

Efficient Equilibria in Polymatrix Coordination Games

Mona Rahn¹ and Guido Schäfer^{1,2}

¹ Centrum Wiskunde & Informatica (CWI), Amsterdam, The Netherlands

² VU University Amsterdam, The Netherlands

Abstract. We consider polymatrix coordination games with individual preferences where every player corresponds to a node in a graph who plays with each neighbor a separate bimatrix game with non-negative symmetric payoffs. In this paper, we study α -approximate k -equilibria of these games, i.e., outcomes where no group of at most k players can deviate such that each member increases his payoff by at least a factor α . We prove that for $\alpha \geq 2$ these games have the finite coalitional improvement property (and thus α -approximate k -equilibria exist), while for $\alpha < 2$ this property does not hold. Further, we derive an almost tight bound of $2\alpha(n-1)/(k-1)$ on the price of anarchy, where n is the number of players; in particular, it scales from unbounded for pure Nash equilibria ($k=1$) to 2α for strong equilibria ($k=n$). We also settle the complexity of several problems related to the verification and existence of these equilibria. Finally, we investigate natural means to reduce the inefficiency of Nash equilibria. Most promisingly, we show that by fixing the strategies of k players the price of anarchy can be reduced to n/k (and this bound is tight).

1 Introduction

In this paper, we are interested in strategic games where the players are associated with the nodes of a graph and can benefit from coordinating their choices with their neighbors. More specifically, we consider *polymatrix coordination games with individual preferences*: We are given an undirected graph $G = (N, E)$ on the set of players (nodes) $N := \{1, \dots, n\}$. Every player $i \in N$ has a finite set of strategies S_i to choose from and an individual preference function $q^i : S_i \rightarrow \mathbb{R}^+$. Each player $i \in N$ plays a separate bimatrix game with each of his neighbors in $N_i := \{j \in N \mid \{i, j\} \in E\}$. In particular, every edge $\{i, j\} \in E$ is associated with a payoff function $q^{ij} : S_i \times S_j \rightarrow \mathbb{R}^+$, specifying a non-negative payoff $q^{ij}(s_i, s_j)$ that both i and j receive if they choose strategies s_i and s_j , respectively. Given a joint strategy $s = (s_1, \dots, s_n)$ of all players, the overall payoff of player i is defined as

$$p_i(s) := q^i(s_i) + \sum_{j \in N_i} q^{ij}(s_i, s_j). \quad (1)$$

These games naturally model situations in which each player has individual preferences over the available options (possibly not having access to all options)

and may benefit in varying degrees from coordinating with his neighbors. For example, one might think of students deciding which language to learn, co-workers choosing which project to work on, or friends determining which mobile phone provider to use. On the other hand, these games also capture situations where players prefer to anti-coordinate, e.g., competing firms profiting equally by choosing different markets.

A special case of our games are *polymatrix coordination games* (without individual preferences, i.e., $q^i = 0$ for all i) which have previously been investigated by Cai and Daskalakis [9]. Among other results, the authors show that pure Nash equilibria are guaranteed to exist, but that finding one is PLS-complete. Polymatrix coordination games capture several other well-studied games among which are party affiliation games [6], cut games [11] and congestion games with positive externalities [12].

Yet another special case which will be of interest in this paper are *graph coordination games*. Here every edge $\{i, j\} \in E$ is associated with a non-negative edge weight w_{ij} and the payoff function q^{ij} is simply defined as $q^{ij}(s_i, s_j) = w_{ij}$ if $s_i = s_j$ and $q^{ij}(s_i, s_j) = 0$ otherwise. Intuitively, in this game every player (node) $i \in N$ chooses a color s_i from the set of colors S_i available to him and receives a payoff equal to the total weight of all incident edges to neighbors choosing the same color. These games have recently been studied by Apt et al. [2] for the special case of unit edge weights.

This paper is devoted to the study of equilibria in polymatrix coordination games with individual preferences. It is not hard to see that these games always admit pure Nash equilibria. However, in general these equilibria are highly inefficient. One of the most prominent notions to assess the inefficiency of equilibria is the *price of anarchy* [13]. It is defined as the ratio in social welfare of an optimal outcome and a worst-case equilibrium. Here the social welfare of a joint strategy s refers to the sum of the payoffs of all players, i.e., $\text{SW}(s) = \sum_{i \in N} p_i(s)$.

