Non-Cooperative Multicast and Facility Location Games

[Extended Abstract]

Chandra Chekuri
Lucent Bell Labs
Murray Hill, NJ 07974
chekuri@lucent.com

Julia Chuzhoy
CSAIL MIT and Dept. of CIS,
Univ. of Pennsylvania
cjulia@csail.mit.edu

Liane Lewin-Eytan
Dept. of Electrical Eng.
Technion, Haifa 32000, Israel
liane@tx.technion.ac.il

Joseph (Seffi) Naor*
Microsoft Research
Redmond, WA 98052
naor@cs.technion.ac.il

Ariel Orda
Dept. of Electrical Eng.
Technion, Haifa 32000, Israel.
ariel@ee.technion.ac.il

ABSTRACT

We consider a multicast game with selfish non-cooperative players. There is a special source node and each player is interested in connecting to the source by making a routing decision that minimizes its payment. The mutual influence of the players is determined by a cost sharing mechanism, which in our case evenly splits the cost of an edge among the players using it. We consider two different models: an *integral* model, where each player connects to the source by choosing a single path, and a *fractional* model, where a player is allowed to split the flow it receives from the source between several paths. In both models we explore the overhead incurred in network cost due to the selfish behavior of the users, as well as the computational complexity of finding a Nash equilibrium.

The existence of a Nash equilibrium for the integral model was previously established by the means of a potential function. We prove that finding a Nash equilibrium that minimizes the potential function is NP-hard. We focus on the price of anarchy of a Nash equilibrium resulting from the *best-response dynamics* of a game course, where the players join the game sequentially. For a game with n players, we establish an upper bound of $O(\sqrt{n}\log^2 n)$ on the price of anarchy, and a lower bound of $\Omega(\log n/\log\log n)$. For the fractional model, we prove the existence of a Nash equilibrium via a potential function and give a polynomial time algorithm for computing an equilibrium that minimizes the potential function. Finally, we consider a weighted extension of the multicast game, and prove that in the fractional model, the game always has a Nash equilibrium.

*On leave from the Computer Science Department, Technion, Haifa 32000, Israel.

EC'06, June 11–15, 2006, Ann Arbor, Michigan.

Categories and Subject Descriptors

F.2 [**Analysis of Algorithms and Problem Complexity**]: Miscellaneous

General Terms

Algorithms, Performance, Theory

Keywords

Multicast game, Nash equilibrium, Price of anarchy, Price of stability

1. INTRODUCTION

In many networking scenarios, including the Internet, network users are free to act according to their individual interests, without taking into account overall network performance. Users thus may make selfish decisions (strategy choices) based on the state of the network, which depends (among other factors) on the behavior of other users, resulting in a non-cooperative game. Naturally, these scenarios call for a game-theoretic approach for studying both the behavior of such non-cooperative users, as well as their impact on the network performance. More specifically, we are interested in the properties of Nash equilibrium solutions [20] that are the stable outcomes of a non-cooperative game. We note that there has been a considerable amount of research dealing with non-cooperative games in networks [11, 14, 21, 24, 25, 27].

We consider a multicast game with selfish non-cooperative players. There is a special source node and each player is interested in connecting to the source by making a routing decision that minimizes its payment. Thus, the strategies of the players in the game correspond to the different paths by which the players can connect to the source. Each player independently chooses a strategy minimizing its payment. The mutual influence of the players is determined by a cost sharing mechanism that stipulates how the cost of each edge in the network is shared among its users. Multicast cost sharing methods have been studied extensively [1, 4, 6, 8, 10]. However, while typical models for non-cooperative games in networks have focused on congestion effects, where a resource utility deteriorates with the number of users that share it, an important class of resource sharing problems occurs when a fixed cost needs to be shared between a set of users. This game was recently introduced by Anshelevich *et al.* [3]. In this paper we study a natural cost sharing mechanism that falls into the above framework, where the cost of an edge is split evenly among all the players using it. More precisely, if k players use edge e of cost c_e, then each player pays c_e/k for this edge. This cost sharing formula has an intuitive appeal and it was investigated in several studies [6, 8]; it is also the outcome of the Shapley value [26].

Further motivation for the multicast game we consider is provided by the facility location problem, which is of fundamental interest in operations research. In a facility location game, we are given a set of facilities, with an opening cost associated with each facility. Additionally, we have a set of clients, and for each client-facility pair, we are given a cost that the client must pay for connecting to the facility. Each client needs to connect to one facility. A natural cost sharing mechanism for facility location is splitting the opening cost of each facility between the clients served by it. Additionally, each client pays for connecting to the facility serving it. Naturally, the clients seek to minimize their total payment, thus defining a non-cooperative game. This game constitutes a special case of the directed multicast game: given an instance of the facility location game, we add a source, connect each facility to the source with an edge of cost equal to the opening cost of the facility, and then connect each client to each facility with a *directed* edge of cost equal to the corresponding connection cost.

We consider two different models: an *integral* model, where each user connects to the source through a single path, and a *fractional* model, where each user is allowed to split (fractionally) its connection to the source into several paths, i.e., one unit of flow is sent fractionally by the source to the user. The fractional model, in addition to being a relaxation of the integral model, is interesting in its own right, as it is a splittable multicast model. The fractional model is also closely related to improving network throughput via network coding [2]. The games resulting from these models are referred to as the *integral multicast game* and the *fractional multicast game*, respectively.

A crucial property of our multicast game is that the per-user cost share on an edge is *non-increasing* in the number of users of the edge. Although, in this respect, the game differs from a classic congestion game, the integral multicast game does belong to the well known class of *congestion games*, that was first defined by Rosenthal [23] and has been widely investigated [9, 17, 19, 26, 28]. Rosenthal showed that a potential function can be defined for each congestion game with the property that the potential decreases if a player makes a move that improves its selfish cost. This shows that every congestion game has a Nash equilibrium. Moreover, there is a one-to-one correspondence between Nash equilibrium solutions and the solutions defining a local minimum of Rosenthal's potential function. Since the integral multicast game belongs to the class of congestion games, it has a Nash equilibrium and a potential function. We note that, for the integral model, the cost sharing mechanism guarantees that a Nash equilibrium induces a tree. The Nash equilibrium of the multicast game raises several natural questions. We focus in this paper on the inefficiency resulting from the selfish behavior of the players, and on the computational complexity of finding a Nash equilibrium.

We quantify the inefficiency resulting from a non-cooperative game through the ratio between the cost of a Nash equilibrium multicast tree and the cost of an optimal Steiner tree spanning the players. In keeping with common terminology [15, 22], this ratio is called the *price of anarchy* and it quantifies the "penalty" incurred by lack of cooperation (or coordination) between the players in a non-cooperative game.

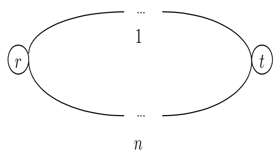

Figure 1: Cost of a Nash equilibrium tree can be n times the cost of an optimum Steiner tree.

