2021 IEEE Intelligent Transportation Systems Conference (ITSC)
Indianapolis, USA. September 19-21, 2021

Real-Time Control of Mixed Fleets in Mobility-on-Demand Systems

Kaidi Yang[1], Matthew W. Tsao[1], Xin Xu[2], and Marco Pavone[1]

Abstract— **Automated vehicles (AVs) are expected to be beneficial for Mobility-on-Demand (MoD), thanks to their ability of being globally coordinated. To facilitate the steady transition towards full autonomy, we consider the transition period of AV deployment, whereby an MoD system operates a mixed fleet of AVs and human-driven vehicles (HVs). In such systems, AVs are centrally coordinated by the operator, and the HVs might strategically respond to the coordination of AVs. We devise computationally tractable strategies to coordinate mixed fleets in MoD systems. Specifically, we model an MoD system with a mixed fleet using a Stackelberg framework where the MoD operator serves as the leader and HVs serve as the followers. We further develop a real-time coordination algorithm for AVs. The proposed approach is validated using a case study inspired by real operational data of an MoD service in Singapore. Results show that the proposed approach can significantly improve system performance.**

I. INTRODUCTION

The past decade has witnessed the widespread deployment of Mobility-on-Demand (MoD) services, thanks to the rapid adoption of smartphones, developments in wireless communication, and the boom of shared economies. These services, e.g., the ride-hailing services provided by Uber, present immense potential to enhance mobility and accessibility while reducing resource usage. One key operational challenge associated with these services is represented by the *vehicle imbalances* due to asymmetric transportation demand: vehicles tend to accumulate in some regions while becoming depleted in others, giving rise to inefficient operations of the MoD system. Currently, MoD systems typically address this challenge by combining dynamic pricing [1] with a real-time heat-map of the passenger demand to rebalance their fleets. However, the rebalancing actions are still performed in a decentralized manner by human drivers who are interested in their own earnings, which may not yield optimal system performance.

The emergence of automated vehicles (AVs) provides opportunities for sophisticated and centralized vehicle control, and thus can be beneficial to MoD systems. By integrating AVs into MoD systems, Autonomous Mobility-on-Demand (AMoD) is expected to be a promising paradigm for future mobility systems, whereby a fleet of autonomous taxi-like vehicles is coordinated by a central operator to provide on-demand mobility services. Compared to traditional MoD services, AMoD offers several advantages. First, by eliminating the driver costs, AMoD can reduce the cost of trips and thus improve the operator's profit. Second, since AVs can always be operational in the system, AMoD can provide continuous and reliable services regardless of the time of day. Third, unlike human drivers who can be ill-informed or self-interested, AVs can be coordinated in a centralized manner to provide better services which enable higher vehicle utilization and operational efficiency. Thanks to these advantages, AMoD has attracted increasing attention in the research fields of transportation and robotics, including demand analysis [2], real-time coordination [3], interactions with public transport [4] and power networks [5], transportation network design [6], and system evaluation [7].

Despite the benefits of AMoD systems, it is evident that AVs will only gradually be technologically mature and adopted in the market. During the transition period, MoD systems will conceivably be operating a mixed fleet of AVs and human-driven vehicles (HVs), whereby HVs might respond to the coordination of AVs strategically, making global optimization challenging. To facilitate the steady transition towards full autonomy, this paper aims to devise computationally tractable strategies to design and coordinate mixed fleets in MoD systems, considering the interactions between AVs and HVs. Specifically, we frame a fleet coordination problem with a mixed equilibrium of AVs that are centrally coordinated and HVs that act according to their own interests. From an operational perspective, such a system could also be interpreted as a mixed system of compliant drivers (e.g., contractor drivers who are paid to strictly follow the instructions given by the operator) and self-interested drivers – thus, the tools and insights derived in this paper could be applied to existing systems (i.e., without AVs) as well.

Related work. To the best of our knowledge, such a mixed fleet system has been rarely studied in the context of MoD services [8]–[10]. Lokhandwala et al. [8] analyzed the ride-sharing serviced provided by autonomous taxis and human-driven taxis based on an agent-based simulation of New York City. Afeche et al. [9] focused on a two-location, four-route loss network and investigated the impact of demand-side admission and supply-side rebalancing control on the spatial vehicle imbalances and the strategical behavior of drivers. Wei et al. [10] explicitly considered the interactions between AVs and HVs and analyzed the steady-state behavior of the mixed fleet system in a transportation network with equidistant nodes. Both [9] and [10] focused on the steady-state analysis of special types of transportation networks. To sum up, it remains unclear how mixed fleet systems with realistic road networks can be controlled, especially in real time.

[1] K. Yang, M. Tsao, and M. Pavone were with the Department of Aeronautics and Astronautics, Stanford University, USA {kaidi.yang, mwtsao, pavone}@stanford.edu
[2] X. Xu was with Grab Holding Inc., Singapore xin.xu@grabtaxi.com
The authors would like to thank S. Hughes and M. Yamato from the Toyota Research Institute and G. Lin, P. Kotamarthi, A. Madelil, J. Zhou, B. Adipradana, Z. K. Wong, and V. Torka from Grab Holding Inc. for sharing the data used for the case study and for helpful discussions. We also would like thank M. Zallio for help with the graphics. This research was supported by the National Science Foundation (NSF) under CAREER Award CMMI1454737 and the Toyota Research Institute (TRI). The first author would like to acknowledge the support of the Swiss National Science Foundation (SNSF) Postdoc.Mobility Fellowship (P2EZP2_184244). This article solely reflects the opinions and conclusions of its authors and not NSF, TRI, SNSF, or any other entity.

2021 IEEE International Intelligent Transportation Systems Conference (ITSC) | 978-1-7281-9142-3/21/$31.00 ©2021 IEEE | DOI: 10.1109/ITSC48978.2021.9564770

Several works analyze the behavior of HVs in the MoD context. Most works focus on the route choice of HVs delivering passengers without considering the rebalancing behavior or the willingness of HVs to accept passengers (e.g., [11]), or analyze the response of HVs to incentives such as surge or dynamic pricing at a single or at a few locations (e.g., [12]). Buchholz [13], on the other hand, empirically analyzed the dynamic spatial equilibrium of taxicabs based on the New York City taxi data. Bimpikis et al. [14] studied the rebalancing behavior of HVs by formulating the equilibria as the solutions to a set of non-linear equations. However, it is assumed that HVs always accept the passengers assigned to them by the operator, which may not be the case in real-world systems. Moreover, it is challenging to leverage the model proposed in [14] for real-time control due to its non-linear structure.

