

Contents lists available at ScienceDirect

Transportation Research Part C

journal homepage: www.elsevier.com/locate/trc





Dynamic pricing and fleet management for electric autonomous mobility on demand systems[☆]

Berkay Turan*, Ramtin Pedarsani, Mahnoosh Alizadeh

Department of Electrical and Computer Engineering, University of California, Santa Barbara, CA, 93106 USA

ARTICLE INFO

Keywords:

Autonomous mobility-on-demand systems Optimization and optimal control Reinforcement learning

ABSTRACT

The proliferation of ride sharing systems is a major drive in the advancement of autonomous and electric vehicle technologies. This paper considers the joint routing, battery charging, and pricing problem faced by a profit-maximizing transportation service provider that operates a fleet of autonomous electric vehicles. We first establish the static planning problem by considering timeinvariant system parameters and determine the optimal static policy. While the static policy provides stability of customer queues waiting for rides even if consider the system dynamics, we see that it is inefficient to utilize a static policy as it can lead to long wait times for customers and low profits. To accommodate for the stochastic nature of trip demands, renewable energy availability, and electricity prices and to further optimally manage the autonomous fleet given the need to generate integer allocations, a real-time policy is required. The optimal real-time policy that executes actions based on full state information of the system is the solution of a complex dynamic program. However, we argue that it is intractable to exactly solve for the optimal policy using exact dynamic programming methods and therefore apply deep reinforcement learning to develop a near-optimal control policy. The two case studies we conducted in Manhattan and San Francisco demonstrate the efficacy of our real-time policy in terms of network stability and profits, while keeping the queue lengths up to 200 times less than the static policy.

1. Introduction

The rapid evolution of enabling technologies for autonomous driving coupled with advancements in eco-friendly electric vehicles (EVs) has facilitated state-of-the-art transportation options for urban mobility. Owing to these developments in automation, it is possible for an autonomous-mobility-on-demand (AMoD) fleet of autonomous EVs to serve the society's transportation needs, with multiple companies now heavily investing in AMoD technology (https://www.cbinsights.com/research/autonomous-driverless-vehicles-corporations-list/).

The introduction of autonomous vehicles for mobility on demand services provides an opportunity for better fleet management. Specifically, idle vehicles can be *rebalanced* throughout the network in order to prevent accumulating at certain locations and to serve induced demand at every location. Autonomous vehicles allow rebalancing to be performed centrally by a platform operator who observes the state of all the vehicles and the demand, rather than locally by individual drivers. Furthermore, EVs provide opportunities for cheap and environment-friendly energy resources (e.g., solar energy). However, electricity supplies and prices differ among the

E-mail addresses: bturan@ucsb.edu (B. Turan), ramtin@ucsb.edu (R. Pedarsani), alizadeh@ucsb.edu (M. Alizadeh).

^{*} This work is supported by the NSF Grant 1847096.

^{*} Corresponding author.

network both geographically and temporally. As such, this diversity can be exploited for cheaper energy options when the fleet is operated by a platform operator that is aware of the electricity prices throughout the whole network. Moreover, a dynamic pricing scheme for rides is essential to maximize profits earned by serving the customers. Coupling an optimal fleet management policy with a dynamic pricing scheme allows the revenues to be maximized while reducing the rebalancing cost and the waiting time of the customers by adjusting the induced demand.

We consider a model that captures the opportunities and challenges of an AMoD fleet of EVs, and consists of complex state and action spaces. In particular, the platform operator has to consider the number of customers waiting to be served at each location (ride request queue lengths), the electricity prices, traffic conditions, and the states of the EVs (locations, battery energy levels) in order to make decisions. These decisions consist of pricing for rides for every origin–destination (OD) pair and routing/charging decision for every vehicle in the network. Upon taking an action, the state of the network undergoes through a stochastic transition due to the randomness in customer behaviour, electricity prices, and travel times.

