Models, Algorithms, and Evaluation for Autonomous Mobility-On-Demand Systems

Marco Pavone
Autonomous Systems Laboratory
Department of Aeronautics and Astronautics
Stanford University

Decision Support for Traffic Workshop
Institute for Pure & Applied Mathematics
UCLA

November 17, 2015

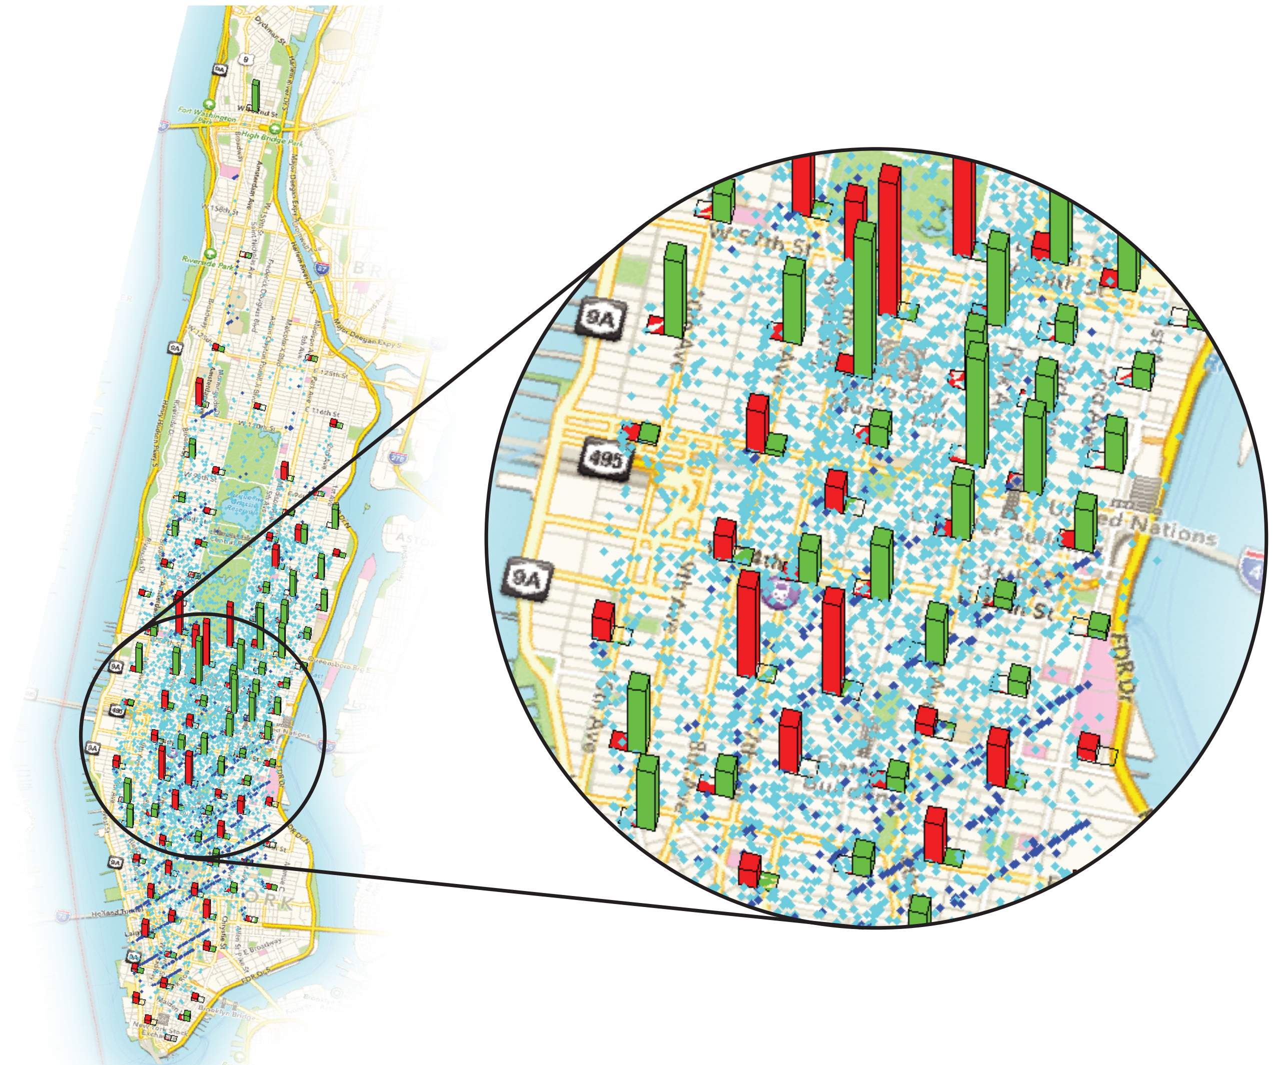

Stanford University

Autonomous Systems Lab

Challenges for Personal Urban Mobility in the 21st Century

Demand for personal urban mobility:

- Urban population will increase by 36% in the next 20 years, reaching the value of 5.4 billions (more than 60% of overall population) [United Nations, ‘14]

Constraint:

- Available urban land for roads and parking is continuously decreasing (limited infrastructure)
- Roadway safety, pollution, etc. (sustainability)

Fact:

- Current urban transportation system very similar to the one conceived by Karl Benz and Henry Ford 100 years ago

Mobility-on-demand (MoD) systems

Key challenge: ensure sustainable personal mobility at a reduced cost

MoD systems: convergence of four emerging technologies [Mitchell, Borroni-Bird, Burns, MIT Press '10]

- Shared vehicles
- The "Mobility Internet"
- Specific-purpose vehicle designs
- Advanced propulsion systems

Key driver: from low (< 10%) to high (> 90%) utilization rates

Limitations of Mobility on Demand systems

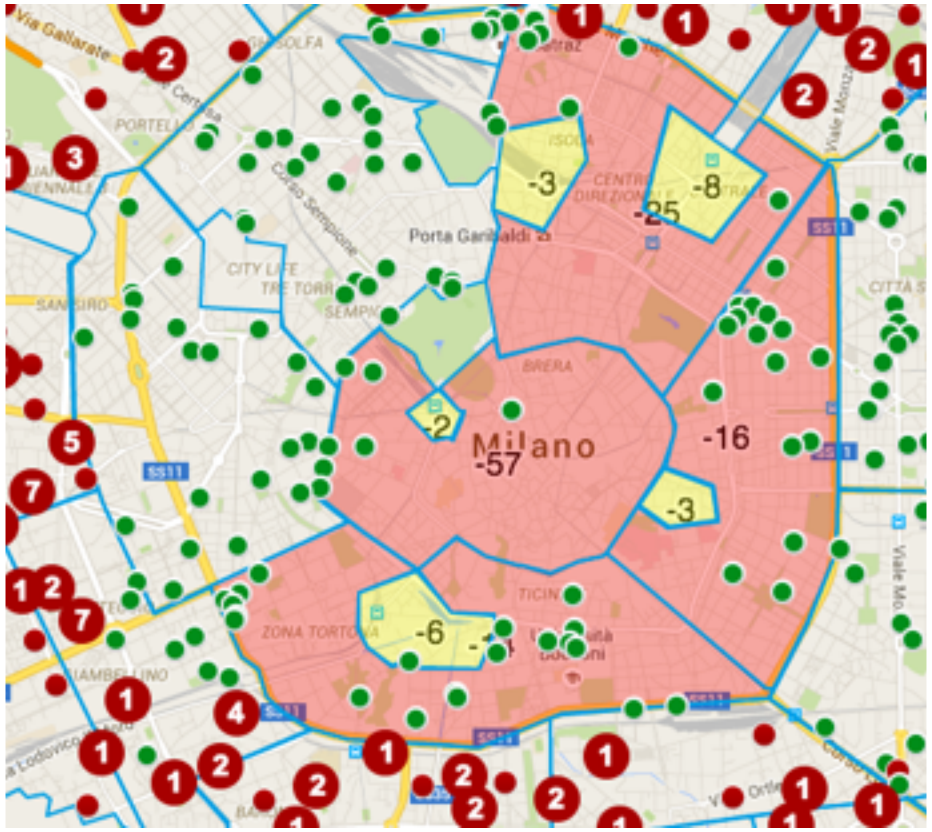

- Honda DIRACC in Singapore closed in 2008 after 6 years: "*Everybody expected cars to be available. But in reality, we could not guarantee*."
- Car2go in the US and elsewhere: Parking, availability and "technology hiccups" were most cited pain points of current MoD model (private communication, but also yelp.com)
- Uber: Higher cost, surge pricing, 80% of the fares go to the drivers.

