

Strategic Facility Location with Clients That Minimize Total Waiting Time

Simon Krogmann¹, Pascal Lenzner¹, Alexander Skopalik²

¹ Hasso Plattner Institute, University of Potsdam

² Department of Applied Mathematics, University of Twente
simon.krogmann@hpi.de, pascal.lenzner@hpi.de, a.skopalik@utwente.nl

Abstract

We study a non-cooperative two-sided facility location game in which facilities and clients behave strategically. This is in contrast to many other facility location games in which clients simply visit their closest facility. Facility agents select a location on a graph to open a facility to attract as much purchasing power as possible, while client agents choose which facilities to patronize by strategically distributing their purchasing power in order to minimize their total waiting time. Here, the waiting time of a facility depends on its received total purchasing power. We show that our client stage is an atomic splittable congestion game, which implies existence, uniqueness and efficient computation of a client equilibrium. Therefore, facility agents can efficiently predict client behavior and make strategic decisions accordingly. Despite that, we prove that subgame perfect equilibria do not exist in all instances of this game and that their existence is NP-hard to decide. On the positive side, we provide a simple and efficient algorithm to compute 3-approximate subgame perfect equilibria.

Introduction

Facility location problems are a classical and popular object of study in AI, Operations Research, Economics, and Theoretical Computer Science. These problems encompass natural problems like placing facilities in a socially optimal way, e.g., the placement of hospitals, fire stations, or schools, or determining the locations of competing facilities, e.g., bars, shops, or gas stations. While the former type of problems are typically modeled as optimization problems and solved with a rich toolbox of combinatorial graph algorithms, this cannot be done for the latter problems, due to the strategic setting. Instead, models and methods from (Algorithmic) Game Theory are necessary to cope with facility location in competitive environments. The focus of this paper will be one such model that explicitly considers strategic behavior by the facilities as well as by the clients that want to patronize these facilities. Most importantly, we thereby consider a very natural type of client behavior that has, to the best of our knowledge, not yet been studied.

The study of strategic facility location models dates back almost a century to the works of Hotelling (1929) who considered how to place two competing shops in a linear mar-

ket, called the “main street”. Later, this model was refined by Downs (1957) and used for the placement of political candidates in a political left-right spectrum. In the Hotelling-Downs model, competing but otherwise identical facilities strategically select a placement in some underlying space, e.g., on the line, to attract as many clients as possible. The clients are modeled in a very simple way: clients are assumed to patronize their nearest opened facility. However, while such simplistic clients are still considered in a wide range of models and recent works, such basic client agents do not seem to convincingly capture the behavior of real-world clients. Abstracting away from more complex aspects like pricing, product quality, or social influences, realistic clients would not only consider the distance to a facility but also the incurred waiting time at the facility. Interestingly, this waiting time depends on the number of other clients that also patronize the same facility. Thus, in a more realistic setting we have that not only the facilities act strategically, but also the clients base their behavior on the behavior of the other clients instead of myopically minimizing distances. Such more realistic clients have first been considered in the work of Kohlberg (1983), who studied clients that minimize a convex combination of distance and waiting time. However, this enhanced realism comes at a cost: while for the original Hotelling-Downs model equilibrium states always exist, except for exactly 3 facilities, equilibria only exist for extreme cases of Kohlberg’s model.

But, as demonstrated by Feldman, Fiat, and Obratzsova (2016), these negative results can be circumvented by considering a natural and interesting variant of the Hotelling-Downs model. In this model, instead of minimizing the distance, every client has a threshold of how far she would travel to a facility and chooses to patronize a facility uniformly at random whose distance is below the threshold. So essentially, the clients uniformly split their purchasing power among all facilities that are close enough. Besides being more realistic than clients committing to a single facility, this model variant always has equilibrium facility placements. However, it also has a downside: it considers simplistic clients that do not take waiting times into consideration. Very recently, this downside was removed in a model proposed by Krogmann et al. (2021). In their two-stage facility location game, a given graph encodes which locations can reach other locations. Clients with a certain purchas-

ing power occupy the nodes of this graph and facility agents strategically select a node for opening their facilities. Given a facility placement, the clients then split their purchasing power among the facilities they can reach to minimize their maximum waiting time. Analogously to classical make-span scheduling, these clients effectively try to distribute their purchasing power to load-balance the reachable facilities.

In this paper, we set out to explore a similar setting as Krogmann et al. (2021) but with a drastically different client behavior. To stay in the analogy with scheduling, we consider sum-of-completion-time scheduling instead of make-span scheduling. I.e., we study facility location problems with clients that distribute their purchasing power to minimize their total waiting time.

Related Work: There is an abundance of models for strategic facility location and we refer to the surveys by Eiselt, Laporte, and Thisse (1993) and ReVelle and Eiselt (2005) for an overview of the classical models. To the best of our knowledge, our model has not yet been studied.

The model by Feldman, Fiat, and Obratzsova (2016) with distance-minimizing clients on a line was generalized by Shen and Wang (2017) to cope with different continuous client distributions and by Cohen and Peleg (2019) to random distance thresholds. Moreover, Fournier, Van der Straeten, and Weibull (2020) study clients that choose the nearest facility if they have multiple options. Also related are Voronoi games (Ahn et al. 2004; Dürr and Thang 2007), which can be understood as generalized Hotelling-Downs models with distance-minimizing clients. For the version on networks by Dürr and Thang (2007), the authors show that equilibria may not exist and that existence is NP-hard to decide. Also a variant on a cycle (Mavronicolas et al. 2008) and in k -dimensional space (de Berg, Kisfaludi-Bak, and Mehr 2019; Ahn et al. 2004; Boppana et al. 2016) was studied.

For Kohlberg’s model, Peters, Schröder, and Vermeulen (2018) prove the existence of equilibria for certain trade-offs of distance and waiting time for small even numbers of facilities and they conjecture that equilibria exist for all cases with an even number of facilities for client utility functions that are heavily tilted towards minimizing waiting times. Feldotto et al. (2019) showed that computing equilibria for Kohlberg’s model can be done by solving a complex system of equations and they investigated the existence of approximate equilibria. They find that under a technical assumption, 1.08-approximate equilibria always exist.

A concept related to our model are utility systems, as introduced by Vetta (2002). There, agents gain utility by selecting a set of acts, which they choose from a collection of subsets of a groundset. Utility is assigned by a function that takes the selected acts of all agents as input. Also covering games (Gairing 2009) are related since they correspond to a one-sided version of our model, where clients distribute their purchasing power uniformly among all facilities in reach. In both settings, utility systems and covering games, pure NE exist and the Price of Anarchy is upper bounded by 2. More general versions are investigated by Goemans et al. (2006) and Brethouwer et al. (2018) in the form of market sharing games. In these models, k agents choose to serve a subset

of n markets. Each market then equally distributes its utility among all agents serving it. Schmand, Schröder, and Skopalik (2019) introduced a model which considers an inherent load balancing problem, however, each facility agent can create and choose multiple facilities and each client agent chooses multiple facilities. An empirical investigation of a two-sided facility location problem was conducted by Schön and Saini (2018), in which a single facility agent opens facilities for strategic clients. The facility agent sets service levels while the client agents determine their strategies based on these levels, but also travel distance and congestion.

