Simple Approximate Equilibria in Large Games

Yakov Babichenko, Siddharth Barman, Ron Peretz

15th ACM Conference on Economics and Computation (EC'14) 753–770.

Speaker: Joseph Chuang-Chieh Lin

Institute of Information Science
Academia Sinica
Taiwan

21 November 2014

Outline

Simplicity of the solution concepts

- Equilibria: central solution concepts in the theory of strategic games.
  - Nash equilibrium [Nash 1951];
  - correlated equilibrium [Aumann 1974];
  - coarse correlated equilibrium [Hanna 1957].
  - ⋆ No player can benefit by unilateral deviation.
- A solution concept is *too complicated* $\Rightarrow$ debatable applicability.
  - Concerns in bounded rationality.

A natural notion of simplicity

- ϵ-approximate equilibrium: a distribution over action files where no player has more than an ϵ incentive to deviate.

A natural notion of simplicity (Lipton *et al.* 2003)

An approximate equilibrium is simple if the equilibrium is a *uniform distribution* on a set of small size.

Contribution of this paper

- Establish the existence of simple approximate Nash, correlated, and coarse correlated equilibria in large games.
  - Simple: uniform distribution over *multisets of small size*.
  - n players and m actions per player.
- Improve the running time of previously best known algorithm for computing a small-support approximate Nash equilibrium in large games.
- Prove that finding an *exact* correlated equilibrium with smallest possible support is NP-hard.

Contribution of this paper (contd.)

ϵ-approx. Equilibrium	support-size upper bound
Nash	$O\left(\frac{\log n + \log m - \log \epsilon}{\epsilon^2}\right)$
Correlated	$O\left(\frac{\log m(\log n + \log m - \log \epsilon)}{\epsilon^4}\right)$
Coarse Correlated	$O\left(\frac{\log n + \log m}{\epsilon^2}\right)$

Support size:

- For Nash: # strategies that are played with positive probability of each player.
- For the other two: # action profiles that are played with positive probability.

Preliminaries

- $[n] = \{1, 2, \ldots, n\}$: the set of players.
- $A_i = [m] = \{1, 2, \ldots, m\}$: the set of actions for player i.
  - $A = [m]^n$: the set of action profiles.
- $N := nm^n$: the size of the game.
- (a_i, a_{-i}): an action profile;
  - a_i: the action of player i.
  - a_{-i}: the actions chosen by players $[n] \setminus \{i\}$.
- $\Delta(B)$: the set of probability distributions over a set B.
- $u_i : A \mapsto [0, 1]$ the payoff function of player i.
  - $u = (u_i)_{i \in [n]}$: the payoff function profile.

⋆ $u_i : \Delta(A) \mapsto [0, 1]$.
  - For $x \in \Delta(A)$, denote by $u_i(x)$ the expected payoff of player i under x.

ϵ-Nash equilibrium

ϵ-Nash equilibrium (ϵ-NE)

A mixed action profile $x = (x_i)_{i \in [n]}$, where $x_i \in \Delta(A_i)$, is an ϵ-Nash equilibrium if

$$\forall i \in [n], \; \forall a_i \in [m], \quad u_i(x) \geq u_i(a_i, x_{-i}) - \epsilon.$$

- product distributions

ϵ-coarse correlated equilibrium

$R^i_j(a) := u_i(j, a_{-1}) - u_i(a)$

- the regret of player i for not playing j at profile a.

ϵ-coarse correlated equilibrium (ϵ-CCE)

$x \in \Delta(A)$ is an ϵ-coarse correlated equilibrium if

$$\forall i \in [n],\ \forall j \in A_i,\quad \mathbf{E}_{a \sim x}[R^i_j(a)] \leq \epsilon.$$

- general (not necessarily product) distributions

ϵ-correlated equilibrium

$R^i_f(a) := u_i(f(a_i), a_{-1}) - u_i(a)$, where $f : A_i \mapsto A_i$ is a switching rule.

- the regret of player i for not implementing f at profile a.

