On the Effectiveness of Connection Tolls in Fair Cost Facility Location Games

Félix Carvalho Rodrigues[1], Guido Schäfer[2,3], and Eduardo Candido Xavier[1]

[1] University of Campinas (UNICAMP), Campinas, Brazil
{felix.rodrigues|eduardo}@ic.unicamp.br
[2] Centrum Wiskunde & Informatica (CWI), Amsterdam, The Netherlands
[3] Vrije Universiteit Amsterdam, Amsterdam, The Netherlands
g.schaefer@cwi.nl

Abstract. We investigate the effectiveness of tolls to reduce the inefficiency of Nash equilibria in the classical *fair cost facility location game*. In this game, every terminal corresponds to a selfish player who wants to connect to some facility at minimum cost. The cost of a player is determined by the connection cost to the chosen facility plus an equal share of its opening cost. We are interested in the problem of imposing tolls on the connections to induce a socially optimal Nash equilibrium such that the total amount of tolls is minimized. It turns out that this problem is challenging to solve even for simple special cases. We provide polynomial-time algorithms for (i) instances with two facilities, and (ii) instances with a constant number of facilities arranged as a star. Our algorithm for (ii) exploits a relation between our tolling problem and a novel *bipartite matching problem without crossings*, which we prove to be NP-hard.

1 Introduction

Facility location problems are one of the fundamental classes of problems in computer science, with practical applications in many fields of industry and services. A common version of the facility location problem can be stated as follows: We are given a set F of facilities that can be opened and a set T of terminals (or clients) that need to be connected to the facilities. Each facility $f \in F$ has a non-negative opening cost c_f that is incurred if it is opened. Further, the cost of connecting terminal $t \in T$ to facility $f \in F$ is given by a non-negative connection cost d_{tf}. The goal is to choose a subset $F' \subseteq F$ of the facilities which are opened and to connect each terminal $t \in T$ to an open facility in F'. The objective is to find a solution that minimizes the total opening cost of all facilities in F' and the connection costs of the terminals to their respective facilities.

The facility location problem has been studied extensively in the literature. However, in most studies a centralized optimization perspective is adopted, i.e., it is assumed that there is a central authority, controlling all facilities and terminals, whose goal is to determine a solution of minimum cost. This assumption is

not justified in settings where several agents are involved who want to minimize the costs of their own facilities or terminals.

In this paper, we study a game-theoretic variant of the facility location problem, which is also known as the *Fair Cost Facility Location Game (FCFLG)*. Here each terminal $t \in T$ corresponds to an independent player (or agent) who wants to connect to some facility. Each player $t \in T$ selfishly chooses a facility in F to which his terminal is assigned to. The opening cost of a facility is shared equally between the players that have chosen it and the connection costs are paid individually by the players. Each player attempts to minimize his individual cost. The social cost objective that we consider throughout this paper is the sum of the individual player costs.

This game is known to suffer from a high inefficiency. In particular, it is not hard to construct instances that show that the social cost ratio between a pure Nash equilibrium and a social optimum can be as large as the number of players.

For example, consider the instance with two facilities f_1 and f_2 and an even number of $n \geq 4$ terminals as depicted in Figure 1. In the social optimum, terminals t_1 to $t_{n/2}$ connect to f_1 and terminals $t_{n/2+1}$ to t_n connect to f_2, yielding a social cost of 4. Suppose all terminals connect to facility f_1. Then the first $\frac{n}{2}$ terminals pay $\frac{2}{n}$ and the second $\frac{n}{2}$ terminals pay $1+\frac{2}{n}$ each. This is a Nash equilibrium because every player who deviates to facility f_2 needs to pay at least the opening cost of $2 > 1+\frac{2}{n}$. Note that this equilibrium is highly inefficient: its social cost is $2 + \frac{n}{2}$, which is $\Omega(n)$ times larger than the optimal social cost.


Fig. 1.

In light of this, it is imperative to seek efficient means to deal with this inefficiency. The idea of designing efficient algorithms, also known as *coordination mechanisms*, to reduce the inefficiency caused by selfish behavior has recently attracted a lot of attention in the algorithmic game theory literature. For example, in the context of network routing games the use of tolling schemes was shown to be an effective way to steer selfish players into more favorable equilibrium outcomes. However, relatively little work has been done considering facility location games.

Suppose that in our facility location game there is a central authority that, while not being able to control the terminals directly, is capable to increase the perceived costs of the players through some form of external costs, such as tolls. This authority is interested to induce Nash equilibria which have optimal social cost, while altering the game as little as possible. Immediately, two possibilities for the placement of tolls come to one's mind: either on the opening costs of the facilities or on the connection costs between terminals and facilities.

