Existence and Computation of Correlated Equilibria

Andreas Hoenselaar

CS 244

January 11, 2006

Outline

Notation

Notation for games

- Players $1, \ldots, n$
- Set of strategies S_p for $1 \leq p \leq n$ (eg $\{\text{stop}, \text{go}\}$)
- Set of strategy profiles $S = \prod_{p=1}^{n} S_p$
- Payoff functions $u^p : S \longrightarrow \mathbf{Z}$

Notation for distributions

- Distribution on S: $x \in [0,1]^{|S|}$ with $\sum_{i=1}^{|S|} x_i = 1$
- Set of all possible distributions Δ
- Equivalently, distribution for player p: $x^p \in [0,1]^{|S_p|}$ with $\sum_{i=1}^{|S_p|} x_i^p = 1$

Definition of correlated equilibria

Definition

A distribution $x \in \Delta$ is a correlated equilibrium if the following is true for all players p and pairs of strategies (i, j): After drawing a strategy profile from x where player p's component is i, there is no incentive to play another strategy j and it holds that

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

Correlated equilibria vs. Nash equilibria

Nash equilibria are special cases of correlated equilibria.

Nash equilibria

- Distribution x on S is a product distribution:
  $x_s = \prod_{p=1}^{n} x^p_{s_p} \quad \forall s \in S$
- $\Rightarrow x$ is completely defined by its marginal distributions x^p

Correlated Equilibria

- As long as x is a valid distribution, it can have any form
- All $|S|$ entries are required for a full description

Intentions of this paper

Existence

- Prove that there exists at least one correlated equilibrium in every game
- Follows from the existence of Nash equilibria

Polynomial runtime

- Construct a polynomial runtime algorithm which computes correlated equilibria for a wide range of games

Primal and dual problem

Primal

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

Dual

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

Strong duality

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

Solutions to the primal and dual problem

Optimization results for linear programs

An LP can

- have an optimal feasible solution x^*
- be infeasible (contradicting constraints)
- be unbounded (ie there are feasible points with arbitrarily large objective value)

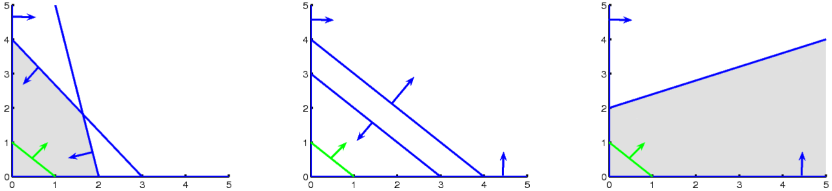

Solutions to the primal and dual problem

The relationship between primal and dual opt. result

Primal \ Dual	Fin. optimum	Unbounded	Infeasible
Fin. optimum	Y	N	N
Unbounded	N	N	Y
Infeasible	N	Y	Y

The Relationship between (CE) and (P)

(CE)

$$\text{Find } x = (x_1, \ldots, x_{|S|})$$
$$\text{s.t.} \quad Ux \geq 0$$
$$\sum_{s \in S} x_s = 1$$
$$0 \leq x \leq 1$$

(P)

$$\max \sum_{s \in S} x_s$$
$$\text{s.t.} \quad Ux \geq 0$$
$$x \geq 0$$

Lemma

Problem (P) is unbounded iff (CE) has a solution.

(D) is always infeasible

Lemma

For every $y \geq 0$ there is a product distribution x such that $xU^\top y = 0$.

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

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

Definition of a Markov Chain

Definition

A Markov chain $\mathcal{C}$ is a stochastic process defined by a tripel $< Q, T, p^0 >$ with

- a non-empty, finite set of states $Q = \{q_1, \ldots, q_n\}$
- a $n \times n$ transition matrix T where

$$\begin{array}{ll} \sum_{k=1}^{n} T_{i,k} = 1 & \\ 0 \leq T_{i,j} \leq 1 & \end{array} \qquad \forall i,j \in \{1, \ldots, n\}$$

- an initial distribution $p^0 = (p_1^0, \ldots, p_n^0)^\top$ such that $\sum_{i=1}^{n} p_i^0 = 1$

that fulfills the Markov property.

