Network-Formation Games with Regular Objectives

Guy Avni[1], Orna Kupferman[1], and Tami Tamir[2]

[1] School of Computer Science and Engineering, The Hebrew University, Jerusalem, Israel
[2] School of Computer Science, The Interdisciplinary Center, Herzliya, Israel

Abstract. Classical network-formation games are played on a directed graph. Players have reachability objectives, and each player has to select a path satisfying his objective. Edges are associated with costs, and when several players use the same edge, they evenly share its cost. The theoretical and practical aspects of network-formation games have been extensively studied and are well understood. We introduce and study *network-formation games with regular objectives*. In our setting, the edges are labeled by alphabet letters and the objective of each player is a regular language over the alphabet of labels, given by means of an automaton or a temporal-logic formula. Thus, beyond reachability properties, a player may restrict attention to paths that satisfy certain properties, referring, for example, to the providers of the traversed edges, the actions associated with them, their quality of service, security, etc.

Unlike the case of network-formation games with reachability objectives, here the paths selected by the players need not be simple, thus a player may traverse some transitions several times. Edge costs are shared by the players with the share being proportional to the number of times the transition is traversed. We study the existence of a pure Nash equilibrium (NE), convergence of best-response-dynamics, the complexity of finding the social optimum, and the inefficiency of a NE compared to a social-optimum solution. We examine several classes of networks (for example, networks with uniform edge costs, or alphabet of size 1) and several classes of regular objectives. We show that many properties of classical network-formation games are no longer valid in our game. In particular, a pure NE might not exist and the Price of Stability equals the number of players (as opposed to logarithmic in the number of players in the classic setting, where a pure NE always exists). In light of these results, we also present special cases for which the resulting game is more stable.

1 Introduction

Network design and formation is a fundamental well-studied problem that involves many interesting combinatorial optimization problems. In practice, network design is often conducted by multiple strategic users whose individual costs are affected by the decisions made by others. Early works on network design focus on analyzing the efficiency and fairness properties associated with different sharing rules (e.g., [23,30]). Following the emergence of the Internet, there has been an explosion of studies employing game-theoretic analysis to explore Internet applications, such as routing in computer networks and network formation [17,1,12,2]. In network-formation games (for a survey, see [35]), the network is modeled by a weighted graph. The weight of an edge indicates the cost of activating the transition it models, which is independent of the number of times the edge is used. Players have reachability objectives, each given by sets of possible source and target nodes. Players share the cost of edges used in order to fulfill their objectives. Since the costs are positive, the runs traversed by the players are simple. Under the common Shapley cost-sharing mechanism, the cost of an edge is shared evenly by the players that use it.

The players are selfish agents who attempt to minimize their own costs, rather than to optimize some global objective. In network-design settings, this would mean that the players selfishly select a path instead of being assigned one by a central authority. The focus in game theory is on the *stable* outcomes of a given setting, or the *equilibrium* points. A Nash equilibrium (NE) is a profile of the players' strategies such that no player can decrease his cost by an unilateral deviation from his current strategy, that is, assuming that the strategies of the other players do not change.[3]

Reachability objectives enable the players to specify possible sources and targets. Often, however, it is desirable to refer also to other properties of the selected paths. For example, in a *communication* setting, edges may belong to different providers, and a user may like to specify requirements like "all edges are operated by the same provider" or "no edge operated by AT&T is followed by an edge operated by Verizon". Edges may also have different quality or security levels (e.g., "noisy channel", "high-bandwidth channel", or "encrypted channel"), and again, users may like to specify their preferences with respect to these properties. In *planning* or in *production systems*, nodes of the network correspond to configurations, and edges correspond to the application of actions. The objectives of the players are sequences of actions that fulfill a certain plan, which is often more involved than just reachability [21]; for example "once the arm is up, do not put it down until the block is placed".

The challenge of reasoning about behaviors has been extensively studied in the context of formal verification. While early research concerned the input-output relations of terminating programs, current research focuses on on-going behaviors of reactive systems [22]. The interaction between the components of a reactive system correspond to a multi-agent game, and indeed in recent years we see an exciting transfer of concepts and ideas between the areas of game theory and formal verification: logics for specifying multi-agent systems [3,9], studies of equilibria in games that correspond to the synthesis problem [8,7,16], an extension of mechanism design to on-going behaviors [25], studies of non-zero-sum games in formal methods [10,6], and more.

In this paper we extend network-formation games to a setting in which the players can specify regular objectives. This involves two changes of the underlying setting: First, the edges in the network are labeled by letters from a designated alphabet. Second, the objective of each player is specified by a *language* over this alphabet. Each player should select a path labeled by a word in his objective language. Thus, if we view the network as a *nondeterministic weighted finite automaton* (WFA) $\mathcal{A}$, then the set of strategies for a player with objective L is the set of accepting runs of $\mathcal{A}$ on some word in L. Accordingly, we refer to our extension as *automaton-formation games*. As in classical network-formation games, players share the cost of edges they use. Unlike

[3] Throughout this paper, we concentrate on pure strategies and pure deviations, as is the case for the vast literature on cost-sharing games.

the classical game, the runs selected by the players need not be simple, thus a player may traverse some edges several times. Edge costs are shared by the players, with the share being proportional to the number of times the edge is traversed. This latter issue is the main technical difference between automaton-formation and network-formation games, and as we shall see, it is very significant.

Many variants of cost-sharing games and congestion games have been studied. A generalization of the network-formation game of [2] in which players are weighted and a player's share in an edge cost is proportional to its weight is considered in [11], where it is shown that the weighted game does not necessarily have a pure NE. In a different type of congestion games, players' payments depend on the resource they choose to use, the set of players using this resource, or both [29,26,27,19]. In some of these variants a NE is guaranteed to exist while in others it is not. All these variants are different from automaton-formation games, where a player needs to select a *multiset* of resources (namely, the edges he is going to traverse) rather than a single one.

We study the theoretical and practical aspects of automaton-formation games. In addition to the general game, we consider classes of instances that have to do with the network, the specifications, or to their combination. Recall that the network can be viewed as a WFA $\mathcal{A}$. We consider the following classes of WFAs: (1) *all-accepting*, in which all the states of $\mathcal{A}$ are accepting, thus its language is prefix closed (2) *uniform costs*, in which all edges have the same cost, and (3) *single letter*, in which $\mathcal{A}$ is over a single-letter alphabet. We consider the following classes of specifications: (1) *single word*, where the language of each player is a single word, (2) *symmetric*, where all players have the same objective. We also consider classes of instances that are intersections of the above classes.

Each of the restricted classes we consider corresponds to a real-life variant of the general setting. Let us elaborate below on single-letter instances. The language of an automaton over a single letter $\{a\}$ induces a subset of $\mathbb{N}$, namely the numbers $k \in \mathbb{N}$ such that the automaton accepts a^k. Accordingly, single-letter instances correspond to settings in which a player specifies possible lengths of paths. Several communication protocols are based on the fact that a message must pass a pre-defined length before reaching its destination. This includes *onion routing*, where the message is encrypted in layers [33], or *proof-of-work* protocols that are used to deter denial of service attacks and other service abuses such as spam (e.g., [15]).

We provide a complete picture of the following questions for various instances (for formal definitions, see Section 2): (i) Existence of a *pure Nash equilibrium*. That is, whether each instance of the game has a profile of pure strategies that constitutes a NE. As we show, unlike the case of classical network design games, a pure NE might not exist in general automaton-formation games and even in very restricted instances of it. (ii) The complexity of finding the *social optimum* (SO). The SO is a profile that minimizes the total cost of the edges used by all players; thus the one obtained when the players obey some centralized authority. We show that for some restricted instances finding the SO can be done efficiently, while for other restricted instances, the complexity agrees with the NP-completeness of classical network-formation games. (iii) An analysis of *equilibrium inefficiency*. It is well known that decentralized decision-making may lead to solutions that are sub-optimal from the point of view of society as a whole. We quantify the inefficiency incurred due to selfish behavior according to the *price of anarchy* (PoA) [24,31] and *price of stability* (PoS) [2] measures. The PoA is the worst-case inefficiency of a Nash equilibrium (that is, the ratio between the worst NE and the SO). The PoS is the best-case inefficiency of a Nash equilibrium (that is, the ratio between the best NE and the SO). We show that while the PoA in automaton-formation games agrees with the one in classical network-formation games and is equal to the number of players, the PoS also equals the number of players, again already in very restricted instances. This is in contrast with classical network-formation games, where the PoS tends to log the number of players. Thus, the fact that players may choose to use edges several times significantly increases the challenge of finding a stable solution as well as the inefficiency incurred due to selfish behavior. We find this as the most technically challenging result of this work. We do manage to find structural restrictions on the network with which the social optimum is a NE.