An Emergent Technology: Self-Driving Cars

Potential Benefits of Self-Driving Cars (US Market)

- **Safety**[1]:
  - Value of a statistical life = $9.1M
  - Economic cost of traffic accidents = $242B
  - Societal harm of traffic accidents (loss in lifetime productivity) = $594B
- **Value of time** (productivity/leisure)[4] = $1.3T
- **Throughput**:
  - Economic cost of congestion (time/fuel wasted)[2] = $160B
  - Health cost of congestion (pollution)[3] = $15B
- **Enabling carsharing on a massive scale**[4] = $402B

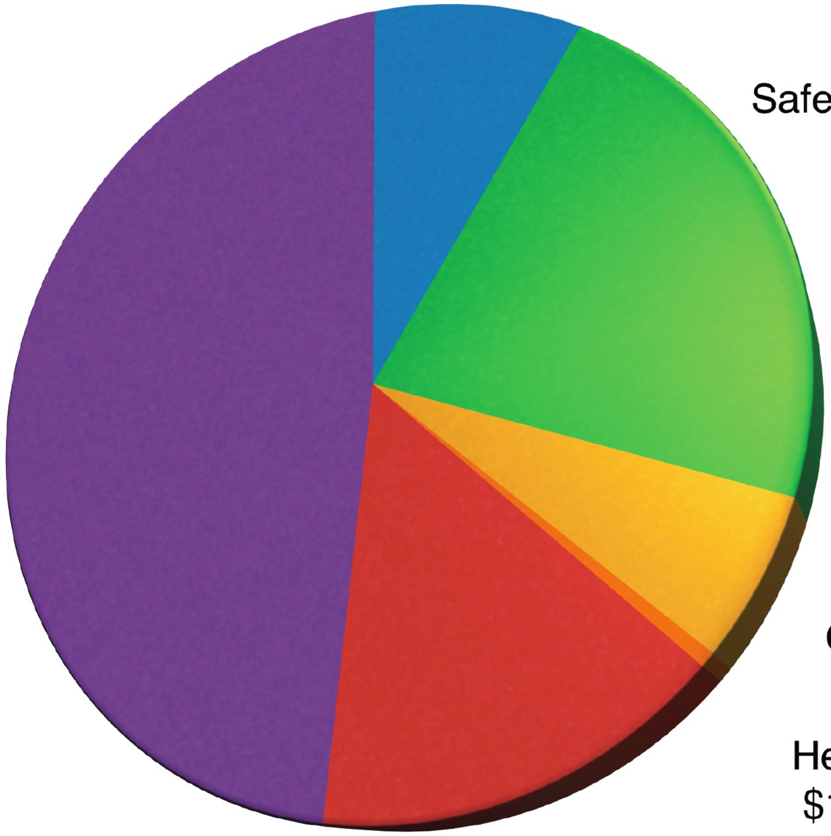

[1] [Blincoe et al., NHTSA Report, ’15]
[2] [Schrank et al., Texas A&M Transportation Institute, ’15]
[3] [Levy et al., Environmental Health, ’10]
[4] [Spieser, Treleaven, Zhang, Frazzoli, Morton, Pavone, Road Vehicle Automation, ’14]

A New Paradigm for Personal Urban Mobility

Vehicle Autonomy

+

Car Sharing

CAR 2GO

UBER

autolib'

Autonomous Mobility-On-Demand (AMoD)

Research objectives:

1. **Modeling**: stochastic models for tractable analyses
2. **Control**: real-time routing of autonomous vehicles at a city-wide scale
3. **Applications**: case studies and technology infusion

How to Control a Fleet of Autonomous Vehicles?

Problem falls under the general class of networked, heterogeneous, stochastic decision problems with uncertain information:

- **Problem Data / Model**: travel demand, road network
- **Control inputs**: vehicle routing, passenger loading/unloading
- **Outputs**: customer waiting times, customer queue lengths, etc.

Key features:

Static version NP-hard

Dynamics add queueing phenomena

Closed system:
cascade feedback effects

Closed-loop control policies aimed at optimal throughput

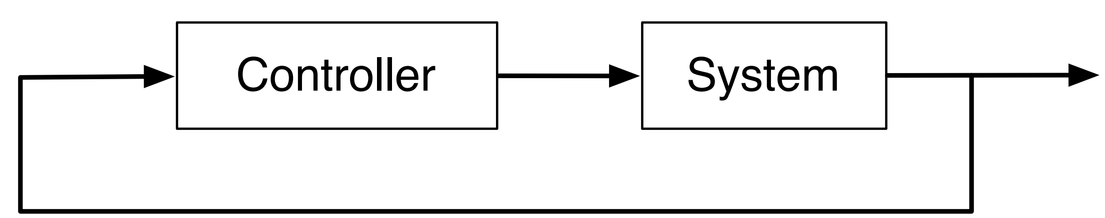

A Family of Models for Control and Evaluation

Macroscopic → Microscopic

Distributed queueing-theoretical models	Lumped queueing-theoretical models	(Stochastic) MPC models
[Pavone, MIT ’10], [Treleaven, Pavone and Frazzoli, TAC ’13]	[Pavone, Smith, Frazzoli, Rus, IJRR ’12], [Zhang, Pavone, IJRR ’15]	[Zhang, Rossi, Pavone, ICRA ’16], [Chow, Yu, Pavone, AAAI, ‘16]

Analytical → Computational

Overarching goals:

1. Theoretical insights and guidelines for system design
2. Real-time control algorithms
3. Formal guarantees for stability and performance

Models and Controls

Distributed Queueing-Theoretical Models

Model:

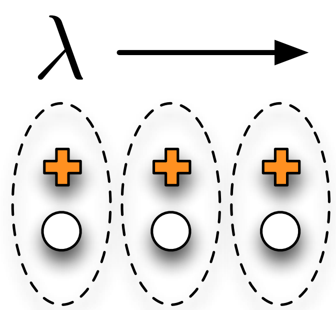

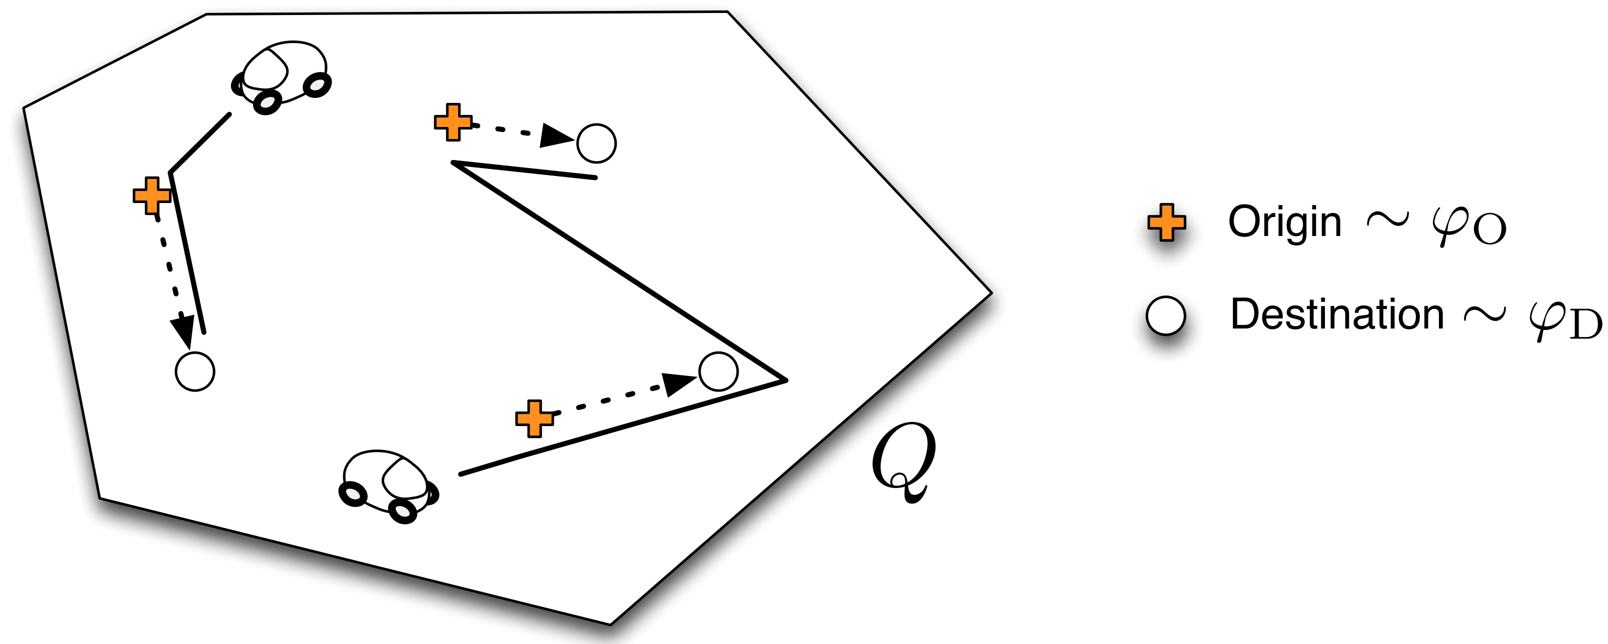

- Poisson process generates origin-destination pairs with rate λ
- Origin-destination pairs distributed with distributions φ_{O} and φ_{D}
- Goal: minimize steady-state expected system time

Challenge: *spatial component introduces correlations among service times*

Analysis and synthesis approach:

Step 1) Queueing model of system and analysis of its structure

Step 2) Fundamental limitations on performance

Step 3) Real-time algorithms with provable performance guarantees

[Bertsimas and Van Ryzin, OR ‘91], [Pavone, MIT ’10], [Bullo, Frazzoli, Pavone, Savla, Smith, IEEE Proc. ’11]

