point $p_i \in P$ by two points sufficiently close to p_i) so the trivial strategy is optimal for \mathcal{P} . Hence, from now on we consider subproblems with $k < \|V_n\|$ and $\ell < 2k$.

We can without loss of generality restrict our attention to strategies for \mathcal{P} that place at most one point in each half-open interval of the form $(v_i, v_{i+1}]$ with $v_i \neq v_{i+1}$, where $0 \leq i \leq n$, $v_0 := -\infty$, and $v_{n^*+1} := \infty$. Indeed, placing more than two points inside an interval $(v_i, v_{i+1}]$ is clearly not useful, and if two points are placed in some interval $(v_i, v_{i+1}]$ then we can always move the leftmost point onto v_i . (If v_i is already occupied by a point in P , then we can just put the point on any unoccupied voter; under our assumption that $k < \|V_n\|$ an unoccupied voter always exists.) We will call a strategy for \mathcal{P} satisfying the property above *reasonable*.

Observation 1 (Banik *et al.* [3]). *Assuming $k < \|V_n\|$ there exist an optimal strategy for \mathcal{P} that is reasonable and has $p_1 \in V$ (that is, p_1 coincides with a voter).*

We can define an ordering on strategies of the same size by sorting them in lexicographical order. More precisely, we say that a strategy $P = \{p_1, \dots, p_k\}$ is *greater than* a strategy $P' = \{p'_1, \dots, p'_k\}$, denoted by $P \succ P'$, if $\langle p_1, \dots, p_k \rangle >_{\text{lex}} \langle p'_1, \dots, p'_k \rangle$, where $>_{\text{lex}}$ denotes the lexicographical order. Using this ordering, the largest reasonable strategy P that is optimal—namely, that realizes $\Gamma_{k,\ell}(V_n)$ —is called the *canonical strategy* of \mathcal{P} .

α -gains, β -gains, and gain sequences. Consider a strategy $P := \{p_1, \dots, p_k\}$. It will be convenient to add two extra points to P , namely $p_0 := -\infty$ and $p_{k+1} := \infty$; this clearly does not influence the outcome of the game. The strategy P thus induces $k + 1$ open intervals of the form (p_i, p_{i+1}) where player \mathcal{Q} may place her points. It is easy to see that there exists an optimal strategy for \mathcal{Q} with the following property: Q contains at most two points in each interval (p_i, p_{i+1}) with $1 \leq i \leq k - 1$, and at most one point in (p_0, p_1) and at most one point in (p_k, p_{k+1}) . From now on we restrict our attention to strategies for \mathcal{Q} with this property.

Now suppose that x and y are consecutive points (with $x < y$) in some strategy P , where x could be $-\infty$ and y could be ∞ . As just argued, \mathcal{Q} either places zero, one, or two points inside (x, y) . When \mathcal{Q} places zero points, then she obviously does not win any of the voters in $V_n \cap (x, y)$. The maximum number of voters \mathcal{Q} can win from $V_n \cap (x, y)$ by placing a single point is the maximum

number of voters in (x, y) that can be covered by an open interval of length $(y - x)/2$; see Banik *et al.* [3]. We call this value the α -gain of \mathcal{Q} in (x, y) and denote it by $\text{gain}_\alpha(V_n, x, y)$. By placing two points inside (x, y) , one immediately to the right of x and one immediately to the left of y , player \mathcal{Q} will win all voters $V_n \cap (x, y)$. Thus the extra number of voters won by the second point in (x, y) as compared to just placing a single point is equal to $|V_n \cap (x, y)| - \text{gain}_\alpha(V_n, x, y)$. We call this quantity the β -gain of \mathcal{Q} in (x, y) and denote it by $\text{gain}_\beta(V_n, x, y)$. Note that for intervals (x, ∞) we have $\text{gain}_\alpha(x, \infty) = |V_n \cap (x, \infty)|$ and $\text{gain}_\beta(x, \infty) = 0$; a similar statement holds for $(-\infty, y)$.

The following observation follows from the fact that $\text{gain}_\alpha(V_n, x, y)$ equals the maximum number of voters in (x, y) that can be covered by an open interval of length $(y - x)/2$.

Observation 2 (Banik *et al.* [3]). *For any x, y we have $\text{gain}_\alpha(V_n, x, y) \geq \text{gain}_\beta(V_n, x, y)$.*

If we let $a := \text{gain}_\alpha(V_n, x, y)$ and $b := \text{gain}_\beta(V_n, x, y)$, then player \mathcal{Q} wins either 0, a , or $a + b$ points depending on whether she plays 0, 1, or 2 points inside the interval. It will therefore be convenient to introduce the notation \oplus_j for $j \in \{0, 1, 2\}$, which is defined as

$$a \oplus_0 b := 0, \quad a \oplus_1 b := a, \quad a \oplus_2 b := a + b.$$

We assume the precedence of these operators are higher than addition.

Let $P := \{p_0, p_1, \dots, p_k, p_{k+1}\}$ be a given strategy for \mathcal{P} , where by convention $p_0 = -\infty$ and $p_{k+1} = \infty$. Consider $\{\text{gain}_\alpha(V_n, p_i, p_{i+1}) : 0 \leq i \leq k\} \cup \{\text{gain}_\beta(V_n, p_i, p_{i+1}) : 0 \leq i \leq k\}$, the multiset of all α -gains and β -gains defined by the intervals (p_i, p_{i+1}) . Sort this sequence in non-increasing order, using the following tie-breaking rules if two gains are equal:

- if one of the gains is for an interval (p_i, p_{i+1}) —that is, the gain is either $\text{gain}_\alpha(V_n, p_i, p_{i+1})$ or $\text{gain}_\beta(V_n, p_i, p_{i+1})$ —and the other gain is for an interval (p_j, p_{j+1}) with $j > i$, then the gain for (p_i, p_{i+1}) precedes the gain for (p_j, p_{j+1}) .
- if both gains are for the same interval (p_i, p_{i+1}) then the α -gain precedes the β -gain.

We call the resulting sorted sequence the *gain sequence* induced by P on V_n . We denote this sequence by $\Sigma_{\text{gain}}(V_n, P)$ or, when P and V_n are clear from the context, sometimes simply by Σ_{gain} .

The canonical strategy of \mathcal{Q} and sequence representations. Given the multiset V_n , a strategy P and value ℓ , player \mathcal{Q} can compute an optimal strategy as follows. First she computes the gain sequence $\Sigma_{\text{gain}}(V_n, P)$ and chooses the first ℓ gains in $\Sigma_{\text{gain}}(V_n, P)$. Then for each $0 \leq i \leq k$ she proceeds as follows. When $\text{gain}_\alpha(V_n, p_i, p_{i+1})$ and $\text{gain}_\beta(V_n, p_i, p_{i+1})$ are both chosen, she places two points in (p_i, p_{i+1}) that win all voters in (p_i, p_{i+1}) ; when only $\text{gain}_\alpha(V_n, p_i, p_{i+1})$ is chosen, she places one point in (p_i, p_{i+1}) that win $\text{gain}_\alpha(V_n, p_i, p_{i+1})$ voters. (By Observation 2 and the tie-breaking rules, when $\text{gain}_\beta(V_n, p_i, p_{i+1})$ is chosen it is always the case that $\text{gain}_\alpha(V_n, p_i, p_{i+1})$ is also chosen.) The resulting optimal strategy Q is called the *canonical strategy* of \mathcal{Q} with ℓ points against P on V_n .

From now on we restrict the strategies of player \mathcal{Q} to canonical strategies. In an optimal strategy, player \mathcal{Q} places at most two points in any interval induced by a strategy $P = \{p_0, \dots, p_{k+1}\}$, and when we know that \mathcal{Q} places a single point (and similarly when she places two points) then we also know where to place the point(s). Hence, we can represent an optimal strategy Q , for given V_n and P , by a sequence $M(V, P, Q) := \langle m_0, \dots, m_k \rangle$ where $m_i \in \{0, 1, 2\}$ indicates how many points \mathcal{Q} plays in the interval (p_i, p_{i+1}) . We call $M(V, P, Q)$ the *sequence representation* of the strategy Q against P on V_n . We denote the sequence representation of the canonical strategy of \mathcal{Q} with ℓ points against P on V_n by $M(V, P, \ell)$. We have the following observation.

Observation 3. *The canonical strategy of \mathcal{Q} with ℓ points against P is the optimal strategy Q with ℓ points against P whose sequence representation is maximal in the lexicographical order.*

2.2 The subproblems for a dynamic-programming solution

For clarity, in the rest of Section 2 we assume the multiset of voters V does not have repetitive entries, i.e we have a set of voters, and not a multiset. While all the results are easily extendible to multisets, dealing with them adds unnecessary complexity to the text.

Our goal is to develop a dynamic-programming algorithm to compute $\Gamma_{k^*, \ell^*}(V)$. Before we can define the subproblems on which the dynamic program is based, we need to introduce the concept of *thresholds*, which is a crucial ingredient in the subproblems.

Strict and loose thresholds. Consider an arbitrary gain sequence $\Sigma_{\text{gain}}(V_n, P) = \langle \tau_1, \dots, \tau_{2k+2} \rangle$. Recall that each τ_i is the α -gain or β -gain of some interval (p_i, p_{i+1}) , and that these gains are sorted in non-increasing order. We call any integer value $\tau \in [\tau_{\ell+1}, \tau_\ell]$ an ℓ -*threshold* for \mathcal{Q} induced by P on V_n , or simply a *threshold* if ℓ is clear from the context. We implicitly assume $\tau_0 := n$ so that talking about 0-threshold is also meaningful. Note that when $\tau_\ell \geq \tau > \tau_{\ell+1}$ then the canonical strategy for \mathcal{Q} chooses all gains larger than τ and no gains smaller or equal to τ . Hence, we call τ a *strict* threshold if $\tau_\ell \geq \tau > \tau_{\ell+1}$. On the other hand, when $\tau = \tau_{\ell+1}$ then gains of value τ may or may not be chosen by the canonical strategy of \mathcal{Q} . (Note that in this case for gains of value τ to be picked, we would actually need $\tau_\ell = \tau = \tau_{\ell+1}$.) In this case we call τ a *loose* threshold.

The idea will be to guess the threshold τ in an optimal solution and then use the fact that fixing the threshold τ helps us to limit the strategies for \mathcal{P} and anticipate the behavior of \mathcal{Q} . Let P_{opt} be the canonical strategy realizing $\Gamma_{k^*, \ell^*}(V)$. We call any ℓ^* -threshold of P_{opt} an *optimal threshold*. We devise an algorithm that gets a value τ as input and computes $\Gamma_{k^*, \ell^*}(V)$ correctly if τ is an optimal threshold, and computes a value not greater than $\Gamma_{k^*, \ell^*}(V)$, otherwise.

Clearly we only need to consider values of τ that are at most n^* . In fact, since each α -gain or β -gain in a given gain sequence corresponds to a unique subset of voters, the ℓ^* -th largest gain can be at most n^*/ℓ^* , so we only need to consider τ -values up to $\lfloor n^*/\ell^* \rfloor$. Observe that when there exists an optimal strategy that induces an ℓ^* -threshold equal to zero, then \mathcal{Q} can win all voters not explicitly covered by P . In this case the trivial strategy is optimal for \mathcal{P} . Our global algorithm is now as follows.