When tolling facilities, it quickly becomes clear that there are instances where no tolling scheme can avoid highly suboptimal Nash equilibria, even if the con-

nection costs constitute a metric. To see this, reconsider the example given above. Suppose we want to avoid inefficient equilibria by imposing non-negative tolls τ_1 and τ_2 on facilities f_1 and f_2, respectively. Assume that $\tau_1 \leq \tau_2$. Then all terminals connecting to facility f_1 still constitutes a Nash equilibrium. To see this, note that the first $\frac{n}{2}$ terminals pay $\frac{2+\tau_1}{n}$ and the second $\frac{n}{2}$ terminals pay $1 + \frac{2+\tau_1}{n}$. If a player deviates to facility f_2 he needs to pay at least the opening cost of $2+\tau_2 > 1+\frac{2+\tau_1}{n}$. If $\tau_1 > \tau_2$ then by using symmetric arguments it follows that all terminals connecting to facility f_2 is a Nash equilibrium. In either case a pure Nash equilibrium remains whose social cost is at least $\Omega(n)$ times larger than the optimal social cost. We conclude that for this instance there is no way to avoid Nash equilibria of high inefficiency merely by tolling the facilities.

Given the above observations, we focus on tolling the connections in this paper. Clearly, it is possible to enforce an arbitrary optimal solution as a Nash equilibrium simply by increasing the costs of all connections which are not part of the solution to infinity. However, an intriguing question that arises is: How large would the tolls need to be in the worst case to ensure that an optimal solution is realized as a Nash equilibrium? And: Can we efficiently compute minimum cost tolls that induce a social optimum as a Nash equilibrium?

Our Contributions. In this paper, we study a model for tolling the connection costs of fair cost facility location games with the objective to steer the players to some desirable strategy profile (e.g., social optimum). We assume that the players start from an arbitrarily given strategy profile and play best response moves sequentially, one player at a time, according to some order (see below for further justification of this assumption).

The main contributions presented in this paper are as follows:

1. We show that if a predefined player order is given, finding optimal tolls which induce a given strategy profile can be solved in polynomial-time (Section 2).

The problem becomes inherently more difficult if the order of the players is not fixed, but can be determined by the central authority. Even for simple special cases, this problem turns out to be very challenging to solve.

2. We identify some properties that optimal tolling schemes have to satisfy for certain restricted types of instances (Section 2).
3. We exploit these properties to derive polynomial-time algorithms for the following two special cases (Sections 3 and 4, respectively):
    (i) instances with two facilities, and
    (ii) instances with a constant number of facilities arranged as a star.

Our algorithm for (ii) is based on a reduction of our tolling problem to a *bipartite matching problem without crossings*: Given an edge-weighted bipartite graph whose nodes on one side are partitioned into consecutive clusters, find a minimum weight perfect matching such that no two edges incident to the same cluster cross. To the best of our knowledge, this problem has not been studied before and is of independent interest.

4. We provide a dynamic programming algorithm for the bipartite matching problem without crossings, which is polynomial if the number of clusters is constant. Further, we prove that this matching problem is NP-hard in general (Section 4).

We conjecture that the tolling problem is NP-hard in general. While we do not have a proof for this yet, we feel that our NP-hardness result for the related bipartite matching problem without crossings lends some support to this.

Related Work. Different versions of facility location games have been studied in the literature. Cardinal and Hoefer [4] consider a non-cooperative facility location game where n players control all terminals, and players do not have rules on how to share the opening costs. They show that in general pure Nash equilibria may not exist. When restricting to instances that admit a pure Nash equilibrium, both the price of anarchy and the price of stability are $\Theta(n)$. The variant of the game with fair cost sharing rules can be reduced to the network design game with fair cost allocation introduced by Anshelevich et al. [1]. It can be shown that the price of anarchy is n, while the price of stability is H_n.

For the metric facility location game, Hansen and Telelis [10] show that constant bounds are possible for the price of stability and the strong price of anarchy. For capacitated facility location games, Rodrigues and Xavier [12] show that the price of anarchy is unbounded even for metric variants, while the price of anarchy becomes bounded when considering a sequential version of the game.

The use of tolling schemes was intensively studied in the context of network routing games. Beckman, McGuire and Winsten [2] prove that marginal cost tolls induce an optimal Nash flow in selfish routing games with non-atomic, homogeneous players. Several works extended this result, for example, to heterogeneous players (see, e.g., [6]) and multi-commodity networks (see, e.g., [9]).

To the best of our knowledge, we are the first to investigate the effectiveness of tolls in facility location games. On the other hand, for the more general network design game with fair cost allocation, there are several works that focus on improving the price of anarchy. Fanelli et al. [8] derive efficient Stackelberg strategies to improve equilibria. Chen et al. [5] study optimal cost sharing protocols for different variants of network design games, such as directed and undirected networks. For cost sharing games in a set cover setting, Buchbinder et al. [3] use a taxation model which offers subsidies to certain sets in order to improve equilibria when using best response dynamics.

