



On Fair Cost Facility Location Games with Non-singleton Players

Félix Carvalho Rodrigues^{1,2} Eduardo C. Xavier^{1,3}

*Institute of Computing
University of Campinas (Unicamp)
Campinas, Brazil*

Guido Schäfer⁴

*Networks and Optimization (N&O) group
Centrum Wiskunde & Informatica (CWI)
Amsterdam, The Netherlands*

Abstract

In the fair cost facility location game, players control terminals and must open and connect each terminal to a facility, while paying connection costs and equally sharing the opening costs associated with the facilities it connects to. In most of the literature, it is assumed that each player control a single terminal. We explore a more general version of the game where each player may control multiple terminals. We prove that this game does not always possess pure Nash equilibria, and deciding whether an instance has equilibria is NP-Hard, even in metric instances. Furthermore, we present results regarding the efficiency of equilibria, showing that the price of stability of this game is equal to the price of anarchy, in both uncapacitated and capacitated settings.

Keywords: price of stability, facility location, algorithmic game theory

1 Introduction

Facility location problems covers a broad range of optimization problems, with practical applications in many different areas such as public policy, urban planning, telecommunications and computer networking. In the general sense, the facility location problem can be stated as follows. Let F be a set of facilities, T a set of terminals, with opening costs c_f for each facility $f \in F$ and connection costs d_{tf}

¹ This work was partially supported by Capes, CNPq and Fapesp.

² Email: felix.rodrigues@ic.unicamp.br

³ Email: ecx@ic.unicamp.br

⁴ Email: g.schaefer@cwi.nl

for connecting terminal $t \in T$ to facility $f \in F$. The problem is to find a subset of facilities to open and establish connections from terminals to this subset such that the sum of all costs are minimized.

Consider a scenario where multiple supermarket or stores share big warehouses that stockpile supplies for them. Each store is located some distance away from a warehouse and they pay the costs associated with maintaining its products in the warehouse and the transportation costs to its specific location. Such scenario can be seen as a facility location problem, where each store is a terminal and each warehouse location is a possible facility. In the classical optimization version, the view point of the warehouse company is given priority, and each terminal can be seen as being controlled by a central authority in order to minimize opening and connection costs globally. When we consider the fact that each supermarket is competing with each other, this approach of global optimization is not possible to be adopted anymore. In order to analyze such scenarios, we can use game theory.

A *non-cooperative game* is a decision scenario where an agent or player selfishly and without coordination other players, chooses a strategy in order to maximize its own utility (or minimize its own cost), which in turn depends on the strategies chosen by the other players. We say that a game is in a pure Nash equilibrium (PNE) if no player has any incentive to unilaterally change its own strategy. In order to compare the social cost of pure equilibria and the social optimum, we use the standard measures found in the literature: the price of anarchy (PoA) and the price of stability (PoS). The PoA of a game is defined as the ratio between the PNE with worst social cost and the social optimal cost, while the PoS is the ratio between the PNE with the best social cost and the social optimum.

In our example scenario, each store is traditionally viewed as a player choosing a facility (warehouses) to connect, such that its own cost is minimized. Each store would then share equally the costs associated with storing the goods in the warehouse they are connected to, as well as individually paying for the connection to such warehouse. This scenario exemplifies the singleton fair cost facility location game, where each store is controlled by a single player. This scenario however fails to accommodate the common case when multiple supermarket or stores are part of the same chain.

When stores are part of the same larger group, each store behaviour cannot be said to be independent and uncoordinated, however one whole group does not coordinate with other competitor stores or chains and therefore it makes sense to analyze this scenario by allowing players to control multiple stores or terminals.

In the fair cost facility location games we analyze, each player controls multiple terminals, and can move them simultaneously to minimize their own costs. For singleton games, terminals sharing opening costs evenly means the same as players sharing these costs evenly, since players and terminals are effectively the same in these scenarios. However when players control multiple terminals, the same is not true. In our paper we use the term fair cost to mean that in any strategy profile each facility has its opening cost shared evenly among all terminals that connect to it, i.e. if in a certain strategy profile a player a has two terminals connected to a

facility f and a different player b has a single terminal connected to it, player a will pay for two thirds of the opening cost of f , while player b pays only a third of the opening cost.

In this work we study how hard it is to find pure Nash equilibria in these fair cost facility location games, as well as how efficient these equilibria can be when compared with the social optimal cost. Finally, we extend our results to weighted and capacitated versions of these games.

2 Related Work and Organization

There have been multiple works examining facility location problems from the outlook of game theory. The class of valid utility games [12, 14, 18] can be viewed as facility games when both facilities and players are controlled by singleton players. There is a whole area in game theory focused in cooperative games, and Goemans and Skutella [8] studied cooperative facility location games.

Facility location has also been extensively studied from the mechanism design perspective, with multiple relevant work in strategy-proof mechanisms for variants of facility location problems [5, 13, 15].

We study in this paper the non-cooperative facility location game where players control terminals, and each terminal t share the opening cost of its chosen facility f equally among all terminals connected to f . These games share great similarities to connection and network design games, and thus multiple results from these games are valid for fair cost facility location games.

For the singleton fair cost facility location game, the results from network design [1] extend to facility location, with a bound of k for the PoA and $H_k = \Theta(\log k)$ for the PoS [17]. For the metric version, where each connection obeys the triangle inequality, Hansen and Telelis [9] proved constant bounds both for the PoS and the strong PoA. Furthermore, when there are weights associated with terminals determining how much of the opening cost is paid by each terminal, they show that there is always an ϵ -approximate equilibrium and that the PoS can be in $\Theta(\log W)$, where W is the sum of all terminal weights. In [3], Chen and Roughgarden prove that there are instances of the weighted network design game where there are no possible PNE, however it is still an open question whether the same applies to weighted fair cost facility location games, even for the non-singleton version.

For facility location games where players do not possess any limitations on how to share opening costs, the price of anarchy and the price of stability have been proven to be in $\Theta(k)$ by Cardinal and Hoefer [2, 11], where k is the number of players in the game. Furthermore, for non-singleton games it is NP-hard to decide whether an instance has a PNE [2].

The capacitated version of the game was considered by Rodrigues and Xavier [16]. They show that for the metric singleton fair cost facility location game, the PoA can be unbounded, while the PoS is H_k , where k is the number of players. Furthermore, they show that it is NP-hard to decide whether there is PNE in singleton capacitated facility location games with arbitrary opening cost sharing.

