On a Network Generalization of the Minmax Theorem

Constantinos Daskalakis[1] and Christos H. Papadimitriou[2*]

[1] Microsoft Research, New England
[2] U.C. Berkeley

Abstract. We consider graphical games in which the edges are zero-sum games between the endpoints/players; the payoff of a player is the sum of the payoffs from each incident edge. Such games are arguably very broad and useful models of networked economic interactions. We give a simple reduction of such games to two-person zero-sum games; as a corollary, a mixed Nash equilibrium can be computed efficiently by solving a linear program and rounding off the results. Our results render polynomially efficient, and simplify considerably, the approach in [3].

1 Introduction

In 1928, von Neumann proved that every two-person zero-sum game has the minmax property [8], and thus a randomized equilibrium — which, we now know, is easily computable via linear programming. According to Aumann, two-person strictly competitive games — that is zero-sum games (see the discussion in the last section) — are "one of the few areas in game theory, and indeed in the social sciences, where a fairly sharp, unique prediction is made" [2]. In this paper, we present a sweeping generalization of this class to multi-player games played on a network.

Networked Interactions. In recent years, with the advent of the Internet and the many kinds of networks it enables, there has been increasing interest in games in which the players are nodes of a graph, and payoffs depend on the actions of a player's neighbors [6]. One interesting class of such games are the *graphical polymatrix games*, in which the edges are two-person games, and, once all players have chosen an action, the payoff of each player is the sum of the payoffs from each game played with each neighbor. For example, games of this sort with coordination games at the edges are useful for modeling the spread of ideas and technologies over social networks [7].

But what if the games at the edges are zero-sum — that is, we have a *network of competitors?* Do von Neumann's positive results carry over to this interesting case? Let us examine a few simple examples. If the network consists of isolated edges, then of course we have many independent zero-sum games and we are

* Supported by NSF grants CCF-0635319, CCF-0515259, a gift from Yahoo! Research, and a MICRO grant.

done. The next simplest case is the graph consisting of two adjacent edges. It turns out that in this case too von Neumann's ideas work: We could write the game, from the middle player's point of view, as a linear program seeking the mixed strategy x such that

$$\begin{aligned} \max\ & z_1 + z_2 \\ \text{subject to}\ & A_1 x \geq z_1 \\ & A_2 x \geq z_2, \end{aligned}$$

where A_1 and A_2 are the middle player's payoff matrices against the two other players. In other words, the middle player assumes that his two opponents will each punish her separately as much as they can, and seeks to minimize the total damage. In fact, a little thought shows that this idea can be generalized to any star network.

But what if the network is a triangle, for example? Now the situation becomes more complicated. For example, if player u plays matching pennies, say, with players v and w (take the stakes of the game with v to be higher than the stakes of the game with w), while v and w play between them, for much higher stakes, a game that rewards v for playing *heads*, then v cannot afford to pay attention to u, and u can steal a positive payoff along the edge (u, v), so that her total payoff is positive — despite the fact that she is playing two matching pennies games. Is there a general method for computing Nash equilibria in such three-player zero-sum polymatrix games? Or is this problem PPAD-complete?

Our main result (Theorem 2) is a reduction implying that *in any zero-sum graphical polymatrix game a Nash equilibrium can be computed in polynomial time,* by simply solving a two-player zero-sum game and rounding off the equilibrium. In other words, we show that there is a very broad and natural class of tractable network games to which von Neumann's method applies rather directly. The basic idea of the reduction is very simple: We create two players whose strategy set is equal to the *union* of the actions of all players, and have both of them "represent" all players. To make sure that the two players randomize evenly between the players they represent, we make them play, on the side, a high-stakes game of generalized rock-paper-scissors. It is not hard to see that any minmax strategy of this two-person zero-sum game can be made (by increasing the stakes of the side game) arbitrarily close to a Nash equilibrium of the original game.[3]

We prove our main result in Section 2. In Section 3 we show an interesting consequence: if the nodes of the network run any distributed iterative learning algorithm of the bounded regret variety known to perform well in many contexts, then the whole game converges to the Nash equilibrium (Theorem 3).

[3] Ilan Adler (private communication, April 2009) pointed out to us a proof of our main result by a direct reduction to linear programming: Formulate the two-player game (without the generalized rock-paper-scissors part) as a linear program, adding constraints which require that each of the two players assigns the same total probability mass to the strategies of each of the players it represents.

Related work. In a very interesting paper [3] (which we discovered after we had proved our results...), Bregman and Fokin present a general approach to solving what they call *separable zero-sum games:* multiplayer games that are zero-sum, and in which the payoff of a player is the sum of the payoffs of the player's interactions with each other player. Their approach is to formulate such games as a linear program with huge dimensions but low rank, and then solve it by a sequence of reductions to simpler and simpler linear programs that can be solved by the column generation version of the simplex method in a couple of special cases, one of which is our zero-sum polymatrix games. Even though their technique does not amount to a polynomial-time algorithm, we believe that it can be turned into one by a sophisticated application of the ellipsoid method and multiple layers of separating hyperplane generation algorithms. In contrast, our method is a very simple and direct reduction to two-player zero-sum games.

