Correlated Equilibria in Graphical Games

Sham Kakade
Gatsby Neuroscience Unit
University College London
sham@gatsby.ucl.ac.uk

Michael Kearns
Computer and Information Science
University of Pennsylvania
mkearns@cis.upenn.edu

John Langford
IBM Research
TJ Watson
jcl@cs.cmu.edu

Luis Ortiz
Computer and Information Science
University of Pennsylvania
leortiz@linc.cis.upenn.edu

ABSTRACT

We examine correlated equilibria in the recently introduced formalism of graphical games, a succinct representation for multiplayer games. We establish a natural and powerful relationship between the graphical structure of a multiplayer game and a certain Markov network representing distributions over joint actions. Our first main result establishes that this Markov network succinctly represents all correlated equilibria of the graphical game up to expected payoff equivalence. Our second main result provides a general algorithm for computing correlated equilibria in a graphical game based on its associated Markov network. For a special class of graphical games that includes trees, this algorithm runs in time polynomial in the graphical game representation (which is polynomial in the number of players and exponential in the graph degree).

Categories and Subject Descriptors

F.2.0 [**Analysis of Algorithms and Problem Complexity**]: General

General Terms

Algorithms, Theory, Economics

Keywords

Game Theory, Correlated Equilibria, Graphical Games, Graphical Models

1. INTRODUCTION

Graphical games are a compact representation of multiplayer games that exploit a graph-theoretic or network structure of strategic interaction among the participants. A number of recent papers have established algorithms for computing Nash equilibria directly on the graph representation, including provably efficient algorithms for computing all approximate Nash equilibria in tree-structured games [Kearns et al.(2001)Kearns, Littman, and Singh] [Littman et al.(2002)Littman, Kearns, and Singh], and convergent heuristics for general graphs that seem to exhibit promising experimental behavior [Ortiz and Kearns(2003)] [Vickrey and Koller(2002)]. Just as graphical models for probabilistic inference (such as Bayesian and Markov networks) have revolutionized the applicability of probabilistic modeling in a wide variety of domains, graphical games are part of an overarching program to develop compact models, and algorithms that exploit them, in order to create a richer computational toolbox for game theory [Koller and Milch()] [La Mura(2000)].

Like much of the history of game theory generally, the study of graphical games so far has been dominated by Nash's classical notion of equilibrium, for which it always suffices to consider product distributions over the players' joint actions. In this paper, we examine *correlated equilibria* [Aumann(1974)] in graphical games, which allow arbitrary joint distributions. Correlated equilibria offer a number of conceptual and computational advantages over Nash equilibria, including the facts that new and sometimes more "fair" payoffs can be achieved, that correlated equilibria can be computed efficiently for games in standard normal form, and that correlated equilibria are the convergence notion for several natural learning algorithms [Foster and Vohra(1999)]. Furthermore, it has been argued that correlated equilibria are the natural equilibrium concept consistent with the Bayesian perspective [Aumann(1987)][Foster and Vohra(1997)]. In this paper, we present a series of fundamental results examining the representational and computational issues that arise when considering correlated equilibria in the compact language of graphical games.

The first issue that arises in this investigation is the problem of *representing* correlated equilibria. Unlike Nash equilibria, even in very simple graphical games there may be correlated equilibria of essentially arbitrary complexity (for instance, any mixture distribution of Nash equilibria is a correlated equilibrium). Since one of our primary goals is to maintain the succinctness of graphical games, some way of addressing this distributional complexity is required. For this we turn to another graphical formalism — namely, undirected graphical models for probabilistic inference, also known as *Markov networks*.

Our main results establish a natural and powerful relationship between a graphical game and a certain associated Markov network. Like the graphical game, the associated Markov network is a graph over the players. While the interactions between vertices in the graphical game are entirely *strategic* and given by local pay-

EC'03, June 9–12, 2003, San Diego, California, USA.

off matrices, the interactions in the associated Markov network are entirely *probabilistic* and given by local potential functions. The graph of the associated Markov network retains the parsimony of the graphical game.

Our first main result shows that the associated Markov network is sufficient for representing *any* correlated equilibria of the graphical game, up to expected payoff equivalence. In other words, the fact that a multiplayer game can be succinctly represented by a graph implies that its entire space of correlated equilibria, up to payoff equivalence, can be represented graphically with comparable succinctness. This basic result establishes a new sense in which graphical games are a powerful formalism, and highlights the natural relationship between computational game theory and modern probabilistic modeling.

