

Tractable Algorithms for Approximate Nash Equilibria in Generalized Graphical Games with Tree Structure

Luis E. Ortiz

Computer and Information Science
University of Michigan - Dearborn
Dearborn, MI 48128

Mohammad T. Irfan

Digital & Comp. Studies and Computer Science
Bowdoin College
Brunswick, ME 04011

Abstract

We provide the first fully polynomial time approximation scheme (FPTAS) for computing an approximate mixed-strategy Nash equilibrium in graphical multi-hypermatrix games (GMhGs), which are generalizations of normal-form games, graphical games, graphical polymatrix games, and hypergraphical games. Computing an exact mixed-strategy Nash equilibria in graphical polymatrix games is PPAD-complete and thus generally believed to be intractable. In contrast, to the best of our knowledge, we are the first to establish an FPTAS for tree polymatrix games as well as tree graphical games when the number of actions is bounded by a constant. As a corollary, we give a quasi-polynomial time approximation scheme (quasi-PTAS) when the number of actions is bounded by a logarithm of the number of players.

Introduction

For over a decade, graphical games have been at the forefront of computational game theory. In a graphical game, a player's payoff is directly affected by her own action and those of her neighbors. This large class of games has played a critical role in establishing the hardness of computing a Nash equilibrium in general games (Daskalakis, Goldberg, and Papadimitriou 2009a). It has also generated a great deal of interest in the AI community since Kearns, Littman, and Singh (2001) drew a parallel with probabilistic graphical models in terms of succinct representation by exploiting the network structure. As a result, this is one of the select topics in computer science that has triggered a confluence of ideas from the theoretical computer science and AI communities.

This paper contributes to this development by providing the first fully polynomial-time approximation scheme (FPTAS) for approximate Nash equilibrium computation in a generalized class of tree graphical games. Tree-structured interactions are natural in hierarchical settings, such as organizational management, supply chains, and even power grids (Dvijotham et al. 2016); see Ortiz and Irfan 2016 for discussion. *Our algorithm also eliminates the exponential dependency on the maximum degree of a node, a problem that has plagued research for 15 years since the inception of graphical games (Kearns, Littman, and Singh 2001).*

Copyright © 2017, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

We consider the problem of *computing approximate MSNE in GMhGs* (see Table 1 for a list of acronyms used throughout this paper). Roughly speaking, in a GMhG, each player's payoff is the summation of several local payoff hypermatrices defined with respect to each individual player's local hypergraph. GMhGs generalize normal-form games, graphical games (Kearns, Littman, and Singh 2001; Kearns 2007), graphical polymatrix games, and hypergraphical games (Papadimitriou and Roughgarden 2008). For approximate MSNE, we adopt the standard notion of ϵ -MSNE (also known as ϵ -approximate MSNE), an additive (as opposed to relative) approximation scheme widely used in algorithmic game theory (Lipton, Markakis, and Mehta 2003; Daskalakis, Mehta, and Papadimitriou 2007; Deligkas et al. 2014; Barman, Ligett, and Piliouras 2015).

In this paper, we provide FPTAS and quasi-PTAS for GMhGs in which the individual player's number of actions m and the hypertree-width w of the underlying game hypergraph are bounded. The key to our solution is the formulation of a CSP such that any solution to this CSP is an ϵ -MSNE of the game. This raises two challenging questions: Will the CSP have any solution at all? In case it has a solution, how can we compute it *efficiently*? Regarding the first question, we discretize *both* the probability space and the payoff space of the game to guarantee that for *any* MSNE of the game (which always exists), the nearest grid point is a solution to the CSP. For the second question, we give a DP algorithm that is an FPTAS when m and w are bounded by a constant. Most remarkably, this algorithm eliminates the exponential dependency on the largest neighborhood size of a node, which has plagued previous research on this problem.

Related Work

In this section, we provide a brief overview of the previous computational complexity and algorithmic results for the problem of ϵ -MSNE computation (additive approximation scheme as most commonly defined in game theory) in general. A full account of all specific sub-classes of GMhGs such as normal-form games and (standard) graphical games is beyond the scope of this paper, just as is the discussion on (a) other types of approximations such as the less common relative approximation; (b) other popular equilibrium-solution concepts such as pure-strategy Nash equilibria and correlated equilibria (Aumann 1974; 1987); and (c) other

CSP	Constraint Satisfaction Problem
DP	Dynamic Programming
FPTAS	Fully Polynomial Time Approx. Scheme
GMhG	Graphical Multi-hypermatrix Game
MSNE	Mixed-Strategy Nash Equilibrium
Quasi-PTAS	Quasi-Polynomial Time Approx. Scheme

Table 1: Acronyms used in this paper.

quality guarantees of solutions, including exact MSNE and “well-supported” approximate MSNE.

The complexity status of normal-form games is well-understood today, thanks to a series of seminal works (Daskalakis, Goldberg, and Papadimitriou 2009a; 2009b) that culminated in the PPAD-completeness of 2-player multi-action normal-form games, also known as *bimatrix games* (Chen, Deng, and Teng 2009). Once the complexity of exact MSNE computation was established, the spotlight naturally fell on approximate MSNE, especially in succinctly representable games such as graphical games. Chen, Deng, and Teng (2009) showed that bimatrix games do not admit an FPTAS unless $\text{PPAD} \subseteq \text{P}$. This result opened up computing a PTAS.

There has been a series of results based on constant-factor approximations. The current best PTAS is a 0.3393-approximation for bimatrix games (Tsaknakis and Spirakis 2008), which can be extended to the cases of three and four-player games with the approximation guarantees of 0.6022 and 0.7153, respectively. Note that sub-exponential algorithms for computing ϵ -MSNE in games with a constant number of players have been known prior to all of these results (Lipton, Markakis, and Mehta 2003). As a result, it is unlikely that the case of constant number of players will be PPAD-complete. Along that line, Rubinfeld (2015) considered the hardness of computing ϵ -MSNE in n -player succinctly representable games such as general graphical games and graphical polymatrix games. He showed that there exists a constant ϵ such that finding an ϵ -MSNE in a 2-action graphical polymatrix game with a bipartite structure and having a maximum degree of 3 is PPAD-complete. Chen, Deng, and Teng (2009) showed the hardness of bimatrix games for a polynomially small ϵ , and Rubinfeld (2015) showed the hardness (in this case, PPAD-completeness) of n -player polymatrix games for a constant ϵ .