The technical challenge of our setting is demonstrated in the seemingly easy instance in which all players have the same objective. Such *symmetric* instances are known to be the simplest to handle in all cost-sharing and congestion games studied so far. Specifically, in network-formation games, the social optimum in symmetric instances is also a NE and the PoS is 1. Moreover, in some games [18], computing a NE is PLS-complete in general, but solvable in polynomial time for symmetric instances. Indeed, once all players have the same objective, it is not conceivable that a player would want to deviate from the social-optimum solution, where each of the k players pays $\frac{1}{k}$ of the cost of the optimal solution. We show that, surprisingly, symmetric instances in AF-games are not simple at all. Specifically, the social optimum might not be a NE, and the PoS is at least $\frac{k}{k-1}$. In particular, for symmetric two-player AF games, we have that $PoS = PoA = 2$. We also show that the PoA equals the number of players already for very restricted instances.

2 Preliminaries

2.1 Automaton-formation games

A *nondeterministic finite weighted automaton* on finite words (WFA, for short) is a tuple $\mathcal{A} = \langle \Sigma, Q, \Delta, q_0, F, c \rangle$, where Σ is an alphabet, Q is a set of states, $\Delta \subseteq Q \times \Sigma \times Q$ is a transition relation, $q_0 \in Q$ is an initial state, $F \subseteq Q$ is a set of accepting states, and $c : \Delta \to \mathbb{R}$ is a function that maps each transition to the cost of its formation [28]. A *run* of $\mathcal{A}$ on a word $w = w_1, \ldots, w_n \in \Sigma^*$ is a sequence of states $\pi = \pi^0, \pi^1, \ldots, \pi^n$ such that $\pi^0 = q_0$ and for every $0 \leq i < n$ we have $\Delta(\pi^i, w_{i+1}, \pi^{i+1})$. The run π is *accepting* iff $\pi^n \in F$. The *length* of π is n, whereas its size, denoted $|\pi|$, is the number of different transitions in it. Note that $|\pi| \leq n$.

An *automaton-formation game* (AF game, for short) between k selfish players is a pair $\langle \mathcal{A}, O \rangle$, where $\mathcal{A}$ is a WFA over some alphabet Σ and O is a k-tuple of regular languages over Σ. Thus, the objective of Player i is a regular language L_i, and he needs to choose a word $w_i \in L_i$ and an accepting run of $\mathcal{A}$ on w_i in a way that minimizes his payments. The cost of each transition is shared by the players that use it in their selected runs, where the share of a player in the cost of a transition e is proportional to the number of times e is used by the player. Formally, The set of strategies for Player i

is $\mathcal{S}_i = \{\pi : \pi$ is an accepting run of $\mathcal{A}$ on some word in $L_i\}$. We assume that $\mathcal{S}_i$ is not empty. We refer to the set $\mathcal{S} = \mathcal{S}_1 \times \ldots \times \mathcal{S}_k$ as the set of *profiles* of the game.

Consider a profile $P = \langle \pi_1, \pi_2, \ldots, \pi_k \rangle$. We refer to π_i as a sequence of transitions. Let $\pi_i = e_i^1, \ldots, e_i^{\ell_i}$, and let $\eta_P : \Delta \to \mathbb{N}$ be a function that maps each transition in Δ to the number of times it is traversed by all the strategies in P, taking into an account several traversals in a single strategy. Denote by $\eta_i(e)$ the number of times e is traversed in π_i, that is, $\eta_i(e) = |\{1 \le j \le \ell_i : e_i^j = e\}|$. Then, $\eta_P(e) = \sum_{i=1\ldots k} \eta_i(e)$. The *cost of player i in the profile P* is

$$cost_i(P) = \sum_{e \in \pi_i} \frac{\eta_i(e)}{\eta_P(e)} c(e). \tag{1}$$

For example, consider the WFA $\mathcal{A}$ depicted in Fig. 1. The label $e_1 : a, 1$ on the transition from q_0 to q_1 indicates that this transition, which we refer to as e_1, traverses the letter a and its cost is 1. We consider a game between two players. Player 1's objective is the language is $L_1 = \{ab^i : i \ge 2\}$ and Player 2's language is $\{ab, ba\}$. Thus, $\mathcal{S}_1 = \{\{e_1, e_2, e_2\}, \{e_1, e_2, e_2, e_2\}, \ldots\}$ and $\mathcal{S}_2 = \{\{e_3, e_4\}, \{e_1, e_2\}\}$. Consider the profile $P = \langle \{e_1, e_2, e_2\}, \{e_3, e_4\} \rangle$, the strategies in P are disjoint, and we have $cost_1(P) = 2 + 2 = 4, cost_2(P) = 1 + 3 = 4$. For the profile $P' = \langle \{e_1, e_2, e_2\}, \{e_1, e_2\} \rangle$, it holds that $\eta_1(e_1) = \eta_2(e_1)$ and $\eta_1(e_2) = 2 \cdot \eta_2(e_2)$. Therefore, $cost_1(P') = \frac{1}{2} + 2 = 2\frac{1}{2}$ and $cost_2(P') = \frac{1}{2} + 1 = 1\frac{1}{2}$.

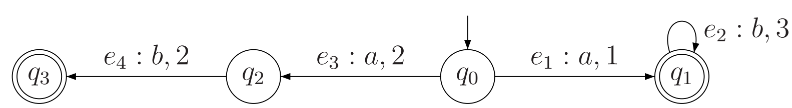

Fig. 1. An example of a WFA.

We consider the following instances of AF games. Let $G = \langle \mathcal{A}, O \rangle$. We start with instances obtained by imposing restrictions on the WFA $\mathcal{A}$. In *one-letter* instances, $\mathcal{A}$ is over a singleton alphabet, i.e., $|\Sigma| = 1$. When depicting such WFAs, we omit the letters on the transitions. In *all-accepting* instances, all the states in $\mathcal{A}$ are accepting; i.e., $F = Q$. In *uniform-costs* instances, all the transitions in the WFA have the same cost, which we normalize to 1. Formally, for every $e \in \Delta$, we have $c(e) = 1$. We continue to restrictions on the objectives in O. In *single-word* instances, each of the languages in O consists of a single word. In *symmetric* instances, the languages in O coicide, thus the players all have the same objective. We also consider combinations on the restrictions. In particular, we say that $\langle \mathcal{A}, O \rangle$ is *weak* if it is one-letter, all states are accepting, costs are uniform, and objectives are single words. Weak instances are simple indeed – each player only specifies a length of a path he should patrol, ending anywhere in the WFA, where the cost of all transitions is the same. As we shall see, many of our hardness results and lower bounds hold already for the class of weak instances.

2.2 Nash equilibrium, social optimum, and equilibrium inefficiency

For a profile P, a strategy π_i for Player i, and a strategy π, let $P[\pi_i \leftarrow \pi]$ denote the profile obtained from P by replacing the strategy for Player i by π. A profile $P \in \mathcal{S}$ is a *pure Nash equilibrium* (NE) if no player i can benefit from unilaterally deviating from his run in P to another run; i.e., for every player i and every run $\pi \in \mathcal{S}_i$ it holds that $cost_i(P[\pi_i \leftarrow \pi]) \geq cost_i(P)$. In our example, the profile P is not a NE, since Player 2 can reduce his payments by deviating to profile P'.

The (social) cost of a profile P, denoted $cost(P)$, is the sum of costs of the players in P. Thus, $cost(P) = \sum_{1 \leq i \leq k} cost_i(P)$. Equivalently, if we view P as a set of transitions, with $e \in P$ iff there is $\pi \in P$ for which $e \in \pi$, then $cost(P) = \sum_{e \in P} c(e)$. We denote by OPT the cost of an optimal solution; i.e., $OPT = \min_{P \in \mathcal{S}} cost(P)$. It is well known that decentralized decision-making may lead to sub-optimal solutions from the point of view of society as a whole. We quantify the inefficiency incurred due to self-interested behavior according to the *price of anarchy* (PoA) [24,31] and *price of stability* (PoS) [2] measures. The PoA is the worst-case inefficiency of a Nash equilibrium, while the PoS measures the best-case inefficiency of a Nash equilibrium. Formally,

Definition 1. *Let $\mathcal{G}$ be a family of games, and let $G \in \mathcal{G}$ be a game in $\mathcal{G}$. Let $\Upsilon(G)$ be the set of Nash equilibria of the game G. Assume that $\Upsilon(G) \neq \emptyset$.*

- *The* price of anarchy *of G is the ratio between the* maximal *cost of a NE and the social optimum of G. That is, $PoA(G) = \max_{P \in \Upsilon(G)} cost(P)/OPT(G)$. The* price of anarchy *of the family of games $\mathcal{G}$ is $PoA(\mathcal{G}) = sup_{G \in \mathcal{G}} PoA(G)$.*
- *The* price of stability *of G is the ratio between the* minimal *cost of a NE and the social optimum of G. That is, $PoS(G) = \min_{P \in \Upsilon(G)} cost(P)/OPT(G)$. The* price of stability *of the family of games $\mathcal{G}$ is $PoS(\mathcal{G}) = sup_{G \in \mathcal{G}} PoS(G)$.*

Uniform Sharing rule: A different cost-sharing rule that could be adopted for automaton-formation games is the uniform sharing rule, according to which the cost of a transition e is equally shared by the players that traverse e, independent of the number of times e is traversed by each player. Formally, let $\kappa_P(e)$ be the number of runs that use the transition e at least once in a profile P. Then, the cost of including a transition e at least once in a run is $c(e)/\kappa_P(e)$. This sharing rule induces a potential game, where the potential function is identical to the one used in the analysis of the classical network design game [2]. Specifically, let $\Phi(P) = \sum_{e \in E} c(e) \cdot H(\kappa_P(e))$, where $H_0 = 0$, and $H_k = 1 + 1/2 + \ldots + 1/k$. Then, $\Phi(P)$ is a potential function whose value reduces with every improving step of a player, thus a pure NE exists and BRD is guaranteed to converge[4]. The similarity with classical network-formation games makes the study of this setting straightforward. Thus, throughout this paper we only consider the proportional sharing rule as defined in (1) above.