We first adopt the common approach of network flow modeling to develop an optimal static pricing, routing, and charging policy that we use as a baseline in this paper. However, flow-based solutions generate fractional flows which can not directly be implemented. Moreover, a static policy executes same actions independent of the network state and is oblivious to the stochastic events that occur in the real setting. Hence, it is not optimal to utilize the static policy in a real dynamic environment. Therefore, a real-time policy that generates integer solutions and acknowledges the network state is required, and can be determined by solving the underlying dynamic program. Due to the continuous and high dimensional state-action spaces however, it is infeasible to develop an optimal real-time policy using exact dynamic programming algorithms. As such, we utilize deep reinforcement learning (RL) to develop a near-optimal policy. Specifically, we show that it is possible to learn a policy via Proximal Policy Optimization (PPO) (Schulman et al., 2017) that increases the total profits generated by jointly managing the fleet of EVs (by making routing and charging decisions) and pricing for the rides. We demonstrate the performance of our policy by using the total profits generated and the queue lengths as metrics.

Our contributions can be summarized as follows:

- 1. We formalize a vehicle and network model that captures the aforementioned characteristics of an AMoD fleet of EVs as well as the stochasticity in demand and electricity prices.
- 2. We analyze the static problem, where we consider a time-invariant environment (time-invariant arrivals, electricity prices, etc.) to characterize the family of policies that guarantee stability of the dynamic system, to gain insight towards the actual dynamic problem, and to further provide a baseline for comparison.
- 3. We employ deep RL methods to learn a joint pricing, routing and charging policy that effectively stabilizes the queues and increases the profits.

We visualize our real-time framework as a schematic diagram in Fig. 1 and preview our results in Fig. 2, showing that a real-time pricing and routing policy can successfully keep the queue lengths 400 times lower than the static policy. This policy is also able to decrease the charging costs by 25% by utilizing smart charging strategies (which will be demonstrated in Section 5).

Related work: Comprehensive research perceiving various aspects of AMoD systems is being conducted in the literature. Studies surrounding fleet management focus on optimal EV charging in order to reduce electricity costs as well as optimal vehicle routing in order to serve the customers and to rebalance the empty vehicles throughout the network so as to reduce the operational costs and the customers' waiting times. Time-invariant control policies adopting queueing theoretical (Zhang et al., 2016), fluidic (Pavone et al.,

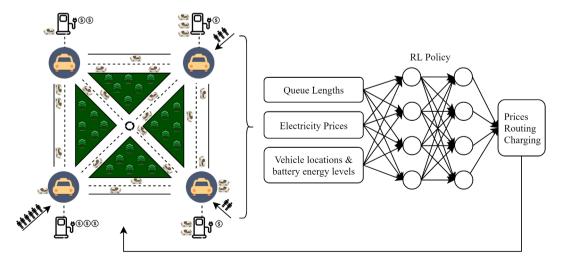


Fig. 1. The schematic diagram of our framework. Our deep RL agent processes the state of the vehicles, queues and electricity prices and outputs a control policy for pricing as well as autonomous EVs' routing and charging.

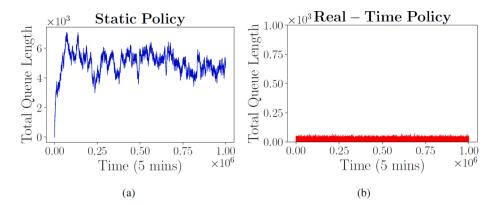


Fig. 2. (a) The optimal static policy manages to stabilize the queues over a very long time period but is unable to clear them whereas (b) RL control policy stabilizes the queues and manages to keep them significantly low (note the scales).

2012), network flow (Rossi et al., 2018), and Markovian (Volkov et al., 2012) models have been developed by using the steady state of the system. The authors of (Wei et al., 2019) consider ride-sharing systems with mixed autonomy. However, the proposed control policies in these papers are not adaptive to the time-varying nature of the future demand. As such, there is work on developing timevarying model predictive control (MPC) algorithms (Zhang et al., 2016; Miao et al., 2016; Iglesias et al., 2017; Miao et al., 2017; Tsao et al., 2018). The authors of (Iglesias et al., 2017; Miao et al., 2017) propose data-driven algorithms and the authors of (Tsao et al., 2018) propose a stochastic MPC algorithm focusing on vehicle rebalancing. In (Zhang et al., 2016), the authors also consider a fleet of EVs and hence propose an MPC approach that optimizes vehicle routing and scheduling subject to energy constraints. Using a fluidbased optimization framework, the authors of (Spieser et al., 2016) investigate tradeoffs between fleet size, rebalancing cost, and queueing effects in terms of passenger and vehicle flows under time-varying demand. The authors in (Swaszek and Cassandras, 2019) develop a parametric controller that approximately solves the intractable dynamic program for rebalancing over an infinite-horizon. Similar to AMoD, carsharing systems also require rebalancing in order to operate efficiently. By adopting a Markovian model, the authors of (Repoux et al., 2019) introduce a dynamic proactive rebalancing algorithm for carsharing systems by taking into account an estimate of the future demand using historical data. In (Boyaci et al., 2017), the authors develop an integrated multi-objective mixed integer linear programming optimization and discrete event simulation framework to optimize vehicle and personnel rebalancing in an electric carsharing system. Using a network-flow based model, the authors of (Warrington and Ruchti, 2019) propose a two-stage approximation scheme to establish a real-time rebalancing algorithm for shared mobility systems that accounts for stochasticity in customer demand and journey valuations.

