Adaptive Discretization Methods for Computing Correlated Equilibria of Polynomial Games

Noah Stein

Joint work with Profs. Pablo Parrilo and

Asuman Ozdaglar

Outline

- Intro to correlated equilibria
- Polynomial games
- Approximating correlated equilibria in poly games
  - Three (two?) discretization methods
- Example
- Implementation

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

- A probability distribution π is a **correlated equilibrium** if

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

  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

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

  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
- Many nice properties
  - Finitely supported equilibria exist; bounds on support size [1950s, SOP 2006]
  - Minimax strategies and values can be computed by semidefinite programming [Parrilo 2006]
  - Can compute outer approximations to set of correlated equilibria by SDP [SOP 2007]

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

$$\sum_{s_{-i} \in \tilde{C}_{-i}} [u_i(s) - u_i(t_i, s_{-i})]\pi(s) + \epsilon_{i,s_i} \geq 0$$

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.

Method I: Static discretization

- Intended as a benchmark to judge other techniques
- Ignore polynomial structure
- Restrict strategy choices (and deviations) to fixed finite sets $\tilde{C}_i \subset C_i$
- Compute exact correlated equilibria of approximate game
- This is a sequence of LPs which converges (slowly!) to the set of correlated equilibria as the discretization gets finer.

Method II: Adaptive discretization attempt (a)

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

$$\begin{aligned} \min \quad & \epsilon \\ \text{s.t.} \quad & \pi \quad \text{is an } \epsilon\text{-correlated equilibrium} \\ & \qquad\qquad \text{supported on } \tilde{C}^k \end{aligned}$$

- Let ϵ^k and π^k be optimal (we're done if $\epsilon^k = 0$)
- Otherwise, compute $\tilde{C}_i^{k+1}$ (next slide) and repeat

Method II: Adaptive discretization attempt (b)

- 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 $\tilde{C}_i^{k+1}$
  - There are finitely many such values by polynomiality
  - For $j \neq i$, let $\tilde{C}_j^{k+1} = \tilde{C}_j^k$
- Intuitively, this adds "good" strategies for player i

Method II: Adaptive discretization attempt (c)

- This often works in practice, but $\epsilon^k \nrightarrow 0$ in general
- Consider the symmetric game with identical utilities

	a	b	c
a	0	1	0
b	1	5	7
c	0	7	0

- If $\tilde{C}_1^0 = \tilde{C}_2^0 = \{a\}$ then $\tilde{C}_i^k = \{a, b\}$ and $\epsilon^k = 1$ for all $k \geq 1$

Method III: Adaptive discretization

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

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

- Let ϵ^k and π^k be optimal (we're done if $\epsilon^k = 0$)
- Compute $\tilde{C}_i^{k+1}$ (same way as above) and repeat
- Convergence theorem: $\epsilon^k \to 0$

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}		$\{-1, 0.60\}$	$\{-0.60, 0.47\}$

Need to solve optimization problems with

- Finitely many decision variables
- Linear objective
- Linear equations and inequalities
- Constraints of the form $p(t) \geq 0$ for all $t \in [-1, 1]$ where $p(t)$ is a univariate polynomial in t whose coefficients are linear in the decision variables

Semidefinite programming

- A **semidefinite program (SDP)** is an optimization problem of the form

$$\begin{aligned} \min \quad & L(S) && \leftarrow L \text{ is a given linear functional} \\ \text{s.t.} \quad & T(S) = v && \leftarrow T \text{ is a given linear transformation,} \\ & && \quad v \text{ is a given vector} \\ & S \succeq 0 && \leftarrow S \text{ is a symmetric matrix} \\ & && \quad \text{of decision variables} \end{aligned}$$

- SDPs generalize linear programs and can be solved efficiently using interior point methods

Sums of squares + SDP

- A polynomial $p(t)$ is ≥ 0 for all $t \in [-1, 1]$ iff there are polynomials $q_k(t), r_k(t)$ such that

$$p(t) \equiv \sum_k q_k^2(t) + (1 - t^2) \sum_k r_k^2(t)$$

- Coefficients of polynomials of this form can be described in an SDP

Closing remarks (a)

- For a completely different approach to computing correlated equilibria of polynomial games that does not use discretization, see [SPO 2007]
- We did not use the polynomial structure of the u_i in the convergence proof, just continuity
- Used polynomiality to convert the optimization problem into an SDP

Closing remarks (b)

- Can also do this conversion if the u_i are rational or even piecewise rational (and continuous)
- Solutions of such games are surprisingly complex – the Cantor measure arises as the unique Nash equilibrium of a game with rational u_i [Gross 1952]
- Now we have a way to approximate these algorithmically!
- Open questions
  - Convergence rate
  - Optimization over the set of correlated equilibria