For a sequential version of facility location, they show that the metric version has bounded PoA and PoS.

In this paper we extend the results from Cardinal and Hoefer [2] to fair cost facility location, proving that it is NP-hard to decide whether there is a PNE in this game, even in the metric case, as long as players are allowed to control multiple terminals. For this, we prove that there are fair cost facility location game instances with no PNE. We base this proof in the instances without PNE in weighted network design games used by Chen and Roughgarden in [3]. Finally, we prove that the PoS of this game can be as inefficient as the PoA for both the uncapacitated metric case and the general capacitated case.

In Section 3, we present the formal definitions needed to understand our paper and define our game. In Section 4, we provide instances with no PNE and prove the NP-hardness of PNE existence, as well as link these instances to weighted games. In Section 5, we prove that the price of stability is as poor as the price of anarchy in both the uncapacitated and capacitated versions of fair cost facility location games. Finally, in Section 6 we present our final remarks and discuss future work.

3 Preliminaries

First we define a few key concepts from non-cooperative game theory.

In game theory, a *non-cooperative game* is a scenario where players choose strategies independently trying to either minimize their costs or maximize their utility. For each player i , there is a set A_i of actions that it can choose to play. A pure strategy S_i consists of one action from A_i , while a mixed strategy corresponds to a probability distribution over A_i . In a pure game each player chooses one action to play, while in a mixed game each player randomizes its action according to the probability distribution. In this paper we assume pure strategies games unless mentioned otherwise.

A set of strategies $S = (S_1, S_2, \dots, S_n)$ consisting of one strategy for each player, is denominated a *strategy profile*. Let $\mathcal{S} = A_1 \times A_2 \times \dots \times A_n$ be the set of all possible strategy profiles and let $c : \mathcal{S} \rightarrow \mathbb{R}^n$ be a cost function that attributes a cost $c_i(S)$ for each player i given a strategy profile S . Define $S_{-i} = (S_1, \dots, S_{i-1}, S_{i+1}, \dots, S_n)$ a strategy profile S without i 's strategy, so that we can write $S = (S_i, S_{-i})$. If all players other than i decide to play S_{-i} , then player i is faced with the problem of determining a best response to S_{-i} . A strategy S_i^* from a player i is a *best response* to S_{-i} , if there is no other strategy which could yield a better outcome for the player, i.e.

$$c_i(S_i^*, S_{-i}) \leq c_i(S_i, S_{-i}) \quad , \forall S_i \in A_i.$$

A strategy profile is in a *pure Nash equilibrium* (PNE) if no player can reduce its cost by choosing a different strategy, i.e. for each player, its strategy in the strategy profile is a best response.

The social welfare or social cost is a function mapping a strategy profile to a real number, indicating a measure of the total cost or payoff of a game. Two of the most important concepts for efficiency analysis are the *Price of Anarchy* (PoA) and

the *Price of Stability* (PoS). The PoA is the ratio between a Nash equilibrium with worst possible social cost and the strategy profile with optimal social cost, while the PoS is the ratio between the best possible Nash equilibrium to the social optimum.

With these ideas formalized, we define the *Facility Location Game with Fair Cost Sharing* (FLG-FC).

Definition 3.1 Let $G = (T \cup F, T \times F)$ be a bipartite graph, with vertex sets F of n facilities and T of m terminals. Each facility $f \in F$ has an opening cost c_f , and connection costs d_{tf} for each terminal $t \in T$. Let $K = \{1, \dots, k\}$ be the set of players.

Each player i controls a subset of terminals $T_i \subseteq T$ (also forming a partition of T), and each terminal must be connected to exactly one opened facility. When a player controls only a single terminal, it is denominated a *singleton* player. A player i chooses a strategy $S_i \subseteq T_i \times F$.

Let $S = (S_1, \dots, S_k)$ be a strategy profile. We abuse notation and use the expression $f \in S$ to represent any facility f connected to a terminal in a strategy profile S and $(t, f) \in S$ to represent any pair of terminal and facility that are connected in S , while $f \in S_i$ represents any facility f player i uses to connect one of its terminals in strategy S_i . Each player tries to minimize its own payment

$$p_i(S) = \sum_{(t,f) \in S_i} \frac{c_f}{x_f(S)} + \sum_{(t,f) \in S_i} d_{tf} ,$$

where $x_f(S) = |\{(t_j, f) \in S_i \mid 1 \leq i \leq k \wedge 1 \leq j \leq m\}|$ is the number of terminals connected to facility f in strategy profile S .

The *social welfare cost* for a strategy S is defined as the sum of all player payments, i.e.,

$$C(S) = \sum_{i \in K} p_i(S) = \sum_{f \in S} c_f + \sum_{(t,f) \in S} d_{tf} .$$

In games with general connection costs, some connections (t, f) should be avoided in any solution, because they do not exist for example. In this case we assume they have a prohibitively large constant cost \mathcal{U}_d . For general costs, if a connection is not shown, it is assumed that it has a cost equal to \mathcal{U}_d , unless mentioned otherwise. For metric connection costs, where connection costs must obey the triangle inequality, it is assumed that any connection (t, f) not shown has a cost equal to the shortest cost path from t to f in the undirected graph formed from the explicitly shown connections.

When dealing with capacitated FLG-FC, we extend this definition to include a capacity for each facility, as well as ways to enforce players to propose valid solutions.

Definition 3.2 Let $G = (T \cup F, T \times F)$ be a bipartite graph, with vertex sets F of n facilities and T of m terminals. Each facility $f \in F$ has an opening cost c_f and a capacity u_f indicating how many terminals can be connected to f at any given time. Furthermore, there are connection costs d_{tf} for each pair (t, f) where $t \in T$ and $f \in F$.

Given a strategy profile S , for *Capacitated Facility Location Games with Fair Cost Sharing (CFLG-FC)*, the definitions of player i payment, $p_i(S)$, and social cost, $C(S)$, are the same as before for the uncapacitated game. However, to ensure that capacity restrictions are respected, if a player i in the solution S has one of its terminals connected to f where $x_f(S) > u_f$, then a prohibitively large constant cost \mathcal{U}_c is added to the payment of player i , i.e., it pays $p_i(S) + \mathcal{U}_c$.