1. For all thresholds $\tau \in \{1, \dots, \lfloor n^*/\ell^* \rfloor\}$, compute an upper bound on the number of voters \mathcal{P} can win with a strategy that has an ℓ^* -threshold τ . For the run where τ is an optimal threshold, the algorithm will return $\Gamma_{k^*, \ell^*}(V)$.
2. Compute the number of voters \mathcal{P} wins in the game $\langle V, k^*, \ell^* \rangle$ by the trivial strategy.
3. Report the best of all solutions found.

The subproblems for a fixed threshold τ . From now on we consider a fixed threshold value $\tau \in \{1, \dots, \lfloor n^*/\ell^* \rfloor\}$. The subproblems in our dynamic-programming algorithm for the game $\langle V, k^*, \ell^* \rangle$ have several parameters.

- A parameter $n \in \{0, \dots, n^*\}$, specifying that the subproblem is on the voter set V_n .
- Parameters $k, \ell \in \{0, \dots, n\}$, specifying that \mathcal{P} can use $k + 1$ points and \mathcal{Q} can use ℓ points.
- A parameter $\gamma \in \{0, \dots, n\}$, specifying the number of voters \mathcal{P} must win.
- A parameter $\delta \in \{\text{strict}, \text{loose}\}$, specifying the strictness of the fixed ℓ -threshold τ .

Intuitively, the subproblem specified by a tuple $\langle n, k, \ell, \gamma, \delta \rangle$ asks for a strategy P where \mathcal{P} wins at least γ voters from V_n and such that P that induces an ℓ -threshold of strictness δ , against an opponent \mathcal{Q} using ℓ points. Player \mathcal{P} may use $k + 1$ points and his objective will be to push his last point, p_{k+1} as far to the right as possible. The value of the solution to such a subproblem, which we denote by $X_{\text{max}}(n, k, \ell, \gamma, \delta)$, will indicate how far to the right we can push p_{k+1} . Ultimately we will

be interested in solutions where \mathcal{P} can push p_{k^*+1} all the way to $+\infty$, which means he can actually win γ voters by placing only k^* points.

To formally define $X_{\max}(n, k, \ell, \gamma, \delta)$, we need one final piece of notation. Let $x \in \mathbb{R} \cup \{-\infty\}$, let $n \in \{1, \dots, n^*\}$, and let a, b be integers. For convenience, define $v_{n^*+1} := \infty$. Now we define the (a, b) -span of x to v_{n+1} , denoted by $\text{span}(x, n, a, b)$, as

$$\text{span}(x, n, a, b) := \begin{cases} \text{the maximum real value } y \in (v_n, v_{n+1}] \text{ such} \\ \text{that } \text{gain}_\alpha(V, x, y) = a \text{ and } \text{gain}_\beta(V, x, y) = b & \text{if } x \neq -\infty \text{ and } y \text{ exists} \\ -\infty & \text{otherwise} \end{cases}$$

Definition 4. For parameters $n \in \{0, \dots, n^*\}$, $k, \ell, \gamma \in \{0, \dots, n\}$, and $\delta \in \{\text{strict}, \text{loose}\}$, we define the value $X_{\max}(n, k, \ell, \gamma, \delta)$ and what it means when a strategy P realizes this value, as follows.

- For $k = 0$, we call it an *elementary subproblem*, and define $X_{\max}(n, k, \ell, \gamma, \delta) = v_{n+1}$ if

1. $\{v_{n+1}\}$ wins at least γ voters from V_n , and
2. $\{v_{n+1}\}$ induces an ℓ -threshold τ with strictness δ on V_n ,

and we define $X_{\max}(n, k, \ell, \gamma, \delta) = -\infty$ otherwise. In the former case we say that $P := \{v_{n+1}\}$ realizes $X_{\max}(n, k, \ell, \gamma, \delta)$.

- For $k > 0$, we call it a *non-elementary subproblem*, and $X_{\max}(n, k, \ell, \gamma, \delta)$ is defined to be equal to the maximum real value $y \in (v_n, v_{n+1}]$ such that there exists a strategy $P := P' \cup \{y\}$ with $P' = \{p_1, \dots, p_k\}$, integer values n', a, b with $0 \leq n' < n$ and $0 \leq a, b \leq n$, an integer $j \in \{0, 1, 2\}$, and a $\delta' \in \{\text{strict}, \text{loose}\}$ satisfying the following conditions:

1. P wins at least γ voters from V_n ,
2. P induces an ℓ -threshold τ with strictness δ on V_n ,
3. $\text{span}(p_k, n, a, b) = y$,
4. P' realizes $X_{\max}(n', k-1, \ell-j, \gamma-n+n'+a \oplus_j b, \delta')$,
5. Let $M(V_{n'}, P', \ell-j) := \langle m'_0, \dots, m'_k \rangle$ and $M(V_n, P, \ell) := \langle m_0, \dots, m_{k+1} \rangle$. Then $m'_i = m_i$ for all $0 \leq i < k$.

When a set P satisfying the conditions exists, we say that P realizes $X_{\max}(n, k, \ell, \gamma, \delta)$. We define $X_{\max}(n, k, \ell, \gamma, \delta) = -\infty$ if no such P exists.

By induction we can show that if the parameters n, k, ℓ are not in a certain range, namely if one of the conditions $\ell < 2(k+1)$ or $k \leq \|V_n\|$ is violated, then $X_{\max}(n, k, \ell, \gamma, \delta) = -\infty$. The next lemma shows we can compute $\Gamma_{k^*, \ell^*}(V)$ from the solutions to our subproblems.

Lemma 5. Let $V = \{v_1, \dots, v_{n^*}\}$ be a set of n^* voters in \mathbb{R}^1 . Let $0 \leq k^* \leq n^*$ and $1 \leq \ell^* \leq n^*$ be two integers such that $\ell^* < 2k^*$ and $k^* < \|V\|$, and let τ be a fixed threshold. Then

$$\Gamma_{k^*, \ell^*}(V) \geq \text{the maximum value of } \gamma \text{ with } 0 \leq \gamma \leq n^* \text{ for which there exist} \quad (1) \\ \text{a } \delta \in \{\text{loose}, \text{strict}\} \text{ such that } X_{\max}(n^*, k^*, \ell^*, \gamma, \delta) = \infty.$$

Moreover, for an optimal threshold $\tau_{\text{opt}} > 0$, the inequality changes to equality.

Proof. It is clear that $\Gamma_{k^*, \ell^*}(V)$ is at least equal to the right-hand side in (1). Indeed, an equation $X_{\max}(n^*, k^*, \ell^*, \gamma, \delta) = \infty$ implies by definition that there is a strategy P' of k^* points such that

$P' \cup \{\infty\}$ wins at least γ voters against an opponent \mathcal{Q} with ℓ^* points, and then P' must win γ voters as well. Next we prove the opposite direction for an optimal threshold τ_{opt} .

Let P_{opt} be the canonical strategy realizing $\Gamma_{k^*, \ell^*}(V)$. By definition, τ_{opt} is an ℓ^* -threshold induced by P_{opt} on V . Let Q_{opt} be the canonical strategy of \mathcal{Q} with ℓ^* points against P_{opt} on V and let $M(V, P_{\text{opt}}, \ell^*) = \langle m_0, \dots, m_{k^*} \rangle$ be its sequence representation. Define $P_k := \{p_1, \dots, p_k, p_{k+1}\}$ and $Q_k := Q_{\text{opt}} \cap [-\infty, p_{k+1}]$, and let n_k be the largest index such that $v_{n_k} < p_{k+1}$. Let $\ell_k = |Q_k|$, and $\gamma_k = |V_{n_k}[P_k \succeq Q_k]|$, that is, γ_k is the number of voters from V_{n_k} won by P_k against Q_k . Note that τ_{opt} is an ℓ_k -threshold of P_k on V_{n_k} . Let $\delta_k \in \{\text{strict}, \text{loose}\}$ indicate whether our fixed threshold τ_{opt} is a strict or loose threshold induced by P_k and ℓ_k on V_{n_k} , where $p_{k^*+1} := \infty$.

We will need the following claim.

Claim. Q_k is the canonical strategy of \mathcal{Q} with ℓ_k points against P_k on V_{n_k} , and for all $0 \leq k \leq k^*$ we have $M(V_{n_k}, P_k, Q_k) = \langle m_0, \dots, m_k, 0 \rangle$.

Proof of claim. The second part of the claim is obvious by the definition of sequence representation and by the definition of Q_k .

The first part of the claim can be shown by contradiction. Suppose Q'_k is the canonical strategy of \mathcal{Q} with ℓ_k points against P_k on V_{n_k} , where $Q'_k \neq Q_k$. Let $M(V_{n_k}, P_k, Q'_k) = \langle m'_0, \dots, m'_k, 0 \rangle$. Consider the strategy Q with sequence representation $\langle m'_0, \dots, m'_k, m_{k+1}, \dots, m_{k^*} \rangle$ of size ℓ^* against P_{opt} . We have two cases.

- Q_k is not an optimal strategy against P_k . This means Q wins more voters than Q_{opt} against P_{opt} , because the winning of any given voter by Q is only dependent on the number of points Q places in the interval (p_i, p_{i+1}) containing that voter. But this contradicts the fact that Q_{opt} is an optimal strategy against P_{opt} .
- Q_k is an optimal strategy against P_k , but Q_k is smaller in the lexicographical order than Q'_k . This means that Q is a strategy of equal gain as Q_{opt} whose sequence representation is lexicographically larger. But this contradicts the fact that Q_{opt} is a canonical strategy against P_{opt} .

◁

Let I_k denote the subproblem given by $\langle n_k, k, \ell_k, \gamma_k, \delta_k \rangle$ for $\tau = \tau_{\text{opt}}$. Below we will show by induction on k that for any $0 \leq k \leq k^*$ we have

$$X_{\max}(I_k) = p_{k+1} \text{ and } X_{\max}(I_k) \text{ is realized by } P_k, \quad (2)$$

where we use $X_{\max}(I_k)$ as a shorthand for $X_{\max}(n_k, k, \ell_k, \gamma_k, \delta_k)$. Note that (2) finishes the proof of the lemma. Indeed, $p_{k^*+1} = \infty$ by definition, and $n_{k^*} = n^*$ which implies $V_{n_{k^*}} = V$. Hence, $\gamma_{k^*} = \Gamma_{k^*, \ell^*}(V)$ and $\ell_{k^*} = \ell^*$, and so there is a $\delta \in \{\text{loose}, \text{strict}\}$ such that $X_{\max}(n^*, k^*, \ell^*, \Gamma_{k^*, \ell^*}(V), \delta_{k^*}) = \infty$, thus finishing the proof. It remains to prove (2).

Base case: $k = 0$. By definition n_0 is such that $v_{n_0} < p_1$ and $v_{n_0+1} \geq p_1$. Moreover, $p_1 \in V$ by Observation 1. Hence, we have $v_{n_0+1} = p_1$, which establishes (2).