Regarding the matching problem we describe in Section 4, the most relevant work is due to Darmann et al. [7], where they prove NP-hardness for a similar maximum matching problem under disjunctive constraints, where each pair of edges has a constraint saying whether they can exist in the same solution or not.

2   Preliminaries

We first formally define the *Fair Cost Facility Location Game (FCFLG)* that we consider in this paper: Let $G = (T \cup F, T \times F)$ be a complete bipartite graph,

where F is the set of facilities and T is the set of terminals. We use $m = |F|$ and $n = |T|$ to refer to the number of facilities and terminals, respectively. Each facility $f \in F$ has a non-negative opening cost c_f and each terminal-facility pair $(t, f) \in T \times F$ has a non-negative connection cost d_{tf}. Each terminal t is a selfish player who chooses to connect to a facility such that t is connected to exactly one opened facility. We use the terms player and terminal interchangeably.

Suppose that player $t \in T$ chooses facility $S_t \in F$. We use $S = (S_1, \ldots, S_n)$ to refer to a strategy profile of all players. We write $f \in S$ to denote that facility f is opened in strategy profile S. Each player $t \in T$ wants to minimize his own payment (or cost) which is defined as $p_t(S) = c_{f_t}/x_{f_t}(S) + d_{tf_t}$, where $f_t = S_t$ is the facility he chose and $x_{f_t}(S) = |\{i \in T : S_i = f_t\}|$ is the number of players using facility f_t in strategy profile S.

Given a strategy profile $S = (S_1, \ldots, S_n)$, a change for player i from a strategy S_i to a different strategy S_i' is called a *move*. Let $S_{-i} = (S_1, \ldots, S_{i-1}, S_{i+1}, \ldots, S_n)$ be the strategy profile resulting from S if we remove i's strategy. We say that S_i is a *best response* of player i with respect to S if $p_i(S_i, S_{-i}) \le p_i(S_i', S_{-i})$ for all $S_i' \in F$. A strategy profile S is a *pure Nash equilibrium* (PNE) if for every player $i \in T$, S_i is a best response with respect to S. The *social cost* is a measure of the overall quality of a particular strategy profile. For facility location games, the social cost for a strategy profile S is defined as the sum of all payments of the players, i.e. $C(S) = \sum_{t \in T} p_t(S) = \sum_{f \in S} c_f + \sum_{t \in T} d_{tf_t}$. In order to analyze the inefficiency of equilibria the *Price of Anarchy* (PoA) [11] and the *Price of Stability* (PoS) [1] are the standard measures used in the literature.

We next define the *modified toll game* that we consider. Let $\mathcal{G} = (G, c, d)$ be an instance of FCFLG. We assume that a central authority can alter $\mathcal{G}$ by placing some tolls $\tau : T \times F \to \mathbb{R}_{\ge 0}$ on every connection. We assume that the players perceive a modified cost function which includes the tolls, while these tolls are excluded from the social cost objective (i.e., tolls are refundable). More formally, define the *modified toll game* $\bar{\mathcal{G}} = \bar{\mathcal{G}}(\tau) = (G, c, d + \tau)$ with respect to tolls τ, where every player i perceives a cost of $\bar{p}_i(S) = p_i(S) + \tau_{if_i}$ with $f_i = S_i$ being the facility that i chooses under S. The social cost $C(S)$ of S in $\bar{\mathcal{G}}$ is the same as the social cost of S in $\mathcal{G}$.

We say that a strategy profile S is *inducible* if there exist tolls τ such that S is a pure Nash equilibrium in the modified toll game $\bar{\mathcal{G}}$. Given a desirable strategy profile S^*, the central authority ideally would like to impose tolls τ such that S^* is inducible. However, the problem with this is that even though the tolls τ imposed on the connections might be sufficient to impose S^* as a Nash equilibrium, S^* might not be reachable because there are multiple Nash equilibria in the modified toll game $\bar{\mathcal{G}}(\tau)$. In fact, even if we start from a fixed strategy profile S there is no guarantee that S^* is reached if the players simply play best response. As a result, we may end up in a PNE whose social cost is significantly higher than the social cost $C(S^*)$ of the desired outcome S^*. On the other hand, enforcing that S^* is the unique PNE is also not desirable, since the amount of tolls needed for this can be very large.

In order to circumvent these problems, conceptually we adopt the following viewpoint: In a first phase, all players arbitrarily play best responses in the (unmodified) facility location game until they reach a pure Nash equilibrium, say S. In a second phase, we then impose tolls τ on the connections such that S^* is reached from S if we let the players play one additional round of best responses according to some order γ. Our goal is to determine tolls τ and an ordering γ such that the total amount of tolls is minimized.

