Algorithmic Persuasion with No Externalities

Haifeng Xu

Postdoctoral Fellow, Harvard (→ UVA CS from Fall 2019)

Joint work with Shaddin Dughmi

Outline

- Introduction and Model
- Persuasion through the Algorithmic Lens
- Conclusion

Model

One sender persuades multiple receivers with no externalities

Example: Recommendation Letters

- Academic advisor vs. two fellowship programs
- 1/3 of the advisor's students are excellent; 2/3 are average
- A fresh graduate is randomly drawn from this population
- Each fellowship:
  - Utility $1+\epsilon$ for awarding excellent student; -1 for average student
  - Utility 0 for no award
  - A-priori, only knows the advisor's student population
  - Student can accept both fellowships

$$\underbrace{(1+\epsilon)\times 1/3 - 1\times 2/3}_{\textit{Awarding}} \quad < \quad \underbrace{0}_{\textit{Not awarding}}$$

Example: Recommendation Letters

- Academic advisor vs. two fellowship programs
- 1/3 of the advisor's students are excellent; 2/3 are average
- A fresh graduate is randomly drawn from this population
- Each fellowship:
  - Utility $1+\epsilon$ for awarding excellent student; -1 for average student
  - Utility 0 for no award
  - A-priori, only knows the advisor's student population
  - Student can accept both fellowships
- Advisor
  - Utility 1 if student gets at least one fellowship, 0 otherwise
  - Knows whether the student is excellent or not

Example: Recommendation Letters

What is the advisor's optimal "recommendation strategy"?

- Attempt 1: always say "excellent" (equivalently, no information)
  - Fellowships ignore the recommendation
  - No fellowship awarded, advisor utility 0

Example: Recommendation Letters

What is the advisor’s optimal “recommendation strategy”?

- Attempt 2: honest recommendation (equivalently, full information)
  - 1/3 of students get both fellowships
  - Advisor expected utility 1/3

Example: Recommendation Letters

What is the advisor’s optimal “recommendation strategy”?

- Attempt 3: noisy information → advisor expected utility 2/3
  - Optimal public scheme

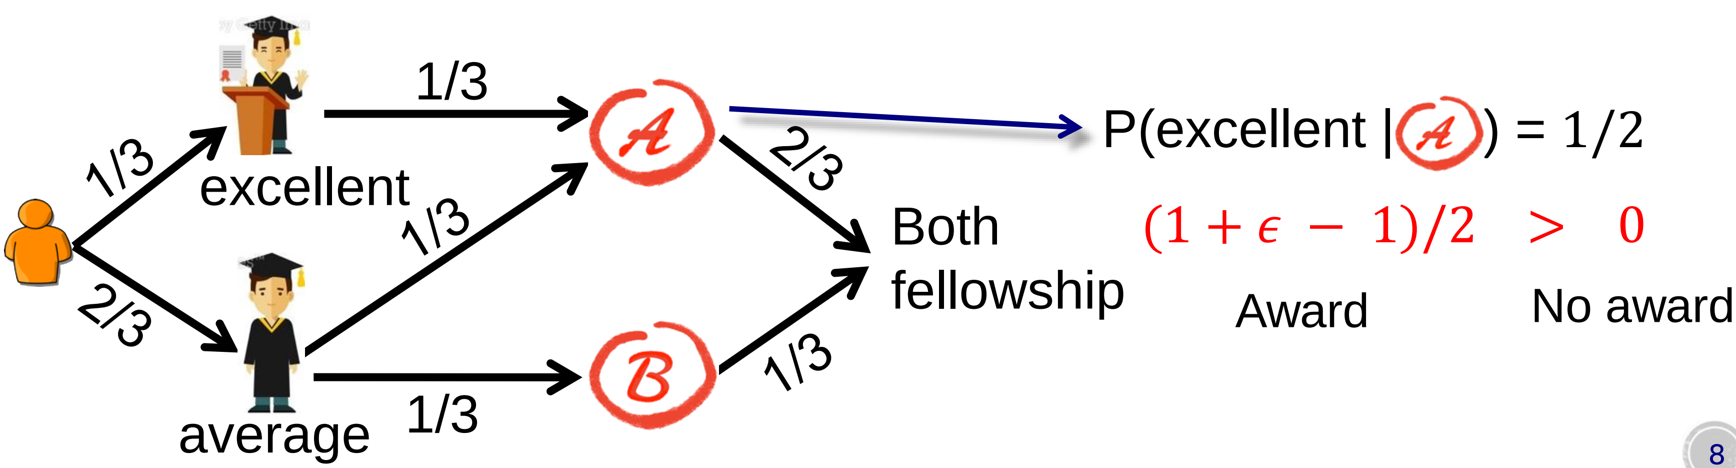

Example: Recommendation Letters

What is the advisor's optimal "recommendation strategy"?

