Distance Polymatrix Coordination Games[⋆]

Alessandro Aloisio[1,*,†], Michele Flammini[2,†], Bojana Kodric[3,†] and Cosimo Vinci[4,†]

[1]Università degli Studi Internazionali di Roma, Via Cristoforo Colombo, 200 – 00147, Roma, Italy
[2]Gran Sasso Science Institute, Viale Francesco Crispi, 7 - 67100 LAquila, Italy
[3]Ca' Foscari University of Venice, Dorsoduro 3246, 30123 Venice, Italy
[4]University of Salerno, Piazza Tancredi, 7 - 73100 Lecce, Italy

Abstract

In *polymatrix coordination games*, each player x is a node of a graph and must select an action in her strategy set. Nodes are playing separate bimatrix games with their neighbors in the graph. Namely, the utility of x is given by the preference she has for her action plus, for each neighbor y, a payoff which strictly depends on the mutual actions played by x and y.

We propose the new class of *distance polymatrix coordination games*, properly generalizing *polymatrix coordination games*, in which the overall utility of player x further depends on the payoffs *arising* from mutual actions of players v, z that are the endpoints of edges at any distance $h < d$ from x, for a fixed threshold value $d \leq n$. In particular, the overall utility of player x is the sum of all the above payoffs, where each payoff is proportionally discounted by a factor depending on the distance h of the corresponding edge.

Under the above framework, which is a natural generalization that is well-suited for capturing positive community interactions, we study the social inefficiency of equilibria resorting to standard measures of Price of Anarchy and Price of Stability. Namely, we provide suitable upper and lower bounds for the aforementioned quantities, both for bounded-degree and general graphs.

Keywords

Polymatrix Games, Strong Nash Equilibrium, Price of Anarchy, Price of Stability

1. Introduction

Polymatrix games [2] are a well-known universal framework for modelling multi-agent games, which considers only pairwise interactions and thus allows a concise representation. They have been thoroughly studied since, both in some classical works [3, 4, 5, 6] and also more recently with a special focus on equilibria [7, 8, 9, 10]. In polymatrix games each player plays a separate bimatrix game with every other player. In the restricted version named *polymatrix coordination games* [7], an outcome of a bimatrix game gives the same payoff $w_{\{x,y\}}(\sigma_x, \sigma_y)$ to

IPS-RCRA-SPIRIT 2023: Italian Workshop on Planning and Scheduling, RCRA Workshop on Experimental evaluation of algorithms for solving problems with combinatorial explosion, and SPIRIT Workshop on Strategies, Prediction, Interaction, and Reasoning in Italy. November 7-9th, 2023, Rome, Italy [1]

*Corresponding author.
†These authors contributed equally.
alessandro.aloisio@unint.eu (A. Aloisio); michele.flammini@gssi.it (M. Flammini); bojana.kodric@unive.it (B. Kodric); cosimo.vinci@unisalento.it (C. Vinci)
0000-0003-3911-4008 (A. Aloisio); 0000-0003-0327-3728 (M. Flammini); 0000-0001-7242-0096 (B. Kodric); 0000-0001-7741-9342 (C. Vinci)

the two players x and y involved in it. Moreover, every player gets also an additional payoff $p_x(\sigma_x)$ that only depends on the strategy she chooses.

In this paper, we generalize polymatrix coordination games by allowing players to receive a further payoff from the interactions in which they are not personally involved. The idea here is that each player benefits not only from good relations with her immediate neighbours but also from the positive environment stemming from good relations between them and their respective immediate neighbours. A further generalization of this thought brings us to a model in which the utility is computed as the sum of the payoffs from the whole connected component of the interaction graph up to a certain maximal distance d, where d is a parameter of the model. Furthermore, it seems reasonable to discount the payoff received from non-neighbouring edges by a factor between zero and one and to make such factors decrease with the distance of the corresponding edge/interaction. In other words, an agent x gets also the payoff $\alpha_{h+1} \cdot w_{\{v,z\}}(\sigma_v, \sigma_z)$ for every edge $\{v, z\}$ at distance $h < d$ from x, where α_{h+1} is the relative discount factor. We call the arising model that generalizes polymatrix coordination games, *distance polymatrix coordination games*.