On a positive note, Deligkas et al. (2014) presented an algorithm for computing a $(0.5 + \delta)$ -MSNE of an n -player polymatrix game. Their algorithm runs in time polynomial in the input size and $\frac{1}{\delta}$. Very recently, Barman, Ligett, and Piliouras (2015) gave a quasi-polynomial time randomized algorithm for computing an ϵ -MSNE in tree-structured polymatrix games. They assumed that the payoffs are normalized so that the local payoff of any player i from any other player j lies in $[0, 1/d_i]$, where d_i is the degree of i . This guarantees, in a strong way, that the total payoff of any player is in $[0, 1]$. In comparison, we do not make the assumption of local payoffs lying in $[0, 1/d_i]$. Also, our algorithm is a deterministic FPTAS when m is bounded by a constant.

Closely related to our work, Ortiz (2014) recently gave

a framework for sparsely discretizing probability spaces in order to compute ϵ -MSNE in tree-structured GMhGs. The time complexity of the resulting algorithm depends on $(\frac{k}{\epsilon})^k$ when m is bounded by a constant. Ortiz’s result is a significant step forward compared to Kearns, Littman, and Singh (2001)’s algorithm in the foundational paper on graphical games. In the latter work, the time complexity depends on $(\frac{2^k}{\epsilon})^k$ when m is bounded by a constant. Both of these algorithms are exponential in the representation size of succinctly representable games such as graphical polymatrix games. Compared to these works, our algorithm eliminates the exponential dependency on k . Furthermore, compared to Ortiz’s work, we discretize *both* probability and payoff spaces in order to achieve an FPTAS. This joint discretization technique is novel for this large class of games and has a great potential for other types of games.

Hardness of Relaxing Key Restrictions. We use two restrictions: (1) Our focus is on GMhGs (e.g., graphical polymatrix games) with tree structure, and (2) our FPTAS for ϵ -MSNE computation hinges on the assumption that the number of actions is bounded by a constant. We next discuss what happens if we relax either of these two restrictions.

Tree-structured polymatrix games with unrestricted number of actions: A bimatrix game is basically a tree-structured polymatrix game with two players. Chen, Deng, and Teng (2009) showed that there exists no FPTAS for bimatrix games with an unrestricted number of actions unless all problems in PPAD are polynomial-time solvable. In this paper, we bound the number of actions by a constant. We should also note the main motivation behind graphical games, as originally introduced by Kearns, Littman, and Singh (2001): compact/succinct representations where the representation sizes do not depend exponentially in n , but are instead exponential in k and linear in n . As Kearns, Littman, and Singh (2001) stated, if $k \ll n$, we obtain exponential gains in representation size. Thus, it is n and k the parameters of main interest in standard graphical games; the parameter m is of secondary interest. Indeed, even Kearns, Littman, and Singh (2001) concentrate on the case of $m = 2$.

Graphical (not necessarily tree-structured) polymatrix games with bounded number of actions: Rubinfeld (2015) showed that for $\epsilon = 10^{-8}$ and $m = 10^4$, computing an ϵ -MSNE for an n -player game is PPAD-hard. This hardness proof involves the construction of graphical (non-tree) polymatrix games. Therefore, the result carries over to n -player graphical polymatrix games. This lower bound result shows that graph structures that are more complex than trees are intractable (under standard assumptions) even for constant m and small but constant ϵ .

Preliminaries, Background, and Notation

Denote by $a \equiv (a_1, a_2, \dots, a_n)$ an n -dimensional vector and by $a_{-i} \equiv (a_1, \dots, a_{i-1}, a_{i+1}, \dots, a_n)$ the same vector without the i -th component. Similarly, for every set $S \subset [n] \equiv \{1, \dots, n\}$, denote by $a_S \equiv (a_i : i \in S)$ the (sub-)vector formed from a using exactly the components of S . $S^c \equiv [n] - S$ denotes the complement of S , and $a \equiv (a_S, a_{S^c}) \equiv (a_i, a_{-i})$ for every i . If A_1, \dots, A_n are

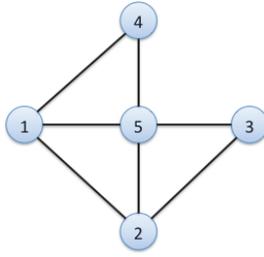


Figure 1: *The primal graph of an example GMhG.* The endpoints of each edge belong to some common hyperedge. The sets of local hyperedges of players 1 to 5 are: $\mathcal{C}_1 = \{\{1, 2\}\}$, $\mathcal{C}_2 = \{\{2, 5\}\}$, $\mathcal{C}_3 = \{\{3, 5\}\}$, $\mathcal{C}_4 = \{\{1, 4\}, \{4, 5\}\}$, and $\mathcal{C}_5 = \{\{1, 5\}, \{5, 2, 3\}\}$. The neighborhood of 5 is $N_5 = \{1, 2, 3, 5\}$ and the set of players affected by 5 is $\mathcal{N}_5 = \{2, 3, 4\}$. The local payoff of 5 is the sum of payoffs from the hyperedges $\{1, 5\}$ and $\{5, 2, 3\}$. For player 5, $\kappa_5 = 2$ and $\kappa'_5 = 3$. For the whole graph, $\kappa = 2$ and $\kappa' = 3$.

sets, denote by $A \equiv \times_{i \in [n]} A_i$, $A_{-i} \equiv \times_{j \in [n] - \{i\}} A_j$ and $A_S \equiv \times_{j \in S} A_j$. To simplify the presentation, whenever we have a difference of a set S with a singleton set $\{i\}$, we often abuse notation and denote by $S - i \equiv S - \{i\}$.