Definitions. An *n-player zero-sum graphical polymatrix game* is defined in terms of an undirected graph $G = (V, E)$, where $V := [n]$ is the set of players, and, for each edge $[u, v] \in E$, an $m_u \times m_v$ real matrix $A^{u,v}$ and another $A^{v,u} = -(A^{u,v})^{\mathrm{T}}$. That is, each player/node u has a set of actions, $[m_u]$, and each edge is a zero-sum game played between its two endpoints. Given any mapping f from V to the natural numbers such that $f(u) \in [m_u]$ for all $u \in V$ — that is, any choice of actions for the players, the payoff of player $u \in V$ is defined as

$$P_u[f] = \sum_{[u,v]\in E} A^{u,v}_{f(u),f(v)}.$$

In other words, the payoff of each player is the sum of all payoffs of the zero-sum games played with the player's neighbors.

In any game, a *(mixed) Nash equilibrium* is a distribution on actions for each player, such that, for each player, all actions with positive probabilities are best responses in expectation. In an *ϵ-Nash equilibrium*, all actions played by a player with positive probability give her expected utility which is within an additive ϵ from the expected utility given by the best response. A weaker but related notion of approximation is the notion of an *ϵ-approximate Nash equilibrium*, in which the mixed strategy of a player gives her expected utility that is within an additive ϵ from the expected utility of the best response. Clearly, an *ϵ-Nash equilibrium* is also an *ϵ-approximate Nash equilibrium*; but the opposite implication is not always true. Nevertheless, the two notions are computationally related as follows.

Proposition 1 *[4] Given an ϵ-approximate Nash equilibrium of an n-player game, we can compute in polynomial time a $\sqrt{\epsilon} \cdot (\sqrt{\epsilon} + 1 + 4(n-1)\alpha_{\max})$-Nash equilibrium, where $\alpha_{\max}$ is the magnitude of the maximum in absolute value possible utility of a player in the game.*

2 Main Result

Theorem 2. *There is polynomial-time reduction from any zero-sum graphical polymatrix game $\mathcal{GG}$ to a symmetric zero-sum bimatrix game $\mathcal{G}$, such that from* *any Nash equilibrium of $\mathcal{G}$ one can recover in polynomial time a Nash equilibrium of $\mathcal{GG}$.*

Proof of Theorem 2: In our construction we use a generalization of the well known rock-paper-scissors game, defined below.

Definition 1 (Generalized Rock-Paper-Scissors). *For an odd integer $n > 0$, the* n-strategy rock-paper-scissors game *is a symmetric zero-sum bimatrix game $(\Gamma, -\Gamma)$ with n strategies per player such that for all $u, v \in [n]$:*

$$\Gamma_{u,v} = \begin{cases} +1, & \text{if } \; v = u+1 \mod n \\ -1, & \text{if } \; v = u-1 \mod n \\ 0, & \text{otherwise.} \end{cases}$$

It is not hard to see that, for every odd n, the unique Nash equilibrium of the n-strategy generalized rock-paper-scissors game is the uniform distribution over both players' strategies. Now let $\mathcal{GG} = \{A^{u,v}\}_{[u,v] \in E}$ be an n-player zero-sum graphical polymatrix game with edge set E, whose u-th player has m_u strategies. Assuming without loss of generality that n is odd, let us define the embedding $\mathcal{G}$ of $\mathcal{GG}$ into the n-strategy rock-paper-scissors game with scaling parameter $M > 0$ as follows: $\mathcal{G} = (R, C)$ is an $\sum_u m_u \times \sum_u m_u$ bimatrix game, whose rows and columns are indexed by pairs $(u : i)$, of players $u \in [n]$ and strategies $i \in [m_u]$, such that, for all $u, v \in [n]$, $i \in [m_u]$, $j \in [m_v]$,

$$R_{(u:i),(v:j)} = M \cdot \Gamma_{u,v} + A^{u,v}_{i,j}$$
$$C_{(u:i),(v:j)} = -M \cdot \Gamma_{u,v} + A^{v,u}_{j,i}.$$

In the above, we take $A^{u,v}$ and $A^{v,u}$ to be the all-zero matrices if $[u, v] \notin E$. Observe that $\mathcal{G}$ is zero-sum and also symmetric, since the generalized rock-paper-scissors game is symmetric.

Lemma 1. *Let $n > 0$ be an odd integer, $\mathcal{GG} = \{A^{u,v}\}_{[u,v] \in E}$ a zero-sum graphical polymatrix game whose largest in absolute value payoff entry has magnitude M/L, and $\mathcal{G} = (R, C)$ the embedding of $\mathcal{GG}$ into the n-strategy rock-paper-scissors game, with scaling parameter M. Then for all $u \in [n]$, in any Nash equilibrium (x, y) of $\mathcal{G}$, $x_u, y_u \in (\frac{1}{n} - \frac{n}{L}, \frac{1}{n} + \frac{n}{L})$, where $x_u = \sum_{i \in [m_u]} x_{u:i}$ and $y_u = \sum_{i \in [m_u]} y_{u:i}$ is the probability mass assigned by x and y to the block of strategies $(u : \cdot)$.*