Induction step: $k > 0$. We first prove that $y := p_{k+1}$ satisfies all five conditions from Definition 4 for $a := \text{gain}_\alpha(V, p_k, p_{k+1})$ and $b := \text{gain}_\beta(V, p_k, p_{k+1})$, and for $j := m_k$ and $\delta' := \delta_k$. This implies that $X_{\max}(I_k) \geq p_{k+1}$. After that we argue there is no larger value of y satisfying all conditions, thus finishing the proof.

1. P_k wins at least γ_k voters from V_{n_k} .

P_k wins γ_k voters against Q_k by definition. Moreover, by the Claim above, Q_k is an optimal strategy for \mathcal{Q} against P_k on V_{n_k} . Hence, P_k can win γ_k voters from V_{n_k} .

2. P_k induces an ℓ -threshold τ_{opt} with strictness δ_k on V_{n_k} .

This is true by definition.

3. $\text{span}(p_k, n_k, a, b) = p_{k+1}$.

Obviously $\text{span}(p_k, n_k, a, b) \geq p_{k+1}$. Now let $y := \text{span}(p_k, n_k, a, b)$ and assume for a contradiction that $y > p_{k+1}$. In the next paragraph, we will argue that this implies that $P := \{p_1, \dots, p_k, y, p_{k+2}, \dots, p_{k^*}\}$ realizes $\Gamma_{k^*, \ell^*}(V)$. This gives the desired contradiction since P is then an optimal strategy that is lexicographically greater than the canonical strategy P_{opt} .

Since the only difference between P and P_{opt} is that p_{k+1} is moved to the right to the position y , the intervals (p_i, p_{i+1}) only change for $i = k$ and for $i = k+1$. Thus the possible gains for \mathcal{Q} in all intervals stay the same, except possibly in (p_k, p_{k+1}) and (p_{k+1}, p_{k+2}) . Since $a = \text{gain}_\alpha(V, p_k, p_{k+1})$ and $b = \text{gain}_\beta(V, p_k, p_{k+1})$ and $y = \text{span}(p_k, n_k, a, b)$ by definition, the gains in (p_k, p_{k+1}) are the same as the gains in (p_k, y) . Finally, since $(y, p_{k+2}) \subset (p_{k+1}, p_{k+2})$, it is clear that \mathcal{Q} cannot win more voters in (y, p_{k+2}) against P than she can win in (p_{k+1}, p_{k+2}) against P_{opt} by playing a single point inside these intervals. Similarly, she cannot win more voters by playing two points in (y, p_{k+2}) , as compared to playing two points in (p_{k+1}, p_{k+2}) . Hence, P wins at least the same number of voters as P_{opt} and so P realizes $\Gamma_{k^*, \ell^*}(V)$.

4. P_{k-1} realizes $X_{\max}(n_{k-1}, k-1, \ell_k - j, \gamma_k - n_k + n_{k-1} + a \oplus_j b, \delta_{k-1})$ for $j = m_k$.

By the induction hypothesis and since $\ell_{k-1} = \ell_k - j$ by definition, P_{k-1} realizes $X_{\max}(n_{k-1}, k-1, \ell_k - j, \gamma_{k-1}, \delta_{k-1})$. By the above Claim, Q_{k-1} is optimal against P_{k-1} on $V_{n_{k-1}}$ and Q_k is optimal against P_k on V_{n_k} . Hence, $\gamma_{k-1} = \gamma_k - n_k + n_{k-1} + a \oplus_j b$, which proves the condition.

5. Let $M(V_{n_{k-1}}, P_{k-1}, \ell_{k-1}) := \langle m'_0, \dots, m'_k \rangle$ and $M(V_{n_k}, P_k, \ell_k) := \langle m_0, \dots, m_{k+1} \rangle$. Then $m'_i = m_i$ for all $0 \leq i < k$.

This is true by the above Claim.

We still need to show that p_{k+1} is the maximum value in $(v_n, v_{n+1}]$ satisfying the above mentioned conditions. This is obvious for $k = k^*$. For $k < k^*$, with a similar reasoning to the proof of the third condition, we can show that p_{k+1} is the maximum value in $(v_n, v_{n+1}]$ satisfying the conditions, as otherwise P_{opt} cannot be the canonical strategy realizing $\Gamma_{k^*, \ell^*}(V)$. Namely, we assume $y \in (v_n, v_{n+1}]$ with $y > p_{k+1}$ is the solution to I_k realized by $P'_k := \{p'_1, \dots, p'_k, y\}$, and then we argue that $P := \{p'_1, \dots, p'_k, y, p_{k+2}, \dots, p_{k^*}\}$ is an optimal solution, thus obtaining a contradiction since P is lexicographically larger than P_{opt} . It remains to prove that P is optimal.

We show that any other strategy Q against P on V cannot win more voters than Q_{opt} wins against P_{opt} , which shows P realizes $\Gamma_{k^*, \ell^*}(V)$ and is therefore optimal. Let Q'_k be the canonical strategy of \mathcal{Q} with ℓ_k points against P'_k on V_{n_k} . As P'_k wins at least γ_k voters and by definition of γ_k , P_k wins exactly γ_k voters, the number of voters Q'_k wins against P'_k cannot be more than the number of voters Q_k , which is part of Q_{opt} , wins against P_k .

Let $a := \text{gain}_\alpha(V, p_{k+1}, p_{k+2})$, and $b := \text{gain}_\beta(V, p_{k+1}, p_{k+2})$. Any strategy Q against P must select its gains from the gain sequence $\Sigma_{\text{gain}}(V, P)$. But, all the gains in $\Sigma_{\text{gain}}(V, P)$ are either the same gains chosen by Q'_k or Q_{opt} of a value at least τ_{opt} , or they have a value of at most τ_{opt} because of the threshold τ_{opt} . The only exceptions are the gains in interval (y, p_{k+2}) where

by moving p_{k+1} to the right a does not increase, but b might increase. But, even in that case when b increases and its value gets bigger than τ_{opt} and it is chosen by Q , it means a is chosen by both Q_{opt} and Q . Hence, any extra voters won by Q through selecting the gain b are stolen from a which means the number of voters won by Q does not increase. □

Remark. Usually in dynamic programming, subproblems have a clean non-recursive definition—the recursion only comes in when a recursive formula is given to compute the value of an optimal solution. Our approach is more complicated: Definition 4 above gives a recursive but “non-constructive” subproblem definition (and Lemma 5 shows how to use it) and Lemma 8 below will then give a different recursive formula to actually compute the solutions to the subproblems.

2.3 Computing solutions to subproblems

The solution to an elementary subproblem follows fairly easily from the definitions, and can be computed in constant time.

Lemma 6. *Assuming the voter set V is given in sorted order, we can find the solution to an elementary subproblem in constant time.*

Proof. Consider an elementary subproblem $I = \langle n, 0, \ell, \gamma, \delta \rangle$. If $\ell = 0$ then $X_{\max}(I) = v_{n+1}$ if and only if the following two conditions hold: (i) $\gamma \leq n$ and (ii) $\tau > n$ or ($\tau = n$ and $\delta = \text{loose}$). If $\ell = 1$ then $X_{\max}(I) = v_{n+1}$ if and only if (i) $\gamma = 0$ and (ii) $\tau < n$ or ($\tau = n$ and $\delta = \text{strict}$). Otherwise $X_{\max}(I) = -\infty$. □

By definition, in order to obtain a strategy P realizing the solution to a non-elementary subproblem $I = \langle n, k, \ell, \gamma, \delta \rangle$ of size k , we need a solution to a smaller subproblem $I' = \langle n', k-1, \ell', \gamma', \delta' \rangle$ of size $k-1$ and add one point $y \in (v_n, v_{n+1}]$ to the strategy $P' = \{p_1, \dots, p_k\}$ realizing I' . Thus by adding y , we extend the solution to I' to get a solution to I . To find the “right” subproblem I' , we guess some values for $n', a, b, j \in \{0, 1, 2\}$, and $\delta' \in \{\text{strict}, \text{loose}\}$; these values are enough to specify I' . We note that there are just a polynomial number of cases and therefore a polynomial number of values for the value $y \in (v_n, v_{n+1}]$ which we want to maximize. Namely, there are $O(n)$ choices for the values n', a , and b , three choices for j , and two choices for δ' . This makes $O(n^3)$ different cases to be considered for each subproblem I , in total. However, not all those subproblems can be extended to I . In the following definition, we list the triples (δ', j, δ) that provide all the valid combinations that guarantee the extendibility of I' to I .

Let a and b be the α -gain and β -gain of the interval (p_k, y) in a strategy $P = \{p_1, \dots, p_k, y\}$ with threshold τ . We define the following sets of triples depending on the relationship between a, b and τ :

$$\Delta(\tau, a, b) := \begin{cases} \{(\text{loose}, 2, \text{loose}), (\text{strict}, 2, \text{strict})\} & \text{if } a > \tau \wedge b > \tau \\ \{(\text{loose}, 1, \text{loose}), (\text{strict}, 1, \text{loose}), (\text{strict}, 2, \text{strict})\} & \text{if } a > \tau \wedge b = \tau \\ \{(\text{loose}, 1, \text{loose}), (\text{strict}, 1, \text{strict})\} & \text{if } a > \tau \wedge b < \tau \\ \{(\text{loose}, 0, \text{loose}), (\text{strict}, 0, \text{loose}), (\text{strict}, 1, \text{loose}), (\text{strict}, 2, \text{strict})\} & \text{if } a = \tau \wedge b = \tau \\ \{(\text{loose}, 0, \text{loose}), (\text{strict}, 0, \text{loose}), (\text{strict}, 1, \text{strict})\} & \text{if } a = \tau \wedge b < \tau \\ \{(\text{loose}, 0, \text{loose}), (\text{strict}, 0, \text{strict})\} & \text{if } a < \tau \wedge b < \tau. \end{cases}$$

Lemma 7. Let $P' = \{p_1, \dots, p_k\}$ and $P := P' \cup \{y\}$, be two reasonable strategies on $V_{n'}$ and V_n , where $n' = \operatorname{argmax}_{1 \leq i \leq n^*} v_i < p_k$, $n = \operatorname{argmax}_{1 \leq i \leq n^*} v_i < y$, and $y \in (v_n, v_{n+1}]$. Let $a = \operatorname{gain}_\alpha(V_n, p_k, y)$ and $b = \operatorname{gain}_\beta(V_n, p_k, y)$, and assume $\tau > 0$ is an $(\ell - j)$ -threshold of strictness δ' for Q induced by P' on $V_{n'}$, where $j \in \{0, 1, 2\}$. Then, there exists a triple $(\delta', j, \delta) \in \Delta(\tau, a, b)$ if and only if

1. P induces an ℓ -threshold τ with strictness δ on V_n ,
2. Let $M(V_{n'}, P', \ell - j) := \langle m'_0, \dots, m'_k \rangle$ and $M(V_n, P, \ell) := \langle m_0, \dots, m_{k+1} \rangle$. Then $m'_i = m_i$ for all $0 \leq i < k$.

Moreover, this triple is unique if it exists.

Proof. We prove the lemma for the case $\tau = a = b$. The proof for the other cases is similar.