Related works on such mixed fleet systems also exist in the context of traffic assignment (i.e., without dispatch or rebalancing), which typically search for the optimal equilibrium in a Stackelberg game where the leader is the group of compliant drivers, and the followers are the self-interested drivers (e.g., [15], [16]). These works, however, focus on congestion games between the two types of vehicles. It is nevertheless unclear how the proposed algorithms in these works can be adapted to the MoD services where rebalancing and passenger assignment are important features.

Statement of contribution. The contribution of this paper is three-fold. First, we initiate research on mixed fleet systems with realistic road networks in an MoD context. We account for the interactions between AVs and HVs using a Stackelberg framework where the MoD operator with its AV fleet serves as the leader and HVs serve as the followers. Second, we develop a Model Predictive Control (MPC)-based approach for the real-time coordination of such mixed fleet systems. Third, we conduct real-world case studies using real data in Singapore to validate the proposed algorithms and provide a guideline to deploy AVs in scenarios with various AV penetration rates.

This paper is organized as follows. Section II presents an overview of MoD systems with a mixed fleet and introduces general notions. Section III develops a Stackelberg game-based MPC formulation for real-time control. Section IV presents simulation results to illustrate the benefits of employing AVs in the system. Section V concludes the paper and proposes future directions.

II. System Description

We consider a city where an operator provides MoD services with a mixed fleet of AVs (denoted as a) and HVs (denoted as h), illustrated in Figure 1. Mathematically, we describe the urban transportation network as a weighted graph $\mathcal{G} = (\mathcal{N}, \mathcal{E})$, where $\mathcal{N}$ is the set of stations (i.e., pick-up or drop-off locations) and $\mathcal{E}$ is a set of directed edges, i.e., shortest paths between pairs of stations. For a typical city, we consider $\mathcal{G}$ to be fully connected such that a directed edge exists for each pair of stations. Let $\mathcal{N}_i = \mathcal{N} \backslash \{i\}$ be the set of stations connected to station i. We discretize the time horizon into a set of discrete intervals $\mathcal{T} = \{1, 2, \cdots, T\}$ of a given length ΔT. Without loss of generality, we assume both types of vehicles can operate on each station of the network. The proposed methodological framework can be readily adapted

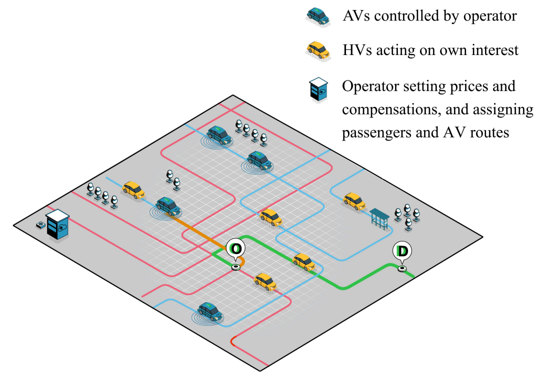

Fig. 1: Illustration of the MoD system, where blue vehicles represent automated vehicles (AVs) and yellow vehicles represent human-driven vehicles (HVs). An operator is able to set prices for request and compensations for HVs, assign passengers, and coordinate the routes of AVs.

to account for the limited driving capacity of AVs during the transition period, whereby AVs can only operate within certain regions or take certain paths.

The travel time for edge $(i, j) \in \mathcal{E}$ is defined as the number of time steps it takes a vehicle to travel along the shortest path between station i and station j, denoted as an integer $\tau_{ij} \in \mathbb{Z}_+$. We make the following remarks for the travel times. First, we assume that the travel times are given and independent of the coordination of the MoD fleet. This assumption applies to cities where the MoD fleet constitutes a relatively small proportion of the entire vehicle population on the transportation network, and thus the impact of the MoD fleet on traffic is marginal. Second, we observe that the travel times typically change slowly with time, and hence for simplicity of notation, we assume τ_{ij} to be constant in time. The proposed modeling framework can be easily adapted to consider time-varying exogenous travel times. Third, we assume the travel times to be independent of the vehicle classes, since vehicles of both classes have similar speeds and will follow the same shortest paths between two stations. The modeling framework can also be easily adapted to incorporate class-dependent travel times. Similar to travel times, we also define the travel distance for edge $(i, j) \in \mathcal{E}$ as the length of the shortest path between station i and station j, denoted as δ_{ij}.

Passengers make transportation requests at each time step. We denote the origin-destination (OD) pairs as tuples (i, j), $i, j \in \mathcal{N}$, and the demand and price for OD pair (i, j) starting at time step $t \in \mathcal{T}$ as $q_{ijt} \geq 0$ and $p_{ijt} \geq 0$, respectively. Clearly, passengers that depart at time step t arrive at their destinations at time step $t + \tau_{ij}$. We make four assumptions for passenger requests. First, we assume the prices are given endogenously, since they are typically derived in many existing MoD systems as the product of a fixed base price and a surge factor defined in another module. Second, we assume that trip prices are the same for AVs and HVs, and we further assume that passengers do not have preference over vehicle classes. These assumptions apply to common scenarios where the goal of passengers is to get transportation services as soon as possible and would accept whichever class of vehicles assigned to them. The relaxation of these assumptions will be investigated in future work. Third, for simplicity of presentation, we consider the assumption as in the existing literature [10] that passengers can only be matched to vehicles located at the same station as their origins. Nevertheless, we can straightforwardly adapt the proposed models to relax this assumption (see Section III-D). Fourth, we assume that the operator can estimate the demand within a time horizon [17]. Discussions of specific estimation algorithms are beyond the scope of the paper.

With the passenger requests, the MoD operator determines its strategies to optimize system performance (e.g., its profits, system-level earnings, social welfare, etc.). Similar to existing ride-sharing systems, the operator dynamically assigns passengers to vehicles (both AVs and HVs). AVs will always accept the assigned passengers, whereas HVs can strategically choose whether to accept the assigned passengers to optimize their own interests. This setting applies to systems where the operator broadcasts the passenger requests to drivers (e.g., in many taxi systems) or systems where the operator makes recommendations, but HVs can reject the assigned requests with minimal penalties. The models proposed in this paper, nevertheless, can be tailored to analyze the scenarios where HVs always accept the assigned requests, or be extended to consider scenarios with heterogeneous HV fleets where different types of HVs (compliant or strategic in passenger assignment) co-exist. After a vehicle (AV or HV) is successfully matched with a request, the vehicle will pick up the matched passengers, and then deliver them to their destinations. HVs will receive a compensation c_{ijt} for serving a trip with OD pair (i, j) from the operator. We assume that compensations are given endogenously, since they are typically calculated in existing MoD systems as a fixed percentage of prices.