The high inefficiency of our games even arises in the special case of graph coordination games as has recently been shown in [2]. To see this, fix an arbitrary graph $G = (N, E)$ with unit edge weights and suppose each player $i \in N$ can choose between a private color c_i (only available to him) and a common color c . Then each player i choosing his private color c_i constitutes a Nash equilibrium in which every player has a payoff of zero. In contrast, if every player chooses the common color c then each player i obtains his maximum payoff equal to the degree of i . As a consequence, the price of anarchy is unbounded. The example demonstrates that the players might be unable to coordinate on the (obviously better) common choice because they cannot escape from a bad initial configuration by unilateral deviations. In particular, observe that the example breaks if two (or more) players can deviate simultaneously. This suggests that one should consider more refined equilibrium notions where deviations of groups of players are allowed.

In our studies, we focus on a general equilibrium notion which allows us to differentiate between both varying sizes of coalitional deviations and different degrees of player reluctance to deviate. More specifically, in this paper we con-

Problem	Complexity
Verification	(α, k) -equilibrium (k constant)
	(α, k) -equilibrium (α fixed)
	α -approximate strong equilibrium
Existence	k -equilibrium ($k \geq 2$ fixed)
	strong equilibrium

Table 1. Complexity of graph coordination games. The parameters α and k are assumed to be part of the input unless they are stated to be fixed. [†] Shown to be efficiently computable for forests.

consider α -approximate k -equilibria as the solution concept, i.e., outcomes that are resilient to deviations of at most k players such that each member increases his payoff by at least a factor of $\alpha \geq 1$. Subsequently, we call these equilibria also (α, k) -equilibria for short. In light of this refined equilibrium notion, several natural questions arise and will be answered in this paper: Which are the precise values of α and k that guarantee the existence of (α, k) -equilibria? What is the price of anarchy of these equilibria as a function of α and k ? How about the complexity of problems related to the verification and existence of such equilibria? And finally, are there efficient coordination mechanisms to reduce the price of anarchy?

Our contributions. We study (α, k) -equilibria of graph and polymatrix coordination games. Our main contributions are summarized below.

1. *Existence:* We prove that for $\alpha \geq 2$ polymatrix coordination games have the finite (α, k) -improvement property, i.e., every sequence of α -improving k -deviations is finite (and thus results in an (α, k) -equilibrium). We also exhibit an example showing that for $\alpha < 2$ this property does not hold in general. For graph coordination games we show that if the underlying graph is a tree then (α, k) -equilibria exist for every α and k . On the other hand, if the graph is a pseudotree (i.e., a tree with exactly one cycle) the existence of (α, k) -equilibria cannot be guaranteed for every $\alpha < \varphi$ and $k \geq 2$, where $\varphi = \frac{1}{2}(1 + \sqrt{5})$ is the golden ratio.
2. *Inefficiency:* We show that the price of anarchy of (α, k) -equilibria for polymatrix coordination games is at most $2\alpha(n-1)/(k-1)$. We also provide a lower bound of $2\alpha(n-1)/(k-1) + 1 - 2\alpha$. In particular, the price of anarchy drops from unbounded for pure Nash equilibria ($k = 1$) to 2α for strong equilibria ($k = n$), both of which are tight bounds.
3. *Complexity:* We settle the complexity of several problems related to the verification and existence of (α, k) -equilibria in graph coordination games. Naturally all hardness results extend to the more general class of polymatrix coordination games with individual preferences. A summary of our results is given in Table 1.
4. *Coordination mechanisms:* We investigate two natural mechanisms that a central coordinator might deploy to reduce the price of anarchy of pure Nash equilibria: (i) asymmetric sharing of the common payoffs q^{ij} and (ii)

strategy imposition of a limited number of players. Concerning (i), we show that there is no payoff distribution rule that reduces the price of anarchy in general. As to (ii), we prove that by (temporarily) fixing the strategies of k players according to an arbitrarily given joint strategy s , the resulting Nash equilibrium recovers at least a fraction of k/n of the social welfare $\text{SW}(s)$ and this is best possible. Exploiting this in combination with a 2-approximation algorithm for the optimal social welfare problem [12], we derive an efficient algorithm to reduce the price of anarchy to at most $2n/k$ for a special class of polymatrix coordination games with individual preferences.

Related work. Apt et al. [2] study k -equilibria in graph coordination games with unit edge weights, which constitute a special case of our games. They identify several graph structural properties that ensure the existence of such equilibria. Interestingly, most of these results do not carry over to our weighted graph coordination games, therefore demanding for the new approach of considering approximate equilibria.