GMhG Representation

Definition 1. A graphical multi-hypermatrix game (GMhG) is defined by a set V of n players and the followings for each player $i \in V$:

- a set of actions or pure strategies A_i ;
- a set $\mathcal{C}_i \subset 2^V$ of local cliques or local hyperedges such that if $C \in \mathcal{C}_i$ then $i \in C$, and two additional sets defined based on \mathcal{C}_i :
 - i 's neighborhood $N_i \equiv \cup_{C \in \mathcal{C}_i} C$ (the set of players, including i , that affect i 's payoff) and
 - $\mathcal{N}_i \equiv \{j \in V \mid i \in N_j, j \neq i\}$ (the set of players, not including i , affected by i);
- a set $\{M'_{i,C} : A_C \rightarrow \mathbb{R} \mid C \in \mathcal{C}_i\}$ of local-clique payoff matrices; and
- the local and global payoff matrices $M'_i : A_{N_i} \rightarrow \mathbb{R}$ and $M_i : A \rightarrow \mathbb{R}$ of i defined as $M'_i(a_{N_i}) \equiv \sum_{C \in \mathcal{C}_i} M'_{i,C}(a_C)$ and $M_i(a) \equiv M'_i(a_{N_i})$, respectively.

We denote by $\kappa_i \equiv |\mathcal{C}_i|$ and $\kappa \equiv \max_i \kappa_i$ the number of hyperedges of player i and the maximum number of hyperedges over all players, respectively. Similarly, we denote $\kappa'_i \equiv \max_{C \in \mathcal{C}_i} |C|$ and $\kappa' \equiv \max_i \kappa'_i$ the size of the biggest hyperedge of player i and the size of the biggest hyperedge over all players, respectively. Also, for consistency with the graphical games literature, we denote by $k_i \equiv |N_i|$ and $k \equiv \max_i k_i$ the size of the neighborhood of the primal graph induced by the local hyperedges of i and the maximum neighborhood size over all players, respectively.

Fig. 1 illustrates some of the above terminology. The GMhG shown there (without the actual payoff matrices) is not a graphical game, because in a graphical game each \mathcal{C}_i must be singleton (i.e., only one local hyperedge for each

node i , which corresponds to N_i). This GMhG is not a polymatrix game either, because not all local hyperedges consist of only 2 nodes. Furthermore, the GMhG is not a hypergraphical game (Papadimitriou and Roughgarden 2008), because the local hyperedges are not symmetric (player 1's local hyperedge has 2 in it, but 2's local hyperedge does not have 1).

The representation sizes of GMhGs, polymatrix games, and graphical games are $O(n \kappa m^{\kappa'})$, $O(n k m^2)$, and $O(n m^k)$, respectively.

Normalizing the Payoff Scale. The dominant mode of approximation in game theory is additive approximation (Lipton, Markakis, and Mehta 2003; Daskalakis, Mehta, and Papadimitriou 2007; Deligkas et al. 2014; Barman, Ligett, and Piliouras 2015). For ϵ to be truly meaningful as a *global* additive approximation parameter, the payoffs of *all* players must be brought to the *same* scale. The convention in the literature (see, e.g., Deligkas et al. 2014) is to assume that (1) each player's local payoffs are spread between 0 and 1, with the local payoff being exactly 0 for some joint action and exactly 1 for another; and (2) the local-clique payoffs (i.e., entries in the payoff matrices) are between 0 and 1. We do not really need the second assumption to obtain our results; that is, we can handle negative values and values larger than 1. Indeed, because of the *additive* nature of the *local* payoffs in GMhGs, the “[0, 1] assumption” on those payoffs may require that some of the local-clique payoffs contain values < 0 or > 1 . This is a *key* aspect of payoff scaling, and in turn the approximation problem, that the previous literature on polymatrix games does not address. We only invoke the second assumption here to simplify the presentation. As we describe in more detail in Ortiz and Irfan (2016), we only need the *maximum spread* of local-clique payoff values to be bounded.

Note that the equilibrium conditions are invariant to affine transformations. In the case of graphical games with local payoff matrices represented in tabular/matrix/normal-form, it is convention to assume, without loss of generality, that the payoff values are such that, for each player $i \in V$, we have $\min_a M_i(a) = \min_{a_{N_i}} M'_i(a_{N_i}) = 0$ and $\max_a M_i(a) = \max_{a_{N_i}} M'_i(a_{N_i}) = 1$. Note that in the case of graphical games using such “tabular” representations, we do not lose generality by assuming the maximum and minimum local payoff values of each player are 0 and 1, respectively, because we can compute them both *efficiently*. While this will *not* be the case for graphical-game generalizations, *in the worst case*, it is also computationally efficient for GMhGs whose local hypergraphs have bounded hypertree-widths. For instance, normalizing the payoffs of a graphical polymatrix game is in P . However, such an approach is intractable in GMhGs in general.

Hence, in general, we do not have much of a choice but to assume that the payoffs of all players are in the same scale, so that using a global ϵ is meaningful.

Some additional notation is necessary regarding the payoff scale. Denote by $u_{i,C} \equiv \max_{a_C \in A_C} M'_{i,C}(a_C)$ and $l_{i,C} \equiv \min_{a_C \in A_C} M'_{i,C}(a_C)$ the max and min payoff val-

ues achieved by the local payoff hypermatrix $M'_{i,C}$, respectively; and by $R_{i,C} \equiv u_{i,C} - l_{i,C}$ its largest range of values.

Discretization Scheme

In contrast with earlier discretization schemes (Kearns, Littman, and Singh 2001), we allow different discretization sizes for different players. Also, in contrast with recent schemes (Ortiz 2014), we discretize *both* the probability space (Definition 2) and the payoff space (Definition 3).

