Algorithmic Sustainability for Resource Allocation in Large-scale Sociotechnical Systems

Devansh Jalota

Stanford → Columbia → Georgia Tech (Postdoc) (ISyE)

> eMERGE Seminar October 6, 2025

The need for market mechanisms in allocating limited resources

Market mechanisms are a necessary tool to allocate limited resources in our resource constrained world

We live in a resource constrained world









Limited Vaccine Availability

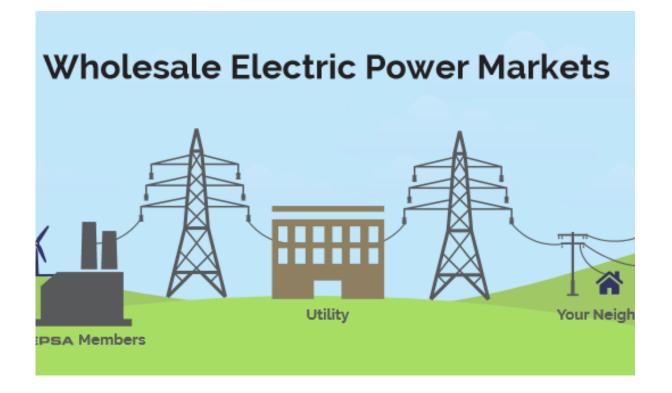
How do we best allocate scarce resources to strategic agents?

Market design involves designing mechanisms to set the rules of the marketplace, e.g., via pricing, to ensure the right resources go to the right people while aligning users' incentives

Pervasiveness of market mechanisms

Market mechanisms are pervasive and are undergoing major transformations in our increasingly data-driven society

Market mechanisms influence our everyday lives



Electricity Markets



Congestion Pricing



Vaccine Allocation

Computing advances have opened new avenues for market design



Online Advertising



Marketplace Platforms



Smart Mobility

Challenges for AI and data-driven market mechanisms

Fairness/Equity



Why congestion pricing is a tough sell, even if it's good policy

By Christian Hetrick June

Market mechanisms such as congestion pricing raise fairness/equity concerns and, hence, often remain unimplemented

Data Privacy and Uncertainty



DOT plans to implement V2X technology raise privacy concerns

Kimberly Adams

Aug 20, 2024

Data is often plagued with uncertainty and raises privacy concerns that render mechanisms relying on complete information inadequate

Security



GAO finds fraud in commuter program

It's a perk of federal employment: a free monthly subsidy that pays for commutes on public transportation. But scores of workers have been taking the government for a ride, selling their

Data abundance enables easy information manipulation, leading to novel security challenges

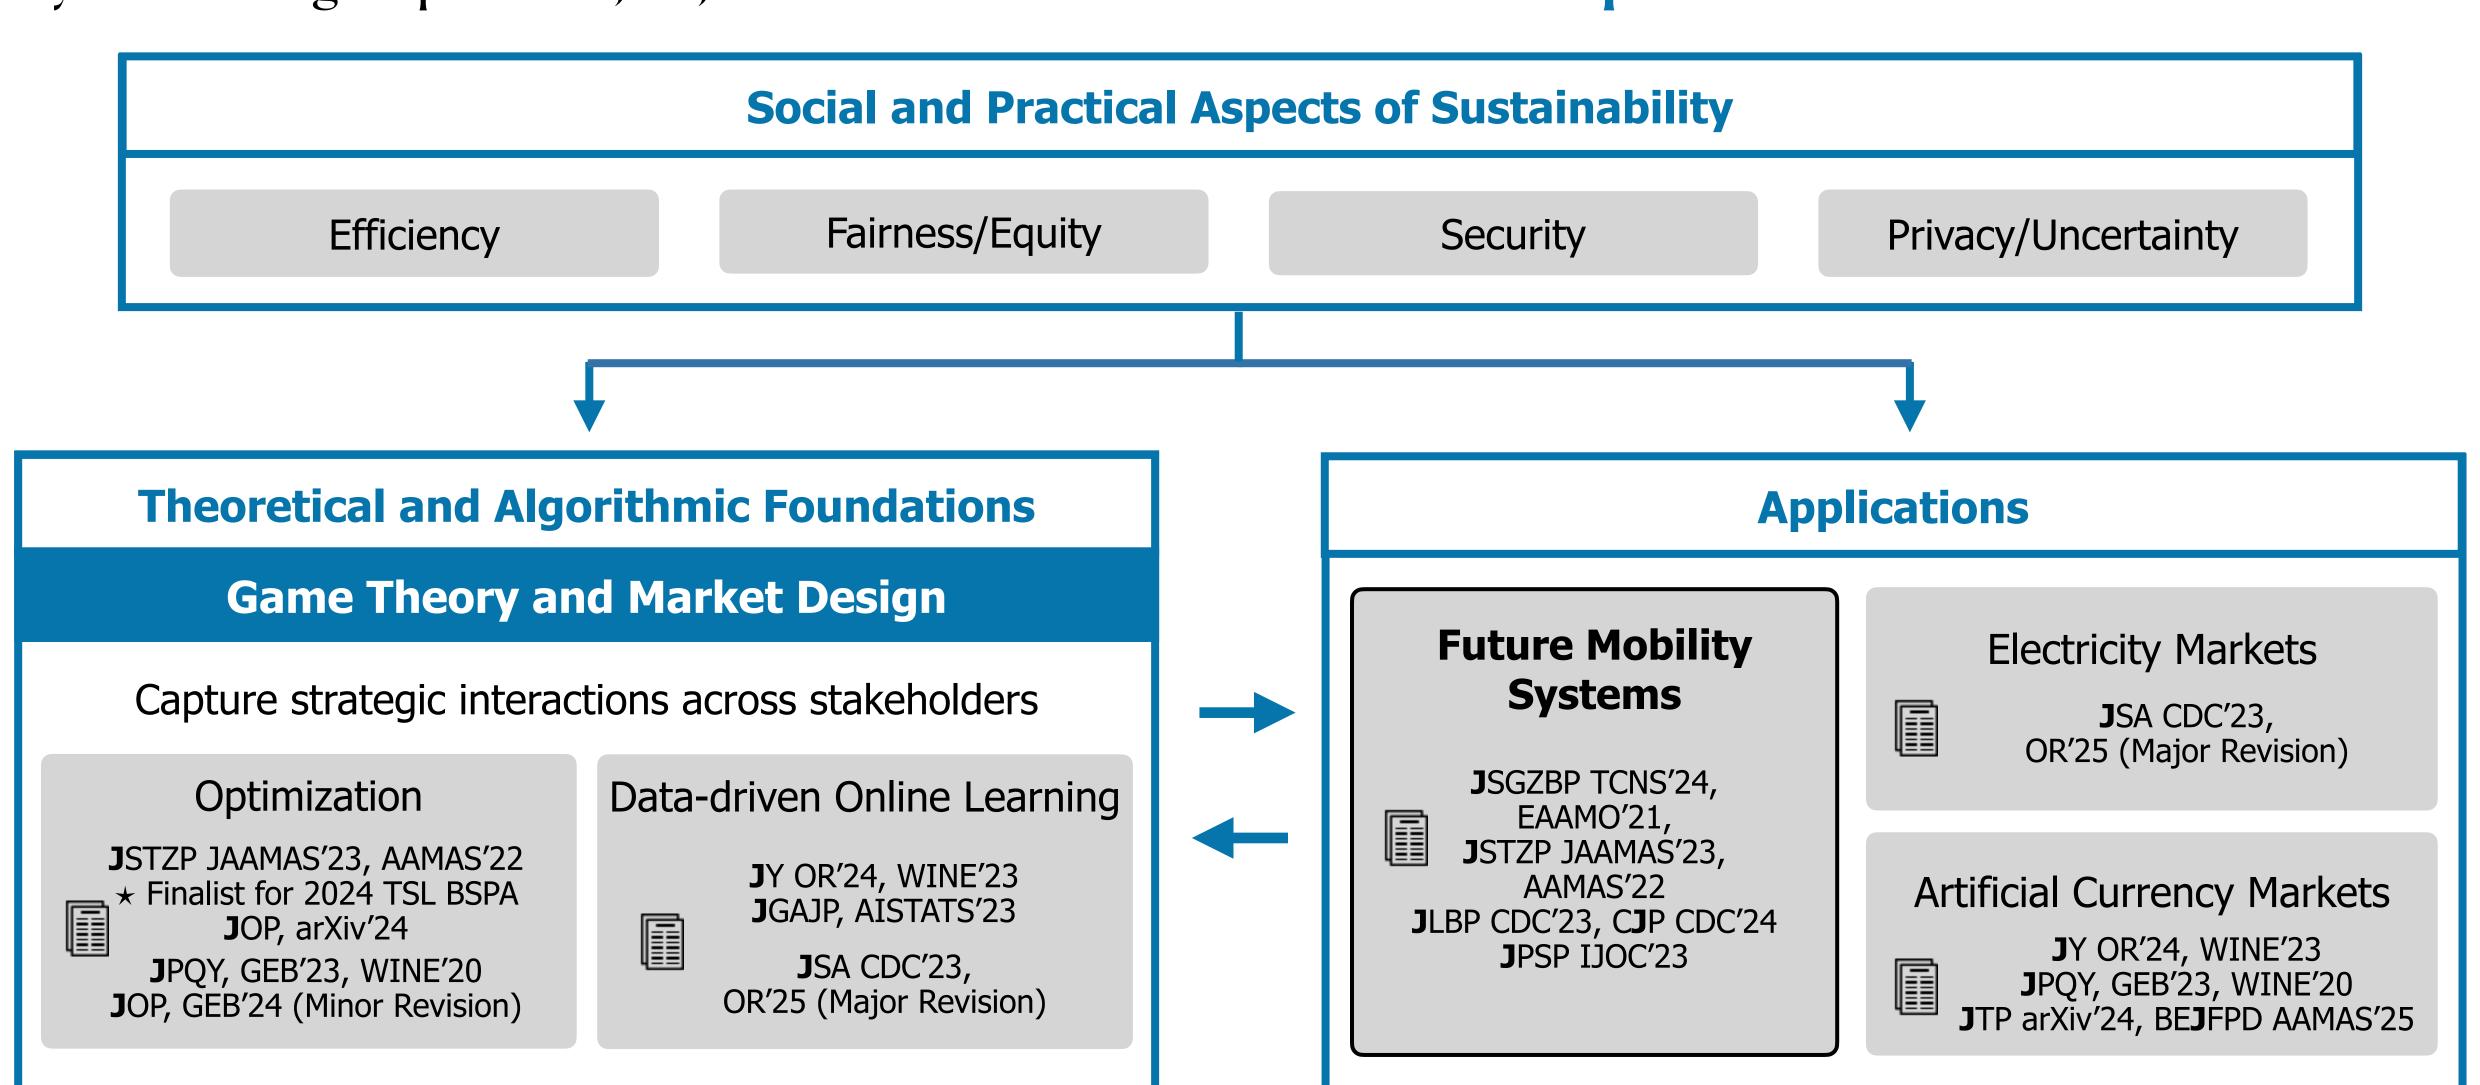
Societal considerations make resource allocation challenging

Traditional Market Mechanisms New Technical Challenges New Societal Considerations Fairness/Equity Efficiency Computing optimal allocations or equilibria is often intractable under new constraints Security Focus of classical resource allocation mechanisms Complete information on user attributes is often necessary for the Privacy/Uncertainty efficacy of market mechanisms

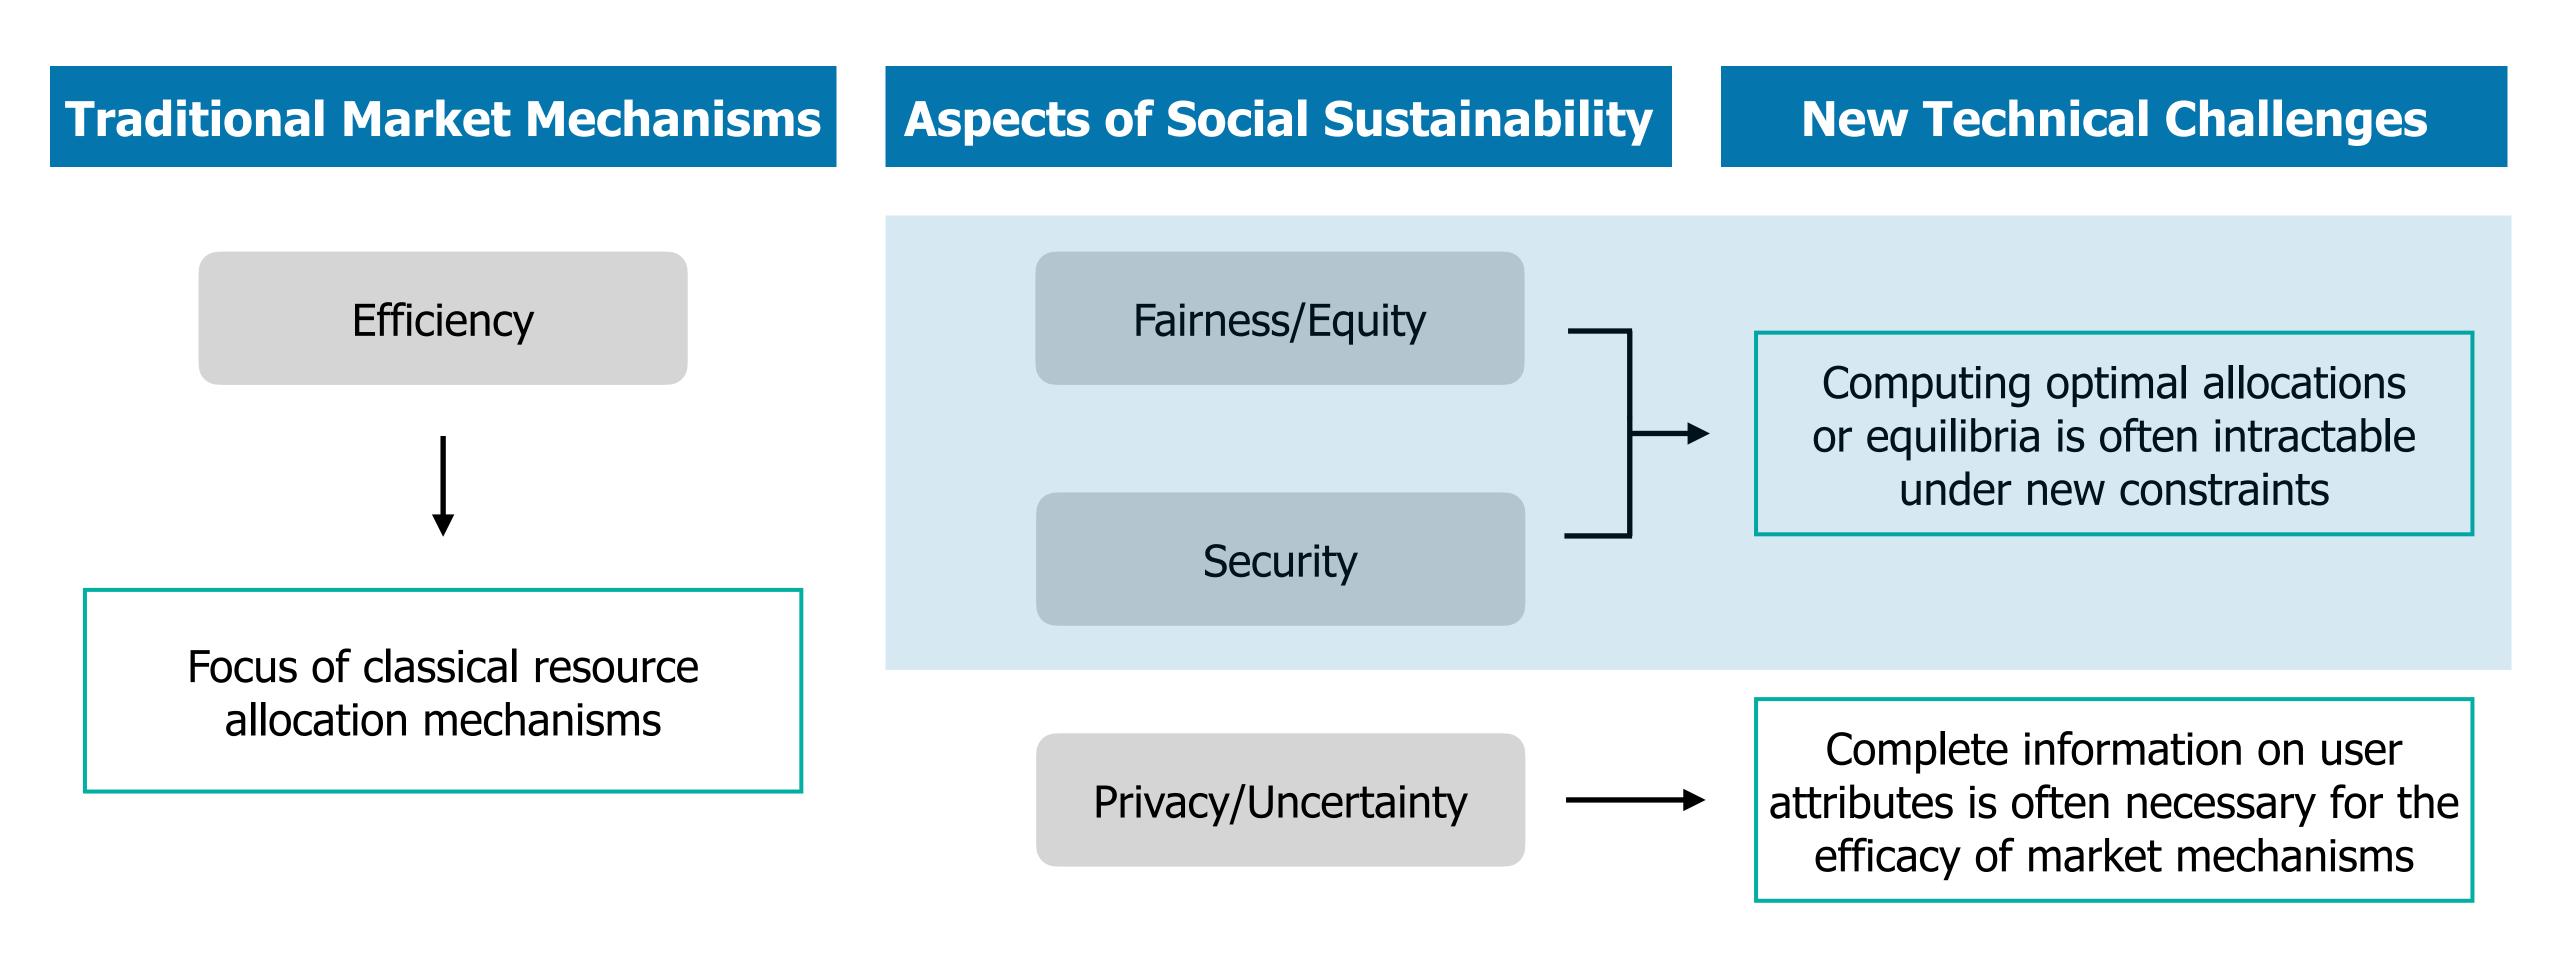
We need novel tools and methods to overcome these technical challenges

