Asymptotically Truthful Equilibrium Selection in Large Congestion Games

Ryan Rogers - Penn
Aaron Roth - Penn

June 12, 2014

Related Work

- **Large Games**
  - Roberts and Postlewaite 1976
  - Immorlica and Mahdian 2005, Kojima and Pathak 2009, Kojima et al 2010
  - Bodoh-Creed 2013
  - Azevedo and Budish 2011
- **Incorporating a Mediator**
  - Monderer and Tennenholtz 2003, 2009
  - Ashlagi et al 2009
- **Work most closely related to ours**
  - Kearns et al 2014.

Routing Game

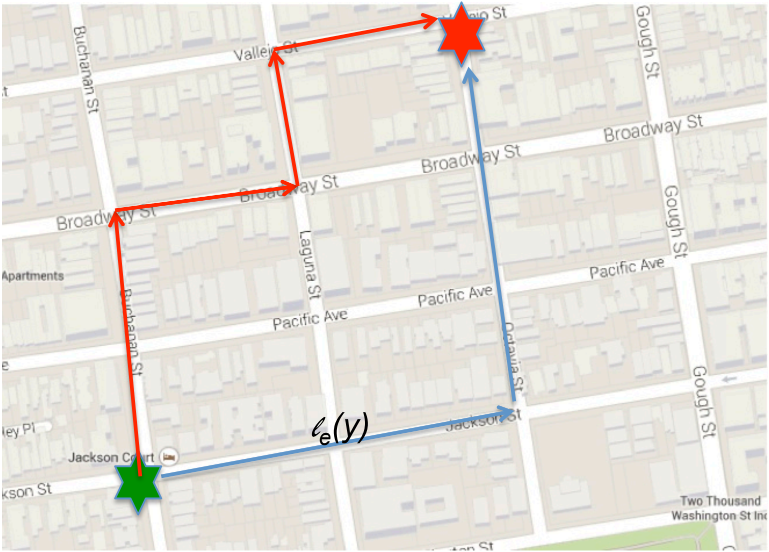

Routing Game

- A game $\mathcal{G}$ is defined by
  - A set of n players
  - A set of types $\mathcal{U} \implies$ **source destination** pair $s_i \in \mathcal{U}$.
  - A set of actions $A \implies$ **routes** for each source destination pair.
  - A cost function $c : \mathcal{U} \times A^n \to \mathbb{R}$

$$c(s_i, \mathbf{a}) = \sum_{e \in a_i} \ell_e(y_e(\mathbf{a}))$$

Routing Game

- A game $\mathcal{G}$ is defined by
  - A set of n players
  - A set of types $\mathcal{U} \implies$ **source destination** pair $s_i \in \mathcal{U}$.
  - A set of actions $A \implies$ **routes** for each source destination pair.
  - A cost function $c : \mathcal{U} \times A^n \to \mathbb{R}$

$$c(s_i, \mathbf{a}) = \sum_{e \in a_i} \ell_e(y_e(\mathbf{a}))$$

- Players may not know each others type.

Routing Game

- A game $\mathcal{G}$ is defined by
  - A set of n players
  - A set of types $\mathcal{U} \implies$ **source destination** pair $s_i \in \mathcal{U}$.
  - A set of actions $A \implies$ **routes** for each source destination pair.
  - A cost function $c : \mathcal{U} \times A^n \to \mathbb{R}$

$$c(s_i, \mathbf{a}) = \sum_{e \in a_i} \ell_e(y_e(\mathbf{a}))$$

- Players may not know each others type.
  - n may be HUGE!!

