

Nash equilibria in Voronoi games on graphs

Christoph DÜRR*

NGUYEN KIM Thang*

Abstract

In this paper we study a game where every player is to choose a vertex (facility) in a given undirected graph. All vertices (customers) are then assigned to closest facilities and a player's payoff is the number of customers assigned to it. We show that deciding the existence of a Nash equilibrium for a given graph is \mathcal{NP} -hard. We also introduce a new measure, the *social cost discrepancy*, defined as the ratio of the costs between the worst and the best Nash equilibria. We show that the social cost discrepancy in our game is $\Omega(\sqrt{n/k})$ and $O(\sqrt{kn})$, where n is the number of vertices and k the number of players.

1 Introduction

Summer is a holiday season for everyone and it is also the season of competition between ice-cream sellers on beaches. On a colorful crowded beach, an ice-cream seller need to choose a “good” location that maximize his profit known that tourists will go to the closest seller to buy ice-creams. A similar competition holds between service providers, enterprises to decide where to open a new facility, a new market in order to attract as much of clients as possible.

The Voronoi Game is a simple geometric model for the competitive facility location. Competitive facility location studies the placement of sites by competing market players. Voronoi game between two players is widely studied on a continuous space, for example on a 2-dimensional rectangle. Players alternatively place points in the space. Then the Voronoi diagram is considered. Every player gains the total surface of the Voronoi cells of his points. The geometric concepts are combined with game theory arguments to study if there exists any winning strategy.

We consider the discrete version of the Voronoi game which plays on a given graph instead on a continuous space. Formally, the discrete Voronoi game plays on a given undirected graph $G(V, E)$ with $n = |V|$ and k players. Every player has to choose a vertex (*facility*) from V , and every vertex (*customer*) is assigned to the closest facilities. If there are more than one closest facility then the vertex is assigned in equal fraction to these closest facilities. One may think that each vertex consists of a group of clients in which one half go to a facility and the other half go to another facility if there are two closest facilities to this group. A player's *payoff* (*utility*) is the number of vertices assigned to his facility. The *social cost* is the total distance that customers go to their closest facilities, i.e. it is defined as the sum of the distances to the closest facility over all vertices. The optimization problem with the objective function as this social cost is the well-known *k-median* problem which is \mathcal{NP} -complete [9, chapter 25].

*LIX, CNRS UMR 7161, Ecole Polytechnique 91128 Palaiseau, France. Supported by ANR Alpage.

Our contributions In the paper, first we show that the existence of Nash equilibria is a graph property for a fixed number of players. There exist Nash equilibria for some graphs but there are none for others. We show that the best-response dynamic does not converge to a Nash equilibria on cycle and line graphs, but does converge for a variant of this game. Moreover, we show that deciding this graph property is an \mathcal{NP} -hard problem.

We introduce a new measure of inefficiency of equilibria in a game, called *social cost discrepancy* defined as the ratio of the costs between the worst and the best Nash equilibria. We show that the social cost discrepancy in Voronoi game is $\Omega(\sqrt{n/k})$ and $O(\sqrt{kn})$. Hence for a constant number of players we have tight bounds. Interestingly, in analyzing this measure, we introduce the notion of Delaunay triangulation on discrete setting and that matches the usual Delaunay triangulation in continuous surface.

Related work Most prior work studies Voronoi game on continuous surface. Ahn et al. [1] were the firsts to demonstrate that there exists a winning strategy for the one dimensional arena case, i.e. a line segment $[0, 1]$. In two dimensional case, the scenario differs significantly. Intuitively, the difficulty is arisen since the Voronoi cells in one dimension are simply lines or curves but in two dimension, they becomes much more complicated. Cheong et al. [3], Fekete and Meijer [5] characterized the winning strategy for the game played on 2D rectangle as a function of width-height ratio of the rectangle.

The closest game related to our work is the facility location game introduced in [10]. The facility location game plays on bipartite weighted graph in which there are k suppliers, each chooses to open a single facility within the set of facilities, and m markets that the suppliers will serve. Given a strategy profile, a supplier serves the markets closest to it and receives a payment from these markets. A market receives a value if it is served. The payment in this game makes it fundamentally different to the Voronoi game. Vetta [10] proved that there always exist a Nash equilibrium in this game and the price of anarchy is at most 2 with respect to the total profit of the game (the formal definition is in [10]).

After the introduction of the Voronoi Games on graphs in the conference version of this paper [4], Mavronicolas et al. [8] study the game on cycle graphs, particularly on the characterization of equilibrium and on the price of anarchy on such graph. Subsequently, the game on transitive graphs is studied in [6]. Inspired by the model of the Voronoi Games on graphs, Zhao et al. [11], Bilò et al. [2] study the isolation game, which may be considered as the dual of the Voronoi game.

Organization In Section 2, we give the formal description of the Voronoi game and the definition of the social cost discrepancy. Next, we warm up by studying the game on cycle graph in Section 3. In Section 4, we prove that it is \mathcal{NP} -complete to decide whether a given game admits an equilibrium. We bound the social cost discrepancy of the game in Section 5

2 The Voronoi game and the social cost discrepancy

2.1 The Voronoi game

For this game we need to generalize the notion of vertex partition of a graph: A *generalized partition* of a graph $G(V, E)$ is a set of n -dimensional non-negative vectors, which sum up to the vector with 1 in every component, for $n = |V|$.