My Work: Sustainable resource allocation in socio-technical systems

My work leverages operations, CS, and economics to advance the science and practice of sustainable resource allocation



This talk: Addressing intractability under fairness & security constraints



We need novel tools and methods to overcome these technical challenges

This talk: Fairness & security from approximation algorithm lens

We highlight two vignettes of this agenda focused on developing approximation algorithms for NP-hard problems arising under fairness and security considerations in mobility applications

Fairness (Congestion Pricing)

Jalota et al. JAAMAS'23

Finalist for the 2024 INFORMS

TSL best student paper award

Source of Hardness	Non-convexity of fairness constraints		
Solution Method	Novel convex programming relaxation that also yields a natural pricing scheme		

Security Games (Parking Enforcement)

Jalota et al. arXiv'24

Collaboration with **Stanford's DPS** to demonstrate efficacy of our algorithms on enforcement data in increasing parking revenues by \$300,000 annually

Source of Hardness	Resource constraints on the available set of security officers
Solution Method	Optimizing a specific concave approximation of bi-level program yielding a natural greedy algorithm

- Fairness in Congestion Pricing
- Security Games for Enforcement in Mobility Applications
- Future Directions



https://sites.google.com/view/devanshjalota/

Fairness in Congestion Pricing

- Introduction and Motivating Example
- Model and Problem Setup
- Algorithmic Approach for Fair Traffic Routing: I-TAP
- Fair Congestion Pricing via I-TAP
- Conclusion
- Security Games for Enforcement in Mobility Applications
- Future Directions

- -> Fairness in Congestion Pricing
 - Introduction and Motivating Example
 - Model and Problem Setup
 - Algorithmic Approach for Fair Traffic Routing: I-TAP
 - Fair Congestion Pricing via I-TAP
 - Conclusion
- Security Games for Enforcement in Mobility Applications
- Future Directions

Traffic congestion is a global challenge

Traffic congestion results in billions of dollars of economic losses, with commuters losing hundreds of hours, every year



Traffic congestion is ubiquitous





A Solution: Congestion Pricing (CP)

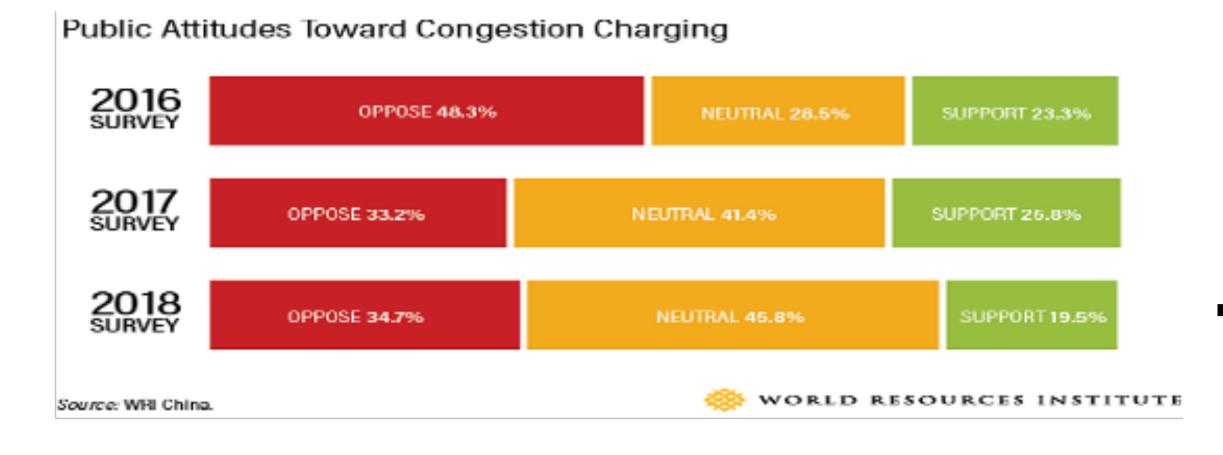




Fairness/Inequity Issues of Congestion Pricing (CP)

However, the real-world deployment of CP has been limited due to the resulting fairness/inequity issue

Regressive Nature of Congestion Pricing



Congestion fees have often been described as a "tax on the working class"

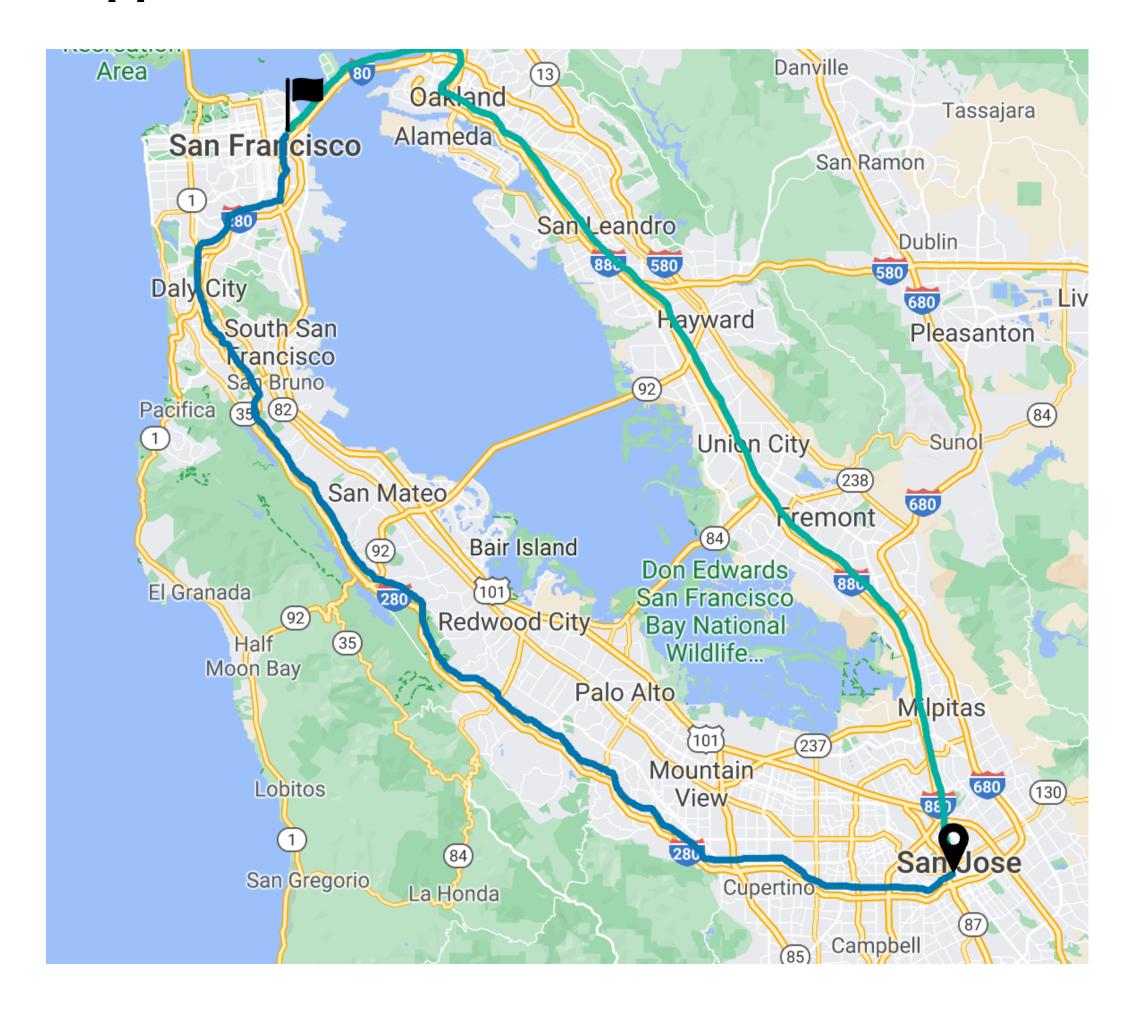
Limited Adoption of Congestion Pricing



Tension between the efficiency gains of congestion pricing and its possible impact on low-income commuters

Example: A Traffic Routing Problem

Suppose we need to route five units of demand between San Jose and San Francisco



Problem Parameters

Travel Demand

 x_1 , x_2 : Flows on route 1 and 2

$$x_1 + x_2 = 5$$

Travel Time

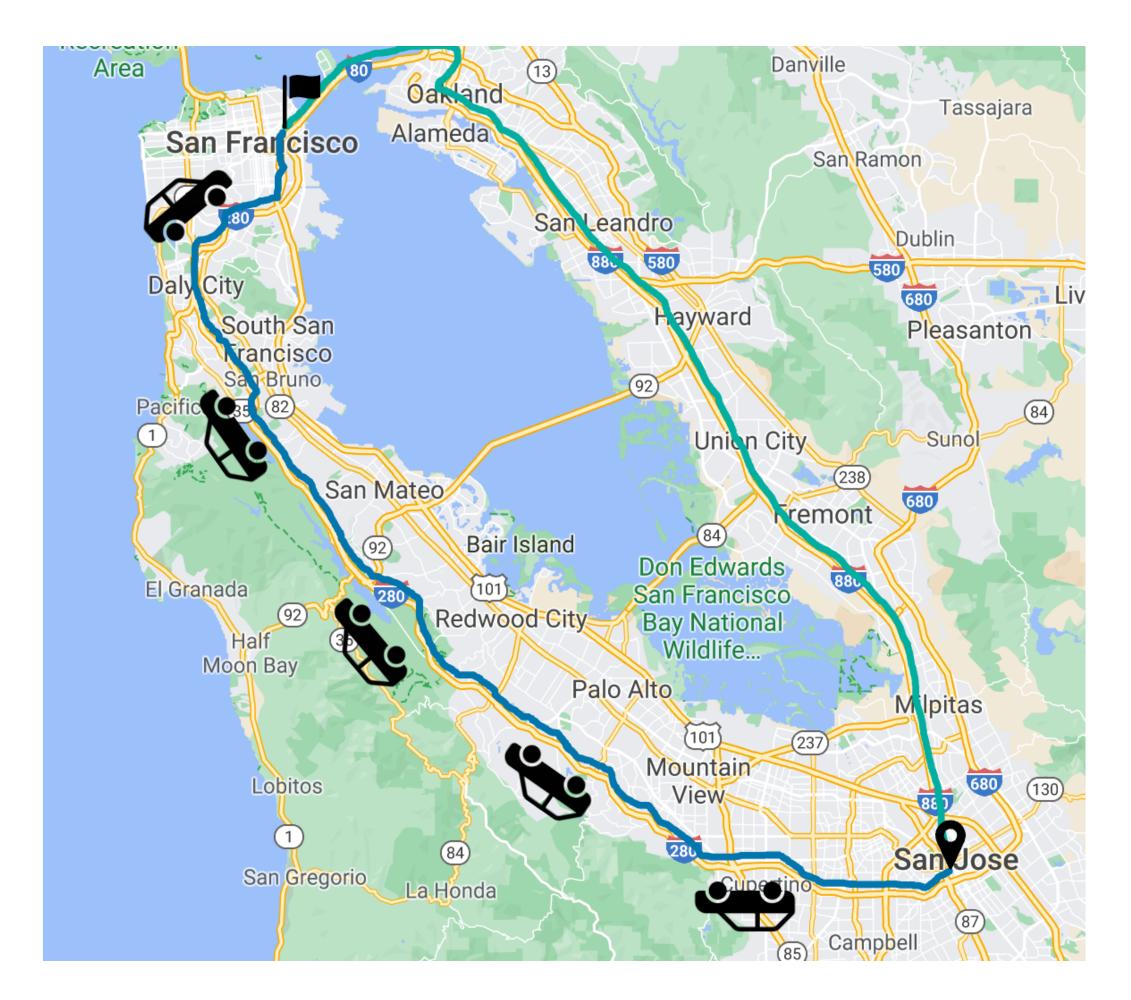
 $t_1(\cdot)$: Travel time function on route 1

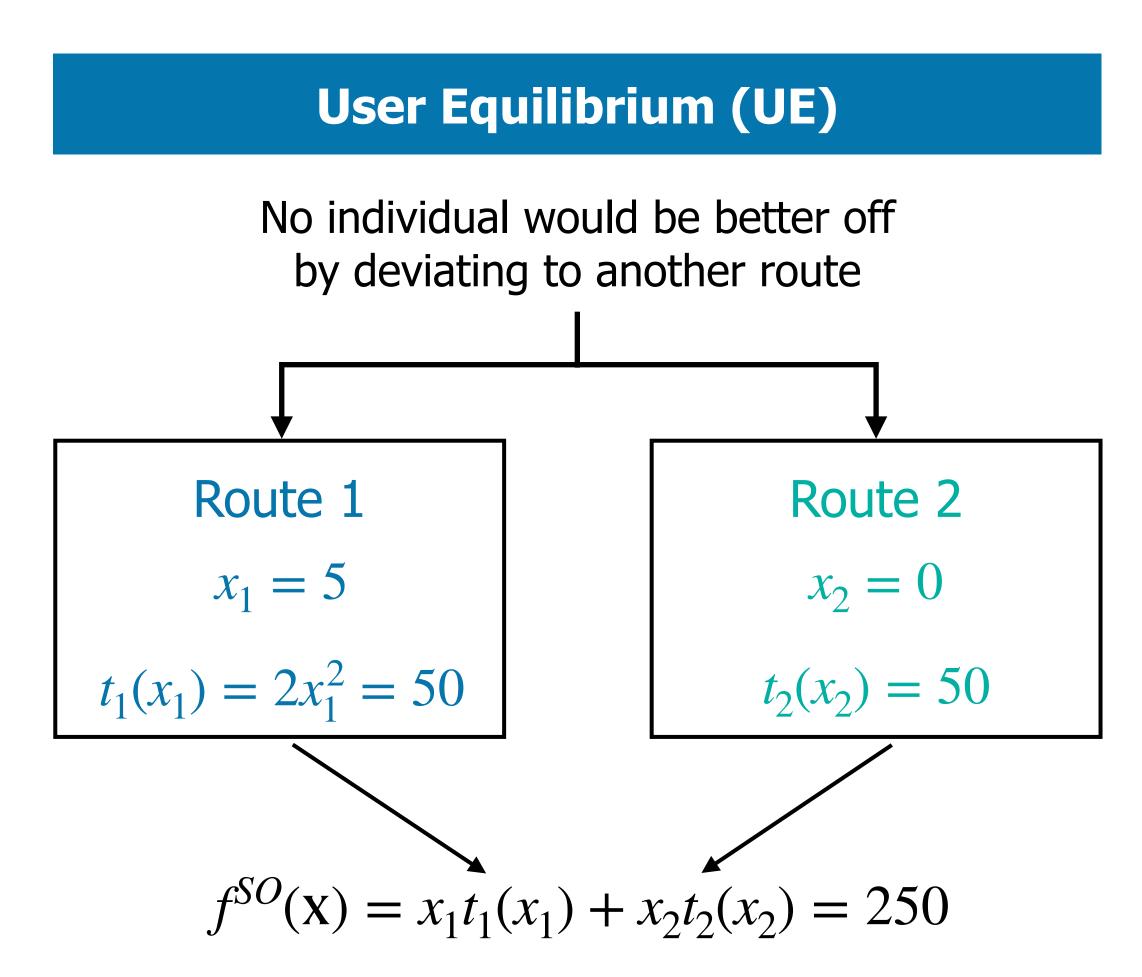
$$t_1(x_1) = 2x_1^2$$

 $t_2(\cdot)$: Travel time function on route 2 $t_2(x_2) = 50$

Example: User Equilibrium (UE) without Tolls

Without tolls, users selfishly minimize their travel times, resulting in potentially large system travel times