Definition 2. (Individually-uniform mixed-strategy discretization scheme) *Let $I = [0, 1]$ be the uncountable set of the possible values of the probability $p_i(a_i)$ of each action a_i of each player i . Discretize I by a finite grid defined by the set $\tilde{I}_i = \{0, \tau_i, 2\tau_i, \dots, (s_i - 1)\tau_i, 1\}$ with interval $\tau_i = 1/s_i$ for some integer $s_i > 0$. Thus the mixed-strategy-discretization size is $|\tilde{I}_i| = s_i + 1$. We only consider mixed strategies q_i such that $q_i(a_i) \in \tilde{I}_i$ for all a_i , and $\sum_{a_i} q_i(a_i) = 1$. The induced mixed-strategy discretized space of joint mixed strategies is $\tilde{I} \equiv \times_{i \in V} \tilde{I}_i^{|A_i|}$, subject to the individual normalization constraints.*

Definition 3. (Individually-uniform expected-payoff discretization scheme) *Let $I = [0, 1]$ denote the possible expected payoff values that each player i can receive from each local-clique payoff matrix $M'_{i,C}(p_C)$, where $p_C \in \tilde{I}_C$ (i.e., p_C is in the grid). Discretize I by a finite grid defined by the set $\tilde{I}'_i = \{0, \tau'_i, 2\tau'_i, \dots, (s'_i - 1)\tau'_i, 1\}$ with interval $\tau'_i = 1/s'_i$ for some integer $s'_i > 0$. Thus the expected-payoff-discretization size is $|\tilde{I}'_i| = s'_i + 1$. Then, for any $B \subset C \in \mathcal{C}_i$, we would only consider an expected-payoff $\tilde{M}'_{i,C}(a_B, q_{C-B})$ in the discretized grid that is closest to the exact local-clique expected payoff $M'_{i,C}(a_B, q_{C-B})$. More formally, $\tilde{M}'_{i,C}(a_B, q_{C-B}) = \arg \min_{r \in \tilde{I}'_i} |r - M'_{i,C}(a_B, q_{C-B})| \equiv \text{Proj}(M'_{i,C}(a_B, q_{C-B}))$. The induced expected-payoff discretized space over all local-cliques of all players is $\tilde{I}' \equiv \times_{i \in V} (\tilde{I}'_i)^{|\mathcal{C}_i|}$.*

A GMhG-Induced CSP

Consider the following CSP induced by a GMhG:

- **Variables:** for all i and a_i , a variable p_{i,a_i} corresponding to the mixed-strategy/probability that player i plays pure strategy a_i and, for all $C \in \mathcal{C}_i$, a variable S_{i,C,a_i} corresponding to some partial sum of the expected payoff of player i based on an ordering of the local hyper-edge elements of \mathcal{C}_i . Formally, if $\mathcal{P}_i \equiv \bigcup_{a_i} \{p_{i,a_i}\}$ and $\mathcal{S}_{i,C} \equiv \bigcup_{a_i} \{S_{i,C,a_i}\}$, then the set of all variables is $\bigcup_i (\mathcal{P}_i \cup_{C \in \mathcal{C}_i} \mathcal{S}_{i,C})$.
- **Domains:** the domain of each variable $p_i(a_i)$ is \tilde{I}_i , while that of each partial-sum variable S_{i,C,a_i} is \tilde{I}'_i .
- **Constraints:** for each i :
 1. **Best-response and partial-sum expected local-clique payoff:** We first compute a hyper-tree decomposition of the local hypergraph induced by hyperedges \mathcal{C}_i . We

then order the set of local-cliques \mathcal{C}_i of each player i such that $\mathcal{C}_i \equiv \{C_i^1, C_i^2, \dots, C_i^{\kappa_i}\}$. The superscript denotes the corresponding order of the local-cliques of player i . We make sure that the order is consistent with the hypertree decomposition of the local hypergraph, in the standard (non-serial) DP-sense used in constraint and probabilistic graphical models (Dechter 2003; Koller and Friedman 2009). For each a_i :

- (a) $\sum_{a'_i} p_{i,a'_i} S_{i,C_i^{\kappa_i},a'_i} \geq S_{i,C_i^{\kappa_i},a_i} - \epsilon/2$;
- (b) $S_{i,C_i^l,a_i} = \tilde{M}_{i,C_i^l}(a_i, p_{C_i^l-i})$, and for $l = 2, \dots, \kappa_i$,

$$S_{i,C_i^l,a_i} = \tilde{M}_{i,C_i^l}(a_i, p_{C_i^l-i}) + S_{i,C_i^{l-1},a_i}.$$

2. **Normalization:** $\sum_{a_i} p_{i,a_i} = 1$.

The number of variables of the CSP is $O(n m \kappa)$. The size of each domain \tilde{I}_i is $O(s)$, where $s \equiv \max_i s_i$. The size of each domain \tilde{I}'_i is $O(s')$, where $s' \equiv \max_i s'_i$. The computation of each $\tilde{M}_{i,C_i^l}(a_i, p_{C_i^l-i})$ in 1(b) above, which takes time $O(s^{\kappa'-1})$, dominates the running time to build the constraint set. The total number of constraints is $O(n m \kappa)$. The maximum number of variables in any constraint is $O(m \kappa')$. Given a hyper-tree decomposition, the amount of time to build the constraint set using a tabular representation is $O(n m \kappa s^m \kappa' (s')^m)$, which is the representation size of the GMhG-induced CSP.

Correctness of the GMhG-Induced CSP

We use the following Lemma of Ortiz (2014). Note that our results do not follow directly from this Lemma, since we also discretize the payoff space. Furthermore, for tree-structured polymatrix games, Ortiz (2014)'s running time depends on $(\frac{k}{\epsilon})^k$ when m is bounded by a constant, whereas ours is polynomial in the maximum neighborhood size k .