- ➢ Attempt 4: optimal private scheme → advisor utility 1
  - ❖ When student is excellent, "strong" to both fellowships
  - ❖ Otherwise: "strong" to one fellowship, chosen randomly
- ➢ Conditioned on "strong", excellent with prob 1/2
- ➢ Always at least one fellowship recommended "strong"

Example: Recommendation Letters

Generalize this example to n fellowships:

advisor utility of optimal private scheme

$\geq \frac{n+1}{2}$ advisor utility of optimal pubic scheme

Conceptual Message

Being able to persuade privately may have a huge advantage

Model : Persuasion with No Externalities

- One sender, n receivers
- Receiver i takes a binary action $a_i \in \{0,1\}$, resulting in utility $r_i(a_i|\theta)$
  - No externality: $r_i(a_i|\theta)$ does not depends on a_j for j $\neq i$

A (random) state of nature from discrete set Θ

Model : Persuasion with No Externalities

- One sender, n receivers
- Receiver i takes a binary action $a_i \in \{0,1\}$, resulting in utility $r_i(a_i|\theta)$
  - No externality: $r_i(a_i|\theta)$ does not depends on a_j for j $\neq i$
- Sender utility is a set function $f(S)$, where $S = \{\text{receivers taking action 1}\}$
  - Assume $f(S)$ is monotone non-decreasing
- All receivers and the sender share a *common* prior belief of θ
- Additionally, sender can observe realized θ
- Before θ is realized, sender commits to a signaling scheme (i.e., a randomized map from *states of nature* to signals)
  - Private scheme: different (possibly correlated) signals to different receivers
  - Public scheme: the same signal to each receiver
- After θ realized, sender sample signals and then communicate them to receivers

Model : Persuasion with No Externalities

[Arieli/Babichenko'16] characterizes optimal *private* signaling scheme for *special classes* of $f(S)$ when *two states* of nature.

This work: pin down complexity of optimal private and public persuasion for natural classes of sender objectives

Outline

- Introduction and Model
- Persuasion through the Algorithmic Lens
- Conclusion

Why Algorithms?