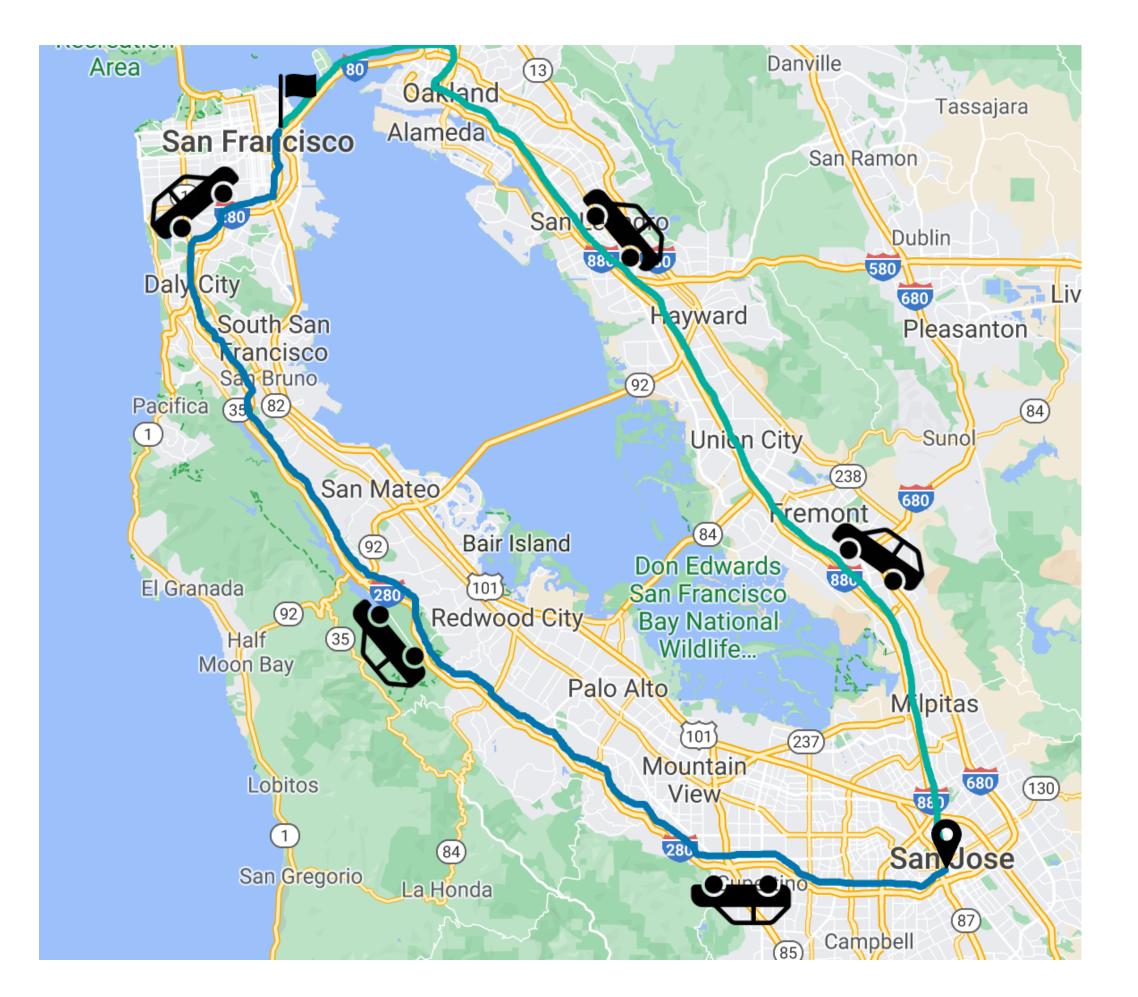


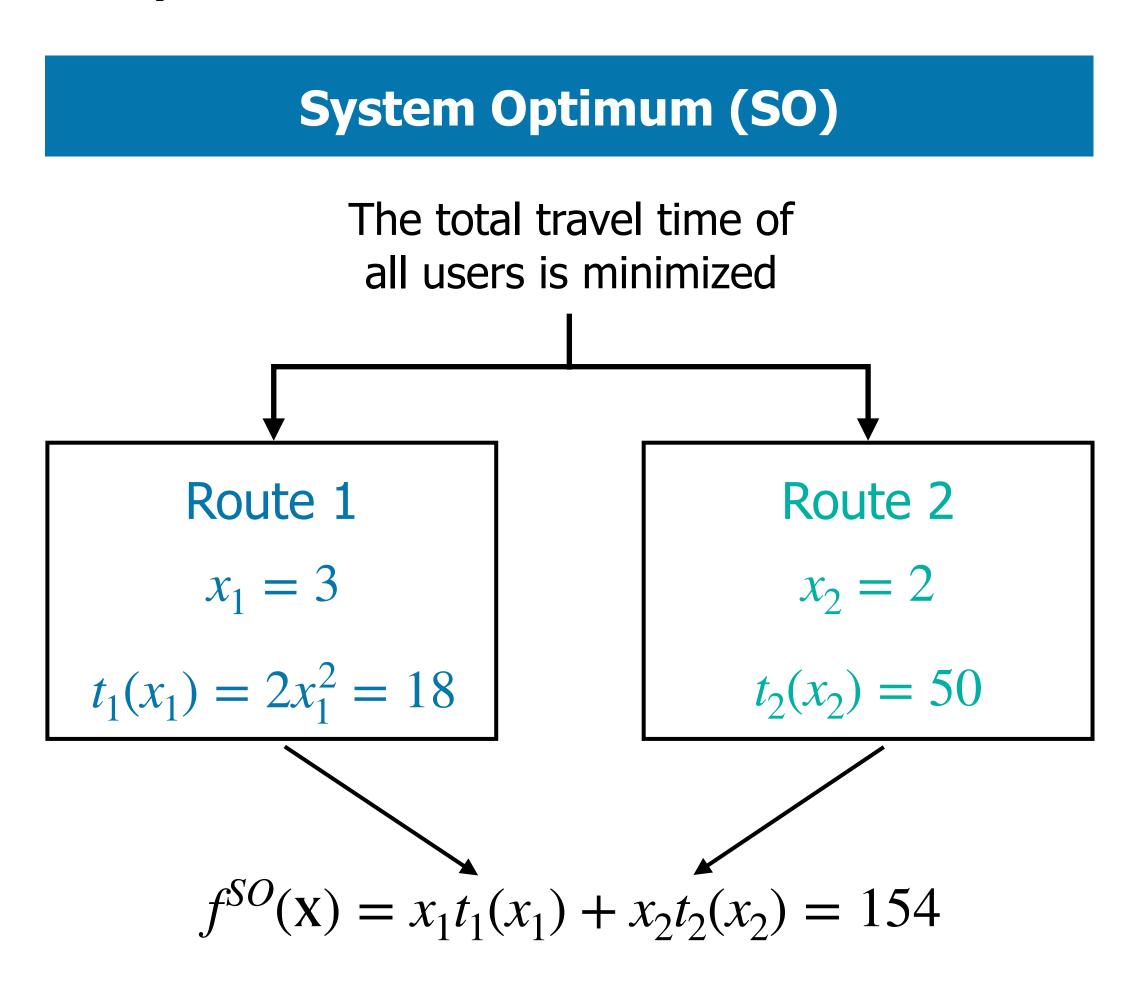


The user equilibrium outcome can result in inefficient outcomes, with high total system travel times

Example: System Optimum (SO) with Congestion Pricing

Marginal cost pricing can be used to steer the traffic to the system optimum





The system optimum outcome can result in unfair outcomes, with some users, typically those with lower incomes, incurring disproportionately high travel times in the pursuit of system efficiency

An Efficiency-Fairness Tradeoff

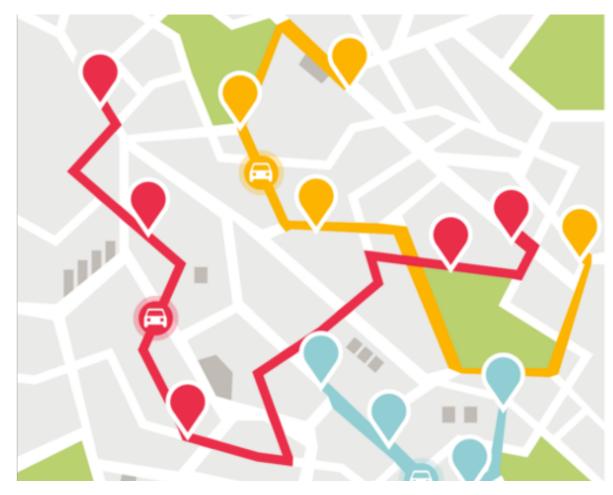
Comparing these two extreme solutions reveals an efficiency-fairness tradeoff

	UE (No Tolls)	SO (Marginal Cost Pricing)		
Total Travel Time	250	154		
Ratio of User Travel Times	1	≈ 3		
	Fair but Inefficient	Efficient but Unfair		

Contributions

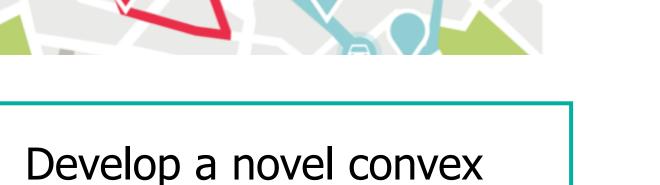
We resolve this tradeoff through an novel traffic routing algorithm and price roads to enforce these flows

Computationally Efficient Routing Algorithm



program (I-TAP) that trades off

efficiency and fairness





Present the first congestion pricing schemes for fair routing in general networks

Fair Congestion Pricing Scheme



Balance Fairness and Efficiency in Traffic Routing



Achieve sustainable transportation's efficiency and fairness objectives

- Fairness in Congestion Pricing
 - Introduction and Motivating Example
 - Model and Problem Setup
 - Algorithmic Approach for Fair Traffic Routing: I-TAP
 - Fair Congestion Pricing via I-TAP
 - Conclusion
- Security Games for Enforcement in Mobility Applications
- Future Directions

Performance Metrics

We evaluate the performance of a traffic routing policy in general road networks through two metrics

Efficiency

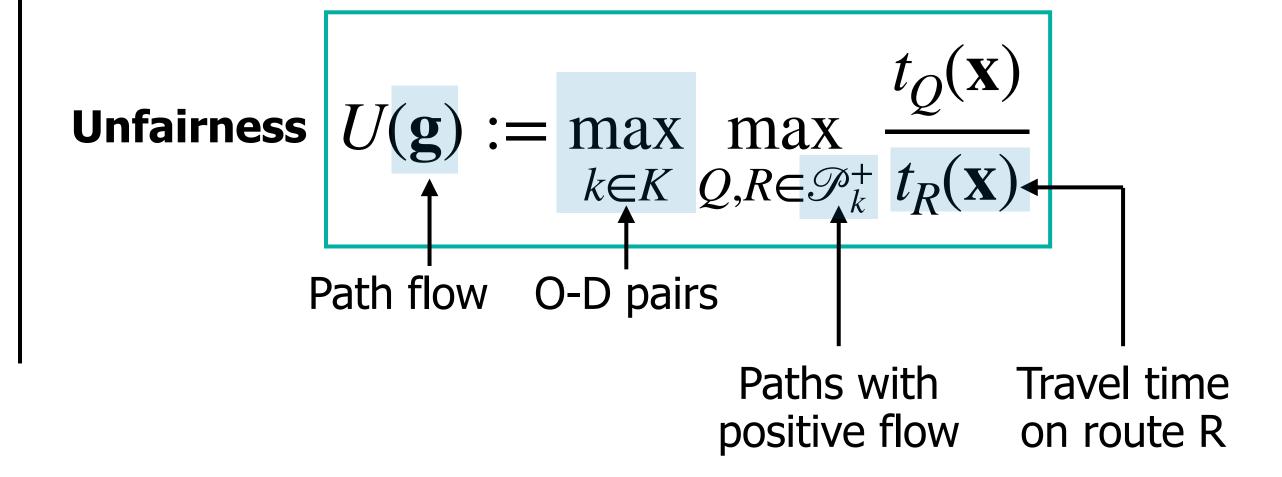
Ratio of total travel time of a traffic assignment \mathbf{x} to that of the system optimum solution \mathbf{x}^{SO}

Inefficiency Ratio

$$\rho(\mathbf{x}) := \frac{f^{SO}(\mathbf{x})}{f^{SO}(\mathbf{x}^{SO})}$$

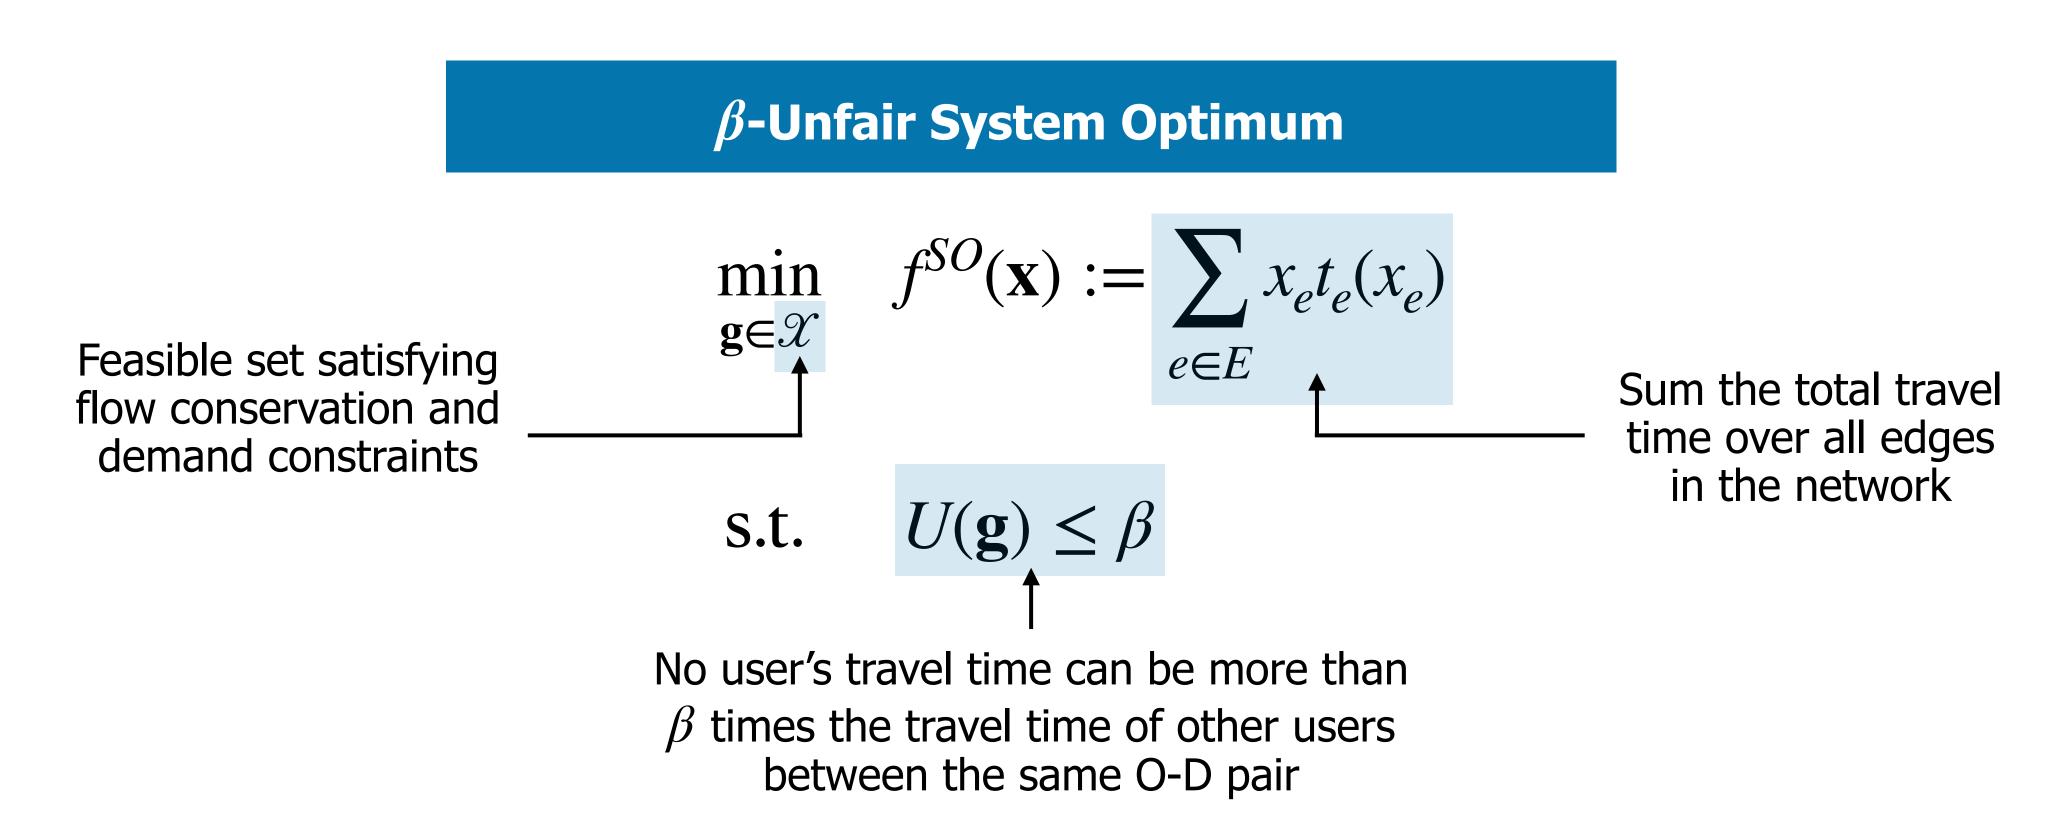
Fairness

Maximum ratio of the travel times of users traveling between the same O-D pair



Unfairness constrained system optimum

To balance efficiency and fairness, we study an unfairness constrained system optimum problem

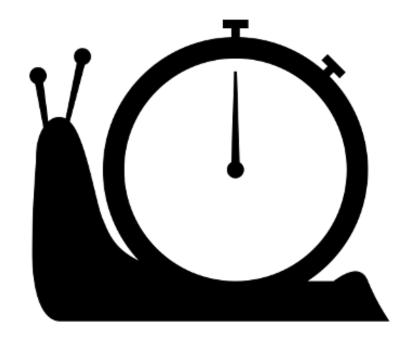


 β -Unfair SO is NP-hard [Basu et al. 2017]: The unfairness constraints are non-convex

Related Literature

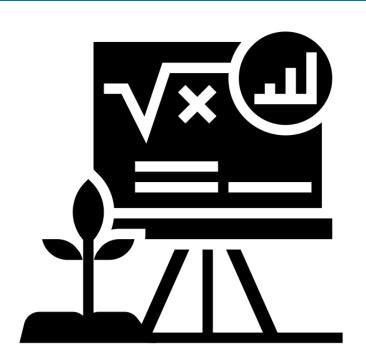
Past literature on fair routing to solve the unfairness constrained system optimum suffer from several drawbacks

