# Approximating Correlated Equilibria using Relaxations on the Marginal Polytope

**Hetunandan Kamisetty** HETU@CS.CMU.EDU
**Eric P. Xing** EPXING@CS.CMU.EDU
**Christopher James Langmead** CJL@CS.CMU.EDU
Carnegie Mellon University 5000 Forbes Ave, Pittsburgh, PA 15213 USA

## Abstract

In game theory, a Correlated Equilibrium (CE) is an equilibrium concept that generalizes the more well-known Nash Equilibrium. If the game is represented as a graphical game, the computational complexity of computing an optimum CE is exponential in the tree-width of the graph. In settings where this exact computation is not feasible, it is desirable to approximate the properties of the CE, such as its expected social utility and marginal probabilities.

We study outer relaxations of this problem that yield approximate marginal strategies for the players under a variety of utility functions. Results on simulated games and in a real problem involving drug design indicate that our approximations can be highly accurate and can be successfully used when exact computation of CE is infeasible.

## 1. Introduction

Game theory is the study of strategic scenarios wherein the utility of each individual's actions depends on the actions made by the other players. One of the primary interests in game theory is computing *equilibrium strategies* for specific games. An equilibrium strategy is a joint probability distribution over actions where no player has an incentive to deviate from the distribution. There are a number of different kinds of equilibrium strategies. A Nash Equilibrium (NE), for example, corresponds to a strategy that factorizes completely over the set of players. A *Correlated Equilibrium* (Aumann, 1974) generalizes the NE by allowing for arbitrary joint distributions over actions. The goal of this paper is to introduce a new algorithm for computing approximate Correlated Equilibria in large multi-player games.

Correlated Equilibria (CE) are of practical and theoretical importance because they have a number of attractive properties. For example, they can lead to higher expected utilities than NE, and are also easier to compute. CE can also be viewed as a Bayesian alternative to NE (Aumann, 1987). Finally, there exist natural algorithms that allow the players to iteratively converge to a CE (Foster & Vohra, 1997); no such algorithms are known to exist for NE, in general.

There are, however, a number of challenges associated when computing CE, especially for games with many players and/or where each player has many actions from which to choose. The first challenge is representing the game itself. In this paper, we address this problem using graphical games (Kearns et al., 2001). A graphical game is a compact representation of a multiplayer game. By analogy to the way in which a probabilistic graphical model exploits the conditional independencies in a multi-variate distribution, a graphical game exploits the structure of the interactions between the players to compactly represent the payoffs in multiplayer games.

While graphical games significantly reduce the cost of representation, analyzing the CE of such games can be computationally expensive, even if the graph structure is sparse. In particular, while there is a polynomial-time algorithm for finding *some* equilibrium strategy in such games, determining an equilibrium strategy with specific properties is PPAD-complete (Papadimitriou & Roughgarden, 2008). This isn't surprising because finding such an equilibrium involves optimizing an objective function over the marginal polytope. For most realistic graph topologies, this problem can be computationally demanding and even infeasible.

Appearing in *Proceedings of the 28$^{th}$ International Conference on Machine Learning*, Bellevue, WA, USA, 2011.

Motivated by the success of outer relaxations in probabilistic inference, we study the accuracy of analogous relaxations to the problem of approximating marginal strategy profiles of the CE. We study a relaxations to the original problem that are easier to compute. These relaxations facilitate approximating the CE of graphical games in time polynomial to its cost of representation. Such relaxations provide an approximate picture of the equilibrium strategy. In addition, they also provide useful bounds (e.g., bounds on the expected payoff, bounds on the social utility of the game). We note that this problem also has natural applications in learning graphical games from data, and in designing games with desired outcomes.

Using randomly generated games, we show that outer relaxations are surprisingly accurate: the median absolute error in the strategy profile of two-action games even when using the weakest approximation on very dense graphs is less then 0.2, with the best approximations giving near-zero error at a marginally higher computational cost. The corresponding bounds are often less than 2% away from the exact value. We then analyze the CE of a game motivated by computer-aided drug design. By computing a lower bound on the binding affinity of a drug, our method is able to quickly determine whether the drug is prone to resistance-conferring mutations.

## 2. Introduction to Game Theory

**Definition 1 (Normal form Game)** *An n-person normal-form game $G$ is a tuple $(X, A, u)$ where*

- *$X$ is a finite set of $n$ players;*
- *$A = \prod A_i$ where $A_i$ is the set of actions available to player $i$;*
- *$u = (u_1, \dots u_i, \dots u_n)$ where $u_i : A \to \mathcal{R}$ is a* utility *function that player $i$ seeks to maximize.*

Given the set of possible actions available, a player's *strategy* $\mu_i : A^i \to [0, 1]$ is a probability distribution over choice of action. Actions are sampled from this distribution and independently of the other players. A strategy that puts all the probability mass on a single action is called a *pure strategy*, otherwise it is called a *mixed strategy*.