- ➢ Enable automated application

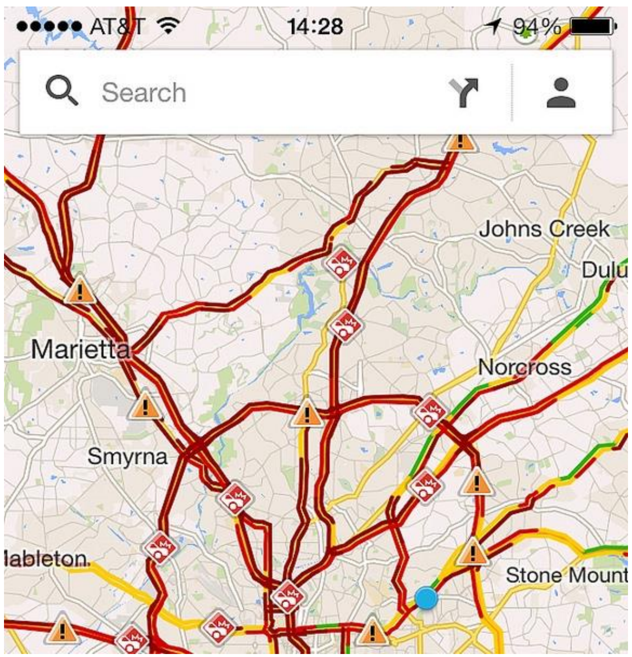

Persuading selfish drivers

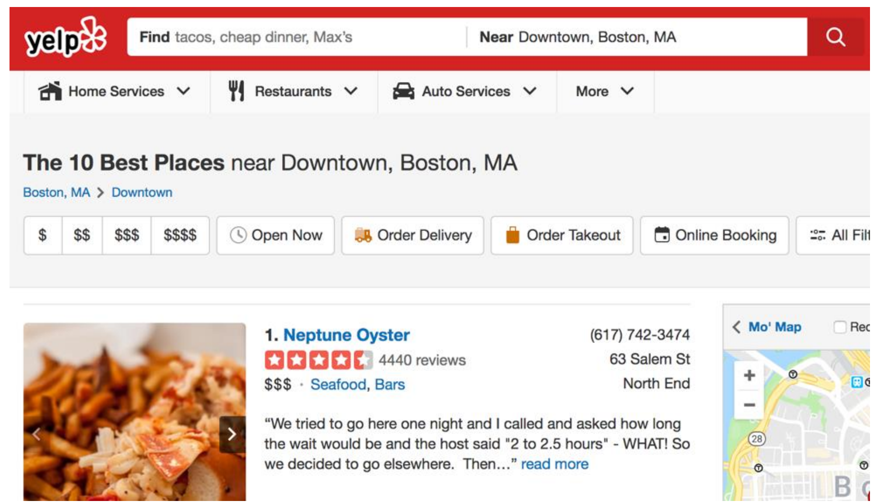

Persuading users of recommendation systems

Why Algorithms?

- Enable automated application
- Understand complexity and limitation of the model
  - Efficient computability is an important modeling prerequisite
  - Some settings are combinatorial by nature
- Lead to economic/structural insights

"If your laptop cannot find it (the equilibrium), then neither can the market."

– Kamal Jain

The Algorithmic Lens

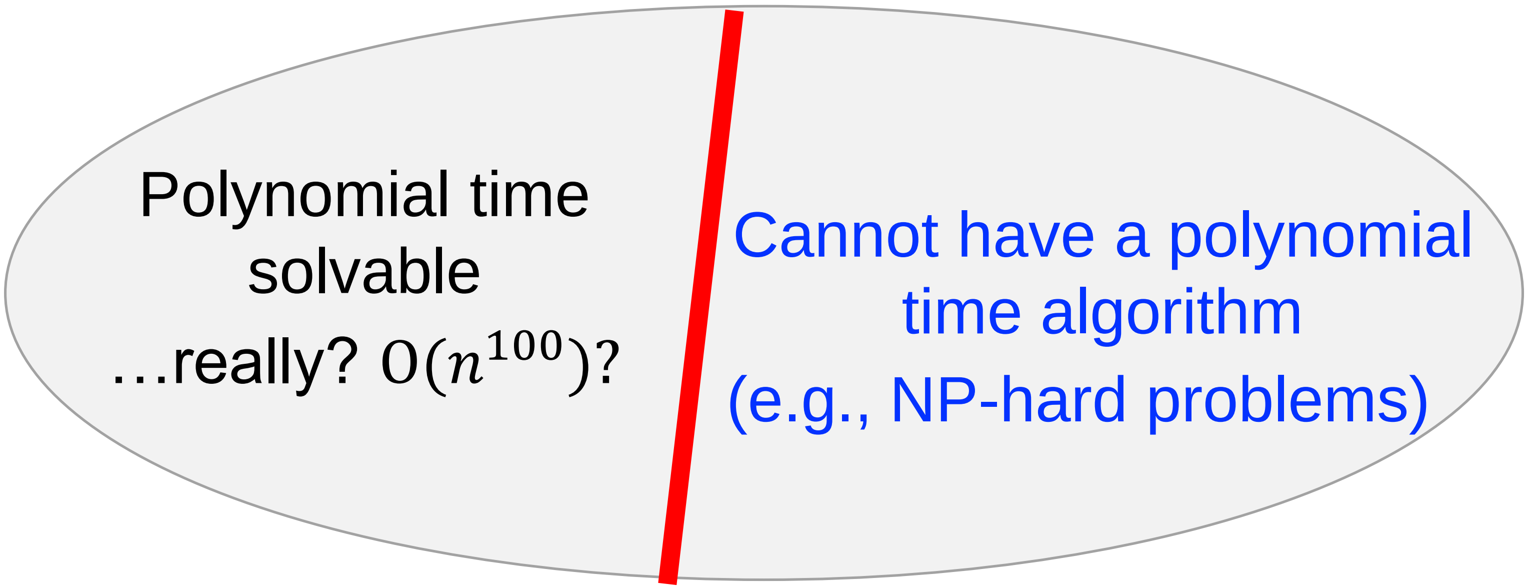

Computational Problems

- Algorithmic study seeks to understand where a problem lies

Private Persuasion

- An exponential-size linear program
- Variable $\pi(\theta, S)$ = prob of recommending action 1 to receivers in set S, given state θ
  - Each signal = an action recommendation