Many of the mentioned games have been studied from a computational complexity point of view. In particular, Cai and Daskalakis [9] show that the problem of finding a pure Nash equilibrium in a polymatrix coordination game is PLS-complete. Further, they show that finding a mixed Nash equilibrium is in $\text{PPAD} \cap \text{PLS}$. While this suggests that the latter problem is unlikely to be hard, it is not known whether it is in P. It is easy to see that these results also carry over to our polymatrix coordination games with individual preferences.³

For the special case of party affiliation games efficient algorithms to compute an approximate Nash equilibrium are known [7, 10]. The current best approximation guarantee is $3 + \varepsilon$, where $\varepsilon > 0$, due to Caragiannis, Fanelli and Gavim [10]. The algorithm crucially exploits that party affiliation games admit an exact potential whose relative gap (called *stretch*) between any two Nash equilibria is bounded by 2. The latter property is not satisfied in our games, even for graph coordination games (as the example outlined in the Introduction shows).

A class of games that is closely related to our graph coordination games are *additively separable hedonic games* [8]. As in our games, the players are embedded in a weighted graph. Every player chooses a coalition and receives as payoff the total weight of all edges to neighbors in the same coalition. These games were originally studied in a cooperative game theory setting. More recently, researchers also address computational issues of these games (see, e.g., [4]). It is important to note that in hedonic games every player can choose every coalition, while in our graph coordination games players may only have limited options.

Anshelevich and Sekar [1] study coordination games with individual preferences where the players are nodes in a graph and profit from neighbors choosing the same color. However, in their setting the edge weight between two neighbors can be distributed asymmetrically and all players are assumed to have the same

³ In [9] the bimatrix games on the edges may have negative payoffs and this is exploited in the PLS-completeness proof. However, we can accommodate this in our model by adding a sufficiently large constant to each payoff.

strategy set. Among other results, they give an algorithm to compute a $(2, n)$ -equilibrium and show how to efficiently compute an approximate equilibrium that is not too far from the optimal social welfare.

Concerning the social welfare optimization problem, a 2-approximation algorithm is given in [12] for the special case of polymatrix coordination games with individual preferences where the bimatrix game of each edge has positive entries only on the diagonal.

Our techniques. Most of our existence results use a generalized potential function argument for coalitional deviations. In our proof of the upper bound on the inefficiency of (α, k) -equilibria we first argue locally for a fixed coalition of players and then use a sandwich bound in combination with a counting argument to derive the upper bound. Most of our lower bounds and hardness results follow by exploiting specific properties and deep structural insights of graph coordination games with edge weights.

It is worth mentioning that our algorithm to compute a strong equilibrium for graph coordination games on trees reveals a surprising connection to a sequential-move version of the game. In particular, we show that if we fix an arbitrary root of the tree and consider the induced sequential-move game then every subgame perfect equilibrium corresponds to a strong equilibrium of the original game. As a consequence, strong equilibria exist and can be computed efficiently. Further, this in combination with our strong price of anarchy bound shows that the *sequential price of anarchy* [14] for these induced games is at most 2, which is a significant improvement over the unbounded price of anarchy for the strategic-form version of the game. This result is of independent interest.

We also note that the k/n bound on the social welfare which is guaranteed by our strategy imposition algorithm is proven via a *smoothness argument* [16]. Besides some other consequences, this implies that our bound also holds for more permissive solution concepts such as correlated and coarse correlated equilibria (see [16] for more details).

2 Preliminaries

Let $\mathcal{G} = (G, (S_i)_{i \in N}, (q^i)_{i \in N}, (q^{ij})_{\{i,j\} \in E})$ be a polymatrix coordination game with individual preferences (w.i.p.) where $G = (N, E)$ is the underlying graph. Recall that we identify the player set N with $\{1, \dots, n\}$. We first introduce some standard game-theoretic concepts.

We call a subset $K := \{i_1, \dots, i_k\} \subseteq N$ of players a *coalition of size k* . We define the set of joint strategies of players in K as $S_K := S_{i_1} \times \dots \times S_{i_k}$ and use $S := S_N$ to refer to the set of joint strategies of all players. Given a joint strategy $s \in S$, we use s_K to refer to $(s_{i_1}, \dots, s_{i_k})$ and s_{-K} to refer to $(s_i)_{i \notin K}$. By slightly abusing notation, we also write (s_K, s_{-K}) instead of s . If there is a strategy x such that $s_i = x$ for every player $i \in K$, we also write $s = (x_K, s_{-K})$.