Lemma 1. (Sparse MSNE Representation Theorem) *For any GMhG and any ϵ such that*

$$0 < \epsilon \leq 2 \min_{i \in V} \frac{\sum_{C \in \mathcal{C}_i} R_{i,C} (|C| - 1)}{\max_{C' \in \mathcal{C}_i} |C'| - 1},$$

a (uniform) discretization with $s_i =$

$$\left\lceil \frac{2 |A_i| \max_{j \in \mathcal{N}_i} \sum_{C \in \mathcal{C}_j} R_{j,C} (|C| - 1)}{\epsilon} \right\rceil = O\left(\frac{m \kappa \kappa'}{\epsilon}\right)$$

for each player i is sufficient to guarantee that for every MSNE of the game, its closest (in ℓ_∞ distance) joint mixed strategy in the induced discretized space is also an ϵ -MSNE.

We next present our sparse-representation theorem, which discretizes the partial sums of expected local-clique payoffs.

Theorem 1. (Sparse Joint MSNE and Expected-Payoff Representation Theorem) *Consider any GMhG and any ϵ ,*

$$0 < \epsilon \leq 2 \min_{i \in V} \frac{\sum_{C \in \mathcal{C}_i} R_{i,C} (|C| - 1)}{\max_{C' \in \mathcal{C}_i} |C'| - 1}.$$

Setting, for all players i , the pair (τ_i, τ'_i) defining the joint (individually-uniform) mixed-strategy and expected-payoff

discretization of player i such that

$$\tau_i = \frac{\epsilon}{8|A_i| \max_{j \in \mathcal{N}_i} \sum_{C \in \mathcal{C}_j} R_{j,C} (|C| - 1)}$$

and

$$\tau'_i = \frac{\epsilon}{4\kappa_i},$$

so that the discretization sizes $s_i =$

$$\left\lceil \frac{8|A_i| \max_{j \in \mathcal{N}_i} \sum_{C \in \mathcal{C}_j} R_{j,C} (|C| - 1)}{\epsilon} \right\rceil = O\left(\frac{m \kappa \kappa'}{\epsilon}\right)$$

and

$$s'_i = \left\lceil \frac{4\kappa_i}{\epsilon} \right\rceil = O\left(\frac{\kappa}{\epsilon}\right)$$

for each mixed-strategy probability and expected payoff value, respectively, is sufficient to guarantee that for every MSNE of the game, its closest (in ℓ_∞ distance) joint mixed strategy in the induced discretized space is a solution of the GMhG-induced CSP, and that any solution to the GMhG-induced CSP (in discretized probability and payoff space) is an ϵ -MSNE of the game.

Proof. Let p' be an MSNE of the GMhG. Let p be the mixed strategy closest, in ℓ_∞ , to p' in the grid induced by the combination of the discretizations that each τ_i generates. For all i and a_i , set $p_{i,a_i}^* = p_i(a_i)$; and for all i and a_i , first set $S_{i,C_i^1,a_i}^* = \widetilde{M}_{i,C_i^1}(a_i, p_{C_i^1-i}^*)$, and then recursively for $l = 2, \dots, \kappa_i$, set $S_{i,C_i^l,a_i}^* = \widetilde{M}_{i,C_i^l}(a_i, p_{C_i^l-i}^*) + S_{i,C_i^{l-1},a_i}^*$. The resulting assignment satisfies the normalization constraint of the CSP, by the definition of a mixed strategy. The assignment also satisfies the partial-sum expected local-clique payoffs by construction. By the setting of τ_i and Lemma 1, we have that p is an $(\epsilon/4)$ -MSNE, and thus also an ϵ -MSNE. In addition, for all i and a_i , we have the following sequence of inequalities:

$$\sum_{a'_i} p_i(a'_i) \sum_{C \in \mathcal{C}_i} M'_{i,C}(a'_i, p_{C-i}) \geq \sum_{C \in \mathcal{C}_i} M'_{i,C}(a_i, p_{C-i}) - \frac{\epsilon}{4}.$$

$$\sum_{a'_i} p_{i,a'_i}^* \sum_{l=1}^{\kappa_i} M'_{i,C_i^l-i}(a'_i, p_{C_i^l-i}^*) \geq \sum_{l=1}^{\kappa_i} M'_{i,C_i^l-i}(a_i, p_{C_i^l-i}^*) - \frac{\epsilon}{4}.$$

By the definition of $\widetilde{M}_{i,C}$, for all i and $C \in \mathcal{C}_i$, we have that for all a_i and $l = 1, \dots, \kappa_i$,

$$\begin{aligned} \widetilde{M}_{i,C_i^l-i}(a'_i, p_{C_i^l-i}^*) - \frac{\tau'_i}{2} &\leq M'_{i,C_i^l-i}(a'_i, p_{C_i^l-i}^*) \\ &\leq \widetilde{M}_{i,C_i^l-i}(a_i, p_{C_i^l-i}^*) + \frac{\tau'_i}{2}. \end{aligned}$$

Combining the previous two inequalities, rearranging the terms, and plugging in $\kappa_i \tau'_i = \epsilon/4$ we get:

$$\sum_{a'_i} p_{i,a'_i}^* S_{i,C_i^{\kappa_i},a'_i}^* \geq S_{i,C_i^{\kappa_i},a'_i}^* - \frac{\epsilon}{2}.$$

Hence, the assignment (p^*, S^*) also satisfies the best-response constraints (1(a) of CSP) and is a solution to the GMhG-induced CSP.