Consider the graph in Figure 1 consisting of a source r and a node t with two parallel paths connecting them. The cost of one path is n, while the cost of the other path is 1. There are n players at t who want to connect to the source r. A solution where all players use the expensive path, each paying one unit, is a Nash equilibrium with a cost of n. A different and much cheaper Nash equilibrium is the one in which the players use the path of cost 1. Note that this second equilibrium is also the minimum cost Steiner tree connecting the players to the source. Thus, the price of anarchy for this game can be very large. Notice however that the expensive solution cannot be reached if the players join an initially empty game one-by-one, each of them choosing the cheapest path to connect to the source. In this paper, we investigate the price of anarchy of the integral multicast game for such scenarios. Motivated by the existence of large-cost Nash equilibria, the notion of *price of stability* was introduced in [3]: it is defined as the ratio between the cost of a Nash equilibrium of minimum cost and the cost of an optimal Steiner tree. In the above example, the price of stability is 1 in contrast to the price of anarchy which is n. For directed graphs, it was shown in [3] that the price of stability is $\Theta(\log n)$; for undirected graphs, an upper bound of $O(\log n)$ on the price of stability is known [3], however, no non-trivial lower bounds are known.

Even if the price of stability in undirected graphs is small, we still have two important questions to answer. Can a Nash equilibrium achieving (or approximating) the price of stability be computed in polynomial time? Second, can a good equilibrium be achieved as a consequence of *best-response dynamics*? That is, a course of the game where each player, in its turn, makes a routing decision that minimizes its cost. The price of anarchy of such a solution strongly depends on the initial configuration from which the players start. For example, if the starting solution is a Nash equilibrium with a large price of anarchy, as in the example in Figure 1, then best-response dynamics would not alter the solution. It is shown in [3] that even in directed graphs, if the initial configuration is a Steiner tree of cost C, then the best-response dynamics would lead to a Nash equilibrium of cost at most $O(C \log n)$. This is shown using Rosenthal's potential function, which can only decrease with each best-response move. This is also a constructive proof that the price of stability is $O(\log n)$. In [3], the above argument is used to suggest a mechanism in which a central authority starts the process by first computing a near-optimal Steiner tree on the receivers, and then allows the users to follow their best-response dynamics.

In this paper we take an approach that does not rely on a central trusted authority starting the game in a specific starting configuration. There are several situations in which having such an authority is expensive or infeasible. Further, not all the players might be available at the same time. In an online setting, players might arrive one by one to join a multicast service from the source. Motivated by these issues, we explore in this paper the following setting. Players first join the game sequentially starting from an "empty" configuration. Upon arrival, each player picks a path selfishly. Once reaching the solution constructed by the players joining one by one, the natural game course induced by best-response dynamics continues until a Nash equilibrium is reached. Our model is inspired by the *round* model considered by Mirrokhni and Vetta [18] to analyze convergence issues in competitive games. The one round question is the following. What is the quality of the solution to the game after each player has made one selfish move, assuming that the starting solution is the empty one? We believe that this gives insight into the realistic scenario where players typically do not make multiple moves. A positive result in this model is a good indication that the price of anarchy will be reasonable in practice. Note that we assume that the arrival of the players is adversarial.

Our Results: We focus on undirected graphs. For the integral multicast game, we establish an upper bound of $O(\sqrt{n}\log^2 n)$ on the price of anarchy of the best-response dynamics in the setting where the players join the game sequentially starting from an "empty" configuration. We then present a lower bound of $\Omega(\frac{\log n}{\log\log n})$ on the price of anarchy of this game. It is an interesting open question whether a polylogarithmic upper bound can be shown in this setting. We also prove that the problem of computing a Nash equilibrium minimizing Rosenthal's [23] potential function is NP-hard. It remains an open question whether a Nash equilibrium of the integral multicast game can be computed in polynomial time. We note that Fabrikant *et al.* [5] investigated the complexity of computing a pure Nash equilibrium for the class of congestion games, where the cost of a facility is a non-decreasing function of the number of its users, and showed that it is PLS-complete for general network congestion games. However, their proof heavily depends on the non-decreasing property of the cost sharing mechanism, and therefore does not seem to hold in our model.

For the fractional multicast game, we prove the existence of a Nash equilibrium by extending Rosenthal's potential function. Our main result for this model is that a Nash equilibrium that minimizes the potential function can be computed in polynomial time using linear programming. We observe that the fractional Nash equilibrium minimizing the potential function has a price of anarchy of $O(\log n)$. The results obtained for the fractional model hold also for more general settings, where the cost sharing mechanisms are *cross monotone*, which intuitively means that the share of a player on an edge cannot increase when additional players use it. Furthermore, the results also hold in the setting where there are multiple sources and each player needs to connect to at least one source. We note that the fact that our cost sharing mechanism is non-increasing in the number of players using an edge allows us to define a fractional extension. This does not seem possible with a non-decreasing cost sharing mechanism.

Finally, we consider a weighted extension of the multicast game, where each player has a *weight*, and the cost sharing mechanism splits the cost of an edge among its downstream receivers proportionally to their weights. For this game we prove that a Nash equilibrium exists in the fractional model. Initial results for the integral weighted game were obtained in [3]. However, establishing whether a Nash equilibrium for weighted multicast games exists in the integral model remains an open problem.

2. THE MODEL

We model our network by an undirected graph $G = (V, E)$. Let $c : E \to \mathcal{R}^+$ be a non-negative edge-cost function, and we denote by c_e the cost of edge $e \in E$. There is a special vertex $r \in V$ called root (or source) and a subset of n vertices $N = \{t_1, t_2, \ldots, t_n\}$ representing multicast *users* (also called *players* or *terminals*).

In the integral model, the goal of each user is to choose a single path P connecting it to the root, while minimizing its payment, which consists of the sum of the payments for the edges along P. A course of action chosen by player i at any time is called its *strategy* and is denoted by s^i. In the integral model, a strategy of player i is a path connecting t_i to the root. The strategy space of player i (i.e., the set of all its possible strategies) is denoted by S^i, and in our integral game, and it is the set of all the possible paths between t_i and the root. At any given moment, a strategy profile (or a configuration) of the game is the vector of all the strategies of the players, $s = (s^1, \ldots, s^n)$. We use s^{-i} to denote vector s without its ith coordinate, and $(s^{-i}, \tilde{s}^i)$ to denote the strategy profile identical to s, except that the ith coordinate is replaced by $\tilde{s}^i$. Given a strategy profile s, $c^i(s) \equiv c(s^i)$ denotes the payment of player i (the cost of its path s^i), and $n_e(s)$ denotes the number of players using edge e. Payment of user i for edge e is denoted by $c^i_e(s)$ and is determined by the cost sharing mechanism. We consider a natural cost sharing mechanism, where the cost of every edge is split evenly between the players sharing it. Thus, the payment of player i for edge e is $c^i_e(s) = \frac{c_e}{n_e(s)}$. We denote by $c(s)$ the sum of the costs of the edges participating in s (we say that edge e participates in a strategy profile s, iff at least one player chooses a path containing s to connect to the source). Let $H(k)$ denote the Harmonic number $\sum_{j=1}^{k} \frac{1}{j}$.

A strategy profile $s \in S$ is at *Nash equilibrium* if no player has an incentive to change its routing strategy, assuming that the strategies of the other players are fixed. We assume that a player changes its routing choice if and only if it reduces its payment. A change of strategy by any player is called a *Nash defection* and the corresponding player is called *Nash defector*. We assume that at each step the acting player chooses a strategy that minimizes the cost of its path, given the strategies of the other players. We therefore say that at each step the strategy of the current player is a *best response* to the other players' strategies.

