A General Framework for Computing Optimal Correlated Equilibria in Compact Games

Albert Xin Jiang Kevin Leyton-Brown
Department of Computer Science
University of British Columbia

December, 2011

Correlated Equilibrium

- correlated equilibrium (CE) [Aumann, 1974; Aumann, 1987]
  - generalization of Nash equilibrium
  - players can coordinate their behavior based on signals from an intermediary

Correlated Equilibrium

- correlated equilibrium (CE) [Aumann, 1974; Aumann, 1987]
  - generalization of Nash equilibrium
  - players can coordinate their behavior based on signals from an intermediary

  - natural learning dynamics converge to CE

Correlated Equilibrium

- correlated equilibrium (CE) [Aumann, 1974; Aumann, 1987]
  - generalization of Nash equilibrium
  - players can coordinate their behavior based on signals from an intermediary

  - natural learning dynamics converge to CE
  - tractable to compute: LP
    - polynomial in the size of the normal form

Compact Game Representations

Compact representations are necessary for large games with structured utility functions

- symmetric games / anonymous games
- graphical games [Kearns, Littman & Singh, 2001]
- congestion games [Rosenthal, 1973]
- action-graph games [Jiang, Leyton-Brown & Bhat, 2011]

Compact Game Representations

Compact representations are necessary for large games with structured utility functions

- symmetric games / anonymous games
- graphical games [Kearns, Littman & Singh, 2001]
- congestion games [Rosenthal, 1973]
- action-graph games [Jiang, Leyton-Brown & Bhat, 2011]

Computation of CE in compact games

Compact Game Representations

Compact representations are necessary for large games with structured utility functions

- symmetric games / anonymous games
- graphical games [Kearns, Littman & Singh, 2001]
- congestion games [Rosenthal, 1973]
- action-graph games [Jiang, Leyton-Brown & Bhat, 2011]

Computation of CE in compact games

- Poly-time to find a CE [Papadimitriou & Roughgarden, 2008; Jiang & Leyton-Brown, 2011]

Compact Game Representations

Compact representations are necessary for large games with structured utility functions

- symmetric games / anonymous games
- graphical games [Kearns, Littman & Singh, 2001]
- congestion games [Rosenthal, 1973]
- action-graph games [Jiang, Leyton-Brown & Bhat, 2011]

Computation of CE in compact games

- Poly-time to find a CE [Papadimitriou & Roughgarden, 2008; Jiang & Leyton-Brown, 2011]
- However, there can be an infinite number of CE

Computing Optimal CE

Computing optimal CE according to some linear objective given a compact game

- social welfare

Computing Optimal CE

Computing optimal CE according to some linear objective given a compact game

- social welfare
- max-min welfare

Computing Optimal CE

Computing optimal CE according to some linear objective given a compact game

- social welfare
- max-min welfare
- provide bounds on learning dynamics, price of anarchy, price of stability

Computing Optimal CE

Computing optimal CE according to some linear objective given a compact game

- social welfare
- max-min welfare
- provide bounds on learning dynamics, price of anarchy, price of stability

- Remove incentive constraints → computation of optimal outcome
  - already nontrivial

Computing Optimal CE

Computing optimal CE according to some linear objective given a compact game

- social welfare
- max-min welfare
- provide bounds on learning dynamics, price of anarchy, price of stability

- Remove incentive constraints $\rightarrow$ computation of optimal outcome
  - already nontrivial
- How does adding the CE incentive constraints affect the computational complexity?

Related Work

- Papadimitriou & Roughgarden [SODA 2005, JACM 2008]
  - NP-hard for many representations
    - graphical games
    - congestion games
    - polymatrix games

Related Work

- Papadimitriou & Roughgarden [SODA 2005, JACM 2008]
  - NP-hard for many representations
    - graphical games
    - congestion games
    - polymatrix games
  - a sufficient condition for tractable computation of optimal CE
    - tractable classes: anonymous games, tree graphical games

Related Work

- Papadimitriou & Roughgarden [SODA 2005, JACM 2008]
  - NP-hard for many representations
    - graphical games
    - congestion games
    - polymatrix games
  - a sufficient condition for tractable computation of optimal CE
    - tractable classes: anonymous games, tree graphical games
    - limited to reduced forms; does not apply to e.g. polymatrix games, congestion games

