Generalized Distance Polymatrix Games⋆

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

[1]University of International Studies of Rome, Via Cristoforo Colombo, 200 – 00147, Roma, Italy
[2]Gran Sasso Science Institute, Viale Francesco Crispi, 7 - 67100 LAquila, Italy
[3]University of Salento, Piazza Tancredi, n.7 - 73100 Lecce, Italy

Abstract

We propose a new class of *generalized distance polymatrix games*, which extends distance polymatrix coordination games by allowing each subgame to be played by more than two agents. These games can be effectively modeled using hypergraphs, where each hyperedge represents a subgame played by its agents. Similar to distance polymatrix coordination games, the overall utility of a player x depends on the payoffs of subgames involving players within a certain distance from x. As in the original model, these payoffs are discounted proportionally by factors that depend on the distance of the corresponding hyperedges. After formalizing and motivating our model, we investigate the existence of exact and approximate strong equilibria. We also examine the degradation of social welfare using the standard measures of the Price of Anarchy and the Price of Stability, both for general and bounded-degree hypergraphs.

Keywords

Polymatrix Games, Price of Anarchy, Price of Stability

1. Introduction

Polymatrix games [1] are well-known *graphical games* [2] where each player chooses a pure strategy from a finite set, which she will play in all the binary games she is involved in. In the subclass of *polymatrix coordination games* [3], the interaction graph is undirected since the outcome of a binary game is the same for both players.

In this paper, we present and study a new, more general model called *generalized distance polymatrix games*, where each local game can involve more than two players, and the utility of an agent x can depend on games at a distance bounded by d. In this new model, the interaction graph is represented as an undirected hypergraph, with each hyperedge corresponding to a game played by the players it includes. Following the idea proposed in [4], the utility of an agent x is the sum of the outcomes of the games they participate in, plus a fraction of the outcomes of games played by other players within a distance of at most d from x. Additionally, each agent x receives an extra payoff that is a function of their chosen strategy.

Our model is related to polymatrix coordination games [3, 5] and the more recent distance polymatrix coordination games [4, 6], where the authors introduced the idea of distances.

ICTCS'24: Italian Conference on Theoretical Computer Science, September 11–13, 2024, Torino, Italy
*Corresponding author.
†These authors contributed equally.
alessandro.aloisio@unint.eu (A. Aloisio); michele.flammini@gssi.it (M. Flammini); cosimo.vinci@unisalento.it (C. Vinci)
0000-0003-3911-4008 (A. Aloisio); 0000-0003-0327-3728 (M. Flammini); 0000-0001-7741-9342 (C. Vinci)

Some preliminary results can be found in [7, 8]. Our studies are also related to *(symmetric) additively separable hedonic games* [9] and *hypergraph hedonic games* [10]. Another closely related model is the *group activity selection problem* [11, 12, 13]. Our model is also connected to *social context games* [14, 15]. The idea of obtaining utility from non-neighbouring players has also been analyzed for *distance hedonic games* [16]. They generalize *fractional hedonic games* [17, 18, 19, 20, 21] similar to how distance polymatrix games and our model do with polymatrix games.