The Voronoi game on graphs consists of:

- A graph $G(V, E)$ and k players. We assume $k < n$ for $n = |V|$, otherwise the game has a trivial structure. The graph induces a distance between vertices $d : V \times V \rightarrow \mathbb{N} \cup \{\infty\}$, which is defined as the minimal number of edges of any connecting path, or infinite if the vertices are disconnected.
- The strategy set of each player is V . A strategy profile of k players is a vector $f = (f_1, \dots, f_k)$ associating each player to a vertex.
- For every vertex $v \in V$ — called *customer* — the distance to the closest facility is denoted as $d(v, f) := \min_{f_i} d(v, f_i)$. Customers are assigned in equal fractions to the closest facilities as follows. The strategy profile f defines the generalized partition $\{F_1, \dots, F_k\}$, where for every player $1 \leq i \leq k$ and every vertex $v \in V$,

$$F_{i,v} = \begin{cases} 1/|\arg \min_j d(v, f_j)| & \text{if } d(v, f_i) = d(v, f) \\ 0 & \text{otherwise.} \end{cases}$$

We call F_i the *Voronoi cell* of player i . The *radius* of the Voronoi cell of player i is defined as $\max_v d(v, f_i)$ where the maximum is taken over all vertices v such that $F_{i,v} > 0$.

- The *payoff (utility)* of player i is the (fractional) amount of customers assigned to it, that is $p_i := \sum_{v \in V} F_{i,v}$. (See figure 1 for an example.)
- The *social cost* of strategy profile f is $\text{cost}(f) := \sum_{v \in V} d(v, f)$.

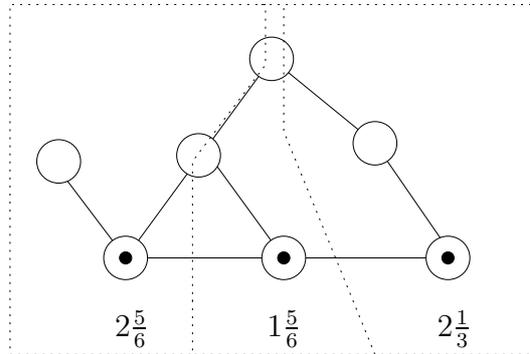


Figure 1: A strategy profile of a graph (players are dots) and the corresponding payoffs.

The *best response* for player i in the strategy profile f is a vertex $v \in V$ maximizing the player i 's payoff in the strategy profile which equals f except that the strategy of player i is v . If player i can strictly improve his payoff by choosing another strategy, we say that player i is *unhappy* in f , otherwise he is *happy*. The *best response dynamic* is the process of repeatedly choosing an arbitrary unhappy player, and change it to an arbitrary best response. A *pure Nash equilibrium* is defined as a fixed point to the best response dynamic, or equivalently as a strategy profile where all players are happy. In this paper we consider only pure Nash equilibria, so we omit from now on the adjective “pure”.

We defined players' payoffs in such a way, that there is a subtle difference between the Voronoi game played on graphs and the Voronoi game played on a continuous surface. Consider a situation where a player i moves to a location already occupied by a single player j who is not neighbor of i . In the continuous case player i gains exactly a half of the previous payoff of player j (since it is now shared with i). However, in our setting (the discrete case), player i can sometimes gain more than a half of the previous payoff of player j (see figure 3).

A simple observation leads to the following bound on the players payoff.

Lemma 1 *In a Nash equilibrium the payoff p_i of every player i is bounded by $n/2k < p_i < 2n/k$.*

Proof: If a player gains p and some other player moves to the same location then both payoffs are at least $p/2$. Therefore the ratio between the largest and the smallest payoffs among all players can be at most 2. If all players have the same payoff, it must be exactly n/k , since the payoffs sum up to n . Otherwise there is at least one player who gains strictly less than n/k , and another player who gains strictly more than n/k . This concludes the proof. \square

2.2 The social cost discrepancy

When consider the inefficiency of equilibria in a game, typical measures are the price of anarchy and the price of stability. The price of anarchy (stability) is defined as the ratio between the worst (best) objective value of an equilibrium of the game and that of an optimal solution.

We introduce the *social discrepancy* defined as the ratio between the worst and the best pure equilibrium. The idea is that a small social cost discrepancy guarantees that the social costs of Nash equilibria do not differ too much, and measures a degree of choice in the game. Additionally, in some settings it may be unfair to compare the cost of a Nash equilibrium with the optimal solution, which may not be attained by selfish players or may not be an outcome of the game.

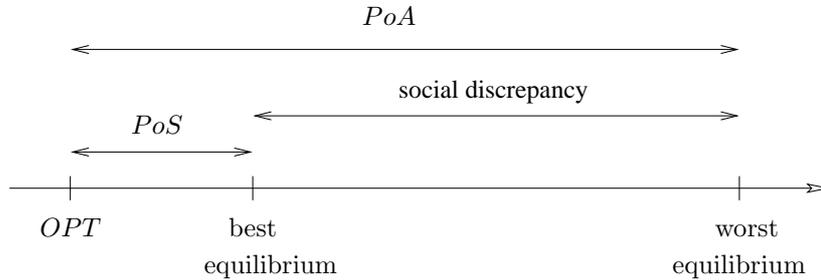


Figure 2: Illustration of different measures of inefficiency.