We formalize the above idea. Let S be a pure Nash equilibrium and S^* be an arbitrary strategy profile of $\mathcal{G}$. Let $\gamma : T \to [1, \ldots, n]$ be an order according to which the players play best response in the game $\bar{\mathcal{G}}$, where the interpretation is that player i is the $\gamma(i)$-th player to move. If, starting at S and playing according to the order γ, there is a best response for every player such that S^* is reachable, then we say that S^* *is reachable from S through γ*. The *Minimum Toll Problem (MTP)* considered in this paper is defined as follows:

Minimum Toll Problem (MTP):
Given: FCFLG instance $\mathcal{G} = (G, c, d)$, pure Nash equilibrium S, arbitrary strategy profile S^*
Goal: Determine tolls $\tau : T \times F \to \mathbb{R}_{\geq 0}$ and an ordering $\gamma : T \to [1, \ldots, n]$ such that S^* is reachable from S through γ and the total amount of tolls $\mathcal{T} = \sum_{(t,f) \in T \times F} \tau_{tf}$ is minimized.

The theorem below shows that if the order of the players is fixed, then determining the optimal tolls is easy. Due to lack of space, several proofs are omitted from this extended abstract and will be given in the full version of the paper.

Theorem 1. *Given an order γ, a starting strategy profile S and a final strategy profile S^*, there is a polynomial-time algorithm that finds the minimum tolls such that S^* is reachable from S through γ.*

Below we establish some useful properties for instances of MTP where the set of facilities used under S and S^* are disjoint.

We say that two terminals $t, t' \in T$ are *similar* if (i) $S_t = S_{t'}$, (ii) $S_t^* = S_{t'}^*$, and (iii) t and t' do not have any other possible connections other than to S_t and S_t^*. Further, we say that $A \subseteq T$ is a *similar set* if for all terminals $t, t' \in A$, t and t' are similar.

Lemma 1 (Monotonicity). *Let T be partitioned into similar sets $T_1, \ldots, T_p$. Let $\tau_t(x, y)$ be the toll that is needed to move terminal $t \in T_j$, when it is the x-th terminal to move to facility S_t^* and the y-th terminal to move from facility S_t. Then, $\tau_t(x, y)$ is monotonically decreasing in x and y.*

Using this, we can infer an optimal terminal order for each similar set. For a terminal $t \in T$, let $\Delta_d(t) = d_{tf^*} - d_{tf}$ be the difference in connection costs between $S_t = f$ and $S_t^* = f^*$.

Lemma 2 (Sorting similar sets). *Let T be partitioned into similar sets $T_1, \ldots, T_p$. Then there is an ordering γ that induces minimum toll costs, where for every two terminals $t, t' \in T_j$: if $\Delta_d(t) < \Delta_d(t')$, then $\gamma(t) < \gamma(t')$.*

3   MTP with two facilities

We derive an algorithm for the special case of MTP with two facilities only.

Theorem 2 (2-MTP). *When restricted to two facilities, there is a polynomial-time algorithm to solve MTP.*

Proof. Let $F = \{f_1, f_2\}$. Consider a terminal t and suppose $S_t^* = f_1$ (the case $S_t^* = f_2$ follows similarly). Because t should not have an incentive to deviate under S^*, it must hold that $\tau_{tf_2} + d_{tf_2} + \frac{c_{f_2}}{x_{f_2}(S^*)+1} \geq d_{tf_1} + \frac{c_{f_1}}{x_{f_1}(S^*)}$. This imposes a restriction on the toll τ_{tf_2} which must be satisfied in any feasible solution and we can (implicitly) add this minimum toll to the connection cost d_{tf_2}. We can thus assume that S^* is a PNE with respect to d.

Let A and B be the sets of all terminals connected to f_1 and f_2 in S, respectively, and define $a = |A|$ and $b = |B|$. Note that all terminals $t \in A$ with $S_t = S_t^* = f_1$ and $t \in B$ with $S_t = S_t^* = f_2$ do not require any tolls on their connections because S^* is a PNE. We can thus let them be the last terminals in the order γ.

Let A_i and B_i be the sets of terminals that are connected to f_1 and f_2, respectively, after the i-th terminal has moved to its facility in S^*. Let $a_i = |A_i|$ and $b_i = |B_i|$. In particular, a_n (resp. b_n) is the number of terminals in A (in B) at the final strategy S^*. Note also that $a_j + b_j = a_i + b_i$, for any $i, j \in [1, n]$. Among the terminals in A (resp. B) we denote by A' (resp. B') the terminals that have to move to the other facility, i.e, each $t_a \in A'$ (resp. $t_b \in B'$) is such that $S_{t_a} = f_1$ and $S_{t_a}^* = f_2$ (resp. $S_{t_b} = f_2$ and $S_{t_b}^* = f_1$).