The $n$ player normal form game is commonly represented as $n$, $n$ dimensional matrices, one for each $u_i$. In subsequent subsections, we will use structured representations that exploit sparsity to efficiently store the game more compactly and develop algorithms that compute equilibria.

*Table 1.* The Game of Chicken. Two individuals challenge each other to either *Dare* or *Chicken out*. Each individual has an incentive to dare, but only if the other chickens out. This game has 3 NE and one additional CE.

| | D | C |
|---|---|---|
| D | (0,0) | (7,2) |
| C | (2,7) | (6,6) |

### 2.1. Equilibria

A *strategy profile* $\mu$ encodes the strategies for all players. A strategy profile is said to be a Nash Equilibrium (NE) of a game if no player has an incentive to unilaterally deviate from this profile. A seminal theorem of Nash (Nash, 1950) proved that every game with a finite number of players and action profiles has at least one Nash Equilibrium, such that $\forall i, \forall a^i, a'^i \in A^i$:

$$\sum_{a^{-i}} \mu(a^i)\mu(a^{-i})(u_i(a^i, a^{-i}) - u_i(a'^i, a^{-i})) \geq 0,$$

where $\mu(a^{-i}) = \mu(a^1) \times \dots \mu(a^{i-1})\mu(a^{i+1}) \dots \mu(a^n)$, since the definition of a mixed strategy profile specifies that each player samples from $\mu_i$ independent of other players.

Relaxing this requirement of independence results in equilibrium called Correlated Equilibria (CE)(Aumann, 1974). A CE is any joint distribution $\mu$ over the player's actions such that $\forall i, \forall a^i, a'^i \in A^i$:

$$\sum_{a^{-i}} \mu(a^i, a^{-i})(u_i(a^i, a^{-i}) - u_i(a'^i, a^{-i})) \geq 0$$

Tab. 1 shows the utilities of a classic two-player non-zero sum game, "Chicken", and Fig. 1 illustrates the NE and CE for this game. Each point in the green polytope corresponds to a valid CE. The hyperbolic surface represents the points that satisfy the constraint $\mu = \mu_1\mu_2$. Any point that is both a CE and on the hyperbolic surface is a NE.

Note that for an $n$-player game where each player has $m$ actions an explicit matrix representation of the game requires $O(n \times m^n)$ space. However, some games can be represented far more efficiently as *graphical games*, which are discussed in the following section.

### 2.2. Graphical Games

A Graphical Game $(G, M)$ is a tuple consisting of a graph $G$ and a set of local utility functions $M = (M_1, ..., M_n)$ that compactly represents a multi-player game. Each player is represented by a vertex in the graph $G$ and edges are drawn between vertices $i$ and

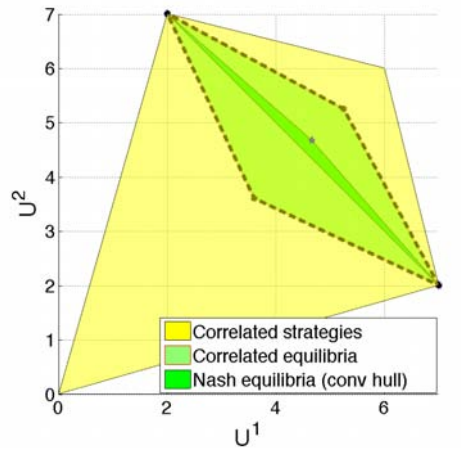


*Figure 1.* The constraints on NE and CE illustrated for the game "Chicken". The yellow linear polytope is the set of valid distributions. The green linear polytope is the set of CE. The constraints for a NE form a non-convex surface in dark green.

$j$ if the utility function of either depends on the actions of the other. If $M_i$ depends on the actions of all remaining players, this representation offers no advantages; however, when $M_i$ only depends on a subset of players, the graphical game representation can be more compact than the matrix form representation. In particular, the graphical game requires space exponential in the largest degree $d$ of the graph.