Note that the social discrepancy is not the ratio between the price of anarchy and the price of stability since each of these measures may be attained by different instances of a game.

3 The cycle graph

Let $G(V, E)$ be the cycle on n vertices with $V = \{v_i : i \in \mathbb{Z}_n\}$ and $E = \{(v_i, v_{i+1}) : i \in \mathbb{Z}_n\}$, where addition is modulo n . The game plays on the undirected cycle, but it will be convenient to fix an arbitrary orientation. Let $u_0, \dots, u_{\ell-1}$ be the distinct facilities chosen by k players in a strategy profile f with $\ell \leq k$, numbered according to the orientation of the cycle. For every $j \in \mathbb{Z}_\ell$, let $c_j \geq 1$

be the number of players who choose the facility u_j and let $d_j \geq 1$ be the length of the directed path from u_j to u_{j+1} following the orientation of G . Now the strategy profile is defined by these 2ℓ numbers, up to permutation of the players. We decompose the distance into $d_j = 1 + 2a_j + b_j$, for $0 \leq b_j \leq 1$, where $2a_j + b_j$ is the number of vertices between facilities u_j and u_{j+1} . So if $b_j = 1$, then there is a vertex in midway at equal distance from u_j and u_{j+1} .

With these notations the payoff of player i located on facility u_j is

$$p_i := \frac{b_{j-1}}{c_{j-1} + c_j} + \frac{a_{j-1} + 1 + a_j}{c_j} + \frac{b_j}{c_j + c_{j+1}}.$$

All Nash equilibria of the game on the cycle graph are explicitly characterized. The intuition is that the cycle is divided by the players into segments of different length, which roughly differ at most by a factor 2. The exact statement is more subtle because several players can be located at a same facility and the payoff is computed differently depending on the parity of the distances between facilities.

Proposition 1 ([8]) *For a given strategy profile, let γ be the minimal payoff among all players, i.e., $\gamma := \min\{p_i | 1 \leq i \leq k\}$. Then this strategy profile is a Nash equilibrium if and only if, for all $j \in \mathbb{Z}_\ell$:*

P1. $c_j \leq 2$.

P2. $d_j \leq 2\gamma$.

P3. If $c_i \neq c_{i+1}$ then $\lfloor 2\gamma \rfloor$ is odd.

P3. If $c_j = 1$ and $d_{j-1} = d_j = 2\gamma$ then 2γ is odd.

P4. If $c_j = c_{j+1} = 1$ and $d_{i-1} + d_i = d_{i+1} = 2\gamma$ then 2γ is odd.

If $c_j = c_{j-1} = 1$ and $d_{i-1} = d_i + d_{i+1} = 2\gamma$ then 2γ is odd.

A method to find Nash equilibrium in some games is to apply the *best-response dynamic* from an initial strategy profile. However, in our game, even in the simple cycle graph in which all Nash equilibria can be exactly characterized, in general there is no hope to use the best-response dynamic to find Nash equilibria in the game.

Proposition 2 *On the cycle graph, the best response dynamic does not converge.*

Proof: Figure 3 shows an example of a graph, where the best response dynamic can iterate forever. □

Nevertheless there is a slightly different Voronoi game in which the best response dynamic converges : The *Voronoi game with disjoint facilities* is identical with the previous game, except that players who are located on the same facility now gain zero.

Proposition 3 *On the cycle graph, for the Voronoi game with disjoint facilities, the best response dynamic does converge to pure Nash equilibria.*

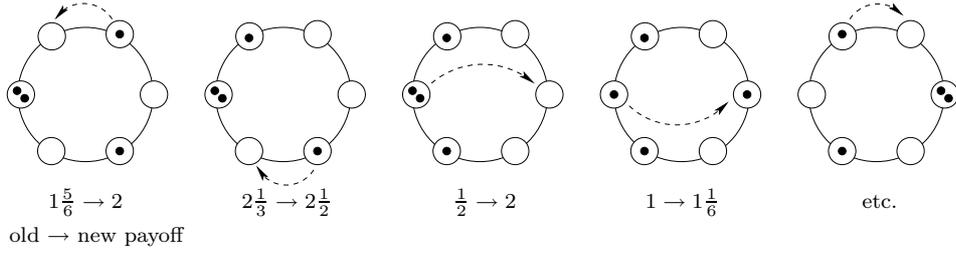


Figure 3: The best response dynamic does not converge on this graph.

Proof: To show convergence we use a potential function. For this purpose we define the *dominance order*: Let A, B be two multisets. If $|A| < |B|$ then $A \succ B$. If $|A| = |B| \geq 1$, and $\max A > \max B$ then $A \succ B$. If $|A| = |B| \geq 1$, $\max A = \max B$ and $A \setminus \{\max A\} \succ B \setminus \{\max B\}$ then $A \succ B$. This is a total order.