Let Q' and Q be canonical strategies associated with $M(V_{n'}, P', \ell - j)$ and $M(V_n, P, \ell)$, respectively, and let $\Sigma_{\operatorname{gain}}(V_{n'}, P') = \langle \tau'_1, \dots, \tau'_{2k+2} \rangle$. We note that $\Sigma_{\operatorname{gain}}(V_n, P)$ can be constructed by inserting a and b after the last τ value in $\Sigma_{\operatorname{gain}}(V_{n'}, P')$. There are six cases,

- $\delta' = \text{loose}$ and $j = 0$.

In this case $(\text{loose}, 0, \text{loose})$ is the only corresponding triple in $\Delta(\tau, a, b)$. As $\delta' = \text{loose}$, we have $\tau'_{\ell-j+1} = \tau$, and therefore a and b are not selected by Q . Therefore $Q' = Q$ and the strictness status also does not change. Hence both conditions are satisfied and $\delta = \text{loose}$.

- $\delta' = \text{strict}$ and $j = 0$.

In this case $(\text{strict}, 0, \text{loose})$ is the only corresponding triple in $\Delta(\tau, a, b)$. As $\delta' = \text{strict}$, we have $\tau'_{\ell-j+1} < \tau$, and therefore a and b are not selected by Q ; however, the strictness of the ℓ -threshold τ for P is different. Hence both conditions are satisfied and $\delta = \text{loose}$.

- $\delta' = \text{loose}$ and $j = 1$.

As $\delta' = \text{loose}$, we have $\tau'_{\ell-j+1} = \tau$, and therefore a and b are not selected by Q . However, the gain $\tau'_{\ell-j+1}$ is selected by Q which violates the condition 2. We also note that there is no corresponding triple in $\Delta(\tau, a, b)$ for this case.

- $\delta' = \text{strict}$ and $j = 1$.

In this case $(\text{strict}, 1, \text{loose})$ is the only corresponding triple in $\Delta(\tau, a, b)$. As $\delta' = \text{strict}$, we have $\tau'_{\ell-j+1} < \tau$, and therefore a is selected by Q and b is not selected, and the strictness of the ℓ -threshold τ for P is different. Hence both conditions are satisfied and $\delta = \text{loose}$.

- $\delta' = \text{strict}$ and $j = 2$.

In this case $(\text{strict}, 2, \text{strict})$ is the only corresponding triple in $\Delta(\tau, a, b)$. As $\delta' = \text{strict}$, we have $\tau'_{\ell-j+1} < \tau$, and therefore both a and b are selected by Q , and the strictness of the ℓ -threshold τ for P is the same. Hence both conditions are satisfied and $\delta = \text{strict}$.

- $\delta' = \text{loose}$ and $j = 2$.

As $\delta' = \text{loose}$, we have $\tau'_{\ell-j+1} = \tau$, and therefore just a is selected by Q , while b is not selected. However, the gain $\tau'_{\ell-j+1}$ is selected by Q which violates the condition 2. We also note that there is no corresponding triple in $\Delta(\tau, a, b)$ for this case.

□

Lemma 8. For a non-elementary subproblem $I = \langle n, k, \ell, \gamma, \delta \rangle$, we have

$$X_{\max}(n, k, \ell, \gamma, \delta) = \max_{0 \leq n' < n} \max_{0 \leq a \leq n} \max_{0 \leq b \leq n} \max_{(\delta', j, \delta) \in \Delta(\tau, a, b)} \text{span}(X_{\max}(n', k-1, \ell-j, \gamma-n+n'+a \oplus_j b, \delta'), n, a, b).$$

Proof. As stated earlier, when one of the conditions $\ell < 2(k+1)$ or $k \leq \|V_n\|$ is violated, the left hand side evaluates to $-\infty$. When this happens, then the corresponding condition for each subproblem $X_{\max}(n', k-1, \ell-j, \gamma-n+n'+a \oplus_j b, \delta')$ is also violated, i.e. $(\ell-j) < 2((k-1)+1)$ or $(k-1) \leq \|V_{n'}\|$. Hence, each of these subproblems and the right hand side evaluates to $-\infty$, too. When $\ell < 2(k+1)$ and $k \leq \|V_n\|$ are both satisfied, we prove the lemma as follows.

Let y be the solution to I and let this solution be realized by the strategy $P := P' \cup \{y\}$ for some $P' = \{p_1, \dots, p_k\}$. By definition, P' realizes $X_{\max}(n', k-1, \ell-j, \gamma-n+n'+a \oplus_j b, \delta')$ for some $\delta' \in \{\text{loose, strict}\}$ and $j \in \{0, 1, 2\}$ where $a = \text{gain}_\alpha(V_n, p_k, y)$, $b = \text{gain}_\beta(V_n, p_k, y)$, and $n' = \text{argmax}_i v_i < p_k$. By the second and fifth conditions of a subproblem and Lemma 7, we have $(\delta', j, \delta) \in \Delta(\tau, a, b)$, which shows

$$X_{\max}(n, k, \ell, \gamma, \delta) \leq \max_{0 \leq n' < n} \max_{0 \leq a \leq n} \max_{0 \leq b \leq a} \max_{(\delta', j, \delta) \in \Delta(\tau, a, b)} \text{span}(X_{\max}(n', k-1, \ell-j, \gamma-n+n'+a \oplus_j b, \delta'), n, a, b).$$

Conversely, consider a value $y = \text{span}(y', n, a, b) > -\infty$, where $y' = X_{\max}(n', k-1, \ell-j, \gamma-n+n'+a \oplus_j b, \delta')$ is realized by $P' = \{p_1, \dots, p_k\}$ for some values n', a, b, δ' , and j such that $(\delta', j, \delta) \in \Delta(\tau, a, b)$. Let $P := P' \cup \{y\}$. It is not hard to verify that the conditions for y to be a solution to I realized by P are satisfied by these values. Indeed, we have:

1. P wins at least γ voters from V_n .

By the second and the fifth conditions and the fact that $\tau > 0$, we have $M_k = j$ which shows \mathcal{P} can win at least γ voters.

2. P induces an ℓ -threshold τ with strictness δ on V_n .

This condition is satisfied by Condition 1 of Lemma 7 and by the uniqueness of δ .

3. $\text{span}(p_k, n, a, b) = y$.

This condition is satisfied by the assumption.

4. P' realizes $X_{\max}(n', k-1, \ell-j, \gamma-n+n'+a \oplus_j b, \delta')$.

This condition is satisfied by the assumption.

5. Let $M(V_{n'}, P', \ell-j) = \langle m'_0, \dots, m'_k \rangle$ and $M(V_n, P, \ell) = \langle m_0, \dots, m_{k+1} \rangle$. Then $m'_i = m_i$ for all $0 \leq i < k$.

This condition is satisfied by Condition 2 of Lemma 7.

As all the conditions are satisfied, we have

$$X_{\max}(n, k, \ell, \gamma, \delta) \geq \max_{0 \leq n' < n} \max_{0 \leq a \leq n} \max_{0 \leq b \leq a} \max_{(\delta', j, \delta) \in \Delta(\tau, a, b)} \text{span}(X_{\max}(n', k-1, \ell-j, \gamma-n+n'+a \oplus_j b, \delta'), n, a, b),$$

which completes the proof. \square

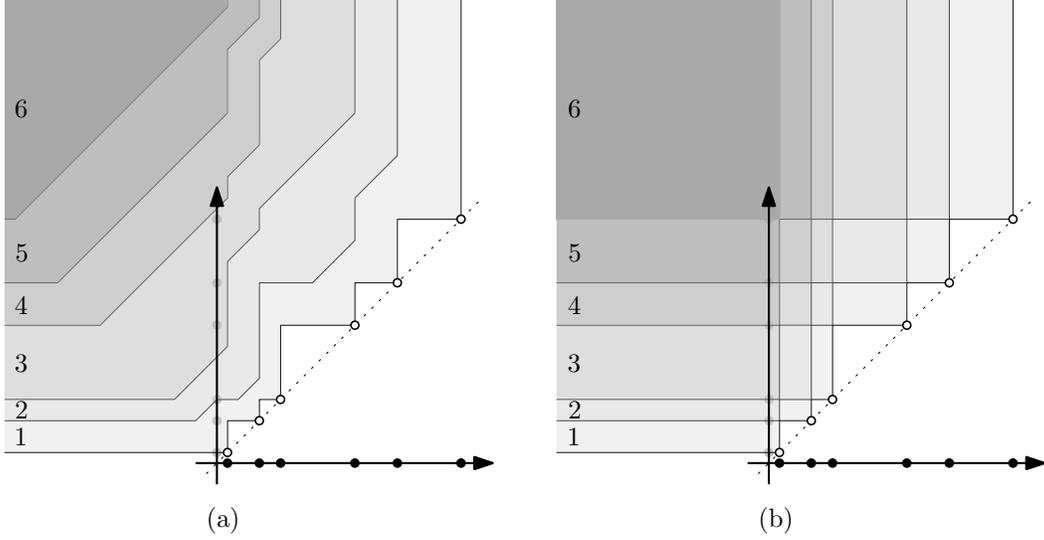


Figure 1: **a)** A -map of $V = \{1, 4, 6, 13, 17, 23\}$ with the corresponding α -gain for each region. **b)** B -map of V with the corresponding β -gain for each region.

If we can compute the span function efficiently, we can compute all the solutions by dynamic programming and solve the problem. However, a solution based on a trivial dynamic-programming algorithm will be of running time complexity $\lfloor n^*/\ell^* \rfloor \cdot O(k^*\ell^*(n^*)^2) \cdot O((n^*)^3 f(n^*)) = O(k^*(n^*)^6 f(n^*))$ where $\lfloor n^*/\ell^* \rfloor$ is the total number of choices for the threshold, $O(k^*\ell^*(n^*)^2)$ is the number of subproblems for each threshold, and $O((n^*)^3 f(n^*))$ is the time needed to solve each subproblem where $f(n)$ is the time needed to compute the $\text{span}(x, n, a, b)$ function. This algorithm is quite slow. More importantly it is not easy to compute the span function. In the following, we introduce some new concepts to compute the span function and also get a better running time.

2.4 Computing the span function using gain maps

Before we give the algorithm we introduce the *gain map*, which we need to compute the span function. Consider an arbitrary strategy P of \mathcal{P} on V , and recall that such a strategy induces open intervals of the form (p_i, p_{i+1}) where \mathcal{Q} can place her points. We can represent any possible interval (x, y) that may arise in this manner as a point (x, y) in the plane. Thus the locus of all possible intervals is the region $R := \{(x, y) : x < y\}$. We will define two subdivisions of this region, the A -map and the B -map, and the gain map will then be the overlay of the A -map and the B -map.