Suppose we are considering the first terminal to move and suppose that $a_1 > a_n$ (then $b_1 < b_n$). Since $a_1 > a_n$ and $b_1 < b_n$, moving any terminal $t \in B'$ from f_2 to $S_t^* = f_1$ does not require tolls currently since at turn 1 we have $\frac{c_{f_1}}{a_1+1} + d_{tf_1} \leq \frac{c_{f_1}}{a_n} + d_{tf_1}$ and $\frac{c_{f_2}}{b_n} + d_{tf_2} < \frac{c_{f_2}}{b_1} + d_{tf_2}$. Since S^* is a PNE we must have $\frac{c_{f_1}}{a_n} + d_{tf_1} \leq \frac{c_{f_2}}{b_n+1} + d_{tf_2}$, which results in

$$\frac{c_{f_1}}{a_1 + 1} + d_{tf_1} \leq \frac{c_{f_1}}{a_n} + d_{tf_1} \leq \frac{c_{f_2}}{b_n + 1} + d_{tf_2} < \frac{c_{f_2}}{b_1} + d_{tf_2}. \qquad (1)$$

So at turn 1 terminal $t \in B'$ has an incentive to move from f_2 to f_1 and no tolls are required. Suppose we move terminals $t \in A'$ from f_1 to f_2 until a turn j such that $a_j = a_n + 1 > a_n$ and $b_j = b_n - 1 < b_n$. It is not hard to see that for all these turns equation (1) remains valid and for any terminal $t \in B'$ no tolls would be required to move it to f_1.

The algorithm constructs an order where first we move terminals $t \in A'$ from f_1 to $S_t^* = f_2$ until a time j where we have $a_j = a_n$ and $b_j = a_n$. All these moves require positive tolls, but after we reach the point where $a_j = a_n$ and $b_j = a_n$, we will show that no tolls are required to move the remaining terminals $t \in A'$ or $t \in B'$. In particular, for the terminals in B' no tolls are required. The optimal solution necessarily moves first terminals from A' until $a_j = a_n$ and $b_j = a_n$. To see this, note that moving terminals from B' first would only increase the total

cost of moving terminals of A' later, since the terminals in B' have always toll cost equal to zero.

Until a turn j where $a_j = a_n$ and $b_j = a_n$, only terminals from A' move from f_1 to f_2 and it is not difficult to prove a result similar to Lemma 2 showing that the optimal order among terminals in A' is to move them in decreasing order of $\Delta_d(t)$.

Now suppose we are at turn j where $a_j = a_n$ and $b_j = a_n$. At this point we can move a terminal $t \in B'$ from f_2 to f_1 with tolls cost equal to zero. To see this, note that $\frac{c_{f_2}}{b_j} + d_{tf_2} = \frac{c_{f_2}}{b_n} + d_{tf_2} > \frac{c_{f_2}}{b_n+1} + d_{tf_2} \geq \frac{c_{f_1}}{a_n} + d_{tf_1} > \frac{c_{f_1}}{a_j+1} + d_{tf_1}$, where the third inequality ($\geq$) is valid since S^* is a PNE. So at this point t has an incentive to move from f_2 to f_1 and no tolls are required.

In the next turn $j + 1$ we have $a_{j+1} = a_n + 1$ and $b_{j+1} = b_n - 1$, and to move a terminal $t \in A'$ no tolls are required either, since S^* is a PNE and $\frac{c_{f_2}}{b_{j+1}+1} + d_{tf_2} = \frac{c_{f_2}}{b_n} + d_{tf_2} \leq \frac{c_{f_1}}{a_n+1} + d_{tf_1} = \frac{c_{f_1}}{b_{j+1}} + d_{tf_1}$. So the amount t is paying for being connected to f_1 is greater than or equal to the amount paid to be connected to f_2. So the order follows a move from a terminal in B' and a terminal in A' until all terminals have moved. The cases where $b_1 > b_n$ or $b_1 = b_n$ are similar to the case $a_1 > a_n$ discussed above. □
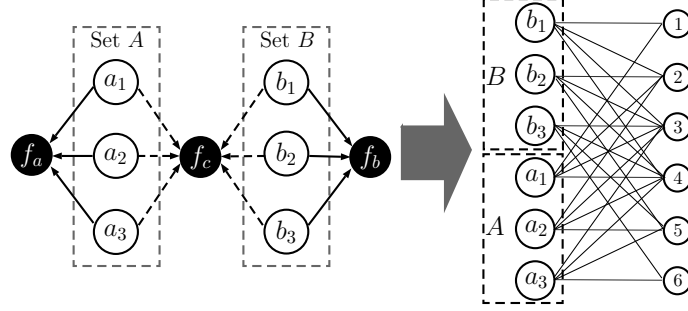
4   Star-MTP

