Minimizing Rosenthal Potential in Multicast Games*

Fedor V. Fomin[1], Petr Golovach[2], Jesper Nederlof[3], and Michał Pilipczuk[1]

[1] Department of Informatics, University of Bergen, Bergen, Norway
{fomin,michal.pilipczuk}@ii.uib.no
[2] School of Engineering and Computing Science, Durham University, Durham, UK
petr.golovach@durham.ac.uk
[3] Utrecht University, Utrecht, the Netherlands
j.nederlof@uu.nl

Abstract. A multicast game is a network design game modelling how selfish non-cooperative agents build and maintain one-to-many network communication. There is a special source node and a collection of agents located at corresponding terminals. Each agent is interested in selecting a route from the special source to its terminal minimizing the cost. The mutual influence of the agents is determined by a cost sharing mechanism, which evenly splits the cost of an edge among all the agents using it for routing. The existence of a Nash equilibrium for the game was previously established by the means of Rosenthal potential. Anshelevich et al. [FOCS 2004, SICOMP 2008] introduced a measure of quality of the best Nash equilibrium, the price of stability, as the ratio of its cost to the optimum network cost. While Rosenthal potential is a reasonable measure of the quality of Nash equilibra, finding a Nash equilibrium minimizing this potential is NP-hard.

In this paper we provide several algorithmic and complexity results on finding a Nash equilibrium minimizing the value of Rosenthal potential. Let n be the number of agents and G be the communication network. We show that

- For a given strategy profile s and integer $k \geq 1$, there is a local search algorithm which in time $n^{O(k)} \cdot |G|^{O(1)}$ finds a better strategy profile, if there is any, in a k-exchange neighbourhood of s. In other words, the algorithm decides if Rosenthal potential can be decreased by changing strategies of at most k agents;
- The running time of our local search algorithm is essentially tight: unless $FPT = W[1]$, for any function $f(k)$, searching of the k-neighbourhood cannot be done in time $f(k) \cdot |G|^{O(1)}$.

The key ingredient of our algorithmic result is a subroutine that finds an equilibrium with minimum potential in $3^n \cdot |G|^{O(1)}$ time. In other words, finding an equilibrium with minimum potential is fixed-parameter tractable when parameterized by the number of agents.

* This work is supported by EPSRC (EP/G043434/1), Royal Society (JP100692), European Research Council (ERC) grant "Rigorous Theory of Preprocessing", reference 267959, and Nederlandse Organisatie voor Wetenschappelijk Onderzoek (NWO), project: 'Space and Time Efficient Structural Improvements of Dynamic Programming Algorithms'.

A. Czumaj et al. (Eds.): ICALP 2012, Part II, LNCS 7392, pp. 525–536, 2012.

1 Introduction

Modern networks are often designed and used by non-cooperative individuals with diverse objectives. A considerable part of Algorithmic Game Theory focuses on optimization in such networks with selfish users [2,6,9,13,14,16,21,22].

In this paper we study the conceptually simple but mathematically rich cost-sharing model introduced by Anshelevich et al. [3,4], see also [15, Chapter 12]. In a variant of the cost-sharing game, which was called by Chekuri et al. the *multicast game* [5], the network is represented by a weighted directed graph with a distinguished source node r, and a collection of n agents located at corresponding terminals. Each agent is trying to select a cheapest route from r to its terminal. The mutual influence of the players is determined by a cost sharing mechanism identifying how the cost of each edge in the network is shared among the agents using this edge. When h agents use an edge e of cost c_e, each of them has to pay c_e/h. This is a very natural cost sharing formula which is also the outcome of the Shapley value.

The multicast game belongs to the widely studied class of congestion games. This class of games was defined by Rosenthal [20], who also proved that every congestion game has a Nash equilibrium. Rosenthal showed that for every congestion game it is possible to define a potential function which decreases if a player improves its selfish cost. Best-response dynamics in these games always lead to a set of paths that forms a Nash equilibrium. Furthermore, every local minimum of Rosenthal potential corresponds to a Nash equilibrium and vice versa. However, while we know that the multicast game always has a Nash equilibrium, the number of iterations in best-response dynamics achieving an equilibrium can be exponential (see [3, Theorem 5.1]), and it is an important open question if any Nash equilibrium can be found in polynomial time. The next step in the study of Rosenthal potential was done by Anshelevich et al. [3], who showed that Rosenthal potential can be used not only for proving the existence of a Nash equilibrium but also to estimate the quality of equilibrium. Anshelevich et al. defined the price of stability, as the ratio of the best Nash equilibrium cost and the optimum network cost, the *social optimum*. In particular, the cost of a Nash equilibrium minimizing Rosenthal potential is within $\log n$-factor of the social optimum, and thus the global minimum of the potential brings to a cheap equilibrium. The computational complexity of finding a Nash equilibrium achieving the bound of $\log n$ relative to the social optimum is still open, while computing the minimum of the Rosenthal potential is NP-hard [3,5].

