Correlated Equilibria and Probabilistic Inference in Graphical Models *

Luis E. Ortiz

Department of Computer Science
Stony Brook University
Stony Brook, NY 11794-4400, USA
leortiz@cs.sunysb.edu
http://www.cs.sunsyb.edu/~leortiz

Abstract. This paper explores the connection between belief inference and game-theoretic equilibria within the context of graphical models. While a lot of the work in graphical models for game theory has mirrored that in probabilistic graphical models, the paper also considers the opposite direction: Taking advantage of recent advances in equilibrium computation for probabilistic inference. In particular, the paper presents formulations of inference problems in Markov random fields as computation of equilibria in a certain class of game-theoretic graphical models. To do so, the paper introduces *graphical multi-hypermatrix games (GMGs)*, a new class of graphical models for game theory that generalize graphical and polymatrix games. It also introduces graphical versions of potential games, which play a key role in the formulations, and presents results derived from the study of their properties. Finally, the paper includes a study of correlated equilibria in GMGs and provides both a characterization of the equilibria in such games and the design of a simple linear feasibility system for computing the equilibria based on the extension of a standard approach within the probabilistic and constraint graphical models community.

1 Introduction

Most of the work in graphical games has borrowed heavily from analogies to probabilistic graphical models. The work of Kakade et al. [2003] established what may be the strongest connections today. However, over-reliance on those analogies and previous standard approaches to exact inference might have led the approach to face the same computational roadblocks that plagued most exact inference techniques.

As is common knowledge within the graphical models community, exact inference is tractable in probabilistic graphical models whose graphs have bounded treewidth. Except for a few special cases such as Ising models with planar graphs (see Istrail [2000] for a recent, reasonably comprehensive account of the computational literature in this area) the problem is for the most part considered intractable outside those confines [Cooper, 1990, Shimony, 1994, Istrail, 2000].Papadimitriou and Roughgarden [2005] appeared to start to cement the same state of affairs for graphical games when they showed the intractability of computing the "social-welfare" optimum correlated equilibria in arbitrary graphical games (see also Papadimitriou and Roughgarden [2008]). In short, everything seemed to point toward an eventual resignation that the approach of Kakade et al. [2003], along with any other approach to the problem for that matter, had hit the "bounded treewidth threshold wall."

* First draft: August 25, 2009

Yet, within a year, Papadimitriou [2005] took what seems like a radically different approach to the problem, and surprised us with an efficient algorithm for computing correlated equilibria in not only graphical games, but also almost all known compactly representable games.

An immediate question that arises from this result is, what is so fundamentally different between the problem of exact inference in graphical models, or similar problems such as constraint satisfaction in constraint networks, and equilibrium computation that made this result possible in the context of graphical games? Of course, correlated equilibria, probabilistic inference and their variants are different problems, even within the same framework of graphical models. The question is, how different are they really?

It is well-known that *pure strategy Nash equilibrium (PNE)* is inherently a *constraints satisfaction problem (CSP)*. It is also well-known that any CSP can be cast as a most-likely, or equivalently, a *maximum a posteriori (MAP)* assignment estimation problem in *Markov random fields (MRFs)*. [1] Through this connection, it is clear that there exists a MAP formulation of PNE. But what about other, more general forms of equilibria?

As presented here, equilibrium computation can be seen as a kind of local conditions for different approximations to belief inference. Similarly, some special games can be seen as defining an *equivalent* MRF whose "locally optimal" solutions correspond to *arbitrary* equilibria of the game. Hence, Papadimitriou's result also opened up the possibility that at least new classes of problems in probabilistic graphical models could be solved exactly and efficiently. The question is, which classes?

On a related note, Papadimitriou's algorithm is based on the ellipsoid method, an approach that, while provably efficient in theory, is often seen as less practical as other alternatives such as interior-point methods. This is in contrast to the simple linear programs that are possible for graphical games. Are there simpler and practically effective variants of Papadimitriou's algorithm?

This work starts to address the questions above, and reports partial progress.

The preliminary section introduces basic notation, terminology and concepts from graphical models and game theory. It also introduces a further generalization of graphical games called *graphical multi-hypernatrix games (GMGs)*. Then, a class of games called *graphical potential games*, a graphical version of classical potential games[Monderer and Shapley, 1996] in game theory, is introduced. Graphical potential games play a key role in establishing a stronger connection between probabilistic inference in MRFs and equilibria in games than previously noted. Several results characterizing such games are derived. The section ends with a discussion of other preliminary results that follow from the connection between games and inference, including the implication of convergence of certain kinds of "playing" processes in games based on connections to the Gibbs sampler [Geman and Geman, 1984] via the Hammersley-Clifford Theorem [Hammersley and Clifford, 1971, Besag, 1974].

The presentation will then shift to showing a reduction of different problems in belief inference in MRFs as computing equilibria in a special class of graphical potential games called *Gibbs potential games*. The reductions presented here vary in generality from MAP assign-

[1] Assuming a solution exists, of course; otherwise the resulting MRF is not well-defined.

ment, marginals and full-joint estimation to pure *Nash equilibria (NE)*, mixed strategy NE and correlated equilibria, respectively. The section ends by discussing a connection between Papadimitriou's algorithm and the work of Jaakkola and Jordan [1997] on approximations via mean-field mixtures to the problem of probabilistic inference in MRFs. Connections to previous work in computer vision on the problem of relaxation labeling and recent work on game-theoretic approaches to (Bayesian) statistical estimation are presented.

The last technical section proposes a simple linear feasibility system for computing correlated equilibria in GMGs. The approach applies to *any* GMG, and leads to guaranteed polynomial-time algorithms in certain subclasses. The paper ends with some concluding remarks on the potential for cross-fertilization between the game theory and graphical models community.

2 Preliminaries

This section introduces basic notation and concepts in graphical models and game theory used throughout the paper. It also includes brief statements on current state-of-the-art mathematical and computational results in the area.

Basic Notation. Denote by $x \equiv (x_1, x_2, \ldots, x_n)$ an n-dimensional vector and by $x_{-i} \equiv (x_1, \ldots, x_{i-1}, x_{i+1}, \ldots, x_n)$ the same vector without component i. Similarly, for every set $S \subset [n] \equiv \{1, \ldots, n\}$, denote by $x_S \equiv (x_i : i \in S)$ the (sub-)vector formed from x using only components in S, such that, letting $S^c \equiv [n] - S$ denote the complement of S, we can denote $x \equiv (x_S, x_{S^c}) \equiv (x_i, x_{-i})$ for every i. If $A_1, \ldots, A_n$ are sets, denote by $A \equiv \times_{i \in [n]} A_i$, $A_{-i} \equiv \times_{j \in [n]-\{i\}} A_j$ and $A_S \equiv \times_{j \in S} A_j$.

Graph Terminology and Notation. Let $G = (V, E)$ be an undirected graph, with finite set of n *vertices* or *nodes* $V = \{1, \ldots, n\}$ and a set of (undirected) edges E. For each node i, let $\mathcal{N}(i) \equiv \{j \mid (i, j) \in E\}$ be the set of neighbors of i in G, *not including* i, and $N(i) \equiv \mathcal{N}(i) \cup \{i\}$ the set *including* i. A *clique* C of G is a set of nodes with the property that they are all mutually connected: for all $i, j \in C$, $(i, j) \in E$; in addition, C is *maximal* if there is no other node k outside C that is also connected to each node in C, i.e., for all $k \in V - C$, $(k, i) \notin E$ for some $i \in C$.

Another useful concept in the context of this paper is that of hypergraphs, which are generalizations of regular graphs. A *hypergraph graph* $\mathcal{G} = (V, \mathcal{E})$ is defined by a set of nodes V and a set of *hyperedges* $\mathcal{E} \subset 2^V$. We can think of the hyperedges as cliques in a regular graph. Indeed, the *primal graph* of the hypergraph is the graph induced by the node set V and where there is a edge between two nodes if they both belong to the some hyperedge; in other words, the primal graph is the graph induced by taking each hyperedge and forming cliques of nodes in a regular graph.

Given a directed graph $G^d = (V, E^d)$, denote by $\mathrm{Pa}(i) \equiv \{j \in V \mid (j, i) \in E^d\}$ the set of *parents* of node i in the graph (i.e., the set of nodes from which there exists an arc pointing to node i in the graph). A *directed acyclic graph (DAG)* has no directed cycles (i.e., no directed path, or sequence of edges, $(i_0, i_1), (i_1, i_2), \ldots, (i_{l-1}, i_l) \in E^d$ of any length $l > 0$, such that

$i_0 = i_l$). For each node i in a DAG, denote by $\mathrm{Nd}(i)$ the set of *non-descendants* of a node i, i.e., the set of nodes not reachable by following a *directed path* starting from i.

2.1 Graphical Models

Graphical models are an elegant marriage of statistics and graph theory that has had tremendous impact in the theory and practice of modern statistics. It has permitted effective modeling of large, structured high-dimensional complex systems found in the real world. The language of graphical models allows us to capture the probabilistic structure of complex interactions between individual entities in the system. The core component of the model is a graph in which each node i corresponds to a random variable X_i and the edges express conditional independence assumptions about those random variables in the probabilistic system.

Markov Random Fields, Gibbs Distributions and the Hammersley-Clifford Theorem A joint probability distribution P is called a *Markov random field (MRF)* with respect to an undirected graph G if for all x, for every node i, $P(X_i = x_i \mid X_{-i} = x_{-i}) = P(X_i = x_i \mid X_{\mathcal{N}(i)} = x_{\mathcal{N}(i)})$. In that case, the neighbors/variables $X_{\mathcal{N}(i)}$ form the *Markov blanket* of node/variable X_i.

A joint distribution P is called a *Gibbs distribution* with respect to a an undirected graph G if it can be expressed as $P(X = x) = \prod_{C \in \mathcal{C}} \Phi_C(x_C)$ for some functions Φ_C indexed by a clique $C \in \mathcal{C}$, the set of all (maximal) cliques in G, and mapping every possible value x_C that the random variables X_c associated with the nodes in C can take to a non-negative number.

Let us say that a joint probability distribution P is *positive* if it has full support (i.e., $P(x) > 0$ for all x).

Theorem 1. [Hammersley-Clifford] *Let P be a positive joint probability distribution. Then, P is an MRF with respect to G if and only if P is a Gibbs distribution with respect to G.*

In the context of the theorem, the functions Φ_C are positive, which allows us to define MRFs in terms of *local potential functions* $\{\phi_C\}$ over each clique C in the graph. Define the function $\Psi(x) \equiv \sum_{C \in \mathcal{C}} \phi_C(x_C)$. Let us refer to any function of this form as a *Gibbs potential* with respect to G. A more familiar expression of an MRF is $P(X = x) \propto \exp(\sum_{C \in \mathcal{C}} \phi_C(x_C)) = \exp(\Psi(x))$.

Probabilistic Hypergraphical Models Let us introduce a class of graphical models better suited to the context of this paper. In particular, it provides finer grain representations and facilitates more accurate characterizations of structural and computational properties of the models, problems and solutions considered here. [2]

[2] This model can be represented as a probabilistic instantiation of a *factor graph* [Kschischang et al., 2001]. Factor graphs provide a extremely general framework in which one can represent at a finer level of detail almost any

Definition 1. *Let us say that a joint probability distribution P is Gibbs with respect to* hypergraph $\mathcal{G} = (V, \mathcal{E})$ *if it can be expressed as* $P(X = x) = \prod_{C \in \mathcal{E}} \Phi_C(x_C)$ *for some functions* Φ_C *indexed by a hyperedge* $C \in \mathcal{E}$*, as before.*

It is not hard to see that a Gibbs distribution with respect to hypergraph $\mathcal{G}$ is still an MRF with respect to the primal graph of $\mathcal{G}$, assuming the conditionals are well-defined. Similarly, a positive MRF with graph G is a Gibbs distribution with respect to a hypergraph with hyperedges being the maximal cliques of G.

Directed Graphical Models: Bayesian Networks The concept of directed graphical models is relevant to the last section of the paper (Section 5), which presents simple linear programs for computing correlated equilibria in a game-theoretic graphical model. In particular, the resulting equilibria is representable as a Bayesian network, which provides efficient implementation of the equilibria.

A joint probability distribution is a *Bayesian network (BN)* with respect to a *directed acyclic graph (DAG)* $G^d = (V, E^d)$ if for all x, and every node $i \in V$, we have [3]

$$P(X = x) = \prod_{i=1}^{n} P(X_i = x_i \mid X_{\mathrm{Pa}(i)} = x_{\mathrm{Pa}(i)}) \ .$$

When all the random variables are finite, each factor $P(X_i = x_i \mid X_{\mathrm{Pa}(i)} = x_{\mathrm{Pa}(i)})$ is called the *conditional probability table (CPT)* of node i.

Thus, a BN is a Gibbs distribution with respect to a hypergraph $\mathcal{G}$ with one hyperedge $\mathrm{Pa}(i) \cup \{i\}$ for each node $i \in V$. The primal graph of $\mathcal{G}$ is called the *moralized graph* of directed graph G^d (i.e., an undirected graph that results from "marrying" or connecting every set of parents of every node and removing the directionality of the arcs in G^d). The implicit conditional independence assumption in a BN is that, for every node $i \in V$, the random variable corresponding to i is conditionally independent of the variables corresponding to its non-descendants $\mathrm{Nd}(i)$ in the graph, given the values of its parents $\mathrm{Pa}(i)$ (i.e., for all $x_i, x_{\mathrm{Pa}(i)}, x_{\mathrm{Nd}(i)}$, we have $P(X_i = x_i \mid X_{\mathrm{Pa}(i)} = x_{\mathrm{Pa}(i)}, X_{\mathrm{Nd}(i)} = x_{\mathrm{Nd}(i)}) = P(X_i = x_i \mid X_{\mathrm{Pa}(i)} = x_{\mathrm{Pa}(i)})$).

One important property of the BN representation is that, in contrast to arbitrary MRFs, no normalization is needed. This permits efficient *exact sampling* or *stochastic simulation* from any BN-represented distribution via a simple, well-known process. [4] In particular, if

graphical model in a wide variety of contexts, not just probabilistic, but also including those for constraint satisfaction and decision theoretic problems, as well as the game-theoretic models used throughout this paper. The factor-graph framework also offers generalizations of algorithms and computational techniques for many different problems within their respective context. The term "probabilistic hypergraphical model" seems more appropriate in the context of this paper and is used here because it emphasizes both the specific probabilistic nature and the underlying hypergraph structure of the model.

[3] If node i is a source node (i.e., $\mathrm{Pa}(i) = \emptyset$), then we use its marginal probability $P(X_i = x_i)$ in the expression.

[4] First note that because the BN graph is directed, it defines a partial order of the random variables, starting from those variables that are sources in the graph (i.e., nodes with no parents). We can draw samples easily from them by using the standard inverse-CDF technique. For each source node i, we can order the possible values of the random variable $\Omega_i = \{x_i^1, \ldots, x_i^m\}$ and (recursively) compute a corresponding *cumulative distribution*

m_i denotes the number of possible values that variable X_i can take and $m \equiv \max_i m_i$ is the largest such number, we can obtain a random sample according to the distribution represented by the BN in time $\sum_{i=1}^{n} m_i = O(nm)$ using $O(n+m)$ space.

This paper concentrates on connections to inference problems in undirected graphical models, which are formally presented next. In some sense, this is without loss of generality, given that directed graphical models such as Bayesian networks are still related to Gibbs distributions and thus can be seen as a special type of undirected graphical models (i.e., a BN is a MRF with respect to the moralized graph of the BN, assuming that the conditional probabilities are well-defined, of course).

Some Inference-Related Problems in MRFs One problem of interest in an MRF is to compute a *most likely assignment* $x^* \in \arg\max_x P(X = x) = \arg\max_x \sum_{C \in \mathcal{C}} \phi_C(x_C))$; that is, the most likely outcome with respect to the MRF P. Another problem is to compute the *individual marginal probabilities* $P(X_i = x_i) = \sum_{x_{-i}} P(X_i = x_i, X_{-i} = x_{-i}) \propto \sum_{x_{-i}} \exp(\sum_{C \in \mathcal{C}} \phi_C(x_C)))$ for each variable X_i. A related problem is to compute the normalizing constant $Z = \sum_x \exp(\sum_{C \in \mathcal{C}} \phi_C(x_C)))$ (also known as the *partition function* of the MRF).