3. THE INTEGRAL MULTICAST GAME

As we mentione earlier, the integral multicast game is a special case of a *congestion game*, formulated by Rosenthal [23], who defined a potential function to show that every congestion game possesses a Nash equilibrium. For our multicast game, given a strategy profile s, the potential function Φ of [23] is

$$\Phi(s) = \sum_{e} \sum_{k=1}^{n_e(s)} \frac{c_e}{k} \quad .$$

It is easy to see that for every instance of our game, a Nash equilibrium solution is a tree rooted at r spanning N.

We now analyze the price of anarchy of a multicast game in an undirected graph. We are interested in a Nash equilibrium that is a consequence of *best-response dynamics*, where each Nash defector, in its turn, chooses a path to the source minimizing its payment. Initially, the players join the game one by one starting from an "empty" configuration and picking a path to the root that minimizes their payment. Once all players are connected to the root, they continue playing until reaching Nash equilibrium. Note that we assume that the order by which the players play is adversarial. In Section 3.1 we establish an upper bound of $O(\sqrt{n}\log^2 n)$ on the price of anarchy for this game course, and in Section 3.2 we prove a lower bound of $\Omega(\frac{\log n}{\log\log n})$ on the price of anarchy. We also prove that finding a Nash equilibrium minimizing Rosenthal's [23] potential function is NP-hard in Section 3.3.

3.1 Upper Bound

In this section we establish an upper bound of $O(\sqrt{n}\log^2 n)$ on the price of anarchy of a Nash equilibrium obtained from best-response dynamics. Our analysis is performed in two steps. We first analyze in Section 3.1.1 the *first round* of the game in which the players connect one-by-one to the root via a cheapest path (best response). The first round finishes when all players are connected to the root. In order to bound the price of anarchy of the strategy profile T obtained from the first round, we define the notion of a *level tree* that serves as a basis of reference for proving the upper bound.

A *greedy online Steiner tree* [12] is defined to be the tree obtained when terminals arrive online and connect one by one via a cheapest path (i.e, the ith terminal connects by a cheapest path to the tree induced by terminals $1,\dots,i-1$). The cost of the greedy online Steiner tree is known to be at most a factor of $O(\log n)$ away from the cost of an optimal Steiner tree [12]. We consider the greedy online Steiner tree obtained from the same sequence of arrivals as in the first round of the game. Our goal is to prove that the cost of the solution obtained by the selfish moves of the players is related to the cost of the online Steiner tree. We are, however, unable to show this directly. We overcome this difficulty by first transforming the online Steiner tree to a level tree with reduced height using a procedure of Zelikovsky [29]. The height reduction increases the cost, but maintains ancestor relationships that are critical to maintain the online nature of the problem. We prove that the cost of the solution obtained from the first round is at most $O(\sqrt{n}\log n)$ times the cost of an optimal Steiner tree.

We complete our analysis in Section 3.1.2. Given the solution obtained from the first round, we then follow the natural game course until a Nash equilibrium is reached. The Nash defections performed until reaching an equilibrium can only decrease the potential of the first round solution, and thus we lose at most another factor of $O(\log n)$ with respect to the cost of the solution obtained from the first round.

We paraphrase below the height reduction lemma of Zelikovsky that we need. A bound claimed in [29] proved to be incorrect and a weaker correct bound is established in [7].

LEMMA 3.1. *Let $T=(V,A)$ be an in-tree rooted at $r\in V$ and let $c: A\to\mathcal{R}^+$ be a non-negative cost function on A. Let $G=(V,A_G)$ be the transitive closure of T and let $c': A_G\to\mathcal{R}^+$ be such that $c'(u,v)$ is the shortest c-path from u to v in T. Then, given integer $h>1$, there exists an in-tree tree $T'=(V,A')$ in G of height at most h such that $\sum_{a\in A'} c'(a)\le h\cdot|V|^{1/h}\sum_{a\in A} c(a)$.*

3.1.1 The First Round

We begin by analyzing the first round of the game in which players arrive one by one and pick a path selfishly. Let the sequence of arrivals of the terminals be $t_1,t_2,\dots,t_n$ (renumber if necessary), and let T be the resulting solution. We assume that the players start from an empty configuration.

Definition 1. A *level tree* T' on the vertex set $\{r(=t_0),t_1,t_2,\dots,t_n\}$, with a cost function $d: E\to\mathbb{R}$, is defined to be a tree having the following properties for each terminal t_i. (i) For $1\le i\le n$, the ancestor of terminal t_i in T' belong to $t_0,t_1,t_2,\dots,t_{i-1}$, i.e., terminals that have arrived before t_i. (ii) Let t and t_i be two terminals in T', such that t is the parent of t_i. The cost of the edge (t_i,t) in T', denoted by $d(i)$, is no less than the cost of the cheapest path between t_i and t in G.

Define $c(T')=\sum_{i=1}^n d(i)$. Let $T(i)$ denote the state of T after the arrival of $t_1,\dots,t_i$. Let P_i denote the path of t_i to the root r in T. We denote by $B(i)$ the set of new edges that are added to T when t_i joins $T(i-1)$. Let $b(i)=\sum_{e\in B(i)} c_e$. Clearly, $c(T)=\sum_{i=1}^n b(i)$. Let $c(i)$ be the cost paid by t_i when it joins T. Clearly, $c(i)\ge b(i)$. Note that in the single round case the cost paid by a player can only decrease during the round.

Given an edge e, let $n_e(i)$ denote the number of paths (terminals) using e in $T(i)$. We use $c_e(i)$ to denote the cost of e as seen by a selfish player in $T(i)$, i.e., $c_e/n_e(i)$. We use $c_e^+(i)$ to denote $c_e/(n_e(i)+1)$ which is the cost per player for using edge e if an additional player were to use e in $T(i)$. We define $c^+(i)$ to be $\sum_{e\in P_i} c_e^+(i)$. The following is immediate.

FACT 3.1. $c^+(i)\le (c(i)-b(i))+b(i)/2 = c(i)-b(i)/2.$

The edge set of T is partitioned by the sets $B(i)$, $1\le i\le n$. We now show how we charge the cost of edges in $B(i)$ to $d(1),\dots,d(i)$. Assume that we are given a level tree T' rooted at r having height 2. Let $t_{i_1},\dots,t_{i_m}$ be the first level terminals, i.e., the children of r in T'. The second level terminals are the children of $t_{i_1},\dots,t_{i_m}$, i.e., leaves of T'. Denote by $A(t_{i_j})$ the children of first level terminal t_{i_j}.

We first analyze the cost of the edges added to T by the first level terminals.

LEMMA 3.2. *For the first level terminals,*

$$\sum_{j=1}^{m} b(i_j) \le \sum_{j=1}^{m} c(i_j) \le \sum_{j=1}^{m} d(i_j).$$

PROOF. By Definition 1, for each first level terminal t_{i_j}, $1\le j\le m$, there is a path to the root r of cost at most $d(i_j)$ (without taking into account cost sharing). Therefore, $b(i_j)\le c(i_j)\le d(i_j)$. $\square$

We now analyze the cost of the edges added to T by the second level terminals.