Given a joint strategy s and a coalition K , we say that $s' = (s'_K, s_{-K})$ is a *deviation of coalition K from s* if $s'_i \neq s_i$ for every player $i \in K$; we also denote

this by $s \rightarrow_K s'$. If we constrain to deviations of coalitions of size at most k , we call such deviations k -deviations. We call a deviation α -improving if every player in the coalition improves his payoff by at least a factor of $\alpha \geq 1$, i.e., for every $i \in K$, $p_i(s') > \alpha p_i(s)$; we also call such deviations (α, k) -improving. We omit the explicit mentioning of the parameters if $\alpha = 1$ or $k = 1$. A joint strategy s is an α -approximate k -equilibrium (also called (α, k) -equilibrium for short) if there is no (α, k) -improving deviation from s . If $k = 1$ or $k = n$ then we also refer to the respective equilibrium notion as α -approximate Nash equilibrium and α -approximate strong equilibrium [3].

We say that a finite strategic game has the *finite (α, k) -improvement property* (or (α, k) -FIP for short) if every sequence of (α, k) -improving deviations is finite. This notion generalizes the finite improvement property introduced by Monderer [15] for $\alpha = k = 1$. A function $\Phi : S \rightarrow \mathbb{R}$ is called an (α, k) -generalized potential if for every joint strategy s , for every (α, k) -improving deviation $s' := (s'_K, s_{-K})$ from s it holds that $\Phi(s') > \Phi(s)$. It is not hard to see that if a finite game admits an (α, k) -generalized potential then it has the (α, k) -FIP.

The *social welfare* of a joint strategy s is defined as $\text{SW}(s) := \sum_{i \in N} p_i(s)$. For $K \subseteq N$, we define $\text{SW}_K(s) := \sum_{i \in K} p_i(s)$. A joint strategy s^* of maximum social welfare is called a *social optimum*. Given a finite game that has an (α, k) -equilibrium, its (α, k) -price of anarchy (POA) is the ratio $\text{SW}(s^*)/\text{SW}(s)$, where s^* is a social optimum and s is an (α, k) -equilibrium of smallest social welfare. In the case of division by zero, we interpret the outcome as ∞ . Note that if $\alpha' \geq \alpha$ and $k' \leq k$, then every (α, k) -equilibrium is an (α', k') -equilibrium. Hence the (α, k) -PoA lower bounds the (α', k') -PoA.

Due to lack of space, several proofs or parts thereof are omitted from this extended abstract and will be given in the full version of the paper.

3 Existence

We first give a characterization of the values α and k for which our polymatrix coordination games with individual preferences have the (α, k) -FIP.

Theorem 1. *Let \mathcal{G} be a polymatrix coordination game w.i.p. Then:*

1. \mathcal{G} has the $(\alpha, 1)$ -FIP for every α .
2. \mathcal{G} has the (α, k) -FIP for every $\alpha \geq 2$ and for every k .

Proof. Observe that every α -improving deviation is also α' -improving for $\alpha \geq \alpha'$. It is thus sufficient to prove the claims above for $\alpha = 1$ and $\alpha = 2$, respectively.

The proof idea for the first claim is to show that the game admits an exact potential and thus has the FIP.

We prove the second claim for $\alpha = 2$ by showing that $\Phi(s) := \text{SW}(s)$ is a $(2, k)$ -generalized potential. Given a joint strategy s and two sets $K, K' \subseteq N$, define

$$Q_s(K, K') := \sum_{i \in K, j \in N_i \cap K'} q^{ij}(s) \quad \text{and} \quad Q_s(K) := \sum_{i \in K} q^i(s).$$

Consider a $(2, k)$ -improving deviation $s' = (s'_K, s_{-K})$ from s . Let \bar{K} be the complement of K . We have $\text{SW}_K(s) = Q_s(K, K) + Q_s(K, \bar{K}) + Q_s(K)$. Note that $\text{SW}_K(s') > 2\text{SW}_K(s)$ because the deviation is 2-improving. Thus,

$$Q_{s'}(K, K) + Q_{s'}(K, \bar{K}) + Q_{s'}(K) > 2(Q_s(K, K) + Q_s(K, \bar{K}) + Q_s(K)). \quad (2)$$

The social welfare of s can be written as

$$\text{SW}(s) = Q_s(K, K) + 2Q_s(K, \bar{K}) + Q_s(\bar{K}, \bar{K}) + Q_s(K) + Q_s(\bar{K}).$$