Aside from these, there are studies on applications of RL methods in transportation such as adaptive routing (Mao and Shen, 2018), traffic management (Zhu and Ukkusuri, 2014; Walraven et al., 2016), traffic signal control (Zhu et al., 2015; Li et al., 2016), and dynamic routing of autonomous vehicles with the goal of reducing congestion in mixed autonomy traffic networks (Lazar et al., 2019). Relevant studies to our work aim to develop dynamic policies for rebalancing as well as ride request assignment via decentralized reinforcement learning approaches (Han et al., 2016; Guériau and Dusparic, 2018; Wen et al., 2017; Lin et al., 2018). In these works however, the policies are developed and applied locally by each autonomous vehicle and this decentralized approach may sacrifice system level optimality. A centralized deep RL approach tackling the rebalancing problem is proposed in (Mao et al., 2020), which is closest to the approach we adopt in this paper. Although their study adopts a centralized deep RL approach similar to our paper, they have a different system model and solely focus on the rebalancing problem and do not consider pricing for rides as a control variable for the queues nor the charging problem of EVs as reviewed next.

Regarding charging strategies for large populations of EVs, (Veldman and Verzijlbergh, 2015; Su et al., 2012; Mukherjee and Gupta, 2015) provide in-depth reviews and studies of smart charging technologies. An agent-based model to simulate the operations of an AMoD fleet of EVs under various vehicle and infrastructure scenarios has been examined in (Chen et al., 2016). By augmenting optimal battery management of autonomous electric vehicles to the classic dial-a-ride problem (DARP), the authors of (Bongiovanni et al., 2019) introduce the electric autonomous DARP that aims to minimize the total travel time of all the vehicles and riders. The authors of (Tucker et al., 2019) propose an online charge scheduling algorithm for EVs providing AMoD services. By adopting a static network flow model in (Turan et al., 2019), the benefits of smart charging have been investigated and approximate closed form expressions that highlight the trade-off between operational costs and charging costs have been derived. Furthermore, (Rossi et al., 2018) studies interactions between AMoD systems and the power grid. In addition, (Chen and Kockelman, 2016) studies the implications of pricing schemes on an AMoD fleet of EVs. In (Guan et al., 2019), the authors propose a dynamic joint pricing and routing strategy for non-electric shared mobility on demand services. (Sheppard et al., 2019) studies a quadratic programming problem in order to jointly optimize vehicle dispatching, charge scheduling, and charging infrastructure, while the demand is defined exogenously.

To the best of our knowledge, there is no existing work on centralized real-time management for electric AMoD systems addressing the joint optimization scheme of vehicle routing and charging as well as pricing for the rides. In this paper we aim to highlight the benefits of a real-time controller that jointly: (i) routes the vehicles throughout the network in order to serve the demand for rides as well as to relocate the empty vehicles for further use, (ii) executes smart charging strategies by exploiting the diversity in the electricity

prices (both geographically and temporally) in order to minimize charging costs, and (iii) adjusts the demand for rides by setting prices in order to stabilize the system (i.e., the queues of customers waiting for rides) while maximizing profits.

Paper Organization: The remainder of the paper is organized as follows. In Section 2, we present the system model and define the platform operator's optimization problem. In Section 3, we discuss the static planning problem associated with the system model and characterize the optimal static policy. In Section 4, we propose a method for developing a near-optimal real-time policy using deep reinforcement learning. In Section 5, we present the numerical results of the case studies we have conducted in Manhattan and San Francisco to demonstrate the performance of our real-time control policy. Finally, we conclude the paper in Section 6.