Vehicles not matched with any passengers may either stay at the same station or rebalance to other stations. The rebalancing decisions for AVs are made by the operator, whereas HVs make such decisions on their own. For both vehicle classes, there is an operational cost (e.g., the energy cost, vehicle depreciation, and vehicle maintenance) σ per vehicle per unit distance driving. To mathematically describe the system, let us define the passenger flow $x^m_{ijt} \geq 0$ and rebalancing flow $y^m_{ijt} \geq 0$ as the number of vehicles of class $m \in \mathcal{M}$ that start moving along edge (i, j) in time step t with and without passengers, respectively.

Passengers that are not matched can either stay in the system and enter the next round of assignment or leave. Let w_{ijt} be the number of passengers with OD pair $(i, j) \in \mathcal{E}$ waiting to be assigned at the beginning of time step $t \in \mathcal{T}$. Let $\epsilon \in [0, 1]$ be the per time step probability that unserved passengers choose to stay in the system, which characterizes the impatience of passengers. We assume that the probability ϵ is determined exogenously and does not depend on the real-time operation of the service (or equivalently, the coordination algorithms).

To summarize, the operator aims to optimize system performance by (1) assigning passengers to AVs and HVs and (2) designing routes for AVs. AVs will strictly follow instructions, whereas HVs will strategically determine (1) whether to accept passengers and (2) their routes. Following the sequential property of such systems, we model the operations of mixed fleet systems using a Stackelberg game framework, where the leader is the MoD operator, and the followers are the HVs, and devise an MPC-based approach to control such systems in real time. We will present the details of the proposed Stackelberg game-based MPC approach in the rest of the paper.

III. Real-Time Control of MoD Systems With Mixed Fleets

In this section, we present a Stackelberg game-based Model Predictive Control (MPC) approach to coordinate MoD systems with mixed fleets in real-time. The proposed MPC approach relies on an embedded optimization model based on a Stackelberg game, where the leader is the MoD operator which optimizes passenger and rebalancing routes for the AVs in order to improve system-level earnings, while the followers are the HVs which strategically respond to the operator's decisions. At each time step, the controller takes as an input the predicted passenger demand and the vehicle states (i.e., the number of idle vehicles or the vehicles en-route for passenger pickup or delivery), and solve an optimization problem to compute passenger and rebalancing routes for the AVs that maximize system-level earnings over a receding time horizon. Notice that HVs then make their decisions based on the assignment of the AVs accordingly. As is typical for MPC-style algorithms, only the passenger and rebalancing routes of AVs at the current time step are executed, and the process is repeated. This mechanism has the advantage of taking future system performance into account when optimizing current actions.

We characterize the rebalancing strategy of HVs as rebalancing probabilities P_{ijt} for HVs to move from station $i \in \mathcal{N}$ to station $j \in \mathcal{N}_i$ at time $t \in \mathcal{T}$. We consider such probabilities to be determined *externally* in the MPC formulation, i.e., they are predicted from experience (e.g., from historical data). The reasons for this modeling choice are as follows. First, we expect HVs not to have real-time global information about passenger demand at other stations or positions of other vehicles, and thus they can only slowly adapt their rebalancing strategy after experiencing difficulty in getting passengers at some stations. As the MPC considers a relatively short time horizon (e.g., 20 min), it appears reasonable to assume that HVs stick to the planned rebalancing strategies over such a horizon. Second, we expect the MPC algorithm, due to its repeated optimizations, to be robust to small errors in the prediction–robustness will be experimentally evaluated in Section IV-C. Third, from a computational standpoint, this choice significantly reduces the number of decision variables, and thus makes the algorithm much more scalable.

We now present the MPC formulation. Denote by r^m_{it} the number of vacant vehicles of class $m \in \mathcal{M}$ located at station $i \in \mathcal{N}$ at time step t. Let K be the length of the planning horizon, $\mathcal{T}_0 = \{t_0, t_0+1,\ t_0+K-1\}$ be the set of time steps in the planning horizon from time step t_0, and $\mathcal{T}_-$ be the set of time steps prior to time step t_0. We use bolded variables to represent vectors, i.e., $(\boldsymbol{q}_t, \boldsymbol{p}_t, \boldsymbol{c}_t, \boldsymbol{w}_t, \boldsymbol{x}^a_t, \boldsymbol{x}^h_t, \boldsymbol{y}^a_t, \boldsymbol{y}^h_t) = \{q_{ijt}, p_{ijt}, c_{ijt}, w_{ijt}, x^a_{ijt}, x^h_{ijt}, y^a_{ijt}, y^h_{ijt}\}_{(i,j)\in\mathcal{E}}$, $(\boldsymbol{r}^a_t, \boldsymbol{r}^h_t) = \{r^a_{it}, r^h_{it}\}_{i\in\mathcal{N}}$, $\forall t \in \mathcal{T}$, and $(\boldsymbol{w}, \boldsymbol{x}^a, \boldsymbol{x}^h, \boldsymbol{y}^a, \boldsymbol{y}^h, \boldsymbol{r}^a, \boldsymbol{r}^h) = \{\boldsymbol{w}_t, \boldsymbol{x}^a_t, \boldsymbol{x}^h_t, \boldsymbol{y}^a_t, \boldsymbol{y}^h_t, \boldsymbol{r}^a_t, \boldsymbol{r}^h_t\}_{t\in\mathcal{T}}$.

We next establish the leader model and follower model in the MPC formulation as Section III-A and Section III-B, respectively. We further propose a relaxation model to efficiently solve the MPC formulation in Section III-C. Section III-D refines the proposed MPC formulation to more explicitly consider the pickup process.

A. Leader model in the MPC formulation

Given a tuple $\{\boldsymbol{q}_t, \boldsymbol{p}_t\}_{t\in\mathcal{T}_0}$ of predicted demand and prices, as well as the initial conditions $(\boldsymbol{r}_0^a, \{\boldsymbol{r}_t^h, \boldsymbol{x}_t^h, \boldsymbol{x}_t^a, \boldsymbol{y}_t^a\}_{t\in\mathcal{T}_-})$ representing the number of idle vehicles and the vehicles en-route for passenger delivery and rebalancing, respectively, the operator determines the passenger and rebalancing flow of AVs $\{\boldsymbol{x}_t^a, \boldsymbol{y}_t^a\}_{t\in\mathcal{T}_0}$ by solving the optimization problem:

$$\max_{\substack{\boldsymbol{x}^a,\boldsymbol{x}^h,\boldsymbol{y}^a,\\ \boldsymbol{w},\boldsymbol{r}^a,\boldsymbol{r}^h}} \quad J_L^{\text{TV}} = \sum_{t\in\mathcal{T}_0}\sum_{(i,j)\in\mathcal{E}} \Big(p_{ijt}\sum_{m\in\mathcal{M}} x_{ijt}^m - \sigma\delta_{ij}\sum_{m\in\mathcal{M}} x_{ijt}^m - \sigma\delta_{ij}(y_{ijt}^a + P_{ijt}r_{it}^h) - \psi w_{ij}^t\Big) \tag{1a}$$