4 On the Existence of Pure Equilibria

Pure Nash equilibria is one of the most well known solution concepts in game theory. While it is not guaranteed to exist in all games, in several practical scenarios it reflects to a greater degree the behaviour of players. The singleton version of FLG-FC is a potential game, and therefore there always exists a PNE [17]. We show that this does not extend to every fair cost facility location game, by showing instances with no PNE when players control multiple terminals and the opening cost sharing happens in relation to terminals. We construct these instances loosely based on an example used to prove the non-existence of PNE in weighted network design games [3]. We first prove this result for weighted fair cost facility location games and then extend this result for unweighted metric games.

4.1 Equilibria Existence in Weighted Games

In the classic facility location game, each terminal is assumed to demand or require the same amount of goods from a facility, and therefore the “fair” way to share costs is to evenly divide opening costs of an opened facility between terminals that connect to it. However this is not always the case. In several practical scenarios some terminals might require more from a facility, and an egalitarian sharing might not reflect fairness.

In the *weighted fair cost facility location game*, each terminal t has an associated positive integer weight $w_t \geq 1$, and each player i pays in a strategy profile S ,

$$p_i(S) = \sum_{(t,f) \in S_i} d_{tf} + \sum_{(t,f) \in S_i} \frac{w_t c_f}{W_{f,S}},$$

where $W_{f,S}$ is the sum of the weights of all terminals connected to f in the strategy profile S .

Here we provide a partial answer to an open question regarding whether there are instances with no PNE for weighted FLG-FC. We provide an instance with no PNE when players are allowed to control more than a single terminal. It remains open whether there are singleton weighted instances with no equilibria.

Theorem 4.1 *There exists a 3-player metric instance for the weighted FLG-FC game with only six terminals where there is no PNE.*

Proof.

Consider the instance in Figure 1, denominated here as I . Let $w > 1$ be a parameter of this instance, and ε be a constant much smaller than $\frac{1}{w^3}$. Let player A control terminals t_2 and t_5 , player B control terminals t_1 and t_3 , and player C control terminals t_4 and t_6 . Allow every terminal controlled by player A to have the same weight $w_A = w^2$, every terminal that B controls to have unitary weight $w_B = 1$ and every terminal that player C controls to have weight $w_C = w$.

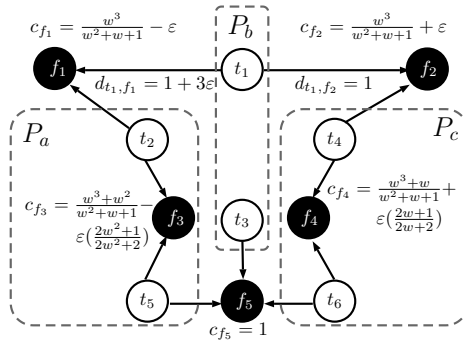


Fig. 1. Game instance of the weighted FLG-FC without PNE. All edges except d_{t_1, f_1} have cost equal to zero.

For player P_a , there are only two feasible strategies: either all terminals connect to f_3 , or t_2 connects to f_1 and t_5 connects to f_5 . The same happens for player P_c : either all terminals connect to f_4 , or t_4 connects to f_2 and t_6 connects to f_5 . For player P_b terminal t_1 has to choose between f_1 and f_2 , while terminal t_3 will always connect to f_5 in any PNE. The proof is based on players P_a and P_c having different facility preferences when presented with mirrored choices from player P_b . To achieve this, we use the fact that player P_a has squared times the weight than P_c does, as well as carefully constructed opening costs.

Note that all direct connection costs are equal for P_a and for P_c , and thus do not interfere in their choices. They only ensure that it is not beneficial for neither P_a nor P_c to connect their terminals to facilities without a direct connection available. For P_b , the only influence that connection costs have is in choosing whether to connect to f_1 or f_2 .

Since the terminals from player P_a have weight w^2 , player A will always pay for the majority of the cost of any facility it helps open. With this in mind, we can sort the five possible scenarios for player P_a by the cost incurred from each in increasing order: (i) player P_a shares f_1 with P_b and f_5 with all players, (ii) it connects to f_1 alone and shares f_5 with all players, (iii) it connects to f_3 alone, (iv) it shares f_1 with P_b and shares f_5 with P_b only and (v) it connects to f_1 alone and shares f_5 with P_b only. The following inequalities show why this is the case:

$$\frac{w^2}{w^2 + 1}c_{f_1} + \frac{w^2}{w^2 + w + 1}c_{f_5} \leq \tag{1}$$

$$c_{f_1} + \frac{w^2}{w^2 + w + 1}c_{f_5} \leq \tag{2}$$

$$\frac{w^3 + w^2}{w^2 + w + 1} - \varepsilon \left(\frac{2w^2 + 1}{2w^2 + 2} \right) = c_{f_3} \leq \tag{3}$$

$$\frac{w^2}{w^2 + 1} \left(\frac{w^3}{w^2 + w + 1} - \varepsilon \right) + \frac{w^2}{w^2 + 1} = \frac{w^2}{w^2 + 1}c_{f_1} + \frac{w^2}{w^2 + 1}c_{f_5} \leq c_{f_1} + \frac{w^2}{w^2 + 1}c_{f_5}. \tag{4}$$

From this we gather that, for player P_a , the preferred scenario is where player P_c connects to f_5 , even if player P_b does not connect to f_1 . For player P_c , we do the same, observing that now the facility which dictates its strategy is f_2 , as even if player P_a does not connect to f_5 , if player P_b does connect to f_2 , the cheapest for P_c is to connect to f_2 and f_5 . Player P_c can, going from least to most expensive, (i) share f_2 with P_b and f_5 with all players, (ii) share f_2 and f_5 with P_b only, (iii) connect to f_4 alone, (iv) connect to f_2 alone and share f_5 with all players and finally (v) connect alone to f_2 and share f_5 with P_b only. The inequalities that show this is the case for player P_c are:

$$\frac{w}{w + 1}c_{f_2} + \frac{w}{w^2 + w + 1}c_{f_5} \leq \tag{5}$$

$$\frac{w}{w + 1}c_{f_2} + \frac{w}{w + 1}c_{f_5} = \frac{w^3 + w}{w^2 + w + 1} + \varepsilon \frac{w}{w + 1} \leq \tag{6}$$

$$\frac{w^3 + w}{w^2 + w + 1} + \varepsilon \left(\frac{2w + 1}{2w + 2} \right) = c_{f_4} \leq \tag{7}$$