Another set of problems concern so called "belief updating." That is, computing information related to the *posterior probability distribution* P' having observed the outcome of some of the variables, also known as the *evidence*. Let $S \subset [n]$ be a set of indexes to the evidence variables and $e_S = (e_j : j \in S)$ be the vector of observed values such that for all $j \in S$, e_j is the observed value of evidence variable X_j. Let $G' = (V', E')$ be the same graph $G = (V, E)$ of the MRF P, *except* that all the vertices/nodes corresponding to variables in S, as well as any edge connected to them, removed (i.e., $V' = V - S$ and $E' = \{(i,j) \in E \mid j \notin S\}$). It turns out that P' is also an MRF with respect to G'. [5] Hence, the corresponding *a posteriori* problems are equivalent to the *a priori* problems described above, but on a different MRF. That is, the *maximum a posteriori (MAP) assignment* of P given $X_S = e_S$ is the most likely assignment of MRF P'. The *posterior marginals* $P(X_i = x_i \mid X_S = e_S)$ for each i are the marginals of $P'(X_i = x_i)$. Computing the probability of the evidence $P(X_S = e_S)$ is

function (CDF) F_i for the corresponding random variable X_i such that $F_i(0) \equiv 0$ and for all $t = 1, \ldots, m$, $F_i(t) \equiv F_i(t-1) + P(X_i = x_i^t) = \sum_{s=1}^{m} P(X_i = x_i^s)$. Then, we can draw a random sample from $R \sim \text{Uniform}[0,1]$ and let $X_i = x_i^t$ if $F_i(t-1) < R \leq F_i(t)$. Note that, we can apply the same approach to any probability distribution. Hence, we can apply the same idea recursively to the non-source nodes once we have assignments for all the variables that are parents to the node in the graph, because the values of the parents of i uniquely determine the distribution over X_i. Also, we can compute the CDFs as a pre-processing step.

[5] To see this, let $\mathcal{C}_S \equiv \{C \in \mathcal{C} \mid C \subset S\}$, $\mathcal{C}_{V'} \equiv \{C \in \mathcal{C} \mid C \subset V'\}$ and $\mathcal{C}_{\mathcal{N}(S)} \equiv \mathcal{C} - (\mathcal{C}_S \cup \mathcal{C}_{V'}) = \{C \in \mathcal{C} \mid C - S \neq \emptyset, C \cap S \neq \emptyset\}$ form a partition of $\mathcal{C}$ such that each set contains the cliques that are subsets of S, subsets of V' and have elements both in S and V', respectively. Let $\mathcal{C}'_{\mathcal{N}(S)} \equiv \{C - S \mid C \in \mathcal{C}_{\mathcal{N}(S)}\}$ and $\mathcal{C}' \equiv \mathcal{C}_{V'} \cup \mathcal{C}'_{\mathcal{N}(S)}$. In addition, for each $C' \in \mathcal{C}_{V'}$, let $\phi'_{C'}(x_{C'}) \equiv \phi_{C'}(x_{C'})$, and for each $C' \in \mathcal{C}'_{\mathcal{N}(S)}$, let $\phi'_{C'}(x_{C'}) \equiv \sum_{C \in \mathcal{C}_{\mathcal{N}(S)}, C' = C - S} \phi_C(x_{C'}, e_{C - C'})$.

related to computing the normalizing constant of P, which in turn is related to computing the normalizing constant $Z' = \sum_{x_{V'}} \exp(\sum_{C' \in \mathcal{C}'} \phi'_{C'}(x_{C'}))$ of P'. [6]

As a result of this "equivalence," the ensuing presentation refers to the problem of computing the most-likely assignment and individual marginals as the *MAP* and *belief inference problems*, respectively.

Brief Overview of Computational Results in Probabilistic Graphical Models Both the exact and approximate versions of most inference-related problems in MRFs are in general intractable (e.g., NP-hard), although polynomial time algorithms do exists for some special cases (see Istrail [2000] and the references therein). The complexity of exact algorithms is usually characterized by structural properties of the graph, and the typical statement is that running times are polynomial only for graphs with bounded treewidth (see, e.g., Russell and Norvig [2003] for more information). Several deterministic and randomize approximation approaches exist (see, e.g., Jordan et al. [1999], Jaakkola [2000], Geman and Geman [1984]).

2.2 Game Theory

Game theory [von Neumann and Morgenstern, 1947] provides a mathematical model of the stable behavior (or outcome) that may result from the interaction of rational individuals.

Then, we have

$$\begin{aligned}
P'(X_{V'} = x_{V'}) \equiv & P(X_{V'} = x_{V'} \mid X_S = e_S) \\
\propto & P(X_{V'} = x_{V'}, X_S = e_S) \\
\propto & \exp(\textstyle\sum_{C \in \mathcal{C}_S} \phi_C(e_C)) \times \\
& \exp(\textstyle\sum_{C \in \mathcal{C}_{\mathcal{N}(S)}} \phi_C(x_{C-S}, e_{C \cap S})) \times \exp(\textstyle\sum_{C \in \mathcal{C}_{V'}} \phi_C(x_C)) \\
\propto & \exp(\textstyle\sum_{C \in \mathcal{C}_{\mathcal{N}(S)}} \phi_C(x_{C-S}, e_{C \cap S})) \times \exp(\textstyle\sum_{C \in \mathcal{C}_{V'}} \phi_C(x_C)) \\
= & \exp(\textstyle\sum_{C' \in \mathcal{C}'_{\mathcal{N}(S)}} \sum_{C \in \mathcal{C}_{\mathcal{N}(S)}, C' = C - S} \phi_C(x_{C'}, e_{C-C'})) \times \\
& \exp(\textstyle\sum_{C' \in \mathcal{C}_{V'}} \phi_{C'}(x_{C'})) \\
= & \exp(\textstyle\sum_{C' \in \mathcal{C}'_{\mathcal{N}(S)}} \phi'_{C'}(x_{C'})) \times \exp(\textstyle\sum_{C' \in \mathcal{C}_{V'}} \phi_{C'}(x_{C'})) \\
= & \exp(\textstyle\sum_{C' \in \mathcal{C}'} \phi'_{C'}(x_{C'}))\,.
\end{aligned}$$

[6] This is because, using some steps from the last derivation (for P'),

$$\begin{aligned}
P(X_S = e_S) = & \textstyle\sum_{x_{V'}} P(X_{V'} = x_{V'}, X_S = e_S) \\
= & \textstyle\sum_{x_{V'}} P(X_{V'} = x_{V'}, X_S = e_S) \\
= & \tfrac{1}{Z} \exp(\textstyle\sum_{C \in \mathcal{C}_S} \phi_C(e_C)) \times \\
& \textstyle\sum_{x_{V'}} \exp(\sum_{C \in \mathcal{C}_{\mathcal{N}(S)}} \phi_C(x_{C-S}, e_{C \cap S})) \times \exp(\textstyle\sum_{C \in \mathcal{C}_{V'}} \phi_C(x_C)) \\
= & \tfrac{1}{Z} \exp(\textstyle\sum_{C \in \mathcal{C}_S} \phi_C(e_C)) \times \textstyle\sum_{x_{V'}} \exp(\sum_{C' \in \mathcal{C}'} \phi'_{C'}(x_{C'})) \\
= & \tfrac{Z'}{Z} \exp(\textstyle\sum_{C \in \mathcal{C}_S} \phi_C(e_C))\,.
\end{aligned}$$

The paper concentrates in *noncooperative* settings: individuals maximize their *own* utility, act *independently* and do not have (direct) control over the behavior of others. [7]

The concept of *equilibrium* is central to game theory. Roughly, an equilibrium in a noncooperative game is a point of strategic stance, where no individual player can gain by *unilaterally* deviating from the equilibrium behavior.

Games and their Representation Let $V = [n]$ denote a finite set of n players in a game. For each player $i \in V$, let A_i denote the set of *actions* or *pure strategies* that i can play. Let $A \equiv \times_{i \in V} A_i$ denote the set of *joint actions*, $x \equiv (x_i, \ldots, x_n) \in A$ denote a joint action, and x_i the individual action of player i in x. Denote by $x_{-i} \equiv (x_1, \ldots, x_{i-1}, x_{i+1}, \ldots, x_n)$ the joint action of all the players *except* i, such that $x \equiv (x_i, x_{-i})$. Let $M_i : A \to \mathbb{R}$ denote the *payoff/utility function* of player i. If the A_i's are finite, then M_i is called the *payoff matrix* of player i. Games represented this way are called *normal-* or *strategic-form games*.

There are a variety of compact representations for large game inspired by probabilistic graphical models in AI and machine learning [La Mura, 2000, Kearns et al., 2001, Koller and Milch, 2003, Leyton-Brown and Tennenholtz, 2003, Jiang and Leyton-Brown, 2008]. The results of this paper are presented in the context of the following generalization of *graphical games* [Kearns et al., 2001], a simple but powerful model inspired by probabilistic graphical models such as MRFs. [8]

Definition 2. *A* graphical multi-hypermatrix game (GMG) *is defined by*

- *a* directed *graph $G = (V, E)$ in which there is a node $i \in V$ in G for each of the n players in the game (i.e., $|V| = n$), and the set of directed edges, or arcs, E defines a set of* neighbors $\mathcal{N}(i) \equiv \{j \mid (j, i) \in E, i \neq j\}$ *whose action affect the payoff function of i (i.e., j is a neighbor of i if and only if there is an arc from j to i); and*
- *for each player $i \in V$,*
  - *a set of actions A_i,*
  - *a hypergraph where the vertex set is its* (inclusive) neighborhood $N(i) \equiv \mathcal{N}(i) \cup \{i\}$ *and the hyperedge set is a set of* cliques *of players $\mathcal{C}_i \subset 2^{N(i)}$, and*
  - *a set $\{M'_{i,C} : A_C \to \mathbb{R} \mid C \in \mathcal{C}_i\}$ of* local-clique payoff (hyper)matrices.

The interpretation of a GMG is that, for each player i, the local *and* global payoff (hyper)matrices $M'_i : A_{N(i)} \to \mathbb{R}$ *and* $M_i : A \to \mathbb{R}$ *of i are (implicitly) defined as* $M'_i(x_{N(i)}) \equiv \sum_{C \in \mathcal{C}_i} M'_{i,C}(x_C)$ *and* $M_i(x) \equiv M'_i(x_{N(i)})$*, respectively.*

If for each player, each clique set is a singleton, we obtain a directed *graphical game*, and the single clique in the set defines the neighborhood of the player (i.e., in that case, $\mathcal{C}_i = \{N(i)\}$ for all i). If, in addition, the graph G is undirected then we obtain the original undirected version of graphical game due to Kearns et al. [2001]. Furthermore, if, in addition, each clique

[7] Individual rationality here means that each player seeks to maximize their own utility. Also note that, while many parlor "win-lose"/zero-sum games involve competition, in general, *noncooperative* $\neq$ *competitive*: each player just wants to do the best for himself, regardless of how useful or harmful his behavior is to others.

[8] Connections have already been established between the different kinds of compact representations [Jiang and Leyton-Brown, 2008], which may facilitate extensions of ideas, frameworks and results to those alternative models.

is the complete set of players, then the game is a standard *normal-form game* (also called strategic- or matrix-form game) (i.e., in that case, $N(i) = V$ for all i). A *hypergraphical game* is a GMG with the property that if, for all i and $j \in \mathcal{N}(i)$, we have that for all $C \in \mathcal{C}_i$, if $i, j \in C$ then $C \in \mathcal{C}_j$. [9] That is, in contrast to hypergraphical games, for example, a GMG does not require that the same "sub-game" be shared among all players in the clique of the "sub-game." In fact, a local-clique payoff hypermatrix involving several players may appear in the summation defining the local payoff hypermatrix of exactly one player. A *polymatrix game* is a hypergraphical game in which all the hyperedge sets are the same and equal to $\{\{i, j\} \mid i, j \in V, j \neq i\}$, which is the set of cliques of pairs of nodes involving the player and every other player. If, instead, the hyperedge set of each player is some (possibly different) *subset* of $\{\{i, j\} \mid j \in V, j \neq i\}$, then let us call the game a *graphical polymatrix game*. Finally, let us say that a hypergraphical game is *hyperedge-symmetric* if, in addition, for every hyperedge C containing players i, j in the game, we have that $M'_{i,C} = M'_{j,C} \equiv M'_C$; if, in particular, the hypergraphical game is a graphical polymatrix game, then let us say that the game is *pairwise-symmetric*.

The representation size of a GMG is $O(\sum_{i \in V} \sum_{C \in \mathcal{C}_i} \prod_{j \in C} |A_j|)$. Hence, the size is dominated by the representation of the local-clique payoff matrices, which are each of size exponential in their respective clique size. [10] However, this representation size could be considerably smaller than for a graphical game, which is exponential in the neighborhood size. Local-clique symmetric GMGs afford even further savings in representation size because we would only need a global set of hyperedges, along with their corresponding hypermatrices.

A special subclass of GMGs was used by Yu and Berthod [1995] to establish an equivalence between *local* maximum-*a-posteriori* (MAP) inference in Markov random fields and Nash equilibria of the game, a topic revisited in Section 4.1.

Equilibria as Solution Concepts. Equilibria is considered *the* solutions of a game. Various notions of equilibria exist. A *pure strategy (Nash) equilibrium (PNE)* of a game is a joint action x^* such that for all players i, and for all actions x_i, $M_i(x_i^*, x_{-i}^*) \geq M_i(x_i, x_{-i}^*)$. That is, no player can improve its payoff by *unilaterally* deviating from its prescribed equilibrium x_i^*, assuming the others stick to their actions x_{-i}^*. Some games, like Prisoner's Dilemma, have PNE; many others, like Rock-Paper-Scissors, *do not*. This is problematic because it will not be possible to "solve" some games using PNE.

A *mixed-strategy* of player i is a probability distribution Q_i over A_i such that $Q(x_i)$ is the probability that i chooses to play action x_i. [11] A *joint mixed-strategy* is a joint probability

[9] Hypergraphical games are very close to GMGs. However, their exact definition, as originally given by Papadimitriou [2005], is unclear to me. Such games may in fact be exactly GMGs. As a side note, a GMG also has the polynomial expectation property and thus a polynomial correlated equilibrium scheme [Papadimitriou, 2005], a topic we revisit in Section 5.

[10] The local-clique payoff hypermatrices can in principle use a variety of representations, including even more succinct parametric representations that exploit context-dependent conditional "action" independence, such as that used in action-graph games [Jiang and Leyton-Brown, 2008]. Here, however, the emphasis is in the tabular or hypermatrix representation, which is assumed throughout the remainder of the paper.

[11] Note that the sets of mixed strategies contain pure strategies, as we can always recover playing a pure strategy exclusively.

distribution Q capturing the players behavior, such that $Q(x)$ is the probability that joint action x is played, or in other words, each player i plays action in component x_i of x. Because we are assuming that the players play *independently*, Q is a product distribution: $Q(x) = \prod_i Q_i(x_i)$. Denote by $Q_{-i}(x_{-i}) \equiv \prod_{j \neq i} Q_j(x_j)$ the joint mixed strategies of all the players except i. The *expected payoff* of a player i when some joint mixed-strategy Q is played is $\sum_x Q(x)M_i(x)$; abusing notation, denote it by $M_i(Q)$. The *conditional expected payoff* of a player i given that he plays action x_i is $\sum_{x_{-i}} Q_{-i}(x_{-i})M_i(x_i, x_{-i})$; abusing notation again, denote it by $M_i(x_i, Q_{-i})$.

A *Nash equilibrium (NE)* is a joint mixed-strategy Q^* that is a product distribution formed by the individual players mixed strategies Q_i^* such that, for all players i, and any other alternative mixed strategy Q_i' for his play, $M_i(Q_i^*, Q_{-i}^*) \geq M_i(Q_i', Q_{-i}^*)$. *Every* game in normal-form has at least one such equilibrium[Nash, 1950, 1951]. Thus, every game has a NE "solution."

One relaxation of NE considers the case where the amount of gain each player can obtain from unilateral deviation is very small. This concept is particularly useful to study approximation versions of the computational problem. Given $\epsilon \geq 0$, an *(approximate) ϵ-Nash equilibrium (NE)* is defined as above, except that the expected gain condition becomes $M_i(Q_i^*, Q_{-i}^*) \geq M_i(Q_i', Q_{-i}^*) - \epsilon$.