2. System model and problem definition

Network and Demand Models: We consider a fleet of AMoD EVs operating within a transportation network characterized by a fully connected graph consisting of $\mathcal{M} = \{1, ..., m\}$ nodes that can each serve as a trip origin or destination. We study a discrete-time system with time periods normalized to integral units $t \in \{0,1,2,...\}$. In this discrete-time system, we model the arrival of the potential riders with OD pair (i,j) as a Poisson process with an arrival rate of $\lambda_{ij}(t)$ in period t, where $\lambda_{ii}(t) = 0$. We adopted a price-responsive rider model studied in (Bimpikis et al., 2019). We assume that the riders are heterogeneous in terms of their willingness to pay. In particular, if the price for receiving a ride from node i to node j in period t is set to $\ell_{ij}(t)$, the induced arrival rate for rides from i to j is given by $\lambda_{ij}(t) = \lambda_{ij}(t)(1 - F(\ell_{ij}(t)))$, where $F(\cdot)$ is the cumulative distribution of riders' willingness to pay with a support of $[0, \ell_{\max}]^{-1}$. Thus, the number of new ride requests in time period t is $A_{ij}(t)$ Pois $(\lambda_{ij}(t))$ for OD pair (i,j).

Vehicle Model: To capture the effect of trip demand and the associated charging and routing (routing also implies rebalancing of the empty vehicles) decisions on the costs associated with operating the fleet (maintenance, mileage, etc.), we assume that each autonomous vehicle in the fleet has a per period operational cost of *β*. Furthermore, as the vehicles are electric, they have to sustain charge in order to operate. Without loss of generality, we assume there is a charging station placed at each node $i \in \mathcal{M}$. To charge at node i during time period i, the operator pays a price of electricity i per unit of energy. We assume that all EVs in the fleet have a battery capacity denoted as i0 max i1 therefore, each EV has a discrete battery energy level i1 more i2 max i3. In our discrete-time model, we assume each vehicle takes one period to charge one unit of energy and i3 periods to travel between OD pair i3 units of energy.

Ride Hailing Model: The platform operator dynamically routes the fleet of EVs in order to serve the demand at each node. Customers that purchase a ride are not immediately matched with a ride, but enter the queue for OD pair (i,j). After the platform operator executes routing decisions for the fleet, the customers in the queue for OD pair (i,j) are matched with rides and served in a first-come, first-served discipline. A measure of the expected wait time is not available to each arriving customer. However, the operator knows that longer wait times will negatively affect their business and hence seeks to minimize the total wait time experienced by users. Denote the queue length for OD pair (i,j) by $q_{ij}(t)$. If after serving the customers, the queue length $q_{ij}(t) > 0$, the platform operator is penalized by a fixed cost of w per person at the queue to account for the value of time of the customers.

Platform Operator's Problem: We consider a profit-maximizing AMoD operator that manages a fleet of EVs that make trips to provide transportation services to customers. The operator's goal is to maximize profits by (1) setting prices for rides and hence managing customer demand at each node; (2) optimally operating the AMoD fleet (i.e., charging and routing) to minimize operational and charging costs. We will study two types of control policies the platform operator utilizes: (1) a static policy, where the pricing, routing and charging decisions are time invariant and independent of the state of the system; (2) a real-time policy, where the pricing, routing and charging decisions are dependent on the system state.

3. Analysis of the static problem

In this section, we establish and discuss the static planning problem to provide a measure for comparison and demonstrate the efficacy of the real-time policy (which will be discussed in Section 4). To do so, we consider the fluid scaling of the dynamic network and characterize the static problem via a network flow formulation. Under this setting, we use the expected values of the variables (arrivals and prices of electricity) and ignore their time dependent dynamics, while allowing the vehicle routing decisions to be flows (real numbers) rather than integers. The static problem is convenient for determining the optimal static pricing, routing, and charging policy, under which the queueing network of the dynamic system is stable (Pedarsani et al., 2017)³.