[4] Best-response-dynamics (BRD) is a local-search method where in each step some player is chosen and plays his best-response strategy, given that the strategies of the other players do not change.

3 Properties of Automaton-Formation Games

In this section we study the theoretical properties of AF games: existence of NE and equilibrium inefficiency. We show that AF games need not have a pure Nash equilibrium. This holds already in the very restricted class of weak instances, and is in contrast with network-formation games. There, BRD converges and a pure NE always exists. We then analyze the PoS in AF games and show that there too, the situation is significantly less stable than in network-formation games.

Theorem 1. *Automaton-formation games need not have a pure NE. This holds already for the class of weak instances.*

Proof. Consider the WFA $\mathcal{A}$ depicted in Fig. 2 and consider a game with $k = 2$ players. The language of each player consists of a single word. Recall that in one-letter instances we care only about the lengths of the objective words. Let these be ℓ_1 and ℓ_2, with $\ell_1 \gg \ell_2 \gg 0$ that are multiples of 12. For example, $\ell_1 = 30000, \ell_2 = 300$. Let C_3 and C_4 denote the cycles of length 3 and 4 in $\mathcal{A}$, respectively. Let D_3 denote the path of length 3 from q_0 to q_1. Every run of $\mathcal{A}$ consists of some repetitions of these cycles possibly with one pass on D_3.

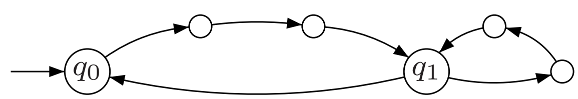

Fig. 2. A weak instance of AF games with no NE.

We claim that no pure NE exists in this instance. Since we consider long runs, the fact that the last cycle might be partial is ignored in the calculations below. We first show that the only candidate runs for Player 1 that might be part of a NE profile are $\pi_1 = (C_4)^{\frac{\ell_1}{4}}$ and $\pi'_1 = D_3 \cdot (C_3)^{\frac{\ell_1}{3}-1}$. If Player 1 uses both C_3 and C_4 multiple times, then, given that $\ell_1 \gg \ell_2$, he must almost fully pay for at least one of these cycles, thus, deviating to the run that repeats this fully-paid cycle is beneficial.

When Player 1 plays π_1, Player 2's best response is $\pi_2 = (C_4)^{\frac{\ell_2}{4}}$. In the profile $\langle \pi_1, \pi_2 \rangle$, Player 1 pays almost all the cost of C_4, so the players' costs are $(4 - \varepsilon, \varepsilon)$. This is not a NE. Indeed, since $\ell_2 \gg 0$, then by deviating to π'_1, the share of Player 1 in D_3 reduces to almost 0, and the players' costs in $\langle \pi'_1, \pi_2 \rangle$, are $(3 + \varepsilon, 4 - \varepsilon)$. This profile is not a NE as Player 2's best response is $\pi'_2 = D_3 \cdot (C_3)^{\frac{\ell_2}{3}-1}$. Indeed, in the profile $\langle \pi'_1, \pi'_2 \rangle$, the players' costs are $(4.5 - \varepsilon, 1.5 + \varepsilon)$ as they share the cost of D_3 and Player 1 pays almost all the cost of C_3. This is not a NE either, as Player 1 would deviate to the profile $\langle \pi_1, \pi'_2 \rangle$, in which the players' costs are $(4-\varepsilon, 3+\varepsilon)$. The latter is still not a NE, as Player 2 would head back to $\langle \pi_1, \pi_2 \rangle$. We conclude that no NE exists in this game. □

The fact a pure NE may not exist is a significant difference between standard cost-sharing games and AF games. The bad news do not end here and extend to equilibrium inefficiency. We first note that the cost of any NE is at most k times the social optimum (as otherwise, some player pays more than the cost of the SO and can benefit from migrating to his strategy in the SO). Thus, it holds that $PoS \leq PoA \leq k$. The following theorem shows that this is tight already for highly restricted instances.

Theorem 2. *The PoS in* AF *games equals the number of players. This holds already for the class of weak instances.*

Proof. We show that for every $k, \delta > 0$ there exists a simple game with k players for which the PoS is more than $k - \delta$. Given k and δ, let r be an integer such that $r > \max\{k, \frac{k-1}{\delta} - 1\}$. Consider the WFA $\mathcal{A}$ depicted in Fig. 3. Let $L = \langle \ell_1, \ell_2, \ldots, \ell_k \rangle$ for $\ell_2 = \ldots = \ell_k$ and $\ell_1 \gg \ell_2 \gg r$ denote the lengths of the objective words. Thus, Player 1 has an 'extra-long word' and the other $k - 1$ players have words of the same, long, length. Let C_r and C_{r+1} denote, respectively, the cycles of length r and $r + 1$ to the right of q_0. Let D_r denote the path of length r from q_0 to q_1, and let D_{kr} denote the 'lasso' consisting of the kr-path and the single-edge loop to the left of q_0.

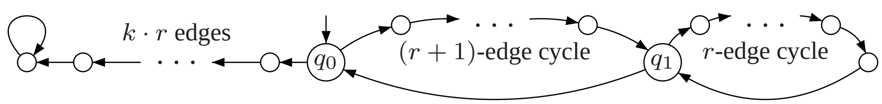

Fig. 3. A weak instance of AF games for which $PoS = k$.

The social optimum of this game is to buy C_{r+1}. Its cost is $r + 1$. However, as we show, the profile P in which all players use D_{kr} is the only NE in this game. We first show that P is a NE. In this profile, Player 1 pays $r + 1 - \varepsilon$ and each other player pays $r + \varepsilon/(k - 1)$. No player will deviate to a run that includes edges from the right side of $\mathcal{A}$. Next, we show that P is the only NE of this game: Every run on the right side of $\mathcal{A}$ consists of some repetitions of C_{r+1} and C_r, possibly with one traversal of D_r. Since we consider long runs, the fact that the last cycle might be partial is ignored in the calculations below.

In the social optimum profile, Player 1 pays $r + 1 - \varepsilon$ and each of the other players pays $\varepsilon/(k - 1)$. The social optimum is not a NE as Player 1 would deviate to $D_r \cdot C_r^*$ and will reduce his cost to $r + \varepsilon'$. The other players, in turn, will also deviate to $D_r \cdot C_r^*$. In the profile in which they are all selecting a run of the form $D_r \cdot C_r^*$, Player 1 pays $r + r/k - \varepsilon > r + 1$ and prefers to return to C_{r+1}^*. The other players will join him sequentially, until the non-stable social optimum is reached. Thus, no NE that uses the right part of $\mathcal{A}$ exists. Finally, it is easy to see that no run that involves edges from both the left and right sides of $\mathcal{A}$ or includes both C_{r+1} and C_r can be part of a NE.

The cost of the NE profile is $kr+1$ and the PoS is therefore $\frac{kr+1}{r+1} = k - \frac{k-1}{r+1} > k - \delta$. □

4 Computational Complexity Issues in AF Games

In this section we study the computational complexity of two problems: finding the cost of the social optimum and finding the best-response of a player. Recall that the social optimum (SO) is a profile that minimizes the total cost the players pay. It is well-known that finding the social optimum in a network-formation game is NP-complete. We show that this hardness is carried over to simple instances of AF games. On the positive side, we identify non-trivial classes of instances, for which it is possible to compute the SO efficiently. The other issue we consider is the complexity of finding the best strategy of a single player, given the current profile, namely, the best-response of a player. In network-formation games, computing the best-response reduces to a shortest-path problem, which can be solved efficiently. We show that in AF games, the problem is NP-complete.

The proofs of the following theorems can be found in the appendix. The reductions we use are from the set-cover problem, where choice of sets are related to choice of transitions.

Theorem 3. *Finding the value of the social optimum in* AF *games is NP-complete. Moreover, finding the social optimum is NP-complete already in single-worded instances that are also uniform-cost and are either single-lettered or all-accepting.*

The hardness results in Theorem 3 for single-word specification use one of two properties: either there is more than one letter, or not all states are accepting. We show that finding the SO in instances that have both properties can be done efficiently, even for specifications with arbitrary number of words.