LEMMA 3.3. *Let t_j be a first-level terminal with children $t_{j_1}, t_{j_2}, \ldots, t_{j_k}$ in T'. Then*

$$\sum_{i=1}^{k} b(j_i) \leq 2c^+(j) + 4\sum_{i=1}^{k} d(j_i).$$

PROOF. Assume that the arrival order is $t_{j_1}, t_{j_2}, \ldots, t_{j_k}$. Consider what happens when t_{j_1} arrives: it can connect to t_j, and then connect to the root via the path connecting t_j to the root. Hence, $c(j_1) \leq d(j_1) + c^+(j)$. Now consider terminal t_{j_i} for $i > 1$. It can connect to $t_{j_{i-1}}$ (paying at most $d(j_{i-1})+d(j_i)$), and then follow $P_{j_{i-1}}$ to the root. Hence, the cost of this path is at most $d(j_{i-1}) + d(j_i) + c^+(j_{i-1})$, which by Fact 3.1 is at most $d(j_{i-1}) + d(j_i) + c(j_{i-1}) - b(j_{i-1})/2$. Thus, we have for $1 < i \leq k$,

$$c(j_i) \leq d(j_{i-1}) + d(j_i) + c(j_{i-1}) - b(j_{i-1})/2.$$

Adding up the above inequalities, we obtain:

$$c(j_k) + \frac{1}{2}(b(j_1) + b(j_2) + \ldots + b(j_{k-1})) \leq$$
$$c^+(j) + d(j_k) + 2(d(j_1) + d(j_2) + \ldots d(j_{k-1})).$$

Since $b(j_k)/2 \leq b(j_k) \leq c(j_k)$, we obtain the desired inequality:

$$\sum_{i=1}^{k} b(j_i) \leq 2c^+(j) + 4\sum_{i=1}^{k} d(j_i).$$

□

We conclude with the next theorem.

THEOREM 1. $c(T) \leq 4c(T')$.

PROOF. We combine Lemmas 3.2 and 3.3 and get:

$$\begin{aligned}
c(T) &= \sum_{i=1}^{n} b(i) = \sum_{j=1}^{m}\left(b(i_j) + \sum_{t_\ell \in A(t_{i_j})} b(\ell)\right)\\
&\leq \sum_{j=1}^{m} b(i_j) + \sum_{j=1}^{m}\left(2c^+(i_j) + \sum_{t_\ell \in A(t_{i_j})} 4d(\ell)\right)\\
&\leq \sum_{j=1}^{m} d(i_j) + \sum_{j=1}^{m}\left(2d(i_j) + \sum_{t_\ell \in A(t_{i_j})} 4d(\ell)\right)\\
&\leq 4\sum_{i=1}^{n} d(i) \leq 4c(T').
\end{aligned}$$

□

An interesting question is whether the use of level trees which have depth greater than two can lead to better bounds on the price of anarchy. The difficulty with this approach is that for trees with more than two levels, a recursive use of Lemma 3.3 is necessary. However, the recursion introduces extra charges, and it is not clear how to bound them.

3.1.2 Completing the Analysis

We first generate the level tree T'. Note that the greedy online Steiner tree obtained from the sequence of arrivals of the first round of the game has all the properties required by a level tree. The difficulty is that the height of the greedy online Steiner tree can be $\Omega(n)$. We generate a new level tree T' from the greedy online Steiner tree by applying Lemma 3.1. The transformation preserves ancestral relationship and thus T' remains a level tree, while allowing us to restrict the height of the tree to be h at the expense of increasing its cost by a factor of $h \cdot n^{1/h}$. By choosing $h = 2$, we get a two level tree T' and cost at most $2\sqrt{n}$ times the cost of the greedy online Steiner tree. As the cost of a greedy online Steiner tree is within a factor of $O(\log n)$ away from the cost of an optimal Steiner tree, we get that $c(T') = O(\sqrt{n}\log n) \cdot c(T^*)$, where T^* is an optimal Steiner tree. Therefore, by Theorem 1, $c(T) = O(\sqrt{n}\log n) \cdot c(T^*)$.

Finally, after reaching the tree T constructed by the terminals in the first round, the natural best response dynamics are followed until a Nash equilibrium is reached. It is easy to see that the value of the potential function of any configuration is at most $\log n$ times the total cost of the edges used in this configuration. Therefore, the potential function value of T is within at most a factor of $O(\log n)$ away from $c(T)$. The potential function value can only decrease throughout the game. As the value of the potential function of a solution is always an upper bound on the cost of the edges participating in the solution, we get that the price of anarchy of our game is $O(\sqrt{n}\log^2 n)$.

3.2 Lower Bound

We present an undirected instance in which best-response dynamics converges to a Nash equilibrium with price of anarchy of $\Omega(\frac{\log n}{\log\log n})$. For the sake of simplicity, we first show how to achieve a lower bound of $2-\epsilon$ on the price of anarchy. The example is then generalized to show $\Omega(\frac{\log n}{\log\log n})$ price of anarchy.

Let q be a large integer. We start from a root vertex r and additional vertex u_1 connected to the root by a unit-length edge. We now add another unit-length path from the root to a new vertex u_2. The edges and the vertices of this path are as follows. Apart from r and u_2, there are $\log q$ vertices $v_1, \ldots, v_{\log q}$ that are placed between r and u_2 in this order, with vertex v_1 adjacent to r. The distance between v_1 and r is $\frac{1}{2}$, and for every $i > 1$, the distance between v_i and v_{i-1} is 2^{-i}, thus the distance between $v_{\log q}$ and u_2 is less than $\frac{1}{q}$. Finally, there is an edge of length $\frac{1}{q}$ between u_1 and u_2.

The idea is as follows. The first q players joining the game are placed on vertex u_1. They connect to r via the unit-length edge (r, u_1) and pay $\frac{1}{q}$ each. In the next step we place q players on vertex v_1. Naturally, they prefer to connect to r via edge (v_1, r) whose cost is $\frac{1}{2}$, instead of connecting via u_2 and u_1. Now the cost of the edge (v_1, r) becomes $\frac{1}{2q}$. When we place the next q players on vertex v_2, they connect via (v_2, v_1, r), as the cost of this path is less than $\frac{1}{4} + \frac{1}{2q}$ while connecting via u_2 and u_1 costs more than $\frac{1}{4} + \frac{1}{q}$. We continue in the same way, placing q users on vertices $v_3, \ldots, v_{\log q}$, where the users placed on vertex v_i all connect via path $(v_i, v_{i-1}, \ldots, v_1, r)$. Finally, we place q players on vertex u_2, who also prefer to connect via path (u_2, r), as its cost is less than $\frac{1}{q}$. It is easy to see that this configuration is a Nash equilibrium. The cost of this solution

