Adaptive Discretization Methods for Computing Correlated Equilibria of Polynomial Games

Noah Stein

Joint work with Profs. Pablo Parrilo and

Asuman Ozdaglar

Game theoretic setting

- Standard strategic (normal) form game
- Players (rational agents) numbered $i = 1, \ldots, n$
- Each has a set C_i of strategies s_i
- Players choose their strategies simultaneously
- Rationality: Each player seeks to maximize his own utility function $u_i : C \to \mathbb{R}$, which represents all his preferences over outcomes

Chicken and correlated equilibria

(u_1, u_2)	Wimpy	Macho
Wimpy	$(4, 4)$	$(1, 5)$
Macho	$(5, 1)$	$(0, 0)$

- Nash equilibria (self-enforcing independent distrib.)
  - (M, W)
  - (W, M)
- Correlated equilibria (self-enforcing joint distrib.)
  - $\frac{1}{2}(W, M) + \frac{1}{2}(M, W)$
  - $\frac{1}{3}(W, W) + \frac{1}{3}(W, M) + \frac{1}{3}(M, W)$

Correlated equilibria in games with finite strategy sets

- $u_i(t_i, s_{-i}) - u_i(s)$ is change in player i's utility when strategy t_i replaces s_i in $s = (s_1, \ldots, s_n)$
- A probability distribution π is a **correlated equilibrium** if

$$\sum_{s \in \{r_i\} \times C_{-i}} [u_i(t_i, s_{-i}) - u_i(s)] \pi(s | s_i = r_i) \leq 0$$

  for all players i and all strategies $r_i, t_i \in C_i$

- No player has an incentive to deviate from his recommended strategy r_i

LP characterization

- A probability distribution π is a correlated equilibrium if and only if

$$\sum_{s_{-i} \in C_{-i}} [u_i(t_i, s_{-i}) - u_i(s)]\pi(s) \leq 0$$

for all players i and all strategies $s_i, t_i \in C_i$

- Set of correlated equilibria of a finite game is a polytope

Polynomial games

- Strategy space is $C_i = [-1, 1]$ for all players i
- Utilities u_i are multivariate polynomials
- Finitely supported equilibria always exist

Finitely supported ϵ-correlated equilibria

- A probability measure π with finite support contained in $\tilde{C} = \prod \tilde{C}_i$ is an ϵ-correlated equilibrium if

$$\sum_{s_{-i} \in \tilde{C}_{-i}} [u_i(t_i, s_{-i}) - u_i(s)]\pi(s) \leq \epsilon_{i,s_i}$$

for all i, $s_i \in \tilde{C}_i$, and $t_i \in C_i$ and

$$\sum_{s_i \in \tilde{C}_i} \epsilon_{i,s_i} \leq \epsilon$$

for all i.

Adaptive discretization

- Given $\tilde{C}_i^k$, optimize the following (as an SDP)

$$\begin{aligned} \min \quad & \epsilon \\ \text{s.t.} \quad & \pi \quad \text{is an } \epsilon\text{-correlated equilibrium} \\ & \qquad \text{which is a correlated equilibrium} \\ & \qquad \text{when deviations are restricted to } \tilde{C}^k \end{aligned}$$

- Let ϵ^k and π^k be an optimal solution
- If $\epsilon^k = 0$ then halt
- Otherwise, compute $\tilde{C}_i^{k+1}$ (next slide) and repeat
- Convergence theorem: $\epsilon^k \to 0$

Adaptive discretization (II)

- Steps to compute $\tilde{C}^{k+1}$
  - For some player i, the ϵ-correlated equilibrium constraints are tight
  - Find values of t_i making these tight (free with SDP duality), add these into $\tilde{C}_i^k$ to get $\tilde{C}_i^{k+1}$
  - For $j \neq i$, let $\tilde{C}_j^{k+1} = \tilde{C}_j^k$

Applying SOS / SDP

- Given a polynomial game and a finite support set $\tilde{C}_i \subset [-1, 1]$ for each player, the condition that π be a probability measure on $\tilde{C}$ and an ϵ-correlated equilibrium can be written in an SDP

- First constraint says a univariate polynomial in t_i with coefficients linear in the $\pi(s)$ and ϵ_{i,s_i} is ≥ 0 on $[-1, 1]$, hence is expressible exactly in an SDP

- Remaining constraints are linear, so usable in SDP

- Always feasible since ϵ can vary and finite games have correlated equilibria

Random example

- Three players, random polynomial utilities (deg. 4)

k	ϵ^k	$\tilde{C}_x^k \setminus \tilde{C}_x^{k-1}$	$\tilde{C}_y^k \setminus \tilde{C}_y^{k-1}$	$\tilde{C}_z^k \setminus \tilde{C}_z^{k-1}$
0	0.99	$\{0\}$	$\{0\}$	$\{0\}$
1	4.16			$\{0.89\}$
2	5.76	$\{-1\}$		
3	0.57		$\{1\}$	
4	0.28	$\{0.53\}$		$\{0.50, 0.63\}$
5	0.16		$\{0.49, 0.70\}$	
6	10^{-7}		$\{-1, 0.60\}$	$\{-0.60, 0.47\}$

Closing remarks

- Can you remove the condition that π is an exact correlated equilibrium when deviations are restricted to $\tilde{C}_i^k$ from the optimization problem?
  - This seems to work well in practice
  - We have explicit counterexamples showing it doesn't work in general
- For a completely different approach to computing correlated equilibria of polynomial games that does not use discretization, see [SPO 2007]