Our Contributions

algorithmic approach for computing optimal CE that applies to all compact representations

- a more general sufficient condition: deviation-adjusted social-welfare problem

Our Contributions

algorithmic approach for computing optimal CE that applies to all compact representations

- a more general sufficient condition: deviation-adjusted social-welfare problem
- identify new tractable classes of compact games:
  - tree polymatrix games

Our Contributions

algorithmic approach for computing optimal CE that applies to all compact representations

- a more general sufficient condition: deviation-adjusted social-welfare problem
- identify new tractable classes of compact games:
  - tree polymatrix games
- also applies to the related solution concept of coarse correlated equilibria (CCE):
  - tractable for singleton congestion games

CE

- simultaneous-move game
  - n players
  - player p's pure strategy $s_p \in S_p$
  - pure strategy profile $s \in S = \prod_{p=1}^{n} S_p$
  - utility for p under pure strategy profile s is integer u_s^p

CE

- simultaneous-move game
  - n players
  - player p's pure strategy $s_p \in S_p$
  - pure strategy profile $s \in S = \prod_{p=1}^{n} S_p$
  - utility for p under pure strategy profile s is integer u_s^p
- a CE is a distribution x over S:
  - a trusted intermediary draws a strategy profile s from this distribution
  - announce to each player p (privately) her own component s_p
  - p will have no incentive to choose another strategy, assuming others follow suggestions

LP formulation for CE

- incentive constraints: for all players p and all $i, j \in S_p$:

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

write as

$$Ux \geq 0.$$

LP formulation for CE

- incentive constraints: for all players p and all $i, j \in S_p$:

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

  write as

$$Ux \geq 0.$$

- x is a distribution: $x \geq 0$, $\sum_s x_s = 1$

LP formulation for CE

- incentive constraints: for all players p and all $i, j \in S_p$:

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

  write as

$$Ux \geq 0.$$

- x is a distribution: $x \geq 0$, $\sum_s x_s = 1$
- m^n variables, nm^2 constraints

LP Formulation for Optimal CE

$$\begin{aligned} \max w^T x \\ Ux \geq 0 \\ x \geq 0 \\ \sum_{s \in S} x_s = 1 \end{aligned} \tag{P}$$

m^n variables, nm^2 constraints

Solving the Dual

Consider the dual of (P),

$$\begin{aligned} \min t \\ U^T y + w \leq t\mathbf{1} \\ y \geq 0. \end{aligned} \qquad (D)$$

- nm^2 variables, about m^n constraints

Solving the Dual

Consider the dual of (P),

$$\begin{aligned} \min t \\ U^T y + w \leq t\mathbf{1} \\ y \geq 0. \end{aligned} \tag{D}$$

- nm^2 variables, about m^n constraints
- ellipsoid method?

Deviation-adjusted Social Welfare

Definition

Given a game, and a vector $y \in \mathbb{R}^N$ such that $y \geq 0$, the deviation-adjusted utility for player p under pure profile s is

$$\hat{u}_s^p(y) = u_s^p + \sum_{j \in S_p} y_{s_p,j}^p \left(u_s^p - u_{js_{-p}}^p \right).$$

The deviation-adjusted social welfare is $\hat{w}_s(y) = \sum_p \hat{u}_s^p(y)$.

Sufficient Condition

deviation-adjusted social welfare problem is the following: given an instance of the representation and rational vector $(y, t) \in \mathbb{Q}^{N+1}$ such that $y \geq 0$, determine if there exists an s such that the deviation-adjusted social welfare $\hat{w}_s(y) > t$; if so output such an s.

Sufficient Condition

deviation-adjusted social welfare problem is the following: given an instance of the representation and rational vector $(y, t) \in \mathbb{Q}^{N+1}$ such that $y \geq 0$, determine if there exists an s such that the deviation-adjusted social welfare $\hat{w}_s(y) > t$; if so output such an s.

Theorem

If the deviation-adjusted social welfare problem can be solved in polynomial time for a game representation, then so can the problem of computing the maximum social welfare CE.