The potential function is the multiset $\{d_0, d_1, \dots, d_{k-1}\}$, that is all distances between successive occupied facilities. Player i 's payoff — renumbered conveniently — is simply $(d_{i-1} + d_i)/2$ since there is at most one player located on a vertex. Now consider a best response for player i moving to a vertex not yet chosen by another player, say between player j and $j+1$. Therefore in the multiset $\{d_0, d_1, \dots, d_{k-1}\}$, the values d_{i-1}, d_i, d_j are replaced by $d_{i-1} + d_i, d', d''$ for some values $d', d'' \geq 1$ such that $d_j = d' + d''$. As player i increases its payoff, $d_j > d_{i-1} + d_i$. Hence, the new potential value is dominated by the previous one. This proves that after a finite number of iterations, the best response dynamic converges to a Nash equilibrium. \square

4 Existence of a Nash equilibrium is \mathcal{NP} -hard

In this section we show that it is \mathcal{NP} -hard to decide whether for a given graph $G(V, E)$ there is a Nash equilibrium for k players. For this purpose we define a more general but equivalent game, which simplifies the reduction.

In the *generalized Voronoi game* $\langle G(V, E), U, w, k \rangle$ we are given a graph G , a set of facilities $U \subseteq V$, a positive weight function w on vertices and a number of players k . Here the set of strategies of each player is only U instead of V . Also the payoff of a player is the weighted sum of fractions of customers assigned to it, i.e., the payoff of player i is $p_i := \sum_{v \in V} w(v) F_{i,v}$.

Lemma 2 *For every generalized Voronoi game $\langle G(V, E), U, w, k \rangle$ there is a standard Voronoi game $\langle G'(V', E'), k \rangle$ with $V \subseteq V'$, which has the same set of Nash equilibria and which is such that $|V'|$ is polynomial in $\sum_{v \in V} w(v)$.*

Proof: To construct G' we will augment G in two steps. Start with $V' = V$.

First, for every vertex $u \in V$ such that $w(u) > 1$, let H_u be a set of $w(u) - 1$ new vertices. Set $V' = V' \cup H_u$ and connect u with every vertex from H_u .

Second, let H be a set of $k(a + 1)$ new vertices where $a = |V'| = \sum_{v \in V} w(v)$. Set $V' = V' \cup H$ and connect every vertex of U with every vertex of H .

Now in $G'(V', E')$ every player's payoff can be decomposed in the part gained from $V' \setminus H$ and the part gained from H . We claim that in a Nash equilibrium every player chooses a vertex from U . If there is at least one player located in U , then the gain from H for any other player is 0 if

located in $V' \setminus (U \cup H)$, is 1 if located in H and is at least $a + 1$ if located in U . Since the total payoff from $V' \setminus H$ over all players is a , this forces all players to be located in U .

Clearly by construction, for any strategy profile $f \in U^k$, the payoffs are the same for the generalized Voronoi game in G as for the standard Voronoi game in G' . Therefore we have equivalence of the set of Nash equilibria in both games. \square

Our \mathcal{NP} -hardness proof will need the following gadget.

Lemma 3 *For the graph G shown in figure 4 and $k = 2$ players, there is no Nash equilibrium.*

Proof: We will simply show that given an arbitrary location of one player, the other player can move to a location where he gains at least 5. Since the total payoff over both players is 9, this will prove that there is no Nash equilibrium.

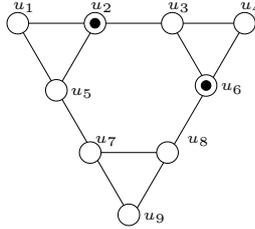


Figure 4: Example of a graph with no Nash equilibrium for 2 players.

By symmetry without loss of generality the first player is located at the vertices u_1 or u_2 . Now if the second player is located at u_6 , his payoff is at least 5. \square

Theorem 1 *Given a graph $G(V, E)$ and a set of k players, deciding the existence of Nash equilibrium for k players on G is \mathcal{NP} -complete for arbitrary k , and polynomial for constant k .*

Proof: The problem is clearly in \mathcal{NP} , since best responses can be computed in polynomial time, therefore it can be verified efficiently if a strategy profile is a Nash equilibrium. Since there are n^k different strategy profiles, for $n = |V|$, the problem is polynomial when k is constant.

For the proof of \mathcal{NP} -hardness, we will reduce 3-PARTITION — which is unary \mathcal{NP} -complete [7] — to the generalized Voronoi game, which by Lemma 2 is itself reduced to the original Voronoi game. In the 3-PARTITION problem we are given integers a_1, \dots, a_{3m} and B such that $B/4 < a_i < B/2$ for every $1 \leq i \leq 3m$, $\sum_{i=1}^{3m} a_i = mB$ and have to partition them into disjoint sets $P_1, \dots, P_m \subseteq \{1, \dots, 3m\}$ such that for every $1 \leq j \leq m$ we have $\sum_{i \in P_j} a_i = B$.

