Mechanism Design in Large Congestion Games

Ryan Rogers, Aaron Roth, Jonathan Ullman, and Steven Wu

July 22, 2015

Routing Game

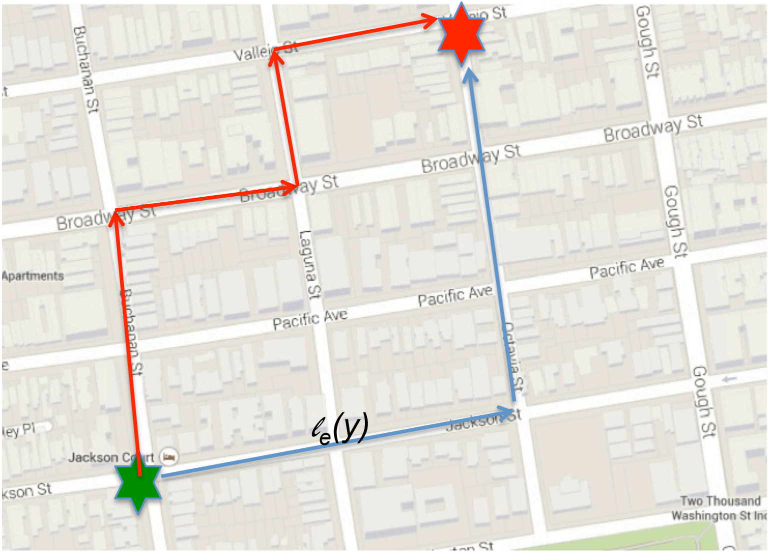

Routing Game

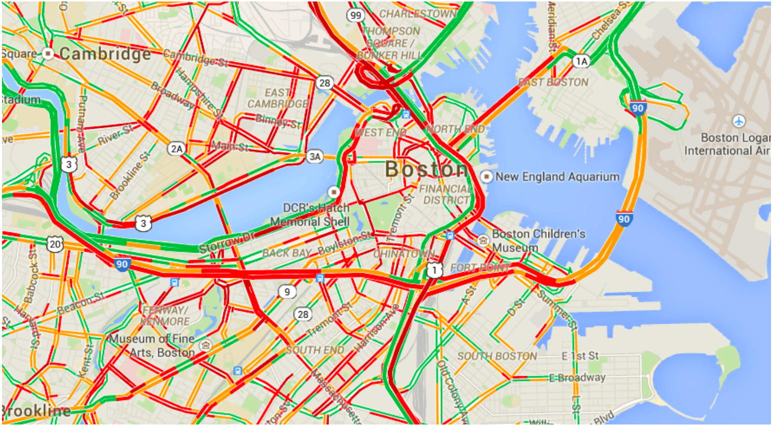

Routing Game

- A routing game $\mathcal{G}$ is defined by

Routing Game

- A routing game $\mathcal{G}$ is defined by
  - A set of n players

Routing Game

- A routing game $\mathcal{G}$ is defined by
  - A set of n players
  - A graph with vertices V and edges E where $|E| = m$ and each edge has a **latency function** $\ell_e : \mathbb{R}_+ \to [0, 1]$.

Routing Game

- A routing game $\mathcal{G}$ is defined by
  - A set of n players
  - A graph with vertices V and edges E where $|E| = m$ and each edge has a **latency function** $\ell_e : \mathbb{R}_+ \to [0, 1]$.
  - A set of types $\mathcal{T} \implies$ **source destination** pair $t_i = (s_i, d_i) \in \mathcal{T}$.

Routing Game

- A routing game $\mathcal{G}$ is defined by
  - A set of n players
  - A graph with vertices V and edges E where $|E| = m$ and each edge has a **latency function** $\ell_e : \mathbb{R}_+ \to [0, 1]$.
  - A set of types $\mathcal{T} \implies$ **source destination** pair $t_i = (s_i, d_i) \in \mathcal{T}$.
  - A set of actions $A \implies$ **routes** for each source destination pair.