Closest to our work is the above-mentioned work by Krogmann et al. (2021), where clients that minimize their maximum waiting time are considered. For this version equilibria always exist due to a potential function argument. Moreover, the authors present a polynomial time algorithm for computing the unique client equilibrium, given a facility placement. In terms of quality of the equilibria, the authors prove an essentially tight bound of 2 on the Price of Anarchy by establishing that this model is a valid utility system. These results on quality apply to a general class of games which also includes the game we study in this paper.

Facility location problems were also studied recently with Mechanism Design, e.g., Procaccia and Tennenholtz (2013); Feldman, Fiat, and Golomb (2016); Aziz et al. (2020); Chan et al. (2021); Harrenstein et al. (2021); Fotakis and Patsilidakos (2021); Deligkas, Eiben, and Goldsmith (2022).

Our Contribution: In this paper, we introduce a two-sided facility location game with clients that minimize their total incurred waiting time. In the first stage, each facility agent individually selects a location on a graph to open a facility, while in the second stage client agents choose which facilities to patronize. Thereby, clients may freely distribute their purchasing power among all opened facilities within their shopping range. In contrast to a similar game by Krogmann et al. (2021), our clients minimize their total expected waiting time instead of their maximum expected waiting time. To the best of our knowledge, we are the first to study this natural behavior in a facility location game.

The client behavior on its own is well studied in the form of *atomic splittable congestion games* and we reduce our client stage to variants of these games to show existence, uniqueness and polynomial time computation of client equilibria. This means that facilities can efficiently predict client behavior to inform their strategic location decisions.

However, on the negative side, we show that for our complete game consisting of both the facility stage and the client stage, subgame perfect equilibria (SPE) are not guaranteed to exist for all instances. This is in surprising contrast to the similar game studied by Krogmann et al. (2021), which is actually a potential game with guaranteed equilibrium existence. In fact, we prove that it is even NP-hard to decide the existence of an SPE for a given instance. This shows that the specific client behavior has a severe impact on the obtained game-theoretic properties.

But, on the positive side, we prove the existence of 3-approximate SPEs and show that they can be computed efficiently by using another facility location game as a proxy.

Model and Preliminaries

We consider a game-theoretic model for non-cooperative facility location, called the *Two-Sided Facility Location Game (2-FLG)*¹, where two types of agents, k facilities and n clients, strategically interact on a given vertex-weighted directed² host graph $H = (V, E, w)$, with $V = \{v_1, \dots, v_n\}$, where $w : V \rightarrow \mathbb{Q}_+$ denotes the vertex weight. Every vertex $v_i \in V$ corresponds to a client with weight $w(v_i)$, which can be understood as her purchasing power, and at the same time, each vertex is a possible location for setting up a facility for any of the k facility agents $\mathcal{F} = \{f_1, \dots, f_k\}$. Any client $v_i \in V$ considers patronizing a facility in her *shopping range* $N(v_i)$, i.e., her direct closed neighborhood $N(v_i) = \{v_i\} \cup \{z \mid (v_i, z) \in E\}$. Moreover, let $w(X) = \sum_{v_i \in X} w(v_i)$, for any $X \subseteq V$, denote the total purchasing power of the client subset X .

In our setting, the strategic behaviors of the facility and the client agents influence each other. Facility agents select a location to attract as much client weight, i.e., purchasing power, as possible, whereas clients strategically decide how to distribute their purchasing power among the facilities in their respective shopping ranges. More precisely, each facility agent $f_j \in \mathcal{F}$ selects a single location vertex $s_j \in V$ for setting up her facility, i.e., the strategy space of any facility agent $f_j \in \mathcal{F}$ is V . Let $\mathbf{s} = (s_1, \dots, s_k)$ denote the *facility placement profile*. And let $\mathcal{S} = V^k$ denote the set of all possible facility placement profiles. We will sometimes use the notation $\mathbf{s} = (s_j, s_{-j})$, where s_{-j} is the vector of strategies of all facilities agents except f_j . Given profile \mathbf{s} , we define the *attraction range* for a facility f_j on location $s_j \in V$ as $A_{\mathbf{s}}(f_j) = \{s_j\} \cup \{v_i \mid (v_i, s_j) \in E\}$. We extend this to sets of facilities $F \subseteq \mathcal{F}$ in the natural way, i.e., $A_{\mathbf{s}}(F) = \{s_j \mid f_j \in F\} \cup \{v_i \mid (v_i, s_j) \in E, f_j \in F\}$. Moreover, let $w_{\mathbf{s}}(\mathcal{F}) = \sum_{v_i \in A_{\mathbf{s}}(\mathcal{F})} w(v_i)$.

We assume that all facilities provide the same service for the same price and arbitrarily many facilities may be co-located at the same location. Each client $v_i \in V$ strategically decides how to distribute her spending capacity $w(v_i)$ among the opened facilities in her shopping range $N(v_i)$. For this, let $N_{\mathbf{s}}(v_i) = \{f_j \mid s_j \in N(v_i)\}$ denote the set of facilities in the shopping range of client v_i under \mathbf{s} .

Let $\sigma : \mathcal{S} \times V \rightarrow \mathbb{R}_+^k$ denote the *client weight distribution function*, where $\sigma(\mathbf{s}, v_i)$ is the weight distribution of client v_i and $\sigma(\mathbf{s}, v_i)_j$ is the weight distributed by v_i to facility f_j . We say that function σ is *feasible* for profile \mathbf{s} , if all clients having at least one facility within their shopping range distribute all their weight to the respective facilities and all other clients distribute nothing. Formally, function σ is feasible for profile \mathbf{s} , if for all $v_i \in V$ we have $\sum_{f_j \in N_{\mathbf{s}}(v_i)} \sigma(\mathbf{s}, v_i)_j = w(v_i)$, if $N_{\mathbf{s}}(v_i) \neq \emptyset$, and $\sigma(\mathbf{s}, v_i)_j = 0$, for all $1 \leq j \leq k$, if $N_{\mathbf{s}}(v_i) = \emptyset$. We use the notation $\sigma = (\sigma_i, \sigma_{-i})$ and (σ'_i, σ_{-i}) to denote the changed client weight distribution function that is identical to σ except for client v_i , which plays $\sigma'(\mathbf{s}, v_i)$ instead of $\sigma(\mathbf{s}, v_i)$.

¹We reuse notation by Krogmann et al. (2021).

²Our results also hold for undirected graphs. Also, our model can be equivalently defined on an undirected bipartite graph.