We construct a weighted graph $G(V, E)$ with the weight function $w : V \rightarrow \mathbb{N}$ and a set $U \subseteq V$ such that for $k = m + 1$ players ($m \geq 2$) there is a Nash equilibrium to the generalized Voronoi game $\langle G, U, w, k \rangle$ if and only if there is a solution to the 3-PARTITION instance. We define the constants $c = \binom{3m}{3} + 1$ and $d = \left\lfloor \frac{Bc - c + c/m}{5} \right\rfloor + 1$. The graph G consists of 3 parts. In the first part V_1 , there is for every $1 \leq i \leq 3m$ a vertex v_i of weight $a_i c$. There is also an additional vertex v_0 of weight 1. In the second part V_2 , there is for every triplet (i, j, k) with $1 \leq i < j < k \leq 3m$ a vertex u_{ijk} of unit weight¹. Every vertex u_{ijk} is connected to v_0, v_i, v_j and v_k . The third part V_3 , consists

¹Ideally we would like to give it weight zero, but there seems to be no simple generalization of the game which allows zero weights, while preserving the set of Nash equilibria.

of the 9 vertex graph of Figure 4 where each of the vertices u_1, \dots, u_9 has weight d . To complete our construction, we define the facility set $U := V_2 \cup V_3$. Note that although the graph for the generalized Voronoi game is disconnected, the reduction of Lemma 2 to the original Voronoi game will connect V_2 with V_3 .

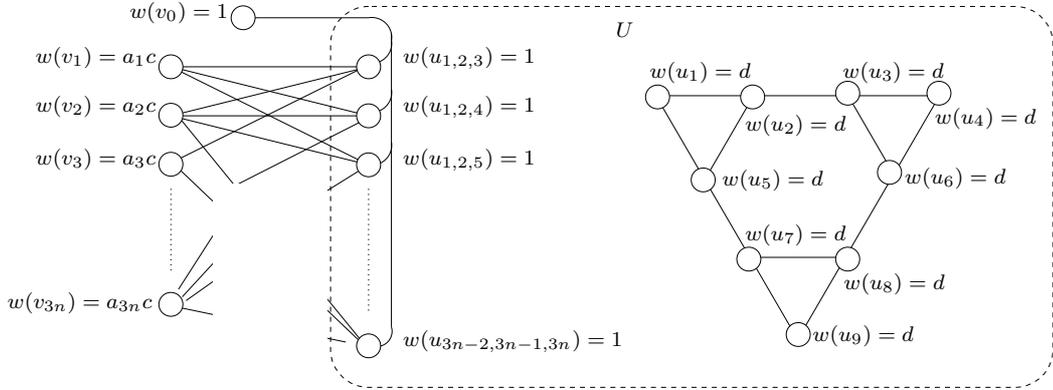


Figure 5: Reduction from 3-PARTITION.

First we show that if there is a solution P_1, \dots, P_m to the 3-PARTITION instance then there is a Nash equilibrium for this graph. Simply for every $1 \leq q \leq m$ if $P_q = \{i, j, k\}$ then player q is assigned to the vertex u_{ijk} . Player $m+1$ is assigned to u_2 . Now player $(m+1)$'s payoff is $9d$, and the payoff of each other player q is $Bc + (c - m - 1)/m + 1/m + 1 = Bc + c/m$. To show that this is a Nash equilibrium we need to show that no player can increase his payoff. There are different cases. If player $m+1$ moves to a vertex u_{ijk} , his payoff will be at most $\frac{3}{4}Bc + c/(m+1) < 9d$, no matter if that vertex was already chosen by another player or not. If player $1 \leq q \leq m$ moves from vertex u_{ijk} to a vertex u_ℓ then his gain can be at most $5d < Bc + c/m$. But what can be his gain, if he moves to another vertex $u_{i'j'k'}$? In case where $i = i', j = j', k \neq k'$, $a_i c + a_j c$ is smaller than $\frac{3}{4}Bc$ because $a_i + a_j + a_k = B$ and $a_k > B/4$. Since $a_{k'} < B/2$, and player q gains only half of it, his payoff is at most $a_i c + a_j c + a_{k'} c/2 + c/m < Bc + c/m$ so he again cannot improve his payoff. The other cases are similar.

Now we show that if there is a Nash equilibrium, then it corresponds to a solution of the 3-PARTITION instance. Consider a Nash equilibrium. First we claim that in this equilibrium, there is exactly one player in V_3 . Clearly if there are 2 players, this contradicts equilibrium by Lemma 3. If there are 3 players or more, then by a counting argument there are vertices v_i, v_j, v_k which are at distance more than one from any player. One of the players located at V_3 gains at most $3d$ and if he moves to u_{ijk} , his payoff would be at least $\frac{3}{4}Bc + c/m > 3d$. Now if there is no player in V_3 , then any player moving to u_2 will gain $9d > \frac{3}{2}Bc + c/m$ which is an upper bound for the payoff of players located in V_2 . So we know that there is a single player in V_3 and the m players in V_2 must form a partition, since otherwise there is a vertex $v_i \in V_1$ at distance at least 2 to any player. Therefore, by the previous argument, there would be a player in V_2 who can increase his payoff by moving to the other vertex in V_2 as well. (He moves in such a way that his new facility is at distance 1 to v_i .) Moreover, in this partition, each player gains exactly $Bc + c/m$ because if one gains less, given all weights in V_1 are multiple of c , he gains at most $Bc - c + c/m$ and he can always augment his payoff by moving to V_3 ($5d > Bc - c + c/m$). \square

5 Social cost discrepancy

In this section, we study how much the social cost of Nash equilibria can differ for a given graph, assuming Nash equilibria exist. Recall that the *social cost* of a strategy profile f is $\text{cost}(f) := \sum_{v \in V} d(v, f)$. Since we assumed $k < n$ the cost is always non-zero. For unconnected graphs, the social cost can be infinite, and so can be the social cost discrepancy. Therefore in this section we consider only connected graphs.