Proof of Lemma 1: Observe first that, since $\mathcal{G}$ is a symmetric zero-sum game, the value of both players is 0 in every Nash equilibrium. We will use this to argue that $x_u \geq x_{(u+2 \mod n)} - \frac{1}{L}$, for all $u \in [n]$, and similarly for y. This is enough to conclude the proof of the lemma. For a contradiction, suppose that, in some Nash equilibrium (x, y), $x_u < x_{(u+2 \mod n)} - \frac{1}{L}$, for some u. Then the payoff to the column player for playing strategy $(u + 1 \mod n : j)$, for any $j \in [m_{u+1 \mod n}]$, is at least

$$M x_{(u+2 \mod n)} - M x_u - \frac{M}{L} > 0.$$

Since (x, y) is an equilibrium, the expected payoff to the column player from y must be at least as large as the expected payoff from $(u + 1 \mod n : j)$, so in particular larger than 0. But this is a contradiction since we argued that in any Nash equilibrium of $\mathcal{G}$ the payoff of each player is 0. ■

We argue next that, given any Nash equilibrium (x, y) of $\mathcal{G}$, we can extract an approximate equilibrium of the game $\mathcal{GG}$ by assigning to each node u of $\mathcal{GG}$ the marginal distribution assigned by x to the block of strategies $(u : i)$, $i \in [m_u]$. For each node u, let us define the distribution $\hat{x}_u$ over $[m_u]$ as follows

$$\hat{x}_u(i) = \frac{x_{u:i}}{x_u}, \qquad \text{for all } i \in [m_u]. \tag{1}$$

Lemma 2. *In the setting of Lemma 1, if (x, y) is a Nash equilibrium of $\mathcal{G}$, then the collection of mixed strategies $\{\hat{x}_u\}_u$ is a $\frac{2M \cdot n^3}{L^2}$–Nash equilibrium of $\mathcal{GG}$.*

Proof of Lemma 2: Notice that, because $\mathcal{G}$ is a symmetric zero-sum game, if x is a minimax strategy of the row player, then x is also a minimax strategy of the column player. Hence, the pair of mixed strategies (x, x) is also a Nash equilibrium of $\mathcal{G}$. Now, for every node u of the polymatrix game, we are going to show that the collection $\{\hat{x}_u\}_u$ satisfies the equilibrium conditions at node u approximately. Indeed, because (x, x) is a Nash equilibrium of $\mathcal{G}$ it must be that, for all $i, j \in [m_u]$:

$$\mathcal{E}\left[\mathcal{P}_{u:i}\right] > \mathcal{E}\left[\mathcal{P}_{u:j}\right] \quad \Rightarrow \quad x_{u:j} = 0, \tag{2}$$

where

$$\mathcal{E}\left[\mathcal{P}_{u:i}\right] = \sum_v M \cdot \mathit{\Gamma}_{u,v} \cdot x_v + \sum_{[u,v] \in E} \sum_{\ell \in [m_v]} A^{u,v}_{i,\ell} \cdot x_{v:\ell}$$

is the expected payoff to the row player of $\mathcal{G}$ for playing strategy $(u : i)$. From Lemma 1 we have

$$\left| \sum_{[u,v] \in E} \sum_{\ell \in [m_v]} A^{u,v}_{i,\ell} \cdot x_{v:\ell} - \frac{1}{n} \sum_{[u,v] \in E} \sum_{\ell \in [m_v]} A^{u,v}_{i,\ell} \cdot \hat{x}_v(\ell) \right| \leq \frac{M \cdot n^2}{L^2}.$$

Hence, (2) implies

$$\frac{1}{n} \sum_{[u,v] \in E} \sum_{\ell \in [m_v]} A^{u,v}_{i,\ell} \cdot \hat{x}_v(\ell) > \frac{1}{n} \sum_{[u,v] \in E} \sum_{\ell \in [m_v]} A^{u,v}_{j,\ell} \cdot \hat{x}_v(\ell) + \frac{2M \cdot n^2}{L^2} \;\Rightarrow\; \hat{x}_u(j) = 0,$$

which is equivalent to

$$\mathcal{E}\left[P_{u:i}\right] > \mathcal{E}\left[P_{u:j}\right] + \frac{2M \cdot n^3}{L^2} \quad \Rightarrow \quad \hat{x}_u(j) = 0, \tag{3}$$

where $\mathcal{E}\left[P_{u:i}\right]$ is the expected payoff of node u in $\mathcal{GG}$ for playing pure strategy i, if the other players play according to the collection of mixed strategies $\{\hat{x}_v\}_{v\neq u}$. Since (3) holds for all $u \in [n]$, $i,j \in [m_u]$, the collection $\{\hat{x}_u\}_u$ is a $\frac{2M\cdot n^3}{L^2}$-Nash equilibrium of $\mathcal{GG}$. ■