Note that $Q_s(\bar{K}, \bar{K}) = Q_{s'}(\bar{K}, \bar{K})$ and $Q_s(\bar{K}) = Q_{s'}(\bar{K})$. Using (2), we obtain

$$\begin{aligned} \Phi(s') - \Phi(s) &= Q_{s'}(K, K) + 2Q_{s'}(K, \bar{K}) + Q_{s'}(K) \\ &\quad - Q_s(K, K) - 2Q_s(K, \bar{K}) - Q_s(K) \\ &> Q_s(K, K) + Q_{s'}(K, \bar{K}) + Q_s(K) \geq 0. \end{aligned}$$

Thus $\Phi(s)$ is a $(2, k)$ -generalized potential which concludes the proof. \square

The next theorem shows that in general our polymatrix coordination games do not have the (α, k) -FIP for $\alpha < 2$.

Theorem 2. *For all $\alpha < 2$ there is a polymatrix coordination game \mathcal{G} that has a cycle of $(\alpha, n - 1)$ -improving deviations.*

We derive some more refined insights for the special case of graph coordination games.

Theorem 3. *The following holds for graph coordination games:*

1. *Let \mathcal{G} be a graph coordination game on a tree. Then \mathcal{G} has a strong equilibrium.*
2. *There is a graph coordination game \mathcal{G} on a graph with one cycle such that no (α, k) -equilibrium exists for every $\alpha < \varphi$ and $k \geq 2$, where $\varphi := \frac{1}{2}(1 + \sqrt{5}) \approx 1.62$ is the golden ratio.*

Note that Theorem 3 shows that for $k \geq 2$ a k -equilibrium may not exist. In contrast, Nash equilibria always exist by Theorem 1. Further, the graph used to show the second claim is a pseudoforest⁴. For graph coordination games with unit edge weights, this guarantees the existence of a strong equilibrium [2].

4 Inefficiency

We analyze the price of anarchy of our polymatrix coordination games. The upper bound in the special case of $(\alpha, k) = (1, n)$ follows from a result in [5].

Theorem 4. *The (α, k) -price of anarchy in polymatrix coordination games w.i.p. is between $2\alpha(n - 1)/(k - 1) + 1 - 2\alpha$ and $2\alpha(n - 1)/(k - 1)$. The upper bound of 2α is tight for α -approximate strong equilibria.*

⁴ A graph is a *pseudoforest* if each of its connected components has at most one cycle.

Proof (upper bound). Let s be an (α, k) -equilibrium (which we assume to exist) and let s^* be a social optimum. Fix an arbitrary coalition $K = \{i_1, \dots, i_k\}$ of size k . Then there is a player $i \in K$ such that $p_i(s_K^*, s_{-K}) \leq \alpha p_i(s)$. Denote by $p_i^K(s^*) := q^i(s^*) + \sum_{j \in N_i \cap K} q^{ij}(s^*)$ the total payoff that i gets from players in K under s^* (including himself). Because all payoffs are non-negative, we have

$$p_i^K(s^*) \leq q^i(s_i^*) + \sum_{j \in N_i \cap K} q^{ij}(s_i^*, s_j^*) + \sum_{j \in N_i \cap \bar{K}} q^{ij}(s_i^*, s_j) = p_i(s_K^*, s_{-K}). \quad (3)$$

Thus, $p_i^K(s^*) \leq \alpha p_i(s)$. Rename the nodes in K such that $i_k = i$ and repeat the arguments above with $K \setminus \{i_k\}$ instead of K . Continuing this way, we obtain that for every player $i_x \in K$, $x \in \{1, \dots, k\}$, $p_{i_x}^{\{i_1, \dots, i_x\}}(s^*) \leq \alpha p_{i_x}(s)$.

We thus have

$$\begin{aligned} \sum_{i \in K} \left(q^i(s^*) + \frac{1}{2} \sum_{j \in N_i \cap K} q^{ij}(s^*) \right) &= \sum_{x=1}^k \left(q^{i_x}(s^*) + \sum_{i_y \in N_{i_x} \cap K: y < x} q^{i_x i_y}(s^*) \right) \\ &= \sum_{x=1}^k p_{i_x}^{\{i_1, \dots, i_x\}}(s^*) \leq \alpha \sum_{i \in K} p_i(s). \end{aligned}$$