Routing Game

- A routing game $\mathcal{G}$ is defined by
  - A set of n players
  - A graph with vertices V and edges E where $|E| = m$ and each edge has a **latency function** $\ell_e : \mathbb{R}_+ \to [0, 1]$.
  - A set of types $\mathcal{T} \implies$ **source destination** pair $t_i = (s_i, d_i) \in \mathcal{T}$.
  - A set of actions $A \implies$ **routes** for each source destination pair.
  - A **cost function** $c : \mathcal{T} \times A^n \to \mathbb{R}$ depends on **congestion** y_e on each edge

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

Incomplete Information Setting

Players may not know each other's type - incomplete information.

Incomplete Information Setting

Players may not know each other's type - incomplete information.

- *Sources and destinations may be sensitive information,*
- n *may be HUGE!!*

First Goal - Equilibrium Selection

Definition

An action profile $\mathbf{a} = (a_1, \ldots, a_n)$ is an η-Nash equilibrium if for every player i of type $t_i \in \mathcal{T}$ and every deviation a'_i we have

$$c(t_i, \mathbf{a}) \leq c(t_i, (a'_i, \mathbf{a}_{-i})) + \eta$$

First Goal - Equilibrium Selection

Definition

An action profile $\mathbf{a} = (a_1, \ldots, a_n)$ is an η-Nash equilibrium if for every player i of type $t_i \in \mathcal{T}$ and every deviation a_i' we have

$$c(t_i, \mathbf{a}) \leq c(t_i, (a_i', \mathbf{a}_{-i})) + \eta$$

Finding a NE requires knowing the types of every player

First Goal - Equilibrium Selection

Definition

An action profile $\mathbf{a} = (a_1, \ldots, a_n)$ is an η-Nash equilibrium if for every player i of type $t_i \in \mathcal{T}$ and every deviation a'_i we have

$$c(t_i, \mathbf{a}) \leq c(t_i, (a'_i, \mathbf{a}_{-i})) + \eta$$

Finding a NE requires knowing the types of every player

- **Goal 1:** Coordinate players to play an approximate Nash equilibrium *as if we knew the types*, even in settings of incomplete information.

Second Goal - Social Welfare

Definition

An action profile **a** is an η-Socially Optimal Routing if

$$C(\mathbf{a}) \leq \min_{\mathbf{a}' \in A^n} C(\mathbf{a}') + \eta$$

where $C(\mathbf{a}) = \sum_{e \in E} y_e \ell_e(y_e)$

Second Goal - Social Welfare

Definition

An action profile $\mathbf{a}$ is an η-Socially Optimal Routing if

$$C(\mathbf{a}) \leq \min_{\mathbf{a}' \in A^n} C(\mathbf{a}') + \eta$$

where $C(\mathbf{a}) = \sum_{e \in E} y_e \ell_e(y_e)$

Finding the social optimal requires knowing the players' types.

Second Goal - Social Welfare

Definition

An action profile **a** is an η-Socially Optimal Routing if

$$C(\mathbf{a}) \leq \min_{\mathbf{a}' \in A^n} C(\mathbf{a}') + \eta$$

where $C(\mathbf{a}) = \sum_{e \in E} y_e \ell_e(y_e)$

Finding the social optimal requires knowing the players' types.

- **Goal 2:** Coordinate selfish players to play an approximate social optimal routing *as if we knew the types*, in settings of incomplete information.

Braess' Paradox

Is it true that the Social Optimal is a NE?

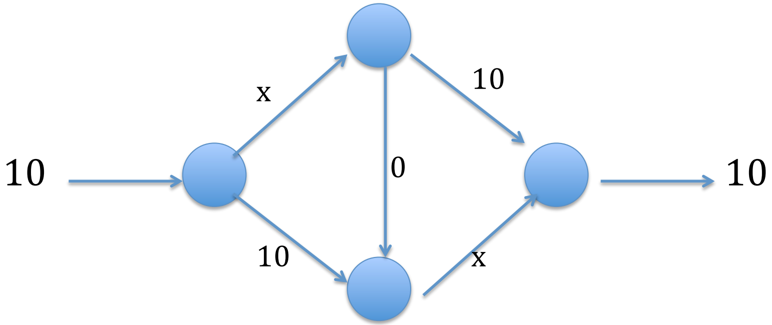

Braess' Paradox

Is it true that the Social Optimal is a NE?