ϵ-correlated equilibrium (ϵ-CE)

$x \in \Delta(A)$ is an ϵ-correlated equilibrium if

$$\forall i \in [n],\ \forall f : A_i \mapsto A_i,\ \ \mathbf{E}_{a \sim x}[R^i_f(a)] \le \epsilon.$$

- general (not necessarily product) distributions

An example: rock-paper-scissors game

	rock	paper	scissors
rock	0, 0	0, 1	1, 0
paper	1, 0	0, 0	0, 1
scissors	0, 1	1, 0	0, 0

k-uniform strategy & distribution

- $x_i \in \Delta(A_i)$ is called a k-uniform strategy if it is a uniform distribution over a multiset of k pure actions from A_i.
  - $x = (x_i)_{i \in [n]}$ is called k-uniform if every x_i is k-uniform.
- $x \in \Delta(A)$ is called k-uniform if it is the uniform distribution over a size-k multiset of action profiles from A.

Part I: Nash Equilibrium

The main theorem

Theorem 3.1

Every n-players m-actions game admits a k-uniform ϵ-NE for every

$$k > \frac{8(\ln m + \ln n - \ln \epsilon + \ln 8)}{\epsilon^2}.$$

Previously best results:

- Two-player games: $O(\log m)$ [Althöfer 1994].
- n-player games: $O(n \log m)$ [Hémon *et al.* 2008].

Time complexity of computing ϵ-Nash equilibria

Corollary 3.2

Let $m = \text{poly}(n)$, $N = nm^n$. For any constant $\epsilon > 0$, there exists an algorithm computing an ϵ-NE of the game in $O(\text{poly}(N^{\log\log N}))$ time.

- $\#$ all possible k-uniform profiles $\leq (\binom{m+k-1}{k})^n \approx m^{nk} = \text{poly}(m^{n \log n})$ $= \text{poly}(m^n)^{\log\log(m^n)} = \text{poly}(N^{\log\log N})$.
- Previous result: $O(\text{poly}(N^{\log N}))$ [Lipton *et al.* 2003].

Corollary 3.3

Let $m = O(1)$, $N = nm^n$. For any constant $\epsilon > 0$, there exists an algorithm computing an ϵ-NE of the game in $O(\text{poly}(N^{\log\log\log N}))$ time.

- Previous result: $O(\text{poly}(N^{\log\log N}))$ [Daskalakis & Papadimitriou 2008].

The fundamental lemma

- Assume that players are playing according to $x = (x_i)_{i \in [n]}$.
- Observe k i.i.d. samples from $x \Rightarrow (a(t))_{t \in [k]}$, $a(t) \in A$.
- s_i^k: the empirical distribution of player i.
  - $s_i^k(\alpha) = \frac{1}{k}|\{t : a_i(t) = \alpha\}|$.
- $s^k = \prod_i s_i^k$, $s_{-i}^k = \prod_{j \neq i} s_j^k$.

Lemma 3.4

For every n-players m-actions game, every player $i \in [n]$, every action $a_i \in A_i = [m]$, and every product distribution $x_{-i} = (x_j)_{j \neq i}$, we have

$$\Pr\left(|u_i(a_i, s_{-i}^k) - u_i(a_i, x_{-i})| \geq \epsilon\right) \leq \frac{4e^{-\frac{\epsilon^2}{2}k}}{\epsilon}.$$

Proof of Lemma 3.4

WLOG, assume $i = 1$, $a_i = 1$.

- For every $\ell \in [k]$, we rewrite the payoff of player 1

$$u_1(1, s_{-1}^k) = \frac{1}{k^{n-1}} \sum_{j_2, j_3, \ldots, j_n \in [k]} u_1(1, a_2(j_2 + \ell), a_3(j_3 + \ell), \ldots, a_n(j_n + \ell)),$$

  where $j_i + \ell$ are taken modulo k.

