Balancing Pickup Time and Idle Time in Ride-sharing Systems

Hongzhang Shao

September 19, 2021

School of Industrial and Systems Engineering
Georgia Institute of Technology
steveshao@gatech.edu

INFORMS 2021

Agenda

What is "*pickup time*"? **What** is the "*trade-off*"?

Why should we care about it?

How can we model pickup time tractably?

What is “pickup time”? What is the “trade-off ”?

Different States of a Driver in Ride-sharing Systems

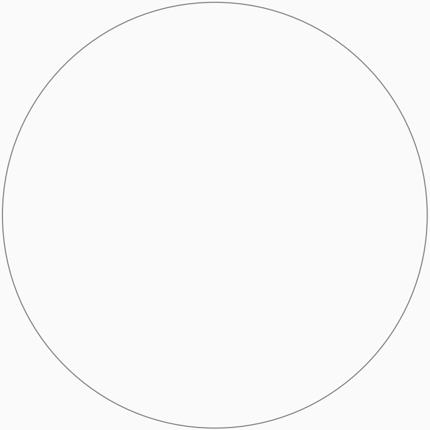

Car-times of a driver:

Different States of a Driver in Ride-sharing Systems

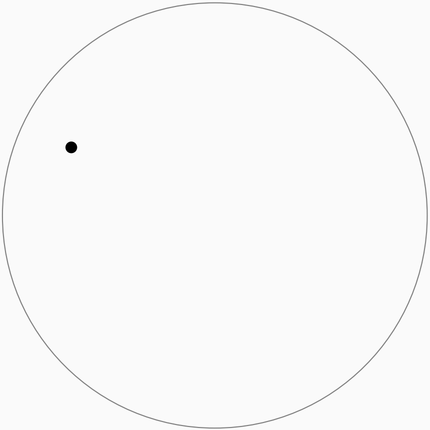

Car-times of a driver:

- Driver is waiting idle.

Different States of a Driver in Ride-sharing Systems

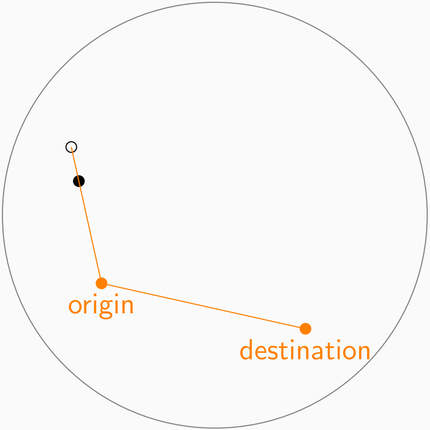

Car-times of a driver:

- Driver is waiting idle.
- Passenger requests a ride.
  The driver is dispatched.

Different States of a Driver in Ride-sharing Systems

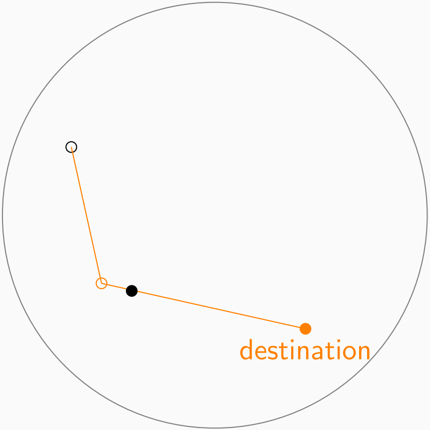

Car-times of a driver:

- Driver is waiting idle.
- Passenger requests a ride. The driver is dispatched.
- Driver picks up the passenger. Delivery starts.

Different States of a Driver in Ride-sharing Systems

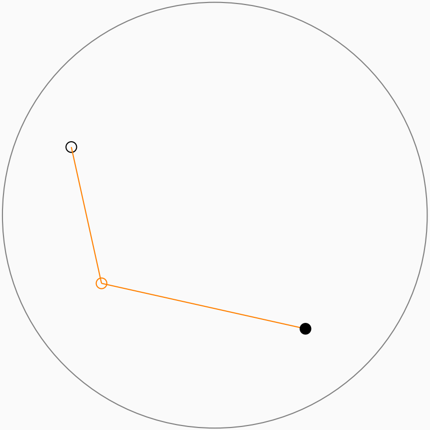

Car-times of a driver:

- Driver is waiting idle.
- Passenger requests a ride. The driver is dispatched.
- Driver picks up the passenger. Delivery starts.
- Passenger is dropped off. Driver becomes idle again.