Our results. In this paper we analyze the following local search problem. Given a strategy profile s, we are interested if a profile with a smaller value of Rosenthal potential can be found in a k-exchange neighbourhood of s, which is the set of all profiles that can be obtained from s by changing strategies of at most k players. Our motivation to study this problem is two-fold.

– If we succeed in finding some Nash equilibrium, say by implementing best-response dynamics, which is still far from the social optimum, it is

an important question if the already found equilibrium can be used to find a better one efficiently. Local search heuristic in this case is a natural approach.
- Since the number of iterations in best-response dynamics scenario can be exponential (see [3, Theorem 5.1]), it can be useful to combine the best-response dynamics with a heuristic that at some moments tries to make "larger jumps", i.e., instead of decreasing Rosenthal potential by changing strategy of one player, to decrease the potential by changing in one step strategies of several players.

Let us remark that the number of paths, and thus strategies, every player can select from, is exponential, so the size of the search space also can be exponential. Since the size of k-exchange neighbourhood is exponential, it is not clear a priori, if searching of a smaller value of Rosenthal potential in a k-exchange neighbourhood of a given strategy profile can be done in polynomial time. We show that for a fixed k, the local search can be performed in polynomial time. The running time of our algorithm is $n^{O(k)} \cdot |G|^{O(1)}$, where n is the total number of players[1]. As a subroutine, our algorithm uses a fixed-parameter algorithm computing the minimum of Rosenthal potential in time $3^n \cdot |G|^{O(1)}$. We find this auxiliary algorithm to be interesting in its own. It is known that for a number of local search algorithms, exploration of the k-exchange neighbourhood can be done by fixed-parameter tractable (in k) algorithms [10,17,23]. We show that, unfortunately, this is not the case for the local search of Rosenthal potential minimum. We use tools from Parameterized Complexity, to show that the running time of our local search algorithm is essentially tight: unless $FPT = W[1]$, searching of the k-neighbourhood cannot be done in time $f(k) \cdot |G|^{O(1)}$ for any function $f(k)$.

2 Preliminaries

Multicast game and Rosenthal potential. A network is modeled by a directed $G = (V, E)$ graph. There is a special *root* or *source* node $r \in V$. There are n multicast users, *players*, and each player has a specified *terminal* node t_i (several players can have the same terminals). A strategy s^i for player i is a path P_i from r to t_i in G. We denote by Π the set of players and by S^i the finite set of strategies of player i, which is the set of all paths from r to t_i. The joint strategy space $S = S^1 \times S^2 \times \cdots \times S^n$ is the Cartesian product of all the possible strategy profiles. At any given moment, a strategy profile (or a configuration) of the game $s \in S$ is the vector of all the strategies of the players, $s = (s^1, \ldots, s^n)$. Notice that for a given strategy profile s, several players can use paths that go through the same edge. For each edge $e \in E$ and a positive integer h, we have a cost $c_e(h) \in \mathbb{R}$ of the edge e for each player who uses a path containing e, provided that exactly h players share e. With each player i, we associate the cost function c^i mapping a strategy profile $s \in S$ to real numbers, i.e., $c^i : S \to \mathbb{R}$. For a

[1] The number of arithmetic operations used by our algorithms does not depend on the size of the input weights, i.e. the claimed running times are in the unit-cost model.

strategy profile $s \in S$, let $n_e(s)$ be the number of players using the edge e in s. Then the cost the i-th player has to pay is

$$c^i(s) = \sum_{e \in E(P_i)} c_e(n_e(s)),$$

and the total cost of s is

$$c(s) = \sum_{i=1}^{n} c^i(s).$$

The *potential* of a strategy profile $s \in S$, or equivalently, the set of paths $(P_1, \ldots, P_n)$, is

$$\Phi(s) = \sum_{e \in \cup_{i=1}^{n} E(P_i)} \sum_{h=1}^{n_e(s)} c_e(h). \tag{1}$$

In this paper, we are especially interested in the case where the cost of every edge is split evenly between the players sharing it, i.e, the payment of player i for edge e is $c_e(h) = \frac{c_e}{h}$ for $c_e \in \mathbb{R}$. Respectively, *Rosenthal potential* of a strategy profile $s \in S$ is

$$\Phi(s) = \sum_{e \in \cup_{i=1}^{n} E(P_i)} c_e \cdot \mathcal{H}(n_e(s)),$$

where $\mathcal{H}(h) = 1 + 1/2 + 1/3 + \cdots + 1/h$ is the h-th Harmonic number.