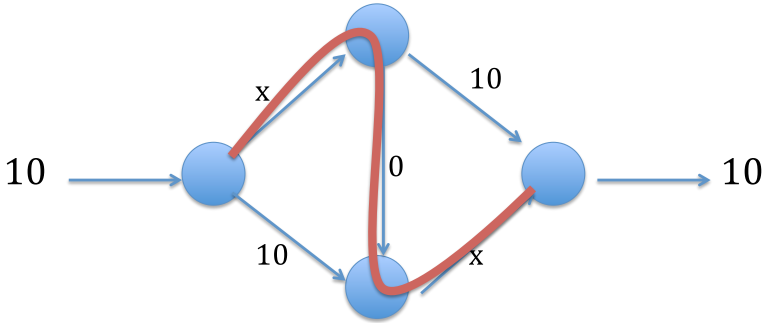

Total Cost

Braess' Paradox

Is it true that the Social Optimal is a NE?

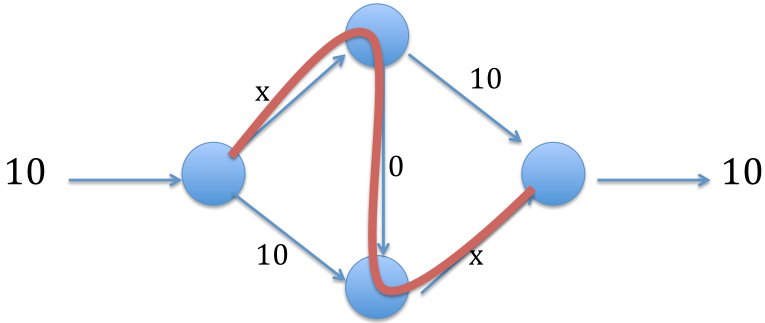

Total Cost $= 10 * 10 + 10 * 10 = 200.$

Braess' Paradox

Is it true that the Social Optimal is a NE?

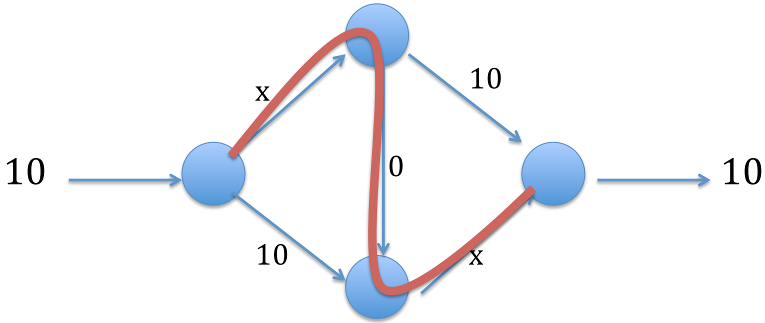

Total Cost $= 10 * 10 + 10 * 10 = 200$.
Equilibrium?

Braess' Paradox

Is it true that the Social Optimal is a NE?

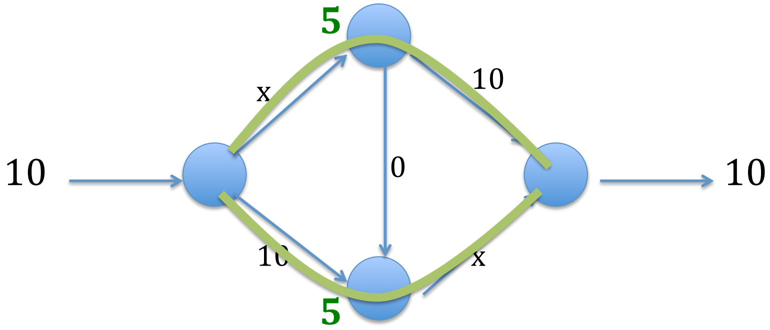

Total Cost

Braess' Paradox

Is it true that the Social Optimal is a NE?