This representation is evocative of a Markov Random Field. While MRFs represent probabilistic interactions, the graphical game represents strategic interactions. An important result due to (Kakade et al., 2003) shows that there is a natural relationship between the structure of the graphical game and the probabilistic relations of a subset of its CE. To describe this result, we first need the following definition:

**Definition 2 (Expected Payoff-Equivalence)** *For a graph $G$, two distributions $P$, $Q$ are equivalent up to expected payoff if for all players $i$, and actions $\vec{a_i}$ over $\Delta_i$, $\mathbb{E}_{a\sim P}u_i(\vec{a_i}) = \mathbb{E}_{a\sim Q}u_i(\vec{a_i})$.*

Here, $\Delta_i$ refers to the set containing $i$ and its neighbors $N(i)$ in $G$, and $\vec{a_i}$ refers to an assignment of actions to $\Delta_i$.

(Kakade et al., 2003) show that any CE of the game that does not factorize according to $G$ is equivalent (up to expected payoff) to some CE of the game that does. The rest of this work therefore limits itself to the CE that factorize according to $G$.

## 3. CE in Graphical Games

Before we discuss algorithms, it will be useful to define some notation and sets of interest. We use indexed superscripts to denote components of these assignments: $\vec{a_i}^j$ corresponds to the assignment of an action to player $j$ according to $\vec{a_i}$. We use the notation $\vec{a_i}[i:a']$ to denote the vector that is obtained by replacing the $i^{th}$ component of $\vec{a_i}$ with action $a'$. $\{\mu_i\}$ will refer to the set of marginals, that each store the marginal probability over variable $i$ and its neighbors.

### 3.1. Exact Correlated Equilibria

If $G$ has no cycles, computing a CE of the graphical game can be solved satisfying a set of linear constraints. An important departure from the standard MRF treatment is that the potentials of a graphical game are usually *not* pairwise. This is due to the fact that the utility is a function from an assignment $\vec{a_i}$ of actions to all neighbors of $i$. Thus, in general, the smallest region over $i$ that would allow incorporation of the CE constraints would need to include all neighbors of $i$.

The marginal polytope $\mathcal{M}_\Delta$ is defined to be the set of marginals $\mu_i(\vec{a_i})$ (over sets of variables of size $\Delta_i$) that are realizable by a joint distribution. This is analogous to the common definition of a marginal polytope over pair-wise marginals commonly used in approximate inference (Wainwright & Jordan, 2008), a set we shall refer to as $\mathcal{M}_2$.

$$\begin{aligned}\mathcal{M}_\Delta(G) &= \{\mu \in \Re^d | \exists p \text{ with marginals } \mu_i(\vec{a_i})\}\\ \mathcal{M}_2(G) &= \{\mu \in \Re^d | \exists p \text{ with marginals } \mu_i(a_i, a_j),\\ &\qquad \forall (i,j) \in E\}\end{aligned}$$

For an acyclic graph, the marginal polytope $\mathcal{M}_\Delta$ can be expressed using the following linear constraints over each $\Delta_i$ and their pair-wise intersections:

$$\begin{aligned}\text{Positivity:} &\quad \forall i, \forall \vec{a_i}, \mu_i(\vec{a_i}) \geq 0\\ \text{Local Normalization:} &\quad \forall i, \sum \mu_i(\vec{a_i}) = 1\\ \text{Local consistency:} &\quad \forall i, j, \forall N_{ij} \in \Delta_i \cap \Delta_j,\\ &\quad \text{and assignments } y^{N_{ij}}\end{aligned}$$

$$\sum_{\vec{a_i}:\vec{a_i}^{N_{ij}} = y^{N_{ij}}} \mu_i(\vec{a_i}) = \sum_{\vec{a_j}:\vec{a_j}^{N_{ij}} = y^{N_{ij}}} \mu_j(\vec{a_j})$$

Since these constraints involve variables that are adjacent to each other in $G$, the set of marginals that satisfy them are referred to as $\text{LOCAL}_\Delta(G)$. In a graph without cycles, $\text{LOCAL}_\Delta(G) = \mathcal{M}_\Delta(G)$, since the joint can always be uniquely reconstructed from the $\mu_i$'s in such graphs.

Since the definition of a CE only imposes constraints on the solution, it is common to determine a CE that optimizes some objective function, for example, the expected social utility $\sum_i \sum_{\vec{a_i}} \mu(\vec{a_i}) u_i(\vec{a_i})$. We will use $f(\mu)$ to refer to any such objective function. While we assume that $f$ is linear, our approach can be naturally extended to any class of functions as long as we can efficiently perform optimization, for example, minimizing convex functions.