Stability Conditions

Key advantage: *Analytical expressions*, e.g., NSC for stability

$$\varrho := \lambda \left[\overbrace{\mathbb{E}_{\varphi_O \, \varphi_D}[D - O]}^{\text{Busy travel}} + \overbrace{\mathrm{EMD}(\varphi_O, \varphi_D)}^{\text{Unavoidable empty travel}} \right] / m < 1$$

where

$$\mathrm{EMD}(\varphi_1, \varphi_2) = \inf_{\gamma \in \Gamma(\varphi_1, \varphi_2)} \int_{X \times X} D(x_1, x_2) d\gamma(x_1, x_2)$$

- Also known as the first Wasserstein Distance
- EMD can be interpreted as the amount of work required to turn the pile of earth described by φ_1 into that described by φ_2

Key drawbacks: Performance results in asymptotic regimes, i.e. $\rho \to 0^+$ and $\rho \to 1^-$, difficult to include road topology

More at http://web.stanford.edu/~pavone/misc/dvr.pdf

[Pavone, MIT ’10], [Treleaven, Pavone and Frazzoli, TAC ’13]

Lumped (Jackson) Queueing-Theoretical Model

Model:

- Stochastic model with passenger loss
- λ_i - arrival rate of customers
- p_{ij} - routing probabilities
- T_{ij} - travel times
- Balance Equations: $\pi_i = \sum_j \pi_j p_{ji}$
- Stationary distribution:

$$\mathbb{P}(X) = \frac{1}{G(m)} \prod_{j=1}^{|\mathcal{N}|} \pi_j^{x_j} \prod_{n=1}^{x_j} \mu_j(n)^{-1}$$

- Relative utilization: $\gamma_i = \pi_i / \lambda_i$

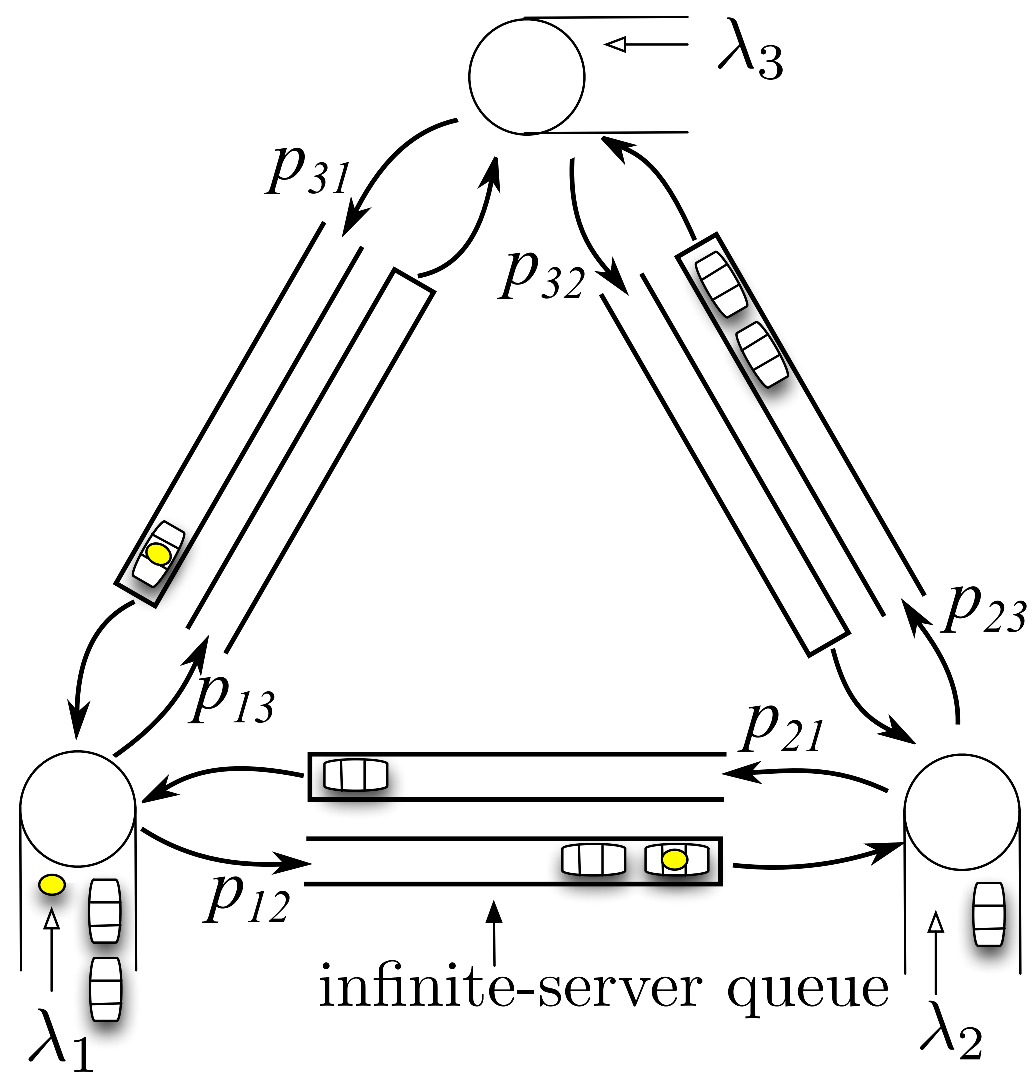

Performance metric: vehicle availability $A_i(m) = \gamma_i \frac{G(m-1)}{G(m)}$

Key advantage: “natural model,” captures network topology

[Zhang, Pavone, IJRR ’15]

Control Strategies for Lumped Model

Key idea: Stochastic "rebalancing-promoting" policy using "virtual" customers with arrival rates ψ_i and routing probabilities α_{ij}

Optimal Rebalancing Problem (ORP): *Given an autonomous MOD system modeled as a closed Jackson network, solve*

$$
\begin{aligned}
\underset{\psi_i,\alpha_{ij}}{\text{minimize}} \quad & \sum_{i,j} T_{ij}\,\alpha_{ij}\psi_i \\
\text{subject to} \quad & \gamma_i = \gamma_j \\
& \sum_j \alpha_{ij} = 1 \\
& \alpha_{ij} \geq 0, \;\; \psi_i \geq 0 \qquad i,j \in \{1,\ldots,N\}
\end{aligned}
$$

where $\gamma_i = \dfrac{\pi_i}{\lambda_i + \psi_i}$

[Zhang, Pavone, IJRR '15]