We consider a special case of MTP which we term the *Star Minimum Toll Problem (Star-MTP)*: In an instance $(\mathcal{G}, S, S^*)$ of Star-MTP all terminals have the same starting facility f_c in strategy profile S and can be partitioned into m similar sets $T_1, \ldots, T_m$ such that every terminal $t \in T_i$ has target facility $S_t^* = f_i$. Furthermore, no terminal in T_i can connect to any facility other than f_c and f_i. We show that Star-MTP admits a polynomial-time algorithm for a constant number of facilities. To this aim, we first reduce the problem to a new matching problem and then present a dynamic programming algorithm for it.

Reduction to Bipartite Matching Without Crossing Edges. We can think of Star-MTP as a bipartite matching problem, where the set of terminals corresponds to the set of nodes on one side of the bipartition and the other side contains the integers $1, 2, \ldots, n$ (n being the number of terminals). These integers represent the order in which each terminal moves from f_c to its final facility. For terminals belonging to the same set T_i, we know from Lemma 2 that they must move according to the order of non-decreasing differences in connection cost between f_c and f_i, denoted by Δ_d. The partition containing the terminals is organized in such a way that we list the vertices from top to bottom grouped by similar sets, and vertices in the same similar set are sorted from top to bottom in increasing order of Δ_d value. For each terminal t and integer $q \in [1, n]$, the edge (t, q) in this bipartite graph has cost equal to the tolls required when t is the q-th terminal to move to its final facility.

For terminals $t, t' \in T_i$ belonging to the same similar set, if $\Delta_d(t) < \Delta_d(t')$ then t must move before t'. We impose this restriction by requiring that in the

Fig. 2. Example of reduction from Star-MTP to PMC. For simplicity costs are omitted.

matching there are no crossing edges between vertices of the same similar set. The problem reduces to finding a perfect matching of minimum cost without crossings.

Definition 1 (Perfect Matching Without Crossing Edges (PMC)).
Given a bipartite graph $G = (A = \cup_{i \in [m]} A_i, B; E)$, let A be a set of n integers from t_1 to t_n, partitioned into subsets $A_i, i = 1, \ldots, m$, and let B be another set of integers from 1 to n. For each pair $t_r \in A$, $j \in B$ there is an edge $e = (t_r, j) \in E$ with cost $w_{t_r j} \in \mathbb{R}$. For vertices $t_r, t_s \in A_i$ where $t_r < t_s$, edges (t_r, q) and (t_s, p) are crossing *if $p < q$. The problem is to find a minimum cost perfect matching without crossing edges.*

Now we present a reduction from Star-MTP to the PMC. Given an instance of the Star-MTP, first sort terminals in each similar set T_i in decreasing order of Δ_d value: if $\Delta_d(t) < \Delta_d(t')$ then $t < t'$, for any $t, t' \in T_i$. Each similar set T_i becomes a set A_i in the PMC instance, and we assume that all $t \in A_i$ have a smaller value then $t' \in A_j$ if $i < j$, i.e, $t < t'$. So in the PMC instance the terminal vertices are sorted from top to bottom from A_1 to A_m, and inside each set A_i, terminals are sorted by Δ_d value. The partition B of the PMC instance just contains the numbers 1 to n, where n is the number of terminals.

Let $q_i = |A_i|$ for $i = 1, \ldots, m$ and let $A_i = \{t_1^i, \ldots, t_{q_i}^i\}$, with terminals sorted from t_1 to t_{q_i} in the order they must move considering just terminals from A_i. For each $t_x^i \in A_i$ and $y \in [x, n]$ we create an edge (t_x^i, y) with cost

$$w_{t_x^i, y} = \max \left(0, \frac{c_{f_i}}{x} + d_{t_x^i, f_i} - d_{t_x^i, f_c} - \frac{c_{f_c}}{n - y + 1} \right)$$

which is equivalent to the toll cost required if t_x^i is the y-th overall terminal to leave f_c, and the x-th to move to f_i among terminals in A_i. An example of the reduction is presented in Figure 2.

A perfect matching of minimum cost to the reduced instance corresponds to an optimal tolling for the MTP instance. To see this, note that no crossing edges are allowed in the matching, so for each T_i, $i = 1, \ldots, m$, terminals move according to the optimal order defined by Lemma 2. Since the costs of any edge

(t, y) represents the toll cost required when t is the y-th overall terminal to move, a minimum perfect matching corresponds to a minimum tolling.