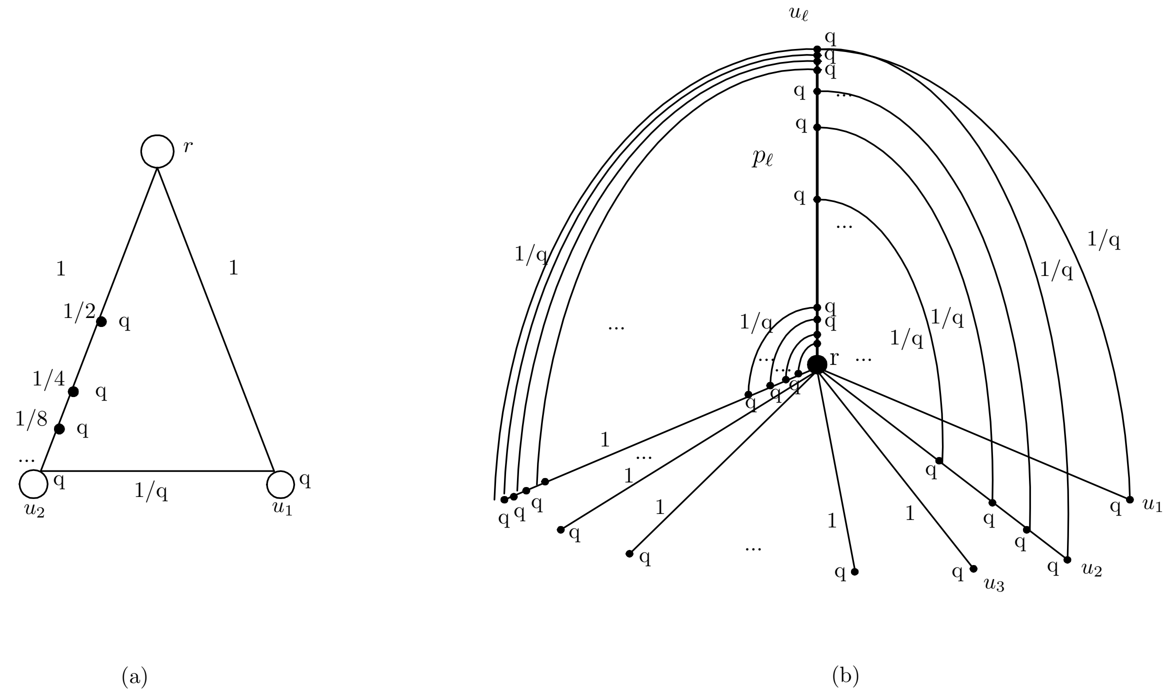

Figure 2: (a) Price of anarchy of $2-\epsilon$. (b) Price of anarchy of $\Omega(\frac{\log n}{\log\log n})$.

is 2, while the cost of the optimal solution is $1+\frac{1}{q}$, which is achieved by connecting all the players on path (u_2, r) via this path and connecting all the players on vertex u_1 via u_2. See Figure 3.1.2(a) for the resulting instance.

We now show how to generalize the above construction to obtain the bound of $\Omega(\frac{\log n}{\log\log n})$ on the price of anarchy. The basic gadget we use in our construction is a log-*division* of an edge.

Definition 2. Suppose we have an edge (a, b) of length c. A log-division of this edge is performed by converting this edge into a path $a, v_1, v_2, \ldots, v_{\log q}, b$ of the same length. The length of the edge (a, v_1) is $\frac{c}{2}$, and for each $i > 1$, the length of edge v_i, v_{i-1} is $\frac{c}{2^i}$. Thus, the length of edge $(v_{\log q}, b)$ is less than $\frac{c}{q}$.

A building block of our construction is a path p defined below. We use $\ell = \Omega(\frac{\log q}{\log\log q})$ copies of p. The construction of path p is as follows. We start from an edge (v, u) of length 1. Vertex u is called a level-1 vertex. We now perform ℓ iterations. In each iteration, we perform a log-division of every edge $e = (w, w')$ on path (v, u). When doing this division, the endpoint of e that is closer to v on the path (say, w) serves as a and the other endpoint (w') serves as b. For each iteration i, we call the vertices added to the path in this iteration "level i vertices". In our construction, we use ℓ copies of path p, denoted by $p_1, p_2, \ldots, p_\ell$. The endpoints v of these paths are merged together and form the root r. The other endpoints of the paths are denoted by $u_1, u_2, \ldots, u_\ell$. Finally, for each i, $1 \le i < \ell$, we connect each of the level-$1, 2, \ldots, i$ vertices on path p_i to the corresponding vertex on path p_ℓ by an edge of length $\frac{1}{q}$.

The players are added to the game as follows. First we add q players on vertex u_1. They connect via path p_1 to the root. Then, we add players on paths $p_2, \ldots, p_\ell$ in this order. For $i > 1$, we add players on all the vertices of levels $1, 2, \ldots, i$ belonging to path p_i, as well as on vertex u_i, in the order by which the vertices appear on the path starting from the root. See Figure 3.1.2(b) for the resulting instance.

Similarly to the case of two paths, for each i, all the players on path p_i connect via the subpath of p_i leading to r. Thus, the cost of the Nash equilibrium is ℓ. In the optimal solution, all players are connected via path p_ℓ. In order to connect some player belonging to path p_i, $i \neq \ell$, we use the edge of length $\frac{1}{q}$ connecting this player to path p_ℓ. Note that the total number of players M is bounded by $2qk$, where k is the number of vertices on path p_ℓ. Clearly, $k \le (\log q)^\ell$. Fixing $\ell = \frac{\log q}{\log\log q}$, we get $k \le q$. The total cost of the optimal solution is less than $1 + \frac{k}{q} \le 2$ and the price of anarchy is therefore $\Omega(\ell) = \Omega(\frac{\log q}{\log\log q})$. As $M \le 2qk \le 2q^2$, the price of anarchy is $\Omega(\ell) = \Omega(\frac{\log M}{\log\log M})$.

3.3 Intractability of Optimizing the Potential Function

We prove that finding a Nash equilibrium that minimizes the potential function is NP-hard. As a building block we use a variation of the Lund-Yannakakis proof [16] of hardness of approximation for the set cover problem.

The input to the set cover problem is a ground set of elements U and a collection $\mathcal{S}$ of subsets of U. The goal is to choose a minimum cardinality collection of sets in $\mathcal{S}$ covering all elements. The reduction of [16] is performed from the 3SAT problem. We use a straightforward deterministic variation of their construction together with a constant number of repetitions in the Raz verifier to obtain the following result:

THEOREM 2. *Given a 3SAT formula φ, an instance of the set cover problem can be constructed in polynomial time, such that:*

- *All sets have equal size (denoted by s).*
- *If φ is satisfiable (yes-instance), then there is a solution to the set cover instance that uses X sets, and each element is covered by exactly one set in this solution.*
- *If φ is not satisfiable (no-instance), then the size of any solution to the set cover instance is at least αX, where $\alpha > 1$ is some constant.*

Suppose we are given a 3SAT formula φ. We construct an integral multicast game based on the corresponding set cover instance, as follows. There is a vertex for each set and each element in the set cover problem, and additionally we have a special vertex r. The players are the vertices that represent the elements. Each vertex representing some set is connected to r with a unit-length edge. Each vertex representing some element i is connected to a vertex representing set S if and only if $i \in$S. The length of the edge is a large integer q, which will ensure that each user (element) chooses a path that contains only one such edge (i.e., connects via a set to which it belongs). Let N denote the total number of users (elements) in the above example. Then it is enough to choose $q \geq N$.