Our second main result establishes the *computational* benefits of this relationship. The fact that correlated equilibria are characterized by a set of linear inequalities is not helpful for unrestricted multiplayer games, since in general there are an exponential number of such inequalities. Here again the graphical representations reap benefits. We show that a graphical game gives rise to a small set of linear inequalities (comparable in size to the game representation itself) with a non-empty feasible region that includes all correlated equilibria. In the case that the associated Markov network is *chordal*, which includes graphical games that are trees as a special case, we prove that any point in the feasible region can be efficiently mapped to a correlated equilibrium on the associated Markov network, thus yielding a polynomial time algorithm for computing correlated equilibria in a large class of graphical games. This algorithm also applies generally to non-chordal graphs, but in the worst case may require an exponential increase in the Markov network size.

2. BACKGROUND AND PRELIMINARIES

2.1 Game Theory and Notions of Equilibria

A multiplayer game consists of n players, each with a finite set of *pure strategies* or actions available to them, along with a specification of the *payoffs* to each player. Throughout the paper, we use A_i as a variable representing the chosen action of player i, and a_i as a specific value of A_i. For simplicity we assume a binary action space, so $a_i \in \{0,1\}$. (The generalization of our results to the multi-action setting is straightforward.) The payoffs to player i are given by a table or matrix M_i, indexed by the joint action $\vec{a} \in \{0,1\}^n$. The value $M_i(\vec{a})$, which we assume without loss of generality to lie in the interval $[0,1]$, is the payoff to player i resulting from the joint action $\vec{a}$. Multiplayer games described in this way are referred to as *normal form* games.

A number of different notions of *equilibria* have been proposed for normal form games, including *Nash equilibria* and *correlated equilibria*. We begin with the latter because it is more general and our main interest.

Correlated equilibria [Aumann(1974)] can be viewed as distributions $P(\vec{a})$ over joint actions satisfying a certain conditional expectation property. Let

$$P|a_i \equiv \mathbf{Pr}_{\vec{A}\sim P}[A_1, ..., A_n | A_i = a_i]$$

denote the conditional distribution over actions given the event that $A_i = a_i$. Let $\vec{a}[i:b]$ be the vector $\vec{a}$, but with the ith component fixed to $b \in \{0,1\}$.

DEFINITION 1. *A* correlated equilibrium (CE) *for a normal form game is a distribution $P(\vec{a})$ over actions satisfying*

$$\forall i \in \{1,...,n\}, \forall a_i, a' \in \{0,1\}: \\ \mathbf{E}_{\vec{a}\sim P|a_i}[M_i(\vec{a})] \geq \mathbf{E}_{\vec{a}\sim P|a_i}[M_i(\vec{a}[i:a'])]$$

Intuitively, in a CE the action played by any player is a best response (in the expected payoff sense) to the conditional distribution over the other players given that action, and thus no player has a *unilateral* incentive to deviate from playing their role in the CE. Note that a CE may be an arbitrarily complex joint distribution. In contrast, a *Nash equilibrium* [Nash(1951)] is a special case of CE in which we demand that P be a product distribution ($P(\vec{a}) = \prod_{i=1}^{n} P_i(a_i)$ for some distributions P_i), so every player acts independently of all others.

Nash equilibria have been extensively studied in the game theory literature, including in the context of graphical games (defined shortly). However, as discussed in the Introduction, CE offer a number of interesting conceptual and computational advantages not shared by Nash equilibria. One of the most interesting aspects of CE is that they broaden the set of "rational" solutions for normal form games without the need to address often difficult issues such as stability of coalitions and payoff imputations [Aumann(1987)]. The traffic signal is often cited as an informal everyday example of CE, in which a single bit of shared information allows a fair split of waiting times [Owen(1995)]. In this example, no player stands to gain greater payoff by unilaterally deviating from the correlated play, for instance by "running a light".

2.2 Graphical Games

In a *graphical game* [Kearns et al.(2001)Kearns, Littman, and Singh], each player i is represented by a vertex in an undirected[1] graph G. We use $N(i) \subseteq \{1,\ldots,n\}$ to denote the *neighborhood* of player i in G — that is, those vertices j such that the edge (i,j) appears in G. By convention $N(i)$ always includes i itself as well. If $\vec{a}$ is a joint action, we use $\vec{a}^{\,i}$ to denote the induced vector of actions just on the players in $N(i)$.