Distance polymatrix coordination games are able to capture many types of interactions in the real world. In fact, several kinds of positive community effects easily fall within their scope. For instance, members of a scientific community obviously benefit from successful collaborations with their colleagues (while at the same time having personal preferences of what they would like to work on). However, any individual also benefits, albeit to a smaller degree, when his close colleagues have successful collaborations that he is not personally a part of. This is quite obvious when thinking about the student-advisor relationship but also noticeable for researchers working at the same university or institution. A further example comes from politics, where a person who belongs to a party profits not only from her direct contacts but also from the contacts of her contacts, etc. At the same time, it is also common that the benefit obtained by relations at second or higher distance level generate less payoff, which is taken into account by our discount factors.

In the setting described above, we will be focusing on the efficiency of the system. Our reference point for stability will be k-strong Nash equilibria, which are action profiles from which no group of up to k agents can simultaneously deviate such that all of them profit from the deviation. Our analysis provides bounds which depend on k and the discounting factors for the part of the utility of the agents coming from non-first-hand interactions.

A full version of our results can be found in [11]. At the same time, a further generalisation to hypergraphs can be found in [12]. Related to our work are also *(symmetric) additively separable hedonic games* [13] and *hypergraph hedonic games* [14], where the players are embedded in a weighted graph, and the utility is computed as the sum of the edges or hyperedges towards members of the same coalition. The difference from our model, however, is that in hedonic games in general, every coalition is a feasible choice for every player, there are no individual preferences, and the weights in each bimatrix are all equal to either 0 or to a fixed value w.

Another model related to our work is the *group activity selection problem* [15, 16, 17], standing between polymatrix coordination games and hedonic games. Also, here, in each bimatrix, all the weights are either 0 or a fixed value w, but there are also individual preferences that depend on the chosen activity.

The idea of obtaining utility from non-neighbouring players has been explored recently for a variant of hedonic games, called *distance hedonic games*, that are not additively separable since the coalition size also plays a role in determining the payoffs [18]. They generalize *fractional hedonic games* [19, 20, 21, 22, 23] similarly as our model does with polymatrix games.

2. Model and Definitions

Distance Polymatrix Coordination Games. Given an integer $d \geq 1$, a *d-distance polymatrix coordination game* $\mathcal{G} = (G, (\Sigma_x)_{x\in V}, (w_e)_{e\in E}, (p_x)_{x\in V}, (\alpha_h)_{h\in[d]})$ is a tuple defined as follows:

- $G = (V, E)$ is an undirected graph, where V is the set of *players* and E the set of *edges* between players.
- For any $x \in V$, Σ_x is a finite set of *strategies* of player x. A *strategy profile* $\sigma = (\sigma_1, \dots, \sigma_n)$ is a configuration in which each player $x \in V$ plays strategy $\sigma_x \in \Sigma_x$.
- For any edge $\{v, z\} \in E$, let $w_{\{v,z\}} : \Sigma_v \times \Sigma_z \to \mathbb{R}_{\geq 0}$ be the *weight function* that assigns, to each pair of strategies σ_v, σ_z played respectively by v and z, a *weight* $w_{\{v,z\}}(\sigma_v, \sigma_z) \geq 0$.
- For any $x \in V$, let $p_x : \Sigma_x \to \mathbb{R}_{\geq 0}$ be the *player-preference function* that assigns, to each strategy profile σ_x played by player x, a non-negative real value $p_x(\sigma_x)$, called *player-preference.*
- Let $(\alpha_h)_{h\in[d]}$ be the *distance-factors sequence* of the game, that is a non-negative sequence of real parameters, called *distance-factors*, such that $1 = \alpha_1 \geq \alpha_2 \geq \ldots \geq \alpha_d \geq 0$.

In what follows, for the sake of brevity, given any strategy profile σ, we will often denote $w_{\{v,z\}}(\sigma_v, \sigma_z)$ and $p_x(\sigma_x)$ simply as $w_{\{v,z\}}(\sigma)$ and $p_x(\sigma)$, respectively. For any $h \in [d]$, let $E_h(x)$ be the set of edges $\{v, z\}$ such that the minimum distance between x and one of the players v and z is exactly $h - 1$. Then, for any $x \in V$, the *utility function* $u_x : \times_{x\in V}\Sigma_x \to \mathbb{R}$ of player x, for any strategy profile σ is defined as $u_x(\sigma) := p_x(\sigma) + \sum_{h\in[d]} \alpha_h \sum_{e\in E_h(x)} w_e(\sigma)$. Given a strategy profile σ, the *social welfare* of σ is defined as $\mathsf{SW}(\sigma) = \sum_{x\in V} u_x(\sigma)$. A *social optimum* of game $\mathcal{G}$ is a strategy profile σ^* that maximizes the social welfare. We denote by $\mathsf{OPT}(\mathcal{G}) = \mathsf{SW}(\sigma^*)$ the corresponding value.