$$\text{s.t.} \quad r_{i,t+1}^a = r_{it}^a + \sum_{j\in\mathcal{N}_i}(y_{ji,t-\tau_{ji}}^a - y_{ijt}^a) + \sum_{j\in\mathcal{N}_i}(x_{ji,t-\tau_{ji}}^a - x_{ijt}^a),\ i\in\mathcal{N},\ t\in\mathcal{T}_0 \tag{1b}$$

$$r_{i,t+1}^h = r_{it}^h + \sum_{j\in\mathcal{N}_i}(P_{ji,t-\tau_{ji}}r_{j,t-\tau_{ji}}^h - P_{ijt}r_{it}^h) + \sum_{j\in\mathcal{N}_i}(x_{ji,t-\tau_{ji}}^h - x_{ijt}^h),\ i\in\mathcal{N},\ t\in\mathcal{T}_0 \tag{1c}$$

$$w_{ij,t+1} = \epsilon(w_{ijt} + q_{ijt} - x_{ijt}^a - x_{ijt}^h),\ (i,j)\in\mathcal{E}, t\in\mathcal{T}_0 \tag{1d}$$

$$\boldsymbol{x}^h = \Phi\Big(\boldsymbol{w}+\boldsymbol{q}-\boldsymbol{x}^a, \boldsymbol{r}^h\Big) \tag{1e}$$

$$\boldsymbol{x}^a, \boldsymbol{y}^a, \boldsymbol{w}, \boldsymbol{r}^a, \boldsymbol{r}^h \geq 0 \tag{1f}$$

where the objective function Eq.(1a) is the system-level earnings, which is defined as the difference between the earnings of the operator from both HVs and AVs (the first term) and costs, including the operational cost for the passenger routes of HVs and AVs (the second term), the operational cost for the rebalancing routes of HVs and AVs (the third term), and the cost associated with passengers waiting to be matched with a driver (the fourth term), where ψ represents a penalty for passenger waiting, which can be chosen, for example, as the VOT of passengers. Here, we consider system-level earnings because the operator can be interested in attracting human drivers to the system, and hence may not want to sacrifice the earnings of HVs. The objective function can also be adapted to consider other criteria (e.g., operator's profit). Constraints Eq.(1b) and Eq.(1c) represent the evolution of vehicle accumulation of AVs and HVs, respectively, at each station. Constraints Eq.(1d) represent the evolution of waiting passengers with respect to each origin-destination pair, where $\epsilon \in [0,1]$ represents the per time step probability that unserved passengers choose to stay in the system. Constraints Eq.(1e) model the behaviors of HVs, where the specific form of function $\Phi(\cdot)$ is specified by the follower model as detailed below in Section III-B. Constraints Eq.(1f) ensure that all variables are nonnegative.

B. Follower model in the MPC formulation

We next specify the follower model. Let us denote v as the system-wide average earnings rate of HVs, and ϕ_i as an estimate of expected earnings for HVs located at station $i\in\mathcal{N}$. Parameters v, $\{\phi_i\}_{i\in\mathcal{N}}$ can be learned from historical data. Then for any $t\in\mathcal{T}_0$, given parameters $\boldsymbol{\phi} = \{\phi_i\}_{i\in\mathcal{N}}$ (modeling expected earnings), the expected earning rate v, compensations $\boldsymbol{c}_t$, remaining demand $\bar{\boldsymbol{q}}_t = \boldsymbol{w}_t + \boldsymbol{q}_t - \boldsymbol{x}_t^a$, vehicle availability $\bar{r}_{it} = r_{it}^h + \sum_{j\in\mathcal{N}_i}(P_{ji,t-\tau_{ji}}r_{j,t-\tau_{ji}}^h - P_{ijt}r_{it}^h) + \sum_{j\in\mathcal{N}_i} x_{ji,t-\tau_{ji}}^h$, $\forall i\in\mathcal{N}$, we derive the passenger flow of HVs at time step t by solving the optimization problem:

$$\max_{\boldsymbol{x}_t^h} \quad J_{F,t}^{\text{TV}} = \sum_{(i,j)\in\mathcal{E}} (c_{ijt} - v\tau_{ij} - \sigma\delta_{ij} + \phi_j - \phi_i)x_{ijt}^h \tag{2a}$$

$$\text{s.t.} \quad 0 \leq x_{ijt}^h \leq \bar{q}_{ij},\ (i,j)\in\mathcal{E}, \tag{2b}$$

$$\sum_{j\in\mathcal{N}_i} x_{ijt}^h \leq \bar{r}_{it},\ i\in\mathcal{N}, \tag{2c}$$

where the objective function Eq.(2a) represents the earning gains for HVs, which is defined as the difference between the expected earnings by taking the trip (i.e., $c_{ijt} - v\tau_{ij} - \sigma\delta_{ij} + \phi_j$) and the expected earning of staying at the current station (i.e., ϕ_i). Constraints Eq.(2b) ensure that the passenger flow is always nonnegative, and reflect the fact that HVs might decline a subset of assigned passenger requests. Constraints (2c) ensure that the resulting passenger flow does not violate the availability of HVs. For details of the modeling of HVs, please refer to an extended version of this paper [18].

C. Relaxation to the MPC formulation

Since the follower model as stated in Problem (2) is a linear programming problem, the leader and follower models can be combined to yield a mixed integer linear programming (MILP). However, this approach is computationally expensive, especially for a large-scale city-level transportation network. To improve scalability, we propose the following relaxation to the MPC model:

$$\max_{\substack{\boldsymbol{x}^a,\boldsymbol{x}^h,\boldsymbol{y}^a,\\ \boldsymbol{w},\boldsymbol{r}^a,\boldsymbol{r}^h}} \quad J_L^{\text{TV}}(\boldsymbol{x}^a, \boldsymbol{x}^h, \boldsymbol{y}^a, \boldsymbol{w}, \boldsymbol{r}^a, \boldsymbol{r}^h) + \lambda\sum_{t\in\mathcal{T}_0} J_{F,t}^{\text{TV}}(\boldsymbol{x}_t^h) \tag{3a}$$

$$\text{s.t.} \quad \text{Eq.(1b)} - \text{Eq.(1f)}, \text{Eq.(2b)} - \text{Eq(2c)}.$$

In Problem (3), the objective function Eq.(3a) combines the objective function of the leader model as stated in Problem (1) and the objective function of the follower model as stated in Problem (2) via a weight parameter $\lambda \geq 0$. Such a parameter can be determined by the operator by using a sensitivity analysis. We will indeed analyze the sensitivity to parameter λ in Section IV. We also performed numerical experiments to evaluate the quality of the relaxation on small-scale case studies (e.g., with 8 stations) with realistic problem data (e.g., in terms of system-level earnings, passenger acceptance rate, etc.). Results showed that the relaxation is on average able to find a solution with an optimality gap within 2%–this motivates the use of Problem (3) in our large-scale case studies.