For a strategy profile $s \in S$ and $i \in \{1, 2, \ldots, n\}$, we denote by s^{-i} the strategy profile of the players $j \neq i$, i.e. $s^{-i} = (s^1, \ldots, s^{i-1}, s^{i+1}, \ldots s^n)$. We use $(s^{-i}, \bar{s}^i)$ to denote the strategy profile identical to s, except that the ith player uses strategy $\bar{s}^i$ instead of s^i. Similarly, for a subset of players Π_0, we define $s^{-\Pi_0}$, the profile of players $j \notin \Pi_0$. For $\sigma \in \times_{i \in \Pi_0} S^i$, we denote by $(s^{-\Pi_0}, \sigma)$ the strategy profile obtained from s by changing the strategies of players in Π_0 to σ.

A strategy profile $s \in S$ is a *Nash equilibrium* if no player $i \in \Pi$ can benefit from unilaterally deviating from his action to another action, i.e.,

$$\forall i \in \Pi \text{ and } \forall \bar{s}^i \in S^i,\ c^i(s^{-i}, \bar{s}^i) \geq c^i(s).$$

The crucial property of Rosenthal potential Φ is that each step performed by a player improving his payoff also decreases Φ. Consequently, if Φ admits a minimal value in strategy profile, this strategy profile is a Nash equilibrium.

Parameterized complexity. We briefly review the relevant concepts of parameterized complexity theory that we employ. For deeper background on the subject see the books by Downey and Fellows [7], Flum and Grohe [12], and Niedermeier [19].

In the classical framework of P *vs* NP, there is only one measurement (the overall input size) that frames the distinction between efficient and inefficient algorithms, and between tractable and intractable problems. Parameterized complexity is essentially a two-dimensional sequel, where in addition to the overall input size n, a secondary measurement k (the *parameter*) is introduced, with the aim of capturing the contributions to problem complexity due to such things as typical input structure, sizes of solutions, goodness of approximation, etc. Here, the parameter is deployed as a measurement of the amount of current solution modification allowed in a local search step. The parameter can also represent an aggregrate of such bounds.

The central concept in parameterized complexity theory is the concept of *fixed-parameter tractability* (FPT), that is solvability of the parameterized problem in time $f(k) \cdot n^{O(1)}$. The importance is that such a running time isolates all the exponential costs to a function of only the parameter.

The main hierarchy of parameterized complexity classes is

$$FPT \subseteq W[1] \subseteq W[2] \subseteq \cdots \subseteq W[P] \subseteq XP.$$

The formal definition of classes $W[t]$ is technical, and, in fact, irrelevant to the scope of this paper. For our purposes it suffices to say that a problem is in a class if it is FPT-reducible to a complete problem in this class. Given two parameterized problems Π and Π', an *FPT reduction* from Π to Π' maps an instance (I, k) of Π to an instance (I', k') of Π' such that

(1) $k' = h(k)$ for some computable function h,
(2) (I, k) is a YES-instance of Π if and only if (I', k') is a YES-instance of Π', and
(3) the mapping can be computed in FPT time.

Hundreds of natural problems are known to be complete for the aforementioned classes, and $W[1]$ is considered the parameterized analog of NP, because the k-Step Halting Problem for nondeterministic Turing machines of unlimited nondeterminism (trivially solvable by brute force in time $O(n^k)$) is complete for $W[1]$. Thus, the statement $FPT \neq W[1]$ serves as a plausible complexity assumption for proving intractability results in parameterized complexity. Independent Set, parameterized by solution size, is a more combinatorial example of a problem complete for $W[1]$. We refer the interested reader to the books by Downey and Fellows [7] or Flum and Grohe [12] for a more detailed introduction to the hierarchy of parameterized problems.

Local Search. Local search algorithms are among the most common heuristics used to solve computationally hard optimization problems. The common method of local search algorithms is to move from solution to solution by applying local changes. Books [1,18] provide a nice introduction to the wide area of local search. Best-response dynamics in congestion games corresponds to local search in 1-exchange neighbourhood minimizing Rosenthal potential Φ; improving moves for players decrease the value of the potential function. For strategy profiles $s_1, s_2 \in S$, we define the Hamming distance $D(s_1, s_2) = |s_1 \triangle s_2|$ between s_1 and s_2, that is the number of players implementing different strategies in s_1 and s_2. We study the following parameterized version of the local search problem for multicast.