Several refinements and generalizations of NE have been proposed. One of the most interesting generalizations is that of a *correlated equilibrium (CE)* [Aumann, 1974]. In contrast to NE, a CE can be a full joint distribution, and thus characterize more complex joint-action behavior by players. Formally, a *correlated equilibrium (CE)* is a joint probability distribution Q over A such that, for all players i, $x_i, x_i' \in A_i$, $x_i \neq x_i'$, and $Q(x_i) > 0$, $\sum_{x_{-i}} Q(x_{-i}|x_i)M_i(x_i, x_{-i}) \geq \sum_{x_{-i}} Q(x_{-i}|x_i)M_i(x_i', x_{-i})$, where $Q(x_i) \equiv \sum_{x_{-i}} Q(x_i, x_{-i})$ is the (marginal) probability that player i will play x_i according to Q and $Q(x_{-i}|x_i) \equiv Q(x_i, x_{-i}) / \sum_{x_i'} Q(x_i', x_{-i})$ is the conditional given x_i. An NE is a product distribution CE. An equivalent expression of the CE condition above is

$$\sum_{x_{-i}} Q(x_i, x_{-i})M_i(x_i, x_{-i}) \geq \sum_{x_{-i}} Q(x_i, x_{-i})M_i(x_i', x_{-i}).$$

As was the case for NE, we can relaxed the condition of deviation to account for potential gains from small deviation. Given $\epsilon > 0$, An *(approximate) ϵ-CE* is defined by replacing the last condition with [12] $\sum_{x_{-i}} Q(x_i, x_{-i})M_i(x_i, x_{-i}) \geq \sum_{x_{-i}} Q(x_i, x_{-i})M_i(x_i', x_{-i}) - \epsilon$.

CE have several conceptual and computational advantages over NE. For instance, better expected payoffs may be achievable by all players [13]; some "natural" forms of play are guaranteed to converge to it [Foster and Vohra, 1997, 1999, Fudenberg and Levine, 1999, Hart and Mas-Colell, 2000, 2003, 2005]; and it is consistent with a Bayesian framework [Aumann, 1987], something not yet possible, and apparently unlikely for NE [Hart and Mansour, 2007].

[12] Note that approximate CE is usually defined based on this unconditional version of the CE conditions [Hart and Mas-Colell, 2000].

[13] The distinction between installing a traffic light at an intersection and leaving the intersection without one is a real-world example of this.

Brief Overview of Results in Computational Game Theory There has been an explosion of computational results on different equilibrium concepts on a variety of game representations and settings over the last few years. The following is a brief summary. The reader is referred to a recent book by [Nisan et al., 2007] for a (partial) introduction to this quickly growing research area.

The problem for two-player *zero-sum* games, where the sum of the entries of both matrix is zero, and therefore only one matrix is needed to represent the game, can be solved in polynomial time: It is equivalent to linear programming [von Neumann and Morgenstern, 1947, Szép and Forgoó, 1985, Karlin, 1959]. After being open for over 50 years, the problems of the complexity of computing NE in games was finally settled recently, following a very rapid sequence of results in the last part of 2005 [Goldberg and Papadimitriou, 2005, Daskalakis et al., 2005, Daskalakis and Papadimitriou, 2005, Daskalakis et al., 2009b, Chen and Deng, 2005b]: Computing NE is likely to be hard in the worst case (i.e., PPAD-complete [Papadimitriou, 1994]), even in games with only two players [Chen and Deng, 2005a, 2006, Chen et al., 2009, Daskalakis et al., 2009a,b]. The result of Fabrikant et al. [2004] suggests that computing PNE in succinctly representable games is also likely to be intractable in the worst case (i.e., PLS-complete [Johnson et al., 1988]). A common statement is that computing NE with "special properties" is hard in the worst case [Gilboa and Zemel, 1989, Gottlob et al., 2003, Conitzer and Sandholm, 2008]. Computing approximate NE is also thought to be hard in the worst case [Chen et al., 2006, 2009].

Most current results for computing exact and approximate NE in graphical games essentially mirror those for MRFs and constraint networks: polynomial time for bounded treewidth graph; intractable in general [Kearns et al., 2001, Gottlob et al., 2003, Daskalakis and Papadimitriou, 2006]. This is unsurprising because they were mostly inspired by analogous version in probabilistic graphical models and constraint networks in AI, and therefore share similar characteristics. Several heuristics exists for dealing with general graphs [Vickrey and Koller, 2002, Ortiz and Kearns, 2002, Daskalakis and Papadimitriou, 2006].

Computing CE on the other hand can be done in polynomial time, both for normal-form games (where the problem reduces to a simple linear feasibility problem) and, as described in the introduction, somewhat surprisingly, even most succinctly representable games known today [Papadimitriou, 2005], including graphical games.

3 Graphical Potential Games

Let us start by introducing a class of games that play a key role in the work presented in the next section. Some structural properties and characterizations of such games are presented. Previous work in probabilistic graphical models facilitate the derivations of the results presented here. Thus, this section serves as a preamble to our pursuit of a stronger connection between inference and equilibria, whose discussion begin in the next section.

Definition 3. *Consider a graph G with (non-inclusive) neighbor sets $\mathcal{N}(i)$ for each payer i. For any graphical game with graph G, and for each player i, consider some function $f_i : \mathbb{R} \times A_{\mathcal{N}(i)} \to \mathbb{R}$. Let us say that the function f_i is a* (conditional) preference-order-preserving

transform *for player i if $f_i(v, x_{\mathcal{N}(i)})$ is a (strictly) monotonically increasing function of v for every $x_{\mathcal{N}(i)} \in A_{\mathcal{N}(i)}$.* [14] *Let us call the transform* (unconditionally) linear *with respect to some positive weight w_i if it takes the form $f_i(v, x_{\mathcal{N}(i)}) = w_i v$. Let us denote by f_i^{-1} the corresponding* (conditional) inverse function *of f_i; that is, $f_i^{-1} : \mathbb{R} \times A_{\mathcal{N}(i)} \to \mathbb{R}$ such that for all $v \in \mathbb{R}$ and $x_{\mathcal{N}(i)} \in A_{\mathcal{N}(i)}$, we have $f_i^{-1}(f_i(v, x_{\mathcal{N}(i)}), x_{\mathcal{N}(i)}) = v$. Denote by $f \equiv \{f_i\}$ and $f^{-1} \equiv \{f_i^{-1}\}$ the set of preference-order-preserving transforms f_i and their inverses, respectively, one for each player i. Let us say that a function Ψ is a* f-transformed potential *for a graphical game with neighbor sets $\{\mathcal{N}(i)\}$ and local payoff matrices $\{M_i'\}$ if for all i, $x_{\mathcal{N}(i)}$, and x_i, x_i',*

$$M_i'(x_i, x_{\mathcal{N}(i)}) - M_i'(x_i', x_{\mathcal{N}(i)}) = f_i(\Psi(x_i, x_{-i}) - \Psi(x_i', x_{-i}), x_{\mathcal{N}(i)}) \ .$$

Let us call a graphical game with a f-transformed potential a graphical f-transformed potential game.

We can generalize the terminology of Monderer and Shapley [1996] to graphical games. Let w be a positive weight vector. If the transformed f_i of each player i is (unconditionally) linear with weight w_i, then Ψ is called a *(weighted) w-potential* for the game. If $w_i = 1$ for all i, then Ψ is called an *(exact) potential*. Finally, let us say that Ψ is an *ordinal potential* for a graphical game with neighbor sets $\{\mathcal{N}(i)\}$ and local payoff matrices $\{M_i'\}$ if it satisfy the following condition:

$$M_i'(x_i, x_{\mathcal{N}(i)}) - M_i'(x_i', x_{\mathcal{N}(i)}) > 0 \text{ if and only if } \Psi(x_i, x_{-i}) - \Psi(x_i', x_{-i}) > 0$$

for all i, x_i, x_i' and $x_{\mathcal{N}(i)}$. Let us refer to a graphical game with a (weighted) w-potential, (exact) potential and ordinal potential as a *graphical (weighted) w-potential, (exact) potential or ordinal potential game*, respectively.

It is not hard to see that *sequential best-response* dynamics, a process in which, at each time step, each player observes the action x_{-i} of others and take an action that maximizes its payoff given that the others played x_{-i}, always converges to a PNE in potential games. Hence, potential games always have PNE, in contrast to arbitrary games.

3.1 Characterizing Graphical Potential Games

The following theorem characterizes graphical f-transformed potential games: the potential function is the sum of local potential functions over each (maximal) clique in the game graph.

Theorem 2. *Every potential of a graphical transformed potential game with graph G is a Gibbs potential with respect to G.*

Proof. The proof is based on an application of the Hammersley-Clifford theorem (Theorem 1). Let Ψ be the f-transformed potential of the graphical game. Define P as a joint probability distribution such that $P(X = x) \propto \exp(\Psi(x))$ for all x. (Thus, P is positive.)

[14] Such transforms were considered in Monderer and Shapley [1996] in the context of (variants of) fictitious play.

The following derivation shows that P is an MRF with respect to the game graph G: for all i, x_i, x_{-i},

$$\begin{aligned} P(X_i = x_i \mid X_{-i} = x_{-i}) &= \frac{\exp(\Psi(x_i, x_{-i}))}{\sum_{x'_i} \exp(\Psi(x'_i, x_{-i}))} \\ &= \frac{1}{\sum_{x'_i} \exp(\Psi(x'_i, x_{-i}) - \Psi(x_i, x_{-i}))} \\ &= \frac{1}{\sum_{x'_i} \exp(f_i^{-1}(M'_i(x'_i, x_{\mathcal{N}(i)}) - M'_i(x_i, x_{\mathcal{N}(i)}), x_{\mathcal{N}(i)}))} \,. \end{aligned}$$

Hence, by the Hammersley-Clifford theorem (Theorem 1), P is also a Gibbs distribution with respect to the graph G of the game. In particular, let $\mathcal{C}$ be the set of (maximal) cliques of G. There exist a local potential function ϕ'_C for each (maximal) clique $C \in \mathcal{C}$ defining a global potential function Ψ' as $\Psi'(x) = \sum_{C \in \mathcal{C}} \phi'_C(x_C)$ such that $P(X = x) \propto \exp(\Psi'(x))$. Let $Z = \sum_x \exp(\Psi(x))$ and $Z' = \sum_x \exp(\Psi'(x))$ be the normalizing constant when expressing P in terms of Ψ and Ψ', respectively, denote by $c \equiv \ln(Z/Z')$, a constant. Then we have that for all x, $\Psi(x) = \Psi'(x) + c$. Defining local potential for Ψ as, for example, $\phi_C(x_C) \equiv \phi'_C(x) + c/|\mathcal{C}|$ completes the proof. □

In view of the strong equivalence established by the last theorem, let us refer to a graphical transformed potential game as a transformed *Gibbs potential game* (with the same graph); and similarly for weighted and exact potentials. As of yet, it is unclear to me whether every graphical ordinal potential game has a (equivalent) ordinal Gibbs potential. So, let us define an *ordinal Gibbs potential game* as a graphical game that has a Gibbs potential with respect to the graph of the game.

The following definitions are useful to present the main corollary of the last theorem.

Definition 4. *Given a n-dimensional positive weight vector w, let us say that a game with payoff matrices $\{M_i^1\}$ is* w-scaled payoff-difference equivalent *to another game with the same players and payoff matrices $\{M_i^2\}$ if for all i, x_{-i}, x_i, x'_i, we have that $M_i^1(x_i, x_{-i}) - M_i^1(x'_i, x_{-i}) = w_i(M_i^2(x_i, x_{-i}) - M_i^2(x'_i, x_{-i}))$. If w is the unit vector then let us simply say that the game is* payoff-difference equivalent *to the other.*

Note that in any two (w-scaled) payoff-difference-equivalent games, every player achieves exactly the same (w-scaled) expected regrets with respect to any (fixed, possibly correlated) joint-mixed strategy. Thus any two such games have the same correlated equilibria.

Definition 5. *Let us say that a graph has* totally disconnected (open) neighborhoods *if for every node of the graph, every subgraph induced by the neighbors of the node is the empty graph; in other words, there is no edge connecting any pair of neighbors of any node in the graph (i.e., formally, if E is the edge set, then for all i and $j, k \in \mathcal{N}(i), j \neq k$, we have $(j, k) \notin E$).*

Some simple examples of graphs with totally disconnected neighborhoods are trees, cycles and grids.

Corollary 1. *Any* w-weighted Gibbs potential *game with graph* G *and potential* Ψ *is* w-*scaled payoff-difference equivalent to a* hyperedge-symmetric hypergraphical *game in which each hyperedge is a (maximal) clique* C *in* G *and has as corresponding hypermatrix the local potential associated to* C *in* Ψ*. If, in addition,* G *has totally disconnected neighborhoods then the equivalence is to* pairwise-symmetric graphical polymatrix *games.*

Proof. The proof follows from the last theorem by applying the respective definitions. For the last statement of the theorem, in particular, note that the maximal cliques of a graph with disconnected neighborhoods are exactly the edges of the graph. Hence, each local potential is pairwise in that case, and the payoff difference is the sum of differences of local pairwise potentials, as is the case in any graphical polymatrix games, by definition. □

It is important to note the implication of the last corollary. In general, there does not seem to be any reason to expect, *a priori*, just from looking at the definition of a graphical potential game with, say for example, a tree graph that the differences in payoff matrix for a player would not be arbitrary functions of the action of the player and those of its neighbors. The corollary tells us that this is not possible in this case: the payoff matrices difference must be sums of simple pairwise matrices, each being a simple 2-dimensional matrix depending on the actions of the player and one of its neighbors. The same holds for cycles, grids and similar structure, and their corresponding generalizations, including those to higher dimensions.

The following proposition completes the (easier) connection in the reverse direction.

Proposition 1. *Any hyperedge-symmetric hypergraphical game is a graphical Gibbs potential game.*

Proof. For each hyperedge C in the hypergraphical game, define a corresponding local potential $\phi_C(x_C) \equiv M'_C(x_C)$ as the local-clique payoff matrix for C. Defining the Gibbs potential using those local potentials and with the graph being the primal graph of the hypergraph, which is an undirected graph with the same vertex set as the hypergraph but where there is an edge between two nodes if there is a hyperedge containing both nodes. [15] □

3.2 Smooth Best-Response Play and Graphical Games

Let us consider a *sequential process of play*, where there is a pre-specified order in which the players play, and at each time step t exactly one player plays by choosing an action x_i^t. [16] In the sequel, let us take the sequence of play to be, without loss of generality, the sequence $1, 2, \ldots, n$. Let us say that a player has a *(time-homogeneous, first-order) Markov playing scheme* if it has a (possibly randomized) policy, or plan, by which the agent selects an action based only on the *last actions* played by the other players; more formally, if the policy p_i is a conditional probability distribution $p_i(x_i \mid x_{-i})$ such that if player i is to play at time

[15] Note that in this case, because we are going in the "reverse" direction, local potential functions do not need to be over *maximal* cliques.

[16] This process may be relaxed to allow certain dynamic variations in the order of play and some kinds of simultaneous moves. The process is also related to (smooth versions of) the *Cournot adjustment process with lock-in* and to *stochastic adjustment models* in the literature on learning in games [Fudenberg and Levine, 1999].

$t+1$ and x_{-i}^t are the *last* joint-actions that player i observed the others take, then player i chooses to play action $x_i^{t+1} = x_i$ with probability $p_i(x_i \mid x_{-i}^t)$ (i.e., $x_i^{t+1} \sim p_i(. \mid x_{-i}^t)$). [17] For a graphical game, let us further say that the playing scheme is *local* if it *only* depends on the last actions of its *neighbors* in the game graph.

Property 1. In a graphical game with graph G, the sequential process of play generated by Markov playing schemes of the type described above, and that are local with respect to G, is equivalent to realizations generated by running the Gibbs sampler with conditional distributions given by the individual player's playing scheme.