For a language L_i over $\Sigma = \{a\}$, let $short(i) = \min_j \{a^j \in L_i\}$ denote the length of the shortest word in L_i. For a set O of languages over $\Sigma = \{a\}$, let $\ell_{max}(O) = \max_i short(i)$ denote the length of the longest shortest word in O. Clearly, any solution, in particular the social optimum, must include a run of length $\ell_{max}(O)$. Thus the cost of the social optimum is at least the cost of the cheapest run of length $\ell_{max}(O)$. Moreover, since the WFA is single-letter and all-accepting, the other players can choose runs that are prefixes of this cheapest run, and no additional transitions should be acquired. We show that finding the cheapest such run can be done efficiently.

Theorem 4. *The cost of the social optimum in a single-letter all-accepting instance $\langle \mathcal{A}, O \rangle$ is the cost of the cheapest run of length $\ell_{max}(O)$. Moreover, this cost can be found in polynomial time.*

We turn to prove the hardness of finding the best-response of a player. Our proof is valid already for a single player that needs to select a strategy on a WFA that is not used by other players (one-player game).

Theorem 5. *Finding the best-response of a player in* AF *games is NP-complete.*

5 Tractable Instances of AF Games

In the example in Theorem 1, Player 1 deviates from a run on the shortest (and cheapest) possible path to a run that uses a longer path. By doing so, most of the cost of the original path, which is a prefix of the new path and accounts to most of its cost, goes to Player 2. We consider *semi-weak* games in which the WFA is uniform-cost, all-accepting, and single-letter, but the objectives need not be a single word. We identify a property of such games that prevents this type of deviation and which guarantees that the social optimum is a NE. Thus, we identify a family of AF games in which a NE exists, finding the SO is easy, and the PoS is 1.

Definition 2. *Consider a semi-weak game $\langle \mathcal{A}, O \rangle$. A* lasso *is a path $u \cdot v$, where u is a simple path that starts from the initial state and v is a simple cycle. A lasso ν is minimal in $\mathcal{A}$ if $\mathcal{A}$ does not have shorter lassos. Note that for minimal lassos $u \cdot v$, we have that $u \cap v = \emptyset$. We say that $\mathcal{A}$ is* resistant *if it has no cycles or there is a minimal lasso $\nu = u \cdot v$ such that for every other lasso ν' we have $|u \setminus \nu'| + |v| \leq |\nu' \setminus \nu|$.*

Consider a resistant weak game $\langle \mathcal{A}, O \rangle$. In order to prove that the social optimum is a NE, we proceed as follows. Let ν be the lasso that is the witness for the resistance of $\mathcal{A}$. We show that the profile S^* in which all players choose runs that use only the lasso ν or a prefix of it, is a NE. The proof is technical and we go over all the possible types of deviations for a player and use the weak properties of the network along with its resistance. By Theorem 4, the cost of the profile is the SO. Hence the following. The full proof can be found in Appendix A.

Theorem 6. *For resistent semi-weak games, the social optimum is a NE.*

A corollary of Theorem 6 is the following:

Corollary 1. *For resistant semi-weak games, we have PoS= 1.*

We note that resistance can be defined also in WFAs with non-uniform costs, with $cost(\nu)$ replacing $|\nu|$. Resistance, however, is not sufficient in the *slightly* stronger model where the WFA is single-letter and all-accepting but not uniform-cost. Indeed, given k, we show a such a game in which the PoS is kx, for a parameter x that can be arbitrarily close to 1. Consider the WFA A in Fig. 5. Note that $\mathcal{A}$ has a single lasso and is thus a resistant WFA. The parameter ℓ_1 is a function of x, and the players' objectives are single words of lengths $\ell_1 \gg \ell_2 \gg \ldots \gg \ell_k \gg 0$. Similar to the proof of Theorem 2, there is only one NE in the game, which is when all players choose the left chain. The social optimum is attained when all players use the self-loop, and thus for a game in this family, $PoS = \frac{k \cdot x}{1}$. Since x tends to 1, we have $PoS = k$ for resistant all-accepting single-letter games. The proof can be found in the full version.

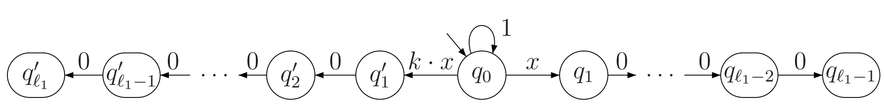

Fig. 4. A resistant all-accepting single-letter game in which the PoS tend to k.

6 Surprises in Symmetric Instances

In this section we consider the class of symmetric instances, where all players share the same objective, that is, there exists a language L, such that for all $1 \leq i \leq k$, we have $L_i = L$. In such instances it is tempting to believe that the social optimum is also a NE, as all players evenly share the cost of the solution that optimizes their common objective. While this is indeed the case in all known symmetric games, we show that, surprisingly, this is not the case for AF-games, in fact already for the class of one-letter, all accepting, unit-cost and single word instances.

To demonstrate the anomaly, let us first consider the two-player game appearing in Fig. 5. All the states in the WFA $\mathcal{A}$ are accepting, and the objectives of both players is a single long word. The social optimum is when both players traverse the loop q_0, q_1, q_0. Its cost is $2 + \epsilon$, so each player pays $1 + \frac{\epsilon}{2}$. This, however, is not a NE, as Player 1 (or, symmetrically, Player 2) prefers to deviate to the run $q_0, q_1, q_1, q_1, \ldots$, where he pays the cost of the loop q_1, q_1 and his share in the transition from q_0 to q_1. We can choose the length of the objective word and ϵ so that this share is smaller than $\frac{\epsilon}{2}$, justifying his deviation. Note that the new situation is not a NE either, as Player 2, who now pays 2, is going to join Player 1, resulting in an unfortunate NE in which both players pay 1.5.

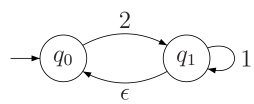

Fig. 5. The WFA $\mathcal{A}$ for which the SO in a symmetric game is not a NE.

Before we continue to the example where the PoS is at least $\frac{k}{k-1}$, let us elaborate on the PoA. It is easy to see that in symmetric AF games, we have $PoA = k$. This bound is achieved, as in the classic network-formation game, by a network with two parallel edges labeled by a and having costs k and 1. The players all have the same specification $L = \{a\}$. The profile in which all players select the expensive path is a NE. We now show that $PoA = k$ is achieved even for weak symmetric instances.

Theorem 7. *The PoA equals the number of players, already for weak symmetric instances.*

Proof. We show a lower bound of k. The example is a generalization of the PoA in cost sharing games [2]. For k players, consider the weak instance depicted in Fig. 6, where all players have the length k. Intuitively, the social optimum is attained when all players use the loop $\langle q_0, q_0 \rangle$ and thus $OPT = 1$. The worst NE is when all players use the run $q_0 q_1 \ldots q_k$, and its cost is clearly k. Formally, there are two NEs in the game:

- The cheap NE is when all players use the loop $\langle q_0, q_0 \rangle$. This is indeed a NE because if a player deviates, he must buy at least the transition $\langle q_0, q_1 \rangle$. Thus, he pays at least 1, which is higher than $\frac{1}{k}$, which is what he pays when all players use the loop.
- The expensive NE is when all players use the run $q_0, q_1, \ldots, q_k$. This is a NE because a player has two options to deviate. Either to the run that uses only the loop, which costs 1, or to a run that uses the loop and some prefix of $q_0, q_1, \ldots, q_k$, which costs at least $1 + \frac{1}{k}$. Since he currently pays 1, he has no intention of deviating to either runs.

Since the cheap NE costs 1 and the expensive one costs k, we get $PoA = k$. $\square$

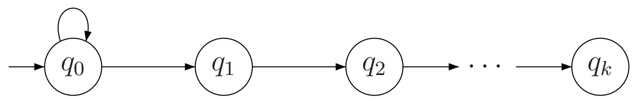

Fig. 6. The WFA $\mathcal{A}$ for which a symmetric game with $|L| = 1$ achieves $PoA = k$.

A more surprising result is that $PoS > 1$. We show a lower bound of $\frac{k}{k-1}$ for the PoS. Thus, for two-player games, we have $PoS = PoA = 2$.