Figure 1: Two instances of the Min-2-FLG with their client equilibria visualized for a given facility placement profile. The clients on each node split their weight (above the nodes) among the facilities in their shopping ranges to minimize their cost. The facilities are marked by dots inside the nodes and receive the loads below in the respective unique client equilibria. The respective cost $L_1(\mathbf{s}, \sigma) = \sigma(\mathbf{s}, v_1)_1 \ell_1(\mathbf{s}, \sigma) + \sigma(\mathbf{s}, v_1)_2 \ell_2(\mathbf{s}, \sigma)$ of client v_1 (and also v_2 on the right) is given below each instance.

Any state (\mathbf{s}, σ) of the 2-FLG is determined by a facility placement profile \mathbf{s} and a feasible client weight distribution function σ . A state (\mathbf{s}, σ) then yields a *facility load* $\ell_j(\mathbf{s}, \sigma)$, with $\ell_j(\mathbf{s}, \sigma) = \sum_{i=1}^n \sigma(\mathbf{s}, v_i)_j$ for facility agent f_j . Hence, $\ell_j(\mathbf{s}, \sigma)$ naturally models the total congestion for the service offered by the facility of agent f_j induced by σ . A facility agent f_j strategically selects a location s_j to maximize her induced facility load $\ell_j(\mathbf{s}, \sigma)$. We assume that the service quality of facilities, e.g., the waiting time, deteriorates with increasing congestion. Hence for a client, the facility load corresponds to the waiting time at the respective facility.

There are many ways how clients could distribute their spending capacity. We investigate the *Min-2-FLG* with *waiting time minimizing clients*, i.e., a natural strategic behavior where client v_i strategically selects $\sigma(\mathbf{s}, v_i)$ to minimize her total waiting time. More precisely, the cost of client i is

$$L_i(\mathbf{s}, \sigma) = \sum_{j=1}^k \sigma(\mathbf{s}, v_i)_j \ell_j(\mathbf{s}, \sigma).$$

Another version of the 2-FLG we use in this paper is the *Uniform-2-FLG*, in which clients distribute their weight equally among all facilities in their range. Formally, for each pair of client v_i and facility f_j , with $f_j \in N_{\mathbf{s}}(v_i)$, the client's weight is $\sigma(\mathbf{s}, v_i)_j = \frac{w(v_i)}{N_{\mathbf{s}}(v_i)}$. Another such model is the load balancing 2-FLG introduced by Krogmann et al. (2021), which we do not consider in this paper.

We say that σ^* is a *client equilibrium weight distribution*, or simply a *client equilibrium*, if for all $v_i \in V$ we have that $L_i(\mathbf{s}, (\sigma_i^*, \sigma_{-i})) \leq L_i(\mathbf{s}, (\sigma'_i, \sigma_{-i}))$ for all feasible weight distributions $\sigma'(\mathbf{s}, v_i)$ of client v_i . See Figure 1 for an illustration of the client behavior in the Min-2-FLG.

We define the *stable states* of the 2-FLG as *subgame perfect equilibria (SPE)*, since we inherently have a two-stage game. First, the facility agents select locations for their facilities and then, given this facility placement, the clients strategically distribute their purchasing power among the facilities in their shopping range. A state (\mathbf{s}, σ) is in SPE, or *stable*, if

- (1) $\forall f_j \in \mathcal{F}, \forall s'_j \in V: \ell_j(\mathbf{s}, \sigma) \geq \ell_j((s'_j, s_{-j}), \sigma)$ and
- (2) $\forall \mathbf{s} \in \mathcal{S}, \forall v_i \in V: L_i(\mathbf{s}, \sigma) \leq L_i(\mathbf{s}, (\sigma'_i, \sigma_{-i}))$ for all feasible weight distributions $\sigma'(\mathbf{s}, v_i)$ of client v_i .

As we will show, SPE do not always exist. Hence, we relax the first condition as follows and obtain the notion of α -approximate subgame perfect equilibria (α -SPE):

$$(1') \quad \forall f_j \in \mathcal{F}, \forall s'_j \in V: \ell_j(\mathbf{s}, \sigma) \geq \alpha \ell_j((s'_j, s_{-j}), \sigma)$$

We study dynamic properties of the 2-FLG. Let an *improving move* by some (facility or client) agent be a strategy change that improves the agent's utility. A game has the *finite improvement property (FIP)* if all sequences of improving moves are finite. The FIP is equivalent to the existence of an *ordinal potential function* (Monderer and Shapley 1996), which implies equilibrium existence.

Client Equilibria

The existence of client equilibria is implied by Kakutani's fixed-point theorem (Kakutani 1941). A reduction of our game to *atomic splittable routing games* by Bhaskar et al. (2015) obtains uniqueness of client equilibria. Note that for the reduction to their model to work, we require the extension defined by the authors where each player may have her own individual delay function for each edge.

Theorem 1. *For a given facility placement profile \mathbf{s} the client equilibrium in the Min-2-FLG is unique.*

Client equilibria can be computed in polynomial time with an algorithm given by Harks and Timmermans (2021).

Theorem 2. *For a given facility placement profile \mathbf{s} the client equilibrium in the Min-2-FLG can be computed in polynomial time.*

Notably, the most involved step to compute a client equilibrium is to determine the pairs of clients and facilities with non-zero weight. This information then yields a set of linear equations solvable by Gaussian elimination.

We refer to Krogmann, Lenzner, and Skopalik (2022) for details of the application of Kakutani's fixed-point theorem and the reductions for Theorems 1 and 2.

We now show how the loads of two facilities with a shared client relate to each other in a client equilibrium.

Lemma 1. *Let f_p and f_q be two facilities that share a client v_i in a given facility placement profile \mathbf{s} . In a client equilibrium σ , if $\sigma(\mathbf{s}, v_i)_p > 0$, then*

$$\ell_p(\mathbf{s}, \sigma) + \sigma(\mathbf{s}, v_i)_p \leq \ell_q(\mathbf{s}, \sigma) + \sigma(\mathbf{s}, v_i)_q.$$

Proof. If v_i transfers weight from f_p to f_q it only affects the terms $\sigma(\mathbf{s}, v_i)_p \ell_p(\mathbf{s}, \sigma)$ and $\sigma(\mathbf{s}, v_i)_q \ell_q(\mathbf{s}, \sigma)$ in her cost and not the other terms of the sum. Let $L_{i,j} = \sigma(\mathbf{s}, v_i)_j \ell_j(\mathbf{s}, \sigma)$, for all $v_i \in V$ and $f_j \in \mathcal{F}$. Since σ is the unique client equilibrium, we know that for a transfer of weight of ϵ with $0 < \epsilon \leq \sigma(\mathbf{s}, v_i)_p$ we have $L_{i,p} + L_{i,q} < (\sigma(\mathbf{s}, v_i)_p - \epsilon)(\ell_p(\mathbf{s}, \sigma) - \epsilon) + (\sigma(\mathbf{s}, v_i)_q + \epsilon)(\ell_q(\mathbf{s}, \sigma) + \epsilon)$. This yields $\ell_p(\mathbf{s}, \sigma) + \sigma(\mathbf{s}, v_i)_p < 2\epsilon + \ell_q(\mathbf{s}, \sigma) + \sigma(\mathbf{s}, v_i)_q$. Since ϵ may be arbitrarily small, but not zero, this finishes the proof. \square

Lemma 1 also implies equality of the two terms, i.e., $\ell_p(\mathbf{s}, \sigma) + \sigma(\mathbf{s}, v_i)_p = \ell_q(\mathbf{s}, \sigma) + \sigma(\mathbf{s}, v_i)_q$, if a client v_i has non-zero weight on both f_p and f_q .