Results

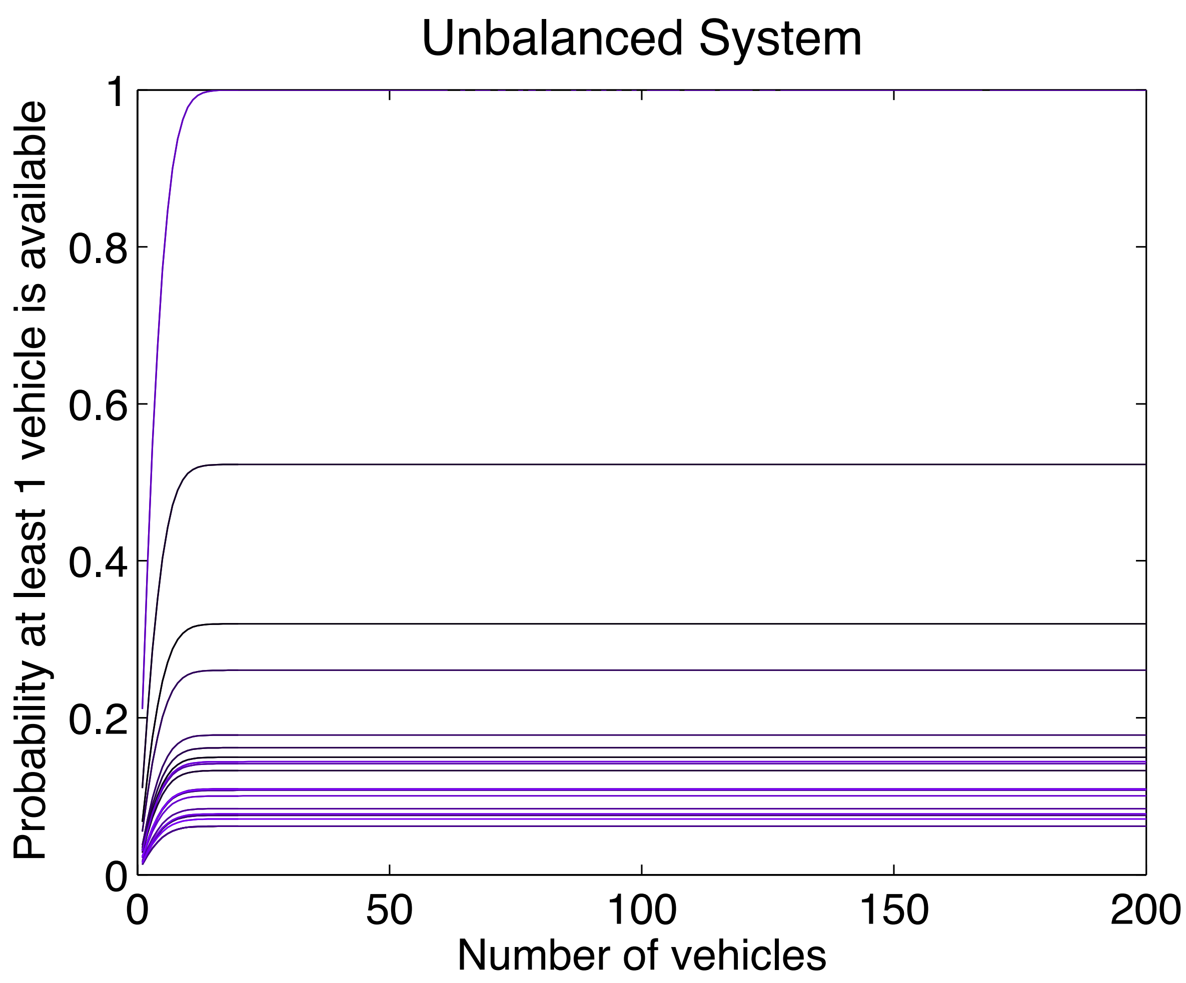

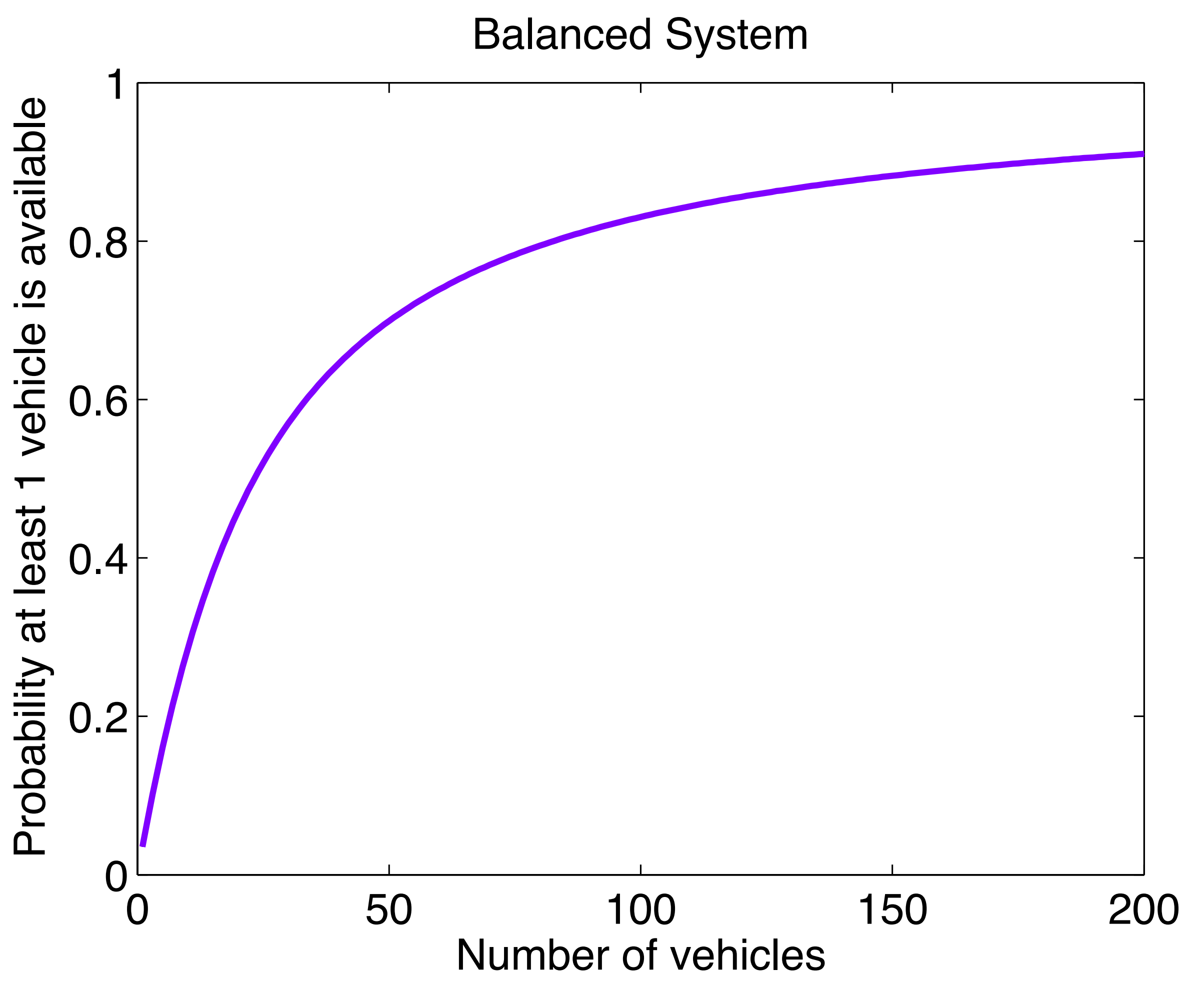

- Provides theoretical justification to earlier fluidic approximations [Pavone, Smith, Frazzoli, Rus, IJRR '12]

[Zhang, Pavone, IJRR '15]

Real-Time Control Algorithms

- Optimal rebalancing policy found by solving a linear program

$$\begin{aligned} \underset{\beta_{ij}}{\text{minimize}} \quad & \sum_{i,j} T_{ij}\beta_{ij} \\ \text{subject to} \quad & \sum_{j\neq i}(\beta_{ij}-\beta_{ji}) = -\lambda_i + \sum_{j\neq i} p_{ji}\lambda_j \\ & \beta_{ij} \geq 0 \end{aligned}$$

- In practice, routing problem solved with an MPC-like algorithm

$$\begin{aligned} \underset{\mathrm{num}_{ij}}{\text{minimize}} \quad & \sum_{i,j} T_{ij}\mathrm{num}_{ij} \\ \text{subject to} \quad & v_i^e(t) + \sum_{j\neq i}(\mathrm{num}_{ji} - \mathrm{num}_{ij}) \geq v_i^d(t) \;\; \text{for all } i \in S \\ & \mathrm{num}_{ij} \in \mathbb{N} \;\; \text{for all } i,j \in S, i \neq j \end{aligned}$$

[Zhang, Pavone, IJRR ’15]

System-Level Control of MoD Systems

[Zhang, Pavone, IJRR ’15]

Model Predictive Control Approach

Key advantages: Wait times can be minimized directly, and can taken into account practical constraints such as battery charging

Model highlights:

- $d_{ij}(t)$ - number of customers waiting at station *i* to go to station *j* at time *t*
- $q^k(t)$ - charge remaining on vehicle *k* at time *t*
- $w_{ij}^k(t)$ - vehicle *k* rebalances from *i* to *j* at time *t*

Objectives:

- $J_x(x(t)) = \sum_{i,j \in \mathcal{N}} d_{ij}(t)$ (minimize number of waiting customers)
- $J_u(u(t)) = \sum_{k \in \mathcal{V}} \sum_{i,j \in \mathcal{N}} T_{ij} w_{ij}^k(t)$ (minimize rebalancing)

$$
\begin{aligned}
\underset{u(t),\ldots,u(t+t_{\text{hor}}-1)}{\text{minimize}} \quad & \sum_{\tau=t}^{t+t_{\text{hor}}-1} J_x(x(\tau+1)) + \rho J_u(u(\tau)) \\
\text{subject to} \quad & x(\tau+1) = Ax(\tau) + Bu(\tau) + c(\tau) \\
& x(\tau+1) \in \mathcal{X} \\
& u(\tau) \in \mathcal{U}(\tau) \\
& \tau = t, \ldots, t+t_{\text{hor}}-1.
\end{aligned}
$$

Performance Comparison

Simulation scenario: 40 vehicles, 15 stations, real taxi data from NYC Financial District

Algorithms evaluated:

- Nearest-neighbor dispatch
- Collaborative dispatch[1]
- Markov redistribution[2]
- Real-time rebalancing algorithm[3]
- MPC (sampled) – arrival statistics computed using historical data, and sampled as Poisson arrivals
- MPC (full information) – actual customer arrivals over the time horizon are fed into the MPC algorithm