Theorem 8. *In a symmetric k-player game, the PoS is at least $\frac{k}{k-1}$.*

Proof. For $k \geq 2$, we describe a family of symmetric games for which the PoS tends to $\frac{k}{k-1}$. For $n \geq 1$, the game $G_{\epsilon,n}$ uses the WFA that is depicted in Figure 7. Note that this is a one-letter instance in which all states are accepting. The players have an identical specification, consisting of a single word w of length $\ell \gg 0$. We choose ℓ and $\epsilon = \epsilon_0 > \ldots > \epsilon_{n-1}$ as follows. Let $C_0, \ldots, C_n$ denote, respectively, the cycles with costs $(k^n + \epsilon_0), (k^{n-1} + \epsilon_1), \ldots, (k + \epsilon_{n-1}), 1$. Let $r_0, \ldots, r_n$ be lasso-runs on w that end in $C_0, \ldots, C_n$, respectively. For every $1 \leq i < n$, consider the profile P_i in which all players choose run r_i. Since all the players need to select an accepting run for the word w, they share the cost of the transitions in r_i equally. Thus, the cost of a player in P_i is $\frac{1}{k} \cdot (\sum_{0 \leq j \leq i} k^{n-j} + \epsilon_i)$. Specifically, for the cycle C_i, the players each pay $\frac{k^{n-i}+\epsilon_i}{k} = k^{n-(i+1)} + \frac{\epsilon_i}{k}$. Let P'_i be the profile in which Player 1 deviates from P_i to r_{i+1}. In P'_i, Player 1 pays the same amount for the path leading to C_i, but his share of the k^{n-i}-valued transition decreases drastically. He now pays for it only $\frac{k^{n-i}}{(k-1)\cdot(\ell-i)/2}$. On the other hand, he now pays the full price for C_{i+1}, which is $k^{n-(i+1)} + \epsilon_{i+1}$ or 1 if $i = n-1$. We choose ϵ_i, ϵ_{i+1}, and ℓ so that $cost_1(P'_i) < cost_1(P_i)$. Also, we choose these values so that players $2, \ldots, k-1$, prefer joining Player 1 in r_{i+1} rather than staying with r_i.

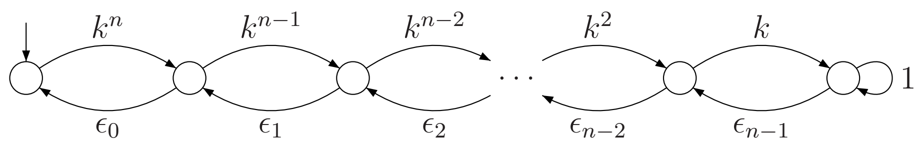

Fig. 7. The network of the identical-specification game $G_{\epsilon,n}$, in which PoS tends to $\frac{k}{k-1}$.

We claim that the only NE is when all players use the run r_n. Indeed, it is not hard to see that every profile in which a player selects a run that is not from $r_0, \ldots, r_n$ cannot be a NE. Also, a profile in which two players select runs r_i and r_j, for $1 \leq i < j \leq n$, cannot be a NE as the player using r_i can decreases his payment by joining the other player in r_j. Finally, by our selection of $\epsilon_1, \ldots, \epsilon_n$, and ℓ, every profile in which all the players choose the run r_i, for $0 \leq i \leq n-1$, is not a NE.

Clearly, the social optimum is attained when all players choose the run r_0, and its cost is $k^n + \epsilon$. Since the cost of the only NE in the game is $\sum_{0 \leq i \leq n} k^{n-i}$, the PoS in this family of games tends to $\frac{k}{k-1}$ as n grows to infinity and ϵ to 0. $\square$

Finally, we note that our hardness result in Theorem 5 implies that finding the social optimum in a symmetric AF-game is NP-complete. Indeed, since the social optimum is the cheapest run on some word in L, finding the best-response in a one-player game is equivalent to finding the social optimum in a symmetric game. This is contrast with other cost-sharing and congestion game (e.g. [18], where the social optimum in symmetric games can be computed using a reduction to max-flow).

7 Conclusions and Future Work

Our results on the stability of AF games are mostly negative. We identified some stable cases and we believe that additional positive results can be derived for restricted classes of instances. As we suggest below, these restrictions can be characterized by the structure of the automaton or by the set of players' objectives.

Ordinary open problems include the study of approximate-NE, networks with profits, capacitated networks, and coordinated deviation. We highlight below several interesting directions for future work that are specific to the study of AF games.

1. Our lower bounds use WFAs with cycles. We believe that for acyclic all-accepting one-letter instances, the PoS for can be bounded by a constant. Specifically, for k players, we conjecture that $PoS = \sum_{i=1}^{k} \frac{1}{2^{i-1}}$, which is bounded by 2. In Appendix B we present a lower bound of this value that is valid already for automata consisting of disjoint paths. Such an analysis will provide a nice distinction between the classical network-formation game, for which $PoS = \Theta(\log k)$, and our game, even when all players use a *simple* path for their run. We note that it is possible to restrict the class of languages in the objectives so that the players have no incentive not to use simple paths for their runs. For example, when the languages are *closed under infix disposal* (that is, if $x \cdot y \cdot z \in L$, for $x, y, z \in \Sigma^*$, then $x \cdot z \in L$).
2. Other presumably more stable games are those in which the range of costs or the ratio between the maximal and the minimal transition costs is bounded, or when the ratio between the longest and the shortest word in the objective languages is bounded. Indeed, bounding these ratios also bounds the proportion in which costs are shared, making the game closer to one with a uniform sharing rule.
3. AF-games are an example of cost-sharing games in which players' strategies are *multisets* of resources. In such games, a player may need multiple uses of the same resource, and his share in the resource cost is proportional to the number of times he uses the resource. Our results imply that, in general, such games are less stable than classical cost-sharing games. It is desirable to study more settings of such games, and to characterize non-trivial instances that arise in practice and for which the existence of pure NE can be shown, and its inefficiency can be bounded. In the context of formal method, an appealing application is that of *synthesis from components*, where the resources are functions from a library, and agents need to

synthesize their objectives using such functions, possibly by a repeated use of some functions.

4. For symmetric AF games, we leave open the problem of NE existence as well as the problem of finding an upper-bound for the PoS for $k > 2$.

Recall that in planning, the WFA models a production system in which transitions correspond to actions. In such cases, the objectives of the players may be languages of *infinite* words, describing desired on-going behaviors. The objectives can be specified by linear temporal logic or nondeterministic Büchi automata, and each player has to select a lasso computation or accepting run for a word in his language. The setting of infinite words involves transitions that are taken infinitely often and calls for new sharing rules. When the sharing rule refers to the frequency in which transitions are taken, we obtain a proportional sharing rule that is similar to the one studied here. One can also follow a sharing rule in which all players that traverse a transition infinitely often share its cost evenly, perhaps with some favorable proportion towards players that use it only finitely often. This gives rise to simpler sharing rules, which seem more stable.

Acknowledgments. We thank Michal Feldman, Noam Nisan, and Michael Schapira for helpful discussions.

References

1. S. Albers, S. Elits, E. Even-Dar, Y. Mansour, and L. Roditty. On Nash Equilibria for a Network Creation Game. In *Proc. 17th SODA*, pages 89-98, 2006.
2. E. Anshelevich, A. Dasgupta, J. Kleinberg, É. Tardos, T. Wexler, and T. Roughgarden. The Price of Stability for Network Design with Fair Cost Allocation. *SIAM J. Comput.* 38(4): 1602–1623, 2008.
3. R. Alur, T.A. Henzinger, and O. Kupferman. Alternating-time temporal logic. *Journal of the ACM*, 49(5):672–713, 2002.
4. B. Aminof, O. Kupferman, and R. Lampert, Reasoning about online algorithms with weighted automata, *ACM Transactions on Algorithms*, 6(2), 2010.
5. B. Alpern and F.B. Schneider. Recognizing safety and liveness. *Distributed computing*, 2:117–126, 1987.
6. T. Brihaye, V. Bruyère, J. De Pril, and H. Gimbert. On subgame perfection in quantitative reachability games. *Logical Methods in Computer Science*, 9(1), 2012.
7. K. Chatterjee. Nash equilibrium for upward-closed objectives. In *Proc. 15th CSL*, LNCS 4207, pages 271–286. Springer, 2006.
8. K. Chatterjee, T. A. Henzinger, and M. Jurdzinski. Games with secure equilibria. *Theoretical Computer Science*, 365(1-2):67–82, 2006.
9. K. Chatterjee, T. A. Henzinger, and N. Piterman. Strategy logic. In *Proc. 18th CONCUR*, pages 59–73, 2007.
10. K. Chatterjee, R. Majumdar, and M. Jurdzinski. On Nash equilibria in stochastic games. In *Proc. 13th CSL*, LNCS 3210, pages 26–40. Springer, 2004.
11. H. Chen and T. Roughgarden. Network Design with Weighted Players, *Theory of Computing Systems*, 45(2), 302–324, 2009.
12. J. R. Correa, A. S. Schulz, and N. E. Stier-Moses. Selfish Routing in Capacitated Networks. *Mathematics of Operations Research* 29: 961–976, 2004.