Lemma 4 *There are connected graphs for which the social cost discrepancy is $\Omega(\sqrt{n/k})$, where n is the number of vertices and k the number of players.*

Proof: We construct a graph G as shown in figure 6. The total number of vertices in the graph is $n = k(2a + b + 2)$. We distinguish two strategy profiles f and f' : the vertices occupied by f are marked with round dots, and the vertices of f' are marked with triangular dots.

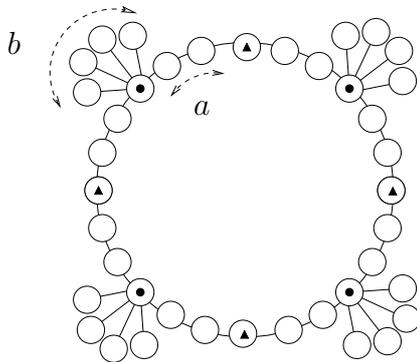


Figure 6: Example of a graph with high social cost discrepancy.

By straightforward verification, it can be checked that both f and f' are Nash equilibria. However the social cost of f is $\Theta(kb + ka^2)$ while the social cost of f' is $\Theta(kab + ka^2)$. The ratio between both costs is $\Theta(a) = \Theta(\sqrt{n/k})$ when $b = a^2$ and thus the cost discrepancy is lower bounded by this quantity. \square

In continuous setting, Delaunay triangulation for a set P of points in the plane is a triangulation such that no point in P is inside the circumcircle of any triangle in the triangulation. Moreover, Delaunay triangulations maximize the minimum angle of all the angles of the triangles in the triangulation; they tend to avoid skinny triangles. The Delaunay triangulation, considered as the dual of Voronoi diagram together with these nice properties, is an important object as well as a tool in mathematics and computational geometry. Here, we give the definition of *Delaunay triangulation* on graph which is a generalization of Delaunay triangulation in continuous setting (see Figure 7 for an illustration). Interestingly, this notion is useful in analyzing the social cost discrepancy of Voronoi game.

Definition 1 *Given a strategy profile f , the Delaunay triangulation corresponding to f is a graph H_f on the k players in profile f . There is an edge (i, j) in H_f either if there is a vertex v in G with $F_{i,v} > 0$ and $F_{j,v} > 0$ or if there is an edge (v, v') in G with $F_{i,v} > 0$ and $F_{j,v'} > 0$.*

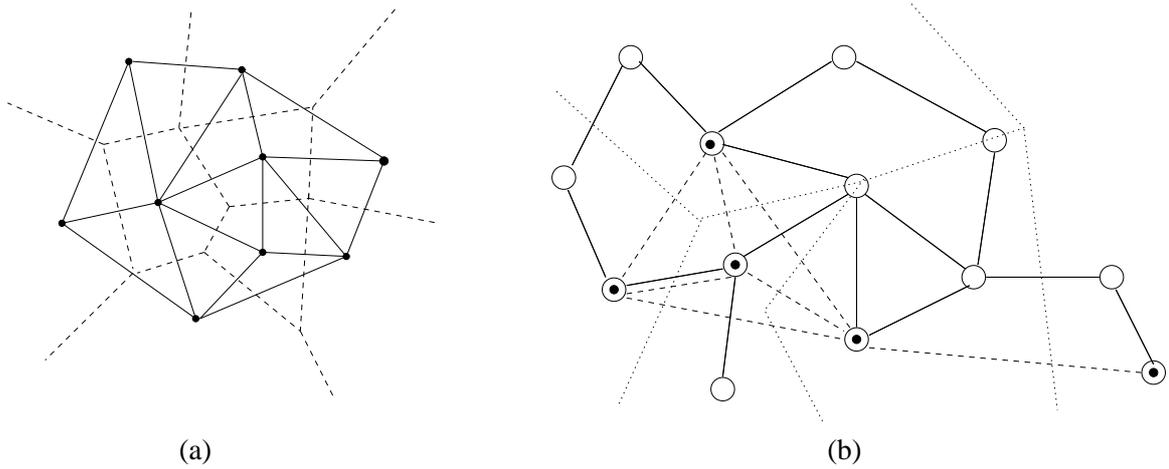


Figure 7: The Voronoi diagram and Delaunay triangulation in continuous and discrete settings. In continuous setting (a), the Voronoi diagram and the Delaunay triangulation are characterized by dashed and continuous lines, respectively. In discrete setting (b), the edges of graph are drawn in continuous and players are dots. The Voronoi diagram and the Delaunay triangulation in this case are characterized by dotted and dashed lines, respectively.