Different States of a Driver in Ride-sharing Systems

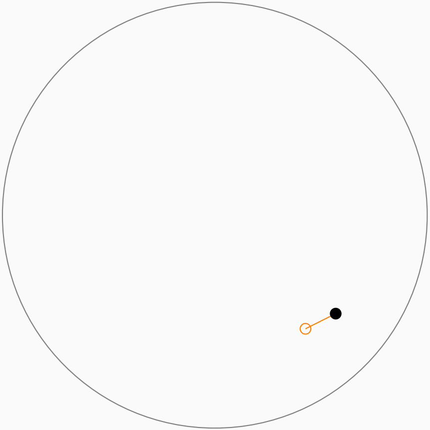

Car-times of a driver:

- Driver is waiting idle.
- Passenger requests a ride. The driver is dispatched.
- Driver picks up the passenger. Delivery starts.
- Passenger is dropped off. Driver becomes idle again.
- Driver can also reposition without a passenger.

Pickup Time and the Trade-off

Assume that:

- the space is partitioned into multiple zones
- within each zone, idle cars are uniformly distributed
- when a request comes, the closest idle car is dispatched

Pickup Time and the Trade-off

Assume that:

- the space is partitioned into multiple zones
- within each zone, idle cars are uniformly distributed
- when a request comes, the closest idle car is dispatched

Then the pickup time is a random variable that depends on the number / density of idle cars.

Pickup Time and the Trade-off

Assume that:

- the space is partitioned into multiple zones
- within each zone, idle cars are uniformly distributed
- when a request comes, the closest idle car is dispatched

Then the pickup time is a random variable that depends on the number / density of idle cars.

Higher idle car density means shorter pickup time – They cannot be both small, and there is a trade-off.

Relationship between Pickup Time and Idle Time

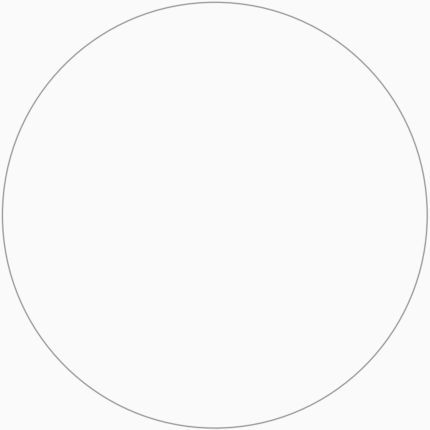

[Arnott, 1996]:

- consider a steady-state, continuous approximation

Relationship between Pickup Time and Idle Time

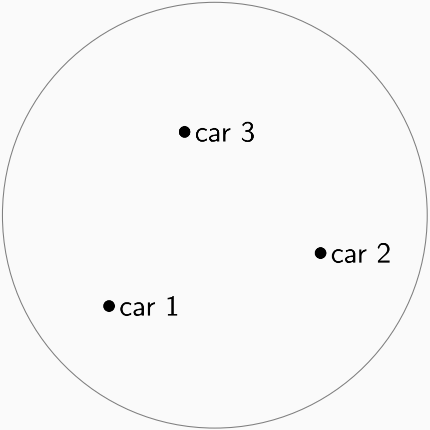

[Arnott, 1996]:

- consider a steady-state, continuous approximation
- idle cars are distributed according to a (2-D) Poisson process with rate γ)

Relationship between Pickup Time and Idle Time

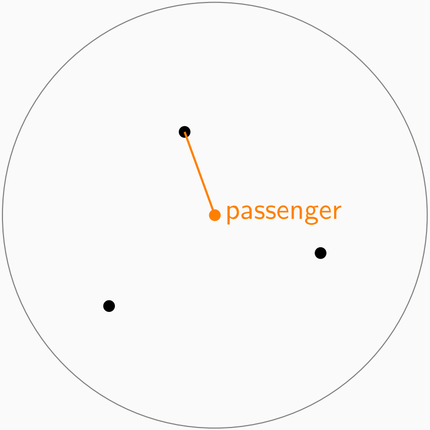

[Arnott, 1996]:

- consider a steady-state, continuous approximation
- idle cars are distributed according to a (2-D) Poisson process with rate γ)
- the closest car is matched

Relationship between Pickup Time and Idle Time

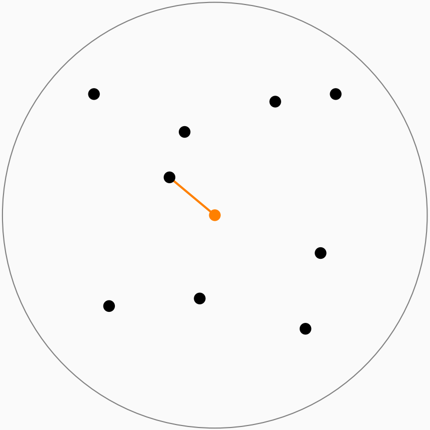

[Arnott, 1996]:

- consider a steady-state, continuous approximation
- idle cars are distributed according to a (2-D) Poisson process with rate γ)
- the closest car is matched
- more cars, shorter distance

Relationship between Pickup Time and Idle Time

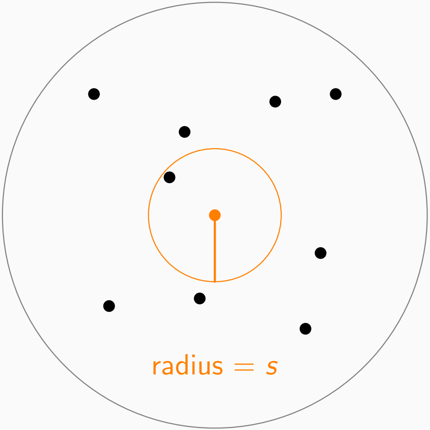

[Arnott, 1996]:

$$
\begin{aligned}
&\mathbb{P}(\text{pickup distance} \leq s) \\
&= 1 - \mathbb{P}(\text{no car within } R) \\
&= 1 - \exp(-\pi s^2 \gamma)
\end{aligned}
$$

Relationship between Pickup Time and Idle Time

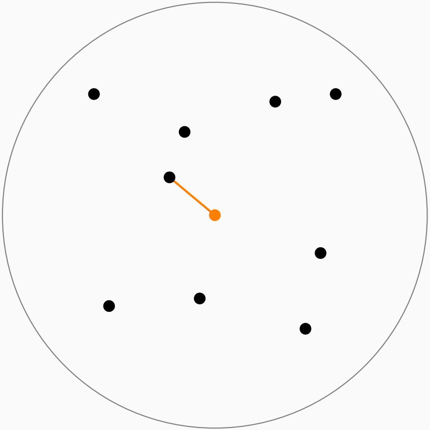

[Arnott, 1996]:

$$
\begin{aligned}
&\text{average pickup time} \\
&= \int_0^\infty \mathbb{P}(\text{distance} > s)\,\mathrm{d}s \\
&= \int_0^\infty \exp(-\pi s^2 \gamma)\,\mathrm{d}s = \frac{1}{2\sqrt{\gamma}}
\end{aligned}
$$

Our Objective

1. Formulate a **stochastic model** of the ride-sharing system that captures the trade-off between pickup time and idle time.

Our Objective

1. Formulate a **stochastic model** of the ride-sharing system that captures the trade-off between pickup time and idle time.
2. Formulate a **tractable** fluid-based **optimization problem** that approximates the stochastic model at steady-state.

Our Objective

1. Formulate a **stochastic model** of the ride-sharing system that captures the trade-off between pickup time and idle time.
2. Formulate a **tractable** fluid-based **optimization problem** that approximates the stochastic model at steady-state.
3. Derive static / state-dependent **policies** for the system.

Our Objective

1. Formulate a **stochastic model** of the ride-sharing system that captures the trade-off between pickup time and idle time.
2. Formulate a **tractable** fluid-based **optimization problem** that approximates the stochastic model at steady-state.
3. Derive static / state-dependent **policies** for the system.
4. **Compare** performance of the policies with existing results, **theoretically** and **numerically**.

Why should we care about it?

What can go wrong if we ignore pickup time?

- In most of the recent papers (e.g. [Bimpikis et al., 2019], [Braverman et al., 2019]), pickup time is assumed to be 0.