Suppose φ is a yes-instance. Then there is a solution $\mathcal{S}'$ of size X to the set cover instance. This solution naturally induces a Nash equilibrium in our game, where each element connects to the set that covers it in $\mathcal{S}'$ and all the sets in $\mathcal{S}'$ are connected to the root. Observe that there are exactly s users on every edge that connects some set in $\mathcal{S}'$ to the root. The value of the potential function in this solution is $qN + X \cdot H(s)$.

Assume now that φ is a no-instance and suppose we are given some Nash equilibrium. This Nash equilibrium defines a solution to the set cover instance, since each element has to connect to one of the sets to which it belongs. However, the number of sets used in this solution is at least αX, and some of the edges connecting these sets to the root are used by less than s users. Thus, the value of the potential function in this solution is strictly greater than $qN + X \cdot H(s)$.

As determining whether a given 3SAT formula is satisfiable is NP-hard, it is NP-hard to find a Nash equilibrium minimizing the value of the potential function.

4. THE FRACTIONAL MULTICAST GAME

In this section we introduce a fractional model of the multicast game, where each user is allowed to split (fractionally) its connection to the source into several paths. The fractional model represents a splittable multicast model. The cost of each flow fraction on an edge is evenly split between its users. We present our results for undirected graphs, yet they hold for directed graphs as well. In the fractional model, each user i has to route one unit of flow from t_i to the source r. User i can split its unit of flow among any number of paths connecting r to t_i. Denote the flow of user i on edge e by $f_{e,i}$ and the number of users on edge e by n_e. Given a strategy profile s, assume without loss of generality that $f_{e,1} \leq f_{e,2} \leq \cdots \leq f_{e,n_e} \leq 1$. Define $f_{e,0} = 0$. Edge e has capacity equal to f_{e,n_e}, and for convenience we think of the capacity of e as defining an "address space" in the range $[0, f_{e,n_e}]$, where user j uses $[0, f_{e,j}]$. The cost of each fraction of the capacity of e is equally split between its users, as follows: $[f_{e,j-1}, f_{e,j}]$ is shared by $n_e - j + 1$ users, where each user pays $c_e \frac{f_{e,j} - f_{e,j-1}}{n_e - j + 1}$. Therefore, the total cost c_e^i paid by user i for the use of edge e is:

$$c_e^i = c_e \cdot \sum_{k=1}^{i} \frac{f_{e,k} - f_{e,k-1}}{n_e - k + 1}.$$

As the total flow fraction sent on edge e is f_{e,n_e}, the total cost of the edge is simply $c_e \cdot f_{e,n_e}$.

We denote by P^i the set of paths used by user i. The cost of a path $p \in P^i$ is the sum of its edge costs, that is $\sum_{e \in p} c_e^i$. The total cost c^i of a user i is the sum of its path costs, that is $\sum_{p \in P^i} \sum_{e \in p} c_e^i$. Each user i aims to establish its flow from the source r to t_i so as to minimize its cost. Thus, a flow f is at Nash equilibrium if no user has any incentive of changing its flow to the root. An instance of the fractional model, consisting of a graph G, a source r, a set of receivers N, and a cost vector c is denoted by $frac(G, r, N, c)$. We introduce a potential function Φ for the fractional multicast game which is based on Rosenthal's potential function [23], as follows:

$$\Phi = \sum_{e \in s} \sum_{j=1}^{n_e(s)} \sum_{i=1}^{n_e+1-j} c_e \frac{f_{e,j} - f_{e,j-1}}{i} \quad .$$

The proof of the following theorem follows from the proof of the potential function defined by Rosenthal [23].

THEOREM 3. *Potential fuction Φ is an exact potential for the fractional multicast game. That is, for every $k \in N$, and every two strategy profiles (s^{-k}, s_1^k) and (s^{-k}, s_2^k), where $c^k(s^{-k}, s_1^k) < c^k(s^{-k}, s_2^k)$ (c^k denoting the total cost of user k), it holds that*

$$c^k(s^{-k}, s_2^k) - c^k(s^{-k}, s_1^k) = \Phi(s^{-k}, s_2^k) - \Phi(s^{-k}, s_1^k).$$

As a fractional flow configuration defining a local minimum of the potential function is at Nash equilibrium, we get:

THEOREM 4. *A Nash equilibrium exists for every instance $frac(G, r, N, c)$.*

4.1 Computing a Minimum Potential Nash Equilibrium

We proceed to describe how a Nash equilibrium of the fractional game can be computed in polynomial time using linear programming. Moreover, the computed Nash equilibrium minimizes the potential function Φ. Compare that with the hardness of finding an integral solution minimizing the potential function.

Given an instance $frac(G, r, N, c)$, we create a new graph $G' = (V, E')$ by replacing each edge e by n copies $e_1, e_2, \ldots, e_n$. The cost of a unit flow on edge e_j is c_e/j. For a path p from t_i to r in G', we denote by f_p^i the amount of flow of commodity i sent on it. Note that different paths can use an edge in opposite directions.

We formulate a linear program with an objective function equal to the potential of the fractional multicast game. The variables of the linear program are the flows of the users sent on the set of paths in G' from the terminals $t_1, \ldots, t_n$ to the root r, and the capacities of the edges in E'. Denote a path from t_i to r by $t_i \rightsquigarrow r$. The capacity of edge e_j is denoted by x_{e_j}, where $0 \leq x_{e_j} \leq 1$. The linear program is as follows.

$$\text{minimize} \quad \sum_{e \in E} \sum_{j=1}^{n} \sum_{i=1}^{j} \frac{c_e \cdot x_{e_j}}{i} \qquad s.t.$$

$$\text{For each commodity } i\text{: } \sum_{p:t_i \rightsquigarrow r} f_p^i \geq 1 \quad (1)$$

For each edge e, copy j, commodity i:

$$\sum_{p:t_i \rightsquigarrow r | e_j \in p} f_p^i \leq x_{e_j} \quad (2)$$

For each edge e, copy j:

$$\sum_{i=1}^{n} \sum_{p:t_i \rightsquigarrow r | e_j \in p} f_p^i = j \cdot x_{e_j} \quad (3)$$

$$0 \leq x_{e_j} \leq 1, \quad f_p^i \geq 0 \quad (4)$$

The total flow of user i, summed up over all paths from t_i to r, is at least 1 (Constraint (1)). Constraint (2), the *non-aggregating* flow constraint, restricts the flow of each user i on edge e_j to be at most its capacity x_{e_j}. The total flow, taken over all commodities on edge e_j, is constrained to be precisely $j \cdot x_{e_j}$, as restricted by Constraint (3), the *aggregating* flow constraint. This constraint is satisfied in the integral case: if j commodities are sent on edge e, then edge e_j is "bought", and the number of users on this edge is j. The sum of the costs of the commodities on e_j is then exactly c_e.

Note that the above linear program uses an exponential number of variables. However, it can be solved in polynomial time via the dual program using the Ellipsoid algorithm. Alternatively, it can be formulated with a polynomial number of variables by using the flows of the users on the different edges in G' as variables.