$$\frac{w^3}{w^2 + w + 1} + \varepsilon + \frac{w}{w^2 + w + 1} = c_{f_2} + \frac{w}{w^2 + w + 1}c_{f_5} \leq \tag{8}$$

$$c_{f_2} + \frac{w}{w + 1}c_{f_5}. \tag{9}$$

Now we prove the theorem by contradiction. Suppose that there exists a PNE for this instance. Terminal t_1 can be connected to either f_1 or f_2 . First assume t_1 is connected to f_1 . Then, player P_c has no incentive to connect to either f_2 or f_5 , as shown by (7) and (8), and connects all its terminals to f_4 . Since player P_c does not connect to f_5 , player P_a also does not have enough incentive to connect to f_1 and f_5 , as shown by (3) and (4), and connects all its terminals to f_3 . With this, player P_b is paying alone for f_1 , and since f_2 is cheaper to connect, it is not in a PNE, and thus t_1 cannot connect to f_1 in any PNE.

Now assume t_1 is connected to f_2 . Player P_c now will connect to f_2 with t_4 , with t_6 connecting to f_5 , as the inequalities in (6) and (7) show. Since player P_c connects to f_5 , now player P_a will also opt to connect to f_5 and therefore will also open f_1 , as shown by (2) and (3). Since f_1 has player P_a connected to it, player

P_b now has enough incentive to connect t_1 to f_1 instead of f_2 , and therefore our assumption is not true.

Since connecting t_1 to neither f_1 nor f_2 results in a PNE, there is a contradiction with our claim that there exists a PNE for this instance, and thus we have proven that there is no pure equilibrium. \square

4.2 Equilibria Existence in FLG-FC

Now we extend the results from Theorem 4.1 for the unweighted case. To accomplish this we first make the key observation that for any weighted fair cost facility location game instance G with m terminals, there is a fair cost facility location game instance G' with W terminals with equivalent PNE, where W is the total sum of weights $w_1 + \dots + w_m$. This is the case since for each terminal t controlled by a player i with weight $w_t > 1$, we can add in G' terminals t_j , for $j \in [2, w_t]$, all controlled by i with the same connections and costs as t . Since they are equal in every way in regards to cost calculation, player i has no reason to split strategies when choosing which facility to connect t and the added terminals t_j .

Theorem 4.2 *There exists a metric 3-player instance for the FLG-FC game where there is no PNE.*

Proof. Consider the instance I depicted in Figure 2. Let $w > 1$ be a parameter on this graph, and ε be a constant much smaller than $\frac{1}{w^3}$, such that player P_a controls terminals $t_2^1, \dots, t_2^{w^2}$ and $t_5^1, \dots, t_5^{w^2}$, for a total of $2w^2$ terminals. Player P_b controls two terminals, t_1 and t_3 , while player P_c controls terminals t_4^1, \dots, t_4^w and t_6^1, \dots, t_6^w , for a total of $2w$ terminals. All connection costs are unitary, with the exception of d_{t_1, f_1} , which has cost $1 + 3\varepsilon$. All opening costs are as shown in Figure 2.




Fig. 2. Game instance of the FLG-FC without a PNE. All edges except d_{t_1, f_1} have cost equal to one. Any edge (t, f) not drawn has cost equal to the shortest path cost from t to f .

We prove the theorem by showing that the instance seen in Theorem 4.1, here denominated I_0 , can be transformed to instance I , seen in Figure 2 without changes to the players overall strategies, and thus the proof for Theorem 4.1 applies for instance I . Start by changing I_0 to I' so that we add terminals t'_2 and t'_5 connected to the same facilities (and the corresponding connection costs) as t_2 and t_5 , respectively, while changing the weight w_A of all terminals of player A to $\frac{w^2}{2}$. In any pure

equilibria for this instance, if we connect any terminal to facility f_3 , all terminals from A will connect to it, since A will pay the full cost of f_3 . Thus, if t_2 connects to f_3 , terminal t'_2 will also connect to f_3 , and if t_2 connects to f_1 , t'_2 will do the same. The same is true for t'_5 in relation to t_5 . Since this is the case, if player A connects to f_1 or f_5 , it will still pay exactly the same share $\frac{w^2}{W(f,S)}$ of the opening cost of facility f as in I , where $W(f,S)$ is the sum of weights of all terminals connected to f in the strategy profile S . Instance I' thus incurs the same decisions from players as instance I_0 . The same can be done to player C and its terminals t_4 and t_6 .

Therefore, in order to transform I_0 into instance I while preserving the same possible equilibria, it suffices to incrementally add terminals with the same connections for players A and C , while dividing the weight of these players by the number of added terminals until the number of terminals A controls is $2w^2$ (w^2 of “ t_2 ” terminals and w^2 of “ t_5 ”) and they all have weight $w_A = 1$, while player C will control $2w$ terminals (w of “ t_4 ” terminals and w of “ t_6 ”), all with weight $w_C = 1$. Thus, Theorem 4.1 applies to the unweighted metric instance I , and there exists a 3-player weighted instance I without any equilibria. \square

Using a version of this instance where $w = 2$ as a gadget, we can prove that deciding whether an instance of FLG-FC has a PNE is NP-Hard.

Theorem 4.3 *It is NP-hard to determine if an instance of the Metric FLG-FC has a PNE or not.*

Proof. First notice that it is NP-hard to verify whether a given solution S to an instance of the FLG-FC is a PNE or not when players control multiple facilities, since for each player we need to solve an optimization facility location problem on the instance restricted to this player alone. To verify that a strategy of a player is a best response, we have to fix all other players strategies, and check if the solution restricted to this player is minimum, which is an NP-hard problem.

In order to prove NP-hardness for the existence of a PNE, we make a reduction from the 3-SAT problem to the problem of deciding whether an instance of the metric FLG-FC game has a PNE or not.

Let I be an instance from 3-SAT, with clauses C_1, \dots, C_q , where each clause C_j consists of a triple of literals from the set of decision variables x_1, \dots, x_p in clausal normal form, so that each literal in a clause can assume form x_i or \bar{x}_i for a decision variable x_i . Then, we form the fair cost facility location instance G as follows: for each decision variable x_i , we create a terminal t_i , controlled by a single player, and facilities f_{x_i} and $f_{\bar{x}_i}$ and link t_i to them, as Figure 3a shows. The opening cost of each facility f_{x_i} and $f_{\bar{x}_i}$ is equal to the number of terminals that can directly connect to them times $1 + \varepsilon$. We use the term directly here to discern from the connections not shown, which are assumed to cost the shortest path from a terminal to the facility.