Figure 2: Best responses (dashed) to game states for the instance G^* of the Min-2-FLG without SPE. The utilities of the facilities before and after the move are given on the right.

Subgame Perfect Equilibria

Unlike other versions of the 2-FLG, the Min-2-FLG does not admit subgame perfect equilibria in all instances. In fact, the counterexample only needs two facility agents with the host graph H being a path of size 4. Therefore, it may be complicated to find non-trivial subclasses of the Min-2-FLG which always admit SPE.

Theorem 3. *There are instances of the Min-2-FLG for which a subgame perfect equilibrium does not exist.*

Proof. Let $G^* = (V^*, E^*)$ with $V^* = \{v_1, v_2, v_3, v_4\}$ and $E^* = \{(v_1, v_2), (v_2, v_3), (v_3, v_4)\}$. Let the client weights be $w(v_1) = 3$, $w(v_2) = 2$, $w(v_3) = 7$, and $w(v_4) = 1$. For two facility agents, this instance does not admit an SPE. In Figure 2 we show the best responses for states that cannot trivially be excluded as SPE. \square

Additionally, determining whether an instance admits an SPE is computationally intractable.

Theorem 4. *Deciding if an instance of the Min-2-FLG admits an SPE is NP-hard.*

Proof. We reduce from INDEPENDENTSET (IS): Given a graph $G = (V, E)$ and an integer $k \leq |V|$, decide whether there exists a subset $I \subseteq V$ with $|I| = k$ such that no two vertices in I share an edge. This problem is NP-hard, even for graphs with maximum degree 3 (Garey and Johnson 1990). In the following, we assume that G has a maximum degree of at most 3.

To prove the theorem, we construct an instance of the Min-2-FLG on a host graph $H = (V', E', w)$ with $k' = 2k$ facilities such that there is an SPE in H if and only if G contains an independent set of size k . We obtain H from G by replacing every edge $e = \{u, v\}$ by a new vertex x_e and two new edges (x_e, u) and (x_e, v) . For every vertex v of degree 1 (or 2) add two (or one) vertices y_v (and z_v) and the edge (y_v, v) (and (z_v, v)). The newly added x -, y - and z -vertices have weight 1.75, vertices that originally belonged to V have weight 0. Hence, we now have in H exactly $|V|$ many vertices that all have weight 0 with exactly 3 neighbors that have weight 1.75 each. All other vertices have weight 1.75 and one or two neighbors with weight 0. To complete the

construction of H , we add k copies of the graph G^* that we used in Theorem 3 (see Figure 2).

Now, if G contains an independent set I of size k , then there is an equilibrium in which k facilities are placed on the vertices of I and one facility each is placed on the second vertex from the right in each of the k copies of G^* . The first k facilities each have a payoff of 5.25 and, hence, play their best response. In particular, it is not an improvement to choose any vertex in any of the copies of G^* as this yields at most 4.75. Each of the k players in the copies of G^* has a payoff of 9 and, hence, is clearly playing the best response.

If there is no independent set of size k in G , then there cannot be more than $k - 1$ players with a payoff of more than 4.375 on vertices outside of the copies of G^* . Note that a higher payoff is only possible on nodes of degree 3 with no other facility within distance 2. However, in an equilibrium no facility would choose a location with a payoff of at most 4.375 as there is at least one of the copies of G^* with at most one other facility on it. Switching to the best vertex in that copy of G^* guarantees a payoff of at least 4.5. Finally, we observe that there is no equilibrium with two players on host graph G^* in Figure 2, hence there is no equilibrium in H . \square

Approximation of SPE

In this section, we show that an SPE in the Uniform-2-FLG is a 3-approximate SPE in the Min-2-FLG. We first give an example instance, where in an SPE in the Uniform-2-FLG a facility can improve by a factor of 2 when treating it as a state of the Min-2-FLG. This instance serves as a lower bound of the approximation quality using the Uniform-2-FLG and is also shown in Figure 3.

Example 1. Let $t > 0$ be a natural number. Let $G = (V, E)$, with $V = \{v_b, v_a\} \cup \bigcup_{i=1}^t V_i$, with $V_i = \{x_i\} \cup \{y_{i,1}, \dots, y_{i,t}\}$ and $E = \bigcup_{i=1}^t E_i$, with $E_i = \{(v_a, x_i)\} \cup \{(x_i, y_{i,1}), \dots, (x_i, y_{i,t})\}$. Let the client weights be $w(v_b) = 1$, $w(v_a) = 0$, $w(x_i) = 1$, for all i and $w(y_{i,j}) = \frac{2t-2}{t}$, for all i, j . Let the number of facilities be $n = t^2 + 1$.

For this example in a client equilibrium of the Uniform-2-FLG, facilities are located at v_b and all of the nodes $y_{i,j}$, with $1 \leq i, j \leq t$, while for the Min-2-FLG the facility agent on v_b can improve by a factor of 2 by moving to location v_a .

Theorem 5. Let \mathbf{s} be an SPE for an instance of the Uniform-2-FLG. When transferred to the Min-2-FLG, it is possible for profile $\mathbf{s}_{\text{Uniform}}$ to be a 2-approximate SPE.

Proof. For the Uniform-2-FLG, the facility placement profile $\mathbf{s} = (v_b, y_{1,1}, \dots, y_{1,t}, \dots, y_{t,1}, \dots, y_{t,t})$ is an SPE for Example 1, since the facility on v_b receives $\frac{2t-2}{2t} + \frac{1}{t+1} < 1$ by switching to any $y_{i,j}$, receives $\frac{1}{t+1} < 1$ by switching to any x_i , and receives $\frac{1}{t+1}$ by switching to v_a . Clearly, the other facilities lose more utility by switching strategies.

We transfer \mathbf{s} to the Min-2-FLG, after which the facility f_j on v_b still receives a utility of 1. We consider a switch by f_j to v_a resulting in the facility placement profile \mathbf{s}' . Since there is only one client equilibrium, by symmetry all clients x_1, \dots, x_t have the same cost and all facilities except f_j have the same utility, so we fix some arbitrary




Figure 3: An instance of the 2-FLG for which an SPE in the Uniform-2-FLG is a 2-approximate SPE in the Min-2-FLG. The improving facility improves by moving from v_b to v_a .

client $v_i \in x_1, \dots, x_t$ and some facility $f_p \neq f_j$ with $f_p \in N_{\mathbf{s}}(x_i)$.