For every round $r = 1, 2, \ldots$ composed of consecutive time steps of length n, denote by $z^r = (z_1^r, \ldots, z_n^r)$ the *play (joint-action) outcome* at round r, so that for all i, $z_i^r = x_i^{(r-1)n+i}$ (i.e., the joint-action generated during r by the sequential process of play). Let us refer to the total sequence generated by z^r for $r = 1, 2, \ldots, T$ as the *empirical play* up to round T (i.e., after nT rounds). Let us say that an empirical play *(conditionally) converges* starting from an initially assigned play x^0 if the empirical joint-probability distribution defined by the empirical play converges (almost surely) to some joint probability distribution as we let the sequential process of play run for an infinite number of rounds (i.e., for every joint-action x, if we denote the empirical distribution of play after T rounds as $\widehat{P}^T(x) \equiv \frac{1}{T}\sum_{r=1}^{T} 1\left[z^r = x\right]$, then $\widehat{P}^\infty(x) \equiv \lim_{T\to\infty} \widehat{P}^T(x)$ exists, with probability one.) Let us refer to a set of playing schemes, one for each player, as a *playing procedure* for the game. Furthermore, let us say that a playing procedure is *consistent* (or *globally convergent*) if the empirical play generated is convergent to the *same* (possibly correlated) joint mixed-strategy $\widehat{P}^\infty(x)$ from *any* initial joint play x^0.

Let us say that a player i with payoff function M_i uses a *smooth best-response (SBR) playing scheme* p_i with respect to a (conditional) preference-order-preserving transform $f_i : \mathbb{R} \times A_{-i} \to \mathbb{R}$, if for all x_{-i}, we have $p_i(x_i \mid x_{-i}) \propto \exp(f_i(M_i(x_i, x_{-i}), x_{-i})$. If, instead, the f_i's above are such that the ratio $p_i(x_i \mid x_{-i})/p_i(x_i' \mid x_{-i}) = \exp(f_i(M_i(x_i, x_{-i}) - M_i(x_i', x_{-i}), x_{-i}))$, then let us call the scheme a *smooth best-response-difference (SBRD) playing scheme.* [18] Recall from the previous discussion that, for a graphical game, a player's (playing) scheme is *local* with respect to the game graph, if it only depends on the actions of its neighbors in that graph. In the case of SBR and SBRD schemes, in general, this requires that the domain of the second argument to f_i be $A_{\mathcal{N}(i)}$. Finally, let us say that the game has a *(local, Markov) playing procedure* if every player uses a (local, Markov) playing scheme.

The following theorems and corollaries provide another characterization of graphical potential games.

Theorem 3. *Any graphical game with graph G that has a consistent (Markov) local SBR playing procedure is an ordinal Gibbs potential game with graph G. If, in particular, the transform f_i used by each player i is (unconditionally) linear with weight $1/w_i$, then the game is a w-weighted Gibbs potential game. If the playing procedure is SBRD, instead of SBR, then*

[17] Although not pursued in this paper, more complex playing scheme could in principle be considered.

[18] This condition implies that $f_i(v, x_{-i}) = -f_i(-v, x_{-i})$ and $p_i(x_i \mid x_{-i}) = 1/(\sum_{x_i'} \exp(f_i(M_i(x_i', x_{-i}) - M_i(x_i, x_{-i}), x_{-i}))$.

the game is a f^{-1}-transformed Gibbs potential game, where f^{-1} denote the inverses of the transforms f used by the players for the SBRD scheme.

Proof. As noted in Property 1, the playing procedure is equivalent to the Gibbs sampler. The Markov chain associated with the Gibbs sampler is regular because all the conditional probabilities associated with the players' schemes are positive. Because the procedure is convergent the limiting empirical distribution of play $\widehat{P}$ is the (unique) MRF consistent with the conditional probability distributions. Because $\widehat{P}$ is positive, by the Hammersley-Clifford theorem (Theorem 1), $\widehat{P}$ is also a Gibbs distribution with respect to the graph of the game G. Denote by $\widehat{\Psi}$ the Gibbs potential of $\widehat{P}$. Also, denote by $\widehat{Z}_i(x_{-i}) \equiv \sum_{x'_i} \exp(\widehat{\Psi}(x'_i, x_{-i}))$ the normalizing constant for the marginal distribution over all the variables except i with respect to $\widehat{P}$. Similarly, denote by $Z_i(x_{\mathcal{N}(i)}) \equiv \sum_{x'_i} \exp(f_i(M'_i(x'_i, x_{\mathcal{N}(i)}), x_{\mathcal{N}(i)}))$ the normalizing constant of the player's scheme. The conditionals of MRF $\widehat{P}$ must satisfy the following condition:

$$\begin{aligned}
\frac{\exp(\widehat{\Psi}(x_i, x_{-i}))}{\widehat{Z}_i(x_{-i})} &= \widehat{P}(X_i = x_i \mid X_{-i} = x_{-i}) \\
&= \widehat{P}(X_i = x_i \mid X_{\mathcal{N}(i)} = x_{\mathcal{N}(i)}) \\
&= p_i(x_i \mid x_{\mathcal{N}(i)}) \\
&= \frac{\exp(f_i(M'_i(x_i, x_{\mathcal{N}(i)}), x_{\mathcal{N}(i)}))}{Z_i(x_{\mathcal{N}(i)})} \, .
\end{aligned}$$

From the last equality, we obtain that for all x_i, x_{-i},

$$f_i(M'_i(x_i, x_{\mathcal{N}(i)})) = \widehat{\Psi}(x_i, x_{-i}) + \ln(Z_i(x_{\mathcal{N}(i)})/\widehat{Z}_i(x_{-i})) \, .$$

Because the second term on the right hand side of the last equation does not depend on x_i, we can obtain that for all x_{-i} and every pair x_i, x'_i,

$$f_i(M'_i(x_i, x_{\mathcal{N}(i)}), x_{\mathcal{N}(i)}) - f_i(M'_i(x'_i, x_{\mathcal{N}(i)}), x_{\mathcal{N}(i)}) = \widehat{\Psi}(x_i, x_{-i}) - \widehat{\Psi}(x'_i, x_{-i}) \, . \tag{1}$$

Because the f_i's are (strictly) monotonically increasing with their first argument, we have

$$\begin{aligned}
\widehat{\Psi}(x_i, x_{-i}) - \widehat{\Psi}(x'_i, x_{-i}) > 0 \iff & f_i(M'_i(x_i, x_{\mathcal{N}(i)}), x_{\mathcal{N}(i)}) - \\
& f_i(M'_i(x'_i, x_{\mathcal{N}(i)}), x_{\mathcal{N}(i)}) > 0 \\
\iff & f_i(M'_i(x_i, x_{\mathcal{N}(i)}), x_{\mathcal{N}(i)}) > \\
& f_i(M'_i(x'_i, x_{\mathcal{N}(i)}), x_{\mathcal{N}(i)}) \\
\iff & M'_i(x_i, x_{\mathcal{N}(i)}) > M'_i(x'_i, x_{\mathcal{N}(i)}) \\
\iff & M'_i(x_i, x_{\mathcal{N}(i)}) - M'_i(x'_i, x_{\mathcal{N}(i)}) > 0 \, .
\end{aligned}$$

This completes the proof of the first statement.

The second statement in the theorem follows by noting that in the case of (unconditionally) linear transforms with weight w_i, we have that the left-hand side of equation 1 above becomes $f_i(M_i'(x_i, x_{\mathcal{N}(i)}), x_{\mathcal{N}(i)}) - f_i(M_i'(x_i', x_{\mathcal{N}(i)}), x_{\mathcal{N}(i)}) = w_i M_i'(x_i, x_{\mathcal{N}(i)}) - w_i M_i'(x_i', x_{\mathcal{N}(i)}) = w_i(M_i'(x_i, x_{\mathcal{N}(i)}) - M_i'(x_i', x_{\mathcal{N}(i)}))$.

The case of SBRD procedure is similar (but simpler). Note that in that case

$$\begin{aligned}\exp(f_i(M_i(x_i, x_{\mathcal{N}(i)}) - M_i(x_i', x_{\mathcal{N}(i)}), x_{\mathcal{N}(i)})) &= \frac{p_i(x_i \mid x_{\mathcal{N}(i)})}{p_i(x_i' \mid x_{\mathcal{N}(i)})} \\ &= \exp(\widehat{\Psi}(x_i, x_{-i}) - \widehat{\Psi}(x_i', x_{-i}))\ .\end{aligned}$$

Hence, we have $M_i(x_i, x_{\mathcal{N}(i)}) - M_i(x_i', x_{\mathcal{N}(i)}) = f_i^{-1}(\widehat{\Psi}(x_i, x_{-i}) - \widehat{\Psi}(x_i', x_{-i}), x_{\mathcal{N}(i)})$, and the proof is completed by the definition of a graphical transformed potential game. □

The next proposition completes the characterization of transformed Gibbs potential games as exactly those that have consistent local Markov SBRD playing procedures.

Proposition 2. *Any f-transformed Gibbs potential game with graph G has a consistent Markov SBRD playing procedure that is local with respect to G and each player uses the inverse transforms of f in their playing scheme. If, in particular, the game is a w-weighted Gibbs potential game, then it also has an SBR procedure with the same properties.*

Proof. Let Ψ be the Gibbs potential of the game. Define the SBRD playing scheme for each player using the inverse f_i^{-1} of the corresponding game transform, such that

$$\begin{aligned}\frac{p_i(x_i \mid x_{\mathcal{N}(i)})}{p_i(x_i' \mid x_{\mathcal{N}(i)})} &\equiv \exp(f_i^{-1}(M_i(x_i, x_{\mathcal{N}(i)}) - M_i(x_i', x_{\mathcal{N}(i)}), x_{\mathcal{N}(i)})) \\ &= \exp(\widehat{\Psi}(x_i, x_{-i}) - \widehat{\Psi}(x_i', x_{-i}))\end{aligned}$$

where the second equality follows from the definition of transformed potential games. The scheme is local by construction. Because Ψ is a Gibbs potential, by the Hammersley-Clifford theorem (Theorem 1), there is a positive joint (global) MRF consistent with the conditional distributions induced by the scheme. Hence, running the Gibbs sampler will always converge, regardless of initial conditions. Because the playing procedure is equivalent to running the Gibbs sampler (Property 1), the corresponding empirical distribution of play will always converge to the same Gibbs distribution with potential Ψ. Thus, the procedure will be consistent.

In the special case of a w-weighted Gibbs potential game, then each player i's scheme takes the form

$$p_i(x_i \mid x_{\mathcal{N}(i)}) = \frac{\exp(M_i(x_i, x_{\mathcal{N}(i)})/w_i)}{\sum_{x_i'} \exp(M_i(x_i', x_{\mathcal{N}(i)})/w_i)}$$

which corresponds to an SBR with a (unconditionally) linear transform with weight $1/w_i$. □

Whether the last theorem extends to graphical *ordinal* potential games is left open.

Concluding Remarks The following are some concluding remarks for this section.

Remark 1. Recent work [Daskalakis et al., 2007] state that if random graphical games with graphs of n nodes are generated according to an Erdös-Rényi model with p being the probability of an edge, and the local payoff matrices are generated uniformly at random according to that graph, then, if $1/n << np < (0.5 - \Omega(1)) \log n$, then it is unlikely that the game will have a PNE. This result suggests then that it is unlikely that games generated at random will have payoffs whose normalized versions correspond to conditional distributions of some full joint distribution with full support. The same is true for large graphical games with bounded degree graphs (e.g., long chains, large trees, etc.). But if the payoff values were generated uniformly at random by, for example, for each joint-action of the neighbors $x_{\mathcal{N}(i)}$, generating, for each action of the player x_i, a random number $U_i(x_i, x_{\mathcal{N}(i)}) \sim \text{Uniform}[0, 1]$ uniformly at random from the interval $[0, 1]$ and then setting the payoff value $M_i(x_i, x_{\mathcal{N}(i)}) = \ln(U_i(x_i, x_{\mathcal{N}(i)}))/\sum_{x_i'} \ln(U_i(x_i, x_{\mathcal{N}(i)}))$, then the normalized payoff can be interpreted as conditional distributions generated uniformly at random from the set of all possible probability distributions. Thus, we conclude that for such graphs, it is unlikely that conditional distributions generated uniformly at random from the set of all possible probability distributions are consistent with an MRF with that graph, with high probability. This in turn implies that running the Gibbs sampler with those conditionals will not converge [Casella and George, 1992].

Remark 2. One could envision modified, *learning-like* versions of the SBR and SBRD playing procedures that would converge to the pure strategy that is equivalent to a global optimum of the corresponding potential function in a graphical potential game. Such a variant would have further connections to more general MCMC methods, such as Metropolis-Hastings, used in a variety of fields (e.g., statistics and physics) and simulated annealing used for search and optimization (see, e.g., Russell and Norvig [2003] for more information).

Remark 3. It is known that the Gibbs sampler may fail to converge if the local conditional probabilities are not consistent with a full joint distribution. I am unaware of any work that establishes how hard it is to test for that global condition. It is not hard to derive a polynomial-time algorithm to test consistency in general bounded-treewidth and other similar tree-like graphs, by applying the same dynamic-programming approach to exact inference in graphical models. In general, however, one would guess this problem is hard. Indeed, it is unclear whether it even belongs to NP: what would be a polynomial-time-checkable solution "certificate" that would verify an answer to the consistency question? The (seemingly related) problem of testing the global consistency of a set of arbitrary, pairwise consistent marginals is NP-complete [Sontag, 2007]. (We will revisit the consistent-marginals problem in Section 5.)

4 Probabilistic Inference and Equilibria

The line of work presented in this section is partly motivated by the following question: *Can we leverage advances in computational game theory for problems in the graphical models*

community? Establishing a strong bilateral connection between both problems may help us answer this question.

The literature on computing equilibria in games has skyrocketed over the last few years. As we discover techniques developed early on within the game theory community, and as new novel results are generated from the extremely active computational game theory community, we may be able to adapt those techniques from solving games to the inference setting. If we can establish a strong bilateral connection between inference problems and the computation of equilibria, we may be able to relate algorithms in both areas and exchange previously unknown results in each.

4.1 Pure-Strategy Nash Equilibrium and Approximate MAP Inference

Consider an MRF P with respect to graph G and Gibbs potential Ψ defined by the set potential functions $\{\phi_C\}$. For each node i, denote by $\mathcal{C}_i \subset \mathcal{C}$ the subset of cliques in G that include i. Note that the (inclusive) neighborhood of player is given by $N(i) = \cup_{C \in \mathcal{C}_i} C$.

Define an *MRF-induced* (hyperedge-symmetric) hypergraphical game with the same graph G, and for each player i, hypergraph with hyperedges $\mathcal{C}_i$ and local-clique payoff hypermatrices $M'_{i,C}(x_C) \equiv \phi_C(x_C)$ for all $C \in \mathcal{C}_i$. A few observations about the game are in order.

Property 2. The representation size of the MRF-induced game is the same as that of the MRF: not exponential in the largest neighborhood size, but the size of the largest clique in G.

Property 3. By Proposition 1, the MRF-induced game is a (exact) Gibbs potential game with graph G and potential Ψ: i.e., for all i, x and x'_i,

$$\begin{aligned}
M_i(x_i, x_{-i}) - M_i(x'_i, x_{-i}) &= M'_i(x_i, x_{\mathcal{N}(i)}) - M'_i(x'_i, x_{\mathcal{N}(i)}) \\
&= \sum_{C \in \mathcal{C}_i} \phi_C(x_i, x_{C-\{i\}}) - \sum_{C \in \mathcal{C}_i} \phi_C(x'_i, x_{C-\{i\}}) \\
&= \sum_{C \in \mathcal{C}_i} \phi_C(x_i, x_{C-\{i\}}) + \sum_{C' \in \mathcal{C}-\mathcal{C}_i} \phi_{C'}(x_{C'}) + \\
&\quad - \sum_{C \in \mathcal{C}_i} \phi_C(x'_i, x_{C-\{i\}}) - \sum_{C' \in \mathcal{C}-\mathcal{C}_i} \phi_{C'}(x_{C'}) \\
&= \Psi(x_i, x_{-i}) - \Psi(x'_i, x_{-i}).
\end{aligned}$$

A few additional remarks follow.