Dynamic Program. We now present a dynamic program algorithm to the PMC problem. Let $G = (A = \cup_{i \in [m]} A_i, B; E)$ be an instance of PMC, with $q_i = |A_i|$ for $i = 1, \ldots, m$, and $n = q_1 + \cdots + q_m$. Let $DP(q_1, q_2, \ldots, q_m)$ be the cost of a minimum cost perfect matching without crossings for that instance. Define $DP(q_1, \ldots, q_m) = 0$ if $q_1 = \cdots = q_m = 0$ and

$$DP(q_1, \ldots, q_m) = \min_{\substack{i = 1, \ldots, m \\ \text{such that } q_i > 0}} \left(DP(q_1, \ldots, q_i - 1, \ldots, q_m) + w_{(t^i_{q_i}, y)} \right) \quad (2)$$

where y in $w_{(t^i_{q_i}, y)}$ is equal to $y = \sum_{i=1}^m q_i$.

Theorem 3. *The recurrence relation above correctly computes the optimal solution $DP(q_1, q_2, \ldots, q_m)$ for an instance of the PMC.*

It is not hard to construct a dynamic program algorithm to solve PMC, since the algorithm only needs to create an m-dimensional table of size $q_1 \times q_2 \times \ldots \times q_m = \Theta(n^m)$ and compute the value of each cell in $\Theta(m)$ time following the recurrence (2). The overall time of the algorithm is then $\Theta(mn^m)$ which is polynomial if m is a constant.

Hardness. We show that given an instance of PMC, it is NP-hard to decide whether it admits a perfect matching without crossings.

Theorem 4. *The problem of deciding whether a given instance of PMC admits a perfect matching without crossings is NP-hard.*

Proof. Let I be an instance of the 3-SAT with m clauses and n variables. We construct an instance of PMC as follows: for each variable x_i we build vertices $a_i^T < a_i^F$ in A and $b_i^T < b_i < b_i^F$ in B, with edges (a_i^T, b_i^T), (a_i^T, b_i), (a_i^F, b_i) and (a_i^F, b_i^F). For any occurrence of the literal x_i or $\overline{x_i}$ in a clause C_j, we add vertices $a_i^{C_j}$ to A and vertices $b_i^{C_j \overline{x_i}}, b_i^{C_j x_i}$ to B, with edges $(a_i^{C_j}, b_i^{C_j \overline{x_i}})$ and $(a_i^{C_j}, b_i^{C_j x_i})$, such that $a_i^T < a_i^{C_j} < a_i^F$ and $b_i^T < b_i^{C_j \overline{x_i}} < b_i < b_i^{C_j x_i} < b_i^F$. We call this gadget X_i. All vertices a from a gadget X_i form a subset A_i, so it is forbidden to have crossing edges in this gadget. Note that each vertex $a \in A_i$ can connect to exactly two vertices in B, where the one with smaller value is denoted by $S(a)$, and the one with greater value $G(a)$.

For each clause C_j, we also construct a vertex $c_j \in A$ which connects to each $b_i^{C_j l_i} \in B$ for each literal l_i occurring in C_j, i.e., if literal x_i (resp. $\overline{x_i}$) appears in C_j we include edge $(c_j, b_i^{C_j x_i})$ (resp. $(c_j, b_i^{C_j \overline{x_i}})$). Each vertex c_j belongs to its own partition $A_{c_j} = \{c_j\}$. See Fig. 3 for an example of this construction.

Finally, note that each variable x_i gives rise to two vertices in A (a_i^T, a_i^F) and three in B (b_i^T, b_i, b_i^F). Each clause C_j adds four vertices to A (one in the clause gadget (c_j) plus one for each literal ($a_i^{C_j}$)) and six in B (two for each

Fig. 3. Example of the reduction where $C_1 = (\overline{x_1} \vee x_2 \vee x_3)$, literal x_1 appears only on C_2 and $\overline{x_1}$ appears only on C_1.

literal $(b_i^{C_j \overline{x_i}}, b_i^{C_j x_i}))$, resulting in a total of $2n + 4m$ vertices in A and $3n + 6m$ in B. We add $n + 2m$ dummy vertices all belonging to the same partition A_k which can connect to any vertex in B except the vertices b_i, for $i \in [1, n]$, where $k = n + m + 1$ is the number of partitions.

Suppose there is a feasible assignment to the instance I of 3-SAT. We construct a perfect matching as follows: for each variable x_i which is true, we assign each vertex $a \in A_i$ to its vertex $S(a) \in B$. All b vertices of X_i larger than b_i are unassigned, and therefore for any clause C_j which contains x_i, its vertex c_j can be assigned to $b_i^{C_j x_i}$. Similarly the opposite is done if x_i is false, i.e., assign each vertex $a \in A_i$ to its vertex $G(a) \in B$, allowing each vertex c_j, from a clause C_j containing $\overline{x_i}$, to connect to $b_i^{C_j \overline{x_i}}$. With this, all vertices from set A which belong to variable and clause gadgets are matched. To complete the perfect matching, just assign the vertices from subset A_k, from smallest to greatest, in this order to the smallest to greatest vertices in B which are still not matched.