What can go wrong if we ignore pickup time?

- In most of the recent papers (e.g. [Bimpikis et al., 2019], [Braverman et al., 2019]), pickup time is assumed to be 0.
- That is, if there is any available car, the passenger will start the ride immediately upon submitting a request.

What can go wrong if we ignore pickup time?

- In most of the recent papers (e.g. [Bimpikis et al., 2019], [Braverman et al., 2019]), pickup time is assumed to be 0.
- That is, if there is any available car, the passenger will start the ride immediately upon submitting a request.
- However, we actually need to keep enough idle cars in order to have a reasonable pickup time. The pickup time and idle time cannot be both very small.

What can go wrong if we ignore pickup time?

- In most of the recent papers (e.g. [Bimpikis et al., 2019], [Braverman et al., 2019]), pickup time is assumed to be 0.
- That is, if there is any available car, the passenger will start the ride immediately upon submitting a request.
- However, we actually need to keep enough idle cars in order to have a reasonable pickup time. The pickup time and idle time cannot be both very small.
- Thus, policies we get from these models may perform badly in a real system.

Key Results

Agenda:

- build a fluid-based optimization problem with pickup time.
- build a fluid-based optimization problem without pickup time (similar to the model from [Braverman et al., 2019]).
- use both optimal solutions to build static policies
- compare performance of the two static policies.

Key Results - Numerical

Model with pickup time:

- optimal revenue from optimization model: 13.82
- true revenue from simulated MDP: 13.72

Model without pickup time:

- optimal revenue from optimization model: 20.02
- true revenue from simulated MDP: 7.76

Key Results - Numerical

Model with pickup time (1000 cars in 5 zones):

- (steady state) number of idle cars in the optimal solution of optimization model: 81.2, 63.1, 69.5, 74.6, 24.4.
- true number of idle cars in 5 zones: see below
- (y-axis: number of cars; x-axis: time)

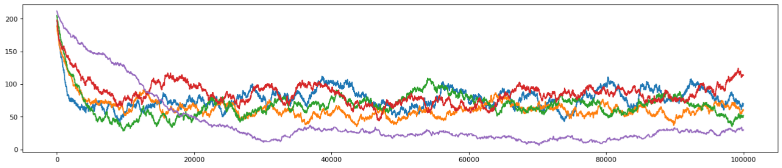

Key Results - Numerical

Model without pickup time (1000 cars in 5 zones):

- (steady state) number of idle cars in the optimal solution of optimization problem: 0, 0, 0, 0, 0.
- true number of idle cars in 5 zones: see below
- (y-axis: number of cars; x-axis: time)

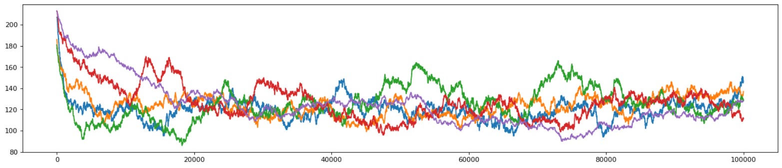

How can we model pickup time tractably?

Discretization of Pickup Time

We found it hard to model the average pickup time in a tractable way

$$\bar{t} = \frac{1}{2\sqrt{\gamma}} = \frac{1}{2\sqrt{a_i/\theta_i}}$$

- a_i: # of idle cars at zone i
- θ_i: area of zone i

Discretization of Pickup Time

We found it hard to model the average pickup time in a tractable way

$$\bar{t} = \frac{1}{2\sqrt{\gamma}} = \frac{1}{2\sqrt{a_i/\theta_i}}$$

- a_i: # of idle cars at zone i
- θ_i: area of zone i

Instead, it is easier to work with this probability that a car is within distance δ

$$\mathbb{P}(\text{dist} \leq \delta) = 1 - \exp(-\frac{\pi\delta^2 a_i}{\theta_i})$$

Discretization of Pickup Time

Discretize of pickup time:

- We choose a sequence $0 = \delta_0 < \delta_1 < \cdots < \delta_{M_i}$ of distances.
- Suppose that distance to the closest car is in $(\delta_{m-1}, \delta_m]$
- Then pickup time is assumed to be exponentially distributed with mean $1/\nu_m$