We define *arena* as a directed graph G with root vertex r, a multiset of target vertices $t_1, \ldots, t_\ell$ and for every edge e of the graph a cost function $c_e : \mathbb{Z}^+ \to \mathbb{R}^+ \cup \{0\}$ such that $c_e(h) \geq c(h+1)$ for $h \geq 1$.

p-LOCAL SEARCH ON POTENTIAL Φ **Parameter:** k
Input: An arena consisting of graph G, vertices $r, (t_1, \ldots, t_\ell)$ and cost functions c_e, a strategy profile s, and an integer $k \geq 0$
Problem: Decide whether there is a strategy profile s' such that $\Phi(s') < \Phi(s)$ and $D(s, s') \leq k$, where Φ is as defined in (1).

3 Minimizing Rosenthal Potential

The aim of this section is to prove the following theorem.

Theorem 1. *The p-LOCAL SEARCH ON POTENTIAL Φ problem is solvable in time*

$$\binom{|\Pi|}{k} \cdot 3^k \cdot |G|^{O(1)}.$$

Let us remark that in particular, if Φ is Rosenthal's potential, and hence the cost functions are of the special type $c_e(h) = \frac{c_e}{h}$, the p-LOCAL SEARCH ON POTENTIAL Φ problems can be solved within the running time of Theorem 1.

We need some additional terminology. Let G be a directed graph. We say that a subdigraph T of G is an *out-tree* if T is a directed tree with only one vertex r of in-degree zero (called the *root*). The vertices of T of out-degree zero are called *leaves*. We also say that a strategy profile s^* is *optimal* if it gives the minimum value of the potential, i.e., for any other strategy profile s, $\Phi(s) \geq \Phi(s^*)$. Let $s = (P_1, \ldots, P_{|\Pi|})$ be a strategy profile and $C \geq 1$ be an integer. We say that s *uses C arcs* if the union T of the paths P_i consists of C arcs.

If edge-sharing is profitable, then we can make the following observation about the structure of optimal strategies.

Lemma 1. *Let C be an integer such that there is a strategy profile using at most C arcs. Let $s = (P_1, \ldots, P_{|\Pi|})$ be a strategy profile using at most C arcs such that*

(i) Among all profiles using at most C arcs, s is optimal. In other words, for any profile s' using at most C arcs, we have $\Phi(s') \geq \Phi(s)$.
(ii) Subject to (i), S uses the minimum number of arcs.

Then the union T of the paths P_i, $i \in \{1, \ldots, |\Pi|\}$, is an out-tree rooted in r.

Proof. Targeting towards a contradiction, let us assume that $T = \cup_{i=1}^{|\Pi|} P_i$ is not an out-tree. Then there are paths P_i, P_j, $i, j \in \{1, \ldots, |\Pi|\}$, that have a common vertex $v \neq r$ such that the (r, v)-subpaths P_i^v and P_j^v of P_i and P_j respectively enter v by different arcs.

We show first that

$$\sum_{e \in E(P_i^v)} c_e(n_e(s)) > \sum_{e \in E(P_j^v)} c_e(n_e(s)). \tag{2}$$

cannot occur. Assume that (2) holds. We claim that then the i-th player can improve his strategy and, consequently, Φ can be decreased, which will contradict the optimality of s. Denote by P the (r, t_i)-walk obtained from P_i by replacing path P_i^v by P_j^v. Notice that P is not necessarily a path. Let P' be a (r, t_i)-path in P and let us construct the new strategy profile $s' = (s^{-i}, P')$. This profile uses at most C arcs. By non-negativity of $c_e(h)$, the new cost for the i-th player is equal to

$$\begin{aligned}\sum_{e \in E(P')} c_e(n_e(s')) = & \sum_{e \in E(P') \cap E(P_i)} c_e(n_e(s)) + \sum_{e \in E(P') \setminus E(P_i)} c_e(n_e(s)+1) \leq \\ \leq & \sum_{e \in E(P) \cap E(P_i)} c_e(n_e(s)) + \sum_{e \in E(P) \setminus E(P_i)} c_e(n_e(s)+1).\end{aligned}$$

Since for each $e \in E$ and $h \geq 1$, we have $c_e(h) \geq c_e(h+1)$,

$$\sum_{e \in E(P) \setminus E(P_i)} c_e(n_e(s)+1) \leq \sum_{e \in E(P) \setminus E(P_i)} c_e(n_e(s)).$$

Therefore,

$$\sum_{e \in E(P')} c_e(n_e(s')) \leq \sum_{e \in E(P)} c_e(n_e(s)).$$

By (2), we have

$$\sum_{e \in E(P)} c_e(n_e(s)) < \sum_{e \in E(P_i)} c_e(n_e(s)),$$

and the claim that player i can improve follows.