We create for each clause C_j two gadgets: one with no equilibria unless stabilized by the second gadget, which in turn links to the decision gadget, as exemplified in Figure 3b. For the bottom gadget, note that it is exactly the same example as the one used in Theorem 4.2, with parameter $w = 2$. We have three players in

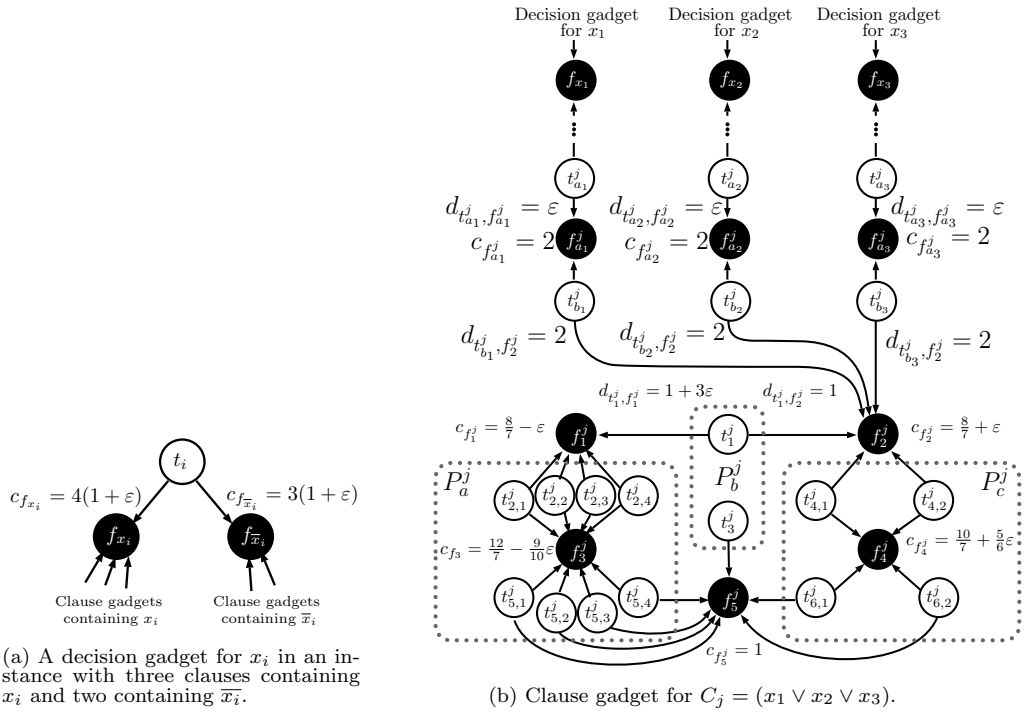


Fig. 3. A decision (a) and a clause (b) gadget. Drawn edges without explicit costs have connection cost one, any edge (t, f) not drawn has cost equal to the shortest path cost from t to f .

this gadget for each clause C_j , where player P_a^j controls terminals $t_{2,1}^j, \dots, t_{2,4}^j$ and $t_{5,1}^j, \dots, t_{5,4}^j$, player P_b^j controls terminals t_1^j and t_3^j and player P_c^j control terminals $t_{4,1}^j, t_{4,2}^j, t_{6,1}^j, t_{6,2}^j$. The opening and connection costs of this gadget follow the same costs presented in Theorem 4.2, when parameter w equals to two.

For the link gadget, each player controls a single terminal, and for each literal x_i or \bar{x}_i in C_j there is a terminal t_{ai}^j with possible direct connection to both the facility in the decision gadget for x_i (either f_{x_i} or $f_{\bar{x}_i}$), with connection cost one, and central facility f_{ai}^j , with connection cost ε and opening cost 2. Furthermore, we create another terminal t_{bi}^j which can connect directly to either the central facility f_{ai}^j , with connection cost one, or to facility f_2^j with connection cost 2. Now we add connections from every terminal t to every facility f obeying the triangle inequality by setting cost d_{tf} to the shortest cost path from t to f in the undirected network formed by the connections between terminals and facilities.

Suppose there is a truth assignment for I . Then we can build the following PNE for instance G : if decision variable $x_i = 1$, assign t_i to f_{x_i} , and otherwise assign t_i to $f_{\bar{x}_i}$. Now for each clause gadget C_j where literal x_i appears connect t_{ai}^j to f_{x_i} if $x_i = 1$, or to f_{ai}^j if $x_i = 0$. Furthermore, connect t_{bi}^j to f_2^j if $x_i = 1$, or to f_{ai}^j if $x_i = 0$. Do the same for clauses where literal \bar{x}_i appears.

Finally, connect the rest of the terminals in C_j as follows: player P_a^j connects its terminals to f_1^j and f_5^j , player P_b^j connects to f_2^j and f_5^j and player P_c^j connects to f_2^j and f_5^j .

First, note that if $x_i = 1$, then f_{x_i} is open and connected with all terminals that can directly connect to it, meaning that each terminal pays $1 + \varepsilon$ opening cost. At the same time, in each clause where x_i appears, $f_{a_i}^j$ is not opened, and thus the cheapest alternative for terminal $t_{a_i}^j$ costs $2 + \varepsilon$, the same amount it pays for connecting to f_{x_i} . Alternatively, if $x_i = 0$, facility f_{x_i} is not open and $t_{a_i}^j$ is connected to its cheapest possible facility.

Since I has a satisfying assignment, at least one terminal $t_{b_i}^j$ will be connected to f_2^j , paying at most $2 + \frac{2}{7} + \frac{\varepsilon}{4}$ and stabilizing the game induced by players P_a^j, P_b^j and P_c^j . Furthermore, when $x_i = 1$, terminal $t_{b_i}^j$ is in equilibrium connected to f_2^j , since the alternative facility $f_{a_i}^j$ would require $t_{b_i}^j$ to pay 3 in opening and connection costs. Alternatively, if $x_i = 0$, terminal $t_{b_i}^j$ is also in equilibrium by connecting to $f_{a_i}^j$, since in this case it shares the opening cost with $t_{a_i}^j$, paying a total of 2. Therefore, we have a PNE for G given a satisfying assignment for I .