Now, for the second part of the theorem, suppose (p^*, S^*) is a solution of the GMhG-induced CSP. Then, by the combination of the best-response and partial-sum expected local-clique payoff constraints, we have that, for all i and a_i ,

$$\sum_{a'_i} p_{i,a'_i}^* S_{i,C_i^{\kappa_i},a'_i}^* \geq S_{i,C_i^{\kappa_i},a_i}^* - \frac{\epsilon}{2},$$

$$S_{i,C_i^1,a_i}^* = \widetilde{M}_{i,C_i^1}(a_i, p_{C_i^1-i}^*),$$

$$S_{i,C_i^l,a_i}^* = \widetilde{M}_{i,C_i^l}(a_i, p_{C_i^l-i}^*) + S_{i,C_i^{l-1},a_i}^*.$$

This in turn implies that for all i and a_i , we can obtain the following sequence of inequalities:

$$\sum_{a'_i} p_{i,a'_i}^* \sum_{l=1}^{\kappa_i} \widetilde{M}_{i,C_i^l}(a'_i, p_{C_i^l-i}^*) \geq \sum_{l=1}^{\kappa_i} \widetilde{M}_{i,C_i^l}(a_i, p_{C_i^l-i}^*) - \frac{\epsilon}{2}.$$

$$\sum_{a'_i} p_{i,a'_i}^* \sum_{C \in \mathcal{C}_i} \widetilde{M}_{i,C}(a'_i, p_{C-i}^*) \geq \sum_{C \in \mathcal{C}_i} \widetilde{M}_{i,C}(a_i, p_{C-i}^*) - \frac{\epsilon}{2}.$$

Using the CSP constraints and after some algebra, we get:

$$\begin{aligned} \sum_{a'_i} p_{i,a'_i}^* \sum_{C \in \mathcal{C}_i} M'_{i,C}(a'_i, p_{C-i}^*) &\geq \\ \sum_{C \in \mathcal{C}_i} M'_{i,C}(a_i, p_{C-i}^*) &- \epsilon \end{aligned}$$

Hence, the corresponding joint mixed-strategy p^* is an ϵ -MSNE of the GMhG. \square

CSP-Based Computational Results

The CSP formulation in the previous section leads us to the following computational results based on well-known algorithms for solving CSPs (Russell and Norvig 2003, Ch. 5), and the application of equally well-known computational results for them (Dechter 2003; Gottlob, Greco, and Scarcello 2014; Gottlob et al. 2016).

Theorem 2. *There exists an algorithm that, given as input a number $\epsilon > 0$ and an n -player GMhG with maximum local-hyperedge-set size κ and maximum number of actions m , and whose corresponding CSP has a hypergraph with hypertree-width w , computes an ϵ -MSNE of the GMhG in time $[n(m \kappa \kappa' / \epsilon)^{m \kappa'}]^{O(w)}$.*

For GMhGs with bounded hypertree width w , the following corollary establishes our main CSP-based result.

Corollary 1. *There exists an algorithm that, given as input a GMhG with bounded w , outputs an ϵ -MSNE in polynomial time in the size of the input and $1/\epsilon$, for any $\epsilon > 0$; hence, the algorithm is an FPTAS. If, instead, we have $w = O(\text{polylog}(n))$, then the algorithm is a quasi-PTAS.*

Theorem 2 also implies that we can compute an ϵ -MSNE of a tree-structured polymatrix game in $O(n(mk/\epsilon)^{2m})$. Note that the running time is polynomial in the maximum neighborhood size k .

The following results are in term of the primal-graph representation of the GMhG-induced CSP.

Theorem 3. *There exists an algorithm that, given as input a number $\epsilon > 0$ and an n -player GMhG with maximum number of actions m , primal-graph treewidth w' of the corresponding CSP, maximum local-hyperedge-set size κ , and maximum local-hyperedge size κ' , computes an ϵ -MSNE of the game in time $2^{O(w')} n \log(n) + n[(m \kappa \kappa' / \epsilon)^m]^{O(w')}$.*

Several corollary FPTAS results follow from the above theorem (Ortiz and Irfan 2016).

DP for ϵ -MSNE Computation

We present a DP algorithm in the context of the special, but still important class of tree-structured *polymatrix* games. This is for simplicity and clarity, and as we later discuss, is without loss of generality. We first designate an arbitrary node as the *root* of the tree and define the notion of *parents* and *children* nodes as follows. For any node/player i , we denote by $\text{pa}(i)$ the *single parent of any non-root node* in the tree and by $\text{Ch}(i)$ the *children* of node i in the root-designated-induced directed tree. If i is the root, then $\text{pa}(i)$ is undefined. If i is a leaf, then $\text{Ch}(i) = \emptyset$.

The two-pass algorithm is similar in spirit to *TreeNash* (Kearns, Littman, and Singh 2001), except that (1) here the messages are $\{-\infty, 0\}$, instead of bits $\{0, 1\}$; and (2) more distinctly, our algorithm implicitly passes messages about the partial-sum of expected payoffs across the siblings.

Collection Pass. For each non-root node i , we denote by $j = \text{pa}(i)$. We order $\text{Ch}(i)$ as $o_1, \dots, o_{|\text{Ch}(i)|}$. We then apply the following DP bottom-up (i.e., from leaves to root). We give an intuition before giving the formal specification. The message $T_{i \rightarrow j}(p_i, p_j)$ is 0 iff it is “OK” for i to play p_i when i ’s parent j plays p_j (the notion of OK recursively makes sure that i ’s children are also OK). The message $B_i(p_i, p_j, S_{o_{|\text{Ch}(i)|}})$ is 0 iff i ’s best response to j playing p_j is p_i , given that i gets a combined payoff of $S_{o_{|\text{Ch}(i)|}}$ from its children. The message $R_{o_l}(p_i, S_{o_l})$ can be thought of as being implicitly passed from i ’s child o_l to the next (and back to i from the last child $o_{|\text{Ch}(i)|}$). $R_{o_l}(p_i, S_{o_l})$ is 0 iff S_{o_l} is the maximum payoff that i can get from its first l children when i plays p_i and those children are OK with that. Fig. 2 illustrates the message passing.