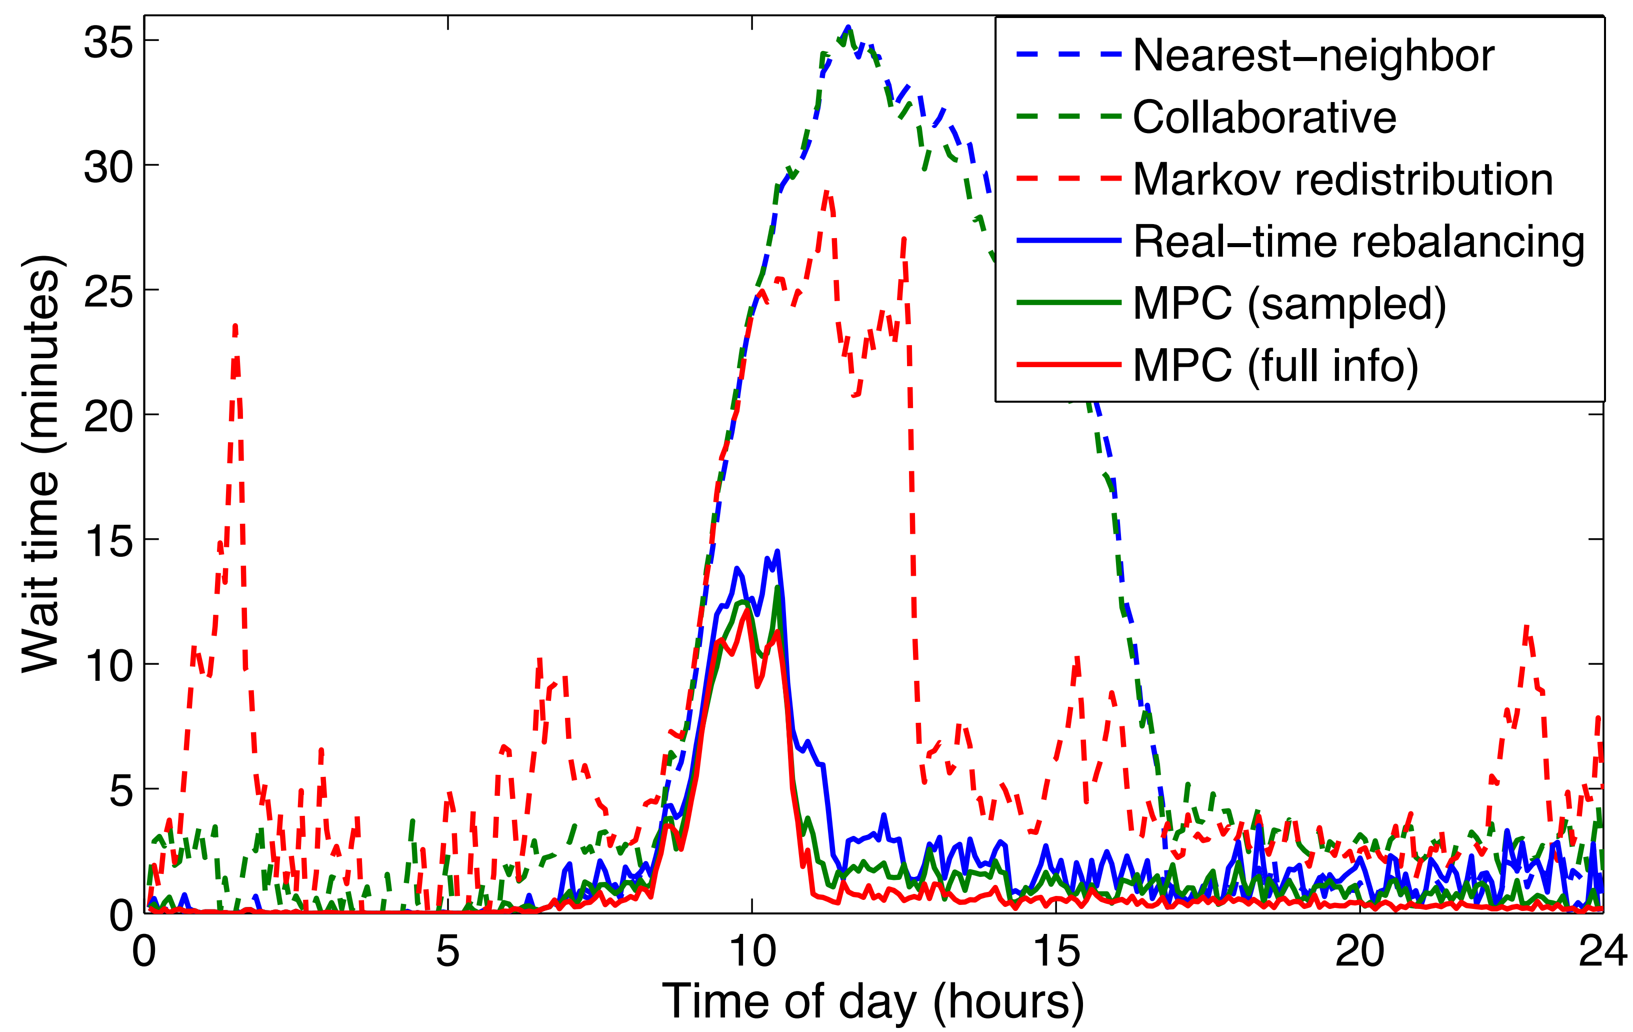

1 [Seow et al., TASE '10]

2 [Volkov et al., CITS '12]

3 [Pavone et al., IJRR '12]

[Zhang, Rossi, Pavone, ICRA '16]

Evaluation and Open Questions

Evaluation: Case Study of Manhattan

- Trip data collected on March 1, 2012, consisting of 439,950 trips within Manhattan
- ≈ 13, 300 taxis
- 100 stations for robotic MoD

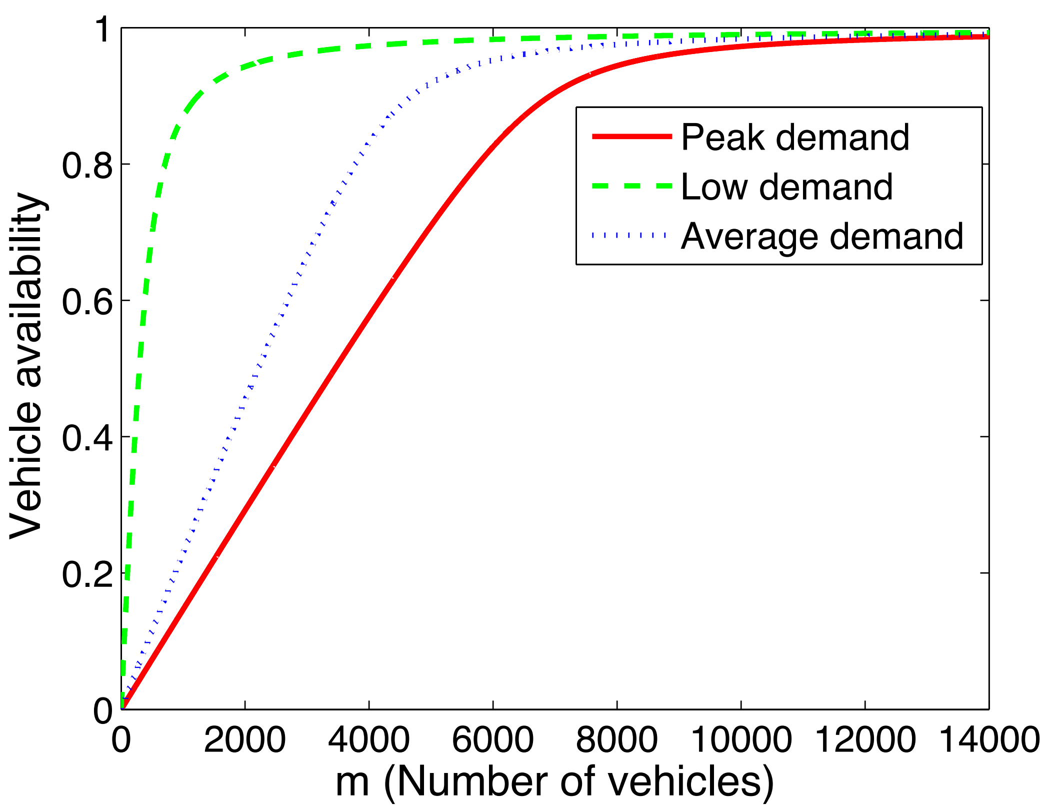

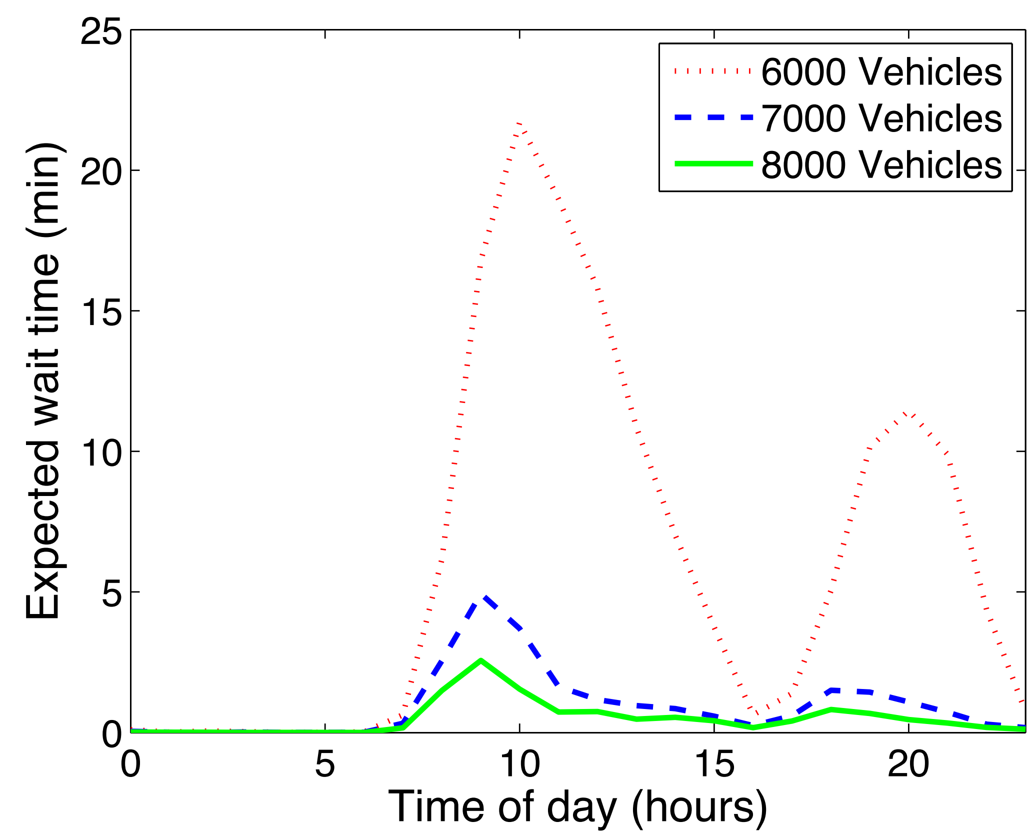

Key result: fleet size reduced by ~40%!