- Further,

$$u_1(1, s_{-1}^k) = \frac{1}{k^{n-1}} \sum_{j_2, j_3, \ldots, j_n \in [k]} \frac{1}{k} \sum_{\ell \in [k]} u_1(1, a_2(j_2 + \ell), a_3(j_3 + \ell), \ldots, a_n(j_n + \ell)).$$

Proof of Lemma 3.4 (contd.)

$$u_1(1, s_{-1}^k) = \frac{1}{k^{n-1}} \sum_{j_2, j_3, \ldots, j_n \in [k]} \frac{1}{k} \sum_{\ell \in [k]} u_1(1, a_2(j_2 + \ell), a_3(j_3 + \ell), \ldots, a_n(j_n + \ell)).$$

- $a_{-1}(j_* + \ell) := (a_2(j_2 + \ell), a_3(j_3 + \ell), \ldots, a_n(j_n + \ell))$.
- Define:

$$d(j_*) = \begin{cases} 0, & \text{if } \left| \frac{1}{k} \sum_{\ell \in [k]} u_1(1, a_{-1}(j_* + \ell)) - u_1(1, x_{-1}) \right| \leq \frac{\epsilon}{2} \\ 1, & \text{otherwise.} \end{cases}$$

$\Rightarrow$

$$d(j_*) + \frac{\epsilon}{2} \geq \left| \frac{1}{k} \sum_{\ell \in [k]} u_1(1, a_{-1}(j_* + \ell)) - u_1(1, x_{-1}) \right|.$$

- $\mathbf{E}[d(j_*)] \leq 2e^{-\frac{\epsilon^2}{2}k}$ by Hoeffding's inequality.

Proof of Lemma 3.4 (contd.)

Then we have

$$
\begin{aligned}
& \Pr(|u_i(1, s_{-1}^k) - u_i(1, x_{-1})| \geq \epsilon) \\
= \; & \Pr\left(\left| \frac{1}{k^{n-1}} \sum_{j_* \in [k]^{n-1}} \frac{1}{k} \sum_{\ell \in [k]} u_1(1, a_{-1}(j_* + \ell)) - u_1(1, x_{-1}) \right| \geq \epsilon \right) \\
\leq \; & \Pr\left(\frac{1}{k^{n-1}} \sum_{j_* \in [k]^{n-1}} \left| \frac{1}{k} \sum_{\ell \in [k]} u_1(1, a_{-1}(j_* + \ell)) - u_1(1, x_{-1}) \right| \geq \epsilon \right) \\
\leq \; & \Pr\left(\frac{1}{k^{n-1}} \sum_{j_* \in [k]^{n-1}} d(j_*) \geq \frac{\epsilon}{2} \right) \\
\leq \; & \frac{4e^{-\frac{\epsilon^2}{2}k}}{\epsilon} \quad \text{(Markov's inequality).}
\end{aligned}
$$

Let's come back to Theorem 3.1.

Theorem 3.1

Every n-players m-actions game admits a k-uniform ϵ-Nash equilibrium for every

$$k > \frac{8(\ln m + \ln n - \ln \epsilon + \ln 8)}{\epsilon^2}.$$

Proof of Theorem 3.1

Let

- $x = (x_i)_{i \in [n]}$: a Nash equilibrium of the game
- s^k: the product empirical distribution of play w.r.t. x

By Lemma 3.4 and the choice of k we have $(\forall i, \forall a_i)$:

$$\Pr\left(|u_i(a_i, s^k_{-i}) - u_i(a_i, x_{-i})| \geq \frac{\epsilon}{2}\right) \leq \frac{8e^{-\frac{\epsilon^2}{8}k}}{\epsilon} < \frac{1}{2mn}.$$

So we have $|u_i(a_i, s^k_{-i}) - u_i(a_i, x_{-i})| < \frac{\epsilon}{2}$ for all i and all a_i with probability $> 1/2$ (union bound).

Proof of Theorem 3.1 (contd.)

$(s_i^k)_{i\in[n]}$ is an ϵ-Nash equilibrium:

$$\begin{aligned}
u_i(a_i, s_{-i}^k) &\leq u_i(a_i, x_{-i}) + \frac{\epsilon}{2} \\
&\leq \sum_{a_i'\in A_i} s_i^k(a_i')u_i(a_i', x_{-i}) + \frac{\epsilon}{2} \\
&\leq \sum_{a_i'\in A_i} s_i^k(a_i')u_i(a_i', s_{-i}^k) + \epsilon \\
&= u_i(s_i^k, s_{-i}^k) + \epsilon.
\end{aligned}$$

k-uniform random sampling algorithm (k-URS)

- k-uniform random sampling algorithm (k-URS):
  - Sample uniformly at random n-tuples of k-uniform strategy profiles.
  - Check whether the profile forms an ϵ-Nash equilibrium.
- poly(m) samples for *small-probability games* [Daskalakis & Papadimitriou 2009].
  - Admitting a Nash equilibrium where each pure action is played with prob. $\leq O(1/m)$.