Routing Game

- A game $\mathcal{G}$ is defined by
  - A set of n players
  - A set of types $\mathcal{U} \implies$ **source destination** pair $s_i \in \mathcal{U}$.
  - A set of actions $A \implies$ **routes** for each source destination pair.
  - A cost function $c : \mathcal{U} \times A^n \to \mathbb{R}$

$$c(s_i, \mathbf{a}) = \sum_{e \in a_i} \ell_e(y_e(\mathbf{a}))$$

- Players may not know each others type.
  - n may be HUGE!!
  - Types may be sensitive information

Routing Game

- A game $\mathcal{G}$ is defined by
  - A set of n players
  - A set of types $\mathcal{U} \implies$ **source destination** pair $s_i \in \mathcal{U}$.
  - A set of actions $A \implies$ **routes** for each source destination pair.
  - A cost function $c : \mathcal{U} \times A^n \to \mathbb{R}$

$$c(s_i, \mathbf{a}) = \sum_{e \in a_i} \ell_e(y_e(\mathbf{a}))$$

- Players may not know each others type.
  - n may be HUGE!!
  - Types may be sensitive information
- **Main Goal** : Have players play a pure strategy Nash equilibrium of the complete information game in settings of partial information.

Mediator for a Routing Game

Mediator for a Routing Game

Introduce a Mediator

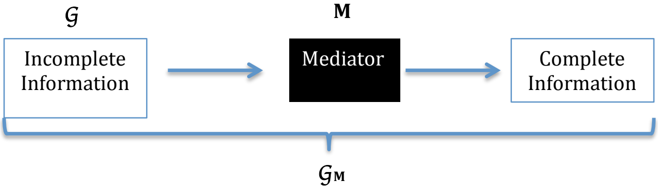

- A mediator is an algorithm $M : (\mathcal{U} \cup \perp)^n \to (A \cup \perp)^n$.

Weak Mediator

Weak Mediator

- Mediator cannot force people to use it

Weak Mediator

- Mediator cannot force people to use it
- Players need not follow its suggested action

Weak Mediator

- Mediator cannot force people to use it
- Players need not follow its suggested action
- Players may lie to the mechanism if they choose to use it.

Augmented Game

- Define the augmented game $\mathcal{G}_M$ (Kearns et al 2014):
  - Action Space:

$$A' = \{(s, f) : s \in \mathcal{U} \cup \perp, f : (A \cup \perp) \to A\}$$

$$g_i = (s_i, f_i) \in A'$$

  - Costs for $\mathbf{g}' = ((s'_i, f_i))_{i=1}^n$:

$$c_M(s_i, \mathbf{g}') = \mathbb{E}_{\mathbf{a} \sim M(\mathbf{s}')} \left[c(s_i, \mathbf{f}(\mathbf{a})) \right]$$

Good Behavior

- Player's should:

Good Behavior

- Player's should:
    - Use the Mediator M

Good Behavior

- Player's should:
  - Use the Mediator M
  - Report her true type to M

Good Behavior

- Player's should:
  - Use the Mediator M
  - Report her true type to M
  - Follow the suggested action of $M \implies f_i =$ identity map.

Joint Differential Privacy

- (Kearns et al 2014) Let $M : D^n \to O^n$. Then M satisfies ϵ-joint differential privacy if for every $\mathbf{s} \in D^n$, for every $i \in [n]$, $s_i' \in D$ and for every $B \subset O^{n-1}$

$$\mathbb{P}[M(\mathbf{s})_{-i} \in B] \leq e^{\epsilon}\mathbb{P}[M(s_i', \mathbf{s}_{-i})_{-i} \in B]$$

Billboard Lemma

Billboard Lemma

- If a mechanism $M : \mathcal{U}^n \to O$ is (ϵ, δ)-differentially private and consider any function $\phi : \mathcal{U} \times O \to A^n$. Define $M' : \mathcal{U}^n \to A^n$ to be

$$M'(\mathbf{s})_i = \phi(s_i, M(\mathbf{s})).$$

Then M' is (ϵ, δ)- joint differentially private.