The LP for computing a CE that maximizes some function $f$ can be expressed as:

$$\max f(\mu)$$

s.t.

$$\forall i, \forall a, a' \in A^i, \sum_{\vec{a_i}:\vec{a_i}^i=a} \mu_i(\vec{a_i})(u_i(\vec{a_i}) - u_i(\vec{a_i}[i:a'])) \geq 0 \quad \text{(CE constraints)}$$

$$\mu = [\mu_1 \dots \mu_n] \in \mathcal{M}_\Delta$$

### 3.2. Relaxations to outer marginal polytopes

If $G$ has cycles, $\text{LOCAL}_\Delta(G) \neq \mathcal{M}_\Delta(G)$. In such cases, one option is to run the algorithm on the junction tree $JT(G)$ instead of the original graph. Since by the junction tree property, $\text{LOCAL}_\Delta(JT(G)) = \mathcal{M}_\Delta(JT(G))$ (Wainwright & Jordan, 2008) this process is guaranteed to give the exact CE. However, the algorithm has to maintain marginals over the size of tree-width of the graph which can be prohibitively expensive in games with many players or actions even if each vertex of the graph has a small degree. For example, a grid-structured graphical game of size $n \times n$ with $k$ actions for each player has a maximum degree of 4 resulting in representation cost of $O(nk^4)$ but its tree-width is $n$.

If we are primarily interested in the value of the objective function $f(\mu)$ or in the marginal distribution induced by a CE, we can trade-off accuracy for time by approximating $\mathcal{M}_\Delta(G)$ with $\text{LOCAL}_\Delta(G)$. This is analogous to a Generalized Belief Propagation approach to inference on an MRF. Since $\text{LOCAL}_\Delta(G) \supset \mathcal{M}_\Delta(G)$, this is an outer relaxation of the problem. Solving it will therefore give a lower bound on the objective function.

The CE constraints require all marginals of size $\Delta_i$ implying that a relaxation looser than $LOCAL_\Delta$ is not possible for general games. However, in cases where the utility function additional structure, it is possible to construct further relaxations.

### 3.3. Pair-wise additive utility functions

Consider a setting where the utility of a particular player $u_i(\vec{a_i})$ can be expressed as the sum of pair-wise functions over the actions of the player's neighbors

$$u_i(\vec{a_i}) = \sum_{j \in N(i)} g_{i,j}(\vec{a_i}^i, \vec{a_i}^j)$$

The expression in the CE constraint can then be expressed as

$$\begin{aligned}
\sum_{\vec{a_i}:\vec{a_i}^i=a} \mu_i(\vec{a_i}) u_i(\vec{a_i}) &= \sum_{\vec{a_i}:\vec{a_i}^i=a} \mu_i(\vec{a_i}) \sum_{j \in N(i)} g_{i,j}(\vec{a_i}^i, \vec{a_i}^j) \\
&= \sum_{j \in N(i)} \sum_{\vec{a_i}:\vec{a_i}^i=a} \mu_i(\vec{a_i}) g_{i,j}(\vec{a_i}^i, \vec{a_i}^j) \\
&= \sum_{j \in N(i)} \sum_{a_j} \mu_i(a, a_j) g_{i,j}(a, a_j)
\end{aligned}$$

Thus, if the utility can be expressed as a sum of pair-wise functions, the constraint $\mu \in \mathcal{M}_\Delta$ in the LP for the exact CE can be replaced with $\mu \in \mathcal{M}_2$. It is possible to then construct an outer relaxation to $\mathcal{M}_2$ using local constraints over edges and vertices of the graph. We will therefore refer to this relaxation as $LOCAL_2$.

### 3.4. Cycle Inequalities

*Cycle inequalities* are among the constraints that are satisfied by $\mathcal{M}_2$ and not explicitly enforced by $LOCAL_2$. They arise from the following simple observation on a graph with binary variables: if we start from a node in a cycle and traverse the cycle to come back to the node this traversal must have seen an even number of edges where adjacent variables had different values (if not, the value at the end of the traversal must be different from at the beginning, a contradiction). The advantage of these constraints is that their violation can be detected and a violated constraint identified in graphs in polytime (Barahona & Mahjoub, 1986) by computing shortest paths in a related graph. Violated cycle inequalities are incorporated incrementally into the constraint set and the LP is re-solved until no more violations occur. (Sontag & Jaakkola, 2007) devise an extension of this idea for graphs over variables with $k$ values that first constructs different binary instantiations of these variables and identifies violated constraints on these instantiations. Since there are $2^k$ possible instantiations for each variable in the graph, we restrict ourselves to the *k-projection graph*(Sontag & Jaakkola, 2007) that considers $k$ different binary instantiations for a variable with $k$ values.