Now suppose there exists a PNE in G . For any clause C_j , at least one terminal $t_{b_i}^j$ must be connected to f_2^j , since otherwise players P_a^j, P_b^j and P_c^j would not be in equilibrium. Then, for each terminal $t_{b_i}^j$ connected to f_2^j , terminal $t_{a_i}^j$ must connect to f_{x_i} (or $f_{\bar{x}_i}$, if representing literal \bar{x}_i). This is the case because if $t_{a_i}^j$ were to connect to $f_{a_i}^j$, terminal $t_{b_i}^j$ would not be in equilibrium connecting to f_2^j .

Note that for terminal $t_{a_i}^j$ in clause C_j connected to facility f_{x_i} , even when all terminals with a direct possible connection apart from t_i are connected to it, terminal $t_{a_i}^j$ still is not in equilibrium, since it would be cheaper for it to open $f_{a_i}^j$ paying ε connection cost and 2 opening cost than paying 1 as connection cost and $\frac{l}{l-1}(1 + \varepsilon)$ opening cost, where l is the number of terminals that can connect directly to f_{x_i} . Terminal t_i in any PNE therefore must connect to either f_{x_i} or $f_{\bar{x}_i}$.

This shows that in any PNE in G , at least one literal of each clause is set to true in the corresponding assignment in I . Furthermore, it also shows that there can be no decision variable x_i with $x_i = 1$ and $\bar{x}_i = 1$. If a decision variable x_i is set as both $x_i = 0$ and $\bar{x}_i = 0$, then it must be irrelevant to the truth assignment of I , since G is in equilibrium, and therefore we can assign either x_i or \bar{x}_i to 1. Thus, any PNE in G is consistent and can be used to build a satisfying assignment for I . \square

Note that while we use several players in our proof, it is possible to reduce the total number of players to only six.

Corollary 4.4 *It is NP-hard to determine if a metric FLG-FC has a PNE, even for games with 6 players.*

Proof. First note that in the instance G in Theorem 4.3, there are three players which control multiple terminals for each clause C_j , as well as six singleton players. Furthermore, for each decision gadget there is one additional player. We form a new instance G' with the same terminals, facilities and costs as G , but only six players.

We remark that each clause gadget is “de facto” isolated from each other in G , since the connection cost necessary for a terminal in a clause gadget C_j to connect to a facility in either a different clause gadget C_{j+1} or in a decision gadget not directly

connected to C_j is greater than opening any possible facility it can directly connect. Thus, we can form players P_a, P_b and P_c in instance G' that control terminals controlled by players P_a^j, P_b^j and P_c^j in any clause C_j . Furthermore, we can make players P_{a1}, P_{a2}, P_{a3} and P_{b1}, P_{b2}, P_{b3} in G' to control every terminal $t_{a1}^j, t_{a2}^j, t_{a3}^j$ and $t_{b1}^j, t_{b2}^j, t_{b3}^j$ in any clause gadget C_j .

Similarly, we can join all players from the decision gadgets into a single player P_x in G' , which controls any terminal t_i in any decision gadget x_i . Now the total number of players in this instance is reduced to ten. To achieve the number of six players, we remark that player P_{a1} has no direct connection to facilities that P_{b2} can directly connect, and thus can be safely merged into a single player P_{a1b2} . Similarly, player P_{a2} can be safely merged with P_{b3} and P_{a3} can be merged with P_{b1} , resulting in players P_{a2b3} and P_{a3b1} in G' . Finally, note that player P_a has no direct connection to f_2^j in any clause gadget C_j , and therefore can safely merge with either P_{a1b2}, P_{a2b3} or P_{a3b1} . Thus, we have one player controlling all decision gadgets and five players controlling the clause gadgets, resulting in a total of six players in instance G' while maintaining the same features and possible PNE as instance G . \square

5 Efficiency of Equilibria in Non-singleton Fair Cost Facility Location Games

In this section we turn our analysis towards loss of efficiency due to player behaviour. The most well known measures of efficiency for non-cooperative games are the PoA and the PoS. The price of anarchy in several games can be high and sometimes unrealistic, due to comparing only the worst possible equilibrium in terms of social cost or welfare to the optimal social welfare or cost. For facility location games, we can show that the price of anarchy is at least k , where k is the number of players of the game. To see that this is the case, see Figure 4. In this example, the strategy where all players are connected to the facility to the right is a PNE, and has social cost k , while in the optimal social solution all players are connected to the facility to the left, with total cost equal to 1 (connection costs are zero).




Fig. 4. A game instance of the FLG-FC with PoA equal to k .

In [1], Anshelevich et al. argue that the concept of PoS, formalized in the same paper, better captures the efficiency loss in network design games and consequently singleton fair cost facility location games. They show a bound for the PoS of

$H_k = 1 + \frac{1}{2} + \dots + \frac{1}{k} = \Theta(\log k)$ for these network design games, where k is the number of players in the game, which can be extended for singleton fair cost facility location games [17].

In this paper we show that the same is not true for fair cost facility location games where players can control multiple terminals, even when connection costs obey the triangle inequality. We prove that there are instances where the PoS is $\Theta(k)$, where k is the number of players, the same as the PoA.

Theorem 5.1 *There are instances of the metric FLG-FC which admit PNE but have price of stability in $\Theta(k)$, where k is the number of players.*

Proof.

Consider the instance shown in Figure 5. Let ε be a constant number much smaller than $\frac{1}{8}$. Each player P_i , for $i \in [1, k]$, controls terminal t'_i . Player P_a controls terminals t_2^1 to t_2^4 and t_5^1 to t_5^4 , player P_b controls terminals t_1 and t_3 , while player P_c controls terminals t_4^1, t_4^2, t_6^1 and t_6^2 . The full black lines have connection cost one, while the dashed lines have connection cost ε , with the exception of $d_{t'_b, f'_b}$ with cost $3\varepsilon^2 + \varepsilon$.




Fig. 5. An instance of metric FLG-FC with PoS in $\Theta(k)$. Dashed lines have connection cost ε (except for d_{t_1, f_1} , which has cost $3\varepsilon^2 + \varepsilon$), full black lines have cost 1. Labels next to terminals indicate which player controls the terminal.