3.1. Static profit maximization problem

We formulate the static optimization problem via a network flow model that aims to maximize the platform operator's profits. The

¹ For brevity of notation, we uniformly set ℓ_{\max} to be the maximum willingness to pay for all OD pairs without loss of generality. Our results can be derived in a similar fashion by replacing ℓ_{\max} with ℓ_{\max}^{ij} , where ℓ_{\max}^{ij} is the maximum willingness to pay for OD pair (i,j).

² In this paper, we consider the travel times to be constant and exogenously defined for the time period the policy is developed for. This is because we assume that the number of AMoD vehicles is much less compared to the rest of the traffic. Also, to consider changing traffic conditions throughout the day, it is possible to train multiple static and real-time control policies for the different time intervals.

The stability condition that we are interested in is rate stability of all queues. A queue for OD pair (i,j) is rate stable if $\lim_{t\to\infty}q_{ij}(t)/t=0$.

platform operator maximizes its profits by setting prices and making routing and charging decisions such that the system remains stable

Let ℓ_{ij} be the prices for rides for OD pair (i,j), x_{ij}^{ν} be the number of vehicles at node i with energy level ν being routed to node j, and x_{ic}^{ν} be the number of vehicles charging at node i starting with energy level ν . We state the platform operator's profit maximization problem as follows:

$$\max_{x_{ic}^{\nu}, x_{ij}^{\nu}, \ell_{ij}^{\nu}} \sum_{i \in \mathcal{M}} \sum_{j \in \mathcal{M}} \lambda_{ij} \ell_{ij} \left(1 - F(\ell_{ij}) \right) - \sum_{i \in \mathcal{M}} \sum_{\nu=0}^{\nu_{\max}-1} (\beta + p_i) x_{ic}^{\nu} - \beta \sum_{i \in \mathcal{M}} \sum_{j \in \mathcal{M}} \sum_{i \neq \nu} x_{ij}^{\nu} \tau_{ij}$$
(1a)

s.t.
$$\lambda_{ij} (1 - F(\ell_{ij})) \leqslant \sum_{v=v}^{v_{\text{max}}} x_{ij}^{v} \quad \forall i, j \in \mathcal{M},$$
 (1b)

$$x_{ic}^{\nu} + \sum_{i \in \mathcal{M}} x_{ij}^{\nu} = x_{ic}^{\nu-1} + \sum_{i \in \mathcal{M}} x_{ji}^{\nu+\nu_{ji}} \quad \forall i \in \mathcal{M}, \ \forall \nu \in \mathcal{V},$$
 (1c)

$$x_{ic}^{v_{\max}} = 0 \quad \forall i \in \mathcal{M},$$
 (1d)

$$x_{ii}^{v} = 0 \quad \forall v < v_{ij}, \ \forall i, j \in \mathcal{M},$$
 (1e)

$$x_{ir}^{\nu} \geqslant 0, \ x_{ir}^{\nu} \geqslant 0 \ \forall i, j \in \mathcal{M}, \ \forall v \in \mathcal{V},$$
 (1f)

$$x_{ir}^v = x_{ii}^v = 0 \quad \forall v \notin \mathscr{V}, \, \forall i, j \in \mathscr{M}.$$
 (1g)

The first term in the objective function in (1a) accounts for the aggregate revenue the platform generates by providing rides for $\lambda_{ij}(1-F(\ell_{ij}))$ number of riders with a price of ℓ_{ij} . The second term is the operational and charging costs incurred by the charging vehicles (assuming that $p_i(t) = p_i \ \forall t$ under the static setting), and the last term is the operational costs of the trip-making vehicles (including rebalancing trips).

The constraint (1b) requires the platform to operate at least as many vehicles to serve all the induced demand between any two nodes i and j (The rest are the vehicles travelling without passengers, i.e., rebalancing vehicles). We will refer to this as the *demand* satisfaction constraint. The constraint (1c) is the *flow balance constraint* for each node and each battery energy level, which restricts the number of available vehicles at node i and energy level v to be the sum of arrivals from all nodes (including idle vehicles) and vehicles that are charging with energy level v-1. The constraint (1d) ensures that the vehicles with full battery do not charge further, and the constraint (1e) ensures the vehicles sustain enough charge to travel between OD pair (i,j).