4.1.1 Characterizing an Optimal Solution

We say that a flow f on instance G' is *canonical* if it has, for every edge $e \in E$, the following structure. Denote by $f_{e,j}$ the sum of the flows of user j on all copies of e ($f_{e,j} = \sum_{k=1}^{n} \sum_{p:t_j \rightsquigarrow r | e_k \in p} f_p^j$). Suppose that without loss of generality $f_{e,1} \leq f_{e,2} \leq \cdots \leq f_{e,n_e} \leq 1$, where n_e denotes the number of users with positive flow. Then, the flows routed on $e_1, e_2, \ldots, e_{n_e}$ are $f_{e,1}, (f_{e,2} - f_{e,1}), \ldots, (f_{e,j} - f_{e,j-1}), \ldots, (f_{e,n_e} - f_{e,n_e-1})$, respectively, and the flow on copies e_i for $i \geq n_e + 1$ is zero. Notice that there is a one-to-one correspondence between canonical flows in the instance G' and fractional multicast flows in G. We now turn to prove that there exists a canonical flow minimizing the potential function.

Let f be the output flow of the linear program. We first consider the flow f_{e_k} on each copy e_k of edge e, and rearrange it to be a canonical flow. Then, we merge these resulting canonical flows into a single canonical flow on e. These two steps are performed for each edge $e \in E$. We show that the resulting potential of the new (canonical) flow is not larger than the potential of the original flow f.

LEMMA 4.1. *Consider edge $e_k \in E'$, $1 \leq k \leq n$. There exists a canonical flow on e_k with potential value not greater than that of the original flow on e_k.*

PROOF. Without loss of generality, suppose that $f_{e_k,1} < f_{e_k,2} < \cdots < f_{e_k,\ell}$ are the different amounts of flow routed on e_k by the users, where $f_{e_k,\ell} = x_{e_k}$. For ease of notation, we denote this ordering as $f_1 < f_2 < \cdots < f_\ell$, where $f_\ell = x_{e_k}$. Assume that the number of users routing a flow value $\leq f_i$ is k_i, and thus $k_1 > k_2 > \cdots > k_\ell$. We rearrange the flow f_{e_k} to be a canonical flow by sending each amount of flow to its proper edge copy, i.e. by "buying" capacity f_1 on edge e_{k_1}, capacity $(f_2 - f_1)$ on edge e_{k_2}, etc.

The potential of the resulting canonical flow derived from f_{e_k} is thus

$$c_e \sum_{i=1}^{\ell} \sum_{j=1}^{k_i} \frac{f_i - f_{i-1}}{j} = c_e \sum_{i=1}^{\ell} (f_i - f_{i-1}) H(k_i),$$

where $f_0 = 0$. On the other hand, the potential of the original flow on e_k is

$$c_e \sum_{i=1}^{k} \frac{x_{e_k}}{i} = c_e \cdot H(k) x_{e_k} = c_e \cdot H(k) f_\ell.$$

The total flow on e_k is constrained to be $k \cdot x_{e_k} = k \cdot f_\ell$ (Constraint (3)), which is equal to the total canonical flow derived from f_{e_k}, and thus

$$k \cdot f_\ell = \sum_{i=1}^{\ell} k_i (f_i - f_{i-1}).$$

Since $0 \leq \frac{(f_i - f_{i-1})}{f_\ell} \leq 1$ and $\sum_{i=1}^{\ell} \frac{(f_i - f_{i-1})}{f_\ell} = 1$, by Jensen's inequality,

$$\sum_{i=1}^{\ell} H(k_i) \frac{f_i - f_{i-1}}{f_\ell} \leq H \left(\sum_{i=1}^{\ell} k_i \cdot \frac{f_i - f_{i-1}}{f_\ell} \right) = H(k),$$

and thus $\sum_{i=1}^{\ell} H(k_i)(f_i - f_{i-1}) \leq H(k) f_\ell$. $\square$

LEMMA 4.2. *Consider edge $e \in E$ and two canonical flows f_e and f'_e. Then f_e and f'_e can be added up yielding a canonical flow with potential value not greater than the sum of the potentials of f_e and f'_e.*

PROOF. Consider two canonical flows f and f', and assume that their respective flows on e_k are x_k and x'_k. That is, x_k (resp., x'_k) is the amount of flow routed by each player using e_k according to f (resp., f'). We denote by G_k and G'_k the sets of players that use e_k according to f and f' respectively, where $|G_k| = |G'_k| = k$. Assume, w.l.o.g., that $x'_k \geq x_k$. By merging these two flows into a single canonical flow, we "buy" capacity x_k on copy $|G_k \cup G'_k|$ of edge e, capacity $(x'_k - x_k)$ on e_k, and capacity x_k on copy $|G_k \cap G'_k|$ of edge e. We thus get a new canonical flow with potential

$$H(|G_k \cup G'_k|) x_k + H(k)(x'_k - x_k) + H(|G_k \cap G'_k|) x_k.$$

On the other hand, the sum of potentials of the original flows on e_k is $H(k) x_k + H(k) x'_k$. As

$$2H(k) \geq H(|G_k \cup G'_k|) + H(|G_k \cap G'_k|),$$

we get that the potential of the new canonical flow is not larger than the sum of potentials of the original flows.

In case $G_k \neq G'_k$, the potential of the new canonical flow is strictly less than the sum of the potentials of the original canonical flows. In this case, capacity has to be bought on other copies of e except from e_k, and thus other merging steps should be performed for each such copy. As each such step strictly decreases the potential, the merging process is finite. □

THEOREM 5. *There exists an optimal solution to the linear program which is a canonical flow.*

PROOF. Let f be the output flow of the linear program. As f is a flow of minimum potential, it is either canonical, or can be easily rearranged as such by performing the two steps described by Lemmas 4.1 and 4.2 on all copies of each edge e. □

The linear program presented for computing the minimum potential Nash equilibrium of the fractional model can be used for more general settings, not necessarily egalitarian, where the cost sharing mechanisms are cross-monotonic, i.e. the cost functions are non-increasing in the number of users. Furthermore, it can also be used for settings where the users are not restricted to have a common source. Recall that finding an integral solution with minimum potential is NP-hard.

There are instances for which there is a gap between the minimum potential fractional Nash equilibrium and the minimum potential integral Nash equilibrium. Consider Figure 3, which depicts an instance $(G, r, \{t_1, t_2, t_3\}, c)$, for which the minimum potential fractional Nash equilibrium is smaller than the minimum potential integral Nash equilibrium. The edge costs are as follows: the cost of each edge (r, v_i) $(i = 1, 2, 3)$ is x, and the cost of each edge from v_i to the terminals connected to it is $q \gg x$. The fractional Nash equilibrium that minimizes the potential is as follows: each terminal t_i sends 1/2 unit of flow through each of the two vertices v_j $(j = 1, 2, 3)$ connecting it to r. Therefore, the fractional potential is $\Phi_{frac} = 3x/2(1+1/2)+6q/2 = 9x/4+3q$. On the other hand, the integral Nash equilibrium that minimizes the potential is as follows: two out of the three terminals send their flow through the same vertex v_i to r, and the third terminal sends its flow through one out of the other two vertices v_j connecting it to r. Therefore, the integral potential is $\Phi_{int} = x(1 + 1/2) + x + 3q = 10x/4 + 3q$.

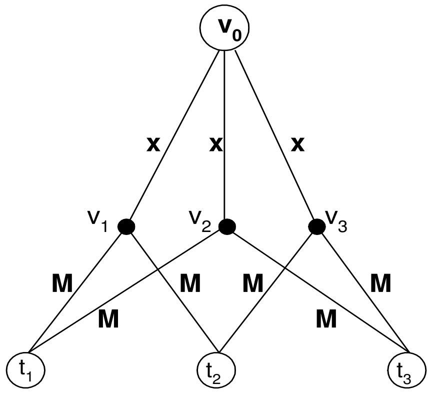

Figure 3: Instance in which the minimum potential fractional Nash equilibrium is smaller than the minimum potential integral Nash equilibrium.