[Zhang, Pavone, IJRR '15]

Evaluation: Case Study of Singapore

- Three complementary data sources: HITS survey, Singapore taxi data, Singapore road network
- 779,890 passenger vehicles operating in Singapore
- 100 stations for robotic MoD

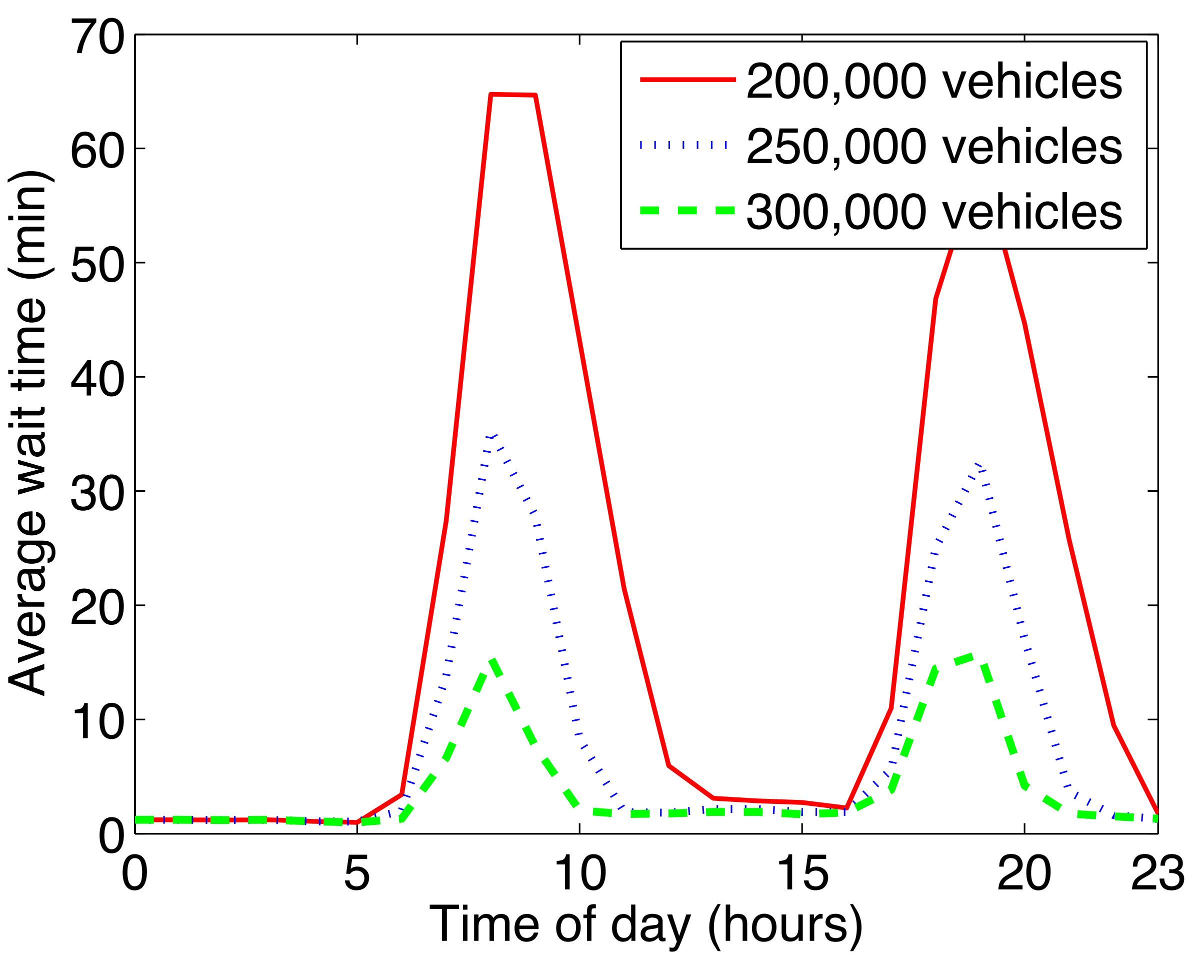

Mobility-related costs (USD/km)

	COS	COT	TMC
Traditional	0.96	0.76	1.72
AMoD	0.66	0.26	0.92

Key result: total mobility cost cut in half!

[Spieser, Treleaven, Zhang, Frazzoli, Morton, Pavone, Road Vehicle Automation, '14]

Does AMoD *Increase* congestion?