Computationally Prohibitive



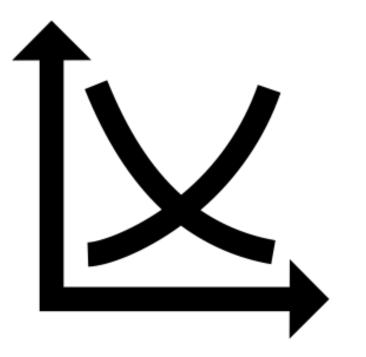
Unlike methods relying on solving MILPs or NP-hard sub-problems, we propose a computationally efficient convex programming approach

Lack of Theoretical Guarantees



We present theoretical bounds on the unfairness and inefficiency ratio to evaluate the performance of our algorithm

No Pricing Mechanism to Enforce Flows



Our convex program yields natural congestion pricing schemes, resulting in the first study of pricing for fair routing in general networks

Our work overcomes several drawbacks of existing approaches to solve the unfairness constrained system optimum

- Fairness in Congestion Pricing
 - Introduction and Motivating Example
 - Model and Problem Setup
 - Algorithmic Approach for Fair Traffic Routing: I-TAP
 - Fair Congestion Pricing via I-TAP
 - Conclusion
- Security Games for Enforcement in Mobility Applications
- Future Directions

Our approach: Interpolated Traffic Assignment

Our approach interpolates between the UE and SO objectives to achieve a balance between efficiency and fairness

System Optimum Traffic Assignment Problem

Minimize Total Travel time

$$\min_{\mathbf{g} \in \mathcal{X}} f^{SO}(\mathbf{x}) := \sum_{e \in E} x_e t_e(x_e)$$
 Inefficiency Ratio: $\rho(\mathbf{x}) = 1$ but Unfair but Unfair $U(\mathbf{g}) \geq 1$

User Equilibrium Traffic Assignment Problem

Users between each O-D pair have same travel time

$$\min_{\mathbf{g} \in \mathcal{X}} f^{UE}(\mathbf{x}) := \sum_{e \in E} \int_0^{x_e} t_e(y) dy \qquad \qquad \text{Inefficiency Ratio: } \rho(\mathbf{x}) \geq 1 \\ \text{Unfairness: } U(\mathbf{g}) = 1 \qquad \qquad \text{Inefficiency Ratio: } \rho(\mathbf{x}) \geq 1 \\ \text{Unfairness: } U(\mathbf{g}) = 1 \qquad \qquad \text{Inefficiency Ratio: } \rho(\mathbf{x}) \geq 1 \\ \text{Unfairness: } U(\mathbf{g}) = 1 \qquad \qquad \text{Inefficiency Ratio: } \rho(\mathbf{x}) \geq 1 \\ \text{Unfairness: } U(\mathbf{g}) = 1 \qquad \qquad \text{Inefficiency Ratio: } \rho(\mathbf{x}) \geq 1 \\ \text{Unfairness: } U(\mathbf{g}) = 1 \qquad \qquad \text{Inefficiency Ratio: } \rho(\mathbf{x}) \geq 1 \\ \text{Unfairness: } U(\mathbf{g}) = 1 \qquad \qquad \text{Inefficiency Ratio: } \rho(\mathbf{x}) \geq 1 \\ \text{Unfairness: } U(\mathbf{g}) = 1 \qquad \qquad \text{Inefficiency Ratio: } \rho(\mathbf{x}) \geq 1 \\ \text{Unfairness: } U(\mathbf{g}) = 1 \qquad \qquad \text{Inefficiency Ratio: } \rho(\mathbf{x}) \geq 1 \\ \text{Unfairness: } U(\mathbf{g}) = 1 \qquad \qquad \text{Inefficiency Ratio: } \rho(\mathbf{x}) \geq 1 \\ \text{Unfairness: } U(\mathbf{g}) = 1 \qquad \qquad \text{Inefficiency Ratio: } \rho(\mathbf{x}) \geq 1 \\ \text{Unfairness: } U(\mathbf{g}) = 1 \qquad \qquad \text{Inefficiency Ratio: } \rho(\mathbf{x}) \geq 1 \\ \text{Unfairness: } U(\mathbf{g}) = 1 \qquad \qquad \text{Inefficiency Ratio: } \rho(\mathbf{x}) \geq 1 \\ \text{Unfairness: } \rho(\mathbf{$$

Interpolated Traffic Assignment Problem (I-TAP)

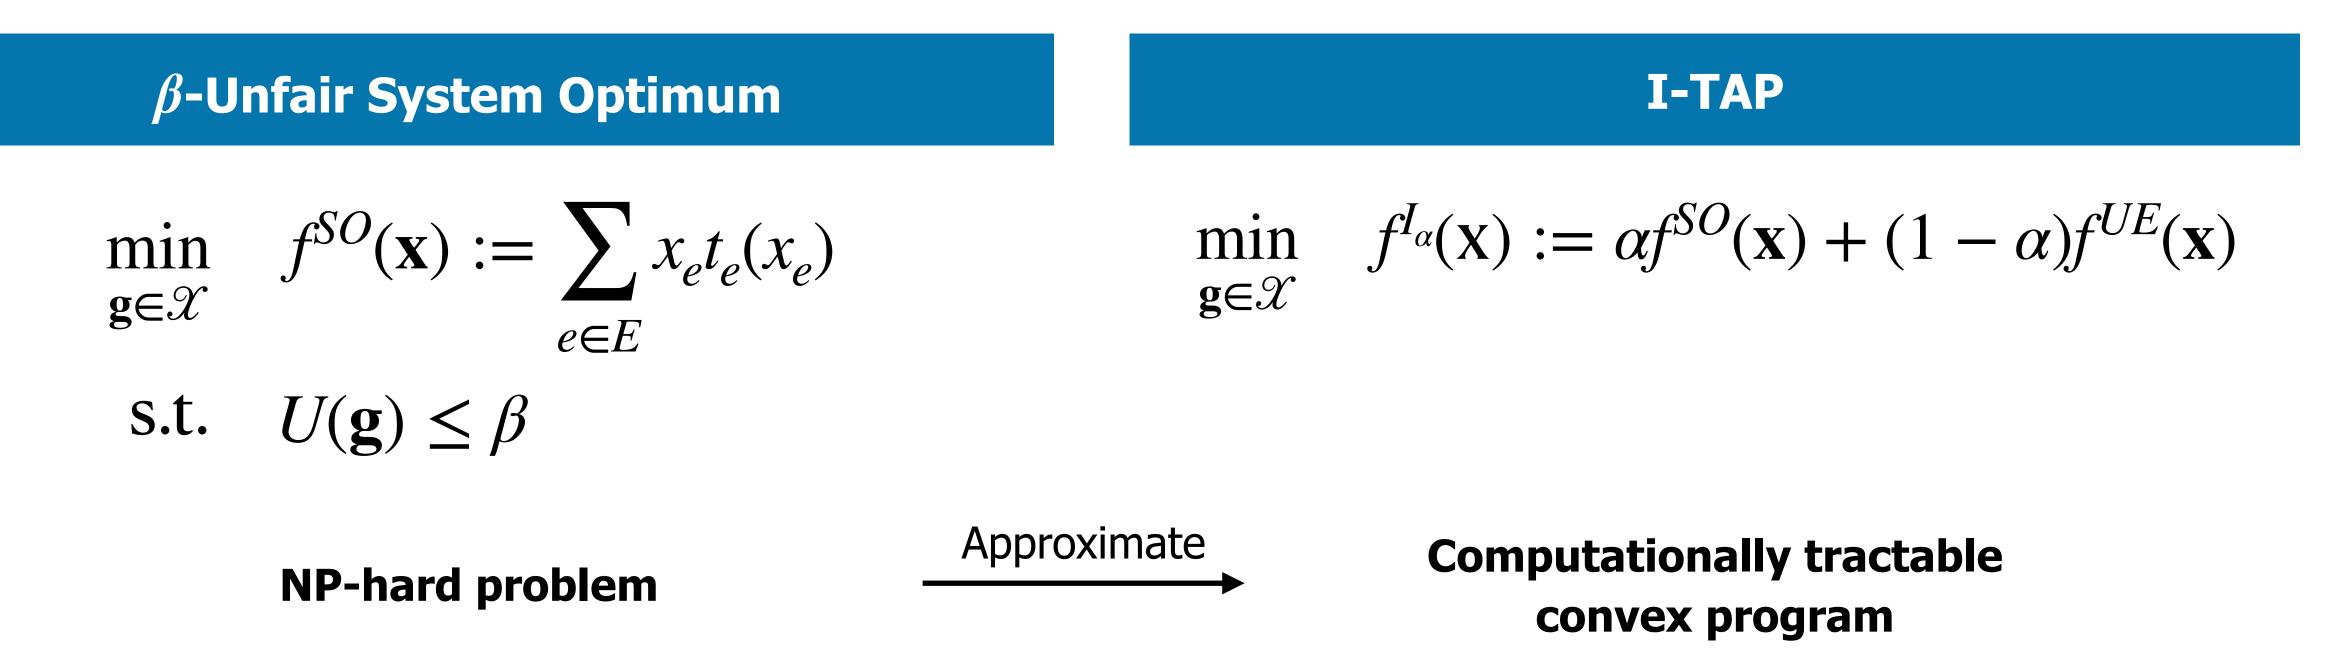
Interpolate between UE and SO objectives

$$\min_{\mathbf{g} \in \mathcal{X}} f^{I_{\alpha}}(\mathbf{x}) := \alpha f^{SO}(\mathbf{x}) + (1 - \alpha) f^{UE}(\mathbf{x})$$

Achieve best of both worlds

Approximation β -SO via I-TAP

I-TAP serves as a tractable approximation to the β -SO problem we seek to solve

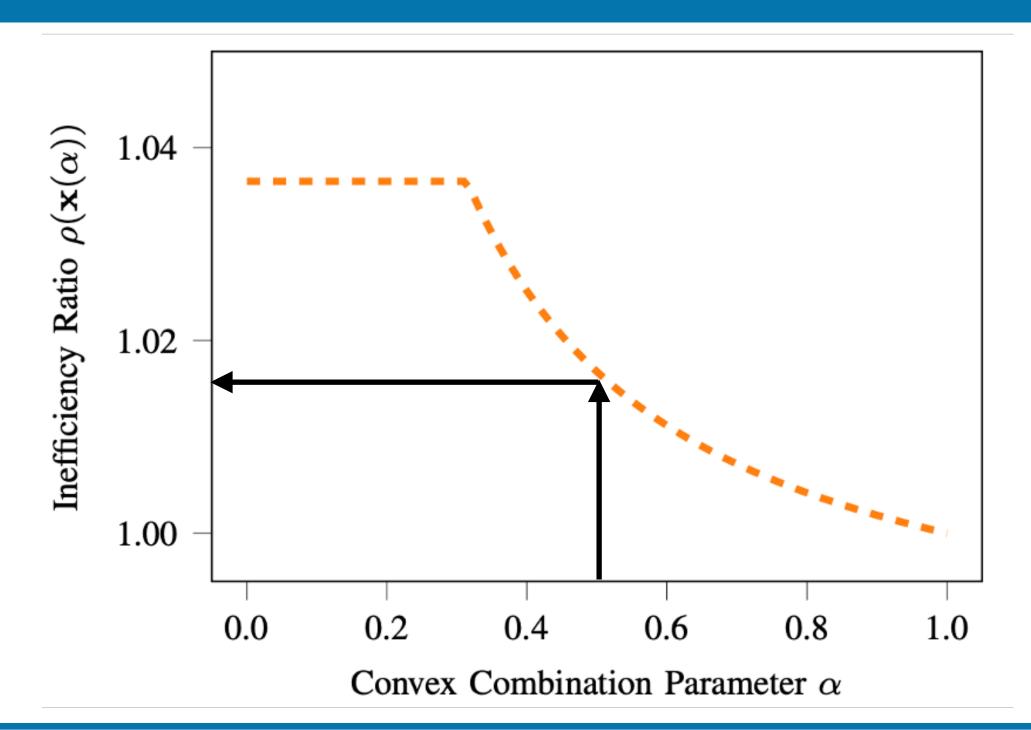


How good an approximation is I-TAP to the β -Unfair SO problem?

Theoretical Bounds for I-TAP

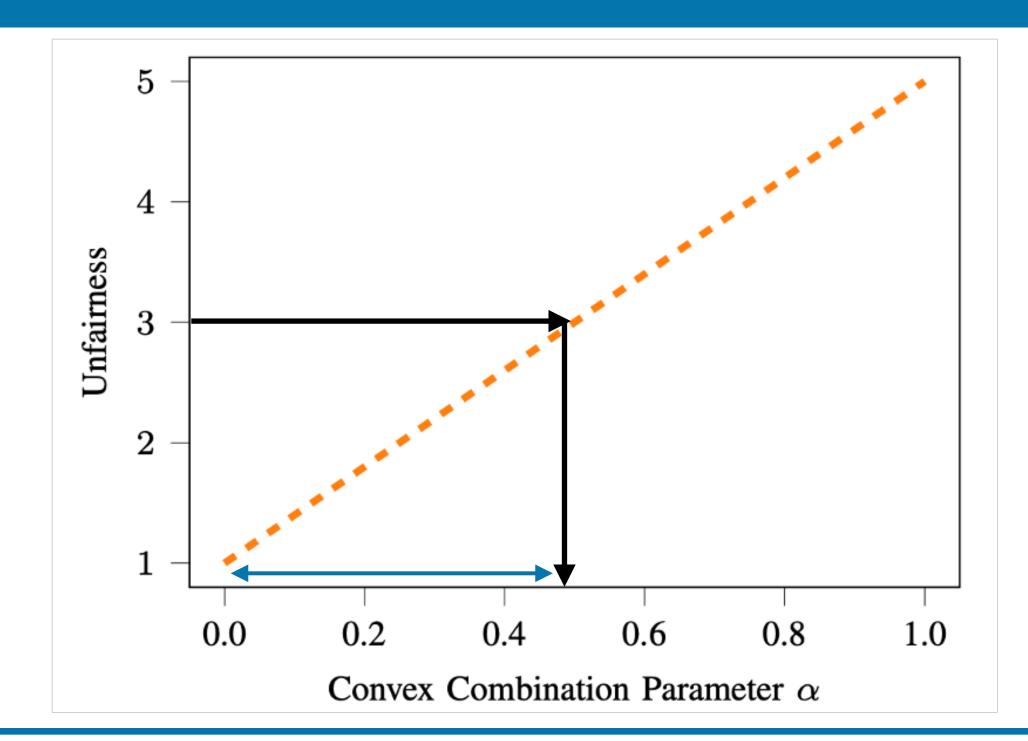
We obtain bounds on the total travel time and unfairness of I-TAP

Inefficiency Ratio



Theorem: For any $\alpha \in [0,1]$, the inefficiency ratio $\rho(\mathsf{X}(\alpha)) \leq \min\{\mathsf{PoA},1 + \frac{(1-\alpha)(f^{\mathit{UE}}(\mathsf{X}(1)) - f^{\mathit{UE}}(\mathsf{X}(0)))}{\alpha f^{\mathit{SO}}(\mathsf{X}(1))}\}.$

Unfairness



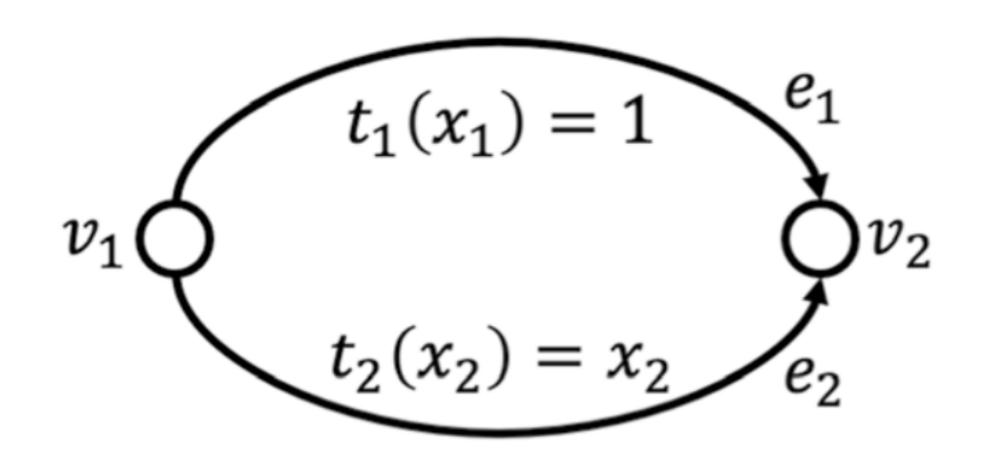
Theorem: Let m be the degree of the polynomial of the travel time functions. Then, $U(\mathbf{g}(\alpha)) \leq \beta$ if $\alpha \leq \frac{\beta-1}{m}$.

This bound is tight.