The A -map is the subdivision of R into regions A^t and B^t , for $0 \leq t \leq n^*$, defined as follows:

$$A^t := \{(x, y) : \text{gain}_\alpha(V, x, y) = t\} \text{ and}$$

$$B^t := \{(x, y) : \text{gain}_\alpha(V, x, y) + \text{gain}_\beta(V, x, y) = t\}.$$

In other words, A^t is the locus of all intervals (x, y) such that, if (x, y) is an interval induced by P , then \mathcal{Q} can win t voters (but no more than t) from $V \cap (x, y)$ by placing a single point in (x, y) . To construct the A -map, let $A^{\geq t}$ denote the locus of all intervals (x, y) such that $\text{gain}_\alpha(V, x, y) \geq t$. Note that $A^t = A^{\geq t} \setminus A^{\geq t+1}$. For $1 \leq i \leq n^* - t + 1$, let $V_i^t := \{v_i, \dots, v_{i+t-1}\}$ and define

$$A_i^{\geq t} := \{(x, y) : V_i^t \subset (x, y) \text{ and}$$

$$\mathcal{Q} \text{ can win all voters in } V_i^t \text{ by placing a single point in } (x, y)\}.$$

Then we have $A_i^{\geq t} = \{(x, y) : x < v_i \text{ and } y > v_{i+t-1} \text{ and } y > x + 2(v_{i+t-1} - v_i)\}$. Here the conditions $x < v_i$ and $y > v_{i+t-1}$ are needed to guarantee that $V_i^t \subset (x, y)$. The condition $y > x + 2(v_{i+t-1} - v_i)$ implies that V_i^t can be covered with an interval of length $(y - x)/2$, which is necessary and sufficient for \mathcal{Q} to be able to win all these voters. Note that each region $A_i^{\geq t}$ is the intersection of three halfplanes, bounded by a vertical, a horizontal and a diagonal line, respectively.

Since \mathcal{Q} can win at least t voters in inside (x, y) with a single point if she can win at least t consecutive voters with a single point, we have $A^{\geq t} = \bigcup_{i=1}^{n^*-t+1} A_i^{\geq t}$. Thus $A^{\geq t}$ is a polygonal region, bounded from below and from the right by a a polyline consisting of horizontal, vertical, and diagonal segments, and the regions A^t are sandwiched between such polylines; see Figure 1a. We call the polylines that form the boundary between consecutive regions A^t *boundary polylines*.

The B -map can be constructed in a similar, but easier manner. Indeed, B^t is the locus of all intervals such that \mathcal{Q} can win t voters (but no more) from $V \cap (x, y)$, and this is the case if and only if $|V \cap (x, y)| = t$. Hence, B^t is the union of the rectangular regions $[v_i, v_{i+1}] \times (v_{i+t}, v_{i+t+1}]$ (intersected with R), for $0 \leq i \leq n^* - t$, where $v_0 := -\infty$ and $v_{n^*+1} := \infty$, as shown in Figure 1b.

As mentioned, by overlaying the A - and B -map, we get the *gain map*. For any given region on this map, the intervals corresponding to points inside this region have equal α -gain and equal β -gain.

Lemma 9. *The complexity of the gain-map is $O((n^*)^2)$.*

Proof. The boundary polylines in the A -map are xy -monotone and comprised of vertical, horizontal, and diagonal lines. The B -map is essentially a grid of size $O((n^*)^2)$ defined by the lines $x = v_i$ and $y = v_i$, for $1 \leq i \leq n^*$. Since each of these lines intersects any xy -monotone polyline at most once—in a point or in a vertical segment—the complexity of the gain map is also $O((n^*)^2)$. \square

Using the gain map, we can compute the values $\text{span}(x_0, n, a, b)$ for a given $x_0 \in \mathbb{R}$ and for all triples n, a, b satisfying $1 \leq n \leq n^*$, and $0 \leq a \leq n^*$ and $0 \leq b \leq n^*$, as follows. First, we compute the intersection points of the vertical line $x = x_0$ with (the edges of) the gain map, sorted by increasing y -coordinates. (If this line intersects the gain map in a vertical segment, we take the topmost endpoint of the segment.) Let $(x_0, y_1), \dots, (x_0, y_z)$ denote this sorted sequence of intersection points, where $z \leq 2n^*$ denotes the number of intersections. Let a_i and b_i denote the α -gain and β -gain of the interval corresponding to the point (x_0, y_i) , and let a_{z+1} and b_{z+1} denote the α -gain and β -gain of the unbounded region intersected by the line $x = x_0$. Define $n_i = \text{argmax}_n v_n < y_i$. Then we have

$$\text{span}(x_0, n, a, b) = \begin{cases} y_i & \text{if } a = a_i \text{ and } b = b_i, \text{ and } n = n_i, \text{ for some } 1 \leq i \leq z \\ +\infty & \text{if } a = a_{z+1} \text{ and } b = b_{z+1} \text{ and } n = n^* \\ -\infty & \text{for all other triples } n, a, b \end{cases} \quad (3)$$

Our algorithm presented below moves a sweep line from left to right over the gain map. During the sweep we maintain the intersections of the sweep line with the gain map. It will be convenient to maintain the intersections with the A -map and the B -map separately. We will do so using two sequences, $A(x_0)$ and $B(x_0)$.

- The sequence $A(x_0)$ is the sequence of all diagonal or horizontal edges in the A -map that are intersected by the line $x = x_0$, ordered from bottom to top along the line. (More precisely, the sequence contains (at most) one edge for any boundary polyline. When the sweep line reaches the endpoint of such an edge, the edge will be removed and it will be replaced by the next non-vertical edge of that boundary polyline, if it exists.)
- The sequence $B(x_0)$ is the sequence of the y -coordinates of the horizontal segments in the B -map intersected by the line $x = x_0$, ordered from bottom to top along the line.

The number of intersections of the A -map, and also of the B -map, with the line $x = x_0$ is equal to $n^* - n_0 + 1$, where $n_0 = \operatorname{argmin}_n v_n > x_0$. Hence, the sequences $A(x_0)$ and $B(x_0)$ have length $n^* - n_0 + 1 \leq n^* + 1$.

If we have the sequences $A(x_0)$ and $B(x_0)$ available then, using Equation (3), we can easily find all triples n, a, b such that $\operatorname{span}(x_0, n, a, b) \neq -\infty$ (and the corresponding y -values) by iterating over the two sequences. We summarize the results of this section in the following observation.

Observation 10. *If we know the sequences $A(x_0)$ and $B(x_0)$ then we can compute all the values $\operatorname{span}(x_0, n, a, b)$, with $1 \leq n \leq n^*$ and $0 \leq a, b \leq n$, that are not equal to $-\infty$ in $O(n^*)$ time in total.*

This observation, together with Lemma 8 forms the basis of our dynamic-programming algorithm.

2.5 The sweep-line based dynamic-programming algorithm

Usually in a dynamic-programming algorithm, the value of a subproblem is computed by looking up the values of certain smaller subproblems. In our case it is hard to determine which smaller subproblems we need, so we take the opposite approach: whenever we have computed the value of a subproblem we determine which other subproblems can use this value, and we update their solutions if necessary. To this end we will use a sweep-line approach, moving a vertical line from left to right over the gain map. We will maintain a table X , indexed by subproblems, such that when the sweep line is at position x_0 , then $X[I]$ holds the best solution known so far for subproblem I , where the effect of all the subproblems with solution smaller than x_0 have been taken into account. When our sweep reaches a subproblem I' , then we check which later subproblems I can use I' in their solution, and we update the solutions to these subproblems.

Recall the algorithm works with a fixed threshold value $\tau \in \{1, \dots, \lfloor n^*/\ell^* \rfloor\}$ and that its goal is to compute the values $X_{\max}(n^*, k^*, \ell^*, \gamma, \delta)$ for all $0 \leq \gamma \leq n^*$ and $\delta \in \{\text{strict}, \text{loose}\}$. Our algorithm maintains the following data structures.

- $A[0..n^*]$ is an array that stores the sequence $A(x_0)$, where x_0 is the current position of the sweep line and $A[i]$ contains the i -th element in the sequence. When the i -th element does not exist then $A[i] = \text{NIL}$.
- Similarly, $B[0..n^*]$ is an array that stores the sequence $B(x_0)$.
- X : This is a table with an entry for each subproblem $I = \langle n, k, \ell, \gamma, \delta \rangle$ with $0 \leq n \leq n^*$, and $0 \leq k \leq k^*$ and $0 \leq \ell \leq \ell^*$, and $0 \leq \gamma \leq n^*$ and $\delta \in \{\text{strict}, \text{loose}\}$. When the sweep line is at position x_0 , then $X[I]$ holds the best solution known so far for subproblem I , where the effect of all the subproblems with solution smaller than x_0 have been taken into account. More precisely, in the right-hand side of the equation in Lemma 8 we have taken the maximum value over all subproblems $I' = \langle n', k - 1, \ell - j, \gamma - n + n' + a \oplus_j b, \delta' \rangle$ with $X_{\max}(I') < x_0$. In the beginning of the algorithm the entries for elementary subproblems are computed using Lemma 6 and all other entries have value $-\infty$.
- E : This is the event queue, which will contain four types of events, as explained below.

The event queue E is a min-priority queue on the x -value of the events. There are four types of events, as listed next, and when events have the same x -value then the first event type (in the list below) has higher priority, that is, will be handled first. When two events of the same type have equal x -value then their order is arbitrary. Note that events with the same x -value are not degenerate cases—this is inherent to the structure of the algorithms, as many events take place at x -coordinates corresponding to voters.

A -map events, denoted by $e_A(a, s, s')$: At an A -map event, the edge s of the A -map ends—thus the x -value of an A -map event is the x -coordinate of the right endpoint of s —and the array A

must be updated by replacing it with the edge s' . Here s' is the next non-vertical edge along the boundary polyline that s is part of, where $s' = \text{NIL}$ if s is the last non-vertical edge of the boundary polyline. The value a indicates that the edge s is on the boundary polyline between A^a and A^{a+1} . In other words s (and s' , if it exist) are the a -th intersection point, $0 \leq a < n^*$, with the A -map along the current sweep line, and so we must update the entry $A[a]$ by setting $A[a] \leftarrow s'$.

B -map events, denoted by $e_B(v_n)$: At a B -map event, a horizontal edge of the B -map ends. This happens when the sweep line reaches a voter v_n —that is, when $x_0 = x_n$ —and so the x -value of this event is v_n . The bottommost intersection of the sweep line with the B -map now disappears (see Figure 1b), and so we must update B by shifting all other intersection points one position down in B and setting $B[n^* - n] \leftarrow \text{NIL}$.

Subproblem events, denoted by $e_X(n', k', \ell', \gamma', \delta')$: At a subproblem event the solution to the subproblem given by $I' = \langle n', k', \ell', \gamma', \delta' \rangle$ is known and the x -value of this event is equal to $X_{\max}(I')$. Handling the subproblem event for I' entails deciding which later subproblems I can use I' in their solution and how they can use it, using the sets $\Delta(\tau, a, b)$, and updating the solutions to these subproblems.

In the beginning of the algorithm all the events associated with elementary subproblems are known. The events associated with non-elementary subproblems are added to the event queue when handling an update event $e_E(v_n)$, as discussed next.

Update events, denoted by $e_E(v_n)$: At the update event happening at x -value v_n , all subproblem events of size n are added to the event queue E . These are simply the subproblems $\langle n, k, \ell, \gamma, \delta \rangle$ for all $k, \ell, \gamma \in \{0, \dots, n\}$ and $\delta \in \{\text{strict}, \text{loose}\}$. The reason we could not add them at the start of the algorithm was that the x -value of such a subproblem I was now known yet. However, when we reach v_n then $X_{\max}(I)$ is determined, so we can add the event to E with $X_{\max}(I)$ as its x -value.