Hence,

$$\sum_{e \in E(P_i^v)} c_e(n_e(s)) \leq \sum_{e \in E(P_j^v)} c_e(n_e(s)).$$

By the same arguments as above, we can replace P_j by a (r, t_j)-path P in the walk obtained from P_j by the replacement of P_j^v by P_i^v without increasing Φ. Moreover, we can repeat this operation for each path P_h, $h \neq i$, that enters v by an arc that is different from the arc in P_i. The number of arcs used by the paths in the modified strategy profiles is at most the number of arcs used in s, and thus is at most C. It remains to observe that we obtain a strategy profile where v has in-degree one in the union of paths. But it contradicts the choice of s, since we obtain a strategy profile that uses less arcs. Hence, T is an out-tree rooted in r. □

We use Lemma 1 to find an optimal strategy profile using the approach proposed by Dreyfus and Wagner [8] for the Steiner Tree problem.

Theorem 2. *Given an arena as input, the minimum value of a potential* Φ *can be found in time* $3^{|\Pi|} \cdot |G|^{O(1)}$. *The algorithm can also construct the corresponding optimal strategy profile* s^* *within the same time complexity.*

Proof. We give a dynamic programming algorithm. For simplicity, we only describe how to find the minimum of Φ, but it is straightforward to modify the algorithm to obtain the corresponding strategy profile.

Let $T = \{t_1, \ldots, t_{|\Pi|}\}$ be the multiset of terminals. We construct partial solutions for subsets $X \subseteq T$. Also, while at the end we are interested in the answer for the source r, our partial solutions are constructed for all vertices of G. For a vertex $u \in V(G)$ and a multiset $X \subseteq T$, let Γ_u^X denote the version of the game, in which only players associated with X build paths from u to their respective terminals. Therefore, we are interested in the game Γ_r^T. For a non-negative integer m, we define $\Psi(u, X, m)$ as the minimum value of the potential $\Phi(s)$ in the game Γ_u^X, taken over all strategy profiles s such that the union of paths in s contains at most m arcs (we say that s *uses* arc e if it is contained in some path from s). We assume that $\Psi(u, X, m) = +\infty$ if there are no feasible strategy profiles. Notice that by Lemma 1, the number of arcs used in an optimal strategy in the original problem is at most $|V(G)| - 1$. Hence, our aim is to compute $\Psi(r, T, |V(G)| - 1)$.

Clearly, $\Psi(u, \emptyset, m) = 0$ for all $u \in V$ and $m \geq 0$. For non-empty X and $m = 0$, $\Psi(u, X, m) = 0$ if all terminals in X are equal to u, and $\Psi(u, X, m) = +\infty$ otherwise. We need the following claim.

Claim. For $X \neq \emptyset$ and $m \geq 1$, $\Psi(u, X, m)$ satisfies the following equation:

$$\Psi(u, X, m) = \min\{\Psi(u, X, m-1), \\ \Psi(u, X \setminus Y, m_1) + \Psi(v, Y, m_2) + \sum_{h=1}^{|Y|} c_{(u,v)}(h)\}, \tag{3}$$

where the minimum is taken over all arcs $(u, v) \in E(G)$, $\emptyset \neq Y \subseteq X$, and $m_1, m_2 \geq 0$ such that $m_1 + m_2 = m - 1$; it is assumed that $\Psi(u, X, m) = \Psi(u, X, m-1)$ if the out-degree of u is zero.

Proof. Let $\psi = \min\{\Psi(u, X, m-1), \Psi(u, X \setminus Y, m_1) + \Psi(v, Y, m_2) + \sum_{h=1}^{|Y|} c_{(u,v)}(h)\}$. We prove that $\Psi(u, X, m) = \psi$ by first showing that $\Psi(u, X, m) \geq \psi$, and then that $\Psi(u, X, m) \leq \psi$. Without loss of generality assume that $X = \{t_1, \ldots, t_\ell\} \subseteq T$, where $\ell = |X|$.

If $\Psi(u, X, m) = +\infty$, then $\Psi(u, X, m) \geq \psi$. Suppose that $\Psi(u, X, m) \neq +\infty$ and consider a strategy $s^* = (P_1, \ldots, P_\ell)$ in the game Γ_u^X which is optimal among those using at most m arcs and, subject to this condition, the number of used arcs is minimum; in particular, s^* has potential $\Psi(u, X, m)$. By Lemma 1, $H = \cup_{i=1}^{\ell} P_i$ is an out-tree rooted in u. If $|E(H)| < m$, then $\Psi(u, X, m) =$

$\Psi(u, X, m-1) \geq \psi$. Assume that $|E(H)| = m$. As $m \geq 1$, vertex u has an out-neighbor v in H. Denote by H_1 and H_2 the components of $H - (u, v)$, where H_1 is an out-tree rooted in u and H_2 is an out-tree rooted in v. Let $Y \subseteq X$ be the multiset of terminals in H_2 and let $m_1 = |E(H_1)|$, $m_2 = |E(H_2)|$. Notice that exactly $|Y|$ players are using the arc (u, v) in s^* and Y is nonempty. Then $\Psi(u, X, m) \geq \Psi(u, X \setminus Y, m_1) + \Psi(v, Y, m_2) + \sum_{h=1}^{|Y|} c_{(u,v)}(h) \geq \psi$.