D. Extension to general pick-up locations

In this subsection, we extend the above framework to account for this possibility that vehicles can be matched with passengers located in other regions. We represent the passenger routes $\xi = (s, o, d) \in \mathcal{P}$ as the hyper-arc that vehicles at station $s \in \mathcal{N}$ take to pick up passengers at station $o \in \mathcal{N}$, and deliver them to their destination $d \in \mathcal{N}$. With such a notation, we can extend Problem (3) to Problem (4) by replacing the variables for passenger flow. Problem (4) reads as follows:

$$
\begin{aligned}
\max_{\substack{\boldsymbol{x}^a,\boldsymbol{y}^a,\boldsymbol{x}^h\\ \boldsymbol{w},\boldsymbol{r}^a,\boldsymbol{r}^h}} J_L^{\mathrm{TV}} = & \sum_{t\in\mathcal{T}_0}\sum_{\xi=(s,o,d)\in\mathcal{P}} p_{odt}(x_{\xi t}^h + x_{\xi t}^a) \\
& - \psi \sum_{t\in\mathcal{T}_0}\sum_{\xi=(s,o,d)\in\mathcal{P}} \tau_{so}(x_{\xi t}^h + x_{\xi t}^a) - \psi \sum_{t\in\mathcal{T}_0}\sum_{(i,j)\in\mathcal{E}} w_{ij}^t \\
& - \sigma \sum_{t\in\mathcal{T}_0}\sum_{\xi=(s,o,d)\in\mathcal{P}} (\delta_{so} + \delta_{od})(x_{\xi t}^h + x_{\xi t}^a) \\
& - \sigma \sum_{t\in\mathcal{T}_0}\sum_{(i,j)\in\mathcal{E}} \delta_{ij}(y_{ijt}^a + P_{ijt} r_{i,t}^h) \\
& + \lambda \sum_{t\in\mathcal{T}_0}\sum_{\xi=(s,o,d)\in\mathcal{P}} (c_{odt} + \phi_d - \upsilon\tau_{so} - \upsilon\tau_{od} \\
& \qquad - \sigma\delta_{so} - \sigma\delta_{od} - \phi_s) x_{\xi t}^h \qquad \text{(4a)}
\end{aligned}
$$

$$
\begin{aligned}
\text{s.t. } r_{i,t+1}^a = & \, r_{it}^a + \sum_{\xi=(s,o,i)\in\mathcal{P}} (x_{\xi,t-\tau_{so}-\tau_{oi}}^a - x_{\xi t}^a) \\
& + \sum_{j\in\mathcal{N}_i} (y_{ji,t-\tau_{ji}}^a - y_{ijt}^a), \ i \in \mathcal{N}, \ t \in \mathcal{T}_0 \qquad \text{(4b)}
\end{aligned}
$$

$$
\begin{aligned}
r_{i,t+1}^h = & \, r_{it}^h + \sum_{j\in\mathcal{N}_i} \left(P_{ji,t-\tau_{ji}} r_{j,t-\tau_{ji}}^h - P_{ijt} r_{it}^h\right) \\
& + \sum_{\xi=(s,o,i)\in\mathcal{P}} (x_{\xi,t-\tau_{so}-\tau_{oi}}^h - x_{\xi t}^h), \ i \in \mathcal{N}, \ t \in \mathcal{T}_0 \qquad \text{(4c)}
\end{aligned}
$$

$$
\begin{aligned}
w_{ij,t+1} = & \, \epsilon\Big(w_{ijt} + q_{ijt} - \sum_{\xi=(s,i,j)\in\mathcal{P}} (x_{\xi t}^a + x_{\xi t}^h)\Big), \\
& (i,j) \in \mathcal{E}, \ t \in \mathcal{T}_0 \qquad \text{(4d)}
\end{aligned}
$$

$$
\boldsymbol{x}^a, \boldsymbol{y}^a, \boldsymbol{x}^h, \boldsymbol{w}, \boldsymbol{r}^a, \boldsymbol{r}^h \geq 0 \qquad \text{(4e)}
$$

where the objective function Eq.(4a) is to maximize system-level earnings, defined as the difference between the total prices of the requests (the first term) and costs, including the passenger waiting cost associated with the pickup process (the second term), the cost for passengers waiting to be matched with a driver (the third term), the operational costs for passenger routes (the fourth term), the operational costs for rebalancing routes (the fifth term), and a penalty term (the sixth term) that characterizes the objective function of the follower problem (Problem (2)) to account for the self-interested behavior of HVs. Notice that the differences between Eq.(4a) and Eq.(1a) are 1) we explicitly penalize the passenger pickup time, and 2) we allow positive passenger pickup times within the same region (i.e., $\tau_{ii} > 0, \ i \in \mathcal{N}$). Constraints Eq.(4b) and Eq.(4c) represent the evolution of AVs and HVs, respectively, and Constraints Eq.(4d) describes the evolution of passengers. Eq.(4e) ensures that all variables are nonnegative.

To summarize this section, we have presented a Stackelberg game-based MPC approach to control the MoD systems with mixed fleet in real-time. The leader problem solved by the operator is formulated as Problem (1) to optimize the system-level earnings, and the follower problem approximating the behavior of HVs is formulated as Problem (2). The Stackelberg game based MPC approach is approximated as Problem (3) to improve scalability, and is extended to consider general pick-up locations as Problem (4).

IV. Case Study and Numerical Evaluation

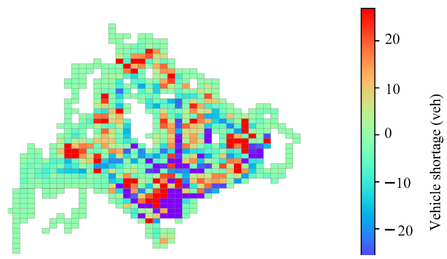

Fig. 2: Vehicle imbalances in Singapore as a function of locations, where the red blocks represent the locations with vehicle shortage and the blue blocks represent the locations with vehicle surplus.