The pseudocode below summarizes the algorithm.

Lemma 11. *Algorithm 1 correctly computes the solutions for subproblems $\langle n, k, \ell, \gamma, \delta \rangle$ for the given value τ , for all $n, k, \ell, \gamma, \delta$ with $0 \leq n \leq n^*$, and $0 \leq k, \ell, \gamma \leq n$, and $\delta \in \{\text{strict}, \text{loose}\}$, and $\ell < 2(k + 1)$. The running time of the algorithm is $O(k^* \ell^* (n^*)^3)$.*

Proof. We handle the A -map and B -map events before a subproblem event so that A and B data structures are up-to-date when we want to compute the span function on handling a subproblem event. We also handle a subproblem event before an update event so that when we want to add a new subproblem event to the event queue on handling an update event, its entry in table X has the correct value. The correctness of the algorithm now follows from the discussion and lemmas above.

The running time is dominated by the handling of the subproblem events. By Observation 10, the algorithm handles each subproblem in $O(n^*)$ time, plus $O(\log n^*)$ for operations on the event queue, and there are $O(k^* \ell^* (n^*)^2)$ subproblems. Hence, the total running time is $O(k^* \ell^* (n^*)^3)$. \square

By Lemmas 5 and 11, the algorithm described at the beginning of Section 2.2 computes $\Gamma_{k^*, \ell^*}(V)$ correctly. Since this algorithm calls `COMPUTESOLUTIONS` $\lfloor n^*/\ell^* \rfloor$ times in Step 1, we obtain the following theorem.

Theorem 12. *There exists an algorithm that computes $\Gamma_{k^*, \ell^*}(V)$, and thus solves the one-dimensional case of the one-round discrete Voronoi game, in time $O(k^* (n^*)^4)$.*

Algorithm 1: COMPUTESOLUTIONS(τ, V, k^*, ℓ^*)

```
1 for  $i \leftarrow 0$  to  $n^* - 1$  do
2    $\lfloor A[i] \leftarrow (v_i, v_{i+1}) - (v_{i+1}, v_{i+1}); \quad B[i] \leftarrow v_{i+1}$   $\triangleright$  define  $v_0 := v_1 - 1$ 
3    $A[n^*] \leftarrow \text{NIL}; \quad B[n^*] \leftarrow \text{NIL}$ 
4   Initialize  $X$  by the solutions to elementary subproblems
5   Initialize  $E$  by all map events, update events, and elementary subproblem events
6   while  $E$  is not empty do
7      $e \leftarrow \text{extractMin}(E); \quad x_0 \leftarrow x\text{-value of } e$ 
8     switch  $e$  do
9       case  $e_A(a, s, s')$  do
10         $\lfloor A[a] \leftarrow s'$ 
11       case  $e_E(v_n)$  do
12         $B[n^* - n] \leftarrow \text{NIL}$ 
13        for  $i \leftarrow 0$  to  $n^* - n - 1$  do
14           $\lfloor B[i] \leftarrow v_{n+i+1}$ 
15       case  $e_X(n', k', \ell', \gamma', \delta')$  do
16        for all  $\text{span}(x_0, n, a, b) = y$  where  $y \neq -\infty$  do
17          for all  $(\delta', j, \delta) \in \Delta(\tau, a, b)$  do
18             $I \leftarrow \langle n, k' + 1, \ell' + j, \gamma' + n - n' - \text{gain}^j(a, b), \delta \rangle$ 
19             $\lfloor X(I) \leftarrow \max(X(I), y)$ 
20       case  $e_E(v_n)$  do
21         $\lfloor$  Add all the events for subproblems of size  $n$  to  $E$ , as explained above
```

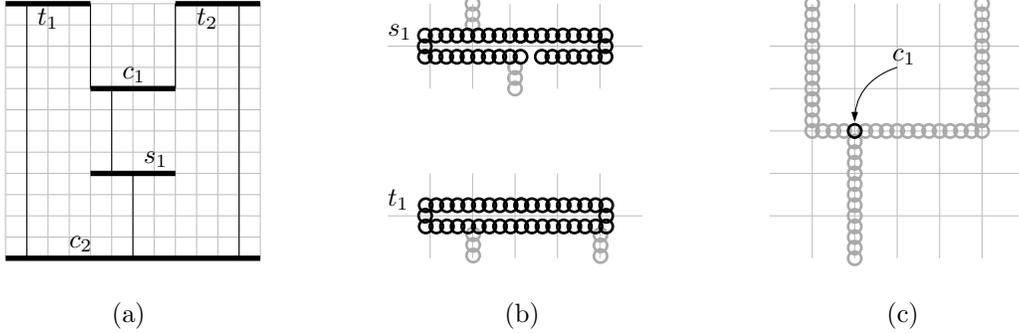


Figure 2: **a)** Bar graph of the boolean expression $\forall s_1 \exists t_1, t_2 : c_1 \wedge c_2$, where $c_1 := \bar{s}_1 \vee t_1 \vee t_2$ and $c_2 := s_1 \vee t_1 \vee \bar{t}_2$. **b)** Representation of variables in the transformed graph. **c)** Representation of a clause in the transformed graph.

Remark. We can also solve the one-dimensional case of the one-round discrete Voronoi game when voters are weighted, i.e. each voter $v \in V$ has an associated weight $\omega(v)$ and the players try to maximize the total weight of the voters they win. In this case, the α -gain and β -gain of an interval is defined as the total weight of voters the second player can win in that interval by placing one point and two points, respectively. The number of possible thresholds is not an integer in range $[0, n^*]$, but the sum of any sequence of consecutive voters define a threshold, which makes a total of $O((n^*)^2)$ different thresholds. The gain map also becomes more complex and in the algorithm we need to spend $O((n^*)^2)$ time (instead of $O(n^*)$) to handle each subproblem event, which results in an algorithm with running time $O(k^* \ell^* (n^*)^5)$.

3 Σ_2^P -Hardness for $d \geq 2$

In this section we prove that the one-round discrete Voronoi game is Σ_2^P -hard in \mathbb{R}^2 , which implies hardness for $d > 2$ as well. To prove this, it suffices to show that deciding if \mathcal{Q} has a winning strategy against every possible strategy of \mathcal{P} is Π_2^P -hard. Our proof will use a reduction from a special case of the quantified Boolean formula problem (QBF), as defined next. Let $S := \{s_1, \dots, s_{n_s}\}$ and $T := \{t_1, \dots, t_{n_t}\}$ be two sets of variables, and let $\bar{S} := \{\bar{s}_1, \dots, \bar{s}_{n_s}\}$ and $\bar{T} := \{\bar{t}_1, \dots, \bar{t}_{n_t}\}$ denote their negated counterparts. We consider Boolean formulas B of the form

$$B := \forall s_1, \dots, s_{n_s} \exists t_1, \dots, t_{n_t} : c_1 \wedge \dots \wedge c_{n_c}$$

where each clause c_i in $C := \{c_1, \dots, c_{n_c}\}$ is a disjunctive combination of at most three literals from $S \cup \bar{S} \cup T \cup \bar{T}$. Deciding if a formula of this form is true is a Π_2^P -complete problem [18].

Consider the undirected graph $G_B := (N, A)$ representing B , where $N := S \cup T \cup C$ is the set of nodes of G_B and $A := \{(c_i, s_j) : s_j \in c_i \vee \bar{s}_j \in c_i\} \cup \{(c_i, t_j) : t_j \in c_i \vee \bar{t}_j \in c_i\}$ is the set of edges of G_B . Lichtenstein [12] showed how to transform an instance of QBF in polynomial time to an equivalent one whose corresponding graph is planar (and of quadratic size). We can use the same technique here. Hence, we may start our reduction from an Boolean formula B such that G_B is planar. We call the resulting problem PLANAR $\forall\exists 3$ -CNF.

In the following, we transform an instance of PLANAR $\forall\exists 3$ -CNF to an instance $\langle V, k, \ell \rangle$ of the Voronoi game problem in the plane such that B is true if and only if \mathcal{Q} has a winning strategy.

The first (and standard) step in our reduction is to construct a specific embedding of the planar graph G_B . More precisely, we use a *bar graph*, where each node is represented by a horizontal segment

and each edge is represented by a vertical segment; see Figure 2a. Rosenstiehl and Tarjan [16] showed that such a representation always exists. Before we describe the specific variable, clause, and edge gadgets that we use, it is useful to make the following observation about when \mathcal{Q} wins a certain voter $v_i \in V$, when the strategy P of \mathcal{P} is fixed. Define D_i , the *disk of v_i* with respect to the given set P , as the disk with center v_i and radius $\text{dist}(v_i, P)$, that is, D_i is the largest disk centered at v_i that has no point from P in its interior. The key to our reduction is the following simple observation.

Observation 13. *Player \mathcal{Q} wins a voter $v_i \in V$ against P if and only if she places a point q in the interior of D_i .*

Next we describe how to transform the bar graph G_B into a voter set V . Because of the above observation, it will be convenient to imagine that each voter is surrounded by a disk, and talk about placing disks. Thus, whenever we say we place a disk somewhere, in our construction we actually place a voter at the center of the disk. Later we will then put more voters that will force most of \mathcal{P} 's strategy, so that these disks become meaningful. The nodes of G_B are replaced by the following gadgets:

Clause gadgets. Each node $c_i \in C$ is transformed to a single disk as shown in Figure 2c. Recall that each such node is a horizontal segment, which has three incoming edges from its constituent literals. The position of the voter for c_i is (roughly) the point where the middle edge arrives at the segment.

Variable gadgets. Each node t_i of degree $\text{deg}(t_i)$ in graph G_B is transformed to a ring of an even number of disks numbered in counterclockwise order starting from any arbitrary disk, which we call a *closed necklace*, containing at least $2 \text{deg}(t_i)$ disks as shown in Figure 2b.

Each node s_i of degree $\text{deg}(s_i)$ in graph G_B is transformed to a ring of disks with one disk missing numbered in counterclockwise order starting from the disk after the missing disk, which we call an *open necklace*, as shown in Figure 2b. An open necklace has odd size and at least $2 \text{deg}(s_i) + 1$ disks. We can assume the distance between the two disks in the place where the necklace is open is exactly 1. Let D_1 and D_2 be the two disks at distance 1, and let $z_1 \in \partial D_1$ and $z_2 \in \partial D_2$ be the closest pair of points on the boundaries of these two disks. (Thus $\text{dist}(z_1, z_2) = 1$.) Then we place a cluster of w voters at z_1 and we place another cluster of w voters at z_2 , where w is a suitable number specified later and a *cluster* of voters is simply a number of coinciding voters.

Edge gadgets. Each edge $\{c_i, s_j\}$ or $\{c_i, t_j\}$ is replaced by a chain of disks of even length that in one end intersects a pair of consecutive disks in a necklace corresponding to the node s_j or t_j , respectively, such that the first disk has an odd position and the next disk in clockwise order has an even position, and in the other end intersects the disk associated with the clause c_i

Similarly, each edge $\{c_i, \bar{s}_j\}$ or $\{c_i, \bar{t}_j\}$ is replaced by a chain of disks of even length that in one end intersects a pair of consecutive disks in a necklace corresponding to the node s_j or t_j , respectively, such that the first disk has an even position and the next disk in clockwise order has an odd position, and in the other end intersects the disk associated with the clause c_i .