2. Preliminaries

Given two integers $r \geq 1$ and $n \geq 1$, let $[n] = \{1, \ldots, n\}$ and define the falling factorial as $(n)_r := n \cdot (n-1) \cdot \ldots \cdot (n-r+1)$. A *weighted hypergraph* is a triple $\mathcal{H} = (V, E, w)$ consisting of a finite set $V = [n]$ of *nodes*, a collection $E \subseteq 2^V$ of *hyperedges*, and a *weight* $w : E \to \mathbb{R}$ associating a real value $w(e)$ with each hyperedge $e \in E$. For simplicity, when referring to weighted hypergraphs, we omit the term "weighted". The *arity* of a hyperedge e is its size $|e|$. An r*-hypergraph* is a hypergraph such that the arity of each hyperedge is at most r, where $2 \leq r \leq n$. A *complete* r*-hypergraph* is a hypergraph (V, E, w) such that $E := \{U \subseteq V : |U| \leq r\}$. A *uniform* r*-hypergraph* is a hypergraph such that the arity of each hyperedge is r. An *undirected graph* is a uniform 2-hypergraph. A hypergraph is said to be Δ-regular if each of its nodes is contained in exactly Δ hyperedges. It is said to be *linear* if any two of its hyperedges share at most one node. A hypergraph is called a *hypertree* if it admits a host graph T such that T is a tree. Given two distinct nodes u and v in a hypergraph $\mathcal{H}$, a $u - v$ *simple path* of length l in $\mathcal{H}$ is a sequence of distinct hyperedges $(e_1, \ldots, e_l)$ of $\mathcal{H}$, such that $u \in e_1$, $v \in e_l$, $e_i \cap e_{i+1} \neq \emptyset$, for every $i \in [l-1]$, and $e_i \cap e_j = \emptyset$ whenever $j > i+1$. The distance from u to v, $d(u, v)$, is the length of the shortest $u - v$ simple path in $\mathcal{H}$. A *cycle* in a hypergraph $\mathcal{H}$ is defined as a simple path $(e_1, \ldots, e_l)$, but the further condition $e_1 \cap e_l \neq \emptyset$ must hold. This definition of cycle is originally due to Berge, and it can be also stated as an alternating sequence of $v_1, e_1, v_2, \ldots, v_n, e_n$ of distinct vertices v_i and distinct hyperedges e_i so that each e_i contains both v_i and v_{i+1}. The *girth* of a hypergraph is the length of the shortest cycle it contains.

Generalized Distance Polymatrix Games. A *generalized distance polymatrix game* (or GDPG) $\mathcal{G} = (\mathcal{H}, (\Sigma_x)_{x \in V}, (w_e)_{e \in E}, (p_x)_{x \in V}, (\alpha_h)_{h \in [d]})$, is a game based on an r-hypergraph $\mathcal{H}$, and defined as follows:

Agents: The set of agents is $V = [n]$, i.e., each node corresponds to an agent. We reasonably assume that $n \geq r \geq 2$.

Strategy profile or outcome: For any $x \in V$, Σ_x is a finite set of *strategies* of player x. A *strategy profile or outcome* $\boldsymbol{\sigma} = (\sigma_1, \ldots, \sigma_n)$ is a configuration in which each player $x \in V$ plays strategy $\sigma_x \in \Sigma_x$.

Weight function: For any hyperedge $e \in E$, let $w_e : \times_{x \in e} \Sigma_x \to \mathbb{R}_{\geq 0}$ be the *weight function* that assigns, to each subset of strategies σ_e played respectively by every $x \in e$, a *weight*

$w_e(\sigma_e) \geq 0$. In what follows, for the sake of brevity, given any strategy profile $\boldsymbol{\sigma}$, we will often denote $w_e(\sigma_e)$ simply as $w_e(\boldsymbol{\sigma})$.

Preference function: 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 σ_x played by player x, a non-negative real value $p_x(\sigma_x)$, called *player-preference*. In what follows, for the sake of brevity, given any strategy profile $\boldsymbol{\sigma}$, we will often denote $p_x(\sigma_x)$ simply as $p_x(\boldsymbol{\sigma})$.

Distance-factors sequence: 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 $1 = \alpha_1 \geq \alpha_2 \geq \ldots \geq \alpha_d \geq 0$.

Utility function: For any $h \in [d]$, let $E_h(x)$ be the set of hyperedges e such that the minimum distance between x and one of the players $v \in e$ 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 $\boldsymbol{\sigma}$ is defined as $u_x(\boldsymbol{\sigma}) := p_x(\boldsymbol{\sigma}) + \sum_{h\in[d]} \alpha_h \sum_{e\in E_h(x)} w_e(\boldsymbol{\sigma})$.

The *social welfare* $\mathsf{SW}(\boldsymbol{\sigma})$ of a strategy profile $\boldsymbol{\sigma}$ is defined as the sum of all the agents' utilities in $\boldsymbol{\sigma}$, i.e., $\mathsf{SW}(\boldsymbol{\sigma}) := \sum_{x\in V} u_x(\boldsymbol{\sigma})$. A *social optimum* of game $\mathcal{G}$ is a strategy profile $\boldsymbol{\sigma}^*$ that maximizes the social welfare. We denote by $\mathsf{OPT}(\mathcal{G}) = \mathsf{SW}(\boldsymbol{\sigma}^*)$ the corresponding value.