Reduced Forms

Definition ([Papadimitriou & Roughgarden, 2008])

Consider a game $G = (\mathcal{N}, \{S_p\}_{p\in\mathcal{N}}, \{u^p\}_{p\in\mathcal{N}})$. For $p = 1, \ldots, n$, let $P_p = \{C_p^1 \ldots C_p^{r_p}\}$ be a partition of S_{-p} into r_p classes. The set $\mathcal{P} = \{P_1, \ldots, P_n\}$ of partitions is a reduced form of G if $u_s^p = u_{s'}^p$ whenever

1. $s_p = s'_p$ and
2. both s_{-p} and s'_{-p} belong to the same class in P_p.

- example: graphical games, anonymous games

Reduced Forms

Definition ([Papadimitriou & Roughgarden, 2008])

Consider a game $G = (\mathcal{N}, \{S_p\}_{p \in \mathcal{N}}, \{u^p\}_{p \in \mathcal{N}})$. For $p = 1, \ldots, n$, let $P_p = \{C_p^1 \ldots C_p^{r_p}\}$ be a partition of S_{-p} into r_p classes. The set $\mathcal{P} = \{P_1, \ldots, P_n\}$ of partitions is a reduced form of G if $u_s^p = u_{s'}^p$ whenever

1. $s_p = s'_p$ and
2. both s_{-p} and s'_{-p} belong to the same class in P_p.

- example: graphical games, anonymous games
- P&R [2008]'s sufficient condition: optimize social welfare of a game with same reduced form but arbitrarily modified utilities

Reduced Forms

Definition ([Papadimitriou & Roughgarden, 2008])

Consider a game $G = (\mathcal{N}, \{S_p\}_{p\in\mathcal{N}}, \{u^p\}_{p\in\mathcal{N}})$. For $p = 1, \ldots, n$, let $P_p = \{C_p^1 \ldots C_p^{r_p}\}$ be a partition of S_{-p} into r_p classes. The set $\mathcal{P} = \{P_1, \ldots, P_n\}$ of partitions is a reduced form of G if $u_s^p = u_{s'}^p$ whenever

1. $s_p = s'_p$ and
2. both s_{-p} and s'_{-p} belong to the same class in P_p.

- example: graphical games, anonymous games
- P&R [2008]'s sufficient condition: optimize social welfare of a game with same reduced form but arbitrarily modified utilities
- We show: given reduced form, deviation-adjusted social welfare problem reduces to P&R [2008]'s sufficient condition.
  - the reduced form structure is preserved under the transformation to deviation-adjusted utilities

Linear Reduced Forms

Linear Reduced Forms

- multiple partitions of S_{-p} for each agent
- an agent's overall utility is a sum over utility functions defined on each of that agent's partitions

Linear Reduced Forms

Linear Reduced Forms

- multiple partitions of S_{-p} for each agent
- an agent's overall utility is a sum over utility functions defined on each of that agent's partitions

Optimal social welfare outcome in poly-time $\Rightarrow$ optimal CE in poly-time

Linear Reduced Forms

Linear Reduced Forms

- multiple partitions of S_{-p} for each agent
- an agent's overall utility is a sum over utility functions defined on each of that agent's partitions

Optimal social welfare outcome in poly-time $\Rightarrow$ optimal CE in poly-time

Corollary

Optimal CE in tree polymatrix games can be computed in polynomial time.

What Types of Structure Is Preserved?

$$\hat{u}_s^p(y) = u_s^p + \sum_{j \in S_p} y_{s_p,j}^p \left(u_s^p - u_{js_{-p}}^p \right)$$

- structure is preserved when partitions do not depend on s_p

What Types of Structure Is Preserved?

$$\hat{u}_s^p(y) = u_s^p + \sum_{j \in S_p} y_{s_p,j}^p \left(u_s^p - u_{js_{-p}}^p \right)$$

- structure is preserved when partitions do not depend on s_p
- representations with action-specific structure: e.g. congestion games, action-graph games
  - $\hat{u}_s^p(y)$ has different structure from u_s^p