Optimality of I-TAP

I-TAP achieves the optimal solution for the β -SO problem for any two-edge Pigou network

Theorem: For any two-edge Pigou network and unfairness parameter β , there exists some convex combination parameter α^* , such that the solutions of I-TAP with parameter α^* and that of the β -SO problem coincide



 β -Unfair SO for Pigou Network



I-TAP for Pigou Network

Experiments: Computational Savings of I-TAP

Solving I-TAP is three orders of magnitude more computationally efficient than a state-of-the-art benchmark

	attributes		runtime (sec.)		
Region Name	V	E	K	Jahn et al.	I-TAP
Sioux Falls (SF)	24	76	528	20.0	0.03
Anaheim (A)	416	914	1406	74.0	0.33
Massachusetts (M)	74	258	1113	24.3	0.09
Tiergarten (T)	361	766	644	18.2	0.20
Friedrichshain (F)	224	523	506	19.8	0.12
Prenzlauerberg (P)	352	749	1406	74.4	0.32

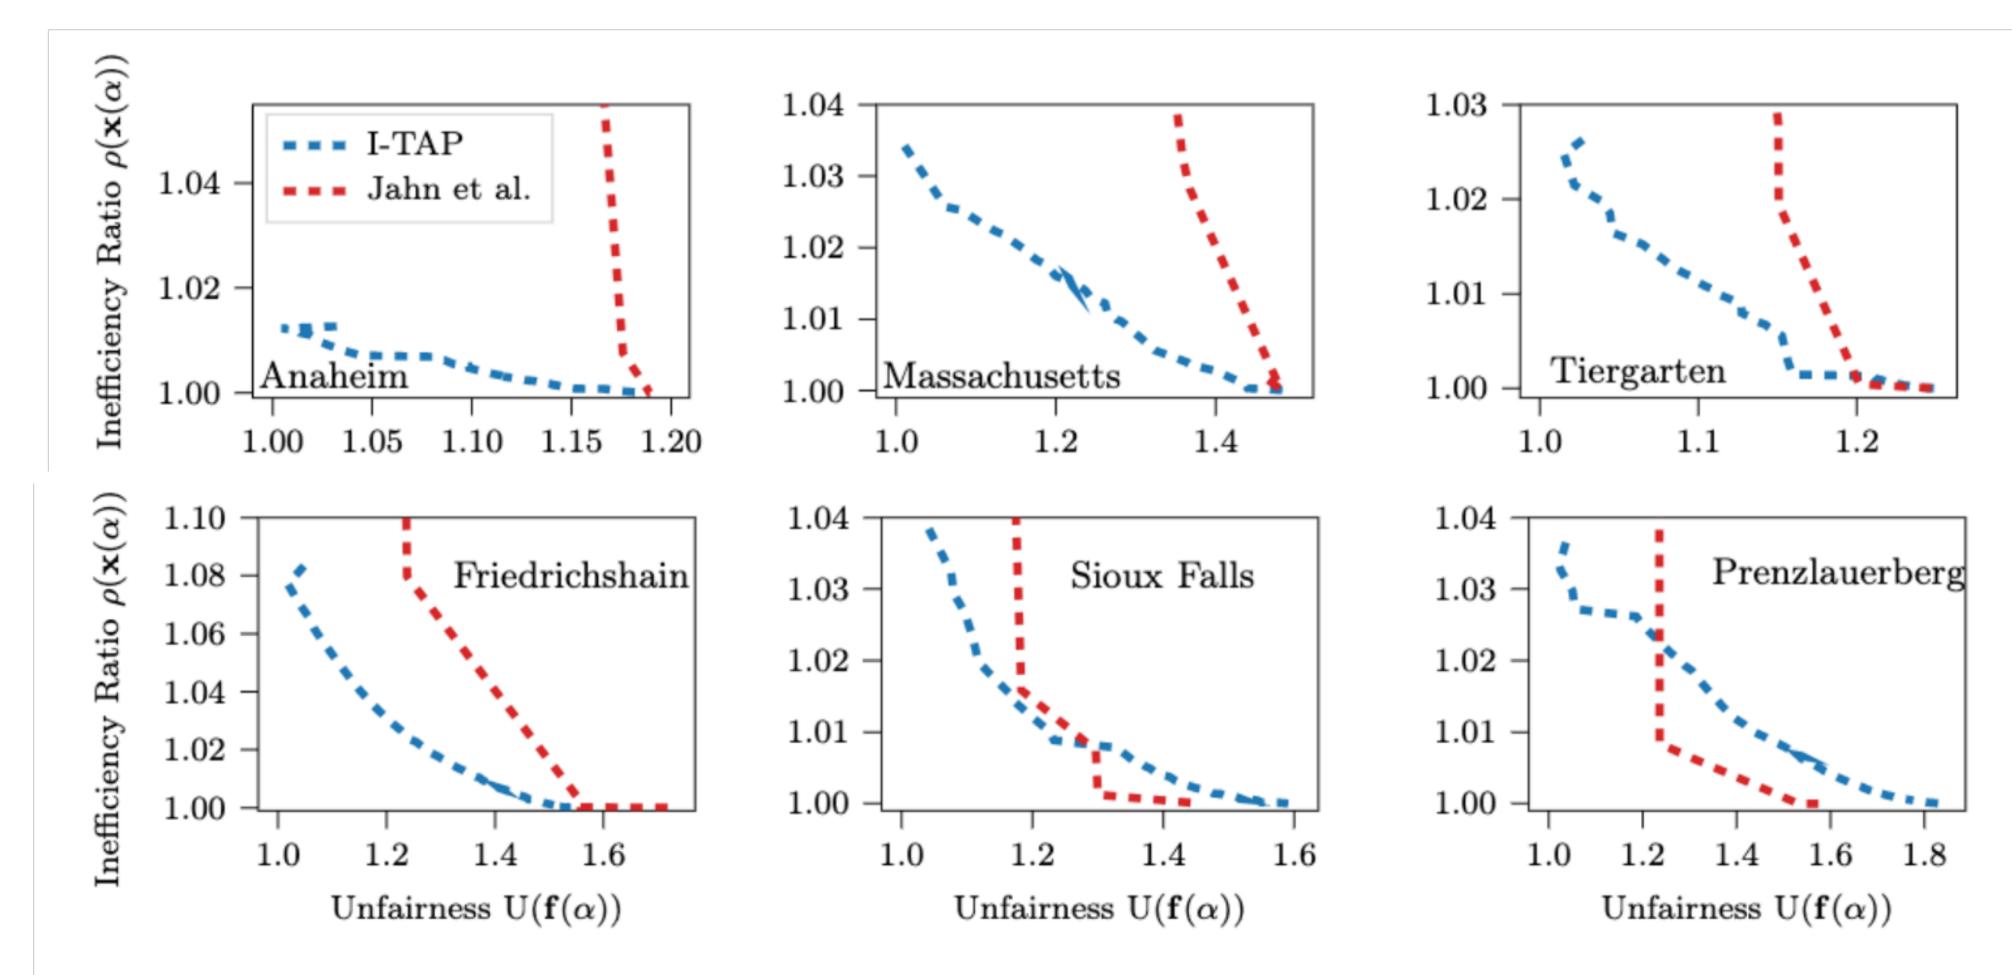
The superior runtime of I-TAP stems from the fact that it can be transformed into a UE-TAP, for which we have computationally efficient algorithms, e.g., Frank-Wolfe

$$f^{I_{\alpha}}(\mathbf{X}) = \alpha f^{SO}(\mathbf{X}) + (1 - \alpha) f^{UE}(\mathbf{X})$$
Fundamental
Theorem of Calculus
$$f^{I_{\alpha}}(\mathbf{X}) = \sum_{e \in E} \int_{0}^{x_e} t_e(y) + \alpha y t'_e(y) dy$$

UE-TAP with modified

Experiments: Efficacy of I-TAP

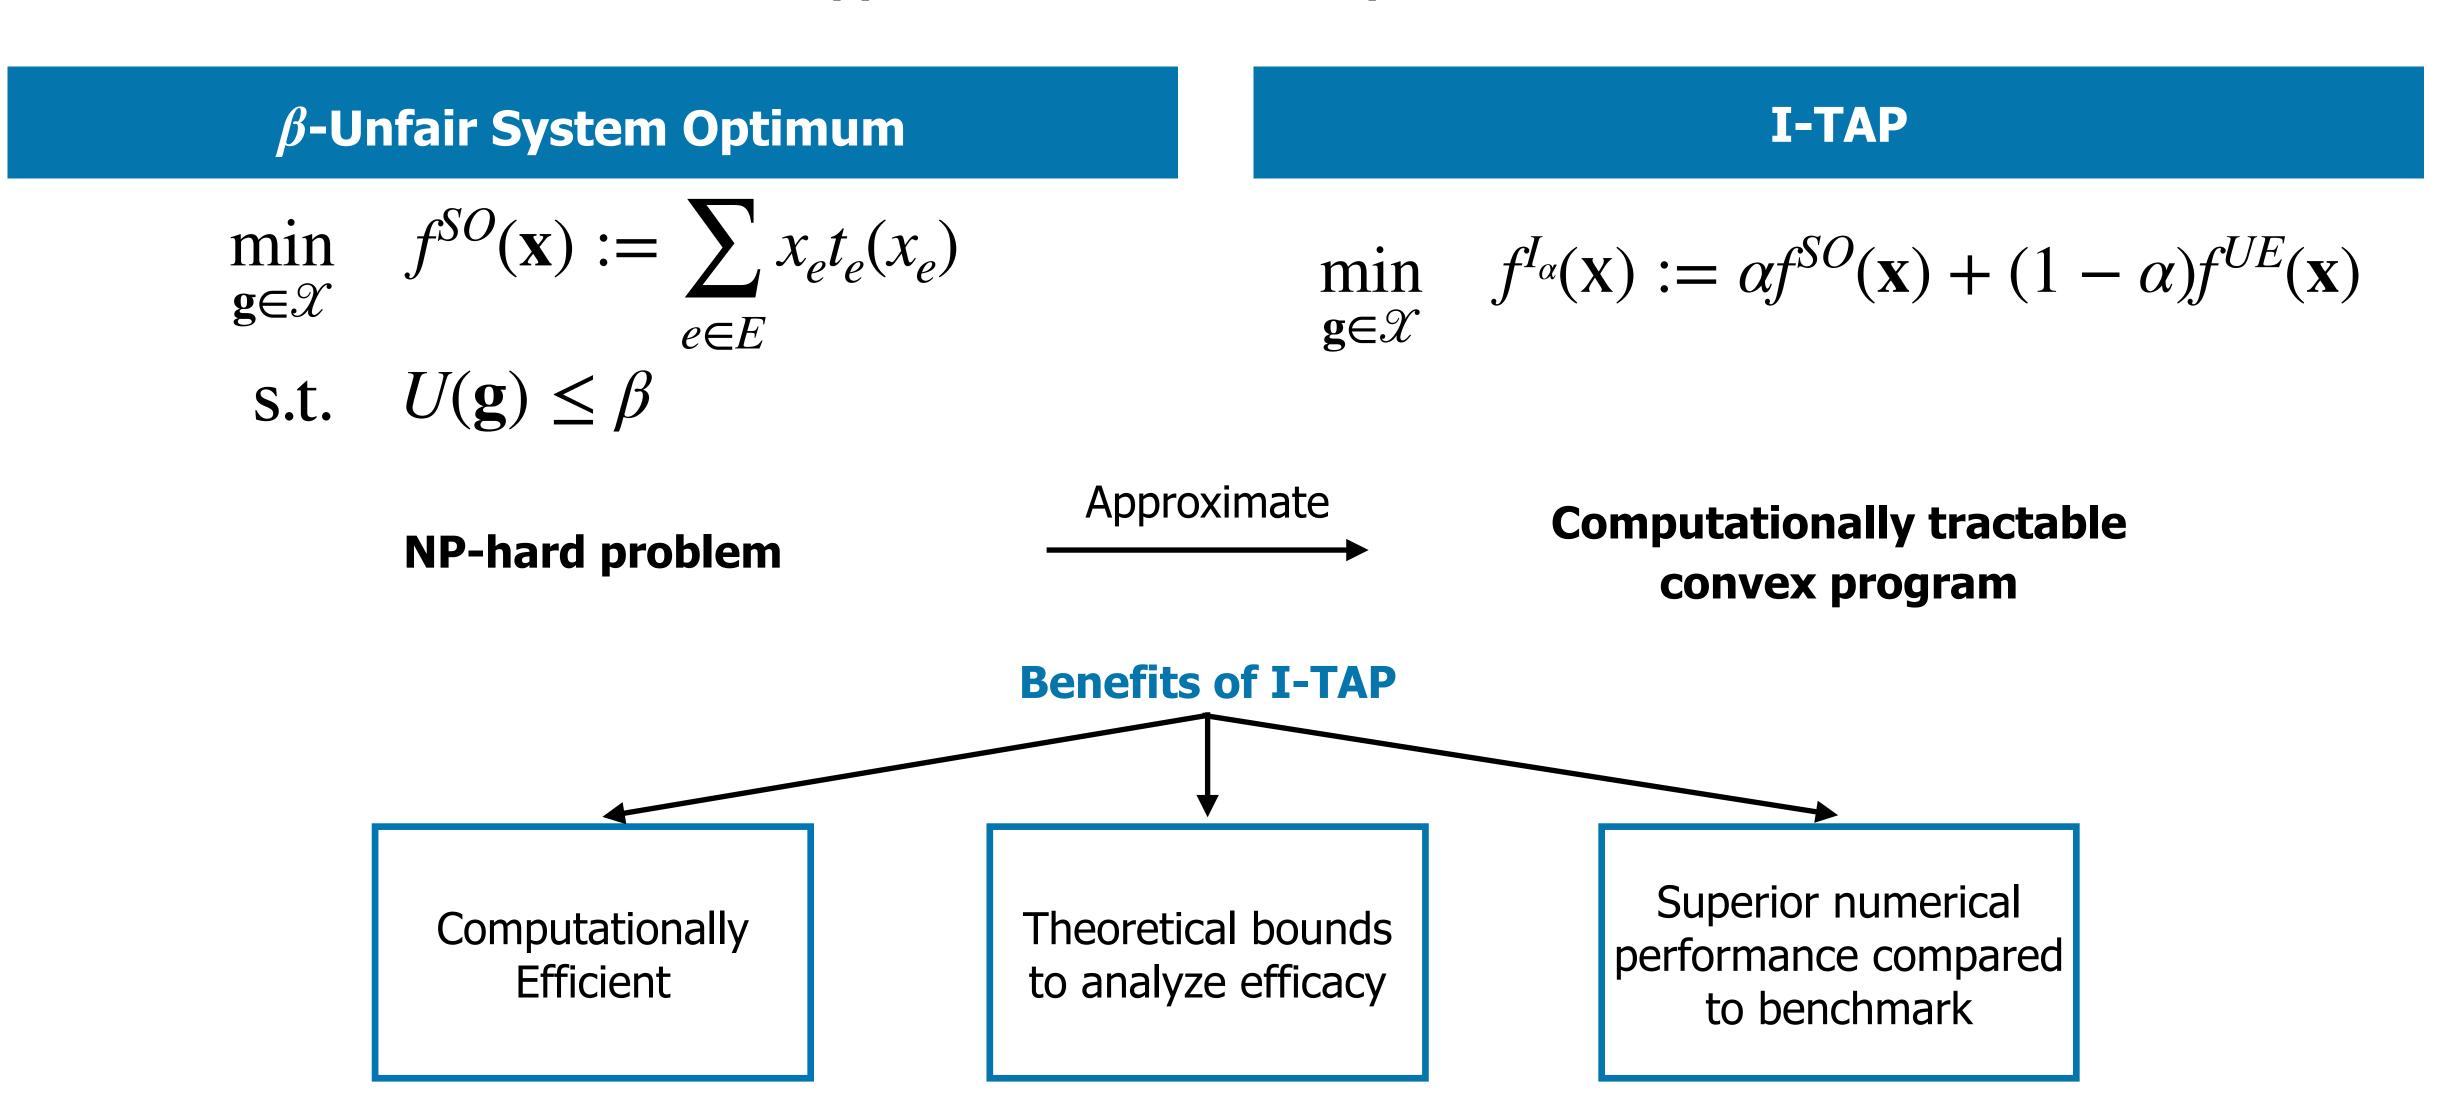
I-TAP outperforms a state-of-the-art benchmark in balancing efficiency and fairness in traffic routing



For all desirable levels of unfairness, i.e., low unfairness, I-TAP achieves lower total travel times than the state-of-the-art benchmark

I-TAP has several advantages

I-TAP serves as a tractable and effective approximation to the β -SO problem we seek to solve



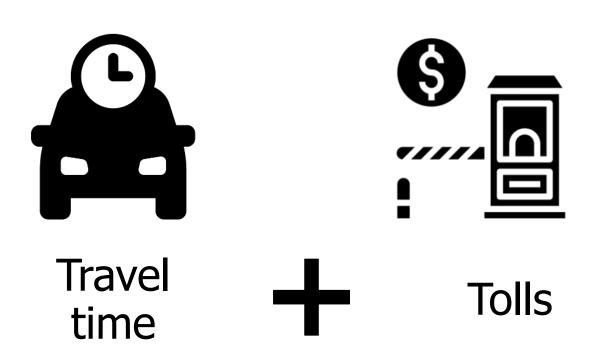
How do we set **prices** to induce fairness constrained flows computed via I-TAP in practice?