A. Experiment design and data

We perform a case study inspired by a typical weekday morning peak between 8 a.m. and 10 a.m. in Singapore, where the road network is discretized into 126 blocks, each with an area of 2.4 km×2.4 km. The Singapore dataset is provided by Grab Holding Inc., which includes the passenger demand, the average travel times between OD station pairs, the pick up times, compensations, and fares. Figure 2 shows the studied area and the vehicle imbalances in the road network in the scenario where all vehicles are human-driven. Here, vehicle imbalances are characterized as the difference between the number of vehicles needed and the vehicle supply. It is clear that vehicles are oversupplied in some of the regions but insufficient in others. The VOT of passengers is assumed to be the average salary of Singapore, i.e., 0.45 SGD/min [19]. The operational cost of vehicles is chosen to be 0.15 SGD/km (estimated from [20]). The expected earning rate of HVs is calibrated to be $\upsilon = 0.25$ SGD/min. The total fleet size is 14,000 vehicles, which is assumed to be fixed and time-invariant throughout the studied time period. We vary the share of AVs in the mixed fleet to be between 0 and 14,000 vehicles. The length of a time step is chosen as 1 min. At the beginning of each time step, passengers make requests, and idle vehicles located at the same or adjacent blocks can be assigned to the requests. The planning horizon for the proposed MPC is chosen as 15 min as a trade-off between computational efficiency, prediction accuracy, and system performance. The weight parameter λ is set via sensitivity analysis equal to 6. We further assume that passengers will leave the system if they are not served within one time step, i.e., $\epsilon = 0$.

We first evaluate a scenario where the proposed MPC approach has perfect information on demand, travel times, and the rebalancing probabilities (Section IV-B), and then test the robustness of the MPC approach against the prediction errors in these variables (Section IV-C and Section IV-D). Simulations are run in Python on a desktop computer with AMD Ryzen 5 3600XT 6-Core Processor and 32 GB memory. The optimization models are solved using CPLEX. The average computation time for each time step is around 5 seconds.

B. Performance of the proposed Stackelberg game-based MPC approach

We evaluate the proposed Stackelberg game-based MPC approach by comparing the following two approaches.

- Baseline approach, defined as an MPC approach that coordinates AVs assuming fully compliant HVs, i.e., without accounting for the strategic interactions with the HVs (Problem (4) with $\lambda = 0$).
- Stackelberg game-based MPC. The control actions are obtained by solving Problem (4) with the $\lambda = 6$ that provides the optimal system-level earnings in the sensitivity analysis.

These approaches are evaluated in scenarios with different fleet sizes of AVs to illustrate the value of introducing AVs in the MoD system. The comparison between these two approaches sheds light on the value of considering the interactions between AVs and HVs. We evaluate the system-level earnings in general, but also look into more detailed performance criteria, e.g., operator's profit, average passenger waiting time, passenger acceptance rate, vehicle utilization, and the empty vehicle kilometers traveled. Here, since we assume passengers do not wait to be assigned, the passenger waiting time refers specifically to the time that a passenger waits to be picked up by the matched vehicle. The passenger acceptance rate refers to the percentage of passengers that are successfully matched with a vehicle, the vehicle utilization represents the percentage of time a vehicle is occupied by a passenger, and the empty vehicle kilometers traveled are defined as the total distance traveled by vehicles to pick up passengers or rebalance themselves. The results are summarized in Figure 3.

a) Value of introducing AVs: We show the value of introducing AVs by comparing the system performance in scenarios with various penetration rates of AVs. We can see from Figure 3 that the MoD system can be improved significantly by replacing HVs with AVs by using both approaches. Specifically, we can improve the system-level earnings by up to 32%, reduce the passenger waiting time by up to 30%, increase the passenger acceptance rate by up to 20%, improve the vehicle utilization by up to 21%, and reduce the empty kilometers traveled by up to 67%. Moreover, the benefits of AVs is more significant at the early stage of deployment. In fact, the system-level earnings can be improved by 13% by replacing 10% of vehicles with AVs. This shows that promising value of AVs in improving MoD services.

b) Value of considering the interactions between AVs and HVs: We show the value of considering the interactions between AVs and HVs by comparing the baseline approach (blue solid lines) with the Stackelberg game-based MPC approach (green dashed lines). We can see from Figure 3 that the proposed Stackelberg game-based MPC can improve system-level earnings by up to 12% (Figure 3a), operator's profit by up to 400% (Figure 3b), passenger acceptance rate by up to 12% (Figure 3c)., and vehicle utilization by up to 20% (Figure 3f) by considering the interactions between HVs and AVs. The reason is two-fold. First, the Stackelberg game-based MPC can predict the behaviors of HVs more accurately, and can match AVs to the passengers that HVs are not willing to serve. Second, by modeling the system more accurately, the Stackelberg game-based MPC is able to coordinate AVs in a more efficient manner, and thus serve more passengers. We further notice that the benefit of considering such interactions is especially evident in scenarios with moderate penetration rates of AVs (i.e., 20% – 60%). This is because the interactions between AVs and HVs are more intense in these scenarios. The passenger waiting time and empty kilometers traveled, however, may not be necessarily improved (Figure3d and Figure 3f). This is expected because the Stackelberg game-based MPC can coordinate AVs to take the passengers with relatively low values and long pickup times that would otherwise not be taken by the HVs. Accepting these passengers may increase the average passenger waiting time and the total empty kilometers traveled. One can increase the weight parameter ψ penalizing passenger waiting times, if passenger waiting times are especially important for the operator.

C. Robustness to prediction errors

The proposed MPC approach relies on the prediction of several variables: the future travel demand, the rebalancing flow of HVs, and the travel times. We analyze how the prediction errors in these variables would affect the performance of the proposed MPC approach. To this end, we allow these variables to be stochastic in the simulation, but use the mean values in the proposed MPC to calculate the actions for AVs. Specifically, the travel demand is assumed to follow a Poisson distribution. The number of rebalancing HVs from each station is assumed to follow a multinomial distribution with the given rebalancing probabilities as the parameters. The prediction error in the travel times is assumed to follow a Gaussian distribution with mean 0 and a standard deviation of 10% and 20% of the value, respectively, to simulate scenarios with moderate and large noises in travel times. Notice that the MPC can obtain the true travel demand at the current time step from passenger requests, but uses the mean value of the Poisson distribution to predict the future demand. We derive the resulting system-level earnings in scenarios with different levels of prediction errors and different fleet sizes of AVs, each with 5 random seeds. The results are as shown in Figure 4.