DEFINITION 2. *A* graphical game *is a pair $(G, \mathcal{M})$, where G is an undirected graph over the vertices $\{1,\ldots,n\}$, and $\mathcal{M}$ is a set of* n local game matrices. *For any joint action $\vec{a}$, the local game matrix $M_i \in \mathcal{M}$ specifies the payoff $M_i(\vec{a}^{\,i})$ for player i, which depends only on the actions taken by the players in $N(i)$.*

Graphical games are a potentially more compact way of representing games than standard normal form. In particular, rather than requiring size exponential in the number of players n, a graphical game requires size exponential only in the size d of the largest local neighborhood. Thus if $d << n$, the graphical representation is exponentially smaller than the normal form. Note that we can represent any normal form game as a graphical game by letting G be the complete graph, but the representation is most useful when a considerably sparser graph can be found.

3. CORRELATED EQUILIBRIA IN GRAPHICAL GAMES: REPRESENTATION

Specifying a CE as a table of joint probabilities over the binary actions requires $2^n - 1$ parameters. One might hope that if we have a compact graphical game then we could also concisely represent

[1]Undirected graphs are used for simplicity. A directed graphical game where each edge denotes "i affects the payoff of j" is more complicated but may result in a sparser graph and further representational savings [Vickrey and Koller(2002)]. The results presented in this paper have natural extentions to directed graphical games.

the CE. Unfortunately, arbitrary high-order correlations might exist in a CE, even for a concisely represented game [2].

However, we shall show that there is a natural *subclass* of the set of all CE, based on *expected payoff equivalence*, whose representation size is linearly related to the representation size of the graphical game. Note that merely finding distributions giving the same payoffs as the CE is not especially interesting *unless* those distributions are themselves CE. In other words, we do not want to only compactly *describe* the payoffs achievable under CE; we want to be able to *prescribe* CE strategies (joint distributions) yielding these payoffs. Our primary tool for accomplishing this goal will be the notion of local neighborhood equivalence, or the preservation of local marginal distributions. Below we establish that local neighborhood equivalence both implies payoff equivalence and preserves CE. In the following subsection, we describe how to represent this natural subclass in a certain *Markov network*, where the structure of the Markov network is closely related to the structure of the graphical game.

3.1 Expected Payoff Equivalence and Local Neighborhood Equivalence

DEFINITION 3. *Two distributions P and Q over joint actions $\vec{a}$ are* expected payoff equivalent*, denoted $P \equiv_{EP} Q$, if P and Q yield the same expected payoff vector: for each i, $\mathbf{E}_{\vec{a}\sim P}[M_i(\vec{a}^{\,i})] = \mathbf{E}_{\vec{a}\sim Q}[M_i(\vec{a}^{\,i})]$.*

Payoff equivalence of two distributions is, in general, dependent upon the reward matrices of a graphical game. Let us consider the following (more stringent) equivalence notion, which is based only on the graph G of a game.

DEFINITION 4. *For a graph G, two distributions P and Q over joint actions $\vec{a}$ are* local neighborhood equivalent *with respect to G, denoted $P \equiv_{LN} Q$, if for all players i, and for all settings $\vec{a}^{\,i}$ of $N(i)$, $P(\vec{a}^{\,i}) = Q(\vec{a}^{\,i})$.*

In other words, the marginal distributions over all local neighborhoods defined by G are identical. Since the graph is always clear from context, we shall just write $P \equiv_{\text{LN}} Q$. The following lemma establishes that local neighborhood equivalence is indeed a more stringent notion of equivalence than expected payoff.

LEMMA 1. *For all graphs G, for all joint distributions P and Q on actions, and for all graphical games with graph G, if $P \equiv_{LN} Q$ then $P \equiv_{EP} Q$. Furthermore, there exists payoff matrices $\mathcal{M}$ such that for the graphical game $(G, \mathcal{M})$, if $P \not\equiv_{LN} Q$ then $P \not\equiv_{EP} Q$.*

PROOF. The first statement follows from the observation that the expected payoff to player i depends only on the marginal distribution of actions in $N(i)$. To prove the second statement, if $P \not\equiv_{\text{LN}} Q$, then there must exist a player i and a joint action $\vec{a}^{\,i}$ for its local neighborhood which has a different probability under P and Q. Simply set $M_i(\vec{a}^{\,i}) = 1$ and $M_i = 0$ elsewhere. Then i has a different payoff under P and Q, and so $P \not\equiv_{\text{EP}} Q$. □

Essentially, local neighborhood equivalence implies payoff equivalence, but the converse is not true in general (though there exists some payoff matrices where the converse is correct).