13. N. Daniele, F. Guinchiglia, and M.Y. Vardi. Improved automata generation for linear temporal logic. In *Proc. 11th CAV*, LNCS 1633, pages 249–260. Springer, 1999.
14. M. Droste, W. Kuich, and H. Vogler (eds.), Handbook of Weighted Automata, Springer, 2009.
15. C. Dwork and M. Naor. Pricing via Processing or Combatting Junk Mail, In *Proc. 12th CRYPTO*, pages 139–177, 1992.
16. D. Fisman, O. Kupferman, and Y. Lustig. Rational synthesis. In *Proc. 16th TACAS*, LNCS 6015, pages 190–204. Springer, 2010.
17. A. Fabrikant, A. Luthra, E. Maneva, C. Papadimitriou, and S. Shenker. On a network creation game. In *Proc. 22nd PODC*, pages 347-351, 2003.
18. A. Fabrikant, C. Papadimitriou, and K. Talwarl, The Complexity of Pure Nash Equilibria, *Proc. 36th STOC*, pages 604–612, 2004.
19. M. Feldman and T. Tamir. Conflicting Congestion Effects in Resource Allocation Games. Journal of *Operations Research* 60(3), pages 529–540, 2012.
20. P. von Falkenhausen and T. Harks. Optimal Cost Sharing Protocols for Scheduling Games. In *Proc. 12th EC*, pages 285-294, 2011.
21. G. de Giacomo and M. Y. Vardi. Automata-Theoretic Approach to Planning for Temporally Extended Goals, In *European Conferences on Planning*, pages 226–238, 1999.
22. D. Harel and A. Pnueli. On the development of reactive systems. In *Logics and Models of Concurrent Systems*, volume F-13 of *NATO Advanced Summer Institutes*, pages 477–498. Springer, 1985.
23. S. Herzog, S. Shenker, and D. Estrin. Sharing the "Cost" of Multicast Trees: An Axiomatic Analysis. *IEEE/ACM Transactions on Networking*, 1997.
24. E. Koutsoupias and C. Papadimitriou. Worst-case Equilibria. *Computer Science Review*,3(2): 65–69, 2009.
25. O. Kupferman and T. Tamir. Coping with selfish on-going behaviors. *Information and Computation*, 210:1–12, 2012.
26. M. Mavronicolas, I. Milchtaich, B. Monien, and K. Tiemann. Congestion Games with Player-specific Constants. In *Proc 32nd MFCS*, pp. 633–644, 2007.
27. I. Milchtaich. Weighted Congestion Games With Separable Preferences. *Games and Economic Behavior*, 67, 750-757, 2009.
28. M. Mohri. Finite-state transducers in language and speech processing. *Computational Linguistics*, 23(2):269–311, 1997.
29. D. Monderer and L. Shapley. Potential Games. *Games and Economic Behavior*, 14:124–143, 1996.
30. H. Moulin and S. Shenker. Strategyproof Sharing of Submodular Costs: Budget Balance Versus Efficiency. *Journal of Economic Theory*, 18: 511–533, 2001.
31. C. Papadimitriou. Algorithms, Games, and the Internet. In *Proc 33rd STOC*, pages 749–753, 2001.
32. R. Paes Leme, V. Syrgkanis, E. Tardos. The curse of simultaneity. *Innovations in Theoretical Computer Science (ITCS), pages 60-67, 2012.*
33. *M. G. Reed, P. F. Syverson, and D. M. Goldschlag. Anonymous Connections and Onion Routing* IEEE J. on Selected Areas in Communication*, Issue on Copyright and Privacy Protection, 1998.*
34. *R. W. Rosenthal. A Class of Games Possessing Pure-Strategy Nash Equilibria.* International Journal of Game Theory*, 2: 65–67, 1973.*
35. *E. Tardos and T. Wexler. Network Formation Games and the Potential Function Method, In* Algorithmic Game Theory*, Cambridge University Press, 2007.*
36. *B. Vöcking.* In N. Nisan, T. Roughgarden, E. Tardos and V. Vazirani, eds., Algorithmic Game Theory. Chapter 20: Selfish Load Balancing. *Cambridge University Press, 2007.*

A Missing Proofs

Proof of Theorem 3: We start with membership in NP. Given a WFA $\mathcal{A}$ with objectives $w_1, \ldots, w_k$ and value $c \in \mathbb{R}$, we can guess a witness profile P and check whether it satisfies $cost(P) \leq c$ in polynomial time. For proving hardness, we show a reduction from the *Set-Cover* (SC) problem. Consider an input $\langle U, S, m\rangle$ to SC. Recall that U is a set of elements, $S = \{C_1, \ldots, C_z\} \subseteq 2^U$ is a collection of subsets of elements of U, and $m \in \mathbb{N}$. Then, $\langle U, S, m\rangle$ is in SC iff there is a subset S' of S of size at most m that covers U. That is, $|S'| \leq m$ and $\bigcup_{C \in S'} C = U$.

Given an input $\langle U, S, m\rangle$ to SC, we construct a uniform-cost single-letter WFA $\mathcal{A}$ and a vector of k integers, where the i-th integer corresponds to the length of the (single) word in L_i. We fix a value y, such that $\langle U, S, m\rangle$ in SC iff the SO value of the game played on $\mathcal{A}$ with the objectives in $\{L_i\}$ is y. We construct $\mathcal{A} = \langle \{a\}, Q, q_0, \Delta, \{q_{acc}\}, c\rangle$ as follows (see an example in the left of Fig. 8). The set Q includes the initial and accepting states, a state for every set in S, and intermediate states required for the disjoint runs defined below. Without loss of generality, we assume that $U = \{1, \ldots, k\}$. Consider an element $i \in U$. For every $C \in S$ such that $i \in C$, there is a disjoint run of length i from C to q_{acc}. Also, for every $C \in S$, there is a transition $\langle q_0, C\rangle$ in Δ. The cost of all transitions in Δ is 1. For every $1 \leq i \leq k$, the length of the word in $|L_i|$ is $i+1$. We define $w = m + (1+2+\ldots+k)$. The size of $\mathcal{A}$ is clearly polynomial in $|U|$ and $|S|$.

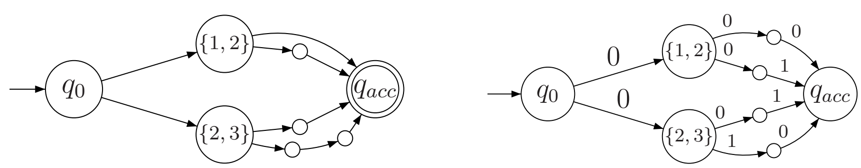

Fig. 8. The WFAs produced by the reduction for $U = \{1, 2, 3\}$ and $S = \{\{1, 2\}, \{2, 3\}\}$.

The construction for uniform-cost all-accepting instances is very similar (see an example in the right of Fig. 8). Let $z = \lceil \log(n) \rceil$ and $\Sigma = \{0, 1\}$. For $C \in S$ and $i \in C$, we have a z-length path from C to q_{acc} that is labeled with the binary representation of $i - 1$ (padded with preceding zeros if needed). The label on all transitions from q_0 to the S states is 0. For $1 \leq i \leq k$, the word for Player i is a single 0 letter followed by the binary representation of $i - 1$. The size of $\mathcal{A}$ is clearly polynomial in $|U|$ and $|S|$. The following claim completes the proof.

Claim. There exists a set-cover of size m iff $OPT \leq m + (1 + 2 + \ldots + k)$ for the uniform-cost single-letter instance and $OPT \leq m + k \cdot z$ for the uniform-cost all-accepting instance.

Proof. We prove the reduction for the uniform-cost single-letter instance. The proof for the uniform-cost all-accepting instance is very similar. For the first direction, let $S' = \{s_{i_1}, \ldots, s_{i_m}\}$ be a set cover. We show a profile $P = \{\pi_1, \ldots, \pi_k\}$ such that

$cost(P) \leq m+(1+2+\ldots+k)$. Recall that the input length for Player i is $i+1$. Since S' is a set cover, there is a set $s \in S'$ with $i \in s$. We define the run π_i to proceed from q_0 to s and from there to q_{acc} on a run of length i. Clearly, the runs $\pi_1, \ldots, \pi_n$ are all legal-accepting runs. Moreover, the runs use m transitions from $\{q_0\} \times S \subseteq E$. Thus, $cost(P) \leq m + (1+2+\ldots+k)$, implying $OPT \leq m + (1+2+\ldots+k)$, and we are done.

For the second direction, assume $OPT = m' + (1+2+\ldots+k) \leq m + (1+2+\ldots+k)$, we prove that there is a set cover of size m'. Let $S^* = \langle \pi_1, \ldots, \pi_k \rangle$. Thus, $OPT = cost(S^*) = m'$. Let $S' \subset S$ be such that $s \in S'$ iff the transition $\langle q_0, s \rangle$ is used in one of the runs in S^*. Note that the run of every player consists of a transition (q_0, s) followed by a disjoint run of length i to q_{acc}. Therefore, $OPT = m' + (1+2+\ldots+k)$, and, $|S'| = m' \leq m$. We claim that S' is a set cover. For every $i \in U$, the first transition in the run is a transition $\langle q_0, s \rangle$ for some $s \in S$, as otherwise, player i can not proceed to q_{acc} along a run of length i. By our definition of S' we have $s \in S'$, thus $i \in U$ is covered.