We will need to partition the Delaunay triangulation into small sets, called stars. For a given graph $G(V, E)$ a *star* is vertex set $A \subseteq V$ such that $|A| \geq 2$, and there is a *center* vertex $v_0 \in A$ such that for every $v \in A, v \neq v_0$ we have $(v_0, v) \in E$. Note that our definition allows the existence of additional edges between vertices from A .

Lemma 5 *For any connected graph $G(V, E)$, V can be partitioned into stars.*

Proof: We define an algorithm to partition V into stars.

As long as the graph contains edges, we do the following. We start choosing an edge: If there is a vertex u with a unique neighbor v , then we choose the edge (u, v) ; otherwise we choose an arbitrary edge (u, v) . Consider the vertex set consisting of u, v as well as of any vertex w that would be isolated when removing edge (u, v) . Add this set to the partition, remove it as well as adjacent edges from G and continue.

Clearly the set produced in every iteration is a star. Also when removing this set from G , the resulting graph does not contain an isolated vertex. This property is an invariant of this algorithm, and proves that it ends with a partition of G into stars. \square

Note that, when a graph is partitioned into stars, the centers of these stars form a dominating set of this graph. Nevertheless, vertices in a dominating set are not necessarily centers of any star-partition of a given graph.

The following lemma states that given two different Nash equilibria f and f' , every player in f is not too far from some player in f' . For this purpose we partition the Delaunay triangulation H_f into stars, and bound the distance from any player of a star to f' by some value depending on the star.

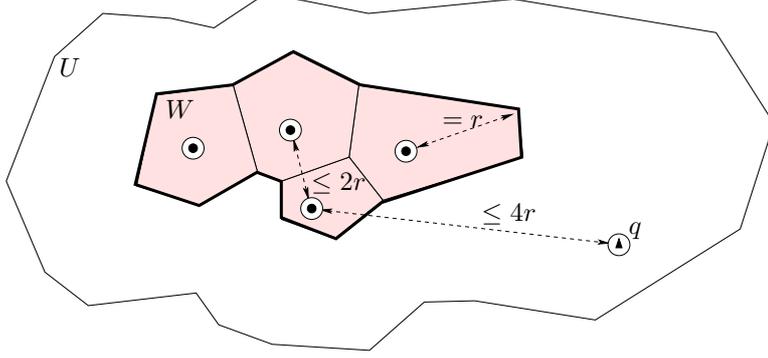


Figure 8: Illustration of Lemma 6.

Lemma 6 *Let f be an equilibrium and A be a star of a star partition of the Delaunay triangulation H_f . Let r be the maximal radius of the Voronoi cells over all players $i \in A$. Then, for any equilibrium f' , there exists a player j such that $d(f_i, f'_j) \leq 6r$ for every $i \in A$.*

Proof: Let $U = \{v \in V : \min_{i \in A} d(v, f_i) \leq 4r\}$. If we can show that there is a facility $f'_j \in U$ we would be done, since by definition of U there would be a player $i \in A$ such that $d(f_i, f'_j) \leq 4r$ and the distance between any pair of facilities of A is at most $2r$. This would conclude the lemma.

So for a proof by contradiction, assume that in the strategy profile f' there is no player located in U . Now consider the player with smallest payoff in f' . His payoff is not more than n/k . However if this player would choose as a facility the center of the star A , then he would gain strictly more: By the choice of r , any vertex in W is at distance at most $3r$ to the center of the star. However, by assumption and definition of U , any other facility of f' would be at distance strictly greater than $3r$ to any vertex in W . So the player would gain at least all vertices around it at distance at most $3r$, which includes W . Since any player's payoff is strictly more than $n/2k$ by Lemma 1, and since a star contains at least two facilities by definition, this player would gain strictly more than n/k , contradicting that f' is an equilibrium. This concludes the proof. \square

Theorem 2 *For any connected graph $G(V, E)$ and any number of players k the social cost discrepancy is $O(\sqrt{kn})$, where $n = |V|$.*

Proof: Let f, f' be arbitrary equilibria on $G(V, E)$. We will consider a generalized partition of V and for each part bound the cost of f' by $c\sqrt{kn}$ times the cost of f for some constant c .

For a non-negative n -dimensional vector W we define the cost restricted to W as $\text{cost}_W(f) = \sum_{v \in V} W_v \cdot d(v, f)$. Now the cost of f would write as the sum of $\text{cost}_W(f)$ over the vectors W from some fixed generalized partition.

Fix a star partition of the Delaunay triangulation H_f . Let A be an arbitrary star from this partition, $a = |A|$, and let W be the sum of the corresponding Voronoi cells, i.e., $W = \sum_{i \in A} F_i$. We will show that $\text{cost}_W(f') = O(\sqrt{kn} \cdot \text{cost}_W(f))$, which would conclude the proof. There will be two cases $k \leq n/4$ and $k > n/4$.

By the previous lemma there is a vertex f'_j such that $d(f_i, f'_j) \leq 6r$ for all $i \in A$, where r is the largest radius of all Voronoi cells corresponding to the star A . So the cost of f' restricted to the

vector W is

$$\begin{aligned}
\text{cost}_W(f') &= \sum_{v \in V} W_v \cdot d(v, f') \leq \sum_{v \in V} W_v \cdot d(v, f'_j) \\
&= \sum_{v \in V} \sum_{i \in A} F_{i,v} \cdot d(v, f'_j) \\
&\leq \sum_{v \in V} \sum_{i \in A} F_{i,v} \cdot (d(v, f_i) + d(f_i, f'_j)) \\
&\leq \text{cost}_W(f) + 6r \cdot |W|,
\end{aligned} \tag{1}$$

where $|W| := \sum_{v \in V} W_v$.