Notice that the subgame composed only by the players P_2 to P_k and their connections, on the left side of Figure 5, has two pure equilibria: either everyone connects to the central facility f_s or each player P_i connects its terminal to facility f'_i . The

subgame composed by players P_a , P_b and P_c and their connections, on the other hand, is equal to the instance used in Theorem 4.2 when $w = 2$ with the costs scaled by ε , and therefore it has no equilibria unless the game is altered to allow some other player to stabilize it.

Player P_1 is the one in this instance that can stabilize these two subgames. Recall from Theorem 4.2 that, as long as some terminal is connected to f_2 , player P_c will connect to f_2 and f_5 . Since there are multiple terminals in f_5 , player P_a will want to connect to f_1 and f_5 as well. This then means that player P_b will want to connect to f_2 and f_5 . Therefore in any PNE, P_1 must connect to f_2 . For this to be the best possible move for P_1 , no player can open the central facility f_s , and therefore the only possible PNE is the one where each P_i connects to f'_i , while P_a , P_b and P_c play as seen above. Note that players P_a , P_b and P_c can be merged with any players P_i, P_j and P_r , for all $i \neq j \neq r \in [2, k]$, since they have disjoint strategy sets. Therefore, we can assume that there are k players in this instance.

The cost of the optimal strategy is the following. In the left side, players P_1 to P_k connect to the central facility f_s , for a cost of $1 + k\varepsilon$. In the right side, players P_b and P_c connect to f_2 and f_5 , while player P_a connects to f_1 and f_5 , paying also a connection cost of ε for each terminal, for a total cost of

$$1 + k\varepsilon + \varepsilon \left(\frac{8}{7} - \varepsilon + 1 + \frac{8}{7} + \varepsilon + 14 \right) = 1 + k\varepsilon + \frac{16}{7}\varepsilon + 15\varepsilon = 1 + \varepsilon \left(k + 15 + \frac{16}{7} \right).$$

In the unique PNE, players P_2 to P_k connect to their single facilities f'_2 to f'_k and player P_1 connects to f_2 , while players P_a , P_b and P_c connect to facilities f_2 , f_5 and f_1 (with P_b using connection d_{t_1, f_1} which costs $3\varepsilon^2 + \varepsilon$). Thus, the total cost for the PNE is

$$\begin{aligned} k + (k - 1)\varepsilon + \varepsilon \left(\frac{8}{7} - \varepsilon + 1 + \frac{8}{7} + \varepsilon + 3\varepsilon + 14 \right) &= k + k\varepsilon - \varepsilon + \frac{16}{7}\varepsilon + 15\varepsilon + 3\varepsilon^2 \\ &= k + \varepsilon \left(k + 14 + \frac{16}{7} + 3\varepsilon \right). \end{aligned}$$

The PoS of this instance is therefore

$$\frac{k + \varepsilon(k + 14 + \frac{16}{7} + 3\varepsilon)}{1 + \varepsilon(k + 15 + \frac{16}{7})} = \Theta(k).$$

Since the PoA of FLG-FC is $\Theta(k)$ (and therefore the PoS is $O(k)$), this instance makes the bound for PoS for FLG-FC asymptotically tight, and thus the PoS for metric FLG-FC is $\Theta(k)$. \square

5.1 Price of Stability in Capacitated Facility Location Games

Until now we have only considered scenarios where there are no limits on the number of terminals that a facility can supply. However this is not always the case in practice. In [16], Rodrigues and Xavier show that even metric facility location games have unbounded PoA, unless sequentiality is considered. Furthermore, they

show that metric singleton capacitated fair cost facility location games have a PoS of at least $\Theta(\log k)$, in comparison to the constant bounds of the uncapacitated metric version [9].

In this section we prove that for capacitated games where players are allowed to control more than a single terminal, there are instances with unbounded PoS. Note that, contrary to most results in this paper, we consider for this result that connection costs do not obey the triangle inequality, and thus any connections not shown in our construction have prohibitively large (and unbounded) cost \mathcal{U}_d .

Theorem 5.2 *There are 3-player instances of the capacitated FLG-FC that admit a PNE but have PoS that is unbounded.*

Proof.

Assume that we are restricted to instances where there is a PNE and players are allowed to control multiple terminals. We combine the instance described in Theorem 4.2 into an instance with unbounded PoA [16] to force that the only possible PNE in the game opens facilities with unbounded opening costs.

Consider the game in Figure 6. Note that the subgraph induced by the terminals that players P_a, P_b, P_c control is the same as in Theorem 4.2 when parameter $w = 2$. Thus, the only way for them to be in a PNE is if terminal t_a connects to f_2 . For this to happen in an equilibrium, terminal t_b must not connect to f_a . If this was the case, then it would be cheaper for t_a to also connect to f_a , as it would pay only $1/4$ which is less than the amount it pays to connect to f_2 . Since f_2 must be occupied by t_a in a PNE, terminal t_c in any equilibrium will connect to f_c , paying the connection cost of \mathcal{U} which can be arbitrarily large, and thus the only possible equilibrium in this game is unbounded. Consequently, the PoS is also unbounded.

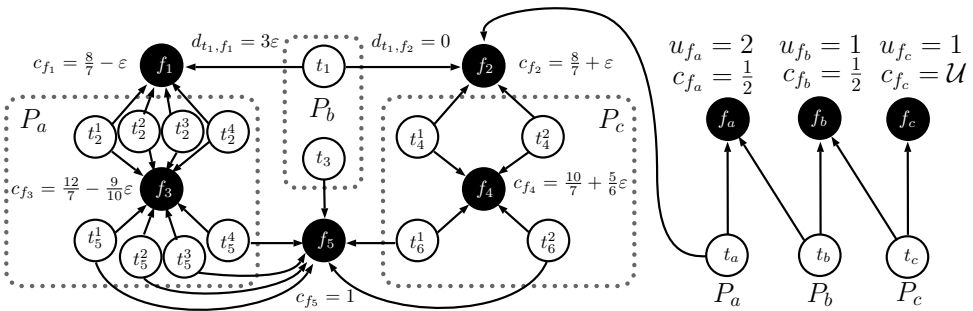


Fig. 6. An instance of capacitated FLG-FC with unbounded PoS. All connections have zero cost (except d_{t_1, f_1} , which has cost 3ϵ). Any facility f without a capacity restriction indicated is unlimited, i.e., $u_f \geq 8$. Labels next to terminals indicate which player controls the terminal.