Choosing $M = 2^{q(|\mathcal{GG}|)} 2n^3 u_{\max}^2$ and $L = \frac{M}{u_{\max}}$, where $q(|\mathcal{GG}|)$ is some polynomial in the size of $\mathcal{GG}$, and $u_{\max}$ the magnitude of the maximum in absolute value entry in the payoff tables of $\mathcal{GG}$, the collection of mixed strategies $\{\hat{x}_u\}_u$ obtained from a Nash equilibrium (x,y) of $\mathcal{G}$, constitutes a $2^{-q(|\mathcal{GG}|)}$-Nash equilibrium of the game $\mathcal{GG}$. If $q(\cdot)$ is a sufficiently large polynomial, then a $2^{-q(|\mathcal{GG}|)}$-Nash equilibrium of $\mathcal{GG}$ can be transformed in polynomial time to an exact equilibrium. To see this, let us consider the following linear program with respect to the variables z and $\{\hat{y}_u\}_u$, where $\hat{y}_u$ is a distribution over $[m_u]$:

$$\min z$$

$$\text{s.t.} \quad \sum_{[u,v]\in E} \sum_{\ell \in [m_v]} A^{u,v}_{i,\ell} \cdot \hat{y}_v(\ell) \geq \sum_{[u,v]\in E} \sum_{\ell \in [m_v]} A^{u,v}_{j,\ell} \cdot \hat{y}_v(\ell) - z, \qquad \begin{array}{r} \forall u \in [n], \\ i \in \operatorname{supp}(\hat{x}_u), \\ j \in [m_u]. \end{array} \tag{4}$$

In LP (4), $\operatorname{supp}(\hat{x}_u)$ denotes the support of the distribution $\hat{x}_u$. Observe in particular that $(2^{-q(|\mathcal{GG}|)}, \{\hat{x}_u\}_u)$ is a solution of LP (4) with objective value $2^{-q(|\mathcal{GG}|)}$. But, let us assume that $q(\cdot)$ has been chosen to be larger than the bit complexity of any optimal solution to LP (4) (for any possible set of supports $\{\operatorname{supp}(\hat{x}_u)\}_u$). It follows then that the optimal solution to LP (4) has objective value $z = 0$, so that the corresponding collection $\{\hat{y}_u\}_u$ is an exact Nash equilibrium of $\mathcal{GG}$. ■

3 Distributed Learning

One of the more subtle advantages of two-person zero-sum games is that a large variety of learning algorithms converge to the Nash equilibrium. Hence in this section we study the behavior arising if every player in a zero-sum graphical polymatrix game runs a no-regret learning algorithm.

Definition 2 (No-Regret Behavior). *Let every node $u \in V$ of a graphical polymatrix game choose a mixed strategy x_u^t, at every time step $t = 1, 2, \ldots$. We say that the sequence of strategies $(x_u^t)_t$ chosen by u is a* no-regret sequence, *if for every mixed strategy x of player u and all times T*

$$\sum_{t=1}^{T} \left(\sum_{[u,v]\in E} (x_u^t)^{\mathrm{T}} \cdot A^{u,v} \cdot x_v^t \right) \geq \sum_{t=1}^{T} \left(\sum_{[u,v]\in E} x^{\mathrm{T}} \cdot A^{u,v} \cdot x_v^t \right) - o(T), \tag{5}$$

where the function $o(T)$ could depend on the number strategies available to player u, the number of neighbors of u and magnitude of the maximum in absolute value entry in the matrices $A^{u,v}$. The function $o(T)$ is called the regret of player u at time T.

Example 1 (Multiplicative Weights-Update Algorithm). In the multiplicative weights-update algorithm (see for example [5]) each player maintains a mixed strategy. At each period, each probability is multiplied by a factor exponential in the utility the corresponding strategy would yield against the opponent's mixed strategy (and the probabilities are renormalized). If every node in a zero-sum graphical polymatrix game runs such an algorithm, then the resulting regret is $O((\sqrt{T \cdot \log m_u} + \log m_u) \cdot d_u \cdot \alpha^u_{\max})$, where m_u is the number of strategies available to player u, d_u is the degree of u, and $\alpha^u_{\max}$ is the magnitude of the largest in absolute value entry in the payoff matrices $\{A^{u,v}\}_{[u,v]\in E}$.

Our main result is the following.

Theorem 3. *Suppose that every node $u \in V$ of a zero-sum graphical polymatrix game $\mathcal{GG}$ plays a no-regret sequence of strategies $(x^t_u)_{t=1,2,\ldots}$, with regret $g(T) = o(T)$. Then, for all T, the set of strategies $\bar{x}^T_u = \frac{1}{T}\sum_{t=1}^{T} x^t_u$, $u \in V$, is a $\left(2.3 \cdot n \cdot \frac{g(T)}{T} + \frac{2}{T}\right)$-approximate Nash equilibrium of the game.*

Proof. Our proof plan is the following: Using the no-regret strategy sequences of the players of $\mathcal{GG}$ we are going to define no-regret strategy sequences for the players of the symmetric zero-sum bimatrix game $\mathcal{G}$ defined in the proof of Theorem 2. We are going to show then that the time-averages of these sequences comprise an approximate equilibrium of the game $\mathcal{G}$, if M is sufficiently large. Going back to the game $\mathcal{GG}$ using the mapping (1), we will then deduce that the time-averages of the original sequences need to also comprise an approximate equilibrium of the game $\mathcal{GG}$.