Remark 4. Through the connection established by the last property, it is easy to see that *sequential* best-response dynamics is guaranteed to converge to a PNE of the game in finite time, regardless of the initial play. [19] In fact, we can conclude that a joint-action x^* is a PNE

[19] Recall that *best-response dynamics* refers to the a process where at each time step, each player observes the action x_{-i} of others and take an action that maximizes its payoff given that the others played x_{-i}. In this case, those dynamics would essentially be implementing an axis-parallel coordinate maximization over the space of assignments for the MRF, which is guaranteed to converge to a local maxima (or critical points) of the MRF.

of the game if and only if x^* is a local maxima or a critical point of the MRF P. Thus, the MRF-induced game, like *all* potential games, always has PNE. [20]

Similarly, for any potential game, one can define a *game-induced MRF* using the potential function of the game whose set of local maxima (and critical points) corresponds exactly to the set of PNE of the potential game. Through this connection we can show that solving the local-MAP problem in MRFs is PLS-complete in general [Fabrikant et al., 2004]. [21]

One question that comes to mind is whether one can say anything about the properties of the globally optimal assignment in the game-induced MRF and the payoff it supports for the players. Or whether it can be characterized by stronger notions of equilibria. For example, are *strong NE*, in which no *coalition* of players could obtain a Pareto dominated set of payoffs by unilaterally deviating, (global) MAP assignments of the MFR? Or more generally, what characteristics can we assign to the MAP assignments of the game-induced MRF?

In short, we can use algorithms for PNE as heuristics to compute locally optimal MAP assignments of P and *vice versa*. [22]

Remark 5. Recent work [Daskalakis et al., 2007] has extended results in game theory characterizing the number of PNE in normal-form games (see Stanford [1995], Rinott and Scarsini [2000], and the references therein) to graphical games, but now taking into consideration the network structure of the game. Information about the number of PNE in games can provide additional insight on the structure of MRFs.

For example, one of the results of Daskalakis et al. [2007] states that for graphs respecting certain expansion properties as the number of nodes/players increases, the number of PNE of the graphical game will have a limiting distribution that is a Poisson with expected value 1. Also according to [Daskalakis et al., 2007], a similar behavior occurs for games with graphs generated according to the Erdös-Rényi model with sufficiently high average-degree (i.e., reasonably high connectivity). Thus, either the set of MRF-induced games has significantly low measure relative to the set of all possible randomly generated games (something that seems likely), or the number of local maxima (and critical points) of the MRF will have a similar distribution, and thus that number is expected to be low. The latter would suggest that local algorithms such as the max-product algorithm may be less likely to get stuck in local maxima (or critical points) of the MRF.

[20] This result should not be surprising given that other researchers have established a one-to-one relationship between the complexity class PLS [Johnson et al., 1988], which characterizes local search problems, of which finding local maxima of the MRF is an instance, and (ordinal) potential games [Fabrikant et al., 2004].

[21] A direct proof of this result follows from [Papadimitriou et al., 1990], and in particular, the result for Hopfield neural networks [Hopfield, 1982]. A Hopfield neural network can be seen as an MRF, and more specifically, and Ising model, when the weights on the edges are symmetric. Similarly, any Hopfield neural network can be seen as a polymatrix game [Miller and Zucker, 1992]; when the weights are symmetric the network can be seen as a potential game (in particular, it is an instance of a *party affiliation game* [Fabrikant et al., 2004]). Indeed, a stable configuration in an arbitrary Hopfield neural network is equivalent to a PNE of a corresponding polymatrix game. (See Papadimitriou et al. [1990] and Miller and Zucker [1992] for the relevant references.)

[22] Note that algorithms for PNE can in principle find critical points of P. In either case, those are the same guarantee that algorithms such as the max-product version of belief propagation (BP) provide, in general.

In addition, there have been several results stating that PNE are unlikely to exist in many graphs, and that, when they do exists, they are not that many [Daskalakis et al., 2007]. [23] MRF-induced games would in that sense represent a very rich class of *non-randomly generated* graphical games for which the results above might not hold.

4.2 Joint Mixed Strategy Equilibria and Belief Inference

Going beyond PNE and MAP estimation, this subsection begins to establish a stronger, and potentially more useful connection between probabilistic inference and equilibria in games.

Let S be a subset of the players (i.e., nodes in the graph) and denote by $Q_S(x_S) \equiv \sum_{x_{V-S}} Q(x)$ the (marginal) probability distribution of Q over possible joint actions of players in S. Consider the condition for correlated equilibria (CE), which for the MRF-induced game we can express as

$$\sum_{x_{\mathcal{N}(i)}} Q_{N(i)}(x_i, x_{\mathcal{N}(i)}) \sum_{C \in \mathcal{C}_i} \phi_C(x_i, x_{C-\{i\}}) \geq \sum_{x_{\mathcal{N}(i)}} Q_{N(i)}(x_i, x_{\mathcal{N}(i)}) \sum_{C \in \mathcal{C}_i} \phi_C(x'_i, x_{C-\{i\}}).$$

Commuting the sums and simplifying we get the following equivalent condition:

$$\sum_{C \in \mathcal{C}_i} \sum_{x_{C-\{i\}}} Q(x_i, x_{C-\{i\}}) \phi_C(x_i, x_{C-\{i\}}) \geq \sum_{C \in \mathcal{C}_i} \sum_{x_{C-\{i\}}} Q(x_i, x_{C-\{i\}}) \phi_C(x'_i, x_{C-\{i\}}). \tag{2}$$

This simplification is important because it highlights that, modulo expected payoff equivalence, we only need distributions over the original cliques, *not* the induced neighbohoods/Markov blankets, to represent CE in this class of games (in contrast to Kakade et al. [2003]); thus, we are able to maintain the size of the representation of the CE to be the same as that of the game. This expression will also be useful in Section 5, which presents simple linear programs to compute CE in arbitrary GMGs.

As an alternative, we can use the fact that the MRF-induced game is a potential game and, via some definitions and algebraic manipulation, get the following sequence of equivalent conditions, which hold for all i, x_i and x'_i.

$$\begin{aligned}
&\textstyle\sum_{x_{-i}} Q(x_i, x_{-i}) \left(M_i(x_i, x_{-i}) - M_i(x'_i, x_{-i})\right) &&\geq 0 \\
&\textstyle\sum_{x_{-i}} Q(x_i, x_{-i}) \left(\Psi(x_i, x_{-i}) - \Psi(x'_i, x_{-i})\right) &&\geq 0 \\
&\textstyle\sum_{x_{-i}} Q(x_i, x_{-i}) \left(\ln P(x_i, x_{-i}) - \ln P(x'_i, x_{-i})\right) &&\geq 0
\end{aligned}$$

Rewriting the last expression, we get the following equivalent condition: for all i, x_i and x'_i,

$$\textstyle\sum_{x_{-i}} Q(x_i, x_{-i})[-\ln P(x_i, x_{-i})] \leq \sum_{x_{-i}} Q(x_i, x_{-i})[-\ln P(x'_i, x_{-i})]\,. \tag{3}$$

The following are some additional remarks on the implications of the last condition.

[23] In particular, the number of PNE has a Poisson distribution with parameter 1.

Remark 6. First, it is useful to introduce the following notation. For any distribution Q', let $H(Q', P) \equiv \sum_x Q'(x)[-\log_2 P(x)]$ be the *cross entropy* between probability distributions Q' and P, with respect to P. [24] Denote by $Q_{-i}(x_{-i}) \equiv \sum_{x_i} Q(x_i, x_{-i})$ the (joint) marginal distribution of play over the joint-actions of all players *except* player i. Denote by $Q'_i Q_{-i}$ the joint distribution defined as $(Q'_i Q_{-i})(x) \equiv Q'_i(x_i) Q_{-i}(x_{-i})$ for all x.

Then, condition 3 implies the following sequence of conditions, which hold for all i.

$$\begin{aligned}
\textstyle\sum_x Q(x)[-\ln P(x)] &\leq \textstyle\sum_{x_{-i}} Q_{-i}(x_{-i})[-\ln P(x'_i, x_{-i})] \text{ for all } x'_i. \\
H(Q, P) &\leq \min_{x'_i} \textstyle\sum_{x_{-i}} Q_{-i}(x_{-i})[-\log_2 P(x'_i, x_{-i})] \\
&= \min_{Q'_i} \textstyle\sum_x Q'_i(x_i) Q_{-i}(x_{-i})[-\log_2 P(x_i, x_{-i})] \\
&= \min_{Q'_i} H(Q'_i Q_{-i}, P)
\end{aligned}$$

Hence, *any CE of the MRF-induced game is a kind of approximate local optimum (or critical point) of an approximation of the MRF based on a special type of cross entropy minimization.*

The following property summarizes this remark.

Property 4. For any MRF P, any correlated equilibria Q of the game induced by P satisfies $H(Q, P) \leq \min_{Q'_i} H(Q'_i Q_{-i}, P)$ for all i.

Remark 7. Let us introduce some additional notation. For any joint distribution of play Q', let $H(Q') \equiv \sum_x Q'(x)[-\log_2 Q'(x)]$ be its entropy. Similarly, for any player i, for any marginal/individual distribution of play Q'_i, let $H(Q'_i) \equiv \sum_{x_i} Q'_i(x_i)[-\log_2 Q'_i(x_i)]$ be its entropy. For any distributions Q', let $\mathrm{KL}(Q' \parallel P) \equiv \sum_x Q'(x) \log_2(Q'(x)/P(x)) = H(Q', P) - H(Q')$ be the *Kullback-Leibler divergence* between Q' and P, with respect to Q'. Denote by $H(Q_{i|-i}) \equiv \sum_{x_i, x_{-i}} Q(x_i, x_{-i}) \log_2(Q(x_i, x_{-i})/Q_{-i}(x_i)) = H(Q_{-i}) - H(Q)$ the conditional entropy of the individual play of player i given the joint play of all the players except i, with respect to Q.

Then, we can express the condition 3 as the following equivalent conditions, which hold for all i.

$$\begin{aligned}
\mathrm{KL}(Q \parallel P) + H(Q) &\leq \min_{Q'_i} \mathrm{KL}(Q'_i Q_{-i} \parallel P) + H(Q'_i Q_{-i}) \ . \\
\mathrm{KL}(Q \parallel P) + H(Q_{i|-i}) &\leq \min_{Q'_i} \mathrm{KL}(Q'_i Q_{-i} \parallel P) + H(Q'_i) \ .
\end{aligned}$$

Hence, *any CE of a MRF-induced game is a kind of approximate local optimum (or critical point) of a special kind of variational approximation of the MRF.* The following property summarizes this remark.

Property 5. For any MRF P, any correlated equilibria Q of the game induced by P satisfies $\mathrm{KL}(Q \parallel P) \leq \left[\min_{Q'_i} \mathrm{KL}(Q'_i Q_{-i} \parallel P) + H(Q'_i)\right] - H(Q_{i|-i})$, for all i.

Note that the last property implies that the approximation Q satisfies the local condition $\mathrm{KL}(Q \parallel P) \leq \min_{Q'_i} \mathrm{KL}(Q'_i Q_{-i} \parallel P) + \log_2 |\Omega_i|$ for all i.

Before continuing exploring connections to CE, it is instructive to first consider mixed-strategy NE.

[24] That is, (a lower bound on) the average number of bits required to transmit "messages/events" generated according to Q but encoded using a scheme based on P.

Mixed Nash Equilibria and Mean-Field Approximations. In the special case of NE, the joint mixed strategy $Q(x) = \prod_i Q_i(x_i)$ is a product distribution. Denote by $Q^{\times}_{-i}(x_{-i}) \equiv \prod_{j \neq i} Q_j(x_j) = \sum_{x_i} Q(x)$ the (marginal) joint action of play over all the players except i, and denote by $(Q'_i Q^{\times}_{-i})$ the probability distribution defined such that the probability of x is $(Q'_i Q^{\times}_{-i})(x) \equiv Q'_i(x_i) Q^{\times}_{-i}(x_{-i})$.

In this special case, the equilibrium conditions imply the following conditions, which hold for all i: for all x_i such that $Q_i(x_i) > 0$,

$$\sum_{x_{-i}} Q_i(x_i) Q^{\times}_{-i}(x_{-i})[-\ln P(x_i, x_{-i})] = \min_{x'_i} \sum_{x_{-i}} Q_i(x_i) Q^{\times}_{-i}(x_{-i})[-\ln P(x'_i, x_{-i})] \; .$$

This implies that

$$\sum_{x_i \text{ s.t. } Q_i(x_i)>0} \sum_{x_{-i}} Q_i(x_i) Q^{\times}_{-i}(x_{-i})[-\ln P(x_i, x_{-i})] = \left(\sum_{x_i \text{ s.t. } Q_i(x_i)>0} Q_i(x_i) \right) \min_{x'_i} \sum_{x_{-i}} Q^{\times}_{-i}(x_{-i})[-\ln P(x'_i, x_{-i})] \; .$$

The last condition is equivalent to

$$\sum_{x_i} \sum_{x_{-i}} Q_i(x_i) Q^{\times}_{-i}(x_{-i})[-\ln P(x_i, x_{-i})] = \min_{x'_i} \sum_{x_{-i}} Q^{\times}_{-i}(x_{-i})[-\ln P(x'_i, x_{-i})] \; ,$$

which, in turn, we can express as $H(Q, P) = \min_{Q'_i} H(Q' Q^{\times}, P)$. The last expression is also equivalent to

$$\mathrm{KL}(Q \parallel P) + H(Q_i) = \min_{Q'_i} \mathrm{KL}(Q'_i Q^{\times}_{-i} \parallel P) + H(Q'_i) \; .$$

Hence, *a NE Q of the game is* almost *a locally optimal mean-field approximation, except for the extra entropic term.* In summary, for NE we have the following tighter condition than for arbitrary CE.

Property 6. For any MRF P, any Nash equilibria Q of the game induced by P satisfies $\mathrm{KL}(Q \parallel P) = \left[\min_{Q'_i} \mathrm{KL}(Q'_i Q^{\times}_{-i} \parallel P) + H(Q'_i)\right] - H(Q_i)$, for all i.

Note that the last property implies that the mean-field approximation Q satisfies the local condition $\mathrm{KL}(Q \parallel P) \leq \min_{Q'_i} \mathrm{KL}(Q'_i Q^{\times}_{-i} \parallel P) + \log_2 |\Omega_i|$ for all i.

One possible way to address the issue of the extra entropic term is to consider instead the *MRF-induced infinite game*, where each player i has the (continuous) utility function [25]

$$\widetilde{M}'_i(Q_i, Q_{\mathcal{N}(i)}) \equiv \sum_{x_i} \sum_{x_{\mathcal{N}(i)}} \left[Q_i(x_i) \prod_{j \in \mathcal{N}(i)} Q_j(x_j) \right] M'_i(x_i, x_{\mathcal{N}(i)}) + H(Q_i)$$

[25] In an *infinite game* the sets of actions or pure strategies are uncountable. Existence of equilibria holds under reasonable conditions (i.e., each set of actions is nonempty compact convex subsets of Euclidean space, and each player utility is continuous and quasi-concave in the player's action), all of which are satisfied by the MRF-induced infinite game considered here. (See Fudenberg and Tirole [1991] for more information.)

and wants to maximize over its mixed-strategy Q_i given the other player mixed-strategies Q_j for all $j \neq i$.

Property 7. The MRF-induced infinite game defined above is an infinite Gibbs potential game with the same graph G and the following potential over the set of individual (product) mixed strategies

$$\Psi(Q) = \sum_{C \in \mathcal{C}} \sum_{x_C} \left[\prod_{j \in C} Q_j(x_j) \right] \phi_C(x_C) + H(Q) = -KL(Q \parallel P) + Z$$

where Z is the normalizing constant for P. From this we can derive that the individual player mixed-strategies $\{Q_i\}$ are a "pure strategy" equilibrium of the infinite game if and only if

$$\mathrm{KL}(Q \parallel P) = \min_{Q'_i} \mathrm{KL}(Q'_i Q^{\times}_{-i} \parallel P).$$

Or, in other words, *if Q is a pure strategy Nash equilibrium of the infinite game, then Q is also a local optimum (or critical point) of the mean-field approximation of P.*

Remark 8. The local payoff function defined above for the infinite game also has connections to the game theory literature on *learning in games* [Fudenberg and Levine, 1999]. This area studies properties of processes by which players "learn" how to play in (usually repeated) games; specially properties related to the existence of convergence of the learning (or playing) dynamics to equilibria. In particular, the local payoff function is similar to that used by *logistic fictitious play*, a special version of a "learning" process called *smooth fictitious play*. The difference is that the last entropy term involving the individual player's mixed strategy has a regularization-type factor $\lambda > 0$ such that players play strict best-response as $\lambda \to 0$. In addition, logistic fictitious play is an instance of a learning process that, if followed by a player, achieves so called approximate *universal consistency* (i.e., roughly, in the limit of infinite play, the average of the payoffs obtained by the player will be close to the best obtained overall during repeated play, *regardless of how the other players behave*), for appropriate values of λ depending on the desired approximation level.