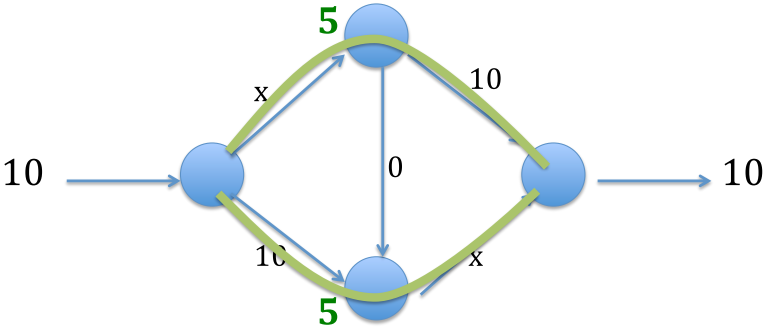

Total Cost $= 5 * (5 + 10) + 5 * (5 + 10) = 150$.

Braess' Paradox

Is it true that the Social Optimal is a NE?

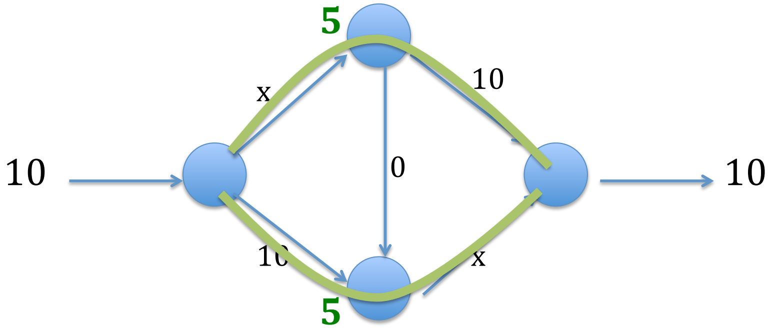

Total Cost $= 5 * (5 + 10) + 5 * (5 + 10) = 150$.
Equilibrium?

Goal One: Equilibrium Selection

- Create a mechanism (or a mediator) that takes reported source destinations as input and then suggests an action (route) for each player to take.
- Mediator for a Routing Game:

Goal One: Equilibrium Selection

- Create a mechanism (or a mediator) that takes reported source destinations as input and then suggests an action (route) for each player to take.
- Mediator for a Routing Game:

Mediated Game

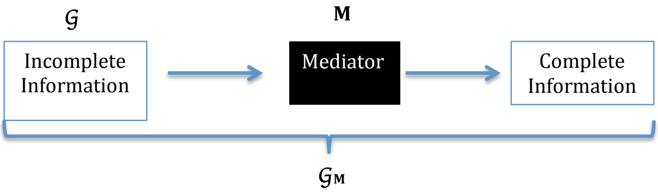

- A mediator is an algorithm $M : (\mathcal{T} \cup \perp)^n \to (A \cup \perp)^n$.

Mediated Game

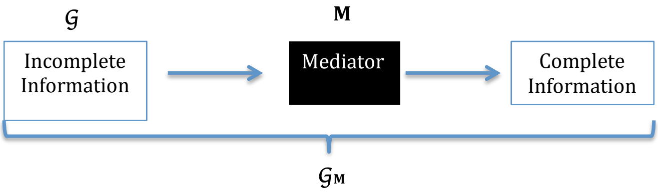

- A mediator is an algorithm $M : (\mathcal{T} \cup \bot)^n \rightarrow (A \cup \bot)^n$.
- The game has changed - Mediated Game $\mathcal{G}_M$.

Mediated Game

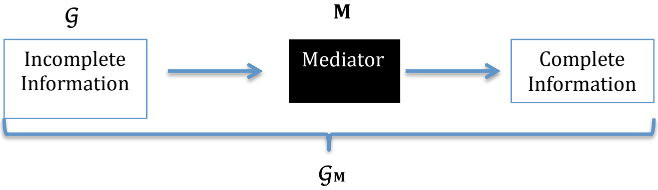

- A mediator is an algorithm $M : (\mathcal{T} \cup \perp)^n \to (A \cup \perp)^n$.
- The game has changed - Mediated Game $\mathcal{G}_M$.
- Players' actions include how they will interact with M.