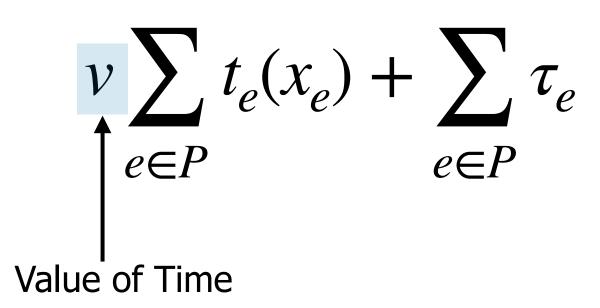
- Fairness in Congestion Pricing
 - Introduction and Motivating Example
 - Model and Problem Setup
 - Algorithmic Approach for Fair Traffic Routing: I-TAP
 - **Fair Congestion Pricing via I-TAP**
 - Conclusion
- Security Games for Enforcement in Mobility Applications
- Future Directions

Congestion Pricing via I-TAP

I-TAP yields a natural congestion pricing scheme to induce fairness-constrained flows in practice

User Behavior Model

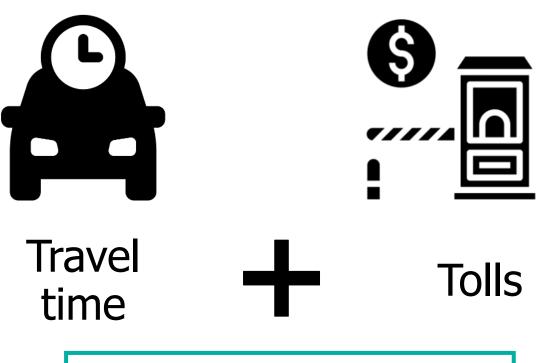




Congestion Pricing via I-TAP

I-TAP yields a natural congestion pricing scheme to induce fairness-constrained flows in practice

User Behavior Model



$$\sum_{e \in P} (t_e(x_e) + \tau_e)$$

Assume homogeneous users

Interpolated Marginal Cost Pricing with I-TAP

$$f^{I_{\alpha}}(\mathbf{x}) := \alpha f^{SO}(\mathbf{x}) + (1 - \alpha) f^{UE}(\mathbf{x})$$

$$f^{I_{\alpha}}(\mathbf{x}) = \sum_{e \in E} \int_{0}^{x_{e}} t_{e}(y) + \alpha y t'_{e}(y) dy$$

If there is positive flow on any path P, the following quantity is minimal for that O-D pair:

 $e \in P$

Interpolated Marginal Cost Tolls

I-TAP can be used to derive a linear programming based congestion pricing scheme to induce fairness constrained flows with heterogeneous users

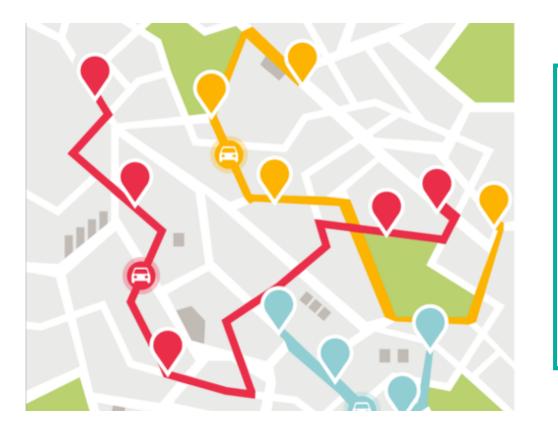
-> Fairness in Congestion Pricing

- Introduction and Motivating Example
- Model and Problem Setup
- Algorithmic Approach for Fair Traffic Routing: I-TAP
- Fair Congestion Pricing via I-TAP
- Conclusion
- Security Games for Enforcement in Mobility Applications
- Future Directions

Conclusion

We balance sustainable transportation's efficiency and equity goals

Computationally Efficient Routing Algorithm



Develop a novel convex program (I-TAP) that trades off efficiency and fairness

Benefits of I-TAP

Computationally Efficient

Theoretical Bounds to analyze efficacy

Superior numerical performance compared to benchmark

Fair Congestion Pricing Scheme



I-TAP yields natural congestion pricing schemes, namely interpolated marginal cost pricing

Extensions (Other Notions of Unfairness)

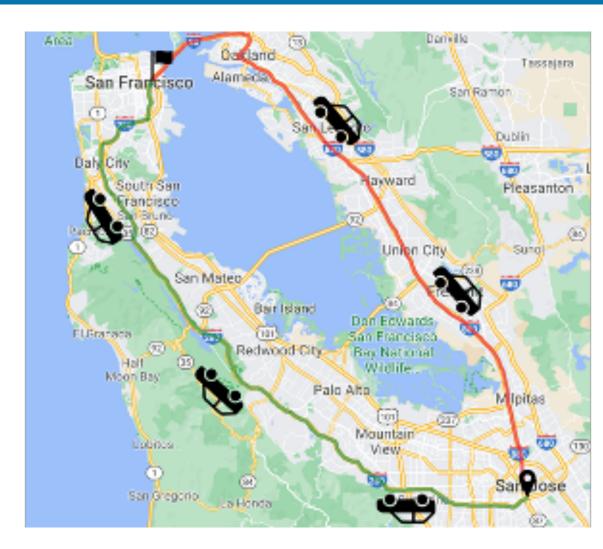
Numerical results show similar performance for other intra-O-D pair unfairness measures

Going **beyond** intra-O-D pair unfairness measures

Broader agenda on equitable congestion pricing

Our work paves the way for the design of sustainable and publicly acceptable congestion pricing schemes

Reduce low-income users' travel-time burden



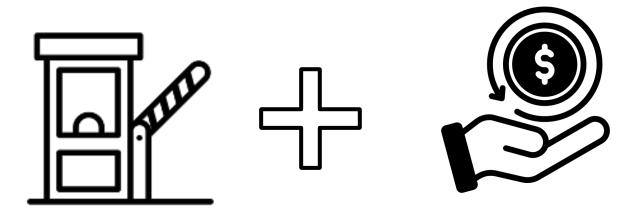
Convex program to overcome inefficiency of methods for fairness-constrained traffic assignment



JSTZP JAAMAS'23, AAMAS'22

Finalist for the **2024 INFORMS TSL best student paper award**

Congestion Pricing with Revenue Refunding





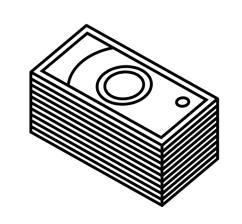
Provably increase system efficiency without worsening inequality



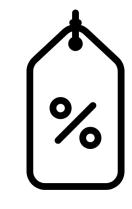
JSGZBP EAAMO'21, TCNS'23

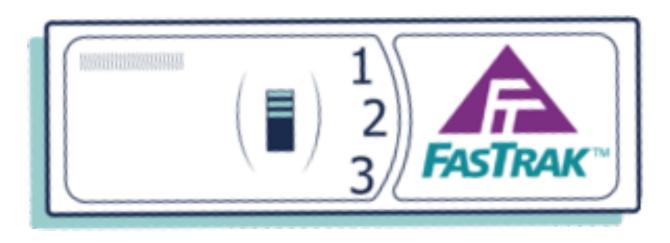
Covered by New York Times

Credit vs. Discount based Congestion Pricing



VS.





Game theoretic bi-level optimization framework to optimally design credit-based and discount-based schemes



JLBP CDC'23, CJP CDC'24

This talk: Fairness & security from approximation algorithm lens

We highlight two vignettes of this agenda focused on developing approximation algorithms for NP-hard problems arising under fairness and security considerations in mobility applications

Fairness (Congestion Pricing)

Jalota et al. JAAMAS'23

Finalist for the **2024 INFORMS TSL best student paper award**

Source of Hardness	Non-convexity of fairness constraints
Solution Method	Novel convex programming relaxation that also yields a natural pricing scheme

Security Games (Parking Enforcement)

Jalota et al. arXiv'24

Collaboration with **Stanford's DPS** to demonstrate efficacy of our algorithms on enforcement data in increasing parking revenues by \$300,000 annually

Source of Hardness	Resource constraints on the available set of security officers
Solution Method	Optimizing a specific concave approximation of bi-level program yielding a natural greedy algorithm

- Fairness in Congestion Pricing
- Security Games for Enforcement in Mobility Applications
 - Introduction
 - Model
 - Analysis of administrator strategies + equilibria
 - Numerical Experiments: Parking Enforcement
 - Conclusion
- Future Directions

- Fairness in Congestion Pricing
- Security Games for Enforcement in Mobility Applications
 - Introduction
 - Model
 - Analysis of administrator strategies + equilibria
 - Numerical Experiments: Parking Enforcement
 - Conclusion
- Future Directions

Fraud is ubiquitous and can be detrimental

Environmental Non-Compliance



SABC News

Several municipalities issued with non-compliance notices over clean water



Minister of Water and Sanitation Senzo Mchunu has described the recent cholera outbreak that claimed 26 lives in Hammanskraal,...

Jun 6, 2023

Water treatment facilities have an incentive to save on compliance costs, which can put resident lives at risk

Traffic/Parking Violations



922,000 parking tickets issued in Los Angeles during first half of 2024



By Jeff Nguyen July 9, 2024 / 10:49 PM PDT / KCAL News f X 🖪

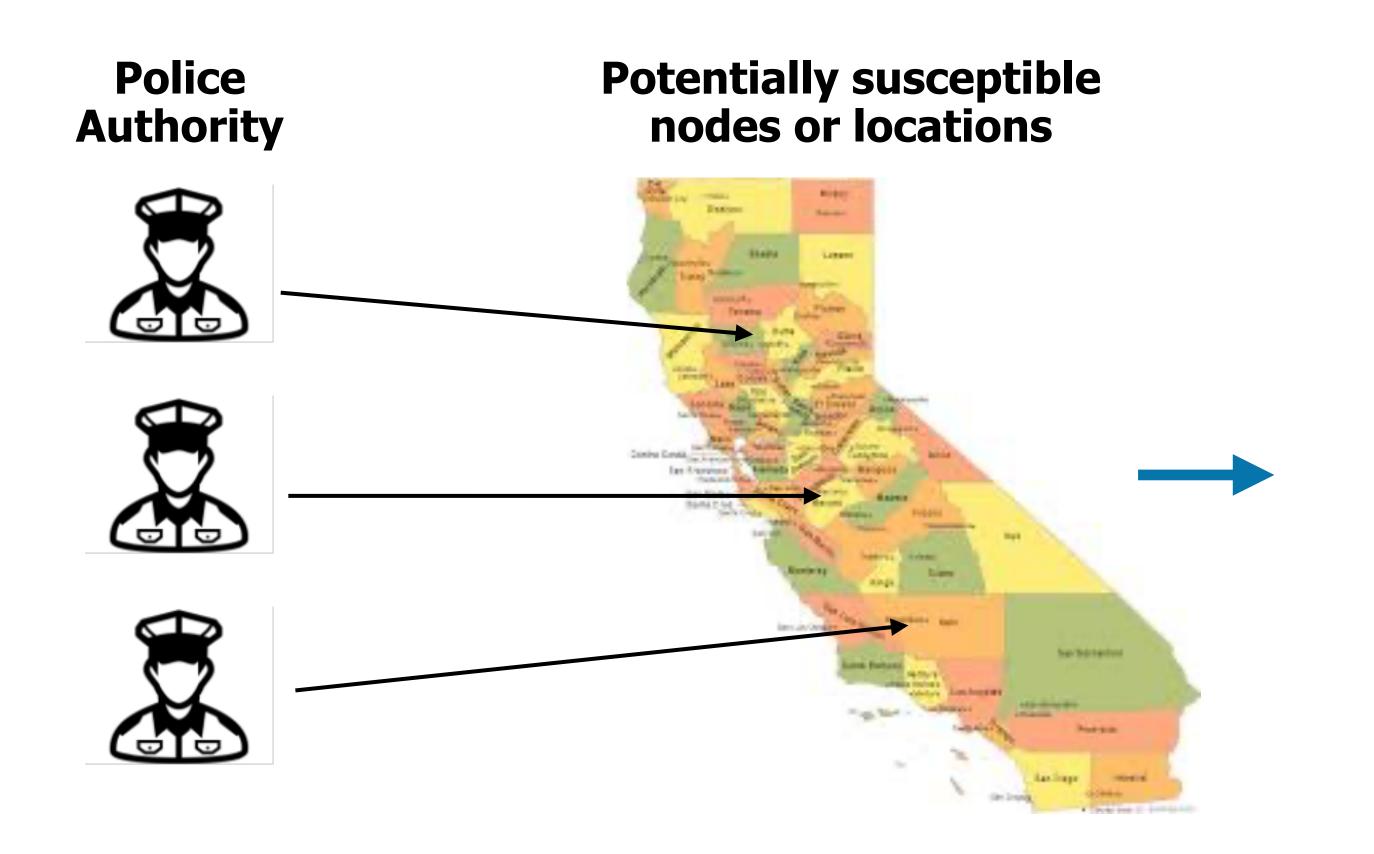
Commuters often engage in illegal parking practices or speeding, which can result in **traffic jams** and **compromise road safety**

Users engage in fraud to **strategically** obtain some benefit that would otherwise be unattainable within the bounds of lawful conduct

The prevalence of fraud raises a key security challenge

Effective enforcement requires policing, which poses a resource allocation problem

How do we best allocate a limited set of security resources to susceptible nodes or locations in a system?



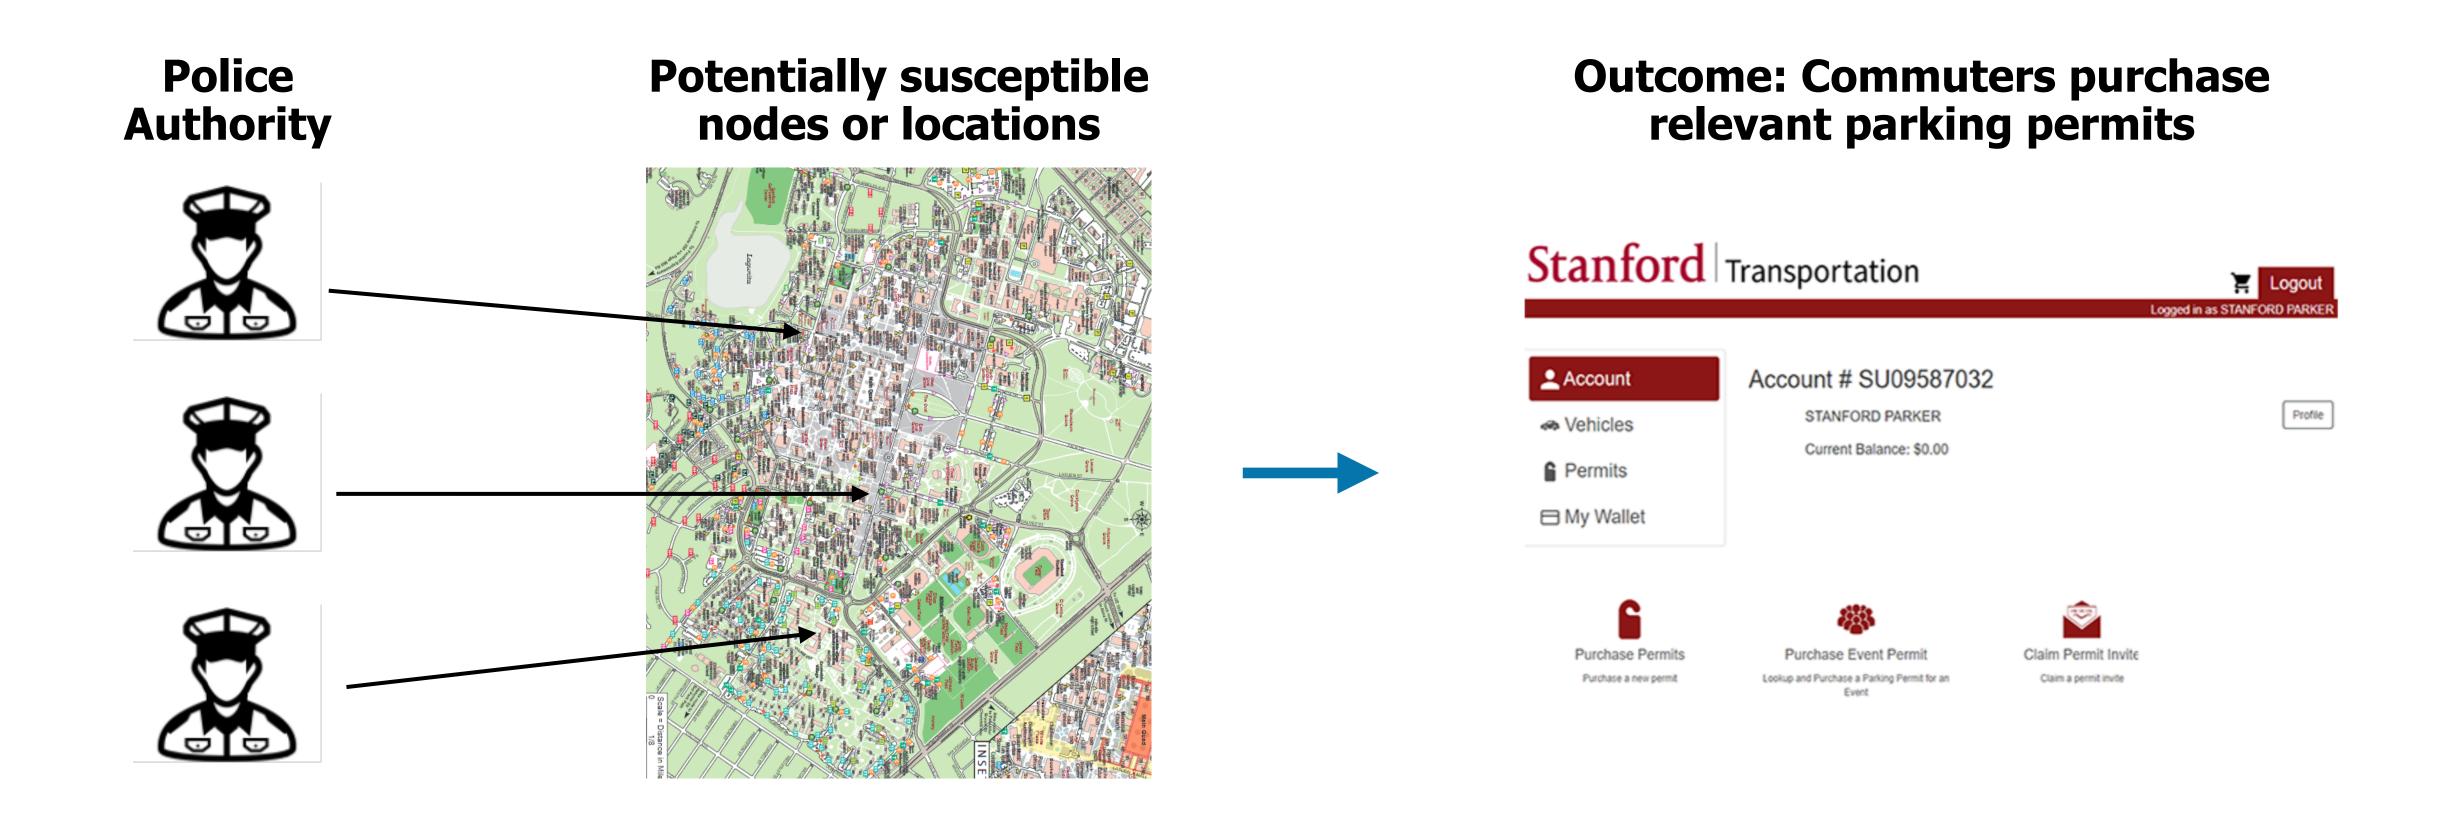
Outcome: More facilities complying with environmental regulations