High entropy helps finding an ϵ-NE

Theorem 3.5

Let

- g: an n-players m-actions game with a NE $x = (x_i)_{i \in [n]}$.
- $k \geq \max\{\frac{16}{\epsilon^2}(\ln n + \ln m - \ln \epsilon + 2), \epsilon^{16/\epsilon^2}\} = O(\log n + \log m)$.
- $H(x) = \sum_{i \in [n]} H(x_i)$: Shannon's entropy of x.

Then the k-URS algorithm finds an ϵ-NE after $\leq 4 \cdot 2^{k(n \lg m - H(x))}$ samples in expectation.

Corollary 3.6

Families of games where $n \lg m - \max_{x \in NE} H(x)$ is bounded and $k = O(\log n + \log m)$ admit a $poly(m, n)$ randomized algorithm for computing an ϵ-NE.

A useful lemma for proving Theorem 3.5

Lemma 3.9

Let y be a random variable that assumes values in a finite set M.
Let $S \subset M$ s.t. $\Pr(y \in S) \geq 1 - \frac{1}{\lg |M|}$.

Then $|S| \geq \frac{1}{4} \cdot 2^{H(y)}$.

- $$\begin{aligned} H(y) &= \Pr(y \in S) \cdot H(y \mid y \in S) + \Pr(y \notin S) \cdot H(y \mid y \notin S) + H(\mathbb{1}_{\{y \in S\}}) \\ &\leq \lg |S| + \Pr(y \notin S) \cdot \lg |M| + 1 \leq \lg |S| + 2. \end{aligned}$$

Proof of Theorem 3.5

- $k \geq \max\{\frac{16}{\epsilon^2}(\ln n + \ln m - \ln \epsilon + 2), e^{16/\epsilon^2}\}$

$$\Rightarrow \quad \Pr\left(|u_i(a_i, s^k_{-i}) - u_i(a_i, x_{-i})| \geq \frac{\epsilon}{2}\right) \leq \frac{8e^{-\frac{k\epsilon^2}{8}}}{\epsilon} \leq \frac{1}{mn} \cdot \frac{1}{nk \lg m}.$$

- Thus, the k-URS algorithm finds an ϵ-NE with prob. $\geq 1 - \frac{1}{nk \lg m} = 1 - \frac{1}{\lg(m^{nk})}$.
- The fraction of the k-uniform strategy profiles that form an ϵ-NE:

$$\geq \frac{\frac{1}{4}2^{kH(x)}}{m^{nk}} = \frac{1}{4}2^{k(H(x) - n \lg m)}.$$

- Expected time: $\leq 4 \cdot 2^{k(n \lg m - H(x))}$.

Part II: Coarse Correlated Equilibrium

The existence of small-support ϵ-CCE

Theorem 4.1

Every n-players m-actions game admits a k-uniform ϵ-coarse correlated equilibrium for every

$$k > \frac{2(\ln n + \ln m)}{\epsilon^2}.$$

- Let $\sigma \in \Delta(A)$ be an *exact coarse correlated equilibrium* of the game.
  - $\mathbf{E}_{a \sim \sigma}[R^i_j(a)] \leq 0, \forall i, \forall j \in A_i$.
- Sample k action files: $a(1), a(2), \ldots, a(k)$ independently at random according to σ.
- Denote by s the uniform distribution over $\{a(1), a(2), \ldots, a(k)\}$.

The existence of small-support ϵ-CCE

Theorem 4.1

Every n-players m-actions game admits a k-uniform ϵ-coarse correlated equilibrium for every

$$k > \frac{2(\ln n + \ln m)}{\epsilon^2}.$$

- Let $\sigma \in \Delta(A)$ be an *exact coarse correlated equilibrium* of the game.
  - $\mathbf{E}_{a\sim\sigma}[R^i_j(a)] \leq 0, \forall i, \forall j \in A_i$.
- Sample k action files: $a(1), a(2), \ldots, a(k)$ independently at random according to σ.
- Denote by s the uniform distribution over $\{a(1), a(2), \ldots, a(k)\}$.