We will refer to the set of solutions that satisfy these constraints as $Cycle - Ineq$. $LOCAL_2 + Cycle - Ineq$ will therefore refer to solutions that are in both $LOCAL_2$ and $CYCLE - Ineq$.

## 4. Simulation Results

We generated 16-node graphs corresponding to games with 16 players. Each player had two actions, and an edge was included between any pair of vertices with probability $\rho$. We generated graphs with the value of $\rho$ varying from 0.01 and 0.5 (corresponding to a graph with half the density of a completely connected graph). We generated 10 graphs for each setting for $\rho$ resulting in a total of 70 graphs. For each player, each element in the utility matrix was randomly generated from $\mathcal{U}(-1, 1)$. We note that our choice of values was limited by a need to be able to compute CE exactly so as to compare the accuracy of our relaxations. The relaxations themselves can be used on games with many more players and actions.

For each of these graphs, we computed exact Correlated Equilibria while optimizing for social utility and computed the marginal distributions of the CE strategy for each player. We then computed approximate marginal distributions that optimized the same objective function but made progressively looser outer relaxations: the first relaxation stored marginal distributions over a player and all its neighbors, for each player while the second relaxation only stored pairwise marginal distributions, one for each edge in graph.

Fig. 2-A shows a boxplot of the error induced in the individual marginal distributions of each player due to approximate inference as $\rho$, the density of the graph, increases. In each group, the first box shows the error when using the $LOCAL_\Delta$ relaxation, the second box shows the errors when using $LOCAL_2$ and the third, when using $LOCAL_2 + Cycle - Ineq$. Increasing $\rho$ has marginal effect on the error of $LOCAL_\Delta$ but increases the error of $LOCAL_2$ . Remarkably, adding the cycle inequality constraints drastically improves the accuracy of $LOCAL_2$ : the error is nearly zero across all settings of $\rho$.

To test the sensitivity of these results to the choice of utility and objective function, we generated 100, 16-node graphs with $\rho = 0.3$. For each of these graphs, we then generated graphical games where the utilities for each action-pair on each edge were sampled from a bivariate normal distribution $\mathcal{N}([0, 0], [1, \sigma; \sigma, 1])$ where $-1 \leq \sigma \leq 1$ resulting in a total of 700 graphical games with wide variety of utilities. At $\sigma = -1$, for example, the utilities for any action pair on each edge sum

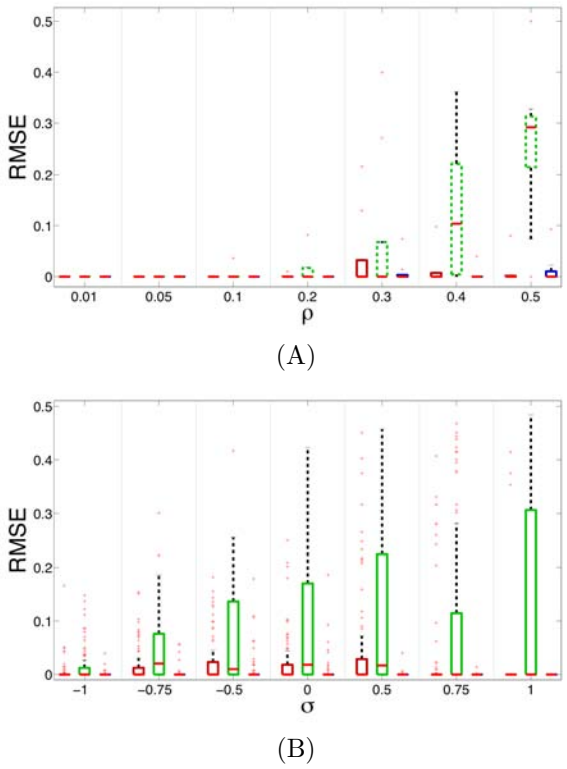


*Figure 2.* Boxplot showing $L_2$ error of marginal distributions on randomly generated graphs with utilities sampled from $\mathcal{U}(-1, 1)$ as the density, $\rho$, increases (A) and from a coupled distribution as the correlation between utilities, $\sigma$, increases(B). In each group, the first boxplot corresponds to $LOCAL_\Delta$, the second to $LOCAL_2$ and the third to $LOCAL_2 + Cycle - Ineq$.

to zero, while for $\sigma = 1$, the utilities on a particular edge are the same for both players. Instead of the social utility, we optimized a random linear function of the marginals. Fig. 2-B shows the error in marginals using the outer relaxations as $\sigma$ increased. $LOCAL_\Delta$ still has consistently low error while the spread of the distribution of errors for $LOCAL_2$ appears to increase as utilities become more coupled. Again incorporating the cycle inequalities to $LOCAL_2$ drastically improves the accuracy to near zero error.

