On a Network Generalization of the Minmax Theorem

Constantinos Daskalakis* Christos H. Papadimitriou †
{costis, christos}@cs.berkeley.edu

February 10, 2009

Abstract

We consider graphical games in which edges are zero-sum games between the endpoints/players; the payoff of a player is the sum of the payoffs from each incident edge. 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. The same is true when the games on the edges belong to the more general class of strictly competitive games. Such games are arguably very broad and useful models of networked economic interactions. Our results extend, make polynomially efficient, and simplify considerably the approach in [2].

*Microsoft Research, New England.
†UC Berkeley.

1 Introduction

In 1928, von Neumann proved that every two-person zero-sum game has the minmax property [5], and thus a randomized equilibrium — which, we now know, is easily computable via linear programming. There is one moderate known generalization of zero-sum games for which von Neumann's ideas apply, namely the two-person *strictly competitive games*, for which, if the two players switch strategies, then their payoffs either stay the same, or one is increased and the other decreased. According to Aumann, two-person strictly competitive games are "one of the few areas in game theory, and indeed in the social sciences, where a fairly sharp, unique prediction is made" [1]. In this paper, we present a sweeping generalization of this class to multi-player games played on a network.

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 [3]. One interesting class of such games are the *graphical polymatrix games*, in which all players have the same strategies, 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 [4].

But what if the games at the edges are strictly competitive — 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 games and we are 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 payoffs 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, while v and w play between them, for much higher stakes, a game that rewards v for playing *heads*, but half as much if u also plays *heads*, then v and w cannot afford to pay much attention to u, and u can steal a positive payoff — despite the fact that she is playing two games of matching pennies. 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 2.1) is a reduction implying that *in any zero-sum polymatrix graphical 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 methods apply rather directly. The basic idea of the reduction is very simple: We create two players whose action set is the *union* of the actions of all players, and have 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. In fact, the same reduction can easily be seen to work for the more general class of strictly competitive games (Theorem 3.1). Another relatively direct corollary of our reduction (whose proof we omit here) is that simple distributed iterative learning algorithms on the network quickly converge to the Nash equilibrium.

Related work. In a very interesting and very short — and very dense — 1998 paper [2] (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. Finally, we do not see how to extend their method to the more general class of strictly competitive polymatrix games handled by our technique.

Definitions. An *n-player zero-sum polymatrix graphical game* is defined in terms of an undirected graph $G = (V, E)$ and, for each edge $[u, v] \in E$, an $n_u \times n_v$ real matrix $A^{u,v}$ and another $(A^{v,u}) = -(A^{u,v})^T$. That is, each player/node u have a set of actions, $[n_u]$, 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 $f(u) \in [n_u]$ for all $u \in V$, — that is, any choice of action 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.

A generalization of this class of games is that of *n-player strictly competitive polymatrix graphical games*, in which for each edge $[u, v]$ the matrices $(A^{v,u})$ and $(A^{u,v})$ satisfy the following relaxed condition: For any $i, i' \in [n_u]$ and $j, j' \in [n_v]$, $A^{u,v}_{i,j} \leq A^{u,v}_{i',j'}$ iff $A^{v,u}_{j,i} \geq A^{v,u}_{j',i'}$.

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.

2 Main Result

Theorem 2.1. *There is polynomial-time reduction from any zero-sum polymatrix game $\mathcal{GG}$ to a symmetric zero-sum bimatrix game $\mathcal{G}$, such that there is a polynomial-time computable surjective mapping from the Nash equilibria of $\mathcal{G}$ to the Nash equilibria of $\mathcal{GG}$.*

Proof of Theorem 2.1: Our construction will be based on a generalization of the rock-paper-scissors game defined below.

Definition 2.2 (Generalized Rock-Paper-Scissors)**.** For an odd integer $n > 0$, the *n-strategy rock-paper-scissors game* is a symmetric zero-sum bimatrix game $(\Gamma, -\Gamma^{\mathrm{T}})$ 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 rock paper scissors game is the uniform distribution over both players' strategies. We are going to establish next a stronger claim. For this, let $\mathcal{GG} = \{A^{u,v}\}_{u,v\in[n]}$ be an n-player zero-sum polymatrix game whose v-th player has m_v strategies, and let us define its embedding $\mathcal{G}$ into the n-strategy rock-paper-scissors game with scaling parameter $M > 0$ as follows: $\mathcal{G} = (R, C)$ is an $\sum_v m_v \times \sum_v m_v$ 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 $u, v \in [n]$, $i \in [m_u]$, $j \in [m_v]$,

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

Observe that $\mathcal{G}$ is a symmetric zero-sum bimatrix game, since the generalized rock-paper-scissors game is symmetric. We characterize next the Nash equilibria of $\mathcal{G}$.

Lemma 2.3. *Let $n > 0$ be an odd integer, $\mathcal{GG} = \{A^{u,v}\}_{u,v\in[n]}$ a zero-sum polymatrix game whose largest payoff entry in absolute value is M/L, and $\mathcal{G} = (R, C)$ the embedding of $\mathcal{G}$ 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 2.3: 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

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

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 zero-sum polymatrix 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

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

Lemma 2.4. *In the setting of Lemma 2.3, if (x, y) is a $\frac{2M \cdot n^2}{L^2}$-Nash equilibrium of $\mathcal{G}$, then the collection of mixed strategies $\{\hat{x}_u\}_u$ is a $1/L$-Nash equilibrium of $\mathcal{GG}$.*