Now we prove that $\Psi(u, X, m) \leq \psi$. If $\psi = \Psi(u, X, m-1)$ then the claim is trivial, so let v, Y, m_1 and m_2 be such that $\psi = \Psi(u, X \setminus Y, m_1) + \Psi(v, Y, m_2) + \sum_{h=1}^{|Y|} c_{(u,v)}(h)$. Assume without loss of generality that $Y = \{t_1, \ldots, t_{\ell'}\}$ for some $\ell' \leq \ell$. If $\Psi(u, X \setminus Y, m_1) = +\infty$ or $\Psi(v, Y, m_2) = +\infty$, then the inequality is trivial. Suppose that $\Psi(u, X \setminus Y, m_1) \neq +\infty$ and $\Psi(v, Y, m_2) \neq +\infty$. Consider a strategy s_1^* in the game $\Gamma_u^{X \setminus Y}$ that is optimal among those using at most m_1 arcs, and a strategy s_2^* in the game Γ_v^Y that is optimal among those using at most m_2 arcs. Of course, the potential of s_1^* is equal to $\Psi(u, X \setminus Y, m_1)$, while the potential of s_2^* is equal to $\Psi(u, Y, m_2)$. We construct the strategy profile s in the game Γ_u^X as follows. For each terminal $t_j \in X \setminus Y$, the players use the (u, t_j)-path from s_1^*. For any $t_j \in Y$, the players use the (v, t_j)-path from s_2^* after accessing v from u via the arc (u, v), unless u already lies on this (v, t_j)-path, in which case they simply use the corresponding subpath of the (v, t_j)-path. Note that s uses at most $m_1 + m_2 + 1 = m$ arcs. Because for every $e \in E(G)$ and every $h \geq 1$, we have that $c_e(h) \geq 0$, and $c_e(h) \geq c_e(h+1)$, we infer that $\Phi(s) \leq \psi$, as possible overlapping of arcs used in s_1^*, s_2^* and the arc (u, v) can only decrease the potential of s. Since $\Psi(u, X, m) \leq \Phi(s)$, this implies that $\Psi(u, X, m) \leq \psi$. □

In order to finish the proof of Theorem 2, we need to observe that using the recurrence (3) one can compute the value $\Psi(r, T, |V(G)| - 1)$ in time $3^{|\Pi|} \cdot |G|^{O(1)}$. □

We use Theorem 2 to construct algorithm for p-Local Search on potential Φ and to conclude with the proof of Theorem 1.

Proof (of Theorem 1). Consider an instance of p-Local Search on potential Φ. Let $T = \{t_1, \ldots, t_{|\Pi|}\}$ be the multiset of terminals and let s be a strategy profile. Recall that p-Local Search on potential Φ asks whether at most k players can change their strategies in such a way that the potential decreases. Observe that we can assume that *exactly* k players are going to change their strategies because some of these players can choose their old strategies. There are $\binom{|\Pi|}{k}$ possibilities to choose a set of k players $\Pi_0 \subseteq \Pi$. We consider all possible choices and for each set Π_0, we check whether the players from this set can apply some strategy to decrease Φ.

Denote by $X \subseteq T$ the multiset of terminals of the players from Π_0, and let $s' = s^{-\Pi_0}$. We compute the potential $\Phi(s')$ for this strategy profile. Now we redefine the cost of edges as follows: for each $e \in E(G)$ and $h \geq 1$, $c'_e(h) = c_e(n_e(s') + h)$. Clearly, $c'_e(h) \geq 0$ and $c'_e(h) \geq c'_e(h+1)$. Let Φ' be the potential for these edge costs. We find the minimum value of $\Phi'(s^*)$ for the set of players Π_0 and the corresponding terminals X. It remains to observe that $\Phi(s') + \Phi'(s^*) = \min\{\Phi(s'') \mid s'' = (s^{-\Pi_0}, \sigma),\ \sigma \in \prod_{i \in \Pi_0} S^i\}$. By Theorem 2, we can find $\Phi'(s^*)$ in time $3^k \cdot |G|^{O(1)}$ and the claim follows. □

4 Intractability of Local Search for Rosenthal Potential

This section is devoted to the proof of the following theorem. Due to space restriction, the most technical part of the proof, i.e., the proof of the completeness of the reduction, will appear in the full version of the paper.

Theorem 3. p-LOCAL SEARCH ON POTENTIAL Φ*, where* Φ *is Rosenthal potential for multicasting game, is W[1]-hard.*

Proof. (Construction and Soundness) We provide an FPT reduction from the MULTICOLOURED CLIQUE problem, which is known to be W[1]-hard [11].