The cost of v_i is $L_i(\mathbf{s}', \sigma) = \sigma(\mathbf{s}', v_i)_j \ell_j(\mathbf{s}', \sigma) + \sigma(\mathbf{s}', v_i)_p \ell_p(\mathbf{s}', \sigma) t = \sigma(\mathbf{s}', v_i)_j (\sigma(\mathbf{s}', v_i)_j + z) + \frac{1 - \sigma(\mathbf{s}', v_i)_j}{t} \left(\frac{1 - \sigma(\mathbf{s}', v_i)_j}{t} + \frac{2t-2}{t} \right) t$, where z is the sum of the weight that f_j receives from all clients except v_i . Because the cost is minimal, we know that for the minimum of $L_i(\mathbf{s}', \sigma)$ has $\sigma(\mathbf{s}', v_i)_j \in [0, w_i]$. Thus, we get $2\sigma(\mathbf{s}', v_i)_j + z - 2 + 2\frac{\sigma(\mathbf{s}', v_i)_j}{t} = 0$. We substitute $z = (t-1)\sigma(\mathbf{s}', v_i)_j$ because of symmetry³ and then get $\sigma(\mathbf{s}', v_i)_j = \frac{2t}{t^2+t+2}$, which is within $[0, w(v_i)]$. Thus, $L_j(\mathbf{s}', \sigma) = \frac{2t^2}{t^2+t+2}$, with $\lim_{t \rightarrow \infty} L_j(\mathbf{s}', \sigma) = 2$. \square

Next, we prove an upper bound on the approximation quality: First, we show that of all facilities in range of a client v_i , in a client equilibrium for the Min-2-FLG, the one with the lowest total load receives at least as much weight from v_i as any other facility.

Lemma 2. In the Min-2-FLG, given a facility placement profile \mathbf{s} and a client equilibrium σ , let f_j be the facility in the attraction range $N_{\mathbf{s}}(v_i)$ of client v_i with the lowest facility load. Then $\sigma(\mathbf{s}, v_i)_j \geq \sigma(\mathbf{s}, v_i)_x$ for any other facility $f_x \in N_{\mathbf{s}}(v_i)$.

Similarly, for $f_p \in N_{\mathbf{s}}(v_i)$ with the highest facility load: $\sigma(\mathbf{s}, v_i)_p \leq \sigma(\mathbf{s}, v_i)_x$ for any other facility f_x .

Proof. Let $f_x \in N_{\mathbf{s}}(v_i)$ be an arbitrary facility with non-zero weight $\sigma(\mathbf{s}, v_i)_x > 0$ on v_i . Using Lemma 1, we get

$$\begin{aligned} \ell_j(\mathbf{s}, \sigma) + \sigma(\mathbf{s}, v_i)_j &\geq \ell_x(\mathbf{s}, \sigma) + \sigma(\mathbf{s}, v_i)_x \\ \ell_j(\mathbf{s}, \sigma) + \sigma(\mathbf{s}, v_i)_j &\geq \ell_j(\mathbf{s}, \sigma) + \sigma(\mathbf{s}, v_i)_x \\ \sigma(\mathbf{s}, v_i)_j &\geq \sigma(\mathbf{s}, v_i)_x. \end{aligned}$$

Thus, f_j receives at least as much weight from v_i as any other facility in the attraction range of v_i . The proof for $f_p \in N_{\mathbf{s}}(v_i)$ with the highest facility load works analogously. \square

With this, we prove that if we move a profile \mathbf{s} from the Min-2-FLG to the Uniform-2-FLG then the facility with the lowest load in the client equilibrium of the Min-2-FLG has an equal or lower load in the Uniform-2-FLG.

³Note that we cannot do this substitution earlier, because doing it before applying the derivative would minimize the sum of utilities of all clients in $\{v_1, \dots, v_t\}$, instead of just v_i .

Lemma 3. In the Min-2-FLG, given a facility placement profile \mathbf{s} and a client equilibrium σ , let $f_j \in \mathcal{F}$ be the facility with the lowest load. Then $\ell_j(\mathbf{s}, \sigma) \geq \ell_j(\mathbf{s}, \sigma_{\text{Uniform}})$, where σ_{Uniform} is the client equilibrium for \mathbf{s} in the Uniform-2-FLG.

Similarly, for the facility $f_p \in \mathcal{F}$ with the highest load, $\ell_p(\mathbf{s}, \sigma) \leq \ell_p(\mathbf{s}, \sigma_{\text{Uniform}})$.

Proof. For each client v_i , f_j is the facility with the lowest load in $N_{\mathbf{s}}(v_i)$. Thus by Lemma 2, f_j receives as least as much weight from v_i as any other facility in the range of v_i . As only clients in $N_{\mathbf{s}}(v_i)$ receive weight from v_i , f_j receives a weight of at least $\sigma(\mathbf{s}, v_i)_j \geq \frac{w(v_i)}{|N_{\mathbf{s}}(v_i)|}$ from each v_i . Thus, for each v_i , it holds that $\sigma(\mathbf{s}, v_i)_j \geq \sigma_{\text{Uniform}}(\mathbf{s}, v_i)_j$ and therefore, $\ell_j(\mathbf{s}, \sigma) \geq \ell_j(\mathbf{s}, \sigma_{\text{Uniform}})$. The proof for $f_p \in V$ with the highest facility load works analogously. \square

We need another lemma to show that removing a facility from an instance of the Min-2-FLG does not result in a utility loss for any other facility.

Lemma 4. Let \mathbf{s} be a facility placement profile and σ be a client equilibrium in the Min-2-FLG. If we remove a facility agent f_j from the placement (and instance) resulting in \mathbf{s}' with client equilibrium σ' , no facility $f_x \neq f_j$ loses utility.

Proof. Let F_L be the set of facilities that lose utility by the removal of f_j . We assume towards contradiction that this set is non-empty. Since the sets of clients $A_{\mathbf{s}}(F_L)$ and $A_{\mathbf{s}'}(F_L)$ in the attraction range of F_L are equal for both facility placement profiles, there must be some client $v_i \in A_{\mathbf{s}}(F_L)$ which allocates more weight outside F_L in σ' than in σ . Thus, there exists a winning facility $f_w \notin F_L$, with $\ell_w(\mathbf{s}', \sigma') \geq \ell_w(\mathbf{s}, \sigma)$ and $\sigma'(\mathbf{s}', v_i)_w > \sigma(\mathbf{s}, v_i)_w$, and a losing facility $f_l \in F_L$, with $\ell_l(\mathbf{s}', \sigma') < \ell_l(\mathbf{s}, \sigma)$ and $\sigma'(\mathbf{s}', v_i)_l < \sigma(\mathbf{s}, v_i)_l$. Thus, we get the two statements

$$\begin{aligned} \ell_w(\mathbf{s}', \sigma') + \sigma'(\mathbf{s}', v_i)_w &> \ell_w(\mathbf{s}, \sigma) + \sigma(\mathbf{s}, v_i)_w \text{ and} \\ \ell_l(\mathbf{s}', \sigma') + \sigma'(\mathbf{s}', v_i)_l &< \ell_l(\mathbf{s}, \sigma) + \sigma(\mathbf{s}, v_i)_l. \end{aligned}$$