β-approximate k-strong Nash equilibrium. Given two strategy profiles $\boldsymbol{\sigma} = (\sigma_1, \ldots, \sigma_n)$ and $\boldsymbol{\sigma}^* = (\sigma_1^*, \ldots, \sigma_n^*)$, and a subset $Z \subseteq V$, let $\boldsymbol{\sigma} \xrightarrow{Z} \boldsymbol{\sigma}^*$ be the strategy profile $\boldsymbol{\sigma}' = (\sigma_1', \ldots, \sigma_n')$ such that $\sigma_x' = \sigma_x^*$ if $x \in Z$, and $\sigma_x' = \sigma_x$ otherwise. Given $k \geq 1$, a strategy profile $\boldsymbol{\sigma}$ is a *β-approximate k-strong Nash equilibrium* (or (β, k)-equilibrium) of $\mathcal{G}$ if, for any strategy profile $\boldsymbol{\sigma}^*$ and any $Z \subseteq V$ such that $|Z| \leq k$, there exists $x \in Z$ such that $\beta u_x(\boldsymbol{\sigma}) \geq u_x(\boldsymbol{\sigma} \xrightarrow{Z} \boldsymbol{\sigma}^*)$. We say that a player x *β-improves* from a deviation $\boldsymbol{\sigma} \xrightarrow{Z} \boldsymbol{\sigma}^*$ if $\beta u_x(\boldsymbol{\sigma}) < u_x(\boldsymbol{\sigma}')$. Informally, $\boldsymbol{\sigma}$ is a (β, k)-equilibrium if, for any coalition of at most k players deviating, there exists at least one player in the coalition that does not β-improve her utility by deviating. We denote the (possibly empty) set of (β, k)-equilibria of $\mathcal{G}$ by $\mathsf{NE}_k^\beta(\mathcal{G})$. Clearly, if $\beta = 1$, $\mathsf{NE}_k^\beta(\mathcal{G})$ contains all the k-strong equilibria, and when $\beta = 1$ and $k = 1$, it contains the classic Nash equilibria.

(β, k)-Price of Anarchy (PoA) and (β, k)-Price of Stability (PoS). The *(β, k)-Price of Anarchy* of a game $\mathcal{G}$ is defined as $\mathsf{PoA}_k^\beta(\mathcal{G}) := \max_{\boldsymbol{\sigma}\in\mathsf{NE}_k^\beta(\mathcal{G})} \frac{\mathsf{OPT}(\mathcal{G})}{\mathsf{SW}(\boldsymbol{\sigma})}$, i.e., it is the worst-case ratio between the optimal social welfare and the social welfare of a (β, k)-equilibrium. The *(β, k)-Price of Stability* of game $\mathcal{G}$ is defined as $\mathsf{PoS}_k^\beta(\mathcal{G}) := \min_{\boldsymbol{\sigma}\in\mathsf{NE}_k^\beta(\mathcal{G})} \frac{\mathsf{OPT}(\mathcal{G})}{\mathsf{SW}(\boldsymbol{\sigma})}$, i.e., it is the best-case ratio between the optimal social welfare and the social welfare of a (β, k)-Nash equilibrium. Clearly, $\mathsf{PoS}_k^\beta(\mathcal{G}) \leq \mathsf{PoA}_k^\beta(\mathcal{G})$, whereas both quantities are not defined if $\mathsf{NE}_k^\beta(\mathcal{G}) = \emptyset$.

3. Our Contribution

First, we analyze the existence of β-approximate k-strong equilibria and investigate the degradation of social welfare when a deviation from the current strategy profile can involve up to k agents. Consequently, we compute tight bounds on the resulting Price of Anarchy and Stability.

3.1. Existence of (β, k)-equilibria

Since (β, k)-equilibria may not exist since they cannot always exist even in polymatrix coordination games [3, 5], we provide some conditions on β that guarantee their existence.

We say that a game $\mathcal{G}$ has a *finite* (β, k)-*improvement property* (or (β, k)-FIP for short) if every sequence of (β, k)-improving deviations is finite. In such a case, we necessarily have that any (β, k)-FIP ends in a (β, k)-equilibrium, which implies the latter's existence, too.