Fig. 3-A shows a boxplot of the error induced in the individual marginal distributions of each player due to approximate inference as $\rho$, the density of the graph, increases. As previously, $LOCAL_\Delta$ and $LOCAL_2 + Cycle - Ineq$ are accurate in all settings while $LOCAL_2$'s error increases as the degree of the vertex increases.

Fig. 3-B shows the CDF of the percentage gap between the upper bound as computed by an outer relaxation

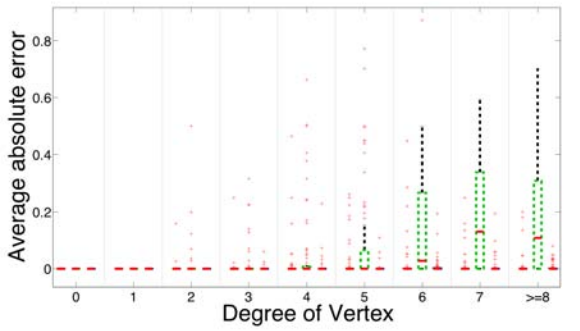


(A)

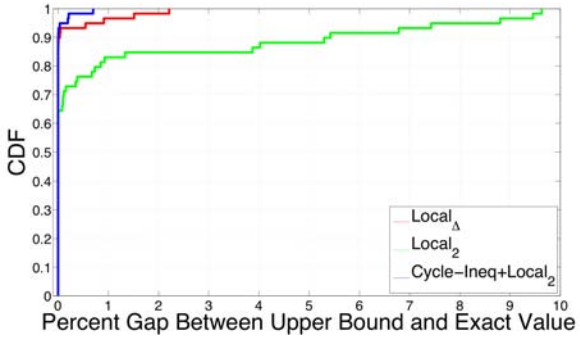


(B)

*Figure 3.* (A) Boxplot showing absolute error of the marginal distribution at a variable across all simulated graphs as a function of the degree of the position. As the degree increases, the error tends to increase although the median error remains very low. Order within each group: $LOCAL_\Delta, LOCAL_2, LOCAL_2 + Cycle - Ineq$. (B) Histogram of the percentage gap between the upper bound of the objective function using an outer approximation and the exact value. Both $LOCAL_2$ and $LOCAL_\Delta$ relaxations produce bounds that are remarkably close to the actual value of the objective function.

and the exact value using the marginal polytope. The bounds provided by all relaxations are $< 10\%$ away from the optimal value for all graphs with the maximum error of $LOCAL_2 + Cycle - Ineq$ being *less than* $1\%$.

The running time of the exact algorithm on our simulated graphs is a few hours even for moderately dense graphs. The $LOCAL_2$ and $LOCAL_2 + Cycle - Ineq$ takes less than ten minutes on the protein-drug game and is even faster on the simulated graphs.

## 5. Application to Drug Design

Rational drug-design is the process of engineering a drug to selectively bind to target molecules and modify their behavior. Traditionally, the target molecule, usually a protein, is assumed to have a fixed, known chemical composition. This assumption is valid for some targets (e.g., asthma drugs), but not others. In particular, some diseases are associated with rapid evolution of the target through mutations (e.g., in cancer, HIV, and bacterial infections). Such mutations can decrease the efficacy of the drug, leading to a phenomenon known as *resistance*.

To address this problem, we are developing a game-theoretic approach to drug design. Here, the game models the interactions between the drug designer and the target protein: the target makes "moves" by introducing mutations, and the drug designer makes moves by choosing one of several candidate drugs. The utility of any given set of moves is the binding affinity of the drug to the (mutated) protein. If these moves result in a protein-drug pair that has high affinity, the drug designer wins; else, the protein wins. The aim of the drug designer, therefore, is to find a drug (or drug cocktail) that binds well against all moves of the protein.

### PDZ

The PDZ domain is a family of related structural motifs found in many signaling proteins. Biologically, it mediates protein-protein interactions. Since some of these interactions are affected in cancerous cells, members of this family are being studied as targets for anti-cancer drugs.

Along with our collaborators, we have developed a sparse linear model that accurately predicts the binding affinity of PDZ domains to short peptides. In our game, these peptides are the drugs and the binding affinities are the utilities. The model (Fig. 4) was learnt using block sparse linear regression from experimental data and consists of thirty eight blocks (corresponding to edges here) and twenty one variable positions (corresponding to sixteen 'protein' players and five 'drug' players). Each protein player has five actions corresponding to the wild-type and the four other most likely amino acid positions at this position, as observed in nature. This restricts the game to mutations that are energetically favorable.