By Lemma 1, we also get

$$\begin{aligned} \ell_w(\mathbf{s}, \sigma) + \sigma(\mathbf{s}, v_i)_w &\geq \ell_l(\mathbf{s}, \sigma) + \sigma(\mathbf{s}, v_i)_l \text{ and} \\ \ell_w(\mathbf{s}', \sigma') + \sigma'(\mathbf{s}', v_i)_w &\leq \ell_l(\mathbf{s}', \sigma') + \sigma'(\mathbf{s}', v_i)_l. \end{aligned}$$

Therefore, we arrive at a contradiction. \square

We use the preceding lemmas to prove a 3-approximation:

Theorem 6. A $(1 + \epsilon)$ -approximate SPE in the Uniform-2-FLG is a $(3 + 2\epsilon)$ -approximate SPE in the Min-2-FLG.

Proof. In an SPE $(\mathbf{s}, \sigma_{\text{Uniform}})$ in the Uniform-2-FLG, the facility $f_{l, \text{Uniform}}$ with the lowest facility load and the facility $f_{h, \text{Uniform}}$ with the highest facility load are separated by at most a factor of $2 + 2\epsilon$, as otherwise $f_{l, \text{Uniform}}$ could improve by more than a factor of $1 + \epsilon$ by deviating to the location of $f_{h, \text{Uniform}}$ and thereby receive at least half her load. When transferring \mathbf{s} to the Min-2-FLG with the corresponding client equilibrium σ , the factor between $f_{l, \text{Min}}$ with the lowest facility load and $f_{h, \text{Min}}$ and with the highest facility load is also at most $2 + 2\epsilon$, by Lemma 3.

Let an arbitrary facility f_p make an improving move regarding the Min-2-FLG, changing the facility placement

Algorithm 1: Approximate Best Response Dynamics

- 1 $\mathbf{s} \leftarrow$ arbitrary facility placement profile;
 - 2 **while** $\exists f_j \in \mathcal{F}$, with its best response $s'_j \in V$ and $\ell_j((s'_j, s_{-j}), \sigma_{\text{Uniform}}) \geq (1 + \epsilon)\ell_j(\mathbf{s}, \sigma_{\text{Uniform}})$ **do**
 - 3 $\mathbf{s} \leftarrow (s'_j, s_{-j});$
-

profile from \mathbf{s} to \mathbf{s}' . Assume that f_p has the highest utility $\ell_p(\mathbf{s}', \sigma_{\text{Min}})$ of all facilities in $(\mathbf{s}', \sigma_{\text{Min}})$ or

$$\ell_p(\mathbf{s}', \sigma_{\text{Min}}) = \max_{f_x \in \mathcal{F}} \ell_x(\mathbf{s}', \sigma_{\text{Min}}). \quad (1)$$

By Lemma 2, the facility f_p receives at most a weight of $\sigma_{\text{Min}}(\mathbf{s}', v_i)_p \leq \frac{w(v_i)}{N_{\mathbf{s}'}(v_i)} = \sigma_{\text{Uniform}}(\mathbf{s}', v_i)_p$ from each client $v_i \in A_{\mathbf{s}'}(f_j)$. Thus, $\ell_p(\mathbf{s}', \sigma_{\text{Min}}) \leq \ell_p(\mathbf{s}', \sigma_{\text{Uniform}}) \leq (2 + 2\epsilon)\ell_{l, \text{Uniform}}(\mathbf{s}, \sigma_{\text{Uniform}}) \leq (2 + 2\epsilon)\ell_p(\mathbf{s}, \sigma_{\text{Min}})$, where the last part holds by Lemma 3. This means that if Equation (1) is true, the gain of facility f_p is limited to a factor of $(2 + 2\epsilon)$. Therefore, f_p can only improve by a factor of more than $(2 + 2\epsilon)$ if another facility f_x exists for which $\ell_p(\mathbf{s}', \sigma_{\text{Min}}) < \ell_x(\mathbf{s}', \sigma_{\text{Min}})$.

With the goal of finding an upper bound on $\ell_x(\mathbf{s}', \sigma_{\text{Min}})$, we investigate the move of f_p in two parts: First, the removal of f_p resulting in \mathbf{s}_r with the client equilibrium σ_r in the Min-2-FLG, second the reinsertion of f_p in her new position. By Lemma 4, all facility utilities in (\mathbf{s}_r, σ_r) have not decreased from the utilities in $(\mathbf{s}, \sigma_{\text{Min}})$. Since the sum of utilities among the non-removed facilities increases by at most $\ell_p(\mathbf{s}, \sigma_{\text{Min}})$, the maximum utility gain of f_x is at most

$$\ell_x(\mathbf{s}_r, \sigma_r) \leq \ell_x(\mathbf{s}, \sigma_{\text{Min}}) + \ell_p(\mathbf{s}, \sigma_{\text{Min}}).$$

By the inverse of Lemma 4, adding f_p in her new position cannot result in a utility gain for f_x and so $\ell_x(\mathbf{s}', \sigma_{\text{Min}}) \leq \ell_x(\mathbf{s}_r, \sigma_r)$. Thus, we have

$$\begin{aligned} \ell_x(\mathbf{s}', \sigma_{\text{Min}}) &\leq \ell_x(\mathbf{s}, \sigma_{\text{Min}}) + \ell_p(\mathbf{s}, \sigma_{\text{Min}}) \\ \ell_p(\mathbf{s}', \sigma_{\text{Min}}) &\leq \ell_x(\mathbf{s}, \sigma_{\text{Min}}) + \ell_p(\mathbf{s}, \sigma_{\text{Min}}) \\ \ell_p(\mathbf{s}', \sigma_{\text{Min}}) &\leq (3 + 2\epsilon)\ell_p(\mathbf{s}, \sigma_{\text{Min}}) \end{aligned}$$

concluding the proof. \square

We will show that there exists a simple FPTAS to compute a $(1 + \epsilon)$ -approximate equilibrium in the Uniform-2-FLG which immediately yields the following theorem.

Theorem 7. A $(3 + 2\epsilon)$ -approximate SPE in the Min-2-FLG can be computed in polynomial time.

As an algorithm to compute an $(1 + \epsilon)$ -approximate equilibrium in the Uniform-2-FLG, we employ approximate best response dynamics (see Algorithm 1). Here we iteratively let facilities switch locations if they improve the payoff by a factor of at least $1 + \epsilon$.

Theorem 8. There is a FPTAS to compute a $(1 + \epsilon)$ -approximate equilibrium in the Uniform-2-FLG.