Let $\mathcal{CE}(G, \mathcal{M})$ denote the set of of all correlated equilibria for a graphical game $(G, \mathcal{M})$. We now establish that local neighborhood equivalence also preserves CE. It is important to note that this result does *not* hold for expected payoff equivalence.

[2]For example, the CE of a game always include all mixture distributions of Nash equilibria, so any game with an exponential number of Nash equilibria can yield extremely complex CE. Such games can be easily constructed with very simple graphs.

LEMMA 2. *For any graphical game $(G, \mathcal{M})$, if $P \in \mathcal{CE}(G, \mathcal{M})$ and $P \equiv_{LN} Q$ then $Q \in \mathcal{CE}(G, \mathcal{M})$.*

PROOF. The lemma follows by noting that the correlated equilibrium expectation equations are only dependent upon the marginal distributions of local neighborhoods, which are preserved in Q. □

While explicitly representing *all* CE is infeasible even in simple graphical games, we next show that we *can* concisely represent, in a single model, all CE *up to local neighborhood (and therefore payoff) equivalence*. The amount of space required is comparable to that required to represent the graphical game itself, and allows us to explore or enumerate the different outcomes achievable in the space of CE.

3.2 Correlated Equilibria and Markov Nets

In the same way that graphical games provide a concise language for expressing local interaction in game theory, *Markov networks* exploit undirected graphs for expressing local interaction in probability distributions. It turns out that (a special case of) Markov networks are a natural and powerful language for expressing the CE of a graphical game, and that there is a close relationship between the graph of the game and its associated Markov network graph. We begin with the necessary definitions.

DEFINITION 5. *A* local Markov network *is a pair $M \equiv (G, \Psi)$ where*

1. *G is an undirected graph on vertices $\{1, \dots, n\}$;*
2. *Ψ is a set of* potential functions*, one for each local neighborhood $N(i)$, mapping binary assignments of values of $N(i)$ to the range $[0, \infty)$:*

$$\Psi \equiv \{\psi_i : \{\vec{a}^{\,i}\} \to [0, \infty).\}$$

Here $\{\vec{a}^{\,i}\}$ is simply the set of all $2^{|N(i)|}$ settings to $N(i)$.

A local Markov network M compactly represents a probability distribution P_M as follows. For any binary assignment $\vec{a}$ to the vertices, define

$$P_M(\vec{a}) \equiv \frac{1}{Z}\left(\prod_{i=1}^{n} \psi_i(\vec{a}^{\,i})\right)$$

where $Z = \sum_{\vec{a}} \prod_{i=1}^{n} \psi_i(\vec{a}^{\,i}) > 0$ is the normalization factor.

Note that any joint distribution can be represented as a local Markov network on a sufficiently dense graph. (Note that if we let G be the complete graph then we simply have a single potential function over the entire joint action space $\vec{a}$.) However, if d is the size of the maximal neighborhood in G, then the representation size of a distribution in this network is $O(n2^d)$, a considerable savings over a tabular representation if $d << n$.

Local Markov networks are a special case of Markov networks, a well-studied probabilistic model in AI and statistics [Pearl(1988)]. A Markov network is typically defined with potential functions ranging over settings of maximal cliques in the graphs. Another approach we could have taken is to transform the graph G to a graph G' which forms cliques of the local neighborhoods $N(i)$, and then used standard Markov networks over G' as opposed to local Markov networks over G. However, this can sometimes result in an unnecessary exponential blow-up of the size of the model when the resulting maximal cliques are much larger than the original neighborhoods. As the following lemma shows, for our purposes, it is sufficient to define the potential functions over just local neighborhoods (as in our definition) rather than maximal cliques in

G', which avoids this potential blow-up. (However, we will sometimes find it useful to discuss G' at various points when connecting with the Markov network literature.)

The following technical lemma, which is the cornerstone of our first main theorem below, establishes that a local Markov network always suffices to represent a distribution up to local neighborhood equivalence.

LEMMA 3. *For all graphs G, and for all joint distributions P over joint actions, there exists a distribution Q that is representable as a local Markov network with graph G such that $Q \equiv_{LN} P$ with respect to G.*

The proof is provided in the appendix. The main result of this section now follows from the previous lemmas, and shows that we can represent any correlated equilibria of a graphical game $(G, \mathcal{M})$, up to payoff equivalence, with a local Markov network (G, Ψ).