In our transformation of G_B to an instance of the Voronoi game, we will use that V can be a multiset, where we will use multiplicities $1, w, w + 1, \lfloor w/2 \rfloor + 1, \lceil w/2 \rceil + 1$, and W for the voters, for suitably chosen values w and W . We call a cluster of i voters an *i -cluster*. We denote the multisets containing clusters of $1, w$, and W voters by V_1, V_w , and V_W , respectively. The values w and W will be chosen

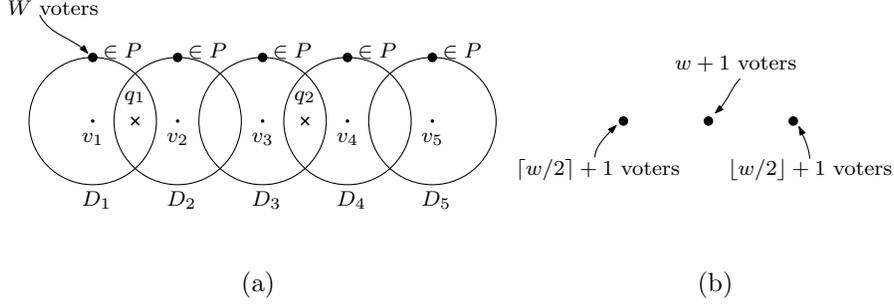


Figure 3: **a)** When all the heavy-weight clusters of W voters are chosen by \mathcal{P} , the best strategy of \mathcal{Q} to win the remaining single voters is to put his points in every other intersection of the disks of the voters. **b)** An example of a balancing gadget.

such that they satisfy

$$w = |V_1| + 1, \tag{4}$$

$$W = |V_w| + |V_1| + 1. \tag{5}$$

We have already described the placement of voters in V_1 by describing the placement of the disks in defining the gadgets. In our construction, for each voter $v \in V_1$, we consider an imaginary disk $z(v, r)$ which is centered in the position of the voter v and has a radius r , which is a real number in range $10 \leq r \leq 20$. Each disk has a W -cluster inside it near its boundary which is denoted by $W(v)$. Two disks $z_1(v_1, r_1)$ and $z_2(v_2, r_2)$ either do not intersect or their interior have a non-empty intersection (see Figure 3a). Similarly when three disks intersect, their interior have a non-empty intersection. We also described the placement of voters in V_w when we defined the variable gadgets for the variables s_i .

We call this construction *the transformed graph*. Now we add $n' := |V_W| - |V_1| + 1$ gadgets, which we call *balancing gadgets*, each comprised of three clusters of sizes $\lfloor w/2 \rfloor + 1$, $\lceil w/2 \rceil + 1$, and $w + 1$ far away from the graph of disks and from each other as shown in Figure 3b. Let V' be the multiset of all these voters.

We define $V := V_1 \cup V_w \cup V_W \cup V'$, $k := \|V_W\| + n' + n_s$, $\ell := (|V_1| - n_c - n_s)/2 + n_s + 2n'$, where $\|V_W\|$ denotes the number of distinct points in V_W .

Lemma 14. \mathcal{Q} has a winning strategy in $\langle V, k, \ell \rangle$ if and only if B is true.

Proof. First we show that in an optimal strategy, \mathcal{P} places his points on the most valuable positions, i.e. $\|V_W\|$ of his points on W -clusters, n' of his points on $(w + 1)$ -clusters, and the remaining n_s points on w -clusters where exactly one w -cluster is selected from each open necklace. Note that with such a strategy for \mathcal{P} , the number of remaining voters for \mathcal{Q} is $n_s w + (w + 2)n' + |V_1|$.

For a contradiction, assume \mathcal{P} does not select some of the W -clusters. In that case \mathcal{P} can either move some of his points from balancing gadgets in which he has two or more points or some of his points from the transformed graph which are not exactly on a W -cluster to the uncovered W -clusters and get a better gain; because each extra point in a balancing gadget has a gain of at most $2w + 3 < W$ and all the voters of the transformed graph in V_1 and V_w together has a gain of at most $|V_w| + |V_1| < W$. Therefore, all the W -clusters will be covered by \mathcal{P} . Moreover, \mathcal{P} must also select all the $(w + 1)$ clusters; because by doing this, all the points of \mathcal{Q} have a gain of at most $2w + 2$ which is less than the gain of an unguarded balancing gadget, i.e. $2w + 3$.

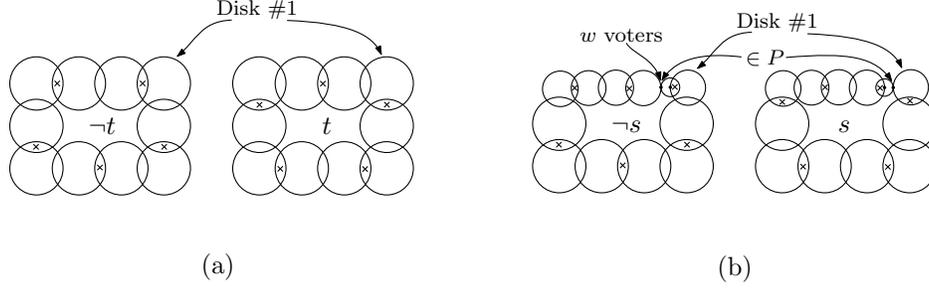


Figure 4: **a)** The two different ways of winning all voters for \mathcal{Q} . **b)** \mathcal{P} can force \mathcal{Q} 's game.

It is obvious that \mathcal{P} does not place more than one point in each balancing gadget as the gain of each extra point in a balancing gadget is at most $\lceil w/2 \rceil + 1$, but the guaranteed gain of a point on a w -cluster is w . Now, we just need to show that \mathcal{P} places all his remaining n_s points on w -clusters. This is also easy to see as if more than n_s w -clusters are uncovered, \mathcal{Q} can easily win $(n_s + 1)w + (w + 2)n'$ voters which is more than the total number of remaining voters $n_s w + (w + 2)n' + |V_1|$ for \mathcal{Q} in the case where at most n_s w -clusters are uncovered by \mathcal{P} .

It is also easy to see that in an optimal strategy, \mathcal{P} selects exactly one w -cluster in each open necklace. As otherwise, at least one pair of w -clusters remains unselected and \mathcal{Q} can gain $2w$ voters using just one point, while if exactly one cluster from each pair of w -clusters is selected by \mathcal{P} , the maximum possible gain of each point of \mathcal{Q} is at most $w + 1$ voters.

Now, assuming \mathcal{P} plays an optimal strategy, the gain of each point of \mathcal{Q} is at most

- $w + 1$ voters for a total of n_s points if he puts his point close to an uncovered w -cluster,
- $\lceil w/2 \rceil + 1$ or $\lfloor w/2 \rfloor + 1$ voters for a total of $2n'$ points if he puts his points in balancing clusters, and
- 2 or 3 voters if he puts his points in the intersection of disks in the transformed graph.

This shows that in her optimal strategy \mathcal{Q} places all her $2n'$ points in the balancing clusters to win all the remaining voters there, and the remaining $(|V_1| - n_c - n_s)/2 + n_s$ points of \mathcal{Q} will be placed on the transformed graph.

To win the remaining voters of a necklace efficiently, \mathcal{Q} has two choices. She can either put her points in odd intersections or in even intersections of the necklace as shown in Figure 4a. In both cases the number points she spends is equal to half the number of disks in the necklace. To win the remaining voters of an open necklace, \mathcal{Q} has just one choice. We note that the w -cluster selected by \mathcal{P} effectively changes the open necklace to a closed necklace with one more disk (see Figure 4b). Depending on the w -cluster chosen by \mathcal{P} , player \mathcal{Q} must put her points in odd or even intersections as shown in Figure 4b. In either case \mathcal{Q} spends $(m + 1)/2$ points, where m is the number of disks in the open necklace. To win all the remaining voters in a chain of m disks, \mathcal{Q} should place some points in every other intersection of the chain in odd positions and he spends $m/2$ points. Alternatively, she can place her $m/2$ points in even positions and also win the point in the clause on the one end of the chain only if the other end of the chain has already been covered by one of his points in the (open) necklace.

The placement of points by \mathcal{Q} in odd intersections of a necklace corresponds to a true value for the associated t_j , and placement in even intersections corresponds to a false value for t_j . Similarly, selection of the starting w -cluster of an open necklace by \mathcal{P} corresponds to a true value for the associated s_j , and selection of the ending w -cluster corresponds to a false value for s_j . When \mathcal{Q} has

a choice to put his points in even intersections of a chain and win the voter of the associated clause at the end of the chain, it corresponds to a true value for the associated literal of the (open) necklace at the other end of the chain which gives the clause a true value. Intuitively, it is clear that given an optimal strategy P , \mathcal{Q} should use all his points to win the remaining voters and he can win the voter associated with each clause if and only if she can satisfy that clause.

More formally, given an optimal strategy P of \mathcal{P} , we set s_j to false if the starting w -cluster in clockwise order in the open necklace of the gadget for s_j is selected in P and set it to true otherwise; we call this set of values $\widehat{s}(P)$. By using the discussion in the two previous paragraphs, we can easily verify that if there exists an assignment \widehat{t} for the variables $t_j (j = 1, \dots, n_t)$ such that $\langle \widehat{s}(P), \widehat{t} \rangle$ satisfies $c_1 \wedge \dots \wedge c_{n_c}$, then \mathcal{Q} can win all the $n_s w + (w + 2)n' + |V_1|$ remaining voters. Conversely, it is not hard to see that if \mathcal{Q} can win all the $n_s w + (w + 2)n' + |V_1|$ remaining voters, then there exists an assignment \widehat{t} for t_j s such that B is true.

Considering the fact that the number of remaining voters is exactly one more than the number of voters directly won by P , i.e.

$$(n_s w + (w + 2)n' + |V_1|) - (\|V_W\| + n_s w + (w + 1)n') = 1,$$

\mathcal{Q} has a winning strategy in $\langle V, k, \ell \rangle$ if and only if B is true. □

Lemma 14 is not enough to show that the Voronoi game is Σ_2^P -hard. We need to show that the reduction can be done in polynomial time and therefore the resulting Voronoi game has polynomial size.

We can easily generate balancing gadgets. Therefore we focus on drawing the transformed graph. As stated earlier, in order to define the exact position of clusters in the transformed graph we use a method devised by Rosenstiehl and Tarjan in [16] to draw a planar graph on a grid of size $O(n) \times O(n)$ where each node is represented by a horizontal line segment and each edge is represented by a vertical line segment (see Figure 2a). As the graph G_B is planar, we can draw it using this method. An (open) necklace associated with a variable t_j or s_j can be easily drawn as shown in Figure 2b. With a big enough cell-size for the grid, say 1000, we can adjust the disk sizes to get the desired properties for the intersections and also for the distance 1 between two ending disks in an open necklace. As a node associated with a clause has a degree of at most three, it can be drawn as in Figure 2c, and similarly with a big enough cell-size for the grid, we can adjust the size of the disks to get an even number of disks for each chain. The W -clusters can be placed at a distance of at most 0.1 at a rational point near the border of each disk, and placing w -clusters is also trivial.