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 mediator if they choose to use it.*

Good Behavior

Players should have little incentive from deviating from:

Good Behavior

Players should have little incentive from deviating from:

- *Using the Mediator M.*

Good Behavior

Players should have little incentive from deviating from:

- *Using the Mediator M.*
- *Reporting their true type to M.*

Good Behavior

Players should have little incentive from deviating from:

- *Using the Mediator M.*
- *Reporting their true type to M.*
- *Following the suggested actions of M.*

Solution Attempt 1

- Naive Solution: Have M output an approximate NE for the game induced by the reported types.

Solution Attempt 1

- Naive Solution: Have M output an approximate NE for the game induced by the reported types.
- If all players report truthfully, then there is little incentive to not follow M's suggestion.

Solution Attempt 1

- Naive Solution: Have M output an approximate NE for the game induced by the reported types.
- If all players report truthfully, then there is little incentive to not follow M's suggestion.
- If one person changes her type, the game has changed and the NE may be very different - costs to players may be very different between different NE.

Solution Attempt 1

- Naive Solution: Have M output an approximate NE for the game induced by the reported types.
- If all players report truthfully, then there is little incentive to not follow M's suggestion.
- If one person changes her type, the game has changed and the NE may be very different - costs to players may be very different between different NE.
- How do we control the impact any one player has on the outcome of M?

Differential Privacy [DMNS'06]

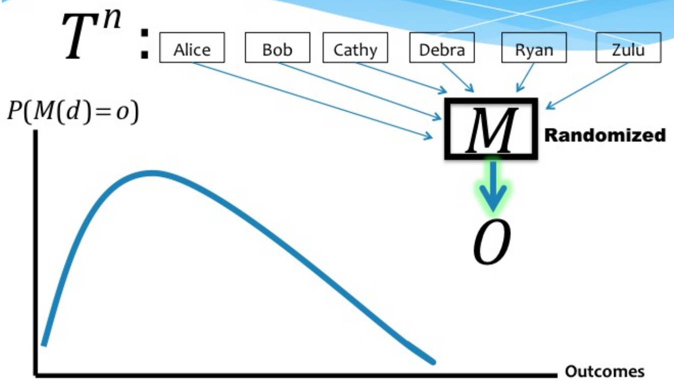

Differential Privacy [DMNS'06]

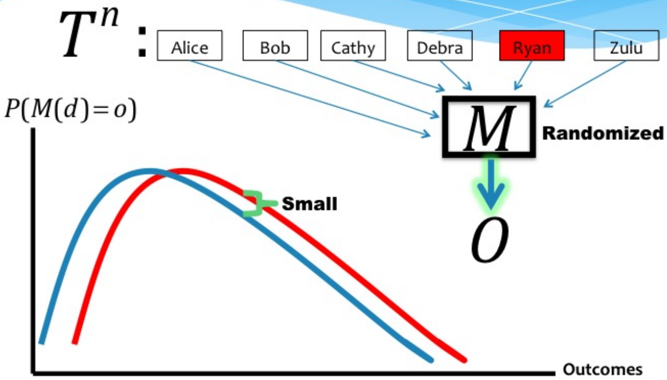

Differential Privacy [DMNS'06]

Definition

A randomized algorithm $M : \mathcal{T}^n \to O$ is ϵ-DP if for all neighboring datasets $\mathbf{t}$ and $\mathbf{t}'$ and all outcome sets $B \subseteq O$ we have

$$\mathbb{P}(M(\mathbf{t}) \in B) \leq e^{\epsilon}\mathbb{P}(M(\mathbf{t}') \in B)$$

Differential Privacy [DMNS'06]

Definition

A randomized algorithm $M : \mathcal{T}^n \to A^n$ is ϵ-DP if for all neighboring datasets $\mathbf{t}$ and $\mathbf{t}'$ and all outcome sets $B \subseteq A^n$ we have

$$\mathbb{P}(M(\mathbf{t}) \in B) \le e^{\epsilon} \mathbb{P}(M(\mathbf{t}') \in B)$$

where $M(\mathbf{t}) = (M_1(\mathbf{t}), \cdots, M_n(\mathbf{t}))$

Differential Privacy [DMNS'06]