Indeed, it is not hard to see that in fact the best-response mixed-strategy Q_i of player i to the mixed strategies $Q_{\mathcal{N}(i)}$ of their neighbors is

$$\begin{aligned} Q_i(x_i) &\propto \exp(\textstyle\sum_{x_{\mathcal{N}(i)}} \left[\prod_{j \in \mathcal{N}(i)} Q_j(x_j) \right] M'_i(x_i, x_{\mathcal{N}(i)})) \\ &= \exp(\textstyle\sum_{C \in \mathcal{C}_i, C \neq \{i\}} \left[\prod_{j \in C - \{i\}} Q_j(x_j) \right] \phi_C(x_i, x_{C-\{i\}})) \ . \end{aligned}$$

Hence, running *sequential* best-response dynamics in the MRF-induced infinite game is equivalent to finding a variational mean-field approximations via recursive updating of the first derivative conditions. [26] The process will then be equivalent to minimizing the function $F(Q) \equiv \mathrm{KL}(Q \parallel P)$ by axis-parallel updates. The resulting sequence of distributions/mixed-strategies monotonically increase the value of F and is guaranteed to converge to a local

[26] In particular, the process is called a *Cournot adjustment with lock-in* in the literature on learning in games [Fudenberg and Levine, 1999].

optimum (or a critical point) of F. Hence, the corresponding learning process is guaranteed to converge to a pure strategy Nash equilibrium of the infinite game, which are in turn approximate mixed-strategy Nash equilibria of the original game. But this is not surprising in retrospect, given the last property (Property 7). That property essentially states a broader property of *all* potential games: they are isomorphic to so called *games with identical interests* [Monderer and Shapley, 1996], which are games where every player has exactly the same payoff function.

Remark 9. The previous discussion suggests that we could use appropriately modified version of algorithms for mixed-strategy NE, such as **NashProp** [Ortiz and Kearns, 2002], as heuristics to obtain a mean-field approximation of the true marginals.

Going in the opposite direction, the discussion above also suggests that, by treating any (graphical) potential game as an MRF, for any fixed $\lambda > 0$, logistic fictitious play in any potential game converges to an approximate $(\lambda / \min_i |A_i|)$-NE of the potential game. Indeed, there has been recent work in this direction, which explores the connection between learning in games and mean-field approximations in machine learning [Rezek et al., 2008]. That work proposes new algorithms based on fictitious play for simple mean-field approximation applied to statistical (Bayesian) estimation.

The game-induced MRF is a λ-temperature Gibbs measure. As we take $\lambda \to 0$, we get the limiting 0-temperature Gibbs measure which is a uniform probability distribution over the set of global maxima of the potential function of the game, and 0 probability everywhere (i.e., the support of the limiting distribution is the set of values that maximize the potential function). The support of the 0-temperature Gibbs measure is a subset of the pure strategy NE of the potential game. But there might be other equilibria corresponding to local optimum (or critical points) of the potential function.

Are there other connections between the Nash equilibria of the game and the support of the limiting distribution?

Correlated Equilibria and Higher-order Variational Approximations. Property 5 suggests that we can use the CE for the MRF-induced game as a heuristic approximation to higher-order variational approximations. In fact, one would argue that in the context of inference, doing so is more desirable because, in principle, it can lead to better approximations that can capture more aspects of the joint distribution that a simple mean-field approximation would alone. For example, mean-field approximations are likely to be poor if the MRF is multi-modal. Motivated by this fact, Jaakkola and Jordan [1997] suggest using mixture of product distributions to improve the simple variational mean-field approximation.

But, consider the algorithm of Papadimitriou [2005] (see also Papadimitriou and Roughgarden [2008]) which we can use to compute a CE of the MRF-induced game in polynomial time. Such a CE will be, by construction, also a *(polynomially-sized) mixture of product distributions*. Hence, Papadimitriou's algorithm provides a means to obtain a heuristic estimate of a local optimum (or critical point) of such a mixture *in polynomial time*. The result would not be exactly the same as that obtained by Jaakkola and Jordan [1997] in general, because of the extra entropic term. *Can we find alternative versions of the payoff matrices, and/or*

alter Papadimitriou's algorithm, so that the resulting correlated equilibria provides an exact answer to the approximate inference problem that uses mixtures of product distributions? Regardless, at the very least one could use the resulting CE to initialize the technique of Jaakkola and Jordan [1997] without specifying an *a priori* number of mixtures!

Having said that, Papadimitriou's algorithm makes a polynomial number of calls to the ellipsoid algorithm to obtain each of the product distribution whose mixture will form the output CE. It is known that the ellipsoid algorithm is slow in practice. Papadimitriou [2005] and Papadimitriou and Roughgarden [2008] leave open the design of more practical algorithms based on interior-point methods. Section 5 presents simple linear programs for arbitrary GMGs, with polynomial time guarantees for certain subclasses.

Finally, this connection also suggests that we can (in principle) use any learning algorithm that guarantees convergence to the set of CE (as described in the section on preliminaries on game theory where the concept was introduced) as a heuristic for approximate inference. Viewed that way, such learning algorithms would be similar in spirit to stochastic simulation algorithms with a kind of "adaptivity" reminiscent of the work on adaptive importance sampling (see, e.g., [Cheng and Druzdzel, 2000, Ortiz and Kaelbling, 2000, Ortiz, 2002] and the references therein). Establishing a possible stronger connections between learning in games, correlated equilibria and probabilistic inference seems like a promising direction for future research. In fact, as previously mentioned (at the end of Remark 8), there has already been some recent work in this direction, but specifically for Nash equilibria and mean-field approximations [Rezek et al., 2008].

Also left for future work is a comprehensive experimental evaluation of the performance of current algorithms in computational game theory in the context of probabilistic inference.

4.3 Other Previous and Related Work

Earlier work on the so called "relaxation labeling" problem in AI and computer vision [Rosenfeld et al., 1976, Miller and Zucker, 1991] has established connections to polymatrix games [Janovskaja, 1968] (see also Hummel and Zucker [1983], although the connection had yet to be recognized at that time). That work also establishes connections to inference in Hopfield network, dynamical systems and polymatrix games [Miller and Zucker, 1991, Zucker, 2001]. A reduction of MAP to pure NE in what we call here a GMG was introduced by [Yu and Berthod, 1995, Berthod et al., 1996] in the same context; although they concentrate on pairwise potentials, which reduce to polymatrix games in this context. Because, in addition, the ultimate goal in MAP inference is to obtain a *global* optimum configuration, Yu and Berthod [1995] proposed a Metropolis-Hastings-style algorithm in an attempt to avoid local minima. Their algorithm is similar to simulated annealing algorithms used for solving satisfiability problems, and other local methods such as WalkSAT [Selman et al., 1996] (See. e.g., Russell and Norvig [2003] for more information). The algorithm can also be seen as a kind of learning-in-games scheme [Fudenberg and Levine, 1999] based on best-response with random exploration (or "trembling hand" best response). That is, at every round, some best response is taken with some probability, otherwise the previous response is replayed. Zucker [2001] presents a modern account of that work. The connection to potential games, and all its well-

known properties (e.g., convergence of best response dynamics) does not seem to have been recognized within that literature. Also, none of the work makes connection to higher-order (i.e., beyond mean-field) inference approximation techniques and the game-theoretic notion of correlated equilibria.

5 Correlated Equilibria in Graphical Multi-Hypermatrix Games

As stated previously, graphical multi-hypermatrix games (GMGs) have the so called *polynomial expectation property* and thus have a polynomial correlated equilibrium scheme [Papadimitriou, 2005]. This implies that the algorithm of Papadimitriou [2005] will compute a sample correlated equilibria (CE) in those games in polynomial time. However, as also discussed previously, the algorithm is based on the ellipsoid algorithm, thus unlikely to be of immediate practical use. Indeed, Papadimitriou [2005] left open the possible designing of an interior point version of his technique.

This section presents a different approach to the problem of correlated equilibria in graphical games based on ideas that have been found effective in the practical use of graphical models for probabilistic inference and constraint satisfaction problems. In particular, this section presents an extension of the work of Kakade et al. [2003] to GMGs. First, it states an extension of the (compact) representation result for CE, which provides a characterization of the equilibria in terms of direct, structural strategic properties of such games. The result states that the additional succinctness in model representation achieved via the special structure of the GMG can in some well-defined sense translate to the representation of the CE. The result also establishes a direct, strong connection between the *strategic* (i.e., payoff) structure of the game and the *behavioral* structure one can achieve when implementing any CE of the game. Then, the design of a simple linear feasibility system for computing CE in GMGs is presented. The linear system formulation takes advantage of the additional special structure of this class of games. The approach is applicable to any GMG. By using and applying well-known graphical model techniques and results, the approach is shown to produce polynomial-time algorithms based on simple linear programs in cases in which the strategic graph of the game has certain structure.

5.1 A Characterization of Correlated Equilibria via Succinct Representations

Just as for Kakade et al. [2003], the following two definitions of equivalence are useful to the discussion in this section.

Definition 6. *Two (possibly correlated) joint mixed strategies for a graphical multi-hypermatrix game are* local-clique equivalent *with respect to the graph structure of the game if, for every player and every hyperedge in its hypergraph, both joint strategies induce the same marginal distribution over the joint-actions of the players in the hyperedge; formally, if P and Q are arbitrary joint mixed strategies, then for every player i and every hyperedge $C \in \mathcal{C}_i$ in its hypergraph, $P_C(x_C) = Q_C(x_C)$. Similarly, two (possibly correlated) joint mixed strategies are* expected-payoff equivalent *with respect to the game if each (individual) player achieves the*

same (individual) expected payoff with both joint strategies; formally, if for all players i, $M_i(P) = M_i(Q)$.

It is not hard to see that expected-payoff equivalence is weaker than local-clique payoff equivalence in that the latter implies the former, but the opposite does not always hold. However, there are games for which local-clique payoff equivalence is a necessary condition for expected-payoff equivalence [Kakade et al., 2003].

To simplify the rest of the presentation in this section, let us introduce the following notation. Denote by $m \equiv \max_{i \in V} |A_i|$ the maximum number of actions of any player. Let $\mathcal{C} \equiv \cup_i \mathcal{C}_i$ be the set of *unique* hyperedges in the game, $\mathcal{C}^{\max} \equiv \mathcal{C} - \{C \in \mathcal{C} \mid C \subset C', C' - C \neq \emptyset \text{ for some } C' \in \mathcal{C}\}$ the subset of those that are maximal, $l \equiv |\mathcal{C}^{\max}|$ be its cardinality, and $k \equiv \max_{C \in \mathcal{C}^{\max}} |C|$ the size of the largest hyperedge. Let us refer to the hypergraph with hyperedge set $\mathcal{C}^{\max}$ the *strategic hypergraph* of the game and cal its primal graph G' the *prime strategic graph*.

The following is the main representation theorem.

Theorem 4. *For every correlated equilibrium of a graphical multi-hypermatrix game there exists a (possibly different) expected-payoff equivalent correlated equilibrium that is a Gibbs distribution with respect to the strategic hypergraph of the game and whose size, in terms of the number of numeric values needed to represent it, is at most the representation size of the game.*

Proof. The following is the main lemma used in the proof of the theorem.

Lemma 1. *For every correlated equilibrium of a graphical multi-hypermatrix game there exists a (possibly different) local-clique equivalent correlated equilibrium that is a Gibbs distribution with respect to the strategic hypergraph of the game.*

Proof. The proof of this lemma parallels the one for a similar lemma presented in Kakade et al. [2003]. Except for the following observations, mostly on aspects in which the proofs differ, the details are omitted. Let Q be a CE of the game. First, the construction considers the marginals $Q_C(x_C) \equiv \sum_{x_{-C}} Q(x_C, x_{-C})$ over every $C \in \mathcal{C}^{\max}$, instead of using those over (inclusive) neighborhoods $N(i)$. The key observation is that the most entropic joint distribution Q^* that matches those marginals has the desired property. There are a few issues that need addressing. The first issue is the (in)feasibility of the resulting optimization problem. This is resolved by recognizing that Q itself satisfy the feasibility constraints! Another issue is marginals with zero probabilities, but because we have access to the marginals, we do not need to include those marginals when setting up the optimization for Q^*. We can account for them in the corresponding local-clique factor afterwards. Doing this also guarantees that, assuming that Q does not correspond to a pure NE (i.e., $Q(x) = 1$ for some x, a case we can handle easilly), the solution of the MAXENT optimization does not lie at a boundary of the feasible space, so that there exists optimal values for the dual variables that are finite/bounded. The last issue is more benign, having to do with the uniqueness of the solution. But because the problem is feasible, the feasible space convex (in fact, only linear constraints) and the entropy function is strictly concave, the solution is unique. In summary,

solving the resulting MAXENT optimization problem by taking the appropriate derivatives, we obtain that Q^* must have the following form:

$$Q^*(x) \propto \prod_{C \in \mathcal{C}^{\max}} 1[Q_C(x_C) > 0] \ \exp(\lambda^*_C(x_C))$$

where $\lambda^*_C(x_C)$ are the optimal values of the dual variables (i.e., Lagrange multipliers, represented here as functions for consistency of presentation) associated to the constraint $Q^*_C(x_C) = Q_C(x_C)$, for each positive local-clique marginal probability, and the factor $1[Q(x_C) > 0]$ denotes the indicator fuction whose value is 1 if $Q(x_C) > 0$; and 0 otherwise. □

The next corollary follows immediately from the proof of the lemma.

Corollary 2. *In the context of the last lemma, the representation of the corresponding Gibbs distribution requires at most* $\sum_{C \in \mathcal{C}^{\max}} \prod_{j \in C} |A_j| \leq l\, m^k$ *numeric values.*

Combining the lemma and the corollary with the fact that local-clique equivalence implies expected-payoff equivalence and that the CE conditions depend only on the expected payoffs, not the individual mixed strategies, leads to the result. □

Structural Properties of Correlated Equilibria The last theorem implies that, within the expected-payoff equivalence (EPE) class, CE of GMGs induce special behavioral/probabilistic structure which is largely determined by the strategic hypergraph of the game. In particular, if the CE is positive, then, by the Hammersley-Clifford theorem (Theorem 1), it is also an MRF with respect to the prime strategic graph of the game. Hence, we can infer a large variety of global and local conditional independence properties of the behavior of players in any CE, modulo EPE, directly from the graph, without having to look at the specific numerical values of the CE (see, e.g., Dawid and Lauritzen [1993] for more information).

As an example, by the definition of an MRF, it follows that conditioning on the behavior of the neighbors of a player, as well as its neighbor's neighbors, a player's behavior in *any* CE can be made independent of that of the rest of the players in the game during the implementation of the CE, without producing any change on the expected payoff they would have achieved in the original CE. Hence, in the same EPE sense, one can make the player's behavior conditionally depends on only those of players that are at most "two-steps away" (or what one might call its "immediately extended neighborhood"). By the equivalence of local and global conditional properties, we can extend that statement to say that, if A, B and C are three disjoint sets of players such that C *separates* A from B in the prime strategic graph G' (i.e. every path between A and B in G' passes through C), then the behavior of players in A can be made (conditionally) independent of those in B given the behavior of players in C.

If the CE is not positive, while we can still represent it compactly,as shown, in general it may not be an MRF. However, we can still infer some of its local, context-sensitive conditional independence properties when the corresponding conditionals are well defined.

In summary, the theorem makes a strong connection between the *strategic/payoff* structure of the game and the *behavioral/probabilistic* structure of its equilibria.

Implementing Correlated Equilibria The connection to Gibbs is also useful for the implementation of the CE. In some cases, depending on structural properties of the prime strategic graph such as the treewidth, this can be done exactly and efficiently (i.e., in polynomial time in the size of the game), by efficiently constructing an equivalent Bayesian network (BN), as shown shortly in the next subsection. We can also use the Gibbs sampler (discussed in a previous section) to obtain samples from Gibbs distributions. In general, the Gibbs sampler provides samples from the CE only in the limit, i.e., after it achieves convergence to the limiting distribution of the corresponding Markov chain. But there are variants of the Gibbs sampler that provide exact samples efficiently in some special cases, such as certain classes of Ising models (see, e.g., Gibbs [2000, 2004], Shiraishi [2007] and the references therein).