Figure 4 shows that the proposed Stackelberg game-based approach is quite robust to these prediction errors. This is expected for three reasons. First, we expect MPC to be robust against a certain level of prediction errors, due to its built-in feedback mechanism. Second, transportation requests at the current time step are submitted by the passengers and are accessible to the operator. Hence, the controller has perfect knowledge of the demand at the current time step, which is used to effectively match passengers with drivers. Third, the knowledge of transportation demand and rebalancing

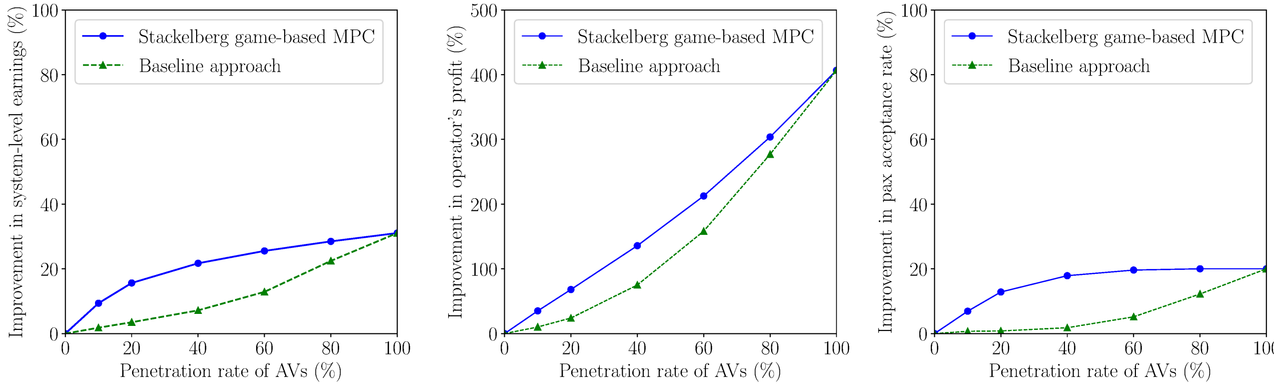

(a) Improvement in system-level earnings (b) Improvement in operator's profit (c) Improvement in pax acceptance rate

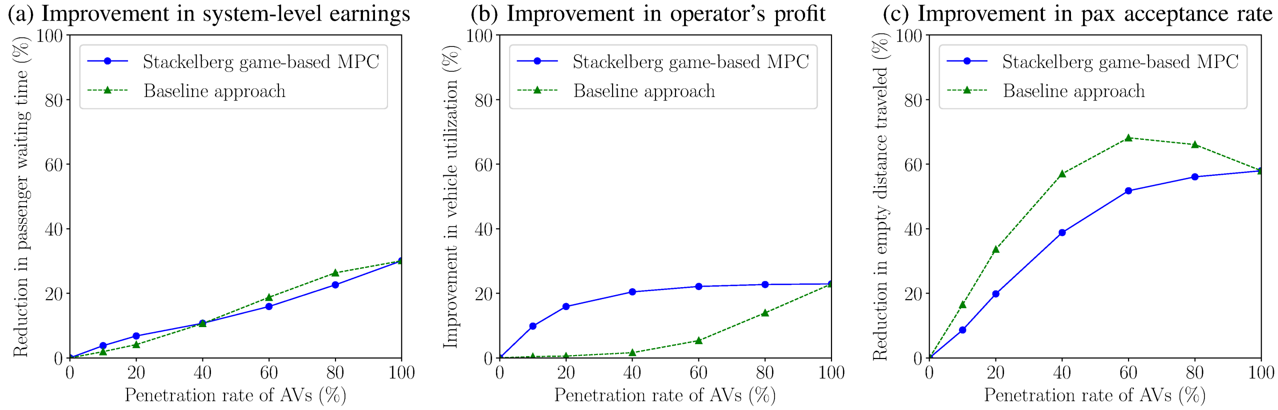

(d) Reduction in pax waiting time (e) Improvement in vehicle utilization (f) Reduction in empty distance traveled

Fig. 3: Performance of the Stackelberg game-based MPC approach (blue solid lines) compared to a baseline approach without considering the interactions between HVs and AVs (green dashed lines).

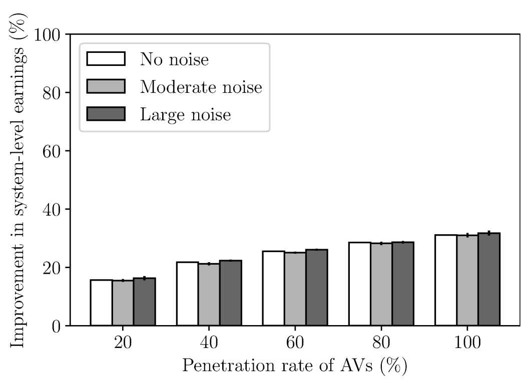

Fig. 4: Robustness to prediction noises, where the vertical axis represents the improvement in system-level earnings compared to the scenario with all HVs; moderate noises and large noises represent the scenarios with 10% and 20% errors in travel time prediction, respectively; and the error bars represent the standard deviation in the improvement across different random seeds.

probabilities enables MPC to frequently rebalance AVs to blocks with vehicle shortages.

D. Sensitivity to model parameters

We analyze the sensitivity of the proposed MPC to the model parameters, i.e., the assumed VOT of HV drivers μ and the weight parameter λ in Problem (4), to evaluate the impact on the performance of the proposed MPC if these parameters deviate from the values that yield the optimal system-level earnings. To this end, we employ a local sensitivity analysis method, i.e., the One-at-a-Time method by varying one parameter in each simulation while keeping the other parameter fixed. Such a local sensitivity analysis suffices for practical applications, because we do not expect the parameter values to deviate significantly from the optimal values. Specifically, we vary the assumed μ between 0.15 SGD/min to 0.35 SGD/min in Problem (4), while keeping $\mu = 0.25$ in the simulation. The value of λ varies between 0 and 6. The results are illustrated in Figure 5.

We can see that the performance of the proposed MPC is not sensitive to either parameter, as long as the VOT of HV drivers μ is within a reasonable range (i.e., between 0.2 SGD/min and 0.3 SGD/min) and the weight parameter λ is greater than 3. The performance of the proposed MPC would drop if the parameters are outside these ranges. If μ lies too far from the true value, the MPC may not model the system accurately, and therefore might sacrifice system performance. On the other hand, if λ is too small, the operator would not account for the behavior of HVs so that certain assigned trips may not be taken by the HVs, and thus worsen the performance of the proposed MPC.

V. Conclusion

We proposed a Stackelberg game-based framework to investigate an MoD system with a mixed fleet of AVs and

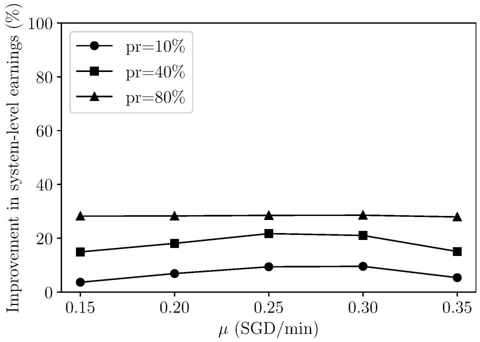

(a) Sensitivity to the assumed VOT of HVs.