k-strong Nash equilibrium. Given two strategy profiles $\sigma = (\sigma_1, \dots, \sigma_n)$ and $\sigma^* = (\sigma_1^*, \dots, \sigma_n^*)$, and a subset $Z \subseteq V$, let $\sigma \xrightarrow{Z} \sigma^*$ be the strategy profile $\sigma' = (\sigma_1', \dots, \sigma_n')$ such that $\sigma_x' = \sigma_x^*$ if $x \in Z$, and $\sigma_x' = \sigma_x$ otherwise. Given $k \geq 1$, a strategy profile σ is a *k-strong Nash equilibrium* of $\mathcal{G}$ if, for any strategy profile σ^* and any $Z \subseteq V$ such that $|Z| \leq k$, there exists $x \in Z$ such that $u_x(\sigma) \geq u_x(\sigma \xrightarrow{Z} \sigma^*)$. We denote the (possibly empty) set of k-strong Nash equilibria of $\mathcal{G}$ by $\mathsf{NE}_k(\mathcal{G})$.

k-strong Price of Anarchy (PoA) and Price of Stability (PoS). The *k-strong Price of Anarchy* of a game $\mathcal{G}$ is defined as $\mathsf{PoA}_k(\mathcal{G}) := \max_{\sigma\in\mathsf{NE}_k(\mathcal{G})} \frac{\mathsf{OPT}(\mathcal{G})}{\mathsf{SW}(\sigma)}$, i.e., it is the worst-case ratio between the optimal social welfare and the social welfare of a k-strong Nash equilibrium. The *k-strong Price of Stability* of game $\mathcal{G}$ is defined as $\mathsf{PoS}_k(\mathcal{G}) := \min_{\sigma\in\mathsf{NE}_k(\mathcal{G})} \frac{\mathsf{OPT}(\mathcal{G})}{\mathsf{SW}(\sigma)}$, i.e., it is the best-case ratio between the optimal social welfare and the social welfare of a k-strong Nash equilibrium.

3. Our Contribution

We study the inefficiency of k-stable Nash equilibria of d-distance polymatrix coordination games and provide suitable bounds on both the Price of Anarchy and the Price of Stability. To the best of our knowledge, there are no previous results of this kind in the literature that would apply to our model. In Section 3.1, we give upper and lower bounds for bounded-degree graphs, with the gap being reasonably small, and in Section 3.2, a tight bound on the Price of Anarchy for general graphs. Finally, in Section 3.3, we show that in general graphs, the Price of Stability is asymptotically equal to the Price of Anarchy, meaning that the inefficiency of k-strong equilibria is fully characterized. We remark that our results also apply to the subclass of the classical polymatrix coordination games, for which, in turn, we get the first upper and lower bounds on the Price of Anarchy for bounded-degree graphs and the first asymptotically tight lower bound on the Price of Stability for general graphs.

3.1. k-strong PoA of Bounded-Degree Graphs

In this section we compute upper and lower bounds on the k-strong Price of Anarchy of bounded-degree graphs. More formally, a game $\mathcal{G}$ is Δ-*bounded-degree* if the degree of each node/player $x \in V$ in graph G is at most Δ.

First we remark that for $k = 1$, $d \geq 1$, and $\Delta = 1$, there exists a simple Δ-bounded-degree d-distance polymatrix coordination game $\mathcal{G}$ such that $\mathsf{PoA}_k(\mathcal{G}) = \infty$ [7]. Thus, we will only consider the case of $k \geq 2$. Furthermore, if $\Delta = 1$, w.l.o.g. we can assume that the graph consists of 2 agents and an edge between them. This special case is encompassed by Section 3.2, so here we will assume that $\Delta \geq 2$.

Theorem 1. *For any integer $k \geq 2$ and any Δ-bounded-degree d-distance polymatrix coordination game $\mathcal{G}$ having a distance-factors sequence $(\alpha_h)_{h\in[d]}$, it holds that*

$$\mathsf{PoA}_k(\mathcal{G}) \leq 2 \sum_{h\in[d]} \alpha_h \cdot \Delta \cdot (\Delta - 1)^{h-1}. \tag{1}$$

Remark 1. From Eq. (1), notice that the k-strong price of anarchy of Δ-bounded-degree d-distance polymatrix coordination games, as a function of d, grows at most as $O((\Delta - 1)^d)$.

In the following theorem we provide a lower bound on the k-strong Price of Anarchy, relying on a nice construction from graph theory.