Model: simple 9-station network

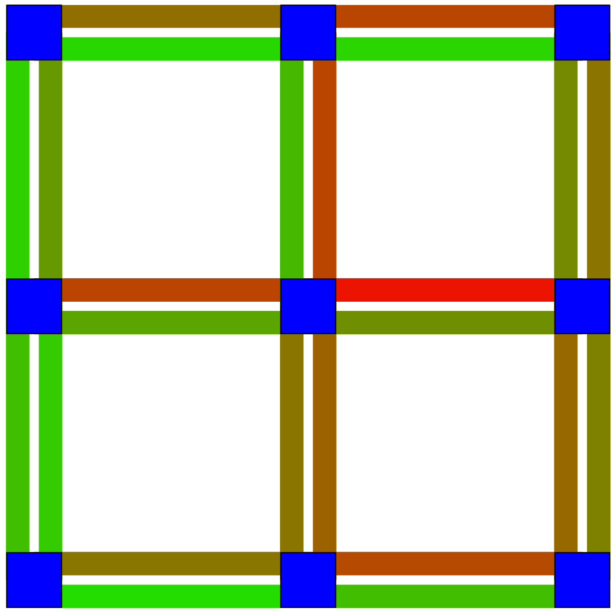

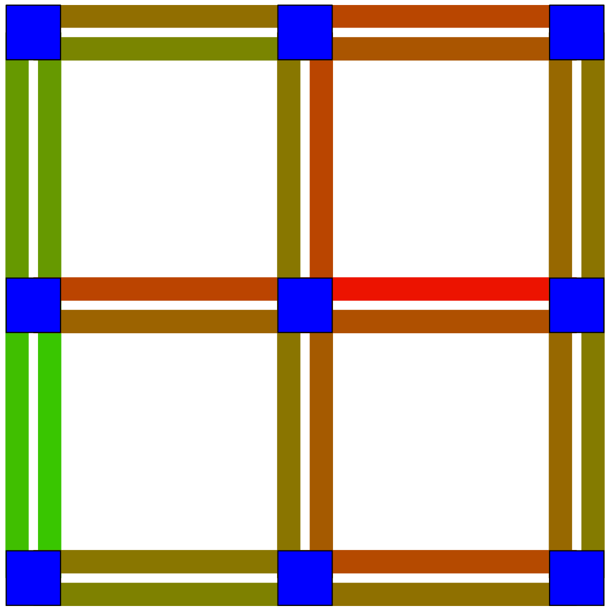

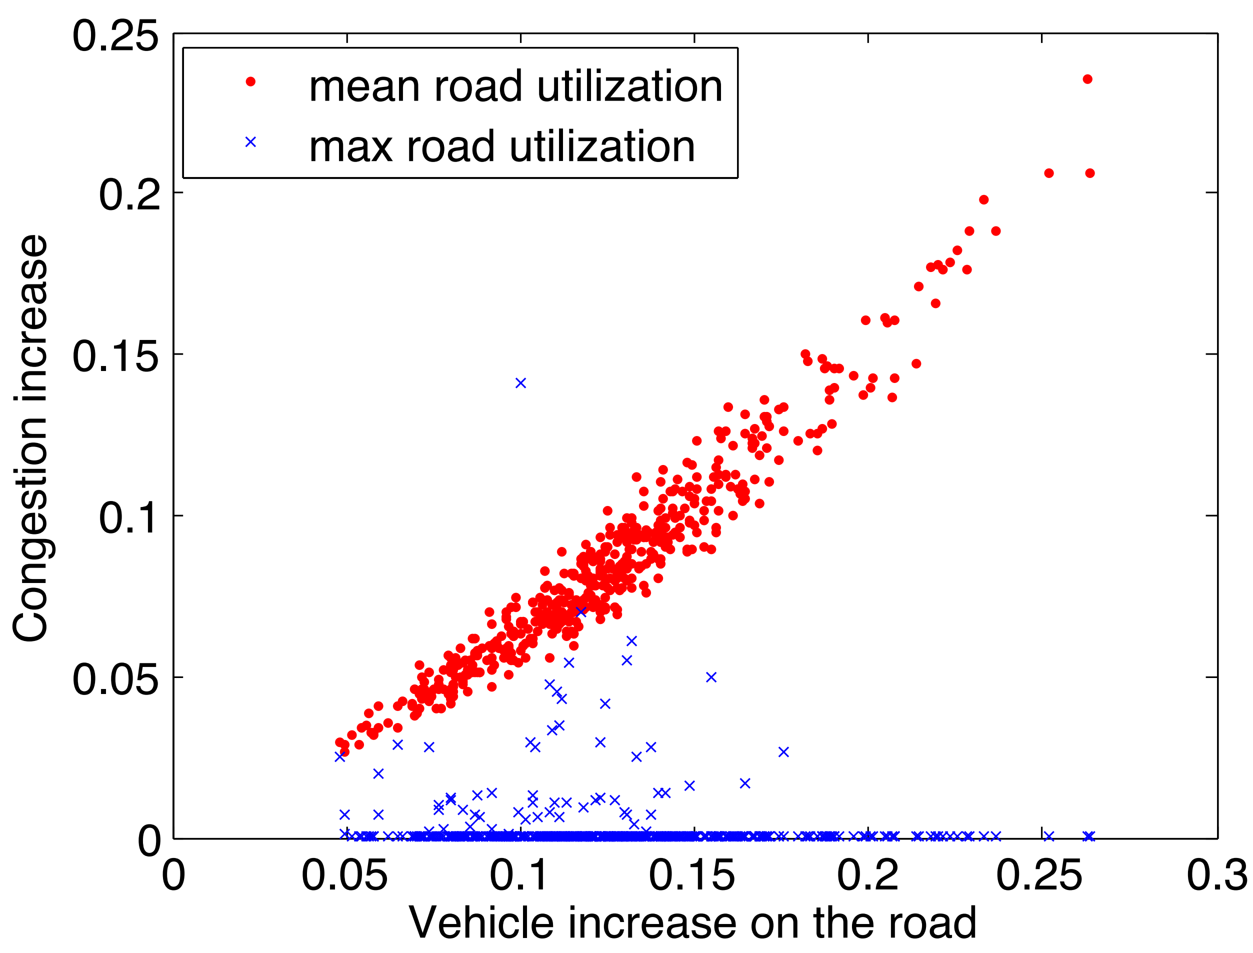

A Network-Theoretic Approach

Network flow model:

- Steady state analysis
- Road network represented by a graph $G(V, E)$ with road capacities $c(u, v) \quad u, v \in V$
- Trip requests represented by a set of sources and sinks $\{s_i, t_i\}$

Theorem: Assume there exists feasible (unbalanced) flows that satisfy the requests. Then it is possible to find a feasible rebalancing flow if and only if for every cut $(S, \bar{S})$

$$F_{\text{out}}(S, \bar{S}) \leq C_{\text{in}}(S, \bar{S})$$

Consequence: Given a *symmetric* road network, it is *always possible* to rebalance an AMoD system *without increasing congestion* in steady state conditions

Example

Procedure:

1. Solve multi-commodity flow problem for the unbalanced flows (NP-hard)
2. Given the unbalanced flows, the rebalancing flows can be solved as a linear program (totally unimodular constraint matrix)

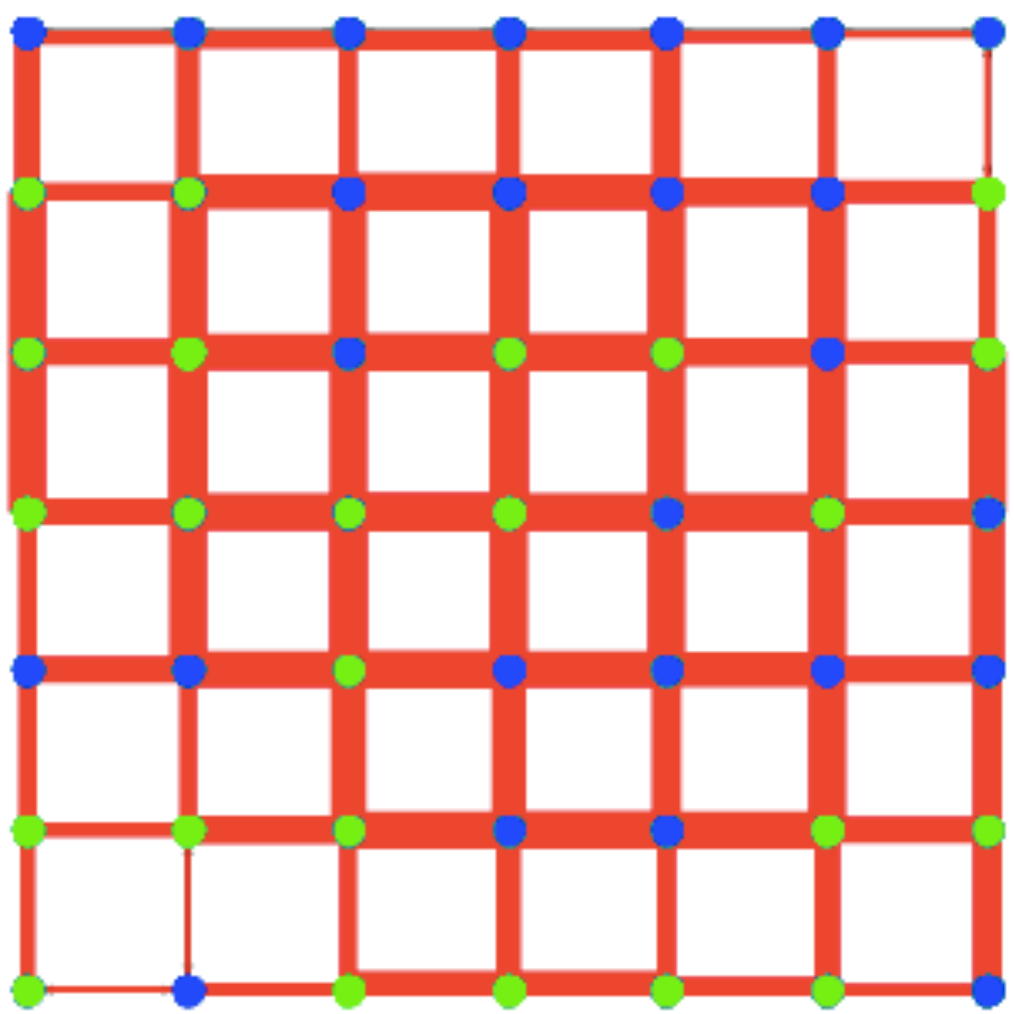

Road utilization from unbalanced flows

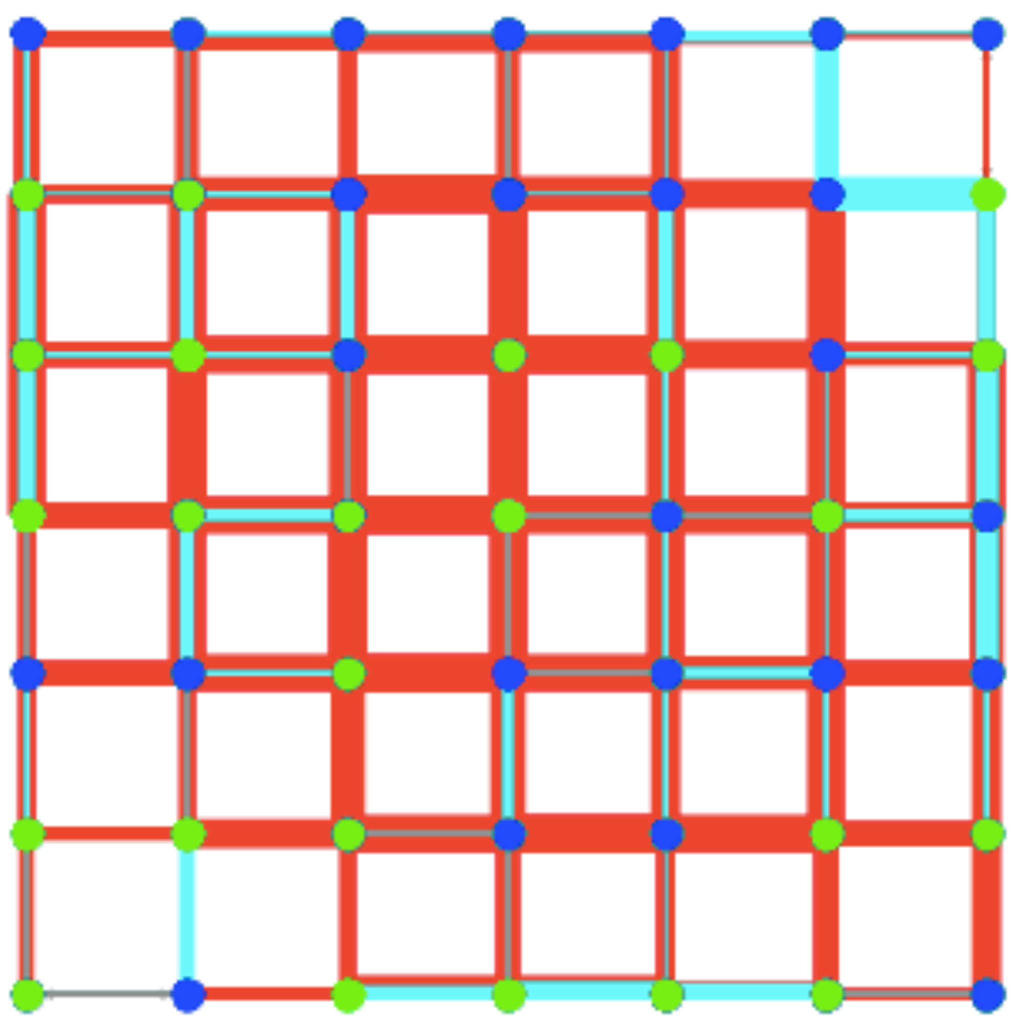

Road utilization after rebalancing

Does AMoD *Decrease* congestion?