THEOREM 4. *(CE Representation Theorem) For all graphical games $(G, \mathcal{M})$, and for all distributions $P \in \mathcal{CE}(G, \mathcal{M})$, there exists a distribution Q such that:*

1. *$Q \in \mathcal{CE}(G, \mathcal{M})$;*
2. *$Q \equiv_{EP} P$;*
3. *Q can be represented as a local Markov network with graph G.*

Note that the representation size for any local Markov network with graph G is linear in the representation size of the graphical game, and thus we can represent the CE of the game parsimoniously.

4. CORRELATED EQUILIBRIA IN GRAPHICAL GAMES: ALGORITHMS

Having established in Theorem 4 that a concise graphical game yields a concise representation of its CE up to payoff equivalence, we now turn our attention to algorithms for *computing* CE. In the spirit of our results thus far, we are interested in algorithms that can efficiently exploit the compactness of graphical games.

It is well-known that it is possible to compute CE via linear programming in time polynomial in the standard *non-compact* normal form. In this approach, one variable is introduced for every possible joint action probability $P(\vec{a})$, and the constraints enforce both the CE condition and the fact that the variables must define a probability distribution. It is not hard to verify that the constraints are all linear and there are $O(2^n)$ variables and constraints in the binary action case. By introducing any linear optimization function, one can get an algorithm based on linear programming for computing a single exact CE that runs in time polynomial in the size of the normal-form representation of the game (that is, polynomial in 2^n).

For graphical games this solution is clearly unsatisfying, since it may require time exponential in the size of the graphical game. What is needed is a more concise way to express the CE and distributional constraints — ideally, linearly in the size of the graphical game representation. We shall show that this is indeed possible for tree graphical games (and more generally for game graphs G where a Markov network on the graph G' — obtained by forming cliques of the neighborhoods of G — can be represented by a local Markov network on G and G' is chordal). The basic idea is to express both the CE and distributional constraints entirely in terms of the local marginals, rather than the global probabilities of joint actions. We begin with a lemma establishing that doing so for the CE constraints is straightforward.

LEMMA 5. *For all graphical games $(G, \mathcal{M})$, and for all action distributions P, $P \in \mathcal{CE}(G, \mathcal{M})$ if and only if for all players i and actions a, a':*

$$\sum_{\vec{a}^{\,i}: a^i_i = a} P(\vec{a}^{\,i}) M_i(\vec{a}^{\,i}) \geq \sum_{\vec{a}^{\,i}: a^i_i = a} P(\vec{a}^{\,i}) M_i(\vec{a}^{\,i}[i : a']).$$

PROOF. Due to the locality of the M_i, the CE constraint condition (Definition 1) simplifies to using only local expectations:

$$E_{\vec{a} \sim P | a_i = a} M_i(\vec{a}) = E_{\vec{a}^{\,i} \sim P(\vec{a}^{\,i} | a_i = a)} M_i(\vec{a}^{\,i})$$

Since $P(a_i = a)$ is a constant, we can multiply this on both sides of the equation and prove the "if" direction. For the "only if" direction, reverse the derivation. $\square$

Furthermore, for the case in which the game graph is a tree, it suffices to introduce linear distributional constraints over only the local marginals, along with *consistency* constraints on the *intersections* of local marginals. We thus have the following three categories of local constraints defining our linear program:

Variables: For every player i and every assignment $\vec{a}^{\,i}$, there is a variable $P_i(\vec{a}^{\,i})$.
LP Constraints:

1. *CE Constraints:* for all players i and actions a, a',

$$\sum_{\vec{a}^{\,i}: a^i_i = a} P_i(\vec{a}^{\,i}) M_i(\vec{a}^{\,i}) \geq \sum_{\vec{a}^{\,i}: a^i_i = a} P_i(\vec{a}^{\,i}) M_i([\vec{a}^{\,i}[i : a'])$$

2. *Neighborhood Marginal Constraints:* for all players i,

$$\forall \vec{a}^{\,i}: \;\; P_i(\vec{a}^{\,i}) \geq 0; \;\; \sum_{\vec{a}^{\,i}} P_i(\vec{a}^{\,i}) = 1$$

3. *Intersection Consistency Constraints:* for all players i and j, and for any assignment $\vec{y}^{\,ij}$ to the *intersection neighborhood* $N(i) \cap N(j)$,

$$\begin{aligned} P_i(\vec{a}^{\,ij}) &\equiv \sum_{\vec{a}^{\,i}: \vec{a}^{\,ij} = \vec{y}^{\,ij}} P_i(\vec{a}^{\,i}) \\ &= \sum_{\vec{a}^{\,j}: \vec{a}^{\,ij} = \vec{y}^{\,ij}} P_j(\vec{a}^{\,j}) \\ &\equiv P_j(\vec{a}^{\,ij}) \end{aligned}$$

Choices for the objective function are discussed below.