For a given hyperedge e and a subset $Z \subseteq V$, let $n_h^Z(e) := |\{x \in Z \ : \ e \in E_h(x)\}|$, i.e., $n_h^Z(e)$ is the number of players $x \in Z$ that are at distance $h-1$ from e.

Theorem 1. *Let $\mathcal{G}$ be a GDPG. Then: i) $\mathcal{G}$ has the $(\beta, 1)$-FIP for every $\beta \geq 1$; ii) $\mathcal{G}$ has the (β, k)-FIP for every $\beta \geq \max_{\substack{Z \subseteq V: \\ |Z|=k}} \{\max_{e \in E}\{\sum_{h \in [d]} \alpha_h n_h^Z(e)\}\}$ and for every k.*

The value $\sum_{h \in [d]} \alpha_h n_h^Z(e)$ strictly depends on d and $n_h^Z(e)$. When $d = 1$, we have $\sum_{h \in [d]} \alpha_h n_h^Z(e) = n_1^Z(e) \leq |e|$ for every $e \in E$ and $Z \subseteq V$, so we can assume $\beta \geq r$. When the hypergraph of a game is a hyperlist, we have $\sum_{h \in [d]} \alpha_h n_h^Z(e) \leq 2r \sum_{h \in [d]} \alpha_h$, for every $e \in E$, and $Z \subseteq V$. When the hypergraph of a game is a hypertree of maximum degree Δ, we have $\sum_{h \in [d]} \alpha_h n_h^Z(e) \leq r \sum_{h \in [d]} \alpha_h r^{h-1} \Delta^{h-1}$, for every $e \in E$, and $Z \subseteq V$.

3.2. (β, k)-PoA and $(1, k)$-PoS of General Hypergraphs

We provide here tight upper and lower bounds for the (β, k)-Price of Anarchy when the hypergraph $\mathcal{H}$ of a game $\mathcal{G}$ is general. We also show a lower bound for the $(1, k)$-Price of Stability asymptotically equal to the upper bound for the $(1, k)$-Price of Anarchy. First we remark that for any integers $\beta \geq 1$, $r \geq 2$, $k < r$, and $n \geq r$, there exists a simple GDPG $\mathcal{G}$ with n agents such that $\mathsf{PoA}_k^\beta(\mathcal{G}) = \infty$. Thus, we will only take into account the estimation of the (β, k)-PoA for $k \geq r \geq 2$ since it is not possible to bound the (β, k)-PoA for $k < r$, not even for bounded-degree graphs and not even when $\Delta = 1$.

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

$$\mathsf{PoA}_k^\beta(\mathcal{G}) \leq \beta \frac{(n-1)_{r-1}}{(k-1)_{r-1}} (r + \alpha_2(n-2))$$

.

We continue by showing the following tight lower bound.

Theorem 3. *For every $\beta \geq 1$, every integers $r \geq 2$, $k \geq r$, $d \geq 1$, $n \geq k$, and every d-distance-factors sequence $(\alpha_h)_{h\in[d]}$, there is a GDPG $\mathcal{G}$ with*

$$\mathsf{PoA}_k^\beta(\mathcal{G}) \geq \beta \frac{(n-1)_{r-1}}{(k-1)_{r-1}} \left(r + \alpha_2(n-r)\right)$$

.

We conclude this section by providing a lower bound for $\mathsf{PoS}_k^1(\mathcal{G})$, which can be used alongside the upper bound in Theorem 2 to characterize the $(1, k)$-Price of Stability.

Theorem 4. *For any $n \geq 6$, there exists a GDPG $\mathcal{G}$ such that*

$$\mathsf{PoS}_k^1(\mathcal{G}) \geq \frac{n-r}{n-1} \frac{(n-1)_{r-1}}{(k-1)_{r-1}} \frac{(r + \alpha_2(n-r))}{2(1+\alpha_2)}$$

3.3. (β, k)-PoA of Bounded-Degree Hypergraphs

In this section, we analyze the (β, k)-Price of Anarchy for games whose hypergraphs have bounded-degree. We also say that a game $\mathcal{G}$ is Δ-bounded degree if the degree of every node in the underlying hypergraph is at most Δ. Here, we will only focus on the cases where $k \geq r$, as already observed, and $\Delta \geq 2$, since the case when $\Delta = 1$ is encompassed by Section 3.2.