Instead, it is easier to work with this probability that a car is within distance δ

$$\mathbb{P}(\text{dist} \leq \delta) = 1 - \exp\left(-\frac{\pi\delta^2 a_i}{\theta_i}\right)$$

Discretization of Pickup Time

Discretize of pickup time:

- We choose a sequence $0 = \delta_0 < \delta_1 < \cdots < \delta_{M_i}$ of distances.
- Suppose that distance to the closest car is in $(\delta_{m-1}, \delta_m]$
- Then pickup time is assumed to be exponentially distributed with mean $1/\nu_m$

Probability that pickup time has mean $1/\nu_m$:

$$Q_{im}(a_i) = \exp(-\frac{\pi(\delta_{m-1})^2 a_i}{\theta_i}) - \exp(-\frac{\pi(\delta_m)^2 a_i}{\theta_i})$$

Discretization of Pickup Time

Intuition

- Customers are split into multiple "streams", whose fraction depend on a_i.
- Average pickup time depends on a_i.
- Customers with mean pickup time longer than $1/\nu_{M_i}$ are lost.

Probability that pickup time has mean $1/\nu_m$:

$$Q_{im}(a_i) = \exp\left(-\frac{\pi(\delta_{m-1})^2 a_i}{\theta_i}\right) - \exp\left(-\frac{\pi(\delta_m)^2 a_i}{\theta_i}\right)$$

Convert Transformation of the Optimization Problem

Consider the following constraints:

$$q_{im} = \exp(-\frac{\pi(\delta_{m-1})^2 a_i}{\theta_i}) - \exp(-\frac{\pi(\delta_m)^2 a_i}{\theta_i}) \quad , \quad m = 1, \cdots, M_i$$

$$\iff \quad \sum_{m'=1}^{m} q_{im'} = 1 - \exp(-\frac{\pi(\delta_m)^2 a_i}{\theta_i}) \quad , \quad m = 1, \cdots, M_i$$

Convex Transformation of the Optimization Problem

Consider the following constraints:

$$q_{im} = \exp(-\frac{\pi(\delta_{m-1})^2 a_i}{\theta_i}) - \exp(-\frac{\pi(\delta_m)^2 a_i}{\theta_i}) \quad , \quad m = 1, \cdots, M_i$$

$$\iff \quad \sum_{m'=1}^{m} q_{im'} = 1 - \exp(-\frac{\pi(\delta_m)^2 a_i}{\theta_i}) \quad , \quad m = 1, \cdots, M_i$$

and the convex relaxation:

$$\sum_{m'=1}^{m} q_{im'} \leq 1 - \exp(-\frac{\pi(\delta_m)^2 a_i}{\theta_i}) \quad , \quad m = 1, \cdots, M_i$$

Convex Transformation of the Optimization Problem

Consider the following constraints:

$$q_{im} = \exp(-\frac{\pi(\delta_{m-1})^2 a_i}{\theta_i}) - \exp(-\frac{\pi(\delta_m)^2 a_i}{\theta_i}) \quad , \quad m = 1, \cdots, M_i$$

$$\Longleftrightarrow \quad \sum_{m'=1}^{m} q_{im'} = 1 - \exp(-\frac{\pi(\delta_m)^2 a_i}{\theta_i}) \quad , \quad m = 1, \cdots, M_i$$

and the convex relaxation:

$$\sum_{m'=1}^{m} q_{im'} \leq 1 - \exp(-\frac{\pi(\delta_m)^2 a_i}{\theta_i}) \quad , \quad m = 1, \cdots, M_i$$

We can show that in the optimization problem, these constraints always hold tight at optimality.

Some topics that are not covered:

Tractability

- We use a choice model to capture relationship between price and demand. That is also part of the convex transformation.
- The feasible region of the resulting convex program is not closed. We show that it can be solved by solving a conic program.

Approximation Guarantees

State-dependent Policies

......

Some Related Works

Arnott, R. (1996).
Taxi Travel Should be Subsidized.
Journal of Urban Economics, 40(3):316–333.

Bimpikis, K., Candogan, O., and Saban, D. (2019).
Spatial Pricing in Ride-Sharing Networks.
Operations Research, 67(3):744–769.

Braverman, A., Dai, J. G., Liu, X., and Ying, L. (2019).
Empty-car Routing in Ridesharing Systems.
Operations Research, 67(5):1437–1452.