Correlated Equilibria in Multi-Player Games

Andreas Hoenselaar
awphoe@cs.brown.edu
Department of Computer Science, Brown University, Providence, RI, USA

May 29, 2007

Abstract

This paper gives a constructive proof for the existence of correlated equilibria in multi-player games and shows how elements of the proof can be used for the computation of correlated equilibria in polynomial time for an important subset of succinct games. None of this work is novel, it merely aims at a more detailed and approachable description of the proof in Papadimitriou's journal article "Computing Correlated Equilibria in Multi-Player Games"[3] and provides short summaries of the fundamental techniques employed in the proof.

1 Equilibria in Multi-Player Games

1.1 Multi-Player Games

A multi-player game G is described by a set of players $p_1, \ldots, p_n$, each of which can choose from a set of strategies S_p. All possible combinations of strategies are elements of the set of strategy profiles $S = \prod_{p=1}^{n} S_p$. To complete the game description, we supply a payoff function $u^p : S \longrightarrow \mathbb{Z}$ for each player.

Given these basic definitions, we can describe the players' behavior by defining distributions over the set of strategies. Every player chooses one of the elements from his strategy set with a certain probability and thus we can build a strategy profile $x^p \in [0,1]^{|S_p|}$ with $\sum_{i=1}^{|S_p|} x_i^p = 1$ as a parsimonious description. Extending this concept to the whole game and all players yields a distribution x on S: $x \in [0,1]^{|S|}$ with $\sum_{i=1}^{|S|} x_i = 1$. Based on the formal game description we have developed, we are going to introduce the concept of game equilibria in the next section.

1.2 Definition of Correlated Equilibria

The formal framework presented so far raises the question whether certain states of optimality[1] exist in multi-player games. If we accept that all players strive to optimize their payoffs, it can be imagined that a certain distribution over the set of strategy profiles will guarantee optimal payoffs for all participants. It is obvious that such a steady-state can only exist if no player has an incentive to deviate from the proposed distribution. As we have already noticed, the choice of strategies is probabilistic and consequently only the expected payoff can be optimized.

Definition 1.1 *A distribution $x \in \Delta(S)$ is a correlated equilibrium if the following is true for each player p and all pairs $(i,j) \in S_p \times S_p$ in his set of strategies: If all players adhere to a strategy profile x that recommends player p to choose strategy i, he has no incentive to play another strategy j instead and it holds that*

$$\sum_{s \in S_{-p}} (u_{is}^p - u_{js}^p) x_{is} \geq 0 \ .$$

To make this definition more approachable, we will calculate correlated equilibria for a simple two-player game in the next section and examine how the concept of correlated equilibria relates to the well-known notion of Nash equilibria [1].

[1] Without supplying a strict definition of optimality, we rely on the reader's intuition of this concept.

1.3 Correlated Equilibria vs. Nash Equilibria

Definition 1.1 does not impose any restriction on the equilibrium distribution x (apart from it being valid of course). One significant concept in game theory that has found wide-spread use in economics and related fields is that of Nash equilibria. Formally, the only additional restriction imposed on Nash equilibria is that the distribution x must be a product distribution such that it is fully defined by its marginal distributions $x^p_{S_p}$:

$$x_s = \prod_{p=1}^{n} x^p_{s_p} \quad \forall s \in S \tag{1.1}$$

	stop	go
stop	$4, 4$	$1, 5$
go	$5, 1$	$0, 0$

Table 1: Payoff matrix for the Chicken game

In order to get a better feeling for equilibria, we will calculate them for the "Chicken game". We use the payoff matrix in Table 1 and the inequalities in Definition 1.1 to assemble a set of constraints that every correlated equilibrium distribution has to fulfill:

$$\begin{aligned}
(4-5)x_{1,1} + (1-0)x_{1,2} &\geq 0 \\
(5-4)x_{2,1} + (0-1)x_{2,2} &\geq 0 \\
(4-5)x_{1,1} + (1-0)x_{2,1} &\geq 0 \\
(5-4)x_{1,2} + (0-1)x_{2,2} &\geq 0 \\
\sum_{i,j \in \{1,2\}} x_{i,j} &= 1 \\
0 \leq x &\leq 1
\end{aligned} \tag{1.2}$$

The first two equations state optimality of the distribution for player 1 (row player) by comparing the payoffs for the cases when player 2 chooses strategy "stop" (first equation) and "go" (second equation). We observe that describing a correlated equilibrium with linear inequalities requires one equation per player and pairs of strategies in his set. The equation system in (1.2) can be rewritten in matrix form:

$$Ux = \begin{pmatrix} -1 & 0 & 1 & 0 \\ 0 & 1 & 0 & -1 \\ -1 & 1 & 0 & 0 \\ 0 & 0 & 1 & -1 \end{pmatrix} \cdot \begin{pmatrix} x_{1,1} \\ x_{2,1} \\ x_{1,2} \\ x_{2,2} \end{pmatrix} \geq 0 \tag{1.3}$$