Formally, for each arc (j, i) in the designated-root-induced directed tree (i.e., j is the parent of i), and (p_i, p_j) a mixed-strategy pair in the induced grid:

$$\begin{aligned} T_{i \rightarrow j}(p_i, p_j) &= \max_{S_{o_{|\text{Ch}(i)|}}} B_i(p_i, p_j, S_{o_{|\text{Ch}(i)|}}) + \\ &\quad R_{o_{|\text{Ch}(i)|}}(p_i, S_{o_{|\text{Ch}(i)|}}) \\ W_{i \rightarrow j}(p_i, p_j) &= \arg \max_{S_{o_{|\text{Ch}(i)|}}} B_i(p_i, p_j, S_{o_{|\text{Ch}(i)|}}) + \\ &\quad R_{o_{|\text{Ch}(i)|}}(p_i, S_{o_{|\text{Ch}(i)|}}) \end{aligned}$$

where $B_i(p_i, p_j, S_{o_{|\text{Ch}(i)|}}) =$

$$\sum_{a_i} \log \left(\mathbb{1} \left[\sum_{a'_i} p_i(a'_i) (\widetilde{M}_{i,j}(a'_i, p_j) + S_{o_{|\text{Ch}(i)|}}(a'_i)) \geq \widetilde{M}_{i,j}(a_i, p_j) + S_{o_{|\text{Ch}(i)|}}(a_i) - \epsilon \right] \right)$$

and, for $l = 1, \dots, |\text{Ch}(i)|$,

$$\begin{aligned} V_{o_l}(S_{o_l}, p_{o_l}, S_{o_{l-1}}) &= \sum_{a_i} \log \left(\mathbb{1} [S_{o_l}(a_i) = \right. \\ &\quad \left. \widetilde{M}_{i,o_l}(a_i, p_{o_l}) + S_{o_{l-1}}(a_i) \right] \end{aligned}$$

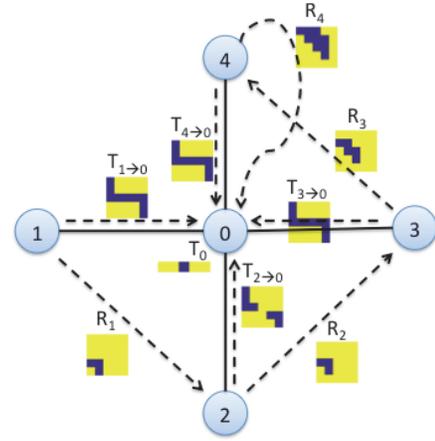


Figure 2: DP on a 5-node star polymatrix game. Solid lines represent edges, broken lines show the final round of message passing. The endpoints of every edge is playing a matching pennies game between them. The visualization of $T_{1 \rightarrow 0}$, for example, plots $p_1(a_1 = 0)$ on x -axis and $p_0(a_0 = 0)$ on y -axis. Dark grid points denote OK (i.e., $T_{1 \rightarrow 0} = 0$) and light grid points not OK. The R_i tables are 3-dimensional. Here we only show one slice of R_i values corresponding to $p_0 = (0.5, 0.5)$. The x -axis represents $S_i(a_0 = 0)$ (partial sum up to the i -th child when player 0 plays 0) and y -axis $S_i(a_0 = 1)$. A 0.1-MSNE computed for this instance is: $p_0 = p_3 = (0.5, 0.5)$, $p_1 = (0.75, 0.25)$, $p_2 = p_4 = (0, 1)$.

$$\begin{aligned} F_{o_l}(p_i, S_{o_l}, p_{o_l}, S_{o_{l-1}}) &= T_{o_l \rightarrow i}(p_{o_l}, p_i) + R_{o_{l-1}}(p_i, S_{o_{l-1}}) \\ &\quad + V_{o_l}(S_{o_l}, p_{o_l}, S_{o_{l-1}}) \end{aligned}$$

$$R_{o_l}(p_i, S_{o_l}) = \max_{p_{o_l}, S_{o_{l-1}}} F_{o_l}(p_i, S_{o_l}, p_{o_l}, S_{o_{l-1}})$$

$$W_{o_l}(p_i, S_{o_l}) = \arg \max_{p_{o_l}, S_{o_{l-1}}} F_{o_l}(p_i, S_{o_l}, p_{o_l}, S_{o_{l-1}}).$$

Following are the boundary conditions: $R_{o_0} \equiv 0$ and $S_{o_0} \equiv 0$, so that $F_{o_1}(p_i, S_{o_1}, p_{o_1}, S_{o_0}) \equiv F_{o_1}(p_i, S_{o_1}, p_{o_1}) = T_{o_1 \rightarrow i}(p_{o_1}, p_i)$. If i is the root, then $T_{i \rightarrow j}(p_i, p_j) \equiv T_i(p_i)$ and $W_{i \rightarrow j}(p_i, p_j) \equiv W_i(p_i)$. If i is a leaf, $T_{i \rightarrow j}(p_i, p_j)$ takes a simpler, non-recursive form.

Assignment Pass. For root i , set $p_i^* \in \arg \max_{p_i} T_i(p_i)$ and $S_{o_{|\text{Ch}(i)|}}^* \in W_i(p_i^*)$. Then recursively apply the following assignment process starting at $o_{|\text{Ch}(i)|}$: for $l = |\text{Ch}(i)|, \dots, 1$, set $(p_{o_l}^*, S_{o_{l-1}}^*) \in W_{o_l}(p_i^*, S_{o_l}^*)$.

The Running Time of the DP Algorithm

A running-time analysis of the DP algorithm presented above yields the following theorem, which is one of our main algorithmic results of this paper.

Theorem 4. *The DP algorithm computes an ϵ -MSNE in a graphical polymatrix game with a tree graph in time $n \left(\frac{mk}{\epsilon}\right)^{O(m)}$.*

Corollary 2. *The DP algorithm is an FPTAS to compute an ϵ -MSNE in an n -player graphical polymatrix game with*

a tree graph and a bounded number of actions m . If $m = O(\text{polylog}(n))$, then the DP algorithm is a quasi-PTAS.

Further Refinement

We describe a more refined alternative to the GMhG-induced CSP that reduces the dependency on κ . The main idea is to evaluate the expressions involving the expected local-clique payoffs matrices $\widetilde{M}_{i,C}(a_i, p_{C-i})$ in a smart way by decomposing the sum involving the expectation, considering one player mixed-strategy at a time, and projecting to the discretized payoff space after evaluating each term in the sum. This approach gives us an FPTAS for tree graphical games (in normal form) and bounded number of actions, for which the best known approximation result to-date is a quasi-PTAS. The resulting alternative CSP is considerably more complex and hence we refer the reader to Ortiz and Irfan (2016) for a detailed, formal presentation. With the above proof sketch, we present the following result.