5.2 Tractable Computation of Correlated Equilibria

The algorithmic approach presented here builds on many of the same operations used to obtain exact solutions in probabilistic and constraint graphical models via dynamic programming. (This should not be surprising given that the problem of computing equilibria in games is inherently a (continuous) CSP.) Given that most of the techniques used are by now well-known and standard within the graphical model community, the following presentation will assume familiarity with such techniques. The reader can find more information in standard books on AI, as well as other more specific books on the subject (see, e.g., [Pearl, 1988, Russell and Norvig, 2003, Dechter, 2003], and the references therein).

Using linearity of expectation, we can simplify the CE conditions in GMGs to obtain the following expression: for all $i, x_i, x_i', x_i \neq x_i'$,

$$\sum_{C \in \mathcal{C}_i} \sum_{x_{C-\{i\}}} Q_C(x_i, x_{C-\{i\}}) \phi_C(x_i, x_{C-\{i\}}) \;\geq\; \sum_{C \in \mathcal{C}_i} \sum_{x_{C-\{i\}}} Q_C(x_i, x_{C-\{i\}}) \phi_C(x_i', x_{C-\{i\}}) \quad . \quad (4)$$

Hence, we only need to find local-clique marginals that are *(globally) consistent* with some full joint distributions. The following theorem, which is an adapted version of Theorem 2.6 in Dawid and Lauritzen [1993] tells us that a *sufficient* condition for this is to find *(locally) pairwise-consistent* marginals over the cliques of a decomposable (i.e., triangulated or chordal) graph of the prime strategic graph of the game.

Theorem 5 ([Dawid and Lauritzen, 1993]). *Let G be an undirected* decomposable *graph and $\mathcal{C}$ be the set of maximal cliques in G. For each clique $C \in \mathcal{C}$, let Q_C be some (arbitrary) probability distribution over the variables in C. If the Q_C's are (locally) pairwise-consistent then there exists a (globally) consistent joint probability distribution P with the following properties.*

1. *P is the unique Markov distribution with respect to G such that, for all $C \in \mathcal{C}$, each marginal $P_C(x_C) \equiv \sum_{x_{-C}} P(x_C, x_{-C})$ of P over the variables in C equals $Q_C(x_C)$.*
2. *Let n be the number of variables, $w \equiv \max_{C \in \mathcal{C}} |C|$ be the size of the largest clique in $\mathcal{C}$, m_i be the size of the set of possible values of the random variable indexed by i, and denote by $m \equiv \max_i m_i$ the largest such size. We can construct an equivalent Bayesian*

network for P whose largest conditional probability table (CPT) is of size at most $s \equiv \max_{C \in \mathcal{C}} \prod_{i \in C} m_i \leq m^w$, in time $O(n\, s) = O(n\, m^w)$.

The proof of the theorem is in fact based on the construction of the stated Bayesian network (BN), although there it is referred to as a "directed decomposable model" which is well-known to be equivalent. [27] Dawid and Lauritzen [1993] describes the process in the appendix. For an alternative description, see Pearl [1988], Section 3.3. Ortiz [2002] also provides a brief description of this process (Section 3.8.5, page 127).

Hence, if the prime strategic graph G' of the game is already decomposable then we are in a sense done: We just introduce a set of variables that would correspond to marginals over each maximal clique in G' and constraints to make them pairwise consistent. Of course, we would also add any necessary constraints to force the marginals for any clique C' in the prime strategic graph to agree with those for any clique C in the strategic hypergraph that is a strict subset.

If, on the other hand, G' is not decomposable, it is well-known that we can make it so by a process of *triangulation*, in which we systematically add extra edges to the graph until it becomes decomposable. Let t be the treewidth of G'. It is well-known that (1) if t is known, then we can compute a decomposable graph for G' whose largest clique has size $t + 1$ in linear time; and (2) if t is unknown, we can find a G' whose largest clique has size $O(t)$ in time $2^{O(t)} n \log(n)$; or one whose largest clique size is $O(t \log t)$ in polynomial time (see, e.g., Becker and Geiger [2001], Reed [1992], and the references therein; see also Bodlaender [2005] for more information on treewidth, including state-of-the-art algorithms and alternative running time bounds).

Thus, let us assume throughout the remainder of the section that G' is a decomposable graph and let $\mathcal{C}'$ be its clique set. (Note that by construction, for every $C \in \mathcal{C}^{\max}$, there must exist a $C' \in \mathcal{C}'$ such that $C \subset C'$.)

We can then set up the following linear feasibility system using information about the decomposable graph G'. [28]

- **Variables:** $Q_C(x_C) \geq 0$ for every clique $C \in \mathcal{C}^{\max} \cup \mathcal{C}'$ and every $x_C \in A_C$
- **Constraints:**
  - *CE conditions:* for all $i \in V$, $x_i, x_i' \in A_i$, $x_i \neq x_i'$,

$$\sum_{C \in \mathcal{C}_i} \sum_{x_{C-\{i\}}} Q_C(x_i, x_{C-\{i\}}) \left(\phi_C(x_i', x_{C-\{i\}}) - \phi_C(x_i, x_{C-\{i\}}) \right) \leq 0$$

[27] This result seems to have been rediscovered by Papadimitriou and Roughgarden [2008] (see Theorem 6.4).

[28] Note that we can reduce the number of variables by considering only cliques in $\mathcal{C}'$ if we embed the submarginal condition within the CE conditions. However, that would complicate the expressions for the CE conditions and make the presentation of those expressions more complex.

- *Pairwise-consistency conditions:* for all $C, C' \in \mathcal{C}'$ such that $S \equiv C \cap C' \neq \emptyset$, and for all $x_S \in A_S$, [29]
$$\sum_{x_{C-C'}} Q_C(x_S, x_{C-S}) = \sum_{x_{C'-C}} Q_{C'}(x_S, x_{C'-S})$$
- *Normalization conditions:* for all $C \in \mathcal{C}'$,
$$\sum_{x_C} Q_C(x_C) = 1$$
- *Submarginal conditions:* for all $C \in \mathcal{C}^{\max}$ such that there exists $C' \in \mathcal{C}'$ satisfying $C \subset C'$ and $C' - C \neq \emptyset$, we need that for all $x_C \in A_C$,
$$\sum_{x_{C'-C}} Q_{C'}(x_C, x_{C'-C}) = Q_C(x_C)$$

In the following discussion, let us denote by $w \equiv \max_{C \in \mathcal{C}'} |C|$ and $c \equiv |\mathcal{C}'|$ the size of the largest clique and the number of cliques in G', respectively.

The number of variables in the linear system is $\sum_{C \in \mathcal{C}^{\max} \cup \mathcal{C}'} \prod_{i \in C} |A_i| \leq (l + c)\, m^w \leq 2l\, m^w$. [30] The number of CE constraints is $\sum_{i \in V} |A_i|(|A_i| - 1) \leq n\, m^2$ and each has size at most $l\, m^{k-1}$. The number of pairwise-consistency constraints is $\sum_{C, C' \in \mathcal{C}', C \cap C' \neq \emptyset} \prod_{i \in C - C'} |A_i| \leq c^2\, m^w$ and each has size at most $2w$. The number of normalization constraints is c and each has size at most m^w. The number of submarginal constraints is $\sum_{C \in \mathcal{C}, C \subset C', C' - C \neq \emptyset, \text{ for some } C' \in \mathcal{C}'} \prod_{i \in C' - C} |A_i| \leq l\, m^k$. The total size of the submarginal constraints is at most $l\, m^w$, and takes time at most $O(l\, c)$ to construct. Hence, the total number of constraints is at most $n\, m^2 + c^2 m^w + c + l\, m^k \leq (n + 2l^2) m^w$. The size of the system is at most $n\, l\, m^{k+2} + 2c^2\, m^{w+1} + c\, m^w + l\, m^w \leq l\, (n + 2(l + 1))\, m^{w+2}$.

The following lemma summarizes the last paragraph.

Lemma 2. *The linear system presented above has at most $2l^2\, m^w$ variables and $(n + 2l^2)\, m^w$ constraints. Furthermore, the system has size at most $l\, (n + 2(l + 1))\, m^{w+2}$ and takes at most the same time to construct.*

We can apply Theorem 5 with any set of (pairwise-consistent) marginals $\{Q^*_C\}$ over the cliques of the decomposable graph G' that are feasible solutions to the linear system above to construct a Bayesian network Q^*. Thus, we obtain the following lemma.

Lemma 3. *Consider any graphical multi-hypermatrix game and assume a decomposable graph for its prime strategic graph is given. Any feasible solution of the linear system constructed above for the game can be used to construct a Bayesian network for a correlated equilibria of the game. Conversely, for any correlated equilibria of the game, its marginals over the cliques of the decomposable graph are feasible solutions of the linear system.*

[29] These conditions do not need to be over all non-empty intersections of cliques in the prime strategic graph, but only over the separators of a junction tree for (a decomposable version of) the prime strategic graph G' of the game. The junction tree can be computed efficiently via a maximum spanning tree computation on a weighted graph constructed from G'.

[30] Actually, if we embed the submarginal condition within the CE conditions, then the number of variables is $\sum_{C \in \mathcal{C}'} \prod_{i \in C} |A_i| \leq c\, m^w$.

We can then use a linear program with the linear feasibility system above to compute *any* linear function over the variables of the system (i.e., the marginals over the cliques of the decomposable graph for the prime strategic graph of the game). More generally, we can also apply other non-linear optimization technique to compute more complex functions (i.e., convex, submodular, etc.) over the variables.

The following theorem is the main computational result of the paper, and follows from the discussion above.

Theorem 6. *Consider any n-player graphical multi-hypermatrix game with at most m actions per player. Let l be the number of hyperedges in the strategic hypergraph of the game and t the treewidth of its prime strategic graph. There exists an algorithm (based on linear programming) to compute a Bayesian-network correlated equilibria of the game of size at most $n\,m^{O(t)}$ in time polynomial in n, l, and $m^{O(t)}$.*

Let us say that a correlated equilibria Q is *efficiently implementable* if we can obtain random samples from Q in polynomial time in the size of the game. [31] Because it is well-known that we can draw samples from a Bayesian Network (BN) in time $O(n\,m)$, any CE represented as a BN is efficiently implementable.

The following is an immediate corollary of the last theorem.

Corollary 3. *There exists a polynomial time algorithm to compute an efficiently implementable correlated equilibria for any n-player m-action graphical multi-hypermatrix game with prime strategic graph of $O(k + \log(n\,l)/\log(m))$ treewidth, where l and k are the maximum number of local-cliques/hyperedges and the size of the largest local-clique/hyperedge of any player in the game, respectively.*

Note that the approach taken here applies to *any* GMG. The running time however is exponential in the treewidth of the prime strategic graph, thus it is mostly useful when the treewidth is bounded or does not grow too fast, as stated in the last corollary. The same is true for many exact algorithms for probabilistic inference and constraint satisfaction problems in graphical models with arbitrary graphs. Yet, some of those algorithms have still been found effective in practice. It is left open whether one can obtain a polynomial-time guarantee for arbitrary GMGs by adapting the approach presented here. [32]

The general approach does suggest possible heuristics. For instance, if the prime strategic graph is not decomposable, and any attempt to find a small decomposable graph for it with a reasonably sized cliques fails, we can just set up the system with a partial set of consistency constraints. [33] We would then obtain a set of local marginals for which, unfortunately, we had no guarantee of global consistency. More problematic is that it is unclear how one would test

[31] This definition assumes access to an efficient random number generator that outputs samples uniformly at random from the interval $[0, 1]$ in the real line in polynomial time.

[32] Clearly, we would have to sacrifice the generality in that it is unlikely we will be able to compute CE with certain properties efficiently, in light of the hardness result of Papadimitriou and Roughgarden [2008]. But perhaps there is a sequence of similar linear programs that one could construct online to compute *some* CE, in the same spirit as the ellipsoid-based algorithm.

[33] This is similar to the idea of using marginal-polytope relaxations for MAP estimation in MRFs (see,e.g., Wainwright and Jordan [2008] for more information).

for such consistency in an efficient way. This problem seems similar to the problem of testing whether a set of arbitrary conditional distributions for an MRF are consistent, a requirement for the convergence of the Gibbs sampler. If we could perform this test reasonably fast, then we can keep adding increasingly more complex consistency constraints to the linear program until we find a set of consistent marginals. Unfortunately, the problem of testing consistent marginals is NP-complete [Sontag, 2007]. The general idea is similar to a recent approach to the problem of MAP estimation in MRFs which as been found effective in large instances of the problem [Sontag, 2007, Sontag and Jaakkola, 2007, Sontag et al., 2008b,a]. Designing new algorithms for CE computation using similar ideas is left for future work.

6 Concluding Remarks

Any insight we can obtain for solving Gibbs potential games, a class of games introduced here, can be translated and may be useful in providing a different perspective on inference problems in MRFs. While a lot of the work in computing equilibria in graphical games has borrowed from work in graphical models, researchers from other fields, specially from classical game theory and economics, may (perhaps indirectly) bring alternative, fresh perspectives to inference problems that would have been hard to view through our current understanding and knowledge of graphical model techniques. Going in the other direction, algorithmic and computational advances for inference problems in MRFs would immediately lead to advances in how we solve a variety of classes of game-theoretic graphical models, including the graphical multi-hypermatrix games introduced here.

Bibliography

R.J. Aumann. Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics*, 1, 1974.

R.J. Aumann. Correlated equilibrium as an expression of Bayesian rationality. *Econometrica*, 55, 1987.

Ann Becker and Dan Geiger. A sufficiently fast algorithm for finding close to optimal clique trees. *Artificial Intelligence*, 125(1-2):3 – 17, 2001. ISSN 0004-3702. doi: DOI:10.1016/S0004-3702(00)00075-8. URL `http://www.sciencedirect.com/science/article/B6TYF-423YMT0-2/2/9b8a9d8205c2bab728e7a464a69264ac`.

Marc Berthod, Zoltan Kato, Shan Yu, and Josiane Zerubia. Bayesian image classification using Markov random fields. *Image and Vision Computing*, 14(4):285 – 295, 1996.

Julian Besag. Spatial interaction and the statistical analysis of lattice systems. *Journal of the Royal Statistical Society. Series B (Methodological)*, 36(2):192–236, 1974. ISSN 00359246. URL `http://www.jstor.org/stable/2984812`.

Hans L. Bodlaender. Discovering treewidth. Technical Report UU-CS-2005-018, Institute of Information and Computing Sciences, Utrecht University, Utrecht, the Netherlands, 2005. URL `http://www.cs.uu.nl/research/techreps/repo/CS-2005/2005-018.pdf`.

George Casella and Edward I. George. Explaining the Gibbs sampler. *The American Statistician*, 46(3):167–174, Aug. 1992. doi: 10.2307/2685208. URL `http://dx.doi.org/10.2307/2685208`.

Xi Chen and Xiaotie Deng. Settling the complexity of 2-player Nash-equilibrium. Technical Report 140, Electronic Colloquium on Computational Complexity (ECCC), 2005a. `http://eccc.hpi-web.de/eccc-reports/2005/TR05-140/index.html`.

Xi Chen and Xiaotie Deng. 3-NASH is PPAD-complete. Technical Report 134, Electronic Colloquium on Computational Complexity (ECCC), 2005b. `http://eccc.hpi-web.de/eccc-reports/2005/TR05-134/index.html`.

Xi Chen and Xiaotie Deng. Settling the complexity of two-player Nash equilibrium. In *Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS'06)*, 2006.

Xi Chen, Xiaotie Deng, and Shang-Hua Teng. Computing Nash equilibria: Approximation and smoothed complexity. In *Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS'06)*, 2006.

Xi Chen, Xiaotie Deng, and Shang-Hua Teng. Settling the complexity of computing two-player Nash equilibria. *J. ACM*, 56(3):1–57, 2009.

Jian Cheng and Marek J. Druzdzel. Ais-bn: An adaptive importance sampling algorithm for evidential reasoning in large Bayesian networks. *Journal of Artificial Intelligence Research*, 13:155–188, 2000.

Vincent Conitzer and Tuomas Sandholm. New complexity results about nash equilibria. *Games and Economic Behavior*, 63(2):621 – 641, 2008. ISSN 0899-8256. doi: DOI: 10.1016/j.geb.2008.02.015. URL `http://www.sciencedirect.com/science/article/`

B6WFW-4SDGR4B-1/2/cf2498492436f6eaa78dd3314a877366. Second World Congress of the Game Theory Society.