Expected sender utility

$$
\begin{array}{lll}
\text{maximize} & \mathbf{E}_{\theta,S}[f(S)] & \\
\text{subject to} & \mathbf{E}_{\theta,S:i\in S}[r_i(1|\theta)] \geq \mathbf{E}_{\theta,S:i\in S}[r_i(0|\theta)], & \text{for any receiver } i. \\
 & \mathbf{E}_{\theta,S:i\notin S}[r_i(0|\theta)] \geq \mathbf{E}_{\theta,S:i\notin S}[r_i(1|\theta)], & \text{for any receiver } i. \\
 & \sum_{S\subseteq[n]} \pi(\theta, S) = 1, & \text{for any state } \theta. \\
 & \pi(\theta, S) \geq 0, & \text{for } \theta, S.
\end{array}
$$

Obedience constraints

Scheme feasibility

Private Persuasion

- An exponential-size linear program
- Variable $\pi(\theta, S)$ = prob of recommending action 1 to receivers in set S, given state θ
  - Each signal = an action recommendation

$$
\begin{array}{lll}
\text{maximize} & \mathbf{E}_{\theta,S}[f(S)] & \\
\text{subject to} & \mathbf{E}_{\theta,S:i\in S}[r_i(1|\theta)] \geq \mathbf{E}_{\theta,S:i\in S}[r_i(0|\theta)], & \text{for any receiver } i. \\
& \mathbf{E}_{\theta,S:i\notin S}[r_i(0|\theta)] \geq \mathbf{E}_{\theta,S:i\notin S}[r_i(1|\theta)], & \text{for any receiver } i. \\
& \sum_{S\subseteq[n]} \pi(\theta, S) = 1, & \text{for any state } \theta. \\
& \pi(\theta, S) \geq 0, & \text{for } \theta, S.
\end{array}
$$

Can private persuasion still be done in poly time?

One approach: examine different classes of $f(S)$

Equivalence Between Private Persuasion and Optimization

Theorem: Optimal private scheme can be computed in poly time *if and only if* (unconstrained) maximization of [$f(S)$ + any modular fnc of S] can be solved in poly time.

Proof: "reduce" these two problems to each other

$$\boxed{\text{Problem A}} \xrightarrow{\text{reduce to}} \boxed{\text{Problem B}}$$

"Rephrase" or "split" problem A as a set of instances of problem B

- E.g., calculating factorial of n reduces to multiplications

Equivalence Between Private Persuasion and Optimization

Theorem: Optimal private scheme can be computed in poly time *if and only if* (unconstrained) maximization of [$f(S) +$ any modular fnc of S] can be solved in poly time.

Proof: "reduce" these two problems to each other

$\Leftarrow$: Solve the dual linear program

$\Rightarrow$: More intricate

- Involve crafting a persuasion instance to encode the set function maximization problem.

Equivalence Between Private Persuasion and Optimization

Theorem: Optimal private scheme can be computed in poly time *if and only if* (unconstrained) maximization of [$f(S)$ + any modular fnc of S] can be solved in poly time.

- Corollary: poly time for supermodular, anonymous (i.e., depend on $|S|$)
- Corollary: NP-hard for submodular, subadditive
- (Algorithmically) unifies/generalizes results from [Arieli/Babichenko '16] and some results of [Babichenko/Barman'17].

Conceptual Message

Without externalities, optimal private persuasion is closely related to directly maximizing the sender's objective without constraints

Private Persuasion: Submodular Objective

Theorem: If $f(S)$ is submodular, a $(1-1/e-\epsilon)$-optimal private scheme can be implemented in poly($n, |\Theta|, 1/\epsilon$) time.

Proof step 1: existence of a "simple" ε-optimal scheme $\{\pi(\theta, S)\}_{\theta,S}$

A Structural Lemma

There always exists an ε-optimal private scheme $\{\pi(\theta, S)\}_{\theta,S}$ such that $\pi(\theta)$ is a *uniform distribution* over poly($n, |\Theta|, 1/\epsilon$) subsets for every θ.

Private Persuasion: Submodular Objective

Theorem: If $f(S)$ is submodular, a $(1-1/e-\epsilon)$-optimal private scheme can be implemented in poly$(n, |\Theta|, 1/\epsilon)$ time.