The prevalence of fraud raises a key security challenge

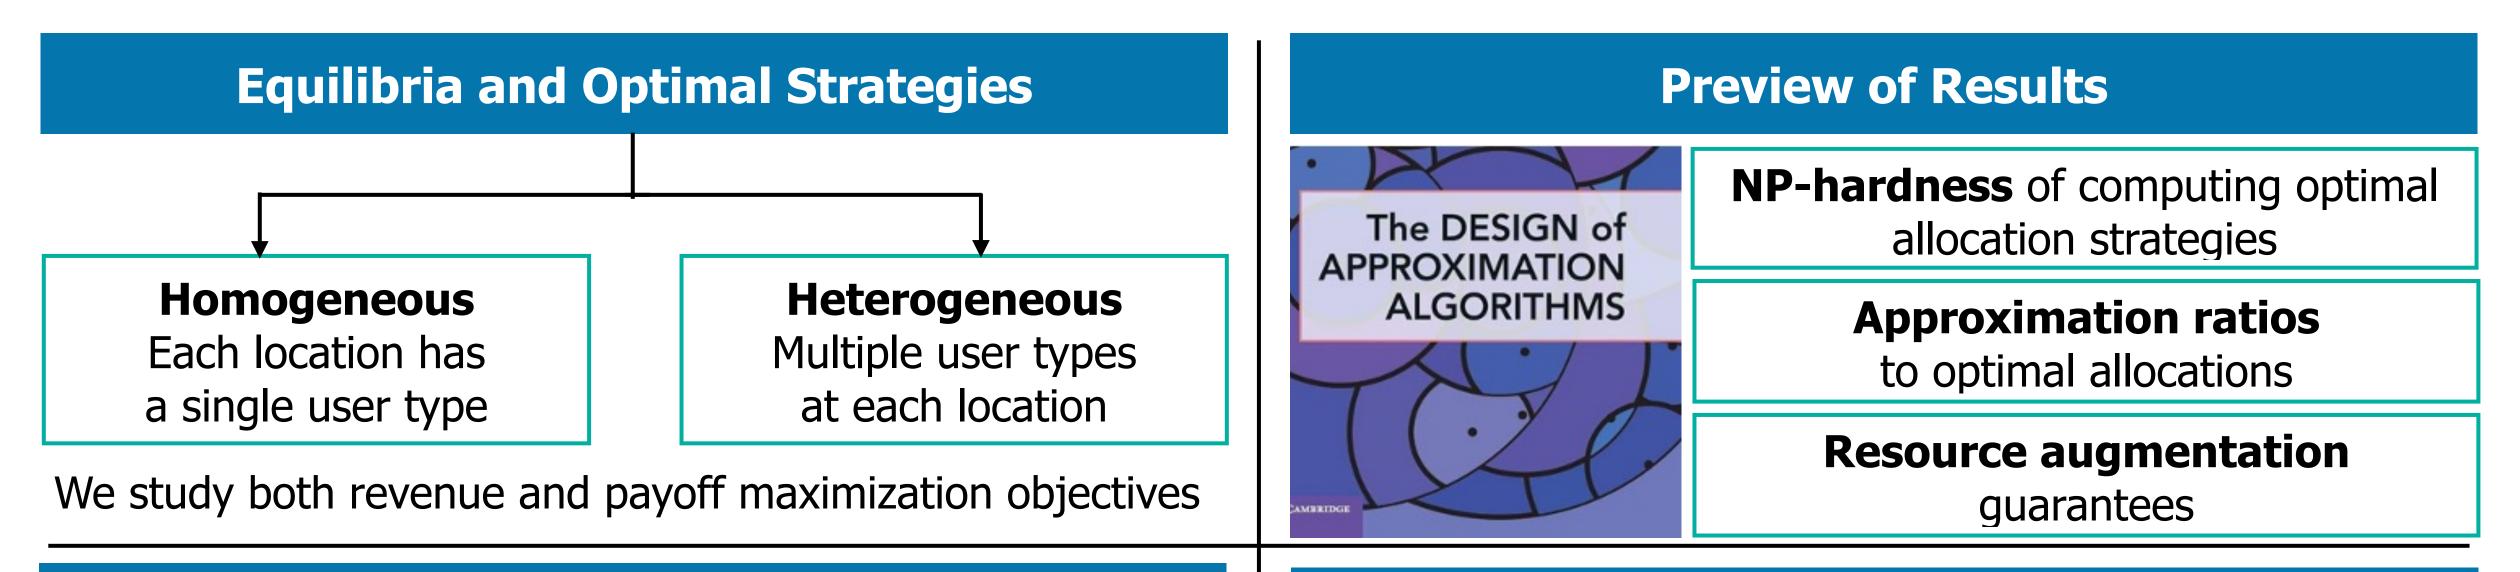
Effective enforcement requires policing, which poses a resource allocation problem

How do we best allocate a limited set of security resources to susceptible nodes or locations in a system?



Contributions

We study a security game between users and an administrator with a budget of security resources, where the administrator levies fines on users engaging in fraud



Numerical Experiments



Case study of Stanford's parking enforcement system

Model Extensions

Contract Game: Payoff and revenue objectives at odds

Additional Constraints over allocation of resources

Variability of Fines: Administrator can optimize fines

Related Literature

Our work provides algorithms with provable guarantees for multi-resource Bayesian SSGs

Prior Work on Security Games

Single-resource
Bayesian Stackelberg
Security Games (SSGs)

Multi-resource single type SSGs

Our Work

Our work provides the first algorithms with provable guarantees for multi-resource Bayesian SSGs

Adversary can have multiple types

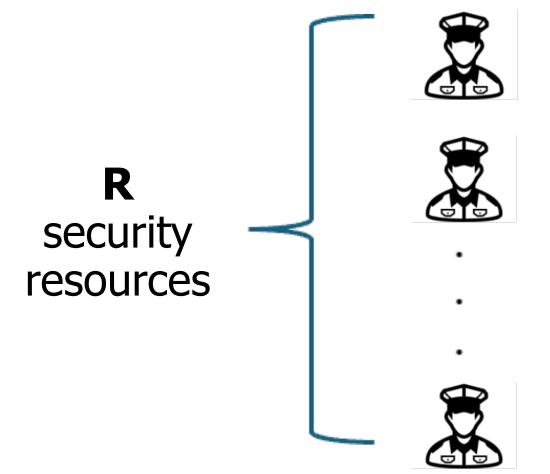
Administrator has a budget of security resources to prevent fraud

- Fairness in Congestion Pricing
- Security Games for Enforcement in Mobility Applications
 - Introduction
 - Model
 - Analysis of administrator strategies + equilibria
 - Numerical Experiments: Parking Enforcement
 - Conclusion
- Future Directions

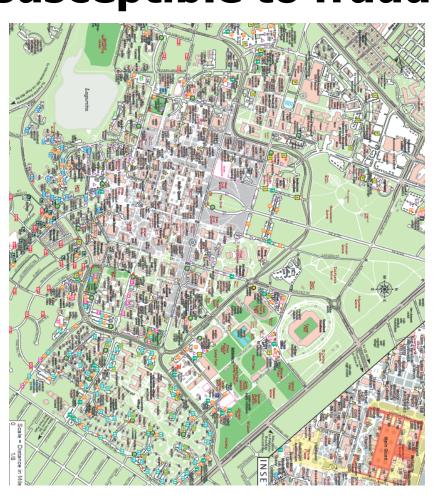
Security Game Model

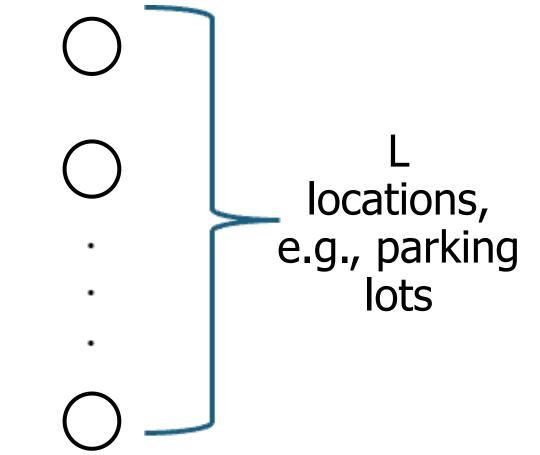
Administrator (Police Authority)





Locations susceptible to fraud





Administrator levies a fine k if users are found engaging in fraud

Location l



$$\Theta_l^i = (\Lambda_l^i, d_l^i, p_l^i)$$

 Λ_{I}^{i} : Number of users

 d_i^i : Benefit from engaging in fraud

 p_l^i : Administrator payoff from reducing fraud

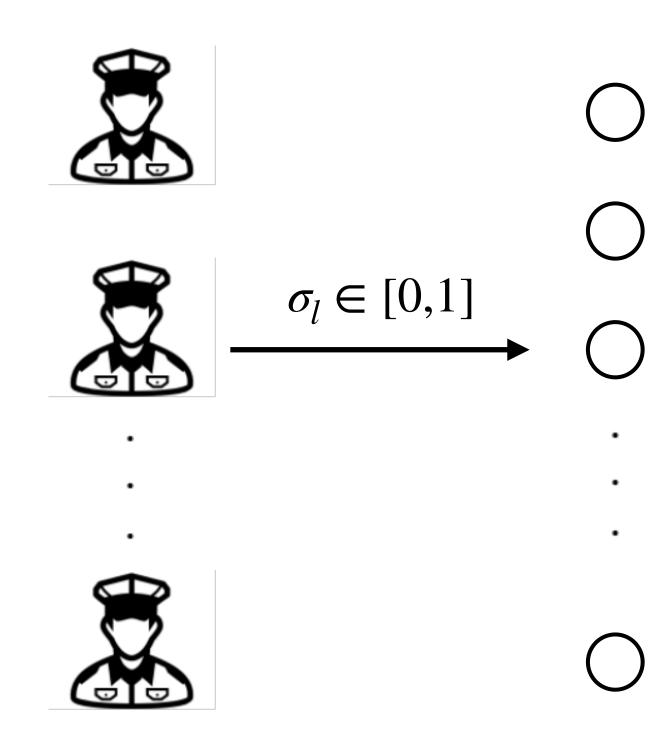
Administrator and User Strategies

Administrator Strategy

Administrator selects a mixedstrategy σ to allocate its budget of R security resources across the locations

 σ_l = Probability security resource is allocated to location l

Feasible strategy set:
$$\Omega_R = \left\{ \boldsymbol{\sigma} = (\sigma_l)_{l \in L} : \sigma_l \in [0,1], \sum_l \sigma_l \leq R \right\}$$



User Strategy

In response to administrator's strategy σ , users decide whether to engage in fraud

 $y_l^i(\boldsymbol{\sigma})$ = Probability user engages in fraud at location l with type igiven σ

 $y_l^l(\boldsymbol{\sigma}) \in [0,1]$ for all locations land types i

Resource Constraint:
$$\sum_{l} \sigma_{l} \leq R$$

Administrator objectives

We study equilibrium outcomes under both revenue and payoff maximization administrator objectives

Revenue Maximization

$$\max_{\substack{\sigma \in \Omega_R \\ y_l^i(\sigma) \in [0,1], \forall l \in L, i \in \mathcal{I}}} Q_R(\sigma) = \sum_{l \in L} \sum_{i \in \mathcal{I}} \sigma_l y_l^i(\sigma) k \Lambda_l^i,$$

$$\mathrm{s.t.} \ \ y_l^i(\sigma) \in \arg\max_{y \in [0,1]} U_l^i(\sigma,y) = y \big[(1-\sigma_l) d_l^i - \sigma_l k \big]$$
 for all $l \in L, i \in \mathcal{I}$
$$\mathrm{Gains} \quad \mathrm{Losses}$$
 from fraud from fraud

Police often maximize revenues from fine collections, e.g., highway speed traps may not target accident prone areas

Payoff Maximization

$$\begin{aligned} \max_{\substack{\sigma \in \Omega_R \\ y_l^i(\sigma) \in [0,1], \forall l \in L, i \in \mathcal{I}}} & P_R(\sigma) = \sum_{l \in L} \sum_{i \in \mathcal{I}} p_l^i (1 - (1 - \sigma_l) y_l^i(\sigma)) \\ \text{s.t.} & y_l^i(\sigma) \in \argmax_{y \in [0,1]} U_l^i(\sigma, y) \\ & \text{for all } l \in L, i \in \mathcal{I} \end{aligned}$$

Administrators, such as owners of parking lots, may seek to minimize fraud, e.g., to maximize parking permit purchases

Goal: Characterize optimal (or near-optimal) solutions to these bi-level programs, i.e., study **equilibrium** formation under both these administrator objectives

- Fairness in Congestion Pricing
- Security Games for Enforcement in Mobility Applications
 - Introduction
 - Model
 - → Analysis of administrator strategies + equilibria
 - Numerical Experiments: Parking Enforcement
 - Conclusion
- Future Directions

Preview of Results

Homogeneous revenue maximization is polynomial time solvable while other settings are NP-hard

	Revenue Maximization	Payoff Maximization
Homogeneous (Each location has a single user type)		NP-hard
	Polynomial time greedy algorithm to compute revenue-maximizing	1/2 Approximation: Greedy algorithm achieves at least half the optimal payoff
	administrator strategy	Resource augmentation: Algorithm with R+1 resources achieves at least the payoff as optimal strategy with R resources
Heterogeneous (Each location has multiple user types)	NP-hard	NP-hard
	1/2 Approximation	1/2 Approximation
	Resource Augmentation	Resource Augmentation

Heterogeneous revenue maximization is NP-hard

No polynomial time algorithm can solve the heterogeneous revenue maximization bi-level program unless P = NP

Theorem: The problem of computing the revenue maximizing strategy with heterogeneous user types is NP-hard

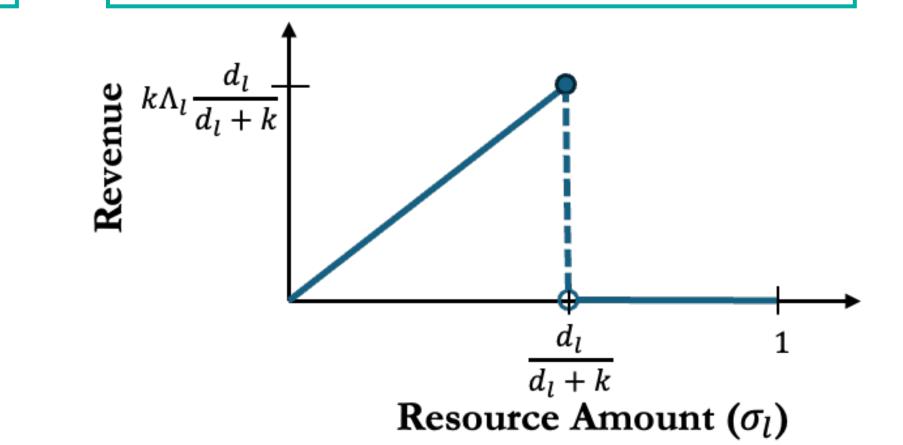


Geometric Intuition

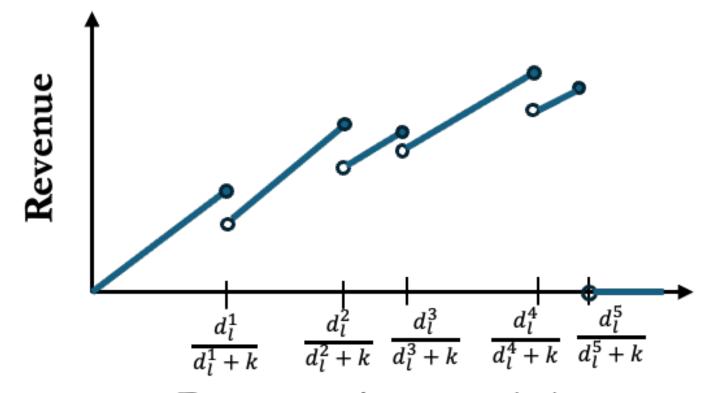
User best-response to revenuemaximizing administrator

$$y_l^i(\boldsymbol{\sigma}) = \begin{cases} 0, & \text{if } \sigma_l > \frac{d_l^i}{d_l^i + k}, \\ 1, & \text{otherwise.} \end{cases}$$

Revenue function at location *l* (**Homogeneous** User Types)



Revenue function at location *l* (**Heterogeneous** User Types)



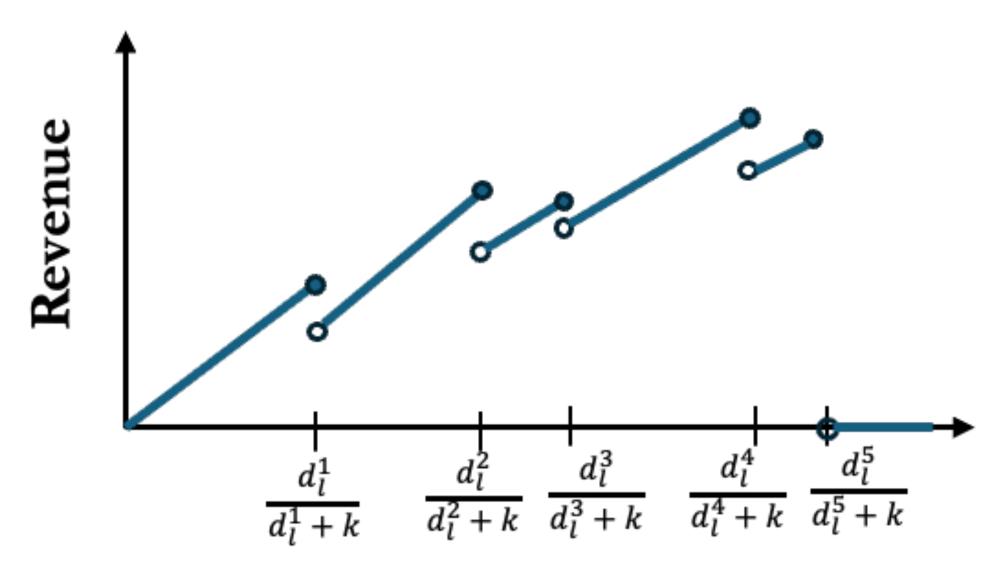
Resource Amount (σ_l)