Proof of Theorem 4.1 (contd.)

$$\Pr_{a(1),a(2),\ldots,a(k)\sim\sigma}(\mathbf{E}_{a\sim s}[R_j^i(a)] \geq \epsilon) = \Pr\left(\frac{1}{k}\sum_{\ell\in[k]} R_j^i(a(\ell)) \geq \epsilon\right)$$
$$\leq e^{-\frac{k\epsilon^2}{2}}. \text{ (Hoeffding's inequality)}$$

- $\mathcal{E}_{ij}$: the event that $\mathbf{E}_{a\sim s}[R_j^i(a)] \geq \epsilon$.
- $k > \frac{2(\log n + \log m)}{\epsilon^2} \Rightarrow \Pr[\mathcal{E}_{ij}] < \frac{1}{nm}$.
- Pr[None of $\mathcal{E}_{ij}$'s happens] > 0.

Complexity of computing an ϵ-CCE

Proposition 4.2

There exists a poly(n, m) time randomized algorithm for computing a k-uniform ϵ-coarse correlated equilibrium for

$$k > \frac{2(\ln n + \ln m + \ln 2)}{\epsilon^2}.$$

- $k > \frac{2(\ln n + \ln m + \ln 2)}{\epsilon^2} \Rightarrow \Pr(\mathbf{E}_{a\sim s}[R^i_j(a)] \geq \epsilon) < \frac{1}{2nm}$.
  - The empirical distribution s is an ϵ-CCE with prob. $\geq 1/2$.
- Find an exact coarse correlated equilibrium $\sigma \in \Delta(A)$ first!
  - Using the polynomial time algorithm in [Jiang & Leyton-Brown 2013].
- Start from σ, sample k action profiles according to σ and make it uniformly distributed, and then check (twice iterations required in expectation).

Complexity of computing an ϵ-CCE

Proposition 4.2

There exists a poly(n, m) time randomized algorithm for computing a k-uniform ϵ-coarse correlated equilibrium for

$$k > \frac{2(\ln n + \ln m + \ln 2)}{\epsilon^2}.$$

- $k > \frac{2(\ln n + \ln m + \ln 2)}{\epsilon^2} \Rightarrow \Pr(\mathbf{E}_{a \sim s}[R_j^i(a)] \geq \epsilon) < \frac{1}{2nm}$.
  - The empirical distribution s is an ϵ-CCE with prob. $\geq 1/2$.
- Find an exact coarse correlated equilibrium $\sigma \in \Delta(A)$ first!
  - Using the polynomial time algorithm in [Jiang & Leyton-Brown 2013].
- Start from σ, sample k action profiles according to σ and make it uniformly distributed, and then check (twice iterations required in expectation).

Complexity of computing an ϵ-CCE

Proposition 4.2

There exists a poly(n, m) time randomized algorithm for computing a k-uniform ϵ-coarse correlated equilibrium for

$$k > \frac{2(\ln n + \ln m + \ln 2)}{\epsilon^2}.$$

- $k > \frac{2(\ln n+\ln m+\ln 2)}{\epsilon^2} \Rightarrow \Pr(\mathbf{E}_{a\sim s}[R^i_j(a)] \geq \epsilon) < \frac{1}{2nm}$.
  - The empirical distribution s is an ϵ-CCE with prob. $\geq 1/2$.
- Find an exact coarse correlated equilibrium $\sigma \in \Delta(A)$ first!
  - Using the polynomial time algorithm in [Jiang & Leyton-Brown 2013].
- Start from σ, sample k action profiles according to σ and make it uniformly distributed, and then check (twice iterations required in expectation).