We define the *price of anarchy* of the fractional game as the ratio between the cost of a Nash equilibrium solution and the cost of an optimal fractional solution to the Steiner tree problem. Similarly to the integral game, it can be shown that the price of anarchy of a minimum potential function Nash equilibrium is $O(\log n)$.

4.2 The Weighted Fractional Multicast Game

We consider a weighted extension of our multicast game, where each user has a positive weight: w_i denotes the weight of user i. The payment of each player is proportional to its weight. Let $W_j = \sum_{i=1}^{j} w_i$. Given a flow vector f, let $f_{e,i}$ denote the flow of user i on e and let n_e be the number of users with non-zero flow on e. Assume that the users are numbered such that $0 = f_{e,0} < f_{e,1} \le f_{e,2} \le \ldots \le f_{e,n_e}$. Consider some $j \le n_e$. User j shares the capacity $f_{e,1}$ with users 1 to n_e, shares $f_{e,2} - f_{e,1}$ with users 2 to n_e and so on. The payment for any share is in proportion to weights. Hence, for the capacity $f_{e,1}$, user $j \le n_e$ pays $c_e \cdot f_{e,1} \cdot w_j / W_{n_e}$. Thus the overall cost paid by $j \le n_e$ on edge e is

$$c_e \cdot w_j \cdot \sum_{i=1}^{j} \frac{f_{e,i} - f_{e,i-1}}{W_{n_e} - W_{i-1}}.$$

The overall payment of a user is the sum of its payments for the flow fractions it uses on all edges in all its paths. Each user j aims to establish its flow from the source r to t_j so as to minimize its cost. Thus, a flow f is at Nash equilibrium if no user has any incentive to change its flow.

An instance of the weighted fractional model, consisting of a graph G, a source r, a set of receivers N with weight vector w, and a cost vector c is denoted by $frac(G, r, N, c, w)$. The proof of the following theorem uses Kakutani's fixed point theorem [13], and it is omitted from this extended abstract. It is not known whether the weighted integral multicast game has a Nash equilibrium.

THEOREM 6. *A Nash equilibrium (in pure strategies) exists for every instance $frac(G, r, N, c, w)$.*

5. REFERENCES

[1] M. Adler, D. Rubenstein. Pricing multicasting in more practical network models. *Proc. of ACM-SIAM SODA*, 981–990, 2002.

[2] A. Agarwal and M. Charikar. On the advantage of network coding for improving network throughput. *Proc. of IEEE Information Theory Workshop*, 2004.

[3] E. Anshelevich, A. Dasgupta, J. Kleinberg, É. Tardos, T. Wexler and T. Roughgarden. The price of stability for network design with fair cost allocation. *Proc. of IEEE FOCS*, 295-304, 2004.

[4] A. Archer, J. Feigenbaum, A. Krishnamurthy, R. Sami and S. Shenker. Approximation and Collusion in Multicast Cost Sharing. *Proc. of ACM EC*, 253–255, 2001.

[5] A. Fabrikant, C. Papadimitriou and K. Talwar. The Complexity of Pure Nash Equilibria. *Proc. of ACM STOC*, 604–612, 2004.

[6] J. Feigenbaum, C. Papadimitriou and S. Shenker. Sharing the cost of multicast transmission. *Proc. of ACM STOC*, 218–227, 2000.

[7] C.H. Helvig, G. Robins, and A. Zelikovsky. Improved approximation scheme for the group Steiner problem. *Networks*, 37(1):8–20, 2001.

[8] S. Herzog, S. Shenker and D. Estrin. Sharing the cost of multicast trees: An axiomatic analysis. *IEEE/ACM Transactions on Networking*, Vol. 5, 847–860, 1997.

[9] R. Holzman and N. Law-Yone. Strong equilibrium in congestion games. *Games and Economic Behavior*, Vol. 21, pp. 85–101, 1997.

[10] K. Jain and V. Vazirani. Applications of approximation algorithms to cooperative games. *Proc. of the ACM STOC*, 364–372, 2001.

[11] R. Johari and J.N. Tsitsiklis. Efficiency loss in a network resource allocation game. To appear in *Math. of Oper. Res.*

[12] M. Imaze and B. Waxman. Dynamic steiner tree problem. *SIAM J. on Discrete Mathematics*, Vol. 4(3):369–384, 1991.

[13] S. Kakutani. A generalization of Brouwer's Fixed Point Theorem. *Duke Mathematical Journal*, Vol. 8, 457–459, 1941.

[14] Y.A. Korilis, A.A. Lazar and A. Orda. Achieving Network Optima Using Stackelberg Routing Strategies. *IEEE/ACM Transactions on Networking*, 5(1):161-173, 1997.

[15] E. Koutsoupias and C. Papadimitriou. Worst-case equilibria. *Proc. of STACS*, LNCS, 404–413, 1999.

[16] C. Lund and M. Yannakakis. On the hardness of approximating minimization problems. *JACM*, 41(5):960–981, 1994.

[17] I. Milchtaich. Congestion games with player specific payoff functions. *Games and Economic Behavior*, Vol. 13, 111–124, 1996.

[18] V. Mirrokni and A. Vetta. Convergence Issues in Competitive Games. *Proc. of APPROX*, LNCS, 183–194, 2004.

[19] D. Monderer and L.S. Shapley. Potential Games. *Games and Economic Behavior*, 14:124–143, 1996.

[20] J. Nash. Non-cooperative games. *Annals of Mathematics*, 2nd Ser., Vol. 54(2), pp. 286-295, 1951.

[21] A. Orda, R. Rom and N. Shimkin. Competitive routing in multi-user communication networks. *IEEE/ACM Transaction on Networking*, Vol. 1, 614–627, 1993.

[22] C. Papadimitriou. Algorithms, games, and the Internet. *Proc. of ACM STOC*, 749–753, 2001.

[23] R.W. Rosenthal. A class of games possessing pure strategy Nash equilibria. *International Journal of Game Theory*, Vol. 2, 65–67, 1973.

[24] T. Roughgarden and É. Tardos. How bad is selfish routing? *JACM*, Vol. 49(2):236–259, 2002.

[25] T. Roughgarden. The price of anarchy is independent of the network topology. *Proc. of ACM STOC*, 428–437, 2002.

[26] L.S. Shapley. A value for n-person games. Contribution to the Theory of Games, pp. 31-40, Princeton University Press, Princeton, 1953.

[27] A.S. Schulz and N.E. Stier Moses. On the performance of user equilibria in traffic networks. *Proc. of ACM-SIAM SODA*, 86–87, 2003.

[28] M. Voorneveld, P. Borm, F. van Megen, S. Tijs and G. Faccini. Congestion games and potentials reconsidered. *International Game Theory Review*, Vol. 1, 283–299, 1999.

[29] A. Zelikovsky. A series of approximation algorithms for the acyclic directed steiner tree problem. *Algorithmica*, Vol. 18, 99–110, 1997.