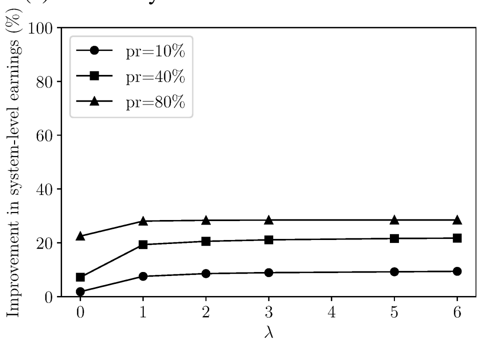

(b) Sensitivity to the weighting parameter λ.

Fig. 5: Sensitivity analysis for the Stackelberg game based MPC approach in scenarios with various penetration rates of AVs, where the prices, demand, and compensations are taken from the Singapore dataset.

HVs where AVs can be fully coordinated by the operator, but HVs act on their own interest. We developed a time-varying formulation to devise a real-time MPC based control algorithm. We further conducted real-world case studies based on a Singapore dataset to validate the proposed formulation and algorithms. Results show that the system-level earnings can be significantly improved (by up to 32%) by introducing AVs into the MoD system. We further show that by considering the interactions between AVs and HVs, the system-level earnings can be improved by up to 12%. Overall, this sheds light on the promising value of the proposed approach.

This research opens the field for several research directions. First, we would like to extend the proposed approaches to facilitate the integration with real-world MoD services by considering formulations with different objectives (e.g., operator's profit), higher spatio-temporal resolution, real-time pricing strategies (e.g., surge pricing), passenger preference over HVs and AVs, and HVs with heterogeneous properties. On the practical side, the strategy proposed in this work has the potential to improve current MoD systems by globally coordinating a small fleet of paid contractor drivers who act "as AVs". Second, it is of interest to account for service externalities, such as congestion effects of vehicle routing, fuel consumption, emissions, etc., which would be important factors to consider as the market share of the MoD service keeps growing. Third, we would like to consider interactions with other MoD operators and with public transport operators. Fourth, we would like to account for the limited driving capability of AVs during the transition period, in the sense that they may not be capable of handling every type of road condition and thus be restricted to operate within a certain subnetwork. Along this line, this work can also be integrated with a co-design framework [6] to jointly optimize the services of the MoD systems with the allocation of dedicated or smart infrastructure for AVs to provide guidelines to traffic authorities on how to manage and advocate shared AVs in cities.

REFERENCES

[1] L. Zha, Y. Yin, and Y. Du. Surge pricing and labor supply in the ride-sourcing market. *Transportation Research Part B: Methodological*, 117:708–722, 2018.

[2] J. Wen, N. Nassir, and J. Zhao. Value of demand information in autonomous mobility-on-demand systems. *Transportation Research Part A: Policy and Practice*, 121:346–359, 2019.

[3] R. Zhang and M. Pavone. Control of robotic Mobility-on-Demand systems: A queueing-theoretical perspective. *Int. Journal of Robotics Research*, 35(1–3):186–203, 2016.

[4] M. Salazar, F. Rossi, M. Schiffer, C. H. Onder, and M. Pavone. On the interaction between autonomous mobility-on-demand and the public transportation systems. In *Proc. IEEE Int. Conf. on Intelligent Transportation Systems*, 2018. Extended Version, Available at `https://arxiv.org/abs/1804.11278`.

[5] F. Boewing, M. Schiffer, M. Salazar, and M. Pavone. A vehicle coordination and charge scheduling algorithm for electric autonomous mobility-on-demand systems. In *American Control Conference*, 2020. Submitted.

[6] G. Zardini, N. Lanzetti, M. Salazar, A. Censi, E. Frazzoli, and M. Pavone. Towards a co-design framework for future mobility systems. In *Annual Meeting of the Transportation Research Board*, 2020.

[7] S. Hörl, C. Ruch, F. Becker, E. Frazzoli, and K.W. Axhausen. Fleet operational policies for automated mobility: A simulation assessment for zurich. *Transportation Research Part C: Emerging Technologies*, 102:20–31, 2019.

[8] M. Lokhandwala and H. Cai. Dynamic ride sharing using traditional taxis and shared autonomous taxis: A case study of nyc. *Transportation Research Part C: Emerging Technologies*, 97:45–60, 2018.

[9] P. Afeche, Z. Liu, and C. Maglaras. Ride-hailing networks with strategic drivers: The impact of platform control capabilities on performance. *Columbia Business School Research Paper*, (18-19), 2018.

[10] Q. Wei, R. Pedarsani, and S. Coogan. Mixed autonomy in ride-sharing networks. *IEEE Transactions on Control of Network Systems*, 7(4):1940–1950, 2020.

[11] L. Li, S. Wang, and F. Wang. An analysis of taxi driver's route choice behavior using the trace records. *IEEE Transactions on Computational Social Systems*, 5(2):576–582, 2018.

[12] S. Banerjee, R. Johari, and C. Riquelme. Pricing in ride-sharing platforms: A queueing-theoretic approach. In *ACM Conf. on Economics and Computation*, 2015.

[13] N. Buchholz. Spatial equilibrium, search frictions and efficient regulation in the taxi industry. Available at `https://scholar.princeton.edu/sites/default/files/nbuchholz/files/taxi_draft.pdf`, 2019.

[14] K. Bimpikis, O. Candogan, and D. Saban. Spatial pricing in ride-sharing networks. *Operations Research*, 2019.

[15] B. Yang, X. Zhang, and Q. Meng. Stackelberg games and multiple equilibrium behaviors on networks. *Transportation Research Part B: Methodological*, 41(8):841–861, 2007.

[16] M. Florian and C. Morosan. On uniqueness and proportionality in multi-class equilibrium assignment. *Transportation Research Part B: Methodological*, 70:173–185, 2014.

[17] D. Gammelli, I. Peled, F. Rodrigues, D. Pacino, and F.C. Pereira. Estimating latent demand of shared mobility through censored gaussian processes. *Transportation Research Part C: Emerging Technologies*, 120, 2020.

[18] K. Yang, M. Tsao, X. Xu, and M. Pavone. Planning and operations of mixed fleets in mobility-on-demand systems, 2020. Avaiable at `https://arxiv.org/abs/2008.08131`.

[19] Singapore Ministry of Manpower. Summary table: Income, 2020. Available at `https://stats.mom.gov.sg/Pages/Income-Summary-Table.aspx`.

[20] P. M. Boesch, F. Becker, H. Becker, and K. W. Axhausen. Cost-based analysis of autonomous mobility services. *Transport Policy*, 64:76–91, 2018.