The maximum degree of this graph was 10; neither exact computation of the CE nor the $LOCAL_\Delta$ approximation had sufficient memory to complete this computation on a standard desktop machine with 4GB RAM. We therefore only report the results of $LOCAL_2 + Cycle$.

In the drug-design game, a CE encodes a distribution over drugs (known as a drug cocktail) *and* a distribu-

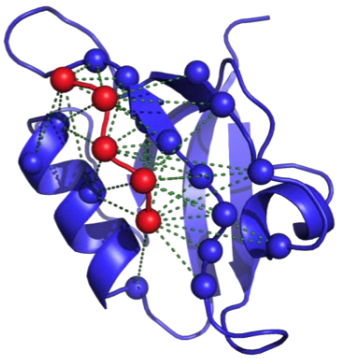

*Figure 4.* The graphical game describing the pdz-drug game is shown overlaid on a protein structure. The players are labeled with spheres (pdz in blue, drug in red) and the edges of the game are shown with dashed lines. The game has 21 players and 38 edges.

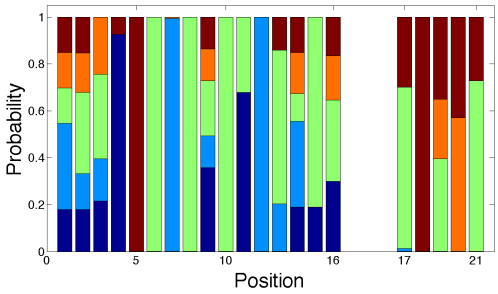


(A)

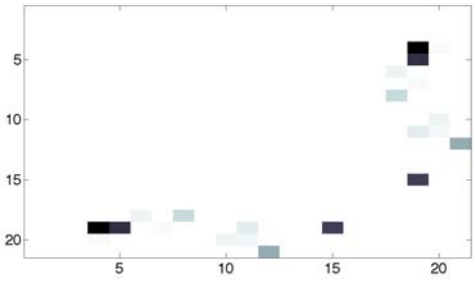


(B)

*Figure 5.* (A) The marginal profile of each player for the worst-case equilibrium scenario. The 16 positions corresponding to the PDZ and the 5 positions corresponding to drug are shown in separate groups. (B) The expected binding energy for each edge in. By choosing mutations that maximize disruption to binding, PDZ is able to shut off interactions on all but eleven edges. However, these eleven edges together ensure that the overall binding energy is negative indicating that the cocktail is successful.

tion over PDZ sequences. The distribution over PDZ mutations can be interpreted as the set of mutations that are likely to arise in response to the cocktail. The key question is whether the expected binding energy of the cocktail and the PDZs is negative (i.e., favorable) or positive (i.e., unfavorable). If positive, then the PDZ is resistant to the cocktail. Thus, it is sufficient to determine whether the expected binding energy is guaranteed to be negative. Therefore, we computed the (approximate) CE that would *maximize* the binding energy (i.e a "pessimistic" CE) and thus obtained an upper bound on the maximum binding energy. Naturally, if this upper bound is negative then the true worst-case expected binding energy is also negative which, in turn, implies that the cocktail is successful at *any* CE of the game.

The expected binding energy was negative (-5.83 kcal/mol) indicating that this cocktail-drug is predicted to successfully bind to all viable mutants of the PDZ. Fig. 5-A shows the marginal profile for each position across its five actions. It is interesting to note that compared to the PDZ, the profile of the drug is limited to a few actions per position, possibly due to the pessimistic nature of our prediction.

Fig. 5-B shows the breakdown of this total binding energy across the 38 edges. Only eleven of the 38 edges had a non-zero contribution to the binding energy. In contrast, the "optimistic" CE of this game (obtained by minimizing binding energy) had 20 edges that had non-zero contribution to the binding energy.

## 6. Discussion and Conclusion

We demonstrated an approach to approximate the properties of a CE in graphical games. Our approach, based on outer relaxations to the marginal polytope computes these approximations efficiently. On a large set of games with different types of utilities, we demonstrated that these relaxations are also remarkably accurate, often giving the exact solution. In addition, our approach bounds the objective function. When used with the social utility for example, our approach can be used to bound the price of anarchy of the game (Koutsoupias & Papadimitriou, 2009).