Proof step 2: approximately compute such a “simple” scheme

- For each θ: pick poly$(n, |\Theta|, 1/\epsilon)$ subsets to maximize sender utility
- Reduce to monotone submodular maximization subject to matroid constraints.
  - (1-1/e) approximation [Calinescu et al. 2011].

Private Persuasion: Submodular Objective

Theorem: If $f(S)$ is submodular, a $(1-1/e-\epsilon)$-optimal private scheme can be implemented in poly($n, |\Theta|, 1/\epsilon$) time.

Remarks

- NP-hard to approximate to within a ratio better than (1-1/e), even with two states of nature [Babichenko/Barman'17]
- With two states, a simple scheme achieves (1-1/e)-approximation: persuade each receiver *independently* to maximize prob of action 1
  - Oblivious to sender objective as long as its submodular!
  - With many states, oblivious schemes will be far from optimality
- Open question: general equivalence between approximate private persuasion and approximate optimization

So…What About Public Scheme?

Sharp contrast to private scheme:

Theorem: For any constant α, it is NP-hard to obtain an α–approximation to optimal public scheme, even for $f(S) = |S|$.

What instances are hard?

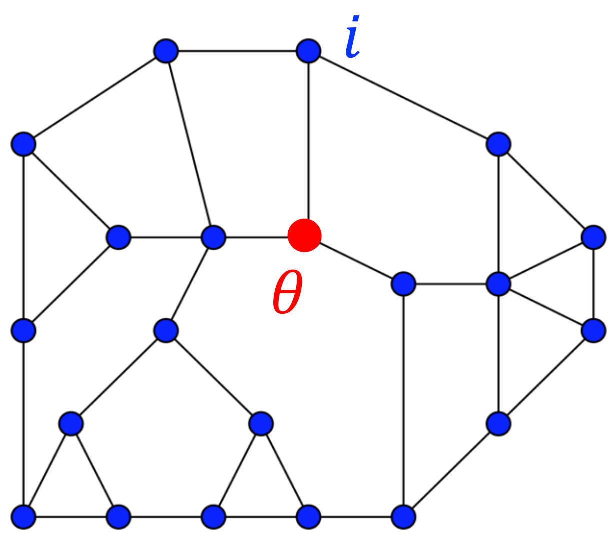

Receivers = vertices

State of nature = a uniformly drawn vertex

Similar receiver payoffs

- Action 0: always 0
- Action 1: 0.5 if $\theta = i$, -1 if θ is a neighbor of i, and 0 otherwise

Sender objective: maximize $|S|$

So…What About Public Scheme?

Sharp contrast to private scheme:

Theorem: For any constant α, it is NP-hard to obtain an α–approximation to optimal public scheme, even for $f(S) = |S|$.

What instances are hard?

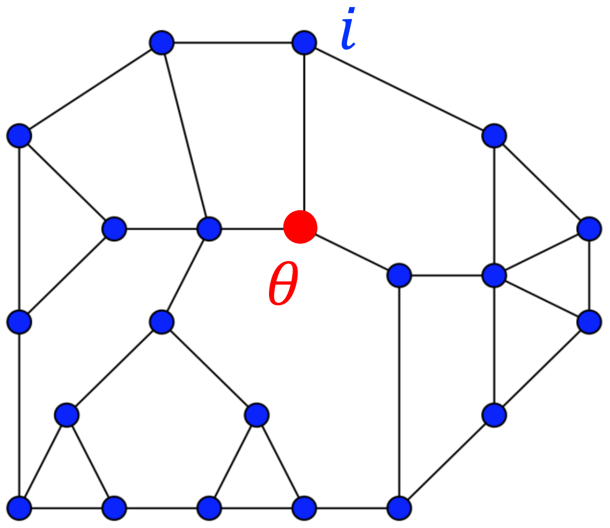

Given a public signal, i takes action 1, if

- With high chance: $\theta = i$
- With low chance, θ is a neighbor of i

So…What About Public Scheme?

Sharp contrast to private scheme:

Theorem: For any constant α, it is NP-hard to obtain an α–approximation to optimal public scheme, even for $f(S) = |S|$.

What instances are hard?