Motivating Theorem

- Let $\mathcal{G}$ be any game with costs in $[0, m]$, and let M be a mediator such that

Motivating Theorem

- Let $\mathcal{G}$ be any game with costs in $[0, m]$, and let M be a mediator such that
  - It is ϵ-joint differentially private

Motivating Theorem

- Let $\mathcal{G}$ be any game with costs in $[0, m]$, and let M be a mediator such that
  - It is ϵ-joint differentially private
  - For any set of reported types $\mathbf{s}$, it outputs an η-approximate pure strategy Nash Equilibrium.

Motivating Theorem

- Let $\mathcal{G}$ be any game with costs in $[0, m]$, and let M be a mediator such that
  - It is ϵ-joint differentially private
  - For any set of reported types $\mathbf{s}$, it outputs an η-approximate pure strategy Nash Equilibrium.
- Then good behavior $\mathbf{g}^*$ is an η'-approximate ex-post equilibrium for the incomplete information game $\mathcal{G}_M$ where

$$\eta' = 2m\epsilon + \eta$$

Main Theorem

- There *exists* such a mechanism from the motivating theorem for **large** congestion games.
- Further, we show that good behavior $\mathbf{g}^*$ is an η'-approximate ex-post equilibrium for the incomplete information game $\mathcal{G}_M$ where

$$\eta' = \tilde{\mathcal{O}}\left(\left(\frac{m^5}{n}\right)^{1/4}\right) \to 0 \text{ as } n \to \infty$$

Large Games

Large Games

- We assume that each player cannot significantly change the cost of another player by changing her route.

$$|\ell_e(y_e) - \ell_e(y_e + 1)| \leq \frac{1}{n} \qquad \text{for } y_e \in [n] \text{ and } e \in E.$$

- The costs then satisfy for $j \neq i$ and $a_j' \neq a_j' \in A$

$$|c(s_i, (a_j, \mathbf{a}_{-j})) - c(s_i, (a_j', \mathbf{a}_{-j}))| \leq \frac{m}{n}.$$

How to Construct Such a Mechanism?

How to Construct Such a Mechanism?

- Simulate Best Response Dynamics

How to Construct Such a Mechanism?

- Simulate Best Response Dynamics
- Compute Best Responses privately

How to Construct Such a Mechanism?

- Simulate Best Response Dynamics
- Compute Best Responses privately
- Limit the number of times a single player can change routes.

Best Responses

- In congestion games, allowing each player to best respond given the other players routes will converge to an approximate Nash Equilibrium.

Best Responses

- In congestion games, allowing each player to best respond given the other players routes will converge to an approximate Nash Equilibrium.
- We will have an algorithm that will have each player move if she can improve her cost by more than α: α-**Best Response**

Best Responses

- In congestion games, allowing each player to best respond given the other players routes will converge to an approximate Nash Equilibrium.
- We will have an algorithm that will have each player move if she can improve her cost by more than α: α-**Best Response**
- There can be no more than $T = \frac{mn}{\alpha}$ best responses.

Best Responses

- In congestion games, allowing each player to best respond given the other players routes will converge to an approximate Nash Equilibrium.
- We will have an algorithm that will have each player move if she can improve her cost by more than α: α-**Best Response**
- There can be no more than $T = \frac{mn}{\alpha}$ best responses.
- We need to only maintain a count of the number of people on every edge to compute α-Best Responses for each player

Binary Mechanism

- Chan et al 2011 and Dwork et al 2010 give a way to obtain an online count of a sensitivity 1 stream $\omega \in \{0,1\}^T$ such that the output $\hat{y}^t$ for any $t = 1, 2, \cdots, T$ is
  - ϵ differentially private
  - Has high accuracy to the exact count y^t for every $t = 1, \cdots, T$

$$|\hat{y}^t - y^t| \leq \tilde{\mathcal{O}}\left(\frac{1}{\epsilon}\right)$$

0	1	1	0	1	1	0	0		47		1	0	1

Generalized Binary Mechanism

0	-1	1	-1	0	0	1	1		21		1	0	-1

Generalized Binary Mechanism

0	-1	1	-1	0	0	1	1		21		1	0	-1
1	0	0	1	0	-1	-1	1		16		1	1	0
0	1	1	1	0	0	0	-1		8		1	1	1
0	0	0	0	1	-1	1	1		37		1	-1	-1