Moreover by definition of the radius, there is a vertex v with $W_v > 0$ such that the shortest path to the closest facility in A has length r . So the cost of f restricted to W is bigger than the cost restricted to this shortest path:

$$\text{cost}_W(f) \geq \left(\frac{1}{k} \cdot 1 + \frac{1}{k} \cdot 2 + \dots + \frac{1}{k} \cdot r\right) \geq \frac{1}{k} \cdot r(r-1)/2.$$

(The fraction $\frac{1}{k}$ appears because a vertex can be assigned to at most k players.)

First we consider the case $k \leq n/4$. We have

$$\text{cost}_W(f) \geq |W| - |A| \geq a(n/2k - 1) \geq an/4k.$$

The first inequality is because the distance of all customers which are not facilities to a facility is at least one. The second inequality is due to Lemma 1 and $|W|$ is the sum of payoffs of all players in A .

Note that $|W| \leq n$ and $2 \leq a \leq k$. Now if $r \leq \sqrt{an}$, then

$$\frac{\text{cost}_W(f')}{\text{cost}_W(f)} \leq 1 + \frac{6r|W|}{\text{cost}_W(f)} \leq 1 + \frac{6r \cdot a \cdot 2n/k}{an/4k} = O(r) = O(\sqrt{kn}).$$

And if $r \geq \sqrt{an}$, then

$$\frac{\text{cost}_W(f')}{\text{cost}_W(f)} \leq 1 + \frac{6r|W|}{\text{cost}_W(f)} \leq 1 + \frac{6r \cdot a \cdot 2n/k}{r(r-1)/2k} = O(an/r) = O(\sqrt{kn}).$$

Now we consider case $n > k > n/4$. In any equilibrium, the maximum payoff is at most $2n/k$. Moreover the radius r of any Voronoi cell is upper bounded by $n/k + 1$, otherwise the player with minimum gain (which is at most n/k) could increase his gain by moving to a vertex which is at distance at least r from every other facility. Therefore $r = O(1)$. Summing (1) over all stars with associated partition W , we obtain $\text{cost}(f') \leq \text{cost}(f) + cn$, for some constant c . Remark that the social cost of any equilibrium is at least $n - k$. Hence, $\frac{\text{cost}(f')}{\text{cost}(f)} = O(n) = O(\sqrt{kn})$. \square

6 Conclusion

It would be interesting to close the gap between the lower and upper bounds for the social cost discrepancy. Furthermore, it is expected that the social cost discrepancy would be considered in other games in order to better understand Nash equilibria in these games.

References

- [1] Hee-Kap Ahn, Siu-Wing Cheng, Otfried Cheong, Mordecai Golin, and Renée van Oostrum. Competitive facility location: the Voronoi game. *Theoretical Computer Science*, 310:457–467, 2004.
- [2] Vittorio Bilò, Michele Flammini, Gianpiero Monaco, and Luca Moscardelli. On the performances of nash equilibria in isolation games. In *Proc. 15th Int. Conf. on Computing and Combinatorics*, pages 17–26, 2009.
- [3] O. Cheong, S. Har-Peled, N. Linial, and J. Matousek. The one-round Voronoi game. *Discrete and Computational Geometry*, 31(1):125–138, 2004.
- [4] Christoph Dürr and Nguyen Kim Thang. Nash Equilibria in Voronoi Games on Graphs. In *Proceedings of the 15th European Symposium on Algorithms (ESA)*, pages 17–28, 2007.
- [5] Sándor P. Fekete and Henk Meijer. The one-round Voronoi game replayed. *Computational Geometry: Theory and Applications*, 30:81–94, 2005.
- [6] Rainer Feldmann, Marios Mavronicolas, and Burkhard Monien. Nash equilibria for voronoi games on transitive graphs. In *Proc. 5th Int. Workshop on Internet and Network Economics (WINE)*, pages 280–291, 2009.
- [7] Michael R. Garey and David S. Johnson. *Computers and Intractability; A Guide to the Theory of NP-Completeness*. W. H. Freeman & Co., New York, NY, USA, 1990.
- [8] Marios Mavronicolas, Burkhard Monien, Vicky G. Papadopoulou, and Florian Schoppmann. Voronoi games on cycle graphs. In *Proceedings of the 33rd International Symposium on Mathematical Foundations of Computer Science 2008 (MFCS)*, pages 503–514, 2008.
- [9] Vijay V. Vazirani. *Approximation Algorithms*. Springer, 2004.
- [10] A. Vetta. Nash equilibria in competitive societies with applications to facility location. In *Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 416–425, 2002.
- [11] Yingchao Zhao, Wei Chen, and Shang-Hua Teng. The isolation game: A game of distances. In *Proceedings of the 19th International Symposium on Algorithms and Computation (ISAAC)*, pages 148–158, 2008.