MULTICOLOURED CLIQUE	**Parameter:** k
Input: An undirected graph H with vertices partitioned into k sets $V_1, V_2, \dots, V_k$, such that H does not contain edges connecting vertices from the same part V_i.	
Problem: Is there a clique C in G of size k?	

Observe that by the assumption on the structure of H, the clique C has to contain exactly one vertex from every part V_i.

We take an instance (H, k) of MULTICOLOURED CLIQUE and construct an instance $(G, s, k(k-1))$ of p-LOCAL SEARCH ON POTENTIAL Φ. First, we provide the construction of the new instance; then, we discuss its soundness and completeness. During the reduction we assume k to be large enough; for constant k we solve the instance (H, k) in polynomial time by a brute-force search and output a trivial YES or NO instance of p-LOCAL SEARCH ON POTENTIAL Φ.

Construction. First create the root vertex r. For every $u \in V_i$, we create k vertices: $\overline{u}$ and $u_1, \dots, u_{i-1}, u_{i+1}, \dots, u_k$. Denote $F_u = \{u_1, \dots, u_{i-1}, u_{i+1}, \dots, u_k\}$. We connect them in the following manner: we construct one arc $(r, \overline{u})$ with cost $R = k^2$, and for all $j \in \{1, 2, \dots, i-1, i+1, \dots, k\}$ we construct arc $(\overline{u}, u_j)$ with cost 0. With every vertex u_j for all $u \in V(H)$ we associate a player that builds a path from r to u_j. In the initial strategy profile s, each of $(k-1)|V(H)|$ players builds a path that leads to his vertex via the corresponding vertex $\overline{u}$. Observe that the potential of this strategy profile is equal to $|V(H)| \cdot R \cdot \mathcal{H}(k-1)$.

We now construct the part of the graph that is responsible for the choice of the clique. We create a *pseudo-root* r' and an arc (r, r') with cost

$$W = \frac{1}{\mathcal{H}(k(k-1))} \left(k \cdot R \cdot \mathcal{H}(k-1) - \frac{3}{2}\binom{k}{2} - \varepsilon \right),$$

where $\varepsilon = \frac{k-1}{k^5}$. Note that $W \geq 1$ for sufficiently large k. For every edge $uv \in E(H)$, where $u \in V_i$ and $v \in V_j$, $i \neq j$, we create a vertex x_{uv}, arc (r', x_{uv}) of cost 1, and arcs (x_{uv}, u_j), (x_{uv}, v_i) of cost 0. This concludes the construction.

Before we proceed with the formal proof of the theorem, let us give some intuition behind the construction. From a clique C in H, we can derive a common strategy of $k(k-1)$ players assigned to vertices from $\bigcup_{u \in V(C)} F_u$, who can agree

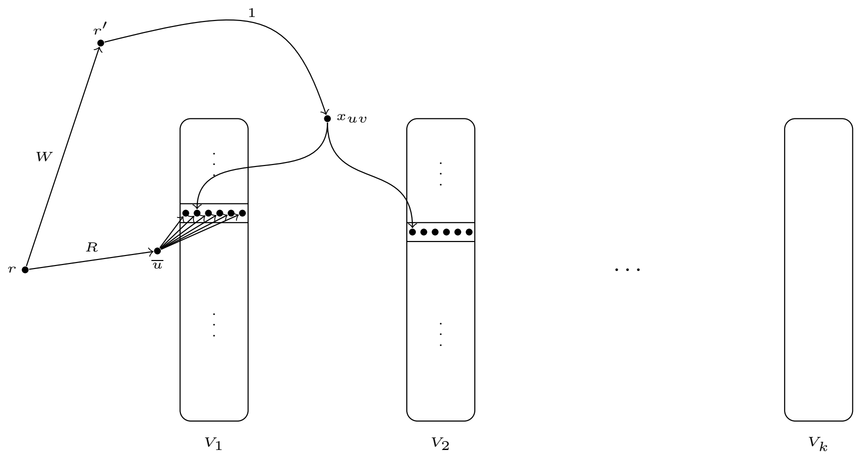

Fig. 1. Graph G

to jointly rebuild their paths via the pseudo-root r'. The "cost of entrance" for remodelling the strategy in this manner is paying for the expensive arc (r, r'); however, this can amortised by sharing cheap arcs (r', x_{uv}) for $uv \in E(C)$. The costs have been chosen so that only the maximum possibility of sharing, which corresponds to a clique in H, can yield a decrease of the potential.