Staggering strategy: if trips are staggered to avoid too many trips at the same time, congestion *may* be reduced

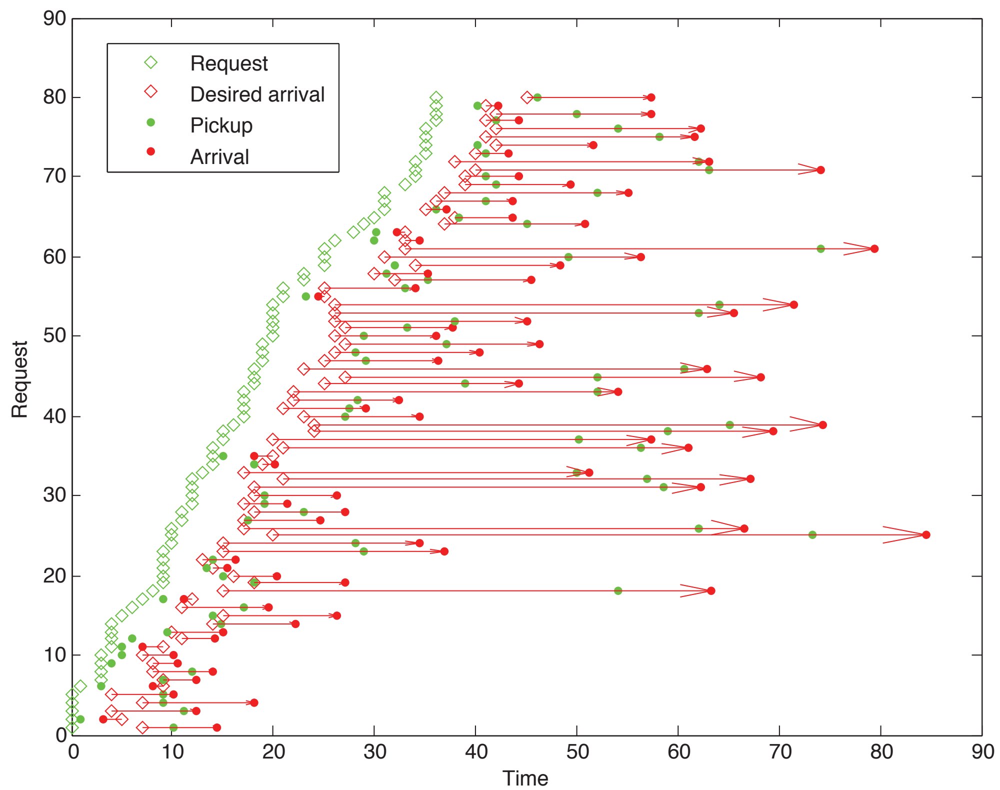

Nearest neighbor

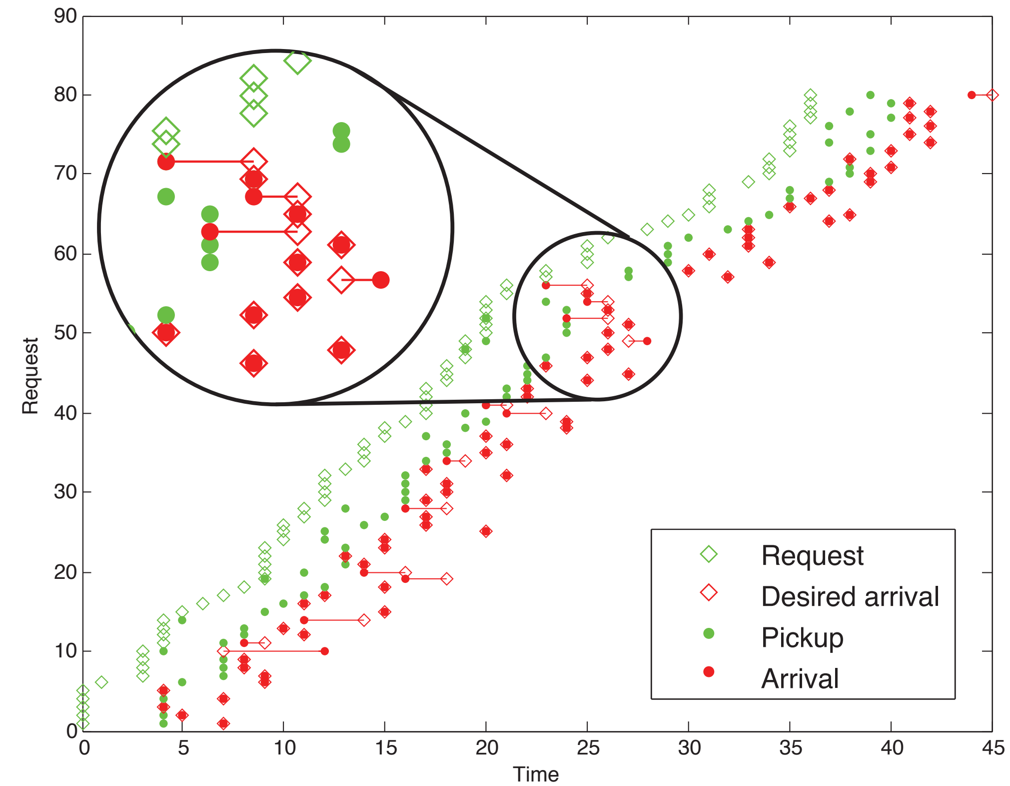

Congestion-aware staggering

- Simultaneously schedule pickup times and plan congestion-free routes
- Passengers can be dropped off before their preferred arrival time
- Passenger satisfaction is maximized with an MPC controller

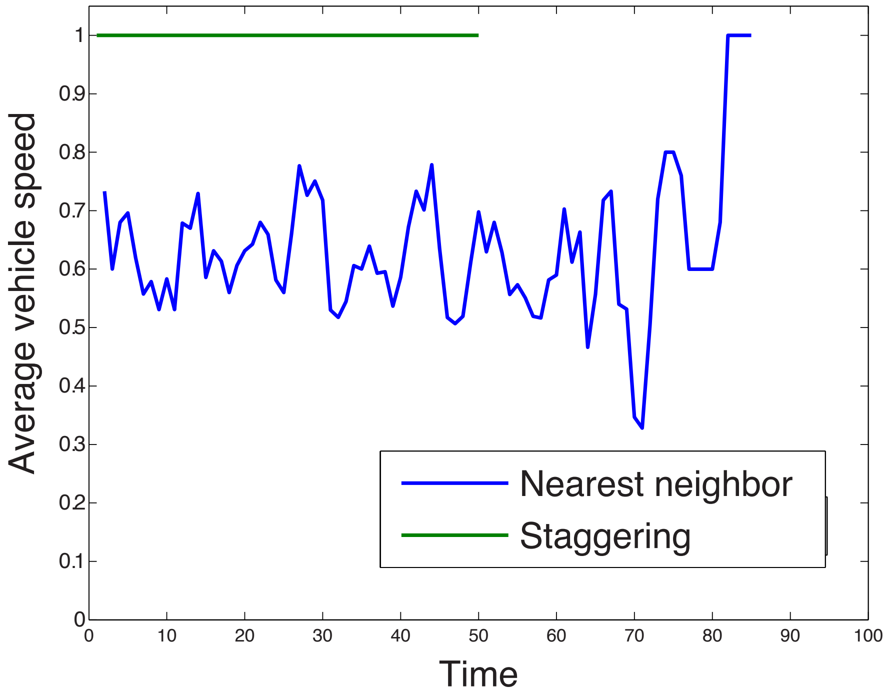

Conclusions

1. Autonomous driving might lead to a transformational paradigm for personal urban mobility (to improve or sustain current mobility needs)
2. Integration of system-wide coordination and autonomous driving gives rise to an entirely new class of problems at the interface of robotics and transportation research
3. Solutions to these problems are key to enable autonomous MoD and to carefully evaluate their value proposition

Current research directions:

- Inclusion of increasingly sophisticated congestion models
- Staggering of customers
- Controlling AMoD systems as part of a multi-modal transportation network
- Technology infusion and advising policy makers

Contact: pavone@stanford.edu

Acknowledgements: NSF CAREER Award, collaborators at Stanford (Rick Zhang and Federico Rossi), MIT, and Singapore