Definition

A randomized algorithm $M : \mathcal{T}^n \to A^n$ is ϵ-DP if for all neighboring datasets $\mathbf{t}$ and $\mathbf{t}'$ and all outcome sets $B \subseteq A^n$ we have

$$\mathbb{P}(M(\mathbf{t}) \in B) \leq e^{\epsilon}\mathbb{P}(M(\mathbf{t}') \in B)$$

where $M(\mathbf{t}) = (M_1(\mathbf{t}), \cdots, M_n(\mathbf{t}))$

Too strong of a definition in our case. We want player's actions to depend on their reported type.

Relaxation of DP - Joint DP [KPRU'14]

Definition

A randomized algorithm $M : \mathcal{T}^n \to A^n$ is ϵ-JDP if for every player i, $\mathbf{t}_{-i} \in \mathcal{T}^{n-1}$, $t_i, t_i' \in \mathcal{T}$, and all outcome sets $B \subseteq A^{n-1}$,

$$\mathbb{P}[\underbrace{M(t_i, \mathbf{t}_{-i})_{-i}}_{\text{everyone except } i} \in B] \leq e^{\epsilon}\mathbb{P}[\underbrace{M(t_i', \mathbf{t}_{-i})_{-i}}_{\text{everyone except } i} \in B]$$

where $M(\mathbf{t})_{-i} = (M_1(\mathbf{t}), \cdots, M_{i-1}(\mathbf{t}), M_{i+1}(\mathbf{t}), \cdots, M_n(\mathbf{t}))$

Relaxation of DP - Joint DP [KPRU'14]

Definition

A randomized algorithm $M : \mathcal{T}^n \to A^n$ is ϵ-JDP if for every player i, $\mathbf{t}_{-i} \in \mathcal{T}^{n-1}$, $t_i, t_i' \in \mathcal{T}$, and all outcome sets $B \subseteq A^{n-1}$,

$$\mathbb{P}[\underbrace{M(t_i, \mathbf{t}_{-i})_{-i}}_{\text{everyone except } i} \in B] \leq e^{\epsilon}\mathbb{P}[\underbrace{M(t_i', \mathbf{t}_{-i})_{-i}}_{\text{everyone except } i} \in B]$$

where $M(\mathbf{t})_{-i} = (M_1(\mathbf{t}), \cdots, M_{i-1}(\mathbf{t}), M_{i+1}(\mathbf{t}), \cdots, M_n(\mathbf{t}))$

Allows outcome for player i to depend on i's report t_i.

JDP Mediators

Key Property: A JDP mediator that also computes an equilibrium of the underlying game is approximately truthful.

JDP Mediators

Key Property: A JDP mediator that also computes an equilibrium of the underlying game is approximately truthful.

Theorem

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

- *It is ϵ-JDP*
- *For any set of reported types $\mathbf{t}$, it outputs an η-approximate pure strategy Nash Equilibrium.*

Then good behavior is an η'-approximate ex-post Nash Equilibrium for the incomplete information game $\mathcal{G}_M$ where

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

Main Result for Equilibrium Selection

- There *exists* such a mediator from the previous theorem for **large** congestion games.

Main Result for Equilibrium Selection

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

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

Main Result for Equilibrium Selection

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

$$\eta' = \tilde{O}\left(\left(\frac{m^5}{n}\right)^{1/4}\right)$$

Resulting play of good behavior is an η' approximate NE of the complete information game.

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 $a_j \neq a'_j \in A$

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

How to Construct Such a Mechanism?

- Simulate Best Response Dynamics $\Longrightarrow$ obtains a NE in routing games [MS'96].

How to Construct Such a Mechanism?

- Simulate Best Response Dynamics $\implies$ obtains a NE in routing games [MS'96].
- Compute Best Responses privately $\implies$ costs only depend on the number of people on each edge.

How to Construct Such a Mechanism?