Theorem 5. *There exists a DP algorithm that computes an ϵ -MSNE in a tree graphical game in time $n \left(\frac{mk}{\epsilon}\right)^{3m+2} O\left((mk)^2\right)$. If m is bounded, then the running time is $\text{poly}\left(n m^k, \frac{1}{\epsilon}\right)$ and the algorithm is an FPTAS. If $m = O(\text{polylog}(n))$, then this algorithm is a quasi-PTAS.*

Concluding Remarks

We have presented tractable algorithms for computing ϵ -MSNE in tree-structured GMhGs when the number of actions is bounded. The implications of our results can best be highlighted by considering a very simple 101-node star polymatrix game with a constant number of actions. For computing an ϵ -MSNE of this game, the algorithm of Kearns, Littman, and Singh (2001) takes $O\left(\left(\frac{2^{k+2}k \log k}{\epsilon}\right)^2 k\right)$ time (here $k = 100$), the algorithm of Ortiz (2014) takes $O\left(\left(\frac{k}{\epsilon}\right)^k\right)$ time, and ours takes $O(\text{poly}\left(\frac{k}{\epsilon}\right))$ time and thereby solves a 15-year-old open problem. We conclude by emphasizing that our DP algorithm is simple to implement and that simplicity is a strength of this work.

References

- Aumann, R. 1974. Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics* 1(1):67–96.
- Aumann, R. J. 1987. Correlated Equilibrium as an Expression of Bayesian Rationality. *Econometrica* 55(1):1–18.
- Barman, S.; Ligett, K.; and Piliouras, G. 2015. Approximating Nash equilibria in tree polymatrix games. In Hofer, M., ed., *Algorithmic Game Theory: 8th International Symposium, SAGT 2015, Saarbrücken, Germany, September 28-30, 2015. Proceedings*, 285–296. Berlin, Heidelberg: Springer Berlin Heidelberg.
- Chen, X.; Deng, X.; and Teng, S.-H. 2009. Settling the complexity of computing two-player Nash equilibria. *J. ACM* 56(3):14:1–14:57.
- Daskalakis, C.; Goldberg, P. W.; and Papadimitriou, C. H. 2009a. The complexity of computing a Nash equilibrium. *Commun. ACM* 52(2):89–97.
- Daskalakis, C.; Goldberg, P. W.; and Papadimitriou, C. H. 2009b. The complexity of computing a Nash equilibrium. *SIAM Journal on Computing* 39(1):195–259.
- Daskalakis, C.; Mehta, A.; and Papadimitriou, C. 2007. Progress in approximate Nash equilibria. In *EC '07: Proceedings of the 8th ACM Conference on Electronic Commerce*, 355–358. New York, NY, USA: ACM.
- Dechter, R. 2003. *Constraint Processing*. Morgan Kaufmann.
- Deligkas, A.; Fearnley, J.; Savani, R.; and Spirakis, P. 2014. Computing approximate Nash equilibria in polymatrix games. In Liu, T.-Y.; Qi, Q.; and Ye, Y., eds., *Web and Internet Economics: 10th International Conference, WINE 2014, Beijing, China, December 14-17, 2014. Proceedings*, 58–71. Cham: Springer International Publishing.
- Dvijotham, K.; Chertkov, M.; Van Hentenryck, P.; Vuffray, M.; and Misra, S. 2016. Graphical models for optimal power flow. *Constraints* 1–26.
- Gottlob, G.; Greco, G.; Leone, N.; and Scarcello, F. 2016. Hypertree decompositions: Questions and answers. In *Proceedings of the 35th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems, PODS '16*, 57–74. New York, NY, USA: ACM.
- Gottlob, G.; Greco, G.; and Scarcello, F. 2014. Treewidth and hypertree width. In Bordeaux, L.; Hamadi, Y.; and Kohli, P., eds., *Tractability: Practical Approaches to Hard Problems*. Cambridge: Cambridge University Press. 3–38.
- Kearns, M. J.; Littman, M. L.; and Singh, S. P. 2001. Graphical models for game theory. In *UAI '01: Proceedings of the 17th Conference in Uncertainty in Artificial Intelligence*, 253–260. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc.
- Kearns, M. 2007. Graphical games. In Nisan, N.; Roughgarden, T.; Tardos, E.; and Vazirani, V. V., eds., *Algorithmic Game Theory*. Cambridge University Press. chapter 7, 159–180.
- Koller, D., and Friedman, N. 2009. *Probabilistic Graphical Models: Principles and Techniques*. MIT Press.
- Lipton, R. J.; Markakis, E.; and Mehta, A. 2003. Playing large games using simple strategies. In *EC '03: Proceedings of the Fourth ACM Conference on Electronic Commerce*, 36–41. New York, NY, USA: ACM.
- Ortiz, L. E., and Irfan, M. T. 2016. FPTAS for mixed-strategy Nash equilibria in tree graphical games and their generalizations. *CoRR* abs/1602.05237.
- Ortiz, L. E. 2014. On sparse discretization for graphical games. *CoRR* abs/1411.3320.
- Papadimitriou, C. H., and Roughgarden, T. 2008. Computing correlated equilibria in multi-player games. *J. ACM* 55(3):14:1–14:29.
- Rubinfeld, A. 2015. Inapproximability of Nash equilibrium. In *Proceedings of the Forty-seventh Annual ACM Symposium on Theory of Computing, STOC '15*, 409–418. New York, NY, USA: ACM.
- Russell, S. J., and Norvig, P. 2003. *Artificial Intelligence: A Modern Approach*. Pearson Education.
- Tsaknakis, H., and Spirakis, P. G. 2008. An optimization approach for approximate Nash equilibria. *Internet Math.* 5(4):365–382.