The solution to the optimization problem in (1) is the optimal static policy that consists of optimal prices as well as optimal vehicle routing and charging decisions. This policy can not directly be implemented in a real environment because it does not yield integer valued solutions. It is possible generate integer-valued solutions to be implemented in a real environment using the fractional flows (e. g., randomizing the vehicle decisions according to the flows, which we do in Section 5), yet the methodology is not the focus of our work. Instead, we highlight a sufficient condition for a realizable policy (generating integer valued actions) to provide stability according to the feasible solutions of (1):

Proposition 1. Let $\{\widetilde{\ell}_{ij}, \widetilde{x}_{ic}^{\nu}, \widetilde{x}_{ic}^{\nu}\}$ be a feasible solution of (1). Let μ be a policy that generates integer actions and can be implemented in the real environment. Then, μ guarantees stability of the system if for all OD pairs (i,j):

- 1. The time average of the induced arrivals equals $(1-F(\widetilde{\ell}_{ij}))$, and
- 2. The time average of the routed vehicles equals $\sum_{\nu=\nu_{ij}}^{\nu_{\max}} \widetilde{\chi}_{ij}^{\nu}$.

The proof of Proposition 1 is provided in Appendix A. According to Proposition 1, for a static pricing policy with the optimal prices ℓ_{ii}^* , there exists an integer-valued routing and charging policy that maintains stability of the system.

Corollary 1.1. An example policy that generates integer-valued actions is randomizing according to the flows. Precisely, given a feasible solution $\{\tilde{\ell}_{ij}, \tilde{\chi}_{ij}^{\nu}, \tilde{\chi}_{ic}^{\nu}\}$ of (1), integer-valued actions can be generated by routing a vehicle at node i with energy level ν to node i with probability

$$\psi_{ij}^{\nu} = \frac{\widetilde{x}_{ij}^{\nu}}{\sum_{k=1}^{m} \widetilde{x}_{ik}^{\nu} + x_{ic}^{\nu}},$$

and charging with probability

$$\psi_{ic}^{\nu} = \frac{\widetilde{x}_{ic}^{\nu}}{\sum_{k=1}^{m} \widetilde{x}_{ik}^{\nu} + x_{ic}^{\nu}},$$

 $\forall i,j \in \mathscr{M} \text{ and } \forall \nu \in \mathscr{V}.$ Combining this randomized policy with a static pricing policy of $\ell_{ij}(t) = \widetilde{\ell}_{ij}, \ \forall t$, results in a policy satisfying the criteria in Proposition 1. The optimization problem in (1) is non-convex for a general $F(\cdot)$. Nonetheless, when the platform's profits are convex in the induced demand $\lambda_{ij}(1-F(\cdot))$, it can be rewritten as a convex optimization problem and can be solved exactly. Hence, we assume that the rider's willingness to pay is uniformly distributed in $[0,\ell_{\text{max}}]$, i.e., $F\left(\ell_{ij}\right) = \frac{\ell_{ij}}{\ell_{\text{max}}} \frac{4}{\ell}$.

Marginal Pricing: The prices for rides are a crucial component of the profits generated. The next proposition highlights how the optimal prices ℓ_{ii}^* for rides are related to the network parameters, prices of electricity, and the operational costs.

Proposition 2. Let ν_{ij}^* be optimal the dual variable corresponding to the demand satisfaction constraint for OD pair (i,j). The optimal prices ℓ_{ii}^* are:

$$\ell_{ij}^* = \frac{\ell_{\text{max}} + \nu_{ij}^*}{2}.$$
 (2)

These prices can be upper bounded by:

$$\mathcal{E}_{ij}^* \leqslant \frac{\ell_{\max} + \beta(\tau_{ij} + \tau_{ji} + \nu_{ij} + \nu_{ji}) + \nu_{ij}p_j + \nu_{ji}p_i}{2}$$
(3)

Moreover, with these optimal prices ℓ_{ii}^* , the profits generated per period is:

$$P = \sum_{i=1}^{m} \sum_{j=1}^{m} \frac{\lambda_{ij}}{\ell_{\max}} \left(\ell_{\max} - \ell_{ij}^*\right)^2. \tag{4}$$