Note that if d is the size of the largest neighborhood, this system involves $O(n2^d)$ variables and $O(n2^d)$ linear inequalities, which is linear in the representation size of the original graphical game, as desired.

It is clear that any CE must satisfy these constraints. Here we must solve the inverse problem: given only a set of marginal assignments $\{P_i(\vec{a}^{\,i})\}$ satisfying the constraints, we would like to construct a CE. Of course, by Lemma 5, if we can find a distribution Q whose marginals match the assignments $(\forall i, \forall \vec{a}^{\,i}, Q(\vec{a}^{\,i}) = P_i(\vec{a}^{\,i}))$, then Q must be a CE. In general finding such a Q may be a difficult problem (and may have no solution), but for tree graphical games, it turns out that any solution to the constraints yields a unique joint distribution that can be represented in a the local Markov network. The following lemma states this.

LEMMA 6. *(Consistent Tree) For all graphical games $(G, \mathcal{M})$ in which $G = (V, E)$ is a tree, and for all assignments $\{P_i(\vec{a}^{\,i})\}$*

satisfying the consistency equations, there is a unique joint distribution Q, defined by

$$Q(\vec{a}) \equiv \frac{\prod_{i \in V} P_i(\vec{a}^{\,i})}{\prod_{(i,j) \in E; i<j} P_i(\vec{a}^{\,ij})} \tag{1}$$

such that $Q \in \mathcal{CE}(G, \mathcal{M})$ and Q is representable as a local Markov network with graph G. Furthermore, the marginals of Q will be consistent with the assignment: $\forall i, \vec{a}^{\,i}, Q_i(\vec{a}^{\,i}) = P_i(\vec{a}^{\,i})$.

PROOF. Given the graph $G = (V, E)$, we construct a graph $G' = (V, E')$ where $E' = \{(i, j) : \exists k \in V : i \neq j, \{i, j\} \subseteq N(k)\}$. Thus, G' has a clique for each local neighborhood in G. G' is a chordal graph[3] and every maximal clique $C \in G'$ is a subset of some neighborhood ($\exists i : \ C \subseteq N(i)$). From this, Theorem 2.6 of [Dawid and Lauritzen(1993)] implies that the joint distribution Q is unique, consistent and is representable by a local Markov network with graph G.

To see that $Q \in \mathcal{CE}(G, \mathcal{M})$, note that $\forall i, \vec{a}^{\,i}, Q(\vec{a}^{\,i}) = P_i(\vec{a}^{\,i})$ implies that the correlated equilibrium constraints hold for Q. $\square$

This result allows us to keep the number of constraints linear in the graphical game size, and an efficient LP algorithm emerges. Before presenting this algorithm, we note that for trees, we have compactly specified the convex polytope of *all* CE up to local neighborhood (and therefore payoff) equivalence. Let $\mathcal{CE}_{local}(G, \mathcal{M})$ be the set of all distributions Q:

$$Q(\vec{a}) \equiv \frac{\prod_{i \in V} P_i(\vec{a}^{\,i})}{\prod_{(i,j) \in E; i<j} P_i(\vec{a}^{\,ij})}$$

obtained by ranging over all solutions $\{P_i(\vec{a}^{\,i})\}$ to the consistency constraints. The following theorem shows that this set is expected payoff equivalent to $\mathcal{CE}(G, \mathcal{M})$.

THEOREM 7. *(Completeness) For all graphical games $(G, \mathcal{M})$ such that G is a tree, we have $\mathcal{CE}_{local}(G, \mathcal{M}) \subseteq \mathcal{CE}(G, \mathcal{M})$; and if $P \in \mathcal{CE}(G, \mathcal{M})$ then there exists a $Q \in \mathcal{CE}_{local}(G, \mathcal{M})$ such that $P \equiv_{EP} Q$.*

PROOF. Lemma 6 gives us $\mathcal{CE}_{local}(G, \mathcal{M}) \subseteq \mathcal{CE}(G, \mathcal{M})$. The remainder of the proof is constructive. Given P, first define:

$$Q_i(\vec{a}^{\,i}) \equiv P(\vec{a}^{\,i})$$

then construct Q using Equation 1. By construction, we have $Q \in \mathcal{CE}_{local}(G, \mathcal{M})$ and $P \equiv_{\mathrm{EP}} Q$ holds by Lemma 6. $\square$