- Simulate Best Response Dynamics $\Longrightarrow$ obtains a NE in routing games [MS'96].
- Compute Best Responses privately $\Longrightarrow$ costs only depend on the number of people on each edge.
- Limit the number of times a single player can change routes $\Longrightarrow$ uses the "largeness" assumption.

Billboard Lemma

- If a mechanism $M : \mathcal{T}^n \rightarrow O$ is (ϵ, δ)-DP and consider any function $\phi : \mathcal{T} \times O \rightarrow A$. Define $M' : \mathcal{T}^n \rightarrow A^n$ to be

$$M'(\mathbf{t})_i = \phi(t_i, M(\mathbf{t})).$$

Then M' is (ϵ, δ)- JDP.

Goal Two - Social Welfare

Recall that we want to minimize the cost to all players:

Definition

An action profile **a** is an η-Socially Optimal Routing if

$$C(\mathbf{a}) \leq \min_{\mathbf{a}' \in A^n} C(\mathbf{a}') + \eta$$

where $C(\mathbf{a}) = \sum_{e \in E} y_e \ell_e(y_e)$

Goal Two - Social Welfare

Recall that we want to minimize the cost to all players:

Definition

An action profile **a** is an η-Socially Optimal Routing if

$$C(\mathbf{a}) \leq \min_{\mathbf{a}' \in A^n} C(\mathbf{a}') + \eta$$

where $C(\mathbf{a}) = \sum_{e \in E} y_e \ell_e(y_e)$

How to get selfish agents to play the socially optimal routing without knowing their types?

Classical Approach - Tolls

Add constant tolls $\tau = (\tau_e)_{e \in E}$ *to the edges such that a NE of the game with tolls (tolled game) is the socially optimal in the game without tolls. However these tolls depend on players' types.*

Classical Approach - Tolls

Add constant tolls $\tau = (\tau_e)_{e \in E}$ to the edges such that a NE of the game with tolls (tolled game) is the socially optimal in the game without tolls. However these tolls depend on players' types.

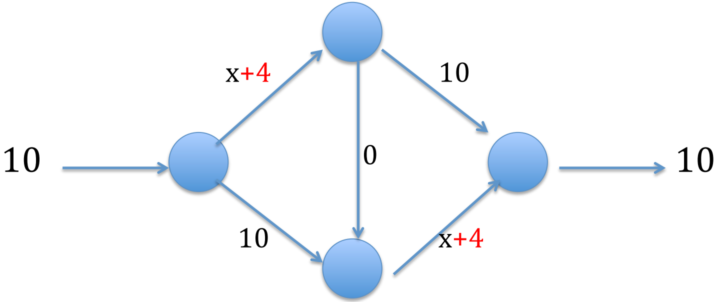

Modified Mediator for Social Welfare

- Mediator still suggests routes to each player $\mathbf{a} = (a_1, \cdots, a_n)$ that they may or may not follow, but it also outputs tolls $\tau = (\tau_e)_{e \in E}$ on each edge, that every player must pay.
- Modified Mediator:

Modified Mediator for Social Welfare

- Mediator still suggests routes to each player $\mathbf{a} = (a_1, \cdots, a_n)$ that they may or may not follow, but it also outputs tolls $\tau = (\tau_e)_{e \in E}$ on each edge, that every player must pay.
- Modified Mediator:

+

E-ZPass SM

JDP Mediators + Tolls

Recall: JDP + NE $\implies$ truthfulness

JDP Mediators + Tolls

Recall: JDP + NE $\implies$ truthfulness

Theorem

Let $M_i : (\mathcal{T} \cup \perp) \to A \times [0, U]^m$ *where* $M_i(\mathbf{t}) = (M_i^A(\mathbf{t}), M^\tau(\mathbf{t}))$ *outputs a suggested route and tolls for each edge. If* $M = (M_1, \cdots, M_n)$ *satisfies both*

- ϵ*-JDP and*
- *for any input types* $\mathbf{t}$*, the action profile* $\mathbf{a} = M^A(\mathbf{t})$ *is an* η*-approximate NE in the modified routing game with*

$$\ell_e^M(\mathbf{y}) = \ell_e(\mathbf{y}) + M_e^\tau(\mathbf{t})$$

then good behavior is an η' *approximate ex-post NE in the mediated tolled game, where*

$$\eta' = \eta + 2m(U + 1)\epsilon$$

Main Result for Social Welfare

- There *exists* a mediator M from the previous theorem for large routing games with convex latencies.