Theorem 2. *For any $k \geq 2$, $\Delta \geq 2$, $d \geq 1$, and any distance-factors sequence $(\alpha_h)_{h\in[d]}$, there exists a Δ-bounded-degree d-distance polymatrix coordination game $\mathcal{G}$ such that*

$$\mathsf{PoA}_k(\mathcal{G}) \geq \frac{\sum_{h\in[d]} \alpha_h \cdot \Delta \cdot (\Delta - 1)^{h-1}}{\sum_{h\in[d]} \alpha_h (\Delta - 1)^{\lfloor h/2 \rfloor}}. \tag{2}$$

Remark 2. Notice that, if all the distance-factors are not lower than a constant $c > 0$, from Eq. (2) we can conclude that the k-strong price of anarchy of Δ-bounded-degree d-distance polymatrix coordination games, as a function of d, can grow as $\Omega((\Delta - 1)^{d/2})$.

3.2. k-strong PoA of General Graphs

In this section, we provide tight bounds on the k-strong Price of Anarchy when there is no particular assumption on the underlying graph of the considered game. Such bounds depend on k, on the number of players n, and on the value α_2 of the distance-factors sequence.

Theorem 3. *For any integer $k \geq 2$ and any d-distance polymatrix coordination game $\mathcal{G}$ having a distance-factors sequence $(\alpha_h)_{h \in [d]}$, we have*

$$\mathsf{PoA}_k(\mathcal{G}) \leq \frac{(2 + \alpha_2 \cdot (n-2)) \cdot (n-1)}{k-1}.$$

In the following theorem, we provide a tight lower bound.

Theorem 4. *For any $k \geq 2$, $d \geq 1$, $n \geq 2$, and any distance-factors sequence $(\alpha_h)_{h \in [d]}$, there is a d-distance polymatrix coordination game $\mathcal{G}$ with*

$$\mathsf{PoA}_k(\mathcal{G}) \geq \frac{(2 + \alpha_2 \cdot (n-2)) \cdot (n-1)}{k-1}.$$

3.3. k-strong PoS of General Graphs

In this section, we show that there exists a d-distance polymatrix coordination game $\mathcal{G}$ such that $\mathsf{PoS}_k(\mathcal{G})$ is asymptotically equal to the upper bound on PoA_k shown in Theorem 3; thus we characterise entirely the inefficiency of d-distance polymatrix coordination games for general graphs. The modus operandi that we use to create the lower bound for PoS_k is to start from the lower bound instance on PoA_k provided in the proof of Theorem 4, in which the optimal outcome is a k-strong Nash equilibrium, and to suitably transform it in such a way that all the outcomes with social welfare close to the optimum cannot be stable.

Theorem 5. *For any $n \geq 6$, there exists a d-distance polymatrix coordination game $\mathcal{G}$ such that*

$$\mathsf{PoS}_k(\mathcal{G}) = \frac{2n - 3 \; + \; \alpha_2(n-2)(n-3/2)}{(1+\alpha_2)k}.$$

4. Conclusion and future works

In this work, we have introduced the class of d-distance polymatrix coordination games and studied their performance (by means of the k-strong Price of Anarchy and Stability). Some open problems left by our work are that of closing the gap between the upper and lower bound on the strong Price of Anarchy for bounded-degree graphs and providing better bounds on the strong Price of Stability specifically for the case of bounded-degree graphs. Another interesting research direction is extending the idea of obtaining utilities from non-neighbouring players (as in [18] and our work) to other graphical games [24, 25], and then studying the social performance of their equilibria in general graphs or specific topologies.

References

[1] R. De Benedictis, M. Castiglioni, D. Ferraioli, V. Malvone, M. Maratea, E. Scala, L. Serafini, I. Serina, E. Tosello, A. Umbrico, M. Vallati, Preface to the Italian Workshop on Planning and Scheduling, RCRA Workshop on Experimental evaluation of algorithms for solving problems with combinatorial explosion, and SPIRIT Workshop on Strategies, Prediction, Interaction, and Reasoning in Italy (IPS-RCRA-SPIRIT 2023), in: Proceedings of the Italian Workshop on Planning and Scheduling, RCRA Workshop on Experimental evaluation of algorithms for solving problems with combinatorial explosion, and SPIRIT Workshop on Strategies, Prediction, Interaction, and Reasoning in Italy (IPS-RCRA-SPIRIT 2023) co-located with 22th International Conference of the Italian Association for Artificial Intelligence (AI*IA 2023), 2023.

[2] E. Janovskaja, Equilibrium points in polymatrix games, Lithuanian Mathematical Journal 8 (1968) 381–384.