Theorem 5. *For every Δ-bounded-degree GDPG $\mathcal{G}$, with distance-factor sequence $(\alpha_h)_{h\in[d]}$, and for every $k \geq r$, it holds that*

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

.

We continue by showing the following tight lower bound.

Theorem 6. *For every $\beta \geq 1$, any integers $k \geq r$, $\Delta \geq 3$, $d \geq 1$, and any distance-factors sequence $(\alpha_h)_{h\in[d]}$, there exists a Δ-bounded-degree GDPG $\mathcal{G}$ such that*

$$\mathsf{PoA}_k^\beta(\mathcal{G}) \geq \frac{\beta \cdot \sum_{h\in[d]} \alpha_h \Delta(\Delta-1)^{h-1} b^{h-1}}{1 + \sum_{h=1}^{d-1} \alpha_{h+1}(2(\Delta-1)^{\lfloor (h+1)/2 \rfloor}(r-1)^{\lfloor (h+1)/2 \rfloor - 1} + 2(\Delta-1)^{\lfloor h/2 \rfloor - 1}(r-1)^{\lfloor h/2 \rfloor})}$$

.

Remark 1. Please note that, if all the distance-factors are not lower than a constant $c > 0$, from Theorem 6 we can conclude that the (β, k)-price of anarchy of Δ-bounded-degree GDPG, as a function of d, can grow as $\Omega(\beta(\Delta-1)^{d/2}(r-1)^{d/2})$.

4. Conclusion and future works

This study leaves some open problems, such as (i) closing the gap between the upper and the lower bound on the Price of Anarchy for bounded-degree hypergraphs; (ii) extending the results on the Price of Stability to values of β greater than one; and (iii) computing a lower bound on the Price of Stability for bounded-degree hypergraphs.

Acknowledgments

This work is partially supported by the project 'Soluzioni innovative per il problema della copertura nelle multi-interfacce e relative varianti', UNINT, GNCS-INdAM, and European Union, PON Ricerca e Innovazione 2014-20 TEBAKA - Fondo Sociale Europeo 2014-20.

References

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

[2] M. J. Kearns, M. L. Littman, S. P. Singh, Graphical Models for Game Theory, in: Proc. 17th Conf. Uncertainty in Artif. Intell. (UAI), 2001, pp. 253–260.

[3] 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.

[4] 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.

[5] K. R. Apt, B. de Keijzer, M. Rahn, G. Schäfer, S. Simon, Coordination games on graphs, Int. J. Game Theory 46 (2017) 851–877.

[6] A. Aloisio, M. Flammini, B. Kodric, C. Vinci, Distance polymatrix coordination games (short paper), in: SPIRIT co-located with 22nd International Conf. AIxIA 2023, November 7-9th, 2023, Rome, Italy, volume 3585, 2023.

[7] A. Aloisio, Distance hypergraph polymatrix coordination games, in: Proc. 22nd Conf. Autonomous Agents and Multi-Agent Systems (AAMAS), 2023, pp. 2679–2681.

[8] A. Aloisio, M. Flammini, C. Vinci, Generalized distance polymatrix games, in: SOFSEM 2024: Theory and Practice of Computer Science - 49th International Conference on Current Trends in Theory and Practice of Computer Science, SOFSEM 2024, Cochem, Germany, February 19-23, 2024, Proceedings, volume 14519, Springer, 2024, pp. 25–39.

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

[10] 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.

[11] A. Darmann, E. Elkind, S. Kurz, J. Lang, J. Schauer, G. J. Woeginger, Group Activity Selection Problem, in: Proc. 8th Intl. Workshop Internet & Network Economics (WINE), volume 7695, 2012, pp. 156–169.

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

[13] 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.

[14] I. Ashlagi, P. Krysta, M. Tennenholtz, Social Context Games, in: Proc. 4th Intl. Workshop Internet & Network Economics (WINE), 2008, pp. 675–683.

[15] V. Bilò, A. Celi, M. Flammini, V. Gallotti, Social context congestion games, Theoret. Comput. Sci. 514 (2013) 21–35.
[16] 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.
[17] 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.
[18] 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.
[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.