Main Result for Social Welfare

- There *exists* a mediator M from the previous theorem for large routing games with convex latencies.
- Further, with this mediator M, good behavior forms an η' approximate ex-post NE in the mediated tolled game, where

$$\eta' = \tilde{O}\left(\frac{m^{3/2}}{n^{1/5}}\right).$$

Main Result for Social Welfare

- There *exists* a mediator M from the previous theorem for large routing games with convex latencies.
- Further, with this mediator M, good behavior forms an η' approximate ex-post NE in the mediated tolled game, where

$$\eta' = \tilde{O}\left(\frac{m^{3/2}}{n^{1/5}}\right).$$

- Resulting play of good behavior is an $\tilde{O}(mn^{4/5})$ socially optimal routing.

Main Result for Social Welfare

- There *exists* a mediator M from the previous theorem for large routing games with convex latencies.
- Further, with this mediator M, good behavior forms an η' approximate ex-post NE in the mediated tolled game, where

$$\eta' = \tilde{O}\left(\frac{m^{3/2}}{n^{1/5}}\right).$$

- Resulting play of good behavior is an $\tilde{O}(mn^{4/5})$ socially optimal routing.
- As long as the optimal solution grows $\omega(n^{4/5})$, then we get a $(1 + o(1))$ multiplicative approximation to the true optimal.

How to Construct such a Mediator?

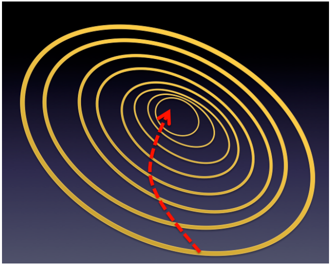

- Compute an approximately optimal flow $\mathbf{a}^{\bullet}$ subject to JDP via a privacy preserving projected gradient descent algorithm.

How to Construct such a Mediator?

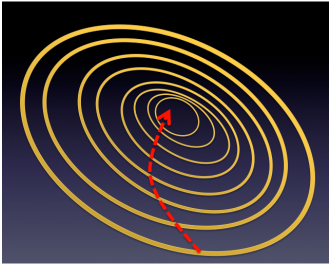

- Compute an approximately optimal flow $\mathbf{a}^\bullet$ subject to JDP via a privacy preserving projected gradient descent algorithm.
- Given target flow $\mathbf{a}^\bullet$, we find the necessary tolls $\hat{\tau}$ so that most players are nearly best responding in this tolled game when playing $\mathbf{a}^\bullet$.

How to Construct such a Mediator

- Allow the few players that are not approximately best responding to then best respond in the tolled game. This will modify $\mathbf{a}^{\bullet}$ only slightly (by largeness assumption) and so will remain nearly optimal in the original game.

How to Construct such a Mediator

- Allow the few players that are not approximately best responding to then best respond in the tolled game. This will modify $\mathbf{a}^{\bullet}$ only slightly (by largeness assumption) and so will remain nearly optimal in the original game.
- Resulting flow $\hat{\mathbf{a}}$ is then nearly optimal in the original game and an approximate NE in the tolled game with tolls $\hat{\tau}$.

Recap

Key Property: NE + JDP $\implies$ Approx Truthful Mediator

Recap

Key Property: NE + JDP $\implies$ Approx Truthful Mediator

Equilibrium Selection Mediator

Main Result:
Output routing that is a $\tilde{o}(1)$-*NE*

Recap

Key Property: NE + JDP $\Longrightarrow$ Approx Truthful Mediator

Equilibrium Selection Mediator

Main Result:

Output routing that is a $\tilde{o}(1)$*-NE*

Social Welfare Mediator

Main Result:

Output routing within $(1 + \tilde{o}(1))$ *OPT*

\+

E-ZPass

QUESTIONS?