Definition of a Markov Chain (2)

The Markov property

For every state sequence $(X^t)_{t=0,1,\ldots}$ ($X^t \in Q$) it holds that

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

This implies that the state at any given point in time only depends on the previous state and the transition matrix.

Example

$$Q = \{q_1, q_2, q_3\}, \quad T = \begin{pmatrix} 0.3 & 0.2 & 0.5 \\ 0.4 & 0 & 0.6 \\ 0.1 & 0.9 & 0 \end{pmatrix}, \quad p^0 = \begin{pmatrix} \frac{1}{3} & \frac{1}{3} & \frac{1}{3} \end{pmatrix}$$

Steady-state of a Markov Chain

Motivation

What can we say about the probabilty to be in a specific state q_m after n steps?

Computation of the steady-state distribution

Given an initial state distribution p^0 we can calculate the state distribution after n steps:

$$p^n = p^0 \cdot (T)^n$$

Surprisingly, this sequence converges towards an equilbirium distribution p^* in the limit of infinitely many steps:

$$\lim_{n \to \infty} \left(p^0 \cdot (T)^n\right) = p^*$$

Steady-state of a Markov Chain

Properties of the steady-state distribution

$$p^* = \lim_{n\to\infty} \left(p^0 \cdot (T)^n\right)$$

- Convergence is independent from p^0
- For p^* it holds that

$$p^* = p^* \cdot T$$

- Can be interpreted as an eigenvalue problem: p^* is the eigenvector corresponding to the eigenvalue 1.
- Explicit formula on a per-state basis:

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

The need for a succinct representation

Exponential size of distributions

A distribution x over the set of strategy profiles S has an exponential number of entries:

$$|x| = |S| = \prod_{i=1}^{n} |S_i| \in \mathcal{O}(s^n) \quad \text{where } s = \max\left(|S_1|, \ldots, |S_n|\right)$$

A first insight

If a correlated equilibrium is to be calculated in polynomial time, the game and thus the set of strategy profiles must have a more succinct representation.

Definitions for Succinct Games

Definition

A succinct game $G = (I, T, U)$ is defined by

- a set of inputs I
- an algorithm T that returns
  - the number of players n
  - the cardinalities of the strategy sets $(t_1, \ldots, t_n)$

  in polynomial time
- an algorithm U that returns the utility $u^i(s)$ for player i under a distribution $s = (s_1, \ldots, s_n)$

The game is of *polynomial type* if n and all t_i's are polynomially bounded in their argument.

Definitions for Succinct Games

Definition

A game G has the *polynomial expectation property* if there is a polynomial algorithm E which, given $z \in I$, $p \leq n$ and a product distribution x_s over S, returns

$$E(z, p, x_s) = \mathbf{E}_{x_s}\left[u_s^p : s \in S\right]$$

in polynomial time.

Remark

The algorithm described in the paper can calculate correlated equilibria only for games that have the polynomial expectation property.

Types of succinct games (1)

Symmetric Games

In symmetric games, players cannot be distinguished and have the same set of strategies:

- The payoff depends on a player's choice and the distribution of the other players' choices over the strategy set
- Requires $\mathcal{O}(n^s)$ space instead of $\mathcal{O}(s^n)$ for its description

Graphical games

in graphical games, players are identified with nodes in a graph and games are played only along the (undirected) edges.

- Description length: $\mathcal{O}(n \cdot s^k)$ for maximum node degree k plus explicit representation of g_p, the game with all players in the neighborhood
- Correlated equilibria can be calculated in polynomial time in graphical games

Polynomial expectation property

- Calculate expectation by iterating over all stratagy profiles in g_p, ignore players not in the neighborhood

Hypergraphical games

Several subsets of players play separate games and the utilities are added. Each player p has the same set of strategies S_p in all games he is involved in.

Polynomial expectation property

- Description length: z describes a hypergraph $H = ([n], E)$ and for each hyperedge $h \in E$ and explicit game that involves all players in h.

s omitted

Congestion Games

The strategies in a congestion game are subsets of a set of resources R.

- Using resources in high demand is penalized via the *delay* (= number of players who use the resource)
- Payoff is equivalent to the negative sum of resource delays in the player's strategy