Now assume there is a perfect matching with no crossing edges for the graph G. First notice that for each gadget X_i either a_i^T or a_i^F is assigned to b_i. If a_i^T is assigned to b_i, then all a vertices of X_i are assigned to their greater vertices $G(a)$ since no crossing edges are allowed, and if a_i^F is assigned to b_i then all a vertices are assigned to their $S(a)$ vertices.

We construct an assignment for the 3-SAT instance by setting x_i to true if all $a \in A_i$ are assigned their smaller vertices $S(a) \in B$, and to false otherwise. Now we show that each clause C_j is satisfiable. Let $c_j \in A$ be the corresponding vertex of a clause C_j. Since we have a perfect matching, c_j must be connected to some $b_i^{C_j l_i}$ corresponding to one of its literals l_i, which is either x_i of $\overline{x_i}$. If $l_i = x_i$ then we know that all a vertices of X_i must be connected to their smaller vertices $S(a)$ and so x_i is true and C_j is satisfiable. Similarly, if $l_i = \overline{x_i}$ then all a vertices of X_i must be connected to their greater vertices $G(a)$, so x_i is false and C_j is satisfiable. □

5   Conclusions

The most natural open problem that our work suggests is to prove that the minimum toll problem is in fact NP-hard. While the NP-hardness result for

the perfect matching problem without crossings provides some support to this, its proof does not translate directly to the more restricted scenario of choosing optimal tolls. Besides this, there are also different possibilities of consideration for the tolling model, such as allowing simultaneous moves or enforcing that the unique possible equilibrium is one with optimal social cost. However for these scenarios, finding optimal tolls often includes finding possible equilibria, which implies that these tolling problems might be even harder than the ones we consider here. Finally, an interesting consideration is to allow for negative tolls on either the connection or opening costs to encourage players to play a specific strategy profile.

References

1. Anshelevich, E., Dasgupta, A., Kleinberg, J., Tardos, E., Wexler, T., Roughgarden, T.: The Price of Stability for Network Design with Fair Cost Allocation. In: 45th Annual IEEE Symposium on Foundations of Computer Science. pp. 295–304 (2004)
2. Beckmann, M., McGuire, C., Winsten, C.: Studies in the Economics of Transportation. Yale University Press (1956)
3. Buchbinder, N., Lewin-Eytan, L., (Seffi) Naor, J., Orda, A.: Non-cooperative cost sharing games via subsidies. Theory of Computing Systems **47**(1), 15–37 (2010)
4. Cardinal, J., Hoefer, M.: Non-cooperative facility location and covering games. Theoretical Computer Science **411**(16–18), 1855–1876 (2010)
5. Chen, H.L., Roughgarden, T., Valiant, G.: Designing networks with good equilibria. In: Proceedings of the Nineteenth Annual ACM-SIAM Symposium on Discrete Algorithms. pp. 854–863 (2008)
6. Cole, R., Dodis, Y., Roughgarden, T.: Pricing network edges for heterogeneous selfish users. In: Proceedings of the Thirty-fifth Annual ACM Symposium on Theory of Computing. pp. 521–530 (2003)
7. Darmann, A., Pferschy, U., Schauer, J., Woeginger, G.J.: Paths, trees and matchings under disjunctive constraints. Discrete Applied Mathematics **159**(16), 1726 – 1735 (2011)
8. Fanelli, A., Flammini, M., Moscardelli, L.: Stackelberg strategies for network design games. In: Proceedings of the 6th International Conference on Internet and Network Economics. pp. 222–233 (2010)
9. Fleischer, L., Jain, K., Mahdian, M.: Tolls for heterogeneous selfish users in multicommodity networks and generalized congestion games. In: 45th Annual IEEE Symposium on Foundations of Computer Science. pp. 277–285 (2004)
10. Hansen, T.D., Telelis, O.: On pure and (approximate) strong equilibria of facility location games. In: Proceedings of the 4th International Conference on Internet and Network Economics. pp. 490–497 (2008)
11. Koutsoupias, E., Papadimitriou, C.: Worst-case equilibria. In: Proceedings of the 16th Annual Conference on Theoretical Aspects of Computer Science. pp. 404–413 (1999)
12. Rodrigues, F.C., Xavier, E.C.: Non-cooperative capacitated facility location games. Information Processing Letters **117**, 45–53 (2017)