Soundness. Assume that C is a clique in H with k vertices. Let us remind, that in the initial strategy profile s each player is using the corresponding arc $(r, \overline{u})$ for his path. We want to show that H contains a clique of size k. We construct the new strategy profile s' by changing strategies of $k(k-1)$ players as follows. For every $uv \in E(C)$, where $u \in V_i$ and $v \in V_j$, $i \neq j$, the players associated with vertices u_j and v_i reroute their paths so that in s' they lead via r' and x_{uv} to respective targets. In comparison to the profile s, the new profile s':

- has congestion withdrawn from arcs $(r, \overline{u})$ for $u \in V(C)$—this decreases the potential by $k \cdot R \cdot \mathcal{H}(k-1)$;
- has congestion introduced to arcs (r, r') and (r', x_{uv}) for $uv \in E(C)$—this increases the potential by $W \cdot \mathcal{H}(k(k-1)) + \frac{3}{2}\binom{k}{2}$.

Therefore, $\Phi(s') = \Phi(s) - k \cdot R \cdot \mathcal{H}(k-1) + W \cdot \mathcal{H}(k(k-1)) + \frac{3}{2}\binom{k}{2} = \Phi(s) - \varepsilon < \Phi(s)$.

The proof of the completeness of the reduction is deferred to the full version of the paper due to space constraints. □

References

1. Aarts, E.H.L., Lenstra, J.K.: Local Search in Combinatorial Optimization. Princeton University Press (1997)
2. Albers, S.: On the value of coordination in network design. SIAM J. Comput. 38, 2273–2302 (2009)

3. Anshelevich, E., Dasgupta, A., Kleinberg, J.M., Tardos, É., Wexler, T., Roughgarden, T.: The price of stability for network design with fair cost allocation. SIAM J. Comput. 38, 1602–1623 (2008)
4. Anshelevich, E., Dasgupta, A., Tardos, É., Wexler, T.: Near-optimal network design with selfish agents. Theory of Computing 4, 77–109 (2008)
5. Chekuri, C., Chuzhoy, J., Lewin-Eytan, L., Naor, J., Orda, A.: Non-cooperative multicast and facility location games. IEEE Journal on Selected Areas in Communications 25, 1193–1206 (2007)
6. Chen, H.-L., Roughgarden, T., Valiant, G.: Designing network protocols for good equilibria. SIAM J. Comput. 39, 1799–1832 (2010)
7. Downey, R.G., Fellows, M.R.: Parameterized complexity. Springer, New York (1999)
8. Dreyfus, S.E., Wagner, R.A.: The Steiner problem in graphs. Networks 1, 195–207 (1972)
9. Epstein, A., Feldman, M., Mansour, Y.: Strong equilibrium in cost sharing connection games. In: Proceedings 8th ACM Conference on Electronic Commerce (EC 2007), pp. 84–92. ACM (2007)
10. Fellows, M., Fomin, F.V., Lokshtanov, D., Rosamond, F., Saurabh, S., Villanger, Y.: Local search: Is brute-force avoidable? In: Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI 2009), pp. 486–491. AAAI (2009)
11. Fellows, M.R., Hermelin, D., Rosamond, F.A., Vialette, S.: On the parameterized complexity of multiple-interval graph problems. Theor. Comput. Sci. 410, 53–61 (2009)
12. Flum, J., Grohe, M.: Parameterized Complexity Theory. Springer, Berlin (2006)
13. Gupta, A., Kumar, A., Pál, M., Roughgarden, T.: Approximation via cost sharing: Simpler and better approximation algorithms for network design. J. ACM 54, 11 (2007)
14. Gupta, A., Srinivasan, A., Tardos, É.: Cost-sharing mechanisms for network design. Algorithmica 50, 98–119 (2008)
15. Kleinberg, J., Tardos, E.: Algorithm design. Addison-Wesley, Boston (2005)
16. Koutsoupias, E., Papadimitriou, C.H.: Worst-case equilibria. Computer Science Review 3, 65–69 (2009)
17. Marx, D.: Local search. Parameterized Complexity Newsletter 3, 7–8 (2008)
18. Michiels, W., Aarts, E.H.L., Korst, J.: Theoretical Aspects of Local Search. Springer (2007)
19. Niedermeier, R.: Invitation to fixed-parameter algorithms. Oxford University Press (2006)
20. Rosenthal, R.W.: A class of games possessing pure-strategy Nash equilibria. Internat. J. Game Theory 2, 65–67 (1973)
21. Roughgarden, T., Sundararajan, M.: Quantifying inefficiency in cost-sharing mechanisms. J. ACM 56 (2009)
22. Roughgarden, T., Tardos, É.: How bad is selfish routing? J. ACM 49, 236–259 (2002)
23. Szeider, S.: The parameterized complexity of k-flip local search for sat and max sat. Discrete Optimization 8, 139–145 (2011)