Finally, to define a concrete LP algorithm we must simply specify an appropriate linear optimization function. One possibility is choosing a correlated equilibrium with the highest total expected payoff over all players:

$$\max_{\text{subject to the constraints on } P_i(\vec{a}^{\,i})} \sum_i \sum_{\vec{a}^{\,i}} P_i(\vec{a}^{\,i}) M_i(\vec{a}^{\,i}).$$

Our main algorithmic result follows.

THEOREM 8. *(Efficient Tree Algorithm) For all tree graphical games $(G, \mathcal{M})$ and all linear objective functions $F(\{P_i(\vec{a}^{\,i})\})$, linear programming computes a CE in time polynomial in the size of the graphical game. The CE computed can be varied by varying the objective function F.*

Hence, the algorithm is polynomial in n and exponential in d for trees. For the multiple action setting, it is straightforward to see that the previous complexity result still holds.

[3] A chord is an edge which connects two nonadjacent nodes in a cycle. A chordal graph is an undirected graph where every cycle larger than 3 has a chord.

5. DISCUSSION AND EXTENSIONS

While Theorem 8 establishes that we can efficiently compute different CE by varying the chosen linear optimization function, an even more powerful result follows by combining the convexity of the class $\mathcal{CE}_{local}(G, \mathcal{M})$ with the results of [Dyer et al.(1991)Dyer, Frieze, and Kannan], who provide an efficient algorithm for sampling nearly uniformly from an arbitrary convex body. This yields a method for sampling CE in $\mathcal{CE}_{local}(G, \mathcal{M})$ nearly uniformly in time polynomial in the size of a tree graphical game.

An important advantage of the Markov net formalism for representing CE (that we will elaborate on in future work) is the ability to immediately infer various conditional independences and probabilistic semantics from the graph alone. Indeed, this is perhaps the main power of graphical models for probabilistic inference, and the results presented here import that power into the domain of game theory. As just one example, let G be the graph of the game, and G' be the graph obtained by forming cliques of local neighborhoods in G. Let i and j be any two players in the game, and let S be any set of players forming a cut of G' such that i and j lie in different components. Then classical Markov net semantics, combined with our results, establish that for any CE, there exists a local neighborhood equivalent CE in which the distribution on a_i is independent of a_j given the actions in S. In other words, the Markov net derived from the game graph allows us to easily "read off" which conditional dependencies are essential, and which are spurious, in representing all CE. We expect that many other interesting and powerful interactions between the strategic and probabilistic aspects of CE will emerge with further study.

6. ACKNOWLEDGEMENTS

We give warm thanks to Dean Foster for numerous helpful discussions, and to Stuart Geman for his help with the proof of Lemma 3.

7. REFERENCES

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.

A. Berger, S. Della Pietra, and V. Della Pietra. A maximum entropy approach to natural language processing. *Computational Linguistics*, 22(1), March 1996.

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.

M. Dyer, A. Frieze, and R. Kannan. A random polynomial time algorithm for approximating the volume of a convex body. *JACM*, 38(1):1–17, 1991.

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.

M. Kearns, M. Littman, and S. Singh. Graphical models for game theory. In *Proceedings of the Conference on Uncertainty in Artificial Intelligence*, pages 253–260, 2001.

Daphne Koller and Brian Milch. Multi-agent influence diagrams for representing and solving games. *Games and Economic Behavior*. To appear.

P. La Mura. Game networks. In *Proceedings of the 16th Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 335–342, 2000.

M. Littman, M. Kearns, and S. Singh. An efficient exact algorithm for singly connected graphical games. In *Neural Information Processing Systems*, 2002.
J. F. Nash. Non-cooperative games. *Annals of Mathematics*, 54:286–295, 1951.
L. Ortiz and M. Kearns. Nash propagation for loopy graphical games. In *Neural Information Processing Systems*, 2003. To appear.
G. Owen. *Game Theory*. Academic Press, UK, 1995.
J. Pearl. *Probabilistic Reasoning in Intelligent Systems*. Morgan Kaufmann, 1988.
D. Vickrey and D. Koller. Multi-agent algorithms for solving graphical games. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, 2002.