Complexity of computing an ϵ-CCE

Proposition 4.2

There exists a poly(n, m) time randomized algorithm for computing a k-uniform ϵ-coarse correlated equilibrium for

$$k > \frac{2(\ln n + \ln m + \ln 2)}{\epsilon^2}.$$

- $k > \frac{2(\ln n + \ln m + \ln 2)}{\epsilon^2} \Rightarrow \Pr(\mathbf{E}_{a \sim s}[R_j^i(a)] \geq \epsilon) < \frac{1}{2nm}$.
  - The empirical distribution s is an ϵ-CCE with prob. $\geq 1/2$.
- Find an exact coarse correlated equilibrium $\sigma \in \Delta(A)$ first!
  - Using the polynomial time algorithm in [Jiang & Leyton-Brown 2013].
- Start from σ, sample k action profiles according to σ and make it uniformly distributed, and then check (twice iterations required in expectation).

Part III: Correlated Equilibrium

The existence of an ϵ-CE with polylogarithmic support

Theorem 5.1

Every n-players m-actions game admits a k-uniform ϵ-correlated equilibrium for every

$$k > \frac{264 \ln m(\ln n + \ln m - \ln \epsilon + \ln 16)}{\epsilon^4} = O\left(\frac{\log m(\log n + \log m - \log \epsilon)}{\epsilon^4}\right).$$

Recall the definition of ϵ-CE

ϵ-correlated equilibrium

$x \in \Delta(A)$ is an ϵ-correlated equilibrium if

$$\forall i \in [n],\ \forall f : A_i \mapsto A_i,\ \ \mathbf{E}_{a \sim x}[R^i_f(a)] \leq \epsilon.$$

- nm^m inequalities of the form $\mathbf{E}_{a \sim x}[R^i_f(a)] \leq \epsilon \Rightarrow$ simply applying previous probabilistic method doesn't work.
- The key idea: sampling from an *approximate* correlated equilibrium to consider only $nm^{(\log n + \log m)}$ inequalities.

Sketch of the proof of Theorem 5.1

- By Theorem 3.1: any n-players m-actions game admits an $(\epsilon/2)$-NE $\sigma = \prod_i \sigma_i$, where each player uses a mixed strategy with support size $\leq b = \left\lceil \frac{32(\ln n + \ln m - \ln \epsilon + \ln 16)}{\epsilon^2} \right\rceil$.
- σ is an $\epsilon/2$-*correlated equilibrium* as well.
- Let $B_i \subset A_i$ be such a support ($|B_i| = |\text{support}(\sigma_i)| \leq b$ for all i).
  - The switching rule becomes $f : B_i \mapsto A_i$.
  - $\mathbf{E}_{a \sim \sigma}[R^i_f(a)] \leq \epsilon/2$.

Sketch of the proof of Theorem 5.1 (contd.)

Apply the probabilistic method again.

- Sample k action profiles $a(1), a(2), \ldots, a(k) \in A$ independently at random according to σ and denote by s the uniform distribution over the samples.
- By Hoeffding's inequality:

$$\Pr(\mathbf{E}_{a\sim s}[R_f^i(a)] \geq \epsilon) = \Pr\left(\frac{1}{k}\sum_{\ell\in[k]} R_f^i(a(\ell)) \geq \epsilon\right) \leq e^{-\frac{k\epsilon^2}{8}}.$$

- Setting $k > \frac{264 \ln m(\ln n + \ln m - \ln \epsilon + \ln 16)}{\epsilon^4}$ guarantees

$$\Pr(\mathbf{E}_{a\sim s}[R_f^i(a(\ell)) \geq \epsilon]) < \frac{1}{nm^b}.$$

Complexity of computing an ϵ-CE

Proposition 5.2

There exists a poly(n, m) time randomized algorithm for computing a k-uniform ϵ-correlated equilibrium for

$$k > \frac{2(m \ln m + \ln n + \ln 2)}{\epsilon^2}.$$

Sketch of the proof:

- The "very specific" σ in the proof of Theorem 5.1:
  - ⋆ Not known to be computable in polynomial time.
- We turn to the polynomial time algorithm finding a CE [Jiang & Leyton-Brown 2013].
- $O\left(\frac{m \log m + \log n}{\epsilon^2}\right)$ samples from the CE are enough.
  - ⋆ Previously best bound ($\epsilon = 0$): $O(nm^2)$ [Jiang & Leyton-Brown 2013; Hart & Mas-Colell 2000].

Thanks for your attention.