Proof of Lemma 3: 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, fixing a node u of the polymatrix game, we are going to justify that the collection $\{\hat{x}_u\}_u$ satisfies the equilibrium conditions approximately. Indeed, because (x, x) is a Nash equilibrium of $\mathcal{G}$ it must be that, for all $i, j \in [m_u]$:

$$\mathcal{E}\left[\mathcal{P}_{u:i}\right] > \mathcal{E}\left[\mathcal{P}_{u:j}\right] \quad \Rightarrow \quad x_{u:j} = 0, \tag{1}$$

where

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

is the expected payoff to the row player for playing strategy $x_{u:i}$. From Lemma 2.3 we have that

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

Hence, (1) implies

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

which is equivalent to

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

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 (2) holds for all $u \in [n]$, $i, j \in [m_u]$, the collection $\{\hat{x}_u\}_u$ is a $\frac{2M \cdot n^2}{L^2}$-Nash equilibrium of $\mathcal{GG}$. ■

Choosing $M = 2^{q(|\mathcal{GG}|)} 2n^2 u_{\max}^2$, where $q(|\mathcal{GG}|)$ is some polynomial of the size of $\mathcal{GG}$, and $u_{\max}$ 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 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, $\{\hat{y}_u\}_u$, $\hat{y}_u$ is a distribution over $[m_u]$:

$$\begin{aligned} & \min z \\ \text{s.t.} \quad & \sum_{[u,v]\in E} \sum_{\ell \in [m_v]} A_{i,\ell}^{u,v} \cdot \hat{y}_v(\ell) > \sum_{[u,v]\in E} \sum_{\ell \in [m_v]} A_{j,\ell}^{u,v} \cdot \hat{y}_v(\ell) - z, \qquad \begin{array}{c} \forall u \in [n], \\ i \in \operatorname{supp}(\hat{x}_u), \\ j \in [m_u]. \end{array} \end{aligned} \tag{3}$$

In LP (3), $\operatorname{supp}(\hat{x}_u)$ denotes the support of the distribution $\hat{x}_u$. Observe that $(2^{-q(\mathcal{GG})}, \{\hat{x}_u\}_u)$ is a solution of LP (3) with objective value $2^{-q(\mathcal{GG})}$. But, let us choose $q(\cdot)$ to be larger than the bit description of any optimal solution of LP (3), for any possible set of supports. It follows that the optimal solution to LP (3) has objective value 0, hence the corresponding collection $\{\hat{y}_u\}_u$ will be an exact Nash equilibrium of $\mathcal{GG}$. ■

3 Extension to Strictly Competitive Games

With a little more care, we can extend our main result to the case in which each game played at the edges is strictly competitive — not necessarily zero-sum:

Theorem 3.1. *There is polynomial-time reduction from any polymatrix game $\mathcal{GG}$ with strictly competitive games to a symmetric zero-sum bimatrix game $\mathcal{G}$, such that there is a polynomial-time computable surjective mapping from the Nash equilibria of $\mathcal{G}$ to the Nash equilibria of $\mathcal{GG}$.*

Proof. (Sketch.) A *symmetric magic square zero-sum game* is a zero-sum game whose matrix (1) is symmetric; (2) all entries are distinct; and (3) all rows and all columns adding up to zero. It is an exercise in linear algebra to see that such games exist, and that the uniform strategy is the only Nash equilibrium. Let us now embed a graphical game to such a game, with payoffs multiplied by a huge number M — exactly as in Lemma 2.3, except we now use a scaled-up magic square game instead of generalized rock-paper-scissors. Then we claim that, if the original graphical game has edge games that are strictly competitive, then the constructed two-person game is also strictly competitive. To see this, consider any pair of pairs of strategies in the new game, say $(i, j); (i', j')$. If both (i, j) and (i', j') are in the same block of the magic square game, then they obey the strictly competitive property (the utility of Player I increases in going from (i, j) to (i', j') iff the utility of Player II decreases), simply because the game embedded in this block was assumed to be strictly competitive. If they are in different blocks of the magic square game, then we know that one player's utility has increased by a lot, and the other's decreased, since the utilities of the original magic square game were all different.

If x is a maxmin of the new game, seen as a zero-sum game, then we know that both players playing x is a Nash equilibrium, and in fact this Nash equilibrium is close to a Nash equilibrium of the original graphical game. By choosing a suitably large multiplier M, one can recover an approximate Nash equilibrium of the original game as in . □

4 Discussion

It can be also shown, as an important byproduct of our reduction, that in any graphical game of the sort considered here, there are natural learning algorithms run by the nodes which converge reasonably fast to the Nash equilibrium.

Our approach suggests that, in any graph of nodes playing competitive games along edges, each node can be assigned a *value*, characterizing its expected payoff at any equilibrium. This brings up the interesting question, which structural properties of the graph and of the position of a node in it — as well as the nature of the competitive games played by all — determines these values. Such investigation could result in important insights into networked economic activity.

References

[1] R. J. Aumann "Game Theory," in *The New Palgrave: A Dictionary of Economics* by J. Eatwell, M. Milgate, and P. Newman, (eds.), pp.460–482, London: Macmillan & Co, 1987.

[2] L. M. Bregman and I. N. Fokin. On Separable non-cooperative zero-sum games. *Optimization*, Volume 44(1):69–84, 1998.

[3] M. J. Kearns, M. L. Littman and S. P. Singh. Graphical Models for Game Theory. *UAI*, 2001.

[4] D. Kempe, J. Kleinberg, É. Tardos "Maximizing the Spread of Influence through a Social Network," *Proc. 9th ACM SIGKDD Intl. Conf. on Knowledge Discovery and Data Mining*, 2003.

[5] J. von Neumann "Zur Theorie der Gesellschaftsspiele," *Math. Annalen, 100,* 295-320 (1928).