8. APPENDIX

Proof of Lemma 3: The objective is to find a single distribution Q that is consistent with the players local neighborhood marginals under P and is also a Markov network with graph G. It is an immediate consequence of previous work on maximum entropy models (see [Berger et al.(1996)Berger, Pietra, and Pietra]) that the maximum entropy distribution Q^*, subject to $P \equiv_{\text{LN}} Q^*$, is a local Markov network.

More formally, we show that the solution to the following constrained maximum entropy problem is representable in G:

$$Q^* = \operatorname*{argmax}_Q H(Q) \equiv \operatorname*{argmax}_Q \sum_{\vec{a}} Q(\vec{a}) \log(1/Q(\vec{a}))$$

subject to

1. $Q(\vec{a}^{\,i}) = P(\vec{a}^{\,i})$, for all $i, \vec{a}^{\,i}$.
2. Q is a proper probability distribution.

Note first that this problem always has a unique answer since $H(Q)$ is strictly concave and all constraints are linear. In addition, the feasible set is non-empty, as it contains P itself.

To get the form of Q^*, we solve the optimization problem by introducing the Lagrange multipliers $\vec{\lambda} \equiv (\lambda_{i,\vec{a}^{\,i}}, \forall i, \vec{a}^{\,i})$ to take care of the neighborhood marginal constraints (condition 1), and β to take care of the normalization constraint (condition 2). The optimization becomes

$$\begin{aligned} Q^* &= \operatorname*{argmax}_{Q,\vec{\lambda},\beta} L(Q, \vec{\lambda}, \beta) \\ &\equiv \operatorname*{argmax}_{Q,\vec{\lambda},\beta} H(Q) + \sum_{i \in V} \sum_{\vec{a}^{\,i}} \lambda_{i,\vec{a}^{\,i}} (Q(\vec{a}^{\,i}) - P(\vec{a}^{\,i})) \\ &\qquad + \beta(\sum_{\vec{a}} Q(\vec{a}) - 1) \end{aligned}$$

where $Q(\vec{a})$ is constrained to be positive. Here, L is the Lagrangian function.

First note that for all $\vec{a}$, if $P(\vec{a}) = 0$, then $Q^*(\vec{a}) = 0$. A necessary condition for Q^* is that $\partial L / \partial Q(\vec{a})|_{Q=Q^*} = 0$, for all $\vec{a}$ such that $P(\vec{a}) > 0$. After taking derivatives and some algebra, this condition implies, for all $\vec{a}$,

$$Q^*_{\vec{\lambda}}(\vec{a}) = (1/Z_{\vec{\lambda}}) \prod_{v=1}^{n} I[P(\vec{a}^{\,i}) \neq 0] \exp(\lambda_{i,\vec{a}^{\,i}})$$

where $I[P(\vec{a}^{\,i}) \neq 0]$ is an indicator function which evaluates to 1 iff $P(\vec{a}^{\,i}) \neq 0$. We use the subscript $\vec{\lambda}$ on $Q^*_{\vec{\lambda}}$ and $Z_{\vec{\lambda}}$ to explicitly denote they are parameterized by the Lagrange multipliers.

It is important to note at this point that regardless of the value of the Lagrange multipliers, each $\lambda_{i,\vec{a}^{\,i}}$ is only a function of the $\vec{a}^{\,i}$s (that are consistent with $\vec{a}$). Let the dual function $F(\vec{\lambda}) \equiv L(Q^*_{\vec{\lambda}}(\vec{a}), \vec{\lambda}, 0)$, and let $\vec{\lambda}^*$ maximize this function. Note that those $\lambda_{i,\vec{a}^{\,i}}$ that correspond to $P(\vec{a}^{\,i}) = 0$ are irrelevant parameters since $F(\vec{\lambda})$ is independent of them. So for all i and $\vec{a}^{\,i}$ such that $P(\vec{a}^{\,i}) = 0$, we set $\lambda^*_{i,\vec{a}^{\,i}} = 0$. For all $i, \vec{a}^{\,i}$, we define the functions $\psi^*_i(\vec{a}^{\,i}) \equiv I[P(\vec{a}^{\,i}) \neq 0] \exp(\lambda^*_{i,\vec{a}^{\,i}})$. Hence, we can express the maximum entropy distribution Q^* as, for all $\vec{a}$,

$$Q^*_{\vec{\lambda}^*} = (1/Z_{\vec{\lambda}^*}) \prod_{i=1}^{n} \psi^*_i(\vec{a}^{\,i})$$

which completes the proof. $\square$