To define the no-regret sequences of the players of the bimatrix game $\mathcal{G}$, it is tempting to take, at every time step, the (uniform) average of the strategies of the players of $\mathcal{GG}$. That is, for every time $t = 1, 2, \ldots$, assign to both players of $\mathcal{G}$ the strategy x^t, such that $x^t_{u:i} = \frac{1}{n} x^t_u(i)$, for all u and i. This, however, may result in large regrets for the players of $\mathcal{G}$ (essentially because the payoffs of the side game are eliminated in this accounting, and the two players will tend to skew their distributions towards the most "lucrative" of the players that they represent). We define instead a non-uniform averaging with weights selected by solving yet another related game, in a manner that depends on the payoffs of the nodes of $\mathcal{GG}$ under the no-regret sequences.

Let us denote the average payoff of player u over the period $t = 1, \ldots, T$ as

$$\bar{P}^T_u := \frac{1}{T}\sum_{t=1}^{T}\left(\sum_{[u,v]\in E} (x^t_u)^{\mathrm{T}} \cdot A^{u,v} \cdot x^t_v\right).$$

Let also $\alpha_{\max}$ be the magnitude of the largest in absolute value entry in the payoff tables of the game $\mathcal{GG}$. We show the following lemma.

Lemma 3. *For any $Z > n^2$ and $M > 2nZ \cdot \alpha_{\max}$, there exist $c > 0$, and positive weights $\{k_u > 0 : u \in V\}$, such that for all u:*

$$M \cdot \left(\frac{1}{k_{(u+1 \bmod n)}} - \frac{1}{k_{(u-1 \bmod n)}}\right) = -\frac{1}{n}\bar{P}^T_u + c; \qquad (6)$$

$$\frac{1}{k_u} \in \left[\frac{1}{n} - \frac{n}{Z}, \frac{1}{n} + \frac{n}{Z}\right] \quad and \quad \sum_u \frac{1}{k_u} = 1. \tag{7}$$

Recall that we have identified the vertices of $\mathcal{GG}$ with the integers in $[n]$, and without loss of generality let us assume that n is odd. Also, for conciseness, in the remaining of the proof of Theorem 3 we are going to omit "mod n." We proceed with the proof of Lemma 3.

Proof of Lemma 3: We define a symmetric bimatrix game $\mathcal{G}'$, with n strategies per player corresponding to the different nodes of the game $\mathcal{GG}$. The payoff matrices (R, C) of the row and column players of $\mathcal{G}'$ are defined as follows. For all $u, v \in V$:

$$R_{u,v} = M \cdot \mathit{\Gamma}_{u,v} + \frac{1}{n}\bar{P}_u^T; \quad C_{u,v} = -M \cdot \mathit{\Gamma}_{u,v} + \frac{1}{n}\bar{P}_v^T;$$

where $\mathit{\Gamma}_{u,v}$ is the payoff matrix defined in the proof of Theorem 2. Since the game $\mathcal{G}'$ is symmetric, there exists a symmetric equilibrium (x, x), where $x = (x_u)_{u \in V}$. We will argue first that $x_u \geq \frac{1}{n} - \frac{1}{Z}$, for all $u \in V$; from this we can easily deduce that $x_u \in \left[\frac{1}{n} - \frac{n}{Z}, \frac{1}{n} + \frac{n}{Z}\right]$, for all u. Take $u \in \arg\min_u\{x_u\}$ and suppose that $x_u < \frac{1}{n} - \frac{1}{Z}$. Then there exists a pair of nodes v and $(v+2)$ such that $x_{v+2} - x_v > \frac{1}{nZ}$. Indeed, if $x_{v+2} - x_v \leq \frac{1}{nZ}$ for all v, it would be easy to deduce (because n is odd) that $x_v \leq x_u + \frac{1}{Z} < \frac{1}{n}$, for all v, which is clearly impossible. Now, given that $x_{v+2} - x_v > \frac{1}{nZ}$, the utility of the players of the game $\mathcal{G}'$ for playing pure strategy $v+1$ is

$$M \cdot (x_{v+2} - x_v) + \frac{1}{n}\bar{P}_{v+1}^T > \frac{M}{nZ} - \alpha_{\max} > \alpha_{\max},$$

since $\alpha_{\max}$ is a bound on the absolute value of every entry in the payoff matrices of the game $\mathcal{GG}$, every node has at most n neighbors, and $\frac{M}{nZ} > 2\alpha_{\max}$. On the other hand, since $u \in \arg\min_u\{x_u\}$ it follows that the payoff of the players of the game $\mathcal{G}'$ for playing pure strategy $u-1$ is

$$M \cdot (x_u - x_{u-2}) + \frac{1}{n}\bar{P}_{u-1}^T \leq \alpha_{\max}.$$

Since (x, x) is an equilibrium, it must be that $x_{u-1} = 0$. It follows that there must exist some w such that $x_w = 0$ and $x_{w-1} \neq 0$. But, the utility of the players of the game $\mathcal{G}'$ for playing pure strategy $w-1$ is

$$M \cdot (x_w - x_{w-2}) + \frac{1}{n}\bar{P}_{w-1}^T \leq \alpha_{\max}.$$

And, using again the fact that the utility for playing $v+1$ is larger than $\alpha_{\max}$, it follows that $x_{w-1} = 0$ (a contradiction). This finishes the proof of Assertion (7) taking $k_u := x_u^{-1}$, for all u.