□

Proof of Theorem 4: Clearly, any solution, in particular the social optimum, must include a run of length $\ell_{max}(O)$. Thus the cost of the social optimum is at least the cost of the cheapest run of length $\ell_{max}(O)$. Moreover, since there are no target vertices, the other players can choose runs that are prefixes of the cheapest run, and no additional transitions should be acquired.

We claim that finding the cheapest such run can be done efficiently. Recall that q_0 is the initial state in $\mathcal{A}$, and let $|Q| = n$. We view $\mathcal{A}$ as a weighted-directed graph $G = \langle V, E, c \rangle$, where the vertices V are the states Q, there is an edge $e \in E$ between two vertices if there is a transition between the two corresponding states, and the cost of the edges is the same as the cost of the transition in $\mathcal{A}$. For $0 \leq i \leq n$, let $d_i : V \times V \to \mathbb{Q}^+$ be the function that, given two vertices $u, v \in V$, returns the value of the cheapest pah of length i from u to v, and ∞ if no such path exists. Note that there is no requirement that the path is simple, and indeed we may traverse cycles in order to accommodate i transitions. The function $d : V \times V \to \mathbb{Q}^+$, returns the value of the cheapest path of any length between two given vertices. Given two vertices $u, v \in V$, computing $d(u, v)$ can be done using Dijkstra's algorithm, and, given an index $i \in \mathbb{N}$, it is possible to compute $d_i(u, v)$ by a slight variation of the Bellman-Ford algorithm.

We distinguish between two cases. If $\ell_{max} > 2n-2$, we claim that the value of the social optimum is $min\{d(q_0, v) + d(v, v) : v \in V\}$. If $\ell_{max} \leq 2n-2$, then we claim that the value of the social optimum is the minimum value of $d_i(q_0, v) + d_j(v, v)$, where $v \in V$, $0 \leq i \leq \ell_{max}$, $0 \leq j \leq \ell_{max} - i$, and if $j = 0$, then $i = \ell_{max}$.

We start with the first case. Assume $\ell_{max} > 2n-2$. Let $ALG = min\{d(q_0, v) + d(v, v) : v \in V\}$. Recall that S^* is the social optimum profile, and $OPT = cost(S^*)$. For the first direction, we claim that $ALG \leq OPT$. Let π be a run in S^* of length ℓ_{max}, where we assume π is a sequence of transitions. Clearly, $OPT \geq cost(\pi)$. Since ALG takes the minimum over all vertices, it suffices to prove that $cost(\pi) \geq d(q_0, v) + d(v, v)$ for some $v \in V$. We view π as a path in the graph G, and we claim that π contains a sub-path that starts in q_0 and ends in v and a sub-path that is a cycle from v to itself, for some $v \in V$. Thus, $OPT \geq cost(\pi) \geq cost(x) + cost(y) \geq d(q_0, v) + d(v, v) \geq ALG$.

We continue to prove the claim. Since $\ell_{max} > n$, there is a vertex v that appears twice in π. We split π into two paths, at the first appearance of v. That is, $\pi = x \cdot y'$, where x is a path that ends in v and v does not appear in x again. Note that if $v = q_0$, then $\pi = y'$. Since π is a legal run, it starts in q_0, and we have that x is a path from q_0 to v. We continue to prove that there is a cycle y from v to itself that is contained in y'. Indeed, since v appears at least twice in π, and since y' is a sequence of transitions that starts in v, we have that v appears in y' at least twice, and we are done.

We continue to prove that $ALG \geq OPT$. Let $v \in V$ be the vertex that attains the minimum in $min\{d(q_0, v) + d(v, v) : v \in V\}$. Let $\tau = \tau_1 \cdot \tau_2$ be a run such that τ_1 is a simple path from q_0 to v with $cost(\tau_1) = d(q_0, v)$ and τ_2 is a simple cycle from v to itself with $cost(\tau_2) = d(v, v)$. We claim that $cost(\tau) \geq OPT$. Since τ_1 and τ_2 are simple, we have $|\tau_1| \leq n - 1$ and $|\tau_2| \leq n - 1$. Thus, $|\tau| < 2n - 2$. We extend τ to a path of length ℓ_{max} by traversing the loop τ_2 many times. Clearly, τ is a legal run of the automaton $\mathcal{A}$ on a word of length ℓ_{max}. Consider the profile S in which the players choose runs that are prefixes of τ'. Since the only transitions used in S are those in τ, we have $cost(S) = cost(\tau)$. Since S^* is the social optimum, we have $ALG = cost(S) \geq cost(S^*) = OPT$, and we are done.

The case in which $\ell_{max} \leq 2n - 2$ is proven in a similar manner. $\square$

Proof of Theorem 5: We start with membership in NP. Given a WFA $\mathcal{A}$ with objectives $L_1, \ldots, L_k$ and value $c \in \mathbb{R}$, we can guess a witness profile P and check whether it satisfies $cost(P) \leq c$ in polynomial time.

For proving hardness, we show a reduction from the *Set-Cover* (SC) problem. Consider an input $\langle U, S, m\rangle$ to SC. Recall that $U = \{1, \ldots, n\}$ is a set of elements, $S = \{C_1, \ldots, C_z\} \subseteq 2^U$ is a collection of subsets of elements of U, and $m \in \mathbb{N}$. Then, $\langle U, S, m\rangle$ is in SC iff there is a subset S' of S of size at most m that covers U. That is, $|S'| \leq m$ and $\bigcup_{C \in S'} C = U$.

Given an input $\langle U, S, m\rangle$ to SC, we construct a game $\langle \mathcal{A}, O\rangle$ such that $\langle U, S, m\rangle$ is in SC iff the SO in the game is at most l. The game is a one-player game. We start by describing the specification L of the player. The alphabet of L is $S \cup U$ and it is given by the regular expression $(C_1 + \ldots + C_m) \cdot 1 \cdot (C_1 + \ldots + C_m) \cdot 2 \cdot \ldots \cdot (C_1 + \ldots + C_m) \cdot n$. The WFA $\mathcal{A}$ is over the alphabet $S \cup U$. There is a single initial state q_{in} and a state for every set in S. For $1 \leq i \leq z$, there is a C_i-labeled transition from q_{in} to the state C_i, and for every $j \in C_i$, there is a j-labeled transition from the state C_i back to q_{in}. The first type of transitions cost 1 and the second cost 0 (for an example see Fig. 9).

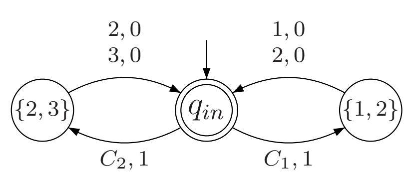

Fig. 9. The WFA produced by the reduction for $U = \{1, 2, 3\}$ and $S = \{\{1, 2\}, \{2, 3\}\}$.

We prove the correctness of the reduction: For the first direction, assume there is a set cover of at most l. Consider the word w in which, for every $1 \leq j \leq n$, the letter that precedes j is $C_i \in S$ such that C_i is in the set cover. Clearly, $w \in L$ and since it uses at most l letters from S, the profile in which the player chooses it, costs at most l. Thus, the SO is also at most l. For the other direction, assume the SO is attained in a profile with the word $w \in L$. It is not hard to see that the letters from S that appear in w form a set cover of size at most l. □

Proof of Theorem 6: Consider a resistant semi-weak game $\langle \mathcal{A}, O \rangle$. Since $\mathcal{A}$ is resistant, it has no cycles or there is a minimal lasso in $\mathcal{A}$ that satisfies the resistance requirements. By Theorem 4, the social optimum is attained when all players use prefixes of the cheapest run of length $\ell_{max}(O)$. For $1 \leq i \leq k$, let ℓ_i be the minimal length of a word in L_i, and assume, without loss of generality, that $\ell_1 \geq \ell_2 \geq \ldots \geq \ell_k$. That is, $\ell_1 = \ell_{max}(O)$. Let $S^* = \langle \pi_1, \ldots, \pi_k \rangle$, where for $1 \leq i \leq k$, the run π_i is of length ℓ_i and π_1 uses the lasso that is the witness for resistence, or an acyclic path if the lasso's length is larger than ℓ_1.

Claim. For every $1 \leq i \leq k$, let π'_i be a run of length ℓ_i. Then, $cost_i(S^*) \leq cost_i(S[i \leftarrow \pi'_i])$.

We show that the claim implies that S^* is a NE, and thus it implies the theorem. Indeed, assume towards contradiction that S^* is not a NE. Let $1 \leq i \leq k$ and π'_i, such that the best response for Player i from S^* is to deviate to π'_i. If the length of π'_i is ℓ_i, then we reach a contradiction to the claim. Otherwise, by deviating to a prefix of π'_i of length ℓ_i, Player i only improves his payment, which is again a contradiction.

We continue to prove the claim. Assume towards contradiction that S^* is not a NE. Thus, there is an index $1 \leq i \leq k$ such that Player i prefers to deviate from π_i to π'_i, and π'_i is of length ℓ_i. We denote by S' the resulting profile, i.e., $S' = S^*[i \leftarrow \pi'_i]$. For $1 \leq j \leq k$, let ν_j be the set of transitions that are used in π_j. Similarly, let ν'_i be the transitions used in π'_i. We sometimes view $\nu_1, \ldots, \nu_k, \nu'_i$ as paths rather than a sets of transitions.