Proof. By its stopping condition, Algorithm 1 clearly computes a $(1 + \epsilon)$ -approximate equilibrium. As for the runtime, each improvement step can be performed in polynomial time

iterating over all facilities and all their strategies. Note that computing the cost of a player for each profile can also be done in polynomial time (cf. Theorem 2). Using the following lemma to bound the overall number of steps completes the proof. \square

Lemma 5. *Every sequence of $(1 + \epsilon)$ -best response steps in the Uniform-2-FLG converges in $\mathcal{O}(\frac{1}{\epsilon} k^2 \log k)$ steps.*

Proof. For the Uniform-2-FLG, we have that $\Phi(\mathbf{s}) = \sum_{v \in V} \sum_{j=1}^{|N_{\mathbf{s}}(v)|} \frac{w(v)}{j}$ is an exact potential function that increases with each improving move of a facility exactly by the difference of the improvement (Rosenthal 1973). That is, if a facility f_j improves from \mathbf{s} by changing from s_j to s'_j with an improvement of $\Delta := \ell_j(\mathbf{s}, \sigma_{\text{Uniform}}) - \ell_j((s_{-j}, s'_j), \sigma_{\text{Uniform}})$, then $\Phi(\mathbf{s}) - \Phi(s_{-j}, s'_j) = \Delta$.

We now prove the lemma by bounding the number of best response steps until we reach an approximate equilibrium. To that end, let \mathbf{s}^* be the equilibrium that maximizes the exact potential function $\Phi(\cdot)$.

Note that an agent could always choose the location of the facility f_p that covers the most client weight in \mathbf{s}^* . That is $p = \arg \max_{\{1, \dots, k\}} w(A_{\mathbf{s}^*}(f_j))$. By an averaging argument, that weight is at least $\frac{1}{k}$ -th of the total weight $w(A_{\mathbf{s}^*}(\mathcal{F}))$ covered in \mathbf{s}^* . Hence, any best response of an agent yields a payoff of at least $\frac{1}{k}$ -th the weight covered by f_p , which is $\frac{1}{k}$ of the total load in \mathbf{s}^* . So, the payoff of a best response is at least $\frac{1}{k^2} w(A_{\mathbf{s}^*}(\mathcal{F}))$.

On the other hand, $\Phi(\mathbf{s}^*)$ is at most H_k times the total covered weight, i.e., $\Phi(\mathbf{s}^*) \leq H_k w(A_{\mathbf{s}^*}(\mathcal{F}))$. Putting both together yields that the payoff of a best response is at least $\frac{1}{k^2 H_k} \Phi(\mathbf{s}^*)$.

As we are considering only best responses that increase a player's payoff by a factor of $1 + \epsilon$, every step improves the payoff (and, hence, the potential function) by at least $\frac{\epsilon}{1 + \epsilon} \frac{1}{k^2 H_k} \Phi(\mathbf{s}^*)$.

Therefore, every sequence of $(1 + \epsilon)$ -best responses reaches an approximate equilibrium after at most $\mathcal{O}(\frac{1}{\epsilon} k^2 \log k)$ steps. \square

Notably the $(1 + \epsilon)$ -factor is unavoidable, as computing an exact equilibrium in the Uniform-2-FLG is PLS-complete.

Theorem 9. *Computing an exact equilibrium in the Uniform-2-FLG is PLS-complete.*

Proof. The problem is in PLS since we can compute the potential function value for each facility placement profile in polynomial time and we can find a better solution in polynomial time by iterating over all unilateral deviations.

To prove hardness, we reduce from LOCALMAXCUT, the local search version of MAXCUT, which is PLS-hard (Schäffer 1991; Elsässer and Tschuschner 2011). That is, given a graph $G = (V, E)$ with edge weights w_e , find a set $C \subseteq V$ such that the value of the cut, i.e., $v(C) := \sum_{u \in C} \sum_{v \in V \setminus C} w_{(u,v)}$ cannot be improved by adding or removing one vertex to or from C .

Given an instance of LOCALMAXCUT with a graph $G = (V, E)$ with edge weights w_e , we construct an instance of

the Uniform-2-FLG on a host graph $H = (V', E', w')$ with $k = |V|$ facilities such that from an equilibrium in H we can easily construct a local optimum of G . For ease of exposition, we let H be an undirected graph. One can easily obtain an equivalent directed graph by duplicating edges.

For every vertex $v \in V$, there is a vertex gadget consisting of the five nodes $left_v$, $right_v$, $dummy1_v$, $dummy2_v$, and $dummy3_v$. The three dummy vertices have weight M , the other two vertices have weight 0. There is an edge from $dummy1$ to $left$, from $dummy2$ to $right$ and edges from $dummy3$ to both, $left$ and $right$.

For every edge $e = (u, v) \in E$, there are two edge vertices $v1_e$ and $v2_e$ each with weight w_e . There is an edge from $v1$ to the $right$ vertex of the vertex gadget for u and the $left$ vertex of the vertex gadget of v . There is an edge from $v2$ to the $right$ vertex of the vertex gadget for v and the $left$ vertex of the vertex gadget of u .

As we have exactly $k = |V|$ facilities, it is easy to verify that in every equilibrium there is exactly one facility on either the $left$ or the $right$ vertex of each edge with a payoff of at least $2M$. Note that more than one player in a vertex gadget or choosing a *dummy* or *edge* vertex gives significantly less payoff.

It remains to show that we can determine a local optimum of the MAXCUT instance from any equilibrium in polynomial time. We interpret an equilibrium profile as a LOCALMAXCUT solution as follows: We define that every vertex $v \in V$ where a facility is on the left node of the corresponding gadget is in C .

In an equilibrium, the payoff of a player on a vertex gadget $u \in C$ is $2M + \sum_{v \in \delta(u) \setminus C} w_{(u,v)}$. For every player on a vertex gadget $u \notin C$, the payoff is $2M + \sum_{v \in \delta(u) \cap C} w_{(u,v)}$, where $\delta(u)$ is the set of neighboring nodes of u .

Since this is an equilibrium, deviating from the right to the left node within a gadget is not an improvement. Hence, the payoff for each player on a vertex gadget $u \in C$ is $2M + \sum_{v \in \delta(u) \setminus C} w_{(u,v)} \geq 2M + \sum_{v \in \delta(u) \cap C} w_{(u,v)}$.

Likewise, for each player on $u \notin C$ the payoff is $2M + \sum_{v \in \delta(u) \cap C} w_{(u,v)} \geq 2M + \sum_{v \in \delta(u) \setminus C} w_{(u,v)}$. \square

Conclusion

We have shown that in our model of two-sided facility location subgame perfect equilibria are not always guaranteed to exist. This is in stark contrast to the model of Krogmann et al. (2021) in which clients exhibit a simpler behavior and merely perform load balancing. To resolve non-existence we studied approximate equilibria and showed the existence and polynomial time computability of approximate equilibria.