Summing over all coalitions K of size k , we obtain

$$\sum_{K: |K|=k} \left(\sum_{i \in K} \left(q^i(s^*) + \frac{1}{2} \sum_{j \in N_i \cap K} q^{ij}(s^*) \right) \right) \leq \alpha \sum_{K: |K|=k} \sum_{i \in K} p_i(s). \quad (4)$$

Consider the right-hand side of (4). Note that every player $i \in N$ occurs in $\binom{n-1}{k-1}$ many coalitions of size k because we can choose $k-1$ out of $n-1$ remaining players to form a coalition of size k containing i . Thus

$$\sum_{K: |K|=k} \sum_{i \in K} p_i(s) = \binom{n-1}{k-1} \sum_{i \in N} p_i(s) = \binom{n-1}{k-1} \text{SW}(s). \quad (5)$$

Similarly, the first term of the left-hand side of (4) yields

$$\sum_{K: |K|=k} \sum_{i \in K} q^i(s^*) = \binom{n-1}{k-1} \sum_{i \in N} q^i(s^*) \geq \frac{1}{2} \binom{n-2}{k-2} \sum_{i \in N} q^i(s^*).$$

Now, consider the second term of the left-hand side of (4). Every pair (i, j) with $i \in N$ and $j \in N_i$ occurs in $\binom{n-2}{k-2}$ many coalitions of size k because we can choose $k-2$ out of $n-2$ remaining players to complete a coalition of size k containing both i and j . Thus for the left-hand side of (4) we obtain

$$\begin{aligned} \sum_{K: |K|=k} \left(\sum_{i \in K} \left(q^i(s^*) + \frac{1}{2} \sum_{j \in N_i \cap K} q^{ij}(s^*) \right) \right) \\ \geq \frac{1}{2} \binom{n-2}{k-2} \left(\sum_{i \in N} q^i(s^*) + \sum_{i \in N} \sum_{j \in N_i} q^{ij}(s^*) \right) = \frac{1}{2} \binom{n-2}{k-2} \text{SW}(s^*). \quad (6) \end{aligned}$$

Combining (5) and (6) with inequality (4), we obtain that the (α, k) -price of anarchy is at most $2\alpha \binom{n-1}{k-1} / \binom{n-2}{k-2} = 2\alpha \frac{n-1}{k-1}$. \square

5 Complexity

In this section, we study the complexity of various computational problems on graph coordination games.

Theorem 5. *Let \mathcal{G} be a graph coordination game. Given a joint strategy s , the problem of deciding whether s is an (α, k) -equilibrium*

1. *is in P , if $k = O(1)$ or $k = n$;*
2. *is co-NP-complete for every fixed α .*

Proof (sketch). We sketch the proof of the first claim for $k = n$. A crucial insight is that if there is an α -improving deviation from s then there is one which is *simple*, i.e., $s' = (s'_K, s_{-K})$ where the subgraph $G[K]$ induced by K is connected and all nodes in K deviate to the same color $s'_K = x$ for some x .

Fix some color x and let $G_x := (N_x, E_x)$ be the subgraph of G induced by the set of nodes N_x that can choose color x but do not do so in s . For each $u \in N_x$ define $d_u := \alpha p_u(s) - w(\{\{u, v\} \in E \mid s_v = x\}) - q^u(x)$. Now, a deviation of a coalition $K \subseteq N_x$ to (x_K, s_{-K}) is α -improving if and only if for every node $u \in K$ the total weight of all incident edges in the induced subgraph $G_x[K]$ is larger than d_u . We prove that an inclusionwise maximal $K \subseteq N_x$ satisfying this property can be found in polynomial time. This way we can verify for every color x whether an α -improving deviation exists. \square

Deciding whether a graph coordination game admits a k -equilibrium is hard for every $k \geq 2$. Note that for unit edge weights 2-equilibria are guaranteed to exist and can be found efficiently, as shown in [2].

Theorem 6. *Let \mathcal{G} be a graph coordination game. Then the problem of deciding whether there is a k -equilibrium is NP-complete for every fixed $k \geq 2$.*

Proof ($k = 2$). We give a reduction from MINIMUM MAXIMAL MATCHING which is known to be NP-complete [17]: Given a graph $G = (V, E)$ and a number l , does there exist an inclusionwise maximal matching of size at most l ?