Now, we need to justify (6). Since $x_u > 0$ for all u, it follows that the expected payoff for playing every u is the same. So, there exists c such that, for all u, $M \cdot (x_{u+1} - x_{u-1}) + \dfrac{1}{n}\bar{P}_u^T = c$. Assertion (6) then follows. ■

Now let us choose $Z = n^2 T \mathit{\Lambda}$ (where $\mathit{\Lambda} > 1$ will be decided later), $M > 2nZ \cdot \alpha_{\max}$, and let us define strategies for the players of $\mathcal{G}$ by averaging the strategies of the nodes of $\mathcal{GG}$ with the weights $\{1/k_u\}_u$ given by Lemma 3. That is, for all t, we define the strategy x^t for each player of $\mathcal{G}$ as follows:

$$x^t_{u:i} = \frac{1}{k_u} x^t_u(i), \text{ for all } u \in [n], i \in [m_u].$$

We show that if both players of $\mathcal{G}$ adopt the sequence of strategies $x^t, t = 1, 2, \ldots$, defined above then the regret of each at time T is at most $\left(\frac{g(T)}{n} + \frac{2\alpha_{\max}}{\Lambda}\right)$.

Lemma 4. *For all mixed strategies z:*

$$\sum_{t=1}^{T} (x^t)^{\mathrm{T}} \cdot R \cdot x^t \geq z^{\mathrm{T}} \cdot R \cdot \left(\sum_{t=1}^{T} x^t\right) - \frac{g(T)}{n} - \frac{2\alpha_{\max}}{\Lambda}, \tag{8}$$

$$\sum_{t=1}^{T} (x^t)^{\mathrm{T}} \cdot C \cdot x^t \geq \left(\sum_{t=1}^{T} x^t\right)^{\mathrm{T}} \cdot C \cdot z - \frac{g(T)}{n} - \frac{2\alpha_{\max}}{\Lambda}. \tag{9}$$

Proof of Lemma 4: Since $\mathcal{G}$ is symmetric it is enough to justify (8). Indeed, for the left hand side we have:

$$\begin{aligned}
\sum_{t=1}^{T} (x^t)^{\mathrm{T}} \cdot R \cdot x^t &= \sum_{t=1}^{T} \sum_{u\in[n]} \frac{1}{k_u} \left(M \cdot \left(\frac{1}{k_{u+1}} - \frac{1}{k_{u-1}} \right) + \sum_{[u,v]\in E} \frac{1}{k_v} (x^t_u)^{\mathrm{T}} A^{u,v} x^t_v \right) \\
&= \sum_{t=1}^{T} \sum_{u\in[n]} \frac{1}{k_u} \left(-\frac{1}{n} \bar{P}^T_u + c + \sum_{[u,v]\in E} \left(\frac{1}{n} \pm \frac{n}{Z} \right) (x^t_u)^{\mathrm{T}} A^{u,v} x^t_v \right) \\
&\geq \sum_{t=1}^{T} \sum_{u\in[n]} \frac{1}{k_u} \left(-\frac{1}{n} \bar{P}^T_u + c + \frac{1}{n} \sum_{[u,v]\in E} (x^t_u)^{\mathrm{T}} A^{u,v} x^t_v - \frac{n}{Z} n\alpha_{\max} \right) \\
&= Tc - \frac{T}{n} \sum_{u\in[n]} \frac{\bar{P}^T_u}{k_u} + \frac{1}{n} \sum_{u\in[n]} \frac{1}{k_u} \sum_{t=1}^{T} \left(\sum_{[u,v]\in E} (x^t_u)^{\mathrm{T}} A^{u,v} x^t_v \right) - \frac{n^2 T}{Z} \alpha_{\max} \\
&= Tc - \frac{T}{n} \sum_{u\in[n]} \frac{\bar{P}^T_u}{k_u} + \frac{1}{n} \sum_{u\in[n]} \frac{1}{k_u} T \bar{P}^T_u - \frac{n^2 T}{Z} \alpha_{\max} \\
&= Tc - \frac{n^2 T}{Z} \alpha_{\max} = Tc - \frac{\alpha_{\max}}{\Lambda}.
\end{aligned}$$

Let us now consider a mixed strategy z such that $z_{u:i} = 1$, for some $u \in [n]$ and $i \in [m_u]$. If we establish (8) for this z, it is easy to see that (8) holds for any z.

$$\begin{aligned}
z^{\mathrm{T}} \cdot R \cdot \left(\sum_{t=1}^{T} x^t \right) &= \sum_{t=1}^{T} \left(M \cdot \left(\frac{1}{k_{u+1}} - \frac{1}{k_{u-1}} \right) + \sum_{[u,v]\in E} \frac{1}{k_v} e^{\mathrm{T}}_{u:i} A^{u,v} x^t_v \right) \\
&= \sum_{t=1}^{T} \left(-\frac{1}{n} \bar{P}^T_u + c + \sum_{[u,v]\in E} \left(\frac{1}{n} \pm \frac{n}{Z} \right) e^{\mathrm{T}}_{u:i} A^{u,v} x^t_v \right)
\end{aligned}$$