Gregory F. Cooper. The computational complexity of probabilistic inference using Bayesian belief networks (research note). *Artif. Intell.*, 42(2-3):393–405, 1990.

Constantinos Daskalakis and Christos H. Papadimitriou. Computing pure Nash equilibria in graphical games via markov random fields. In *EC '06: Proceedings of the 7th ACM conference on Electronic commerce*, pages 91–99, New York, NY, USA, 2006. ACM. ISBN 1-59593-236-4. doi: http://doi.acm.org/10.1145/1134707.1134718.

Constantinos Daskalakis, Alexandros G. Dimakis, and Elchanan Mossel. Connectivity and equilibrium in random games, 2007. http://www.citebase.org/abstract?id=oai:arXiv.org:math/0703902.

Constantinos Daskalakis, Paul W. Goldberg, and Christos H. Papadimitriou. The complexity of computing a Nash equilibrium. *SIAM Journal on Computing*, 39(1):195–259, 2009a.

Constantinos Daskalakis, Paul W. Goldberg, and Christos H. Papadimitriou. The complexity of computing a Nash equilibrium. *Commun. ACM*, 52(2):89–97, 2009b.

Konstantinos Daskalakis and Christos H. Papadimitriou. Three-player games are hard. Technical Report 139, Electronic Colloquium on Computational Complexity (ECCC), 2005. http://eccc.hpi-web.de/eccc-reports/2005/TR05-139/index.html.

Konstantinos Daskalakis, Paul W. Goldberg, and Christos H. Papadimitriou. The complexity of computing a Nash equilibrium. Technical Report 115, Electronic Colloquium on Computational Complexity (ECCC), 2005. http://eccc.hpi-web.de/eccc-reports/2005/TR05-115/index.html.

A. P. Dawid and S. L. Lauritzen. Hyper Markov laws in the statistical analysis of decomposable graphical models. *The Annals of Statistics*, 21(3):1271–1317, September 1993.

Rina Dechter. *Constraint Processing*. Morgan Kaufmann, 2003.

Alex Fabrikant, Christos Papadimitriou, and Kunal Talwar. The complexity of pure Nash equilibria. In *STOC '04: Proceedings of the Thirty-Sixth Annual ACM Symposium on Theory of Computing*, pages 604–612, New York, NY, USA, 2004. ACM.

D. Foster and R. Vohra. Calibrated learning and correlated equilibrium. *Games and Economic Behavior*, 1997.

D. Foster and R. Vohra. Regret in the on-line decision problem. *Games and Economic Behavior*, pages 7 – 36, 1999.

D. Fudenberg and D. Levine. *The Theory of Learning in Games*. MIT Press, 1999.

Drew Fudenberg and Jean Tirole. *Game Theory*. The MIT Press, 1991.

Stuart Geman and Donald Geman. Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, PAMI-6(6):721–741, Nov. 1984.

Alison L. Gibbs. Bounding the convergence time of the Gibbs sampler in Bayesian image restoration. *Biometrika*, 87(4):749–766, 2000. ISSN 00063444. URL http://www.jstor.org/stable/2673608.

Alison L. Gibbs. Convergence in the wasserstein metric for markov chain monte carlo algorithms with applications to image restoration. *Stochastic Models*, 20(4):473 – 492, 2004. ISSN 15326349. URL http://search.ebscohost.com.libproxy.cc.stonybrook.edu/login.aspx?direct=true&db=a9h&AN=14541167&site=ehost-live.

I. Gilboa and E. Zemel. Nash and correlated equilibria: some complexity considerations. *Games and Economic Behavior*, 1:80–93, 1989.

Paul W. Goldberg and Christos H. Papadimitriou. Reducibility among equilibrium problems. Technical Report 090, Electronic Colloquium on Computational Complexity (ECCC), 2005. http://eccc.hpi-web.de/eccc-reports/2005/TR05-090/index.html.

Georg Gottlob, Gianluigi Greco, and Francesco Scarcello. Pure Nash equilibria: Hard and easy games. In *TARK '03: Proceedings of the 9th conference on Theoretical aspects of rationality and knowledge*, pages 215–230, New York, NY, USA, 2003. ACM. ISBN 1-58113-731-1. doi: http://doi.acm.org/10.1145/846241.846269.

J.M. Hammersley and P. Clifford. Markov fields on finite graphs and lattices. Unpublished, 1971.

Sergiu Hart and Yishay Mansour. The communication complexity of uncoupled nash equilibrium procedures. In *STOC '07: Proceedings of the thirty-ninth annual ACM symposium on Theory of computing*, pages 345–353, New York, NY, USA, 2007. ACM. ISBN 978-1-59593-631-8. doi: http://doi.acm.org/10.1145/1250790.1250843. URL http://portal.acm.org/ft_gateway.cfm?id=1250843&type=pdf&coll=Portal&dl=GUIDE&CFID=48027515&CFTOKEN=24230463.

Sergiu Hart and Andreu Mas-Colell. A simple adaptive procedure leading to correlated equilibrium. *Econometrica*, 68(5):1127 – 1150, 2000.

Sergiu Hart and Andreu Mas-Colell. Uncoupled dynamics do not lead to Nash equilibrium. *American Economic Review*, 93(5):1830–1836, 2003.

Sergiu Hart and Andreu Mas-Colell. Stochastic uncoupled dynamics and Nash equilibrium: Extended abstract. In *TARK*, pages 52–61, 2005.

J J Hopfield. Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences of the United States of America*, 79(8):2554–2558, 1982. URL http://www.pnas.org/content/79/8/2554.abstract.

Robert A. Hummel and Steven W. Zucker. On the foundations of relaxation labeling processes. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, PAMI-5(3): 267–287, May 1983.

Sorin Istrail. Statistical mechanics, three-dimensionality and NP-completeness: I. universality of intracatability for the partition function of the Ising model across non-planar surfaces (extended abstract). In *STOC '00: Proceedings of the thirty-second annual ACM symposium on Theory of computing*, pages 87–96, New York, NY, USA, 2000. ACM. ISBN 1-58113-184-4. doi: http://doi.acm.org/10.1145/335305.335316. URL http://portal.acm.org/ft_gateway.cfm?id=335316&type=pdf&coll=Portal&dl=GUIDE&CFID=46800608&CFTOKEN=87597059.

Tommi S. Jaakkola. Tutorial on variational approximation methods. In M. Opper and D. Saad, editors, *Advanced Mean Field Methods: Theory and Practice*, pages 129–159. MIT Press, Cambridge, MA, 2000.

Tommi S. Jaakkola and Michael L. Jordan. Improving the mean field approximation via the use of mixture distributions. In *Proceedings of the NATO ASI on Learning in Graphical Models*. Kluwer, 1997. http://people.csail.mit.edu/tommi/papers/mix.ps.

E. B. Janovskaja. Equilibrium situations in multi-matrix games. *Litovsk. Mat. Sb.*, 8:381–384, 1968.

Albert Xin Jiang and Kevin Leyton-Brown. Action-graph games. Technical Report TR-2008-13, University of British Columbia, September 2008. URL `http://www.cs.ubc.ca/cgi-bin/tr/2008/TR-2008-13`.

David S. Johnson, Christos H. Papadimitriou, and Mihalis Yannakakis. How easy is local search? *Journal of Computer and System Sciences*, 37(1):79 – 100, 1988.

Michael I. Jordan, Zoubin Ghahramani, Tommi S. Jaakkola, and Lawrence K. Saul. An introduction to variational methods for graphical models. *Mach. Learn.*, 37(2):183–233, 1999.

Sham Kakade, Michael Kearns, John Langford, and Luis Ortiz. Correlated equilibria in graphical games. In *EC '03: Proceedings of the 4th ACM Conference on Electronic Commerce*, pages 42–47, New York, NY, USA, 2003. ACM.

Samuel Karlin. *Mathematical Methods and Theory in Games, Programming, and Economics*. Addison Wesley Publishing Company, 1959.

Michael J. Kearns, Michael L. Littman, and Satinder P. Singh. Graphical models for game theory. In *UAI '01: Proceedings of the 17th Conference in Uncertainty in Artificial Intelligence*, pages 253–260, San Francisco, CA, USA, 2001. Morgan Kaufmann Publishers Inc.

Daphne Koller and Brian Milch. Multi-agent influence diagrams for representing and solving games. *Games and Economic Behavior*, 45(1):181–221, 2003.

F.R. Kschischang, B.J. Frey, and H.-A. Loeliger. Factor graphs and the sum-product algorithm. *Information Theory, IEEE Transactions on*, 47(2):498–519, Feb 2001. ISSN 0018-9448. doi: 10.1109/18.910572.

Pierfrancesco La Mura. Game networks. In *Proceedings of the 16th Annual Conference on Uncertainty in Artificial Intelligence (UAI-00)*, 2000.

Kevin Leyton-Brown and Moshe Tennenholtz. Local-effect games. In *Proceedings of the Eighteenth International Joint Conference on Artificial Intelligence (IJCAI)*, pages 772–777, 2003.

Douglas A. Miller and Steven W. Zucker. Copositive-plus lemke algorithm solves polymatrix games. *Operations Research Letters*, 10(5):285 – 290, 1991.

Douglas A. Miller and Steven W. Zucker. Efficient simplex-like methods for equilibria of nonsymmetric analog networks. *Neural Computation*, 4(2):167–190, 1992. doi: 10.1162/neco.1992.4.2.167. URL `http://www.mitpressjournals.org/doi/abs/10.1162/neco.1992.4.2.167`.

Dov Monderer and Lloyd S. Shapley. Potential games. *Games and Economic Behavior*, 14 (1):124 – 143, 1996.

John Nash. Non-cooperative games. *Annals of Mathematics*, 54:286–295, September 1951.

John F. Nash. Equilibrium points in n-person games. *Proceedings of the National Academy of Sciences of the United States of America*, 36(1):48–49, Jan 1950.

Noam Nisan, Tim Roughgarden, Éva Tardos, and Vijay V. Vazirani, editors. *Algorithmic Game Theory*. Cambridge University Press, 2007.

Luis E. Ortiz. *Selecting Approximately-Optimal Actions in Complex Structured Domains*. PhD thesis, Brown University, May 2002. `http://www.cs.sunysb.edu/~leortiz/phd.pdf`.

Luis E. Ortiz and Leslie Pack Kaelbling. Adaptive importance sampling for estimation in structured domains. In *UAI '00: Proceedings of the 16th Conference on Uncertainty in Artificial Intelligence*, pages 446–454, San Francisco, CA, USA, 2000. Morgan Kaufmann Publishers Inc. ISBN 1-55860-709-9.

Luis E. Ortiz and Michael Kearns. Nash propagation for loopy graphical games. In *Neural Information Processing Systems (NIPS)*, 2002.

C. H. Papadimitriou, A. A. Schäffer, and M. Yannakakis. On the complexity of local search. In *STOC '90: Proceedings of the twenty-second annual ACM symposium on Theory of computing*, pages 438–445, New York, NY, USA, 1990. ACM. ISBN 0-89791-361-2. doi: http://doi.acm.org/10.1145/100216.100274.

Christos H. Papadimitriou. Computing correlated equilibria in multi-player games. In *STOC '05: Proceedings of the Thirty-Seventh Annual ACM Symposium on Theory of cComputing*, pages 49–56, 2005.

Christos H. Papadimitriou. On the complexity of the parity argument and other inefficient proofs of existence. *J. Comput. Syst. Sci.*, 48(3):498–532, 1994. ISSN 0022-0000. doi: http://dx.doi.org/10.1016/S0022-0000(05)80063-7.

Christos H. Papadimitriou and Tim Roughgarden. Computing equilibria in multi-player games. In *SODA '05: Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 82–91, 2005.

Christos H. Papadimitriou and Tim Roughgarden. Computing correlated equilibria in multi-player games. *J. ACM*, 55(3):1–29, 2008. ISSN 0004-5411. doi: http://doi.acm.org/10.1145/1379759.1379762. URL `http://portal.acm.org/ft_gateway.cfm?id=1379762&type=pdf&coll=Portal&dl=GUIDE&CFID=48210635&CFTOKEN=34538586`.

Judea Pearl. *Probabilistic Reasoning in Intelligent Systems. Networks of Plausible Inference*. Morgan Kaufmann, Palo Alto, CA, 1988.

Bruce A. Reed. Finding approximate separators and computing tree width quickly. In *STOC '92: Proceedings of the twenty-fourth annual ACM symposium on Theory of computing*, pages 221–228, New York, NY, USA, 1992. ACM. ISBN 0-89791-511-9. doi: http://doi.acm.org/10.1145/129712.129734.

Iead Rezek, David S. Leslie, Steven Reece, Stephen J. Roberts, Alex Rogers, Rajdeep K. Dash, and Nicholas R. Jennings. On similarities between inference in game theory and machine learning. *J. Artif. Intell. Res. (JAIR)*, 33:259–283, 2008. URL `http://dblp.uni-trier.de/db/journals/jair/jair33.html#RezekLRRRDJ08`.

Yosef Rinott and Marco Scarsini. On the number of pure strategy nash equilibria in random games. *Games and Economic Behavior*, 33(2):274 – 293, 2000.

Azreil Rosenfeld, Robert A. Hummel, and Steven W. Zucker. Scene labeling by relaxation operations. *IEEE Transactions of Systems, Man, and Cybernetics*, SMC-6(6), June 1976.

Stuart J. Russell and Peter Norvig. *Artificial Intelligence: A Modern Approach*. Pearson Education, 2003.

Bart Selman, Henry Kautz, and Bram Cohen. Local search strategies for satisfiability testing. In David S. Johnson and Michael A. Trick, editors, *liques, Coloring, and Satisfiability: Second DIMACS Implementation Challenge, October 11-13, 1993*, volume 26 of *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*. AMS, 1996.

Solomon Eyal Shimony. Finding MAPs for belief networks is NP-hard. *Artificial Intelligence*, 68(2):399 – 410, 1994. ISSN 0004-3702. doi: DOI:10.1016/0004-3702(94)90072-8. URL `http://www.sciencedirect.com/science/article/B6TYF-47YY1PM-5/2/464fa2711fc166135138f15751a06bd2`.

Yuichi Shiraishi. An upper bound on the convergence time of the Gibbs sampler in Ising models. *Journal of Statistical Computation & Simulation*, 77(10):839 – 850, 2007. ISSN 00949655. URL `http://search.ebscohost.com.libproxy.cc.stonybrook.edu/login.aspx?direct=true&db=a9h&AN=26705480&site=ehost-live`.

David Sontag and Tommi Jaakkola. New outer bounds on the marginal polytope. In *Neural Information Processing Systems (NIPS) 21*, Dec. 2007.

David Sontag, Amir Globerson, and Tommi Jaakkola. Clusters and coarse partitions in LP relaxations. In *Neural Information Processing Systems (NIPS) 22*, Dec. 2008a.

David Sontag, Talya Meltzer, Amir Globerson, Tommi Jaakkola, and Yair Weiss. Tightening LP relaxations for MAP using message passing. In *Uncertainty in Artificial Intelligence (UAI) 24*, July 2008b.

David Alexander Sontag. Cutting plane algorithms for variational inference in graphical models. Master's thesis, Massachusetts Institute of Technology, May 2007.

William Stanford. A note on the probability of k pure Nash equilibria in matrix games. *Games and Economic Behavior*, 9(2):238 – 246, 1995.

J. Szép and F. Forgoó. *Introduction to the Theory of Games*. D. Reidel Publishing Company, 1985.

D. Vickrey and D. Koller. Multi-agent algorithms for solving graphical games. In *Proceedings of the Eighteenth National Conference on Artificial Intelligence (AAAI-02)*, pages 345–351, 2002.

John von Neumann and Oskar Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, Princeton, NJ, 1947. Second Edition.

Martin J. Wainwright and Michael I. Jordan. Graphical models, exponential families, and variational inference. *Foundations and Trends®in Machine Learning*, 1(1-2):1–305, 2008.

S. Yu and M. Berthod. A game strategy approach for image labeling. *Computer Vision and Image Understanding*, 61(1):32 – 37, 1995.

Steven W. Zucker. Relaxation labeling: 25 years and sill iterating. In L. S. Davis, editor, *Foundations of Image Understanding*, chapter 10, pages 289–322. Kluwer Academic Publishers, Boston, 2001.