Let (G, l) be an instance of this problem with $G = (V, E)$ and $n = |V|$. We add $n - 2l$ gadgets H_1, \dots, H_{n-2l} to G , where an illustration of gadget H_i is given in Figure 1. The dashed edge from v_0^i to G indicates that v_0^i is connected to all vertices in G and each of these edges has weight 3. We assign to each node $v \in V$ the color set $S_v = \{x_v^i \mid i = 1, \dots, n - 2l\} \cup \{y_e \mid e = \{v, w\} \in E\}$, i.e., v can either choose a ‘gadget color’ x_v^i or a color corresponding to some adjacent edge in E . Every edge in E has weight 4. Note that for all joint strategies of nodes in V the set of unicolor edges in E constitutes a matching. The idea is that in every 2-equilibrium $n - 2l$ nodes in V are needed to ‘stabilize’ the gadgets and the $2l$ remaining nodes in V form a maximal matching.

Assume that G has a maximal matching $M \subseteq E$ with $|M| \leq l$. We construct a 2-equilibrium s . For every matched node $v \in V(M)$, choose the color corresponding to the adjacent matching edge. On the unmatched nodes in V and nodes of the form v_0^i , we assign colors in such a way that every gadget has one




Fig. 1. The gadget H_i .

outgoing edge (indicated by the dashed edge) that is unicolor. This is possible because there are at least $n - 2l$ unmatched nodes in V . If there are uncolored nodes in V left, assign arbitrary colors to them. Finally, let v_1^i and v_2^i choose color b for every i . We claim that s is a 2-equilibrium: The matched nodes obtain a payoff of 4, which is the maximal payoff nodes in V can get; so they are not part of any improving deviation. Let $v \in V$ be unmatched. Then v cannot deviate together with another unmatched node to increase the payoff because M is maximal. Further, all gadget nodes are ‘taken’: every v_0^i has a payoff of 3, which a joint deviation with v cannot increase. This implies that v cannot be part of any improving deviation. Lastly, it is easy to see that pairs of gadget nodes cannot profitably deviate. This proves that s is a 2-equilibrium.

Conversely, assume that a joint strategy s is a 2-equilibrium. Let M consist of the unicolor edges in G . By the choice of the color assignment, M is a matching. M is maximal because if there were two unmatched adjacent nodes, then they could form a profitable deviating coalition. It remains to show that $|M| \leq l$. It is not hard to see that if there is a gadget without an outgoing unicolor edge, then there is a 2-improving deviation in H_i . So at least $n - 2l$ nodes choose gadget colors, implying that $|V(M)| \leq 2l$ and thus $|M| \leq l$. \square

On the positive side, we can compute a strong equilibrium in polynomial time if the underlying graph is a tree.

Theorem 7. *Let \mathcal{G} be a graph coordination game on a tree. Then there is a polynomial-time algorithm to compute a strong equilibrium.*

Proof (sketch). The idea is as follows: We fix an arbitrary root r of the tree and consider the induced sequential-move game. This game has a subgame perfect equilibrium s which can be computed in polynomial time by backwards induction. Let \bar{s} be the corresponding joint strategy of \mathcal{G} if every player plays his best response according to s . We can prove that \bar{s} is a strong equilibrium of \mathcal{G} . \square

6 Coordination mechanisms

In this section, we investigate means that a central designer could use to reduce the inefficiency of Nash equilibria.

In our games the common payoff q^{ij} of the bimatrix game on edge $\{i, j\} \in E$ is distributed equally to both i and j . An idea that arises is to use different *payoff sharing rules* to reduce the inefficiency. Unfortunately, it is not hard to see from the example given in the Introduction that the price of anarchy remains unbounded no matter which payoff sharing rule is used.

We therefore consider another natural approach. Suppose the central designer can impose strategies on a subset of the players to reduce the inefficiency. Let \mathcal{G} be a polymatrix coordination game w.i.p. Further, let $K \subseteq N$ be a subset of the players and fix a joint strategy $f_K \in S_K$ for players in K . We define $\mathcal{G}[f_K]$ as the game with players from $N \setminus K$ that arises from \mathcal{G} if we fix the strategies of all players in K according to f_K . We say that f_K *guarantees* social welfare z if $\text{SW}(f_K, s_{-K}) \geq z$ for all Nash equilibria s_{-K} of $\mathcal{G}[f_K]$. We also call f_K a *joint strategy of size $|K|$* .

Suppose that f_K guarantees social welfare z . Then once all players in $\mathcal{G}[f_K]$ have reached a Nash equilibrium we can release all players in K and let them play their best responses too. By Theorem 1, the social welfare can only increase subsequently. As a result, the final Nash equilibrium has social welfare at least z . So we can view f_K as a ‘temporary advice’ for the players in K . A similar idea has been put forward in [6].