[3] J. Howson, Equilibria of polymatrix games, Management Sci. 18 (1972) 312–318.

[4] B. C. Eaves, Polymatrix games with joint constraints, SIAM Journal on Applied Mathematics 24 (1973) 418–423.

[5] J. Howson, R. Rosenthal, Bayesian equilibria of finite two-person games with incomplete information, Management Sci. 21 (1974) 313–315.

[6] D. Miller, S. Zucker, Copositive-plus Lemke algorithm solves polymatrix games, Oper. Res. Lett. 10 (1991) 285–290.

[7] M. Rahn, G. Schäfer, Efficient Equilibria in Polymatrix Coordination Games, in: Proc. 40th Intl. Symp. Math. Foundations of Computer Science (MFCS), 2015, pp. 529–541.

[8] Y. Cai, O. Candogan, C. Daskalakis, C. H. Papadimitriou, Zero-Sum Polymatrix Games: A Generalization of Minmax, Math. Oper. Res. 41 (2016) 648–655.

[9] A. Deligkas, J. Fearnley, R. Savani, P. G. Spirakis, Computing Approximate Nash Equilibria in Polymatrix Games, Algorithmica 77 (2017) 487–514.

[10] A. Deligkas, J. Fearnley, R. Savani, Tree Polymatrix Games Are PPAD-Hard, in: Proc. 47th Intl. Coll. Autom. Lang. Program. (ICALP), 2020, pp. 38:1–38:14.

[11] A. Aloisio, M. Flammini, B. Kodric, C. Vinci, Distance polymatrix coordination games, in: Proc. of the 30th International Joint Conference on Artificial Intelligence, IJCAI-21, 2021, pp. 3–9.

[12] A. Aloisio, Distance Hypergraph Polymatrix Coordination Games, in: Proc. 22nd Conf. Autonomous Agents and Multi-Agent Systems (AAMAS), 2023, pp. 2679–2681.

[13] J. H. Drèze, J. Greenberg, Hedonic Coalitions: Optimality and Stability, Econometrica 48 (1980) 987–1003.

[14] A. Aloisio, M. Flammini, C. Vinci, The Impact of Selfishness in Hypergraph Hedonic Games, in: Proc. 34th Conf. Artificial Intelligence (AAAI), 2020, pp. 1766–1773.

[15] A. Darmann, E. Elkind, S. Kurz, J. Lang, J. Schauer, G. J. Woeginger, Group activity selection problem with approval preferences, Int. J. Game Theory 47 (2018) 767–796.

[16] A. Darmann, J. Lang, Group activity selection problems, in: Trends in computational social choice, 2017, pp. 385–410.

[17] V. Bilò, A. Fanelli, M. Flammini, G. Monaco, L. Moscardelli, Optimality and Nash Stability in Additive Separable Generalized Group Activity Selection Problems, in: Proc. 28th Intl.

Joint Conf. Artif. Intell. (IJCAI), 2019, pp. 102–108.
[18] M. Flammini, B. Kodric, M. Olsen, G. Varricchio, Distance Hedonic Games, in: Proc. 19th Conf. Autonomous Agents and Multi-Agent Systems (AAMAS), 2020, pp. 1846–1848.
[19] H. Aziz, F. Brandl, F. Brandt, P. Harrenstein, M. Olsen, D. Peters, Fractional Hedonic Games, ACM Trans. Economics and Comput. 7 (2019) 6:1–6:29.
[20] E. Elkind, A. Fanelli, M. Flammini, Price of Pareto Optimality in hedonic games, Artif. Intell. 288 (2020) 103357.
[21] G. Monaco, L. Moscardelli, Y. Velaj, Stable outcomes in modified fractional hedonic games, Auton. Agents Multi Agent Syst. 34 (2020) 4.
[22] R. Carosi, G. Monaco, L. Moscardelli, Local Core Stability in Simple Symmetric Fractional Hedonic Games, in: Proc. 18th Conf. Autonomous Agents and Multi-Agent Systems (AAMAS), 2019, pp. 574–582.
[23] V. Bilò, A. Fanelli, M. Flammini, G. Monaco, L. Moscardelli, Nash Stable Outcomes in Fractional Hedonic Games: Existence, Efficiency and Computation, J. Artif. Intell. Res. 62 (2018) 315–371.
[24] M. Kearns, Graphical Games, Cambridge University Press, 2007, p. 159–180.
[25] V. Bilò, A. Fanelli, M. Flammini, L. Moscardelli, When ignorance helps: Graphical multicast cost sharing games, Theoret. Comput. Sci. 411 (2010) 660–671.