$$
\begin{aligned}
&\leq \sum_{t=1}^{T}\left(-\frac{1}{n}\bar{P}_u^T + c + \frac{1}{n}\sum_{[u,v]\in E} e_{u:i}^{\mathrm{T}} A^{u,v} x_v^t + \frac{n}{Z} n\alpha_{\max}\right)\\
&= Tc - \frac{T}{n}\bar{P}_u^T + \frac{1}{n}\sum_{t=1}^{T}\left(\sum_{[u,v]\in E} e_{u:i}^{\mathrm{T}} A^{u,v} x_v^t\right) + \frac{n^2}{Z} T\alpha_{\max}\\
&= Tc - \frac{T}{n}\bar{P}_u^T + \frac{1}{n}\sum_{[u,v]\in E} e_{u:i}^{\mathrm{T}} A^{u,v}\left(\sum_{t=1}^{T} x_v^t\right) + \frac{n^2 T\alpha_{\max}}{Z} \leq Tc + \frac{g(T)}{n} + \frac{\alpha_{\max}}{\Lambda},
\end{aligned}
$$

where for the last derivation we used that the strategy sequence of the node u of $\mathcal{GG}$ has regret at most $g(T)$. Combining the above bounds we get (8). ■

We argue next the following

Lemma 5. *The pair of strategies* $(\frac{1}{T}\sum_{t=1}^{T} x^t, \frac{1}{T}\sum_{t=1}^{T} x^t)$ *is a* $\frac{2}{T}\left(\frac{g(T)}{n} + \frac{2\alpha_{\max}}{\Lambda}\right)$*-approximate Nash equilibrium of the game* $\mathcal{G}$*.*

Proof of Lemma 5: Let $\Phi := \frac{g(T)}{n} + \frac{2\alpha_{\max}}{\Lambda}$, and let us fix a pure strategy z^* for the row player. (8) implies

$$\sum_{t=1}^{T}(x^t)^{\mathrm{T}} \cdot R \cdot x^t \geq z^{*\mathrm{T}} \cdot R \cdot \left(\sum_{t=1}^{T} x^t\right) - \Phi. \tag{10}$$

Recalling that $C = -R$ and setting $z = x^t$ we get from (9) that for all t:

$$-\sum_{t=1}^{T}(x^t)^{\mathrm{T}} \cdot R \cdot x^t \geq -\left(\sum_{t=1}^{T} x^t\right)^{\mathrm{T}} \cdot R \cdot x^t - \Phi. \tag{11}$$

Combining (10) and (11), we get $\left(\sum_{t=1}^{T} x^t\right)^{\mathrm{T}} \cdot R \cdot x^t + \Phi \geq z^{*\mathrm{T}} \cdot R \cdot \left(\sum_{t=1}^{T} x^t\right) - \Phi$, for all t. Summing this for $t = 1, \ldots, T$ we get

$$\left(\sum_{t=1}^{T} x^t\right)^{\mathrm{T}} \cdot R \cdot \left(\sum_{t=1}^{T} x^t\right) \geq T \cdot z^{*\mathrm{T}} \cdot R \cdot \left(\sum_{t=1}^{T} x^t\right) - 2\Phi T.$$

Dividing by T^2 and recalling that that the above holds for all z^* completes the proof. ■

We conclude the proof of Theorem 3 by arguing that the set of strategies $\{\bar{x}_u^T\}_u$, where $\bar{x}_u^T = \frac{1}{T}\sum_{t=1}^{T} x_u^t$, comprise an approximate equilibrium of the game $\mathcal{GG}$. Denote by $\Gamma := \frac{2}{T}\left(\frac{g(T)}{n} + \frac{2\alpha_{\max}}{\Lambda}\right)$ and take $\Xi = \frac{n^2}{1-\frac{n^2}{Z}}\Gamma + \frac{n^3}{Z}2\alpha_{\max} + \frac{1}{T}$.

Lemma 6. *For all $u \in [n]$, and for all mixed strategies z_u of node u:*

$$\sum_{[u,v]\in E} (\bar{x}_u^T)^{\mathrm{T}} \cdot A^{u,v} \cdot \bar{x}_v^T \geq \sum_{[u,v]\in E} z_u^{\mathrm{T}} \cdot A^{u,v} \cdot \bar{x}_v^T - \Xi. \tag{12}$$

Proof of lemma 6: Suppose that (12) is violated for some pair u, z_u. We are going to contradict the assertion of Lemma 5. Indeed, let us define the following strategy q for the row player of $\mathcal{G}$:

$$q_{v:i} = \begin{cases} \frac{1}{k_v}\bar{x}_v^T(i), & \text{if } v \in [n] \setminus \{u\}, i \in [m_v]; \\ \frac{1}{k_u} z_u(i), & \text{if } v = u, i \in [m_u]; \end{cases}$$

and let us consider the change in the row player's payoff if she replaces strategy $\bar{x}^T := \frac{1}{T}\sum_{t=1}^{T} x^t$ by q. Clearly, the payoff from $\bar{x}^T$ is

$$\sum_{v\in[n]} \frac{1}{k_v} (\bar{x}_v^T)^{\mathrm{T}} R_v \bar{x}^T, \tag{13}$$

where we denote by R_v the matrix v restricted to the rows $(v : i)$, $i \in [m_v]$. Similarly, the payoff that the row player gets from q is

$$\frac{1}{k_u}(z_u)^{\mathrm{T}} R_u \bar{x}^T + \sum_{v\in[n]\setminus\{u\}} \frac{1}{k_v} (\bar{x}_v^T)^{\mathrm{T}} R_v \bar{x}^T. \tag{14}$$