A major open problem is whether the approximation factors can be improved. We conjecture that 2-approximate equilibria exist and that our approximation algorithm computes them. On the negative side, a close inspection of our constructions in Theorems 3 and 4 shows already that there do not exist α -approximate equilibria for a suitable small constant α and that the corresponding decision problem is intractable. It would be very interesting to obtain a matching lower bound to the existence result.

References

- Ahn, H.-K.; Cheng, S.-W.; Cheong, O.; Golin, M.; and van Oostrum, R. 2004. Competitive Facility Location: The Voronoi Game. *TCS*, 310(1): 457–467.
- Aziz, H.; Chan, H.; Lee, B.; Li, B.; and Walsh, T. 2020. Facility Location Problem with Capacity Constraints: Algorithmic and Mechanism Design Perspectives. In *AAAI*, 1806–1813.
- Bhaskar, U.; Fleischer, L.; Hoy, D.; and Huang, C.-C. 2015. On the uniqueness of equilibrium in atomic splittable routing games. *Mathematics of Operations Research*, 40(3): 634–654.
- Boppana, M.; Hod, R.; Mitzenmacher, M.; and Morgan, T. 2016. Voronoi Choice Games. In *ICALP*, 23:1–23:13.
- Brethouwer, J.-T.; de Jong, J.; Uetz, M.; and Skopalik, A. 2018. Analysis of Equilibria for Generalized Market Sharing Games. In *WINE*.
- Chan, H.; Filos-Ratsikas, A.; Li, B.; Li, M.; and Wang, C. 2021. Mechanism Design for Facility Location Problems: A Survey. In *IJCAI*, 4356–4365.
- Cohen, A.; and Peleg, D. 2019. Hotelling Games with Random Tolerance Intervals. In *WINE*, 114–128.
- de Berg, M.; Kisfaludi-Bak, S.; and Mehr, M. 2019. On One-Round Discrete Voronoi Games. In *ISAAC*, 37:1–37:17.
- Deligkas, A.; Eiben, E.; and Goldsmith, T. 2022. Parameterized Complexity of Hotelling-Downs with Party Nominees. In *IJCAI*, 244–250.
- Downs, A. 1957. An Economic Theory of Political Action in a Democracy. *Journal of Political Economy*, 65(2): 135–150.
- Dürr, C.; and Thang, N. K. 2007. Nash Equilibria in Voronoi Games on Graphs. In *ESA*, 17–28.
- Eiselt, H. A.; Laporte, G.; and Thisse, J.-F. 1993. Competitive Location Models: A Framework and Bibliography. *Transportation Science*, 27(1): 44–54.
- Elsässer, R.; and Tscheuschner, T. 2011. Settling the complexity of local max-cut (almost) completely. In *ICALP*, 171–182.
- Feldman, M.; Fiat, A.; and Golomb, I. 2016. On Voting and Facility Location. In *EC*, 269–286.
- Feldman, M.; Fiat, A.; and Obratsova, S. 2016. Variations on the Hotelling-Downs Model. In *AAAI*, 496–501.
- Feldotto, M.; Lenzner, P.; Molitor, L.; and Skopalik, A. 2019. From Hotelling to Load Balancing: Approximation and the Principle of Minimum Differentiation. In *AAMAS*, 1949–1951.
- Fotakis, D.; and Patsilinos, P. 2021. Strategyproof Facility Location in Perturbation Stable Instances. In *WINE*, 95–112.
- Fournier, G.; Van der Straeten, K.; and Weibull, J. 2020. Spatial competition with unit-demand functions. arXiv:2001.11422.
- Gairing, M. 2009. Covering Games: Approximation through Non-Cooperation. In *WINE*, 184–195.
- Garey, M. R.; and Johnson, D. S. 1990. *Computers and Intractability; A Guide to the Theory of NP-Completeness*. USA: W. H. Freeman & Co.
- Goemans, M.; Li, E. L.; Mirrokni, V. S.; and Thottan, M. 2006. Market Sharing Games Applied to Content Distribution in Ad hoc Networks. *Journal on Selected Areas in Communications*, 24(5): 1020–1033.
- Harks, T.; and Timmermans, V. 2021. Equilibrium computation in resource allocation games. *Mathematical Programming*, 194: 1–34.
- Harrenstein, P.; Lisowski, G.; Sridharan, R.; and Turrini, P. 2021. A Hotelling-Downs Framework for Party Nominees. In *AAMAS*, 593–601.
- Hotelling, H. 1929. Stability in Competition. *The Economic Journal*, 39(153): 41–57.
- Kakutani, S. 1941. A generalization of Brouwer’s fixed point theorem. *Duke Mathematical Journal*, 8(3): 457–459.
- Kohlberg, E. 1983. Equilibrium Store Locations When Consumers Minimize Travel Time Plus Waiting Time. *Economics Letters*, 11(3): 211–216.
- Krogmann, S.; Lenzner, P.; Molitor, L.; and Skopalik, A. 2021. Two-Stage Facility Location Games with Strategic Clients and Facilities. In *IJCAI*, 292–298.
- Krogmann, S.; Lenzner, P.; and Skopalik, A. 2022. Strategic Facility Location with Clients that Minimize Total Waiting Time. arXiv:2211.14016.
- Mavronicolas, M.; Monien, B.; Papadopoulou, V. G.; and Schoppmann, F. 2008. Voronoi Games on Cycle Graphs. In *MFCS*, 503–514.
- Monderer, D.; and Shapley, L. S. 1996. Potential Games. *Games and Economic Behavior*, 14(1): 124–143.
- Peters, H.; Schröder, M.; and Vermeulen, D. 2018. Hotelling’s location model with negative network externalities. *International Journal of Game Theory*, 47: 811–837.
- Procaccia, A. D.; and Tennenholtz, M. 2013. Approximate Mechanism Design without Money. *ACM Trans. Economics and Comput.*, 1(4): 18:1–18:26.
- ReVelle, C. S.; and Eiselt, H. A. 2005. Location Analysis: A Synthesis and Survey. *European Journal of Operational Research*, 165(1): 1–19.
- Rosenthal, R. W. 1973. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, 2: 65–67.
- Schäffer, A. A. 1991. Simple local search problems that are hard to solve. *SIAM Journal on Computing*, 20(1): 56–87.
- Schmand, D.; Schröder, M.; and Skopalik, A. 2019. Network Investment Games with Wardrop Followers. In *ICALP*, 151:1–151:14.
- Schön, C.; and Saini, P. 2018. Market-oriented service network design when demand is sensitive to congestion. *Transportation Science*, 52(5): 1253–1275.
- Shen, W.; and Wang, Z. 2017. Hotelling-Downs Model with Limited Attraction. In *AAMAS*, 660–668.
- Vetta, A. 2002. Nash Equilibria in Competitive Societies, with Applications to Facility Location, Traffic Routing and Auctions. *FOCS*, 416–425.