Generalized Binary Mechanism

0	-1	1	-1	0	0	1	1		21	1	0	-1
1	0	0	1	0	-1	-1	1		16	1	1	0
0	1	1	1	0	0	0	-1		8	1	1	1
0	0	0	0	1	-1	1	1		37	1	-1	-1

- Each of the m streams are k-sensitive, so we get
  - ϵ differentially private counters
  - With high probability

$$|\hat{y}_e^t - y_e^t| \leq \tilde{O}\left(\frac{km}{\epsilon}\right) \forall e \in E, t = 1, \cdots, T$$

The Gap

- After a player *i* has made an α-private best response, how many times must other players move before *i* can move again? We will call this the **gap** γ.

The Gap

- After a player i has made an α-private best response, how many times must other players move before i can move again? We will call this the **gap** γ.
- Due to the *largeness condition*, each time a player does not move, her cost can increase by at most $\frac{m}{n}$ and can only move once her cost has increased by α

$$\gamma = \tilde{\Omega}\left(\frac{\alpha n}{m}\right)$$

The Gap

- After a player i has made an α-private best response, how many times must other players move before i can move again? We will call this the **gap** γ.
- Due to the *largeness condition*, each time a player does not move, her cost can increase by at most $\frac{m}{n}$ and can only move once her cost has increased by α

$$\gamma = \tilde{\Omega}\left(\frac{\alpha n}{m}\right)$$

- All the players can only make $T = \tilde{\mathcal{O}}\left(\frac{mn}{\alpha}\right)$ (with high probability).

The Gap

- After a player i has made an α-private best response, how many times must other players move before i can move again? We will call this the **gap** γ.
- Due to the *largeness condition*, each time a player does not move, her cost can increase by at most $\frac{m}{n}$ and can only move once her cost has increased by α

$$\gamma = \tilde{\Omega}\left(\frac{\alpha n}{m}\right)$$

- All the players can only make $T = \tilde{\mathcal{O}}\left(\frac{mn}{\alpha}\right)$ (with high probability).
- A player only changes routes k times

$$k = \mathcal{O}\left(\frac{m^2}{\alpha^2}\right)$$

Equilibrium Analysis of our Algorithm

- With high probability, after $T = \tilde{\mathcal{O}}\left(\frac{mn}{\alpha}\right)$ moves by all players, no player will be able to improve her *private* cost by more than α. If we set

$$\alpha = \tilde{\Theta}\left(\left(\frac{m^4}{n\epsilon}\right)^{1/3}\right)$$

then we know no player will be able to improve her *actual* cost by more than

$$\eta \leq \alpha + \text{ Error from BM } = \tilde{\mathcal{O}}\left(\left(\frac{m^4}{n\epsilon}\right)^{1/3}\right)$$

Equilibrium Analysis of our Algorithm

- Is it Joint Differentially Private?

Equilibrium Analysis of our Algorithm

- Is it Joint Differentially Private?
- Recall our motivating theorem that says *good* behavior is an η'-approximate ex-post equilibrium for $\mathcal{G}_M$ and we can set ϵ (which is a parameter we control) to satisfy the following

$$\eta' = \tilde{\mathcal{O}}\left(\left(\frac{m^5}{n}\right)^{1/4}\right) \to 0 \quad \text{as} \quad n \to \infty$$

Open Questions

- Can Nash Equilibria of the complete information game be implemented as exact ex-post or Bayes Nash Equilibria of the incomplete information game?
- Does there exist a jointly differentially private algorithm for computing approximate Nash Equilibria for *general* large games?