□

6 Conclusions and Future work

In this paper we have studied a general version of the fair cost facility location game. In this model, the opened facilities have their opening costs shared evenly among

all terminals that connect to them, and players are allowed to control multiple terminals. We prove that there are instances for these games with no pure Nash equilibrium, and that deciding whether an instance of the game has a PNE or not is NP-hard, even when connection costs are metric. Furthermore, we provide a negative partial answer to an open question whether weighted fair cost facility location games always possess PNE, by showing the connections of these weighted games to scenarios where players control multiple terminals. Finally, we prove negative efficiency bounds related to the price of stability in both uncapacitated and capacitated games, showing that the price of stability is as inefficient as the price of anarchy in this generalized version of fair cost facility location game.

There are several possible directions on future work that relates to our paper. For weighted games, it remains an open problem to find whether there is an instance with no PNE for the singleton case. While our instance with no PNE has only a few terminals, the effect that cooperation between terminals can have on the PNE is crucial for there to be no PNE in this instance. It is not clear whether there is a possible way to adapt this instance for the singleton case. When we consider strong equilibria [6] for the singleton case, where terminals might cooperate to choose a better strategy when it is beneficial to every player in a coalition, a few similarities appear with this general version of fair cost facility location game. It might be possible to extend some results from strong equilibria for singleton fair cost facility location games to our setting, such as the existence of ϵ -approximate strong equilibria in singleton fair cost facility location games [10]. Alternatively, it might as well be possible to adapt some of our results to strong PNE in more restricted settings.

In [2], Cardinal and Hoefer prove a two parameter (α, β) constant approximation to facility location games with arbitrary cost sharing rules, where each player can reduce its cost by unilateral deviation by at most a factor of α , while the social cost of this α -approximated PNE is at most β times from the optimal social cost. An interesting question is whether it is possible to find a similar two parameter approximation for our setting, or if the limited fair cost aspect of our scenario causes one of the parameters to be too large.

Finally, in many practical scenarios there is a central authority with some capacity for interference in the game which has a goal of minimizing the social cost globally. Thus it is important to model in what circumstances this intervention can improve efficiency. The use of stackelberg strategies [7], where a central authority control a few terminals that can play before the other players, might guarantee that a PNE is always possible to find, as well as using tolls [4] in either connections or facilities to ensure only “good” equilibria are chosen by players.

Acknowledgment

The authors thank the ideas and discussions held in the 1st Paulista Workshop on Optimization, Combinatorics and Algorithms (WoPOCA 2017), which contributed for some of the results in this paper. Furthermore, the authors thank the finan-

cial support from CNPq (147141/2016-8, 304856/2017-7, 425340/2016-3), CAPES (PROEX/1406880 and PDSE 99999.006380/2015-08) and Fapesp (2015/11937-9, 2016/23552-7).

References

- [1] Anshelevich, E., A. Dasgupta, J. Kleinberg, E. Tardos, T. Wexler and T. Roughgarden, *The Price of Stability for Network Design with Fair Cost Allocation*, in: *45th Annual IEEE Symposium on Foundations of Computer Science*, 2004, pp. 295–304.
- [2] Cardinal, J. and M. Hoefer, *Non-cooperative facility location and covering games*, *Theoretical Computer Science* **411** (2010), pp. 1855–1876.
- [3] Chen, H. and T. Roughgarden, *Network design with weighted players*, *Theory Comput. Syst.* **45** (2009), pp. 302–324.
- [4] Cole, R., Y. Dodis and T. Roughgarden, *Pricing network edges for heterogeneous selfish users*, in: *Proceedings of the Thirty-fifth Annual ACM Symposium on Theory of Computing*, STOC '03 (2003), pp. 521–530.
- [5] Devanur, N. R., M. Mihail and V. V. Vazirani, *Strategyproof cost-sharing mechanisms for set cover and facility location games*, *Proceedings of the 4th ACM conference on Electronic commerce - EC '03* (2003), pp. 108–114.
- [6] Epstein, A., M. Feldman and Y. Mansour, *Efficient graph topologies in network routing games*, *Games and Economic Behavior* **66** (2009), pp. 115 – 125.
- [7] Fanelli, A., M. Flammini and L. Moscardelli, *Stackelberg strategies for network design games*, in: *Proceedings of the 6th International Conference on Internet and Network Economics*, WINE 2010 (2010), pp. 222–233.
- [8] Goemans, M. X. and M. Skutella, *Cooperative facility location games*, *Journal of Algorithms* **50** (2004), pp. 194–214.
- [9] Hansen, T. D. and O. Telelis, *On pure and (approximate) strong equilibria of facility location games*, in: *Internet and Network Economics, 4th International Workshop, WINE 2008, Shanghai, China, December, 2008. Proceedings*, 2008, pp. 490–497.
- [10] Hansen, T. D. and O. Telelis, *Improved bounds for facility location games with fair cost allocation*, in: *Combinatorial Optimization and Applications, Third International Conference, COCOA 2009, Huangshan, China, June 10-12, 2009. Proceedings*, 2009, pp. 174–185.
- [11] Hoefer, M., *Non-cooperative tree creation*, *Algorithmica* **53** (2007), pp. 104–131.
- [12] Kononov, A. V., Y. A. Kochetov and A. V. Plyasunov, *Competitive facility location models*, *Computational Mathematics and Mathematical Physics* **49** (2009), pp. 994–1009.
- [13] Leonardi, S. and G. Schäfer, *Cross-monotonic cost sharing methods for connected facility location games*, *Theoretical Computer Science* **326** (2004), pp. 431–442.
- [14] Mallozzi, L., *Noncooperative facility location games*, *Operations Research Letters* **35** (2007), pp. 151–154.
- [15] Pal, M. and E. Tardos, *Group strategy proof mechanisms via primal-dual algorithms*, in: *Foundations of Computer Science, 2003. Proceedings. 44th Annual IEEE Symposium on*, 2003, pp. 584–593.
- [16] Rodrigues, F. C. and E. C. Xavier, *Non-cooperative capacitated facility location games*, *Information Processing Letters* **117** (2017), pp. 45 – 53.
- [17] Rodrigues, F. C. and E. C. Xavier, *Non-cooperative facility location games: a survey*, *Revista de Informática Teórica e Aplicada* **24** (2017).
- [18] Vetta, A., *Nash equilibria in competitive societies, with applications to facility location, traffic routing and auctions*, *The 43rd Annual IEEE Symposium on Foundations of Computer Science 2002 Proceedings* **43** (2002), pp. 416–425.