- Revenue function with **homogeneous** user types is **continuous** in the range $\sigma_l \in \left[0, \frac{d_l}{d_l + k}\right]$
- Revenue function with **heterogeneous** user types is **discontinuous** and **non-monotone** at $\frac{d_l^i}{d_l^i + k}$ for all i

Algorithm for heterogeneous revenue maximization

We maximize a monotone concave upper approximation (MCUA) of the revenue function

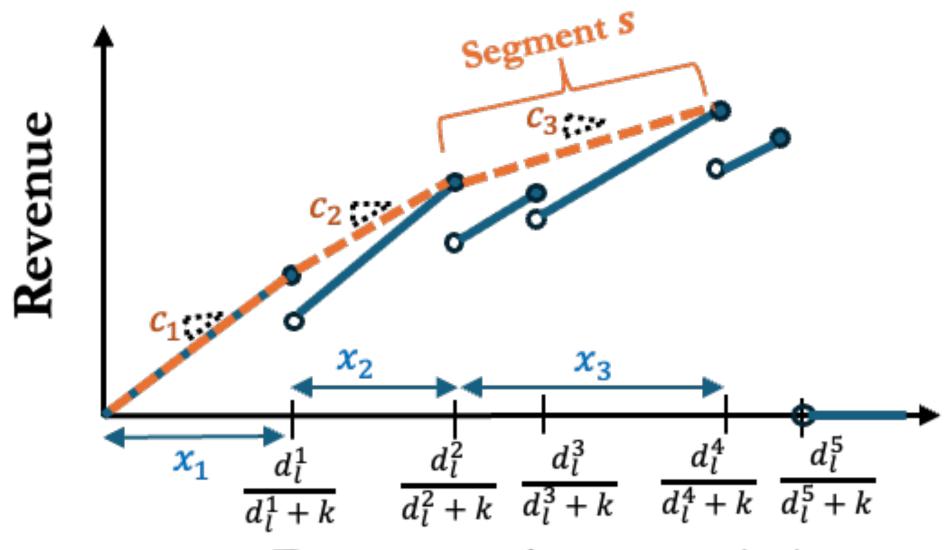
Revenue function



Resource Amount (σ_l)

Revenue function is not continuous and non-monotone

MCUA of revenue function



Resource Amount (σ_l)

MCUA can be **optimized tractably** using a greedy procedure that orders segments in descending order of segment slopes

Guarantees for heterogeneous revenue maximization

Maximizing the MCUA of the revenue function motivates a natural modified greedy algorithm

Output: Resource Allocation Strategy σ_A^* Step 1: Greedy Allocation $\tilde{\sigma}$ Based on Slopes of MCUA of Revenue Function: Define affordability threshold $t_l \leftarrow \min \left\{ R, \max_i \frac{d_l^i}{d_i^i + k} \right\}$ for all locations l; Generate MCUA of the revenue function in range $[0, t_l]$ for each location l; $\tilde{\mathcal{S}} \leftarrow \text{Ordered list of segments } s \text{ across all locations of this MCUA in descending order of slopes } c_s$; Initialize allocation strategy $\tilde{\boldsymbol{\sigma}} \leftarrow \mathbf{0}$; for segment $s \in \tilde{\mathcal{S}}$ do Maximize MCUA of revenue if $x_s \leq R$ then function using greedy procedure $\tilde{\sigma}_{l_s} \leftarrow \tilde{\sigma}_{l_s} + x_s$; Allocate x_s to location l_s ; $R \leftarrow R - x_s$; Update amount of remaining resources; else **break**; Only allocate resources if $x_s \leq R$ end Step 2: Find Solution σ' that maximizes revenue from spending on single location: $\sigma^{\iota} \leftarrow \operatorname{arg} \max_{\sigma \in \Omega_R : \sigma_{l'} = 0, \forall l' \neq l} Q_R(\sigma) \text{ for all locations } l ;$ Optimal spending on single location $\sigma' \leftarrow \operatorname{arg\,max}_{l \in L} Q_R(\sigma^l) ;$ Select best of the two solutions Step 3: Return Solution with a Higher Revenue: $\sigma_A^* \leftarrow \arg\max\{Q_R(\tilde{\sigma}), Q_R(\sigma')\}$;

Theorem: Algorithm 1 achieves **at least half** the revenue as the revenue maximizing solution

Theorem: Algorithm 1 with **R+1** resources achieves **at least** the revenue as the revenue maximizing solution with **R** resources

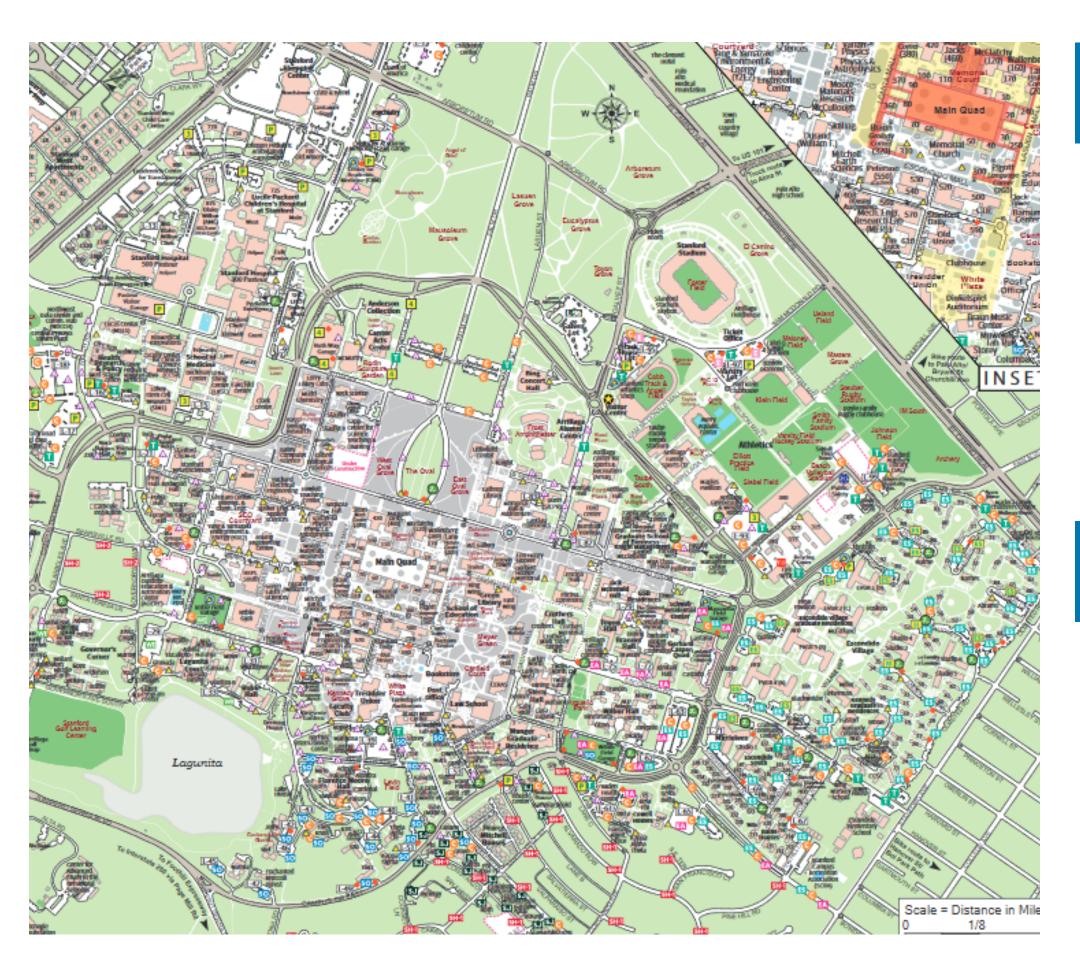
Algorithm 1: Greedy Algorithm for Administrator's Heterogeneous Revenue Maximization Objective

Input: Total Resource capacity R, User Types $\Theta_l^i = (\Lambda_l^i, d_l^i, v_l^i)$ for all locations l and types i

- Fairness in Congestion Pricing
- Security Games for Enforcement in Mobility Applications
 - Introduction
 - Model
 - Analysis of administrator strategies + equilibria
 - Numerical Experiments: Parking Enforcement
 - Conclusion
- Future Directions

Stanford parking enforcement setup and data-set

We obtained 7 months of parking enforcement and citation data from Stanford's Department of Public Safety



Setup

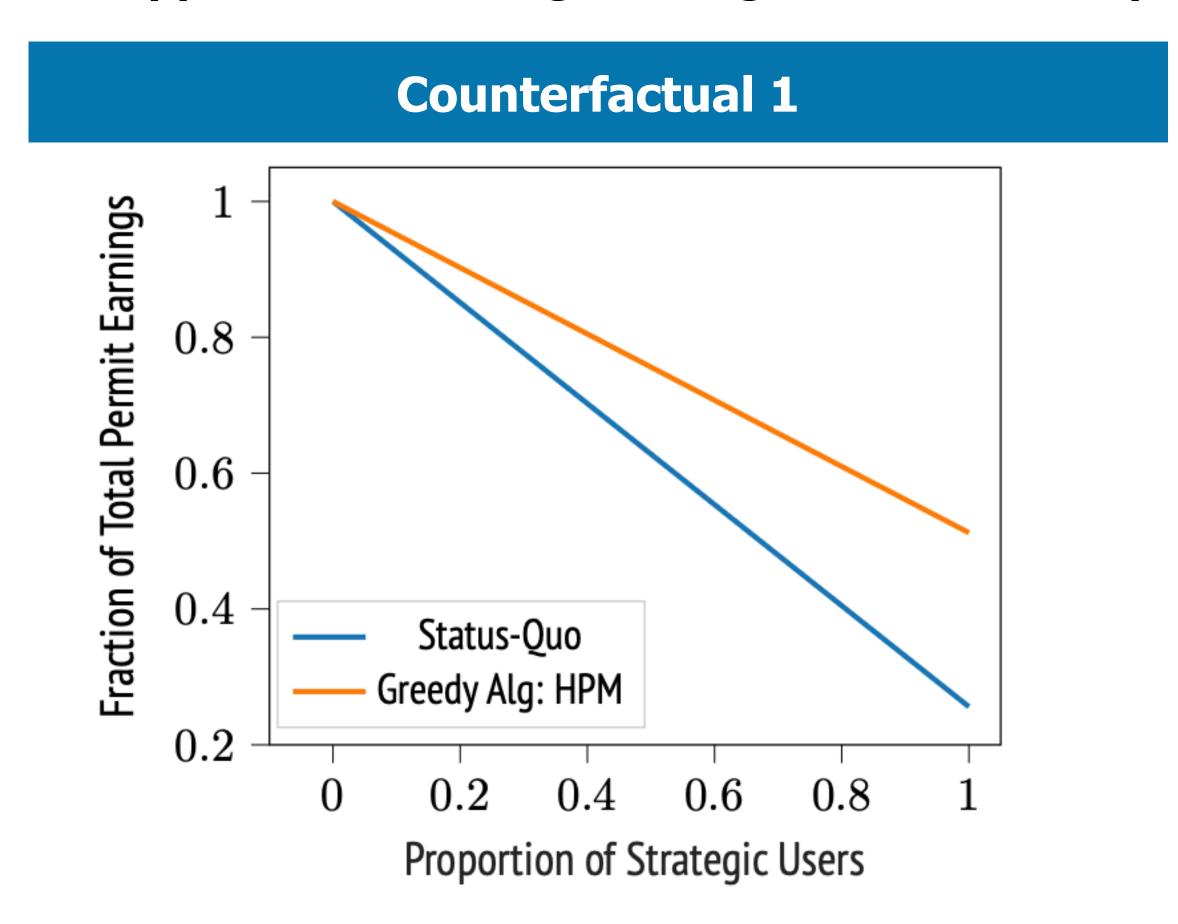
- Commuters are required to purchase permits and are issued citations if found violating regulations
- Citation fees go to Santa Clara County's Sheriff Department but earnings from permit purchases go to Stanford

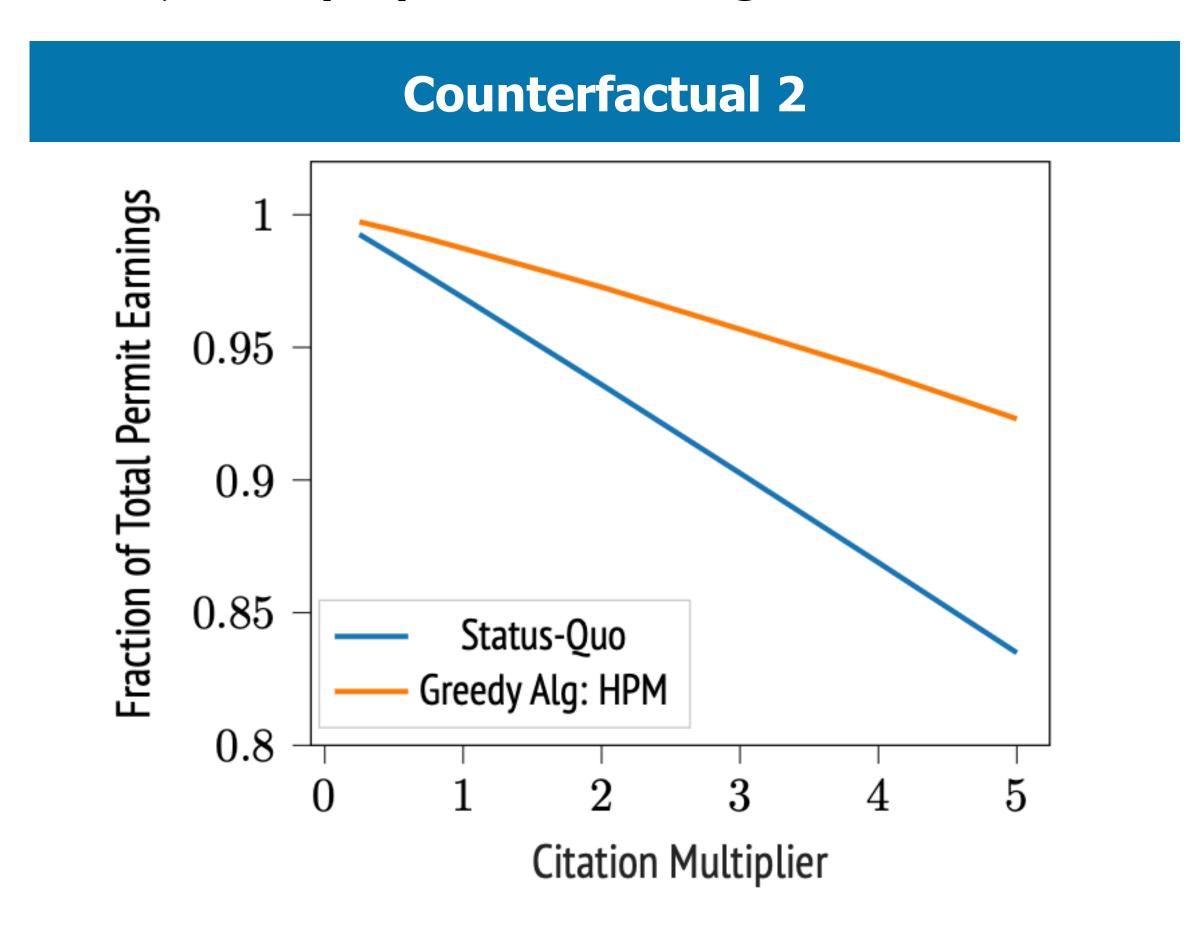
Data-set

- 7 months of data from Sep 2022 Mar 2023
- Enforcement data includes the time and period of enforcement
- Citation data includes the total number of citations issued in each parking lot in a given month
- We formulate Stanford's parking enforcement problem as a heterogeneous payoff maximization problem, where payoffs represent parking permit earnings

Results

Our approach achieves significant gains to the status quo, particularly as the proportion of strategic users increases





- Our approach has significant gains relative to the status quo as the proportion of strategic users increases
- Our approach achieves over \$300,000 additional permit earnings per year relative to status quo

- Fairness in Congestion Pricing
- Security Games for Enforcement in Mobility Applications
 - Introduction
 - Model
 - Analysis of administrator strategies + equilibria
 - Numerical Experiments: Parking Enforcement
 - Conclusion
- Future Directions

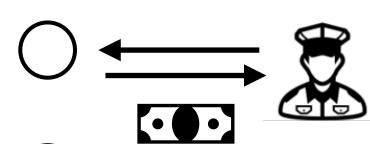
Conclusion

We developed simple yet near-optimal algorithms to compute equilibria in our security game

Equilibria and Optimal Strategies

Studied a multiresource Bayesian
Stackelberg security game
where fines are levied on
fraudulent users







Numerical Experiments



Case study of Stanford's parking enforcement system

Results on computation of administrator's optimal strategies

	Revenue Maximization	Payoff Maximization
Homogeneous	Polynomial time algorithm	NP-hard ½ Approximation, PTAS Resource augmentation
Heterogeneous	NP-hard 1/2 Approximation Resource augmentation	NP-hard 1/2 Approximation Resource augmentation

Model Extensions

Contract Game

Additional Constraints

Variability of Fines

Summary: Fairness & security from approximation algorithm lens

We highlight two vignettes of this agenda focused on developing approximation algorithms for NP-hard problems arising under fairness and security considerations in mobility applications

Fairness (Congestion Pricing)

Jalota et al. JAAMAS'23

Finalist for the **2024 INFORMS TSL best student paper award**

Source of Hardness	Non-convexity of fairness constraints
Solution Method	Novel convex programming relaxation that also yields a natural pricing scheme

Security Games (Parking Enforcement)

Jalota et al. arXiv'24

Collaboration with **Stanford's DPS** to demonstrate efficacy of our algorithms on enforcement data in increasing parking revenues by \$300,000 annually

Source of Hardness	Resource constraints on the available set of security officers
Solution Method	Optimizing a specific concave approximation of bi-level program yielding a natural greedy algorithm

My Work: Sustainable resource allocation in socio-technical systems

My work leverages operations, CS, and economics to advance the science and practice of sustainable resource allocation