We first show that determining the minimum number of players to guarantee a certain social welfare is hard, even for graph coordination games.

Theorem 8. *Let \mathcal{G} be a graph coordination game. Given a joint strategy s , the problem of finding a minimal k such that there is a joint strategy f_K of size k that guarantees social welfare $\text{SW}(s)$ is NP-hard. The claim also holds if f_K is restricted to be s_K .*

In light of the above hardness results, we resort to approximation algorithms.

Theorem 9. *Let \mathcal{G} be a polymatrix coordination game w.i.p. Given a joint strategy s and a number k , we can find in polynomial time a coalition K of size k such that s_K guarantees social welfare $\frac{k}{n}\text{SW}(s)$ and this is tight.*

Using the 2-approximation algorithm in [12] for the social welfare optimization problem, we obtain the following result:

Corollary 1. *Let \mathcal{G} be a polymatrix coordination game w.i.p. where the bimatrix game of every edge has positive entries on the diagonal only. Given a number k , we can compute a joint strategy f_K of size K that guarantees social welfare $\frac{k}{2n}\text{SW}(s^*)$, where s^* is a social optimum.*

References

1. E. Anshelevich and S. Sekar. Approximate equilibrium and incentivizing social coordination. In *Proc. of 28th Conference on Artificial Intelligence (AAAI)*, pages 508–514, 2014.
2. K. R. Apt, M. Rahn, G. Schäfer, and S. Simon. Coordination games on graphs (extended abstract). In *Proc. of the 10th Conference on Web and Internet Economics (WINE)*, volume 8877, pages 441–446, 2014.
3. R. J. Aumann. Acceptable points in general cooperative n-person games. In *Contribution to the theory of game IV, Annals of Mathematical Study 40*, pages 287–324. 1959.
4. H. Aziz, F. Brandt, and H. G. Seedig. Stable partitions in additively separable hedonic games. In *Proc. of the 10th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 183–190, 2011.
5. Y. Bachrach, V. Syrgkanis, E. Tardos, and M. Vojnović. Strong price of anarchy, utility games and coalitional dynamics. In *Proc. of the 7th International Symposium on Algorithmic Game Theory (SAGT)*, pages 218–230, 2014.
6. M. Balcan, A. Blum, and Y. Mansour. Improved equilibria via public service advertising. In *Proc. of the 20th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 728–737, 2009.
7. A. Bhalgat, T. Chakraborty, and S. Khanna. Approximating pure Nash equilibrium in cut, party affiliation, and satisfiability games. In *Proc. of the 11th ACM Conference on Electronic Commerce (EC)*, pages 73–82, 2010.
8. A. Bogomolnaia and M. O. Jackson. The stability of hedonic coalition structures. *Games and Economic Behavior*, 38(2):201–230, 2002.
9. Y. Cai and C. Daskalakis. On minmax theorems for multiplayer games. In *Proc. of the 22nd ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 217–234, 2011.
10. I. Caragiannis, A. Fanelli, and N. Gavin. Short sequences of improvement moves lead to approximate equilibria in constraint satisfaction games. In *Proc. of the 7th International Symposium on Algorithmic Game Theory (SAGT)*, pages 49–60, 2014.
11. G. Christodoulou, V. S. Mirrokni, and A. Sidiropoulos. Convergence and approximation in potential games. *Theoretical Computer Science*, 438:13–27, 2012.
12. B. de Keijzer and G. Schäfer. Finding social optima in congestion games with positive externalities. In *Proc. of the 20th European Symposium on Algorithms (ESA)*, pages 395–406, 2012.
13. E. Koutsoupias and C. Papadimitriou. Worst-case equilibria. *Computer Science Review*, 3(2):65–69, 2009.
14. R. P. Leme, V. Syrgkanis, and E. Tardos. The curse of simultaneity. In *Proc. of the 3rd Conference on Innovations in Theoretical Computer Science (ITCS)*, pages 60–67, 2012.
15. D. Monderer and L. S. Shapley. Potential games. *Games and Economic Behavior*, 14(1):124–143, 1996.
16. T. Roughgarden. Intrinsic robustness of the price of anarchy. In *Proc. 41st ACM Symposium on Theory of Computing (STOC)*, pages 513–522, 2009.
17. M. Yannakakis and F. Gavril. Edge dominating sets in graphs. *SIAM Journal on Applied Mathematics*, 38(3):364–372, 1980.