Subtracting the two payoffs we get that the difference between the payoff from q and the payoff from $\bar{x}^T$ is

$$\begin{aligned} \frac{1}{k_u}\left(z_u^{\mathrm{T}} R_u \bar{x}^T - (\bar{x}_u^T)^{\mathrm{T}} R_u \bar{x}^T\right) &= \frac{1}{k_u}\left(\sum_{[u,v]\in E} \frac{1}{k_v}(z_u - \bar{x}_u^T)^{\mathrm{T}} \cdot A^{u,v} \cdot \bar{x}_v^T\right) \\ &= \frac{1}{k_u}\left(\sum_{[u,v]\in E}\left(\frac{1}{n} \pm \frac{n}{Z}\right)(z_u - \bar{x}_u^T)^{\mathrm{T}} \cdot A^{u,v} \cdot \bar{x}_v^T\right) \\ &\geq \frac{1}{k_u}\sum_{[u,v]\in E}\left(\frac{1}{n}(z_u - \bar{x}_u^T)^{\mathrm{T}} \cdot A^{u,v} \cdot \bar{x}_v^T - \frac{n}{Z}2\alpha_{\max}\right) \\ &\geq \frac{1}{k_u}\cdot\frac{1}{n}\sum_{[u,v]\in E}(z_u - \bar{x}_u^T)^{\mathrm{T}} \cdot A^{u,v} \cdot \bar{x}_v^T - \frac{1}{k_u}\frac{n^2}{Z}2\alpha_{\max} \\ &\geq \frac{1}{k_u}\cdot\frac{1}{n}\Xi - \frac{1}{k_u}\frac{n^2}{Z}2\alpha_{\max} > \Gamma, \end{aligned}$$

and this contradicts the assertion of Lemma 5 that $(\bar{x}^T, \bar{x}^T)$ is a Γ-approximate Nash equilibrium. ■

From Lemma 6 it follows that the strategies $\{\bar{x}_u^T\}_u$ comprise a Ξ-approximate Nash equilibrium of the game $\mathcal{GG}$. Choosing $\Lambda > \max\{8n^2, 5n^2\alpha_{\max}\}$ it follows that $\Xi < 2.3 \cdot n \cdot \frac{g(T)}{T} + \frac{3}{T}$.

4 Discussion

We believe that graphical polymatrix games are useful models of important social phenomena, such as trading or other interaction in social networks. In this paper we focused on the zero-sum variety of such games. Without restrictions on the games played on the edges, it is easy to see that the problem of computing a Nash equilibrium becomes intractable even for two strategies per player [4]. It is interesting to understand what other classes of polymatrix games have nice computational properties. For example, can the Nash equilibrium of such games, with a small number of strategies at each player, be approximated well?

Our main result raises a number of other important questions. Consider a directed graph such that along every edge (u, v) nodes u and v play the same zero-sum game $(A, -A)$. It is easy to see (assuming, for example, that the overall game is non-degenerate) that each node can be assigned a *value*, characterizing its expected payoff at equilibrium. This brings up an interesting question: which structural properties of the graph and of the position of a node in it — as well as the nature of the game A — determines these values? Such investigation could result in important insights into networked economic activity.

Finally, in an earlier version of our paper we had included an extension of our main result to the (ostensibly) more general case in which the games played at the edges are *strictly competitive,* games that share with zero-sum games this property: if both opponents change their mixed strategy, then their utilities either both stay the same, or one increases while the other decreases. In subsequent joint work with Ilan Adler [1], however, we proved that the only examples of such games are zero-sum games (or their trivial affine variants, resulting from a zero-sum game by adding a constant to all payoffs of one player, or multiplying them all by the same positive constant). In other words, this well known, and much discussed in the literature, generalization of zero-sum games is, rather astonishingly, *void!*

References

1. I. Adler, C. Daskalakis and C. H. Papadimitriou, *manuscript*, 2009.
2. R. J. Aumann. Game Theory. *The New Palgrave: A Dictionary of Economics* by J. Eatwell, M. Milgate and P. Newman, (eds.), London: Macmillan & Co, 1987.
3. L. M. Bregman and I. N. Fokin. On Separable Non-Cooperative Zero-Sum Games. *Optimization*, 44(1):69–84, 1998.
4. C. Daskalakis, P. W. Goldberg and C. H. Papadimitriou. The Complexity of Computing a Nash Equilibrium. *STOC*, 2006. To appear in the *SIAM Journal on Computing*, special issue for STOC 2006.
5. Y. Freund and R. E. Schapire. Adaptive Game Playing Using Multiplicative Weights. *Games and Economic Behavior,* 29:79–103, 1999.
6. M. J. Kearns, M. L. Littman and S. P. Singh. Graphical Models for Game Theory. *UAI*, 2001.
7. D. Kempe, J. Kleinberg and É. Tardos. Maximizing the Spread of Influence through a Social Network. *SIGKDD*, 2003.
8. J. von Neumann. Zur Theorie der Gesellschaftsspiele. *Math. Annalen*, 100:295–320, 1928.