A short glance at the matrix U in equation (1.3) reveals that the second and third columns have positive sum while the first and last columns have negative sum. Exploiting this property, several straightforward equilibria distributions can be written down (Table 2). It can also be seen that there is an infinite number of correlated equilibria, an intuitive subset is the convex hull of the pure Nash equilibria in Table 2. The unique mixed Nash equilibrium is depicted in Table 3 and those equilibria that are correlated but not Nash are shown in Table 4

0	0
1	0

0	1
0	0

Table 2: Pure Nash equilibria for the chicken game

$\frac{1}{4}$	$\frac{1}{4}$
$\frac{1}{4}$	$\frac{1}{4}$

Table 3: Mixed Nash equilibrium for the chicken game

0	$\frac{1}{2}$
$\frac{1}{2}$	0

0	$\frac{3}{4}$
$\frac{1}{4}$	0

Table 4: Correlated equilibria that are not Nash equilibria

As the definition of Nash equilibria imposes more restrictions on the distribution, it follows that Nash equilibria are a subset of correlated equilibria. To represent Nash equilibria, the marginal distributions of x (for each player) are sufficient. The mixed Nash equilibrium in Table 3 implies marginal distributions of the form $x^1 = x^2 = \begin{pmatrix} \frac{1}{2} & \frac{1}{2} \end{pmatrix}^\top$. By contrast, such a representation is impossible for any correlated equilibrium that is not a Nash equilibrium (cf. Table 4). Now that terminology and the problem setting have been established, we go on to explain some essential techniques used in the proof.

2 Techniques for the Proof

A correlated equilibrium can be specified by means of linear inequalities (Definition 1.1). It is not too surprising that fundamental theorems from the field of linear programming play an important role in the proof. More surprising is the use of Markov chains for one of the subproblems in Papadimitriou's paper. We chose to cover the technical foundation of the proof at this point but the reader may also skip this section and follow our references in Section 3, if so inclined.

2.1 Linear Programming

Linear programming is a field of optimization theory that covers problems of the type described in Equation 2.1. A vector x is optimized (i.e., minimized

$$
\begin{aligned}
&\min\ c^\top x \\
\text{s.t.}\quad & Ax \geq b \\
& x \geq 0
\end{aligned}
\tag{2.1}
$$

Figure 1: The primal problem

$$
\begin{aligned}
&\max\ \pi^\top b \\
\text{s.t.}\quad & A^\top \pi \leq c \\
& \pi \geq 0
\end{aligned}
\tag{2.2}
$$

Figure 2: The dual problem

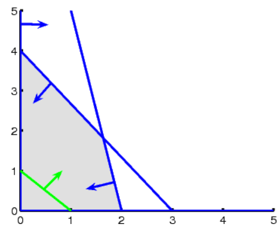

Figure 3: A problem with a feasible optimal solution

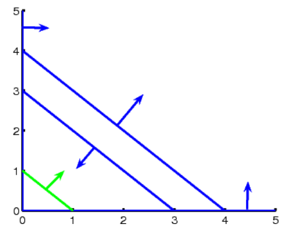

Figure 4: There is no feasible point

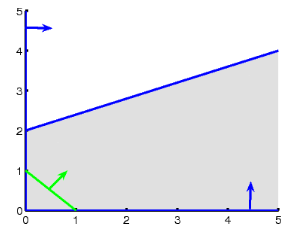

Figure 5: An unbounded linear problem

in this case) with respect to some linear cost function $c^\top x$ while x is restricted by constraints of the form $Ax \geq b$ and has to be non-negative. A fundamental result is that the primal problem in Equation 2.1 is equivalent to the so called dual problem in Equation 2.2. The equivalence of the two problems is based on the idea that linear combinations of constraints can be used to bound the optimal solution x^* for the primal problem from below. We refer to canonical literature about linear programming (e.g., [2]) for a deeper understanding of the mechanics involved. One essential feature of duality is the property of *strong duality* (Theorem 2.1) which states that the objective values of the primal and the dual optimal solutions are identical.

Theorem 2.1 *For optimal x^* and π^* it holds that $c^\top x^* = {\pi^*}^\top b$.*

Of course, not all linear programs have feasible optimal solutions. Constraints can be contradictory in such a way that no non-negative vector x can fulfill all constraints and thus the whole problem becomes *infeasible*. Another scenario ensues when the constraint matrix A allows for feasible solutions at arbitrarily low objective values (or arbitrarily large values in the case of maximization) and the problem is *unbounded*. Figures 3, 4 and 5 show examples of all three problem types discussed. There is a relationship between the characterization of the primal problem as infeasible, unbounded or having a finite optimum and the properties of the dual problem as summarized in Table 5.