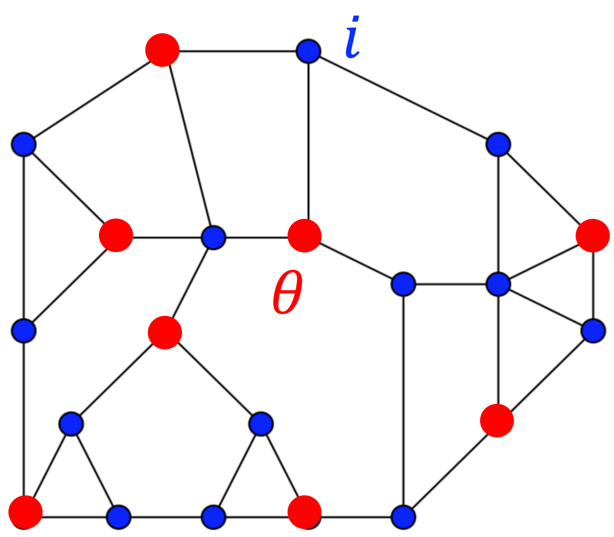

Given a public signal, i takes action 1, if

- With high chance: $\theta = i$
- With low chance, θ is a neighbor of i

In fact, two neighbor receivers will never take 1 simultaneously

A public signal = an “independent set”

So…What About Public Scheme?

Sharp contrast to private scheme:

Theorem: For any constant α, it is NP-hard to obtain an α–approximation to optimal public scheme, even for $f(S) = |S|$.

An intuitive explanation:

- Public scheme coordinates all receiver's actions simultaneously
  - Each signal gives action recommendations to all receivers
  - 2^n possible signal outcomes
- Private scheme coordinates each receiver's decisions separately
  - Each signal recommends an action to an receiver

So…What About Public Scheme?

Sharp contrast to private scheme:

Theorem: For any constant α, it is NP-hard to obtain an α–approximation to optimal public scheme, even for $f(S) = |S|$.

Conceptual Message

Private persuasion is more tractable and effective than public persuasion

Outline

- Introduction and Model
- Persuasion through the Algorithmic Lens
- Conclusion

Summary

- Systematic algorithmic study for a basic model of persuading multiple agents with no externalities

Private Persuasion	Public Persuasion
Tractable, Effective	Intractable, Ineffective

Immediate Open Questions

- Approximate version of the poly-time equivalence between private persuasion and optimization
- Receivers can share their signals
- Externalities

Some Applications of Persuasion

Law enforcement
[XRDT'15, HN'18]

Conservation drones
[XWVT'18]

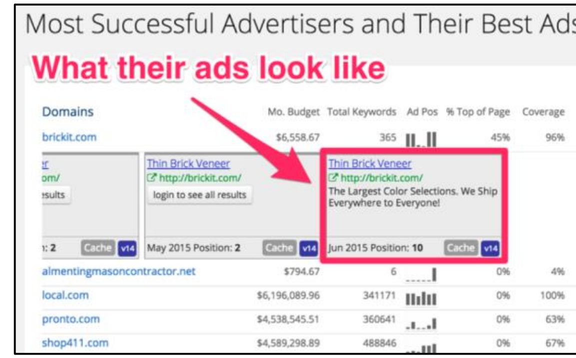

Ad auctions
[EFGLT'12, BBX'18]

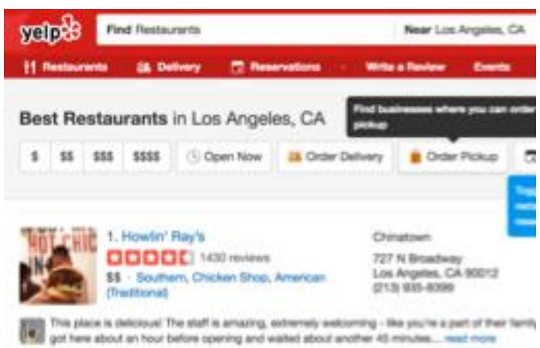

Recommendation systems
[MSS'15, MSSW '16]

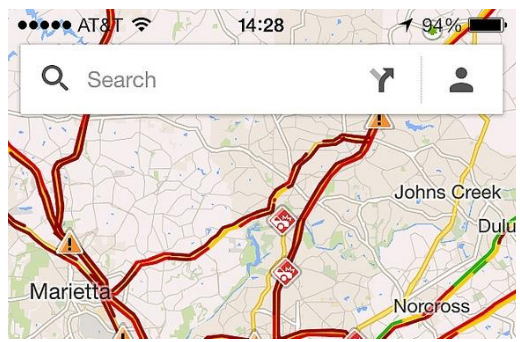

Traffic routing
[VFH'15, BCKS '16]

Queueing systems
[LI'17]

Thank You

Questions?