As the total number of nodes in G_B is $O(n_s + n_t + n_c)$, we need a grid of size $O(n_s + n_t + n_c) \cdot O(n_s + n_t + n_c)$, and in worst case we have a constant number of disks on each edge of this grid. Therefore, the number of disk of $|V_1|$ is bounded from above by the size of the grid, i.e. $|V_1| \in O((n_s + n_t + n_c)^2)$, which means $|V| \in O(|V_1|W + n_s w + |V_1| + (2w + 3)n') = O((n_s + n_t + n_c)^5)$. The following theorem is the result of this discussion.

Theorem 15. *The Voronoi game problem is Σ_2^P -hard for $d > 1$.*

We note that the final construction is rectilinear and therefore this reduction also works for all the L_p norms as disks in the final construction can be replaced by L_p -disks and all the required properties still hold. Moreover, as this reduction is done for the two-dimensional case, this hardness result is also true for the L_∞ norm, because in \mathbb{R}^2 , the disks of the L_∞ norm are similar to those in the L_1 norm, just rotated by $\pi/4$.

4 Containment in $\exists\forall\mathbb{R}$ and the algorithm for $d \geq 2$

We now consider the one-round discrete Voronoi game in the L_p -norm, for some arbitrary p . Then a strategy $P = \{p_1, \dots, p_k\}$ can win a voter $v \in V$ against a strategy $Q = \{q_1, \dots, q_\ell\}$ if and only if the following Boolean expression is satisfied:

$$\text{win}(v) := \bigvee_{i \in [k]} \bigwedge_{j \in [\ell]} (\text{dist}_p(p_i, v))^p \leq (\text{dist}_p(q_j, v))^p,$$

where dist_p is the L_p -distance. This expression has $k\ell$ polynomial inequalities of degree p . The strategy P is winning if and only if the majority of the expressions $\text{win}(v_1), \dots, \text{win}(v_n)$ are true. Having a majority function *Majority* that evaluates to true if at least half of its parameters evaluates to true, player \mathcal{P} has a winning strategy if and only if

$$\begin{aligned} & \exists x_1(p_1), \dots, x_d(p_1), \dots, x_1(p_k), \dots, x_d(p_k) \\ & \forall x_1(q_1), \dots, x_d(q_1), \dots, x_1(q_\ell), \dots, x_d(q_\ell) : \quad \text{Majority}(\text{win}(v_1), \dots, \text{win}(v_n)) \end{aligned}$$

is true, where $x_i(\cdot)$ denotes the i -th coordinate of a point.

Ajtai *et al.* [2] show that it is possible to construct a sorting network, often called the AKS sorting network, composed of comparison units configured in $c \cdot \log n$ levels, where c is a constant and each level contains exactly $\lfloor n/2 \rfloor$ comparison units. Each comparison unit takes two numbers as input and outputs its input numbers in sorted order. Each output of a comparison unit (except those on the last level) feeds into exactly one input of a comparison unit in the next level, and the input numbers are fed to the inputs of the comparison units in the first level. The outputs of the comparison units in the last level (i.e., the outputs of the network) give the numbers in sorted order.

Using AKS sorting networks we can construct a Boolean formula of size $O(n^c)$ for some constant c that tests if the majority of its n inputs are true as follows. Assuming the boolean value *false* is smaller than the boolean value *true* value, we make an AKS sorting network that sorts n boolean values. This is possible using comparison units that get p and q as input, and output $p \wedge q$ and $p \vee q$. It is not hard to verify that the $\lceil n/2 \rceil$ -th output of the network is equal to the value of the majority function on the input boolean values. By construction, we can write the Boolean formula representing the value of this output as *logical and* (\wedge) and *logical or* (\vee) combination of the input boolean values, and the size of the resulting formula is $O(n^c)$.

Thus we can write $\text{Majority}(\text{win}(v_1), \dots, \text{win}(v_n))$ as a Boolean combination of $O(n^c k \ell)$ polynomial inequalities of degree p , where each quantified block has kd and ℓd variables respectively. Basu *et al.* [7] gave an efficient algorithm for deciding the truth of quantified formulas. For our formula this gives an algorithm with $O((n^c k \ell)^{(kd+1)(\ell d+1)} p^{k \ell d^2})$ running time to decide if \mathcal{P} has a winning strategy for a given instance $\langle V, k, \ell \rangle$ of the Voronoi game problem. Note that this is polynomial when k, ℓ and d are constants.

For the L_∞ norm, we can define $F(v)$ as follows:

$$F(v) := \bigvee_{i \in [k]} \bigwedge_{j \in [\ell]} \bigvee_{s' \in [d]} \bigwedge_{s \in [d]} |x_s(p_i) - x_s(v)| \leq |x_{s'}(q_j) - x_{s'}(v)|,$$

By comparing the squared values instead of the absolute values, we have a formula which demonstrates that even with the L_∞ norm, the problem is contained in $\exists\forall\mathbb{R}$ and there exists an algorithm of complexity $O((n^c k \ell d^2)^{(kd+1)(\ell d+1)} 2^{k \ell d^2})$ to solve it.

Theorem 16. *The one-round discrete Voronoi game $\langle V, k, \ell \rangle$ in \mathbb{R}^d with the L_p norm is contained in $\exists\forall\mathbb{R}$. Moreover, for fixed k, ℓ, d there exists an algorithm that solves it in polynomial time.*

De Berg *et al.* [9] introduced the notion of personalized preferences. More precisely, given a natural number p , assuming each axis defines an aspect of the subject voters are voting for, the voter v_i gives different weights to different axes, and v_i has a weighted L_p distance $(\sum_{j \in [d]} w_{ij}(x_j(p) - x_j(v_i))^p)^{1/p}$ from any point $p \in \mathbb{R}^d$. For the weighted L_∞ distance, v_i is at distance $\max_{j \in [d]} (w_{ij}|x_j(p) - x_j(v_i)|)$ from any point $p \in \mathbb{R}^d$. This approach also works when voters have personalized preferences.

5 Concluding remarks

We presented the first polynomial-time algorithm for the one-round discrete Voronoi game in \mathbb{R}^1 . The algorithm is quite intricate, and it would be interesting to see if a simpler (and perhaps also faster) algorithm is possible. Finding a lower bound for the 1-dimensional case is another open problem.

We also showed that the problem is Σ_2^P -hard in \mathbb{R}^2 . Fekete and Meijer [11] conjectured that finding an optimal strategy for the multi-round continuous version of the Voronoi game is PSPACE-complete. We conjecture that in the multi-round version of the discrete version, finding an optimal strategy is PSPACE-hard as well. Note that using the algebraic method presented in this paper, it is easy to show that this problem is contained in PSPACE. While the algebraic method we used is considered a standard technique, it is, as far as we know, the first time this method is combined with polynomial-size boolean formulas for the majority function. We think it should be possible to apply this combination to other problems as well.

References

- [1] Hee-Kap Ahn, Siu-Wing Cheng, Otfried Cheong, Mordecai Golin, and Rene Van Oostrum. Competitive facility location: the voronoi game. *Theoretical Computer Science*, 310(1-3):457–467, 2004.
- [2] Miklós Ajtai, János Komlós, and Endre Szemerédi. Sorting in $c \log n$ parallel steps. *Combinatorica*, 3(1):1–19, 1983.
- [3] Aritra Banik, Bhaswar B Bhattacharya, and Sandip Das. Optimal strategies for the one-round discrete voronoi game on a line. *Journal of Combinatorial Optimization*, 26(4):655–669, 2013.
- [4] Aritra Banik, Bhaswar B Bhattacharya, Sandip Das, and Sreeja Das. Two-round discrete voronoi game along a line. In *Frontiers in Algorithmics and Algorithmic Aspects in Information and Management*, pages 210–220. Springer, 2013.
- [5] Aritra Banik, Bhaswar B Bhattacharya, Sandip Das, and Satyaki Mukherjee. The discrete voronoi game in \mathbb{R}^2 . *Computational Geometry*, 63:53–62, 2017.
- [6] Aritra Banik, Jean-Lou De Carufel, Anil Maheshwari, and Michiel Smid. Discrete voronoi games and epsilon-nets, in two and three dimensions. *Computational Geometry*, 55:41–58, 2016.
- [7] Saugata Basu, Richard Pollack, and Marie-Françoise Roy. On the combinatorial and algebraic complexity of quantifier elimination. *Journal of the ACM*, 43(6):1002–1045, 1996.

- [8] Otfried Cheong, Sariel Har-Peled, Nathan Linial, and Jiří Matoušek. The one-round voronoi game. *Discrete & Computational Geometry*, 31(1):125–138, 2004.
- [9] Mark de Berg, Joachim Gudmundsson, and Mehran Mehr. Faster algorithms for computing plurality points. *ACM Transactions on Algorithms*, 14(3):36:1–36:23, 2018.
- [10] Michael G Dobbins, Linda Kleist, Tillmann Miltzow, and Paweł Rzażewski. $\forall\exists\mathbb{R}$ -completeness and area-universality. In *Proceedings of the 44th International Graph-Theoretic Concepts in Computer Science (WG 2018)*, pages 164–175, 2018.
- [11] Sándor P Fekete and Henk Meijer. The one-round voronoi game replayed. *Computational Geometry*, 30(2):81–94, 2005.
- [12] David Lichtenstein. Planar formulas and their uses. *SIAM Journal on Computing*, 11(2):329–343, 1982.
- [13] Wei-Yin Lin, Yen-Wei Wu, Hung-Lung Wang, and Kun-Mao Chao. Forming plurality at minimum cost. In *Proceedings of the 9th International Workshop on Algorithms and Computation*, pages 77–88, 2015.
- [14] Richard D McKelvey and Richard E Wendell. Voting equilibria in multidimensional choice spaces. *Mathematics of operations research*, 1(2):144–158, 1976.
- [15] Michael Paterson and Uri Zwick. Shallow circuits and concise formulae for multiple addition and multiplication. *Computational Complexity*, 3(3):262–291, 1993.
- [16] Pierre Rosenstiehl and Robert E Tarjan. Rectilinear planar layouts and bipolar orientations of planar graphs. *Discrete & Computational Geometry*, 1(4):343–353, 1986.
- [17] Joachim Spoerhase and H-C Wirth. (r, p) -centroid problems on paths and trees. *Theoretical Computer Science*, 410(47-49):5128–5137, 2009.
- [18] Larry J Stockmeyer. The polynomial-time hierarchy. *Theoretical Computer Science*, 3(1):1–22, 1976.
- [19] Sachio Teramoto, Erik D Demaine, and Ryuhei Uehara. The voronoi game on graphs and its complexity. *Journal of Graph Algorithms and Applications*, 15(4):485–501, 2011.
- [20] Yen-Wei Wu, Wei-Yin Lin, Hung-Lung Wang, and Kun-Mao Chao. Computing plurality points and condorcet points in euclidean space. In *International Symposium on Algorithms and Computation*, pages 688–698, 2013.