The main aim of the protein-drug study is to determine if a set of drugs (drug cocktail) has good worst-case binding properties. Our relaxations are quick and accurate implying that we can determine such a set rapidly. Thus in this application, our main interest is in determining the worst-cast binding energy. We showed the resulting approximate strategy profile to aid in understanding of the game and provide insights

into interactions. While our method produces approximate strategies, we note that the cost of sequencing continues to drop. Thus, it is relatively easy to "observe" the adversarial disease's strategy in the population, or even in a patient-specific fashion. Such information could be incorporated to adapt the cocktail to combat resistance.

In our application of game-theory to a biological problem, we follow a rich body of prior work that model evolutionary behavior in this manner, starting with the pioneering work of (Hamilton, 1964a;b) and (Smith & Price, 1973). Indeed, it can be argued that the process of evolution is best modeled in such a strategic manner: organisms adapt to the behavior of other organisms and the environment to maximize their chances of survival. While these adaptations by themselves are not the output of strategic behavior, the combination of random mutations with non-random selective pressures can, and is often modeled as such.

In (Pérez-Breva et al., 2007), the authors develop a game theoretic model of Protein-DNA binding. The focus of their model is to determine accurately the *configuration* used by various proteins to bind to DNA and the effects of space and resource constraints on the equilibrium configurations. While our model is also used in a biological setting, the focus of our model is on protein-drug binding and the *composition* of these molecules. Thus our model enables understanding of evolutionary changes and equilibria in contrast to spatial changes and is thus fundamentally different.

By modeling the smallest unit of evolutionary change (the mutation of an amino acid) this paper considers strategic behaviors at a much higher-resolution than previous applications of game-theory to the study of evolution. In the future, we hope to extend this promising approach to learning the structure and utilities of graphical games from data and in applying such models to the design of new drugs.

## Acknowledgments

Corresponding author: cjl@cs.cmu.edu. This research was performed under NSF grant IIS-0905193 to CJL, and an allocation on Blacklight provided by the Pittsburgh Supercomputing Center. EPX was supported by an Alfred P. Sloan Fellowship, NIH grant NIH1R01GM093156 and an NSF Career award (NSF DBI-0546594). We thank our collaborator, Prof. Chris Bailey-Kellogg at Dartmouth College for allowing us the use of the protein-drug model used in Section 5.

## References

Aumann, R.J. Subjectivity and correlation in randomized strategies. *J. Math. Econ.*, 1(1):67–96, 1974.

Aumann, R.J. Correlated equilibrium as an expression of Bayesian rationality. *Econometrica*, 55(1):1–18, 1987.

Barahona, F. and Mahjoub, A.R. On the cut polytope. *Mathematical Programming*, 36(2):157–173, 1986.

Foster, D.P. and Vohra, R.V. Calibrated Learning and Correlated Equilibrium* 1. *Games and Economic Behavior*, 21(1-2):40–55, 1997.

Hamilton, W. D. The genetical evolution of social behaviour. I. *J. Theor. Biol.*, 7(1):1–16, 1964a.

Hamilton, W. D. The genetical evolution of social behaviour. II. *J. Theor. Biol.*, 7(1):17–52, 1964b.

Kakade, S., Kearns, M., Langford, J., and Ortiz, L. Correlated equilibria in graphical games. In *Proceedings of the 4th ACM Conference on Electronic Commerce*, pp. 42–47. ACM, 2003.

Kearns, M., Littman, M., and Singh, S. Graphical models for game theory. In *Proceedings of the 17th Conference in Uncertainty in Artificial Intelligence*, pp. 253–260. Citeseer, 2001.

Koutsoupias, E. and Papadimitriou, C. Worst-case equilibria. *Comp. Sci. Rev.*, 3(2):65–69, 2009.

Nash, J.F. Equilibrium points in n-person games. *Proc. Nat. Acad. Sci.*, 36(1):48–49, 1950.

Papadimitriou, C.H. and Roughgarden, T. Computing correlated equilibria in multi-player games. *Journal of the ACM (JACM)*, 55(3):1–29, 2008.

Pérez-Breva, L., Ortiz, L.E., Yeang, C., and Jaakkola, TS. Game theoretic algorithms for Protein-DNA binding. *Advances In Neural Information Processing Systems*, 19:1081–1088, 2007.

Smith, J. Maynard and Price, G. R. The Logic of Animal Conflict. *Nature*, 246(5427):15–18, 1973.

Sontag, D. and Jaakkola, T. New outer bounds on the marginal polytope. *Advances in Neural Information Processing Systems*, 20:1393–1400, 2007.

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