Primal \ Dual	Fin. optimum	Unbounded	Infeasible
Fin. optimum	✓	-	-
Unbounded	-	-	✓
Infeasible	-	✓	✓

Table 5: Combinations of primal and dual problems that can occur together

2.2 Markov Chains

The proof of existence in Section 3 makes heavy use of duality but relies on certain properties of Markov chains to solve a very specific problem at one point. For this reason, we are going to give some basic definitions and properties of Markov chains. For the sake of simplicity, we are not going to present all definitions in the most general form but instead focus on the relevant properties.

Definition 2.1 *A Markov chain $\mathcal{C}$ is a stochastic process that fulfills the Markov property and is defined by a triple $< Q, T, p^0 >$ with a non-empty, finite set of states $Q = \{q_1, \dots, q_n\}$, an $n \times n$ transition matrix T such that*

$$\begin{array}{c} \sum_{k=1}^{n} T_{i,k} = 1 \\ 0 \leq T_{i,j} \leq 1 \end{array} \qquad \forall i, j \in \{1, \dots, n\}, \tag{2.3}$$

an initial distribution $p^0 = (p_1^0, \dots, p_n^0)^\top$ where $\sum_{i=1}^{n} p_i^0 = 1$.

Definition 2.2 *A Markov Chain $< Q, T, p^0 >$ fulfills the Markov property if for every state sequence $(X^t)_{t=0,1,\dots}$ ($X^t \in Q$) it holds that*

$$P(X^{t+1} = q_m | X^0, \dots, X^t) = P(X^{t+1} = q_m | X^t) \quad \forall t \in \mathbf{N}^0, q_m \in Q\ .$$

A Markov chain $\mathcal{C}$ as defined in Definitions 2.1 and 2.2 can be identified with a finite state machine where state transitions are probabilistic with probabilities defined in the transition matrix T. The quintessential feature of the Markov property is simply that the state of a Markov chain at time t only depends on the state at time $(t-1)$. In other words, the past and future of a Markov chain are independent from the point of view of a specific time point in the state sequence.

One can ask the question what can be inferred about the probability to be in a specific state after n steps, given an initial distribution p^0. And in fact, multiplying the distribution p^0 with the transition matrix T yields the new distribution after one step, or more generally $p^n = p^0 \cdot (T)^n$. Surprisingly, the sequence of state distributions $p^i, \quad i = 0, 1, \dots$ converges towards a steady-state distribution $p^* = \lim_{n\to\infty} \left(p^0 \cdot (T)^n\right)$ independently from p^0 under mild conditions. The notion of a steady-state implies two properties:

$$p^* = p^* \cdot T \tag{2.4}$$

$$p^*(q_i) \sum_{j=1}^{n} T_{i,j} = \sum_{j=1}^{n} p^*(q_j) \cdot T_{j,i} \qquad \forall i = \{1, \dots, n\} \tag{2.5}$$

It is Equation 2.5 that will be used in Section 3.

3 Proof of Existence

Papadimitriou's proof is based on transforming linear problems into other linear problems of similar form and heavily relying on the properties of duality. We begin with the problem (CE) which defines correlated equilibria in the form of a constraint program.

$$\begin{aligned} \sum_{s \in S_{-p}} (u^p_{is} - u^p_{js}) x_{is} &\geq 0 \\ \sum_{s \in S} x_s &= 1 \\ 0 \leq x &\leq 1 \end{aligned} \tag{CE}$$

A different formulation is

$$\begin{aligned} \max \quad & \sum_{s \in S} x_s \\ \text{s.t.} \quad & \sum_{s \in S_{-p}} (u^p_{is} - u^p_{js}) x_{is} \geq 0 \\ & x \geq 0 \end{aligned} \tag{P}$$

and the most significant difference between (P) and (CE) is that now an objective function has been added to the problem definition. Furthermore, even though x still has to be non-negative, the restriction that x must be a valid distribution has been removed. The following lemma describes the essential relationship between the two problems.