Coarse Correlated Equilibrium (CCE)

for each player p and each of his actions $j \in S_p$

$$\sum_{(i,s_{-p}) \in S} [u^p_{is_{-p}} - u^p_{js_{-p}}] x_{is_{-p}} \geq 0$$

- a CCE is a CE, not vice versa

Primal and Dual LP Formulations

$$\max w^T x \qquad (CP)$$
$$Cx \geq 0$$
$$x \geq 0$$
$$\sum_{s \in S} x_s = 1$$

Dual:

$$\min t$$
$$C^T y + w \leq t\mathbf{1}$$
$$y \geq 0$$

Coarse Deviation-adjusted Social Welfare

Definition

Given a game, and a vector $y \in \mathbb{R}^{\sum_p |S_p|}$ such that $y \geq 0$, the coarse deviation-adjusted utility for player p under pure profile s is

$$\tilde{u}^p_s(y) = u^p_s + \sum_{j \in S_p} y^p_j (u^p_s - u^p_{js_{-p}})$$

The coarse deviation-adjusted social welfare is $\tilde{w}_s(y) = \sum_p \tilde{u}^p_s(y)$.

Theorem

If the coarse deviation-adjusted social welfare problem can be solved in polynomial time for a game representation, then the problem of computing the maximum social welfare CCE is in polynomial time for this representation.

Singleton Congestion Games

- symmetric; each player choose one from a set of resources $\mathcal{A}$
- utility of choosing $\alpha \in \mathcal{A}$ is a function of $c(\alpha)$, the $\#$ of players choosing α
- social welfare: $w_s = \sum_\alpha c(\alpha) f^\alpha(c(\alpha))$
  - optimal social welfare outcome in polynomial time

Singleton Congestion Games

- symmetric; each player choose one from a set of resources $\mathcal{A}$
- utility of choosing $\alpha \in \mathcal{A}$ is a function of $c(\alpha)$, the $\#$ of players choosing α
- social welfare: $w_s = \sum_\alpha c(\alpha) f^\alpha(c(\alpha))$
  - optimal social welfare outcome in polynomial time
- when y is player-symmetric: $\tilde{w}_s(y) = \sum_\alpha g^\alpha(c(\alpha))$.
  - coarse deviation-adjusted SW problem in polynomial time

CCE for Singleton Congestion Games

How to guarantee symmetric y?

- sufficient to start ellipsoid method with symmetric initial conditions, and
- ensure symmetric cutting planes
  - symmetrize pure-strategy profile

$$(s_1, s_2, s_3) \mapsto [\frac{1}{3}(s_1, s_2, s_3), \frac{1}{3}(s_3, s_1, s_2), \frac{1}{3}(s_2, s_3, s_1)]$$

CCE for Singleton Congestion Games

How to guarantee symmetric y?

- sufficient to start ellipsoid method with symmetric initial conditions, and
- ensure symmetric cutting planes
  - symmetrize pure-strategy profile

$$(s_1, s_2, s_3) \mapsto [\frac{1}{3}(s_1, s_2, s_3), \frac{1}{3}(s_3, s_1, s_2), \frac{1}{3}(s_2, s_3, s_1)]$$

Corollary

Given a singleton congestion game, the optimal social welfare CCE can be computed in polynomial time.

Summary and Open Problems

- sufficient condition for tractable computation of optimal CE and optimal CCE
- new tractable classes of compact games

Summary and Open Problems

- sufficient condition for tractable computation of optimal CE and optimal CCE
- new tractable classes of compact games
- (coarse) deviation-adjusted social welfare problem preserves the structure of original game $\Rightarrow$ optimal (C)CE is no harder than optimal social welfare outcome

Summary and Open Problems

- sufficient condition for tractable computation of optimal CE and optimal CCE
- new tractable classes of compact games
- (coarse) deviation-adjusted social welfare problem preserves the structure of original game $\Rightarrow$ optimal (C)CE is no harder than optimal social welfare outcome

Open Problems

- sufficient & necessary conditions for tractable computation?
- approximations
- learning dynamics

Aumann, R. (1974). Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics*, *1*(1), 67–96.

Aumann, R. (1987). Correlated equilibrium as an expression of Bayesian rationality. *Econometrica: Journal of the Econometric Society*, 1–18.

Jiang, A., & Leyton-Brown, K. (2011). Polynomial computation of exact correlated equilibrium in compact games. *EC: Proceedings of the ACM Conference on Electronic Commerce*. http://arxiv.org/abs/1011.0253.

Papadimitriou, C., & Roughgarden, T. (2008). Computing correlated equilibria in multi-player games. *Journal of the ACM*, *55*(3), 14.

Rosenthal, R. (1973). A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, *2*, 65–67.