We distinguish between four cases. In the first, both ν_i and ν'_i are simple paths. Then, every transition in $\nu_i \cap \nu'_i$ costs the same for Player i in both profiles, and since $\mathcal{A}$ has uniform transition costs, every transition in $\nu'_i \setminus \nu_i$ costs at least as much as any transition in $\nu'_i \setminus \nu_i$. Morevoer, since the runs are simple, the sizes of $\nu_i \setminus \nu'_i$ and $\nu'_i \setminus \nu_i$ are equal. Thus, $cost_i(S^*) \leq cost_i(S')$, and we reach a contradiction to the fact that Player i deviates.

In the second case, ν_i is simple and ν'_i is lasso. Thus, $|\nu'_i| \leq |\nu_i|$. If $|\nu'_i| = |\nu_i|$, we return to the previous case. Otherwise, $|\nu'_i| < |\nu_i|$. But since $|\nu_1| \geq |\nu_i|$, we reach a contradiction to our assumption that $\mathcal{A}$ is resistent. Indeed, if π_1 uses a lasso, then ν'_i is a shorter lasso, contradicting the minimality of the witness lasso for resistence. If π_1 does not use a lasso, then we reach a contradiction to our assumption that the witness lasso has length greater than ℓ_1.

In the third case, ν_i is a lasso and ν'_i is simple. Thus, $\nu_i = \nu_1$. Consider a transition $e \in \nu_i$. Let x_e and x'_e be the number of times Player i uses e in π_i and π'_i, respectively. Let y_e and y'_e be the number of times the other players use e in S^* and in S', respectively. Thus, $x_e > 0$ and $y_e \geq 0$. Also, $x'_e \leq 1$ and $y'_e = y_e$. Consider transitions $e, e' \in \nu_i$ having $x'_e = 1$ and $x_{e'} = 0$. That is, Player i reduces his number of uses of transition e from x_e to 1 and does not use e' at all in π'_i. Since the number of times

Player i uses a transition in π_i' is at most 1, there are $(x_e - 1)$ transitions that are not used by Player i in π_i and are used once in π_i'. Since $\nu_i = \nu_1$, these transitions are all in $\nu_i' \setminus \nu_i$ and Player i pays 1 for each of them. Similarly, there are $x_{e'}$ new transitions in $\nu_i' \setminus \nu_i$ that compensate for the fact that e' is not used in π_i'.

We calculate $cost_i(S^*) - cost_i(S')$. Consider a transition $e \in \nu_i$. Let $cost_i^e(S^*)$ and $cost_i^e(S')$ be the cost Player i pays for transition e in profiles S^* and S', respectively. If $x_e' = 1$, then by the above

$$cost_i^e(S^*) - cost_i^e(S') = \frac{x_e}{y_e + x_e} - (\frac{1}{y_e + 1} + (x_e - 1)) =$$

$$= \frac{x_e y_e + x_e + y_e^2 - y_e x_e^2 - x_e^2 - y_e^2 x_e}{(y_e + x_e) \cdot (y_e + 1)} \leq 0$$

Similarly, if $x_e' = 0$, then the change in cost incurred by e is:

$$cost_i^e(S^*) - cost_i^e(S') = \frac{x_e}{y_e + x_e} - x_e \leq 0$$

Since $cost_i(S^*) - cost_i(S') = \sum_{e \in \Delta} cost_i^e(S^*) - cost_i^e(S')$, we have $cost_i(S^*) - cost_i(S') \leq 0$, and thus $cost_i(S^*) \leq cost_i(S')$, which is a contradiction to the fact that Player i deviates.

We continue to the final case in which both ν_i and ν_i' are lassos. As in the previous case, $\nu_i = \nu_1$. Recall that the lasso ν_1 is the lasso that is the witness for the resistance of $\mathcal{A}$. We show that the lasso ν_i' violates our requirement for ν_1 and thus we reach a contradiction. Let $\nu_1 = u \cdot v$, where u is a simple path from the initial state and v is a simple cycle. Thus,

$$cost_i(S^*) = cost_i(S^*, u) + cost_i(S^*, v) \leq cost_i(S^*, u \cap \nu_i') + |u \setminus \nu_i'| + |v|.$$

Also,

$$cost_i(S') = cost_i(S', u \cap \nu_i') + cost_i(S', \nu_i' \cap v) + |\nu_i' \setminus \nu| \geq cost_i(S^*, u \cap \nu_i') + |\nu_i' \setminus \nu_i|.$$

Subtracting both inequalities we get:

$$cost_i(S^*) - cost_i(S') \leq |u \setminus \nu_i'| + |v| - |\nu_i' \setminus \nu_i|.$$

Since $cost_i(S^*) - cost_i(S') > 0$, we get:

$$|\nu_i' \setminus \nu_i| > |u \setminus \nu_i'| + |v|,$$

which is a contradiction to the resistance of $\mathcal{A}$, and we are done. □

B Directed Acyclic Automata

When the automaton includes no directed cycles, the runs selected by the players must be simple. In this case, we get a potential game, where the potential function is identical to the one used in the analysis of the classical network design game [2]. Specifically, let

$\kappa_P(e)$ be the number of runs that use the transition e at least once in a profile P, and let $\Phi(P) = \sum_{e \in E} c(e) \cdot H(\kappa_P(e))$, where $H(0) = 0$, and $H(k) = 1 + 1/2 + \ldots + 1/k$. Then, $\Phi(P)$ is a potential function whose value reduces with any improving step of a player, thus a pure NE exists and BRD is guaranteed to converge.

Using the analysis of [2], it holds that $PoS(DAG) \leq H_k$. We conjecture that for single-lettered all-accepting instances the price of stability is lower than H_k and is bounded by $\sum_{i=1}^{k} \frac{1}{2^{i-1}}$, which tends to 2 for $k \to \infty$. The lower bound follows from the claim below. Proving the upper bound is an open problem.

Claim 8.1: For any $\epsilon > 0$, the PoS for directed acyclic automata and single-lettered all-accepting instances is at least $\sum_{i=1}^{k} \frac{1}{2^{i-1}} - \epsilon$.

Proof. Consider the automaton given in Fig. 10. Each transition is marked by its cost. Assume that for $1 \leq j \leq k$, the language L_j is a word of length j. Let ϵ_j be a small constant defined such that $\epsilon_j > 2\epsilon_{j-1}$, and $\epsilon = \sum_{j=1}^{k-1} \epsilon_j$. The automaton consists of k disjoint runs. Path j for $1 \leq j \leq k$ has length $k - j + 1$. The first transition of run j has cost $\frac{1}{2^{j-1}} - \epsilon_j$. All other transitions are free.

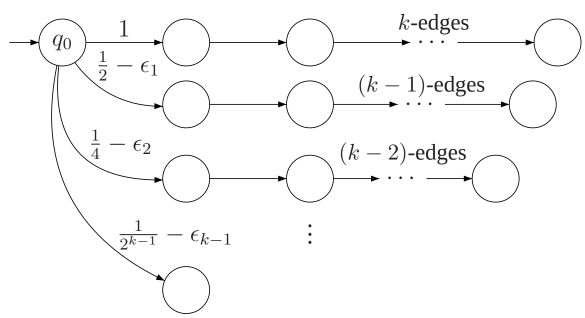

Fig. 10. The WFA for which the price of stability is $\sum_{i=1}^{k} \frac{1}{2^{i-1}} - \epsilon$.

The social optimum for this instance is achieved by choosing for all words the first run - of length k and cost 1. We show that the unique NE is the configuration P_0 in which Player j selects the run of length j. The cost of this NE is $\sum_{j=1}^{k} \frac{1}{2^{j-1}} - \sum_{j=1}^{k-1} \epsilon_j = \sum_{j=1}^{k} \frac{1}{2^{j-1}} - \epsilon$

It is easy to verify that P_0 is a NE: for any j, Player j can only deviate to a longer run. However, every longer run is utilized in P_0 by a single word and its cost is more than double the current cost of Player j. We show that P_0 is the only NE of this instance. Let P' be any other configuration. Since $P_0 \neq P'$, there must be a run in P' whose prefix is shared by two or more players. Assume that a run of length j is shared by $n > 1$ players. The player with the shortest length among these n has length at most $j - n + 1$. He can deviate to the run of length $j - n + 1$ and pay $\frac{1}{2^{k-j+n-1}} - \epsilon_{k-j+n-1}$ (or less if this run is used by additional players in P'). In P' the cost of this player is $\frac{1}{n \cdot 2^{k-j}} - \frac{\epsilon_{k-j}}{n}$. However, for any $n > 1$, $\frac{1}{2^{k-j+n-1}} - \epsilon_{k-j+n-1} < \frac{1}{n \cdot 2^{k-j}} - \frac{\epsilon_{k-j}}{n}$, contradicting the assumption that P' is a NE. □