Lemma 3.1 *Problem* (CE) *has a solution if and only if* (P) *is unbounded.*

Proof 3.1 *"$\Rightarrow$": If* (CE) *has a solution x_{CE}, then x_{CE} is also feasible in Problem* (P) *but not optimal. That is because $\lambda \cdot x_{CE}$ for any $\lambda > 1$ is also feasible in* (P) *but has a larger objective value. It follows that* $\lim_{\lambda \to \infty} \sum (\lambda \cdot x_{CE}) = \infty$ *and* (P) *is unbounded.*
"$\Leftarrow$": If (P) *is unbounded, its set of solutions is, by definition, non-empty. An arbitrary solution $x_P \neq 0$ can be transformed into a solution x_{CE} for* (CE) *by applying a normalization that guarantees that x_{CE} is a valid distribution:*

$$x_{CE} = \frac{x_P}{\|x_P\|_1}$$

The constraints of problem (P) can be written in more compact form by expressing the terms $\sum_{s \in S_{-p}} (u^p_{is} - u^p_{js}) x_{is} \geq 0$ as a matrix multiplication $U \cdot x \geq 0$. Furthermore, the maximization of $\sum_{s \in S} x_s$ is equivalent to the minimization of $-\sum_{s \in S} x_s$ (cf. Figure 6). Equivalent to Equations (2.1) and (2.2), we can formulate the dual problem (D). As all constraints in the primal problem have a right-hand side of 0, the objective function of (D) is independent from the dual variables y and has a constant value of 0. Constraints in the dual problem are of the form "≥ -1" because all variables x in the primal problem have a coefficient of -1 in the objective function.

$$\begin{aligned} \min \quad & -\sum_{s \in S} x_s \\ \text{s.t.} \quad & U \cdot x \geq 0 \\ & x \geq 0 \end{aligned} \qquad \text{(P)}$$

Figure 6: The primal problem

$$\begin{aligned} \max \quad & 0 \\ \text{s.t.} \quad & U^\top \cdot y \leq -1 \\ & y \geq 0 \end{aligned} \qquad \text{(D)}$$

Figure 7: The dual problem

Our intention is to show that (CE) always has a solution. Lemma 3.1 establishes that it is sufficient to prove that (P) is always unbounded. This step will be done indirectly by reasoning about the dual problem (D). Table 5 indicates that (D) is infeasible whenever (P) is unbounded. However, (D) can also be infeasible if the primal problem (P) is infeasible. The latter case is irrelevant in our setting because (P) always has the trivial solution $x = 0$. It follows that proving the infeasibility of (D) is identical to the original claim.

Lemma 3.2 *Problem* (D) *is always infeasible.*

Proof 3.2 *We establish the lemma by proving that for every $y \geq 0$, there exists a product distribution x such that $x \cdot U^\top \cdot y = 0$. For the aspect of existence, it is not necessary to require x to be in product form. However, this property is essential for the development of a polynomial time algorithm in Papadimitriou's article and we adhere to it by defining $x_s = x^1_{s_1} \cdot \dots \cdot x^n_{s_n}$.*

Let us look at the coefficient of a particular reward u^p_{is} for player p, $s \in S_{-p}$ and $i \in S_p$ in the term $x \cdot U^\top \cdot y$:

$$\left[\prod_{q \neq p} x^q_{s_q} \right] \cdot \left[x^p_i \sum_{j \in S_p} y^p_{ji} - \sum_{j \in S_p} x^p_j y^p_{ij} \right] \tag{3.1}$$

If y is interpreted as the transition matrix of a Markov chain (after normalization) and x as a distribution over states, it is guaranteed that a valid x that represents the steady-state distribution of the Markov chain (Equation (2.5)*) can be found. But then $\left[\prod_{q \neq p} x^q_{s_q}\right] \cdot \left[x^p_i \sum_{j \in S_p} y^p_{ji} - \sum_{j \in S_p} x^p_j y^p_{ij}\right] = \prod_{q \neq p} x^q_{s_q} \cdot 0 = 0$.*

If an x *can always be found such that* $x \cdot U^\top \cdot y = 0$*, the condition* $U^\top \cdot y \geq -1$ *cannot hold if* x *is a distribution, i.e., has only non-negative entries. This is true because* $x \cdot U^\top \cdot y$ *represents a convex combination of the left-hand sides of the constraints in* (D)*, all of which are required to be less than or equal* -1*.*

From Lemma 3.2, that states the infeasibility of (D), it can be concluded that (P) is always unbounded (because it can never be infeasible) and that (CE) always has a solution. It follows that all multi-player games, as defined in Section 1.1, have at least one correlated equilibrium.

References

[1] J. Nash. Non-cooperative games. *The Annals of Mathematics*, 54(2):286–295, 1951.

[2] M. Padberg. *Linear Optimization and Extensions*. Springer, 1995.

[3] Christos H. Papadimitriou. Computing correlated equilibria in multi-player games. In *STOC '05: Proceedings of the thirty-seventh annual ACM symposium on Theory of computing*, pages 49–56, New York, NY, USA, 2005. ACM Press.