The proof of Proposition 2 is provided in Appendix B. Observe that the profits in Eq. (4) are decreasing as the prices for rides increase. Thus expensive rides generate less profits compared to the cheaper rides and it is more beneficial if the optimal dual variables ν_{ij}^* are small and prices are close to $\ell_{\text{max}}/2$. We can interpret the dual variables ν_{ij}^* as the cost of providing a single ride between i and j to the platform. In the worst case scenario, every single requested ride from node i requires rebalancing and charging both at the origin and the destination. Hence the upper bound on (3) includes the operational costs of passenger-carrying, rebalancing and charging vehicles (both at the origin and the destination); and the energy costs of both passenger-carrying and rebalancing trips multiplied by the price of electricity at the trip destinations. Similar to the taxes applied on products, whose burden is shared among the supplier and the customer; the costs associated with rides are shared among the platform operator and the riders (which is why the price paid by the riders include half of the cost of the ride).

Although the static policy guarantees stability (by appropriate implementation of integer-valued actions as dictated by Proposition 1), it does not perform well in a real dynamic setting because it does not acknowledge the stochastic dynamics of the system. On the other hand, a real-time policy that executes decisions based on the current state of the environment would likely perform better (e.g., if the queue length for OD pair (i,j) is very large, then it is probably better for the platform operator to set higher prices to prevent the queue from growing further). Accordingly, we present a practical way of implementing a real-time policy in the next section.

4. The real-time policy

The static policy established in the previous section has three major issues:

- 1. Because it is based on a flow model, it generates static fractional flows that are not directly implementable in the real setting.
- 2. It neglects the stochastic events that occur in the dynamic setting (e.g., the induced arrivals), and assumes everything is deterministic. Hence, it does consider the unexpected occurrences (e.g., queues might build in the dynamic setting, whereas the static model assumes no queues) when executing actions.
- 3. It assumes perfect knowledge of the network parameters (arrivals, trip durations, energy consumptions of the trips, and prices of electricity).

Due to the above reasons, it is impractical to implement the static policy in the dynamic environment. A real-time policy that

⁴ It is also possible to use other distributions that might reflect real willingness-to-pay distributions more accurately (such as pareto distribution, exponential distribution, triangular distribution, constant elasticity distribution, and normal distribution). Among these, pareto, exponential, and constant elasticity distributions preserve convexity and therefore the static planning problem can be solved efficiently. Triangular and normal distributions are not convex in their support and therefore the static planning problem is not a convex optimization problem. Nevertheless, it can still be solved numerically for the optimal static policy. Using these distributions however we cannot derive the closed-form results that allow us to interpret the pricing policy of the platform operator. The real-time policy proposed in Section 4 uses model-free Reinforcement Learning and therefore can be applied using other distributions or any other customer price response model.

generates integer solutions and takes into account the current state of the network which is essential for decision making is necessary, and can be determined by solving the dynamic program that describes the system (with full knowledge of the network parameters) for the optimal policy. Such solutions would address issues 1 and 2 outlined above. Inspired by our theoretical model, the state information that describes the network fully consists of the vehicle states (locations, energy levels), queue lengths for each OD pair, and electricity prices at each node. Upon obtaining the full state information, the actions have to be executed for pricing for rides and fleet management (vehicle routing and charging). Consequent to taking actions, the platform operator observes a reward (consisting of revenue gained by arrivals, queue costs, and operational and charging costs), and the network transitions into a new state (Although the transition into the new state is stochastic, the random processes that govern this stochastic transition is known if the network parameters are known). The solution of this dynamic program is the optimal policy that determines which action to take for each state the system is in, and can nominally be derived using classical exact dynamic programming algorithms (e.g., value iteration). However, the complexity and the scale of our dynamic problem presents a difficulty here: Aside from having a large dimensional state space (for instance, $m=10, \nu_{\max}=5, \tau_{ij}=3 \ \forall i,j$: the state has dimension 1240) and action space, the cardinality of these spaces are not finite (queues can grow unbounded, prices are continuous). Considering that the computational complexity per iteration for value iteration is $\mathscr{O}\left(\left|\mathscr{A}\right|\left|\mathscr{S}\right|^2\right)$ and for policy iteration $\mathscr{O}\left(\left|\mathscr{A}\right|\left|\mathscr{S}\right|^2+\left|\mathscr{S}\right|^3\right)$ (Kaelbling et al., 1996), where \mathscr{S} and \mathscr{A} are the state space and the action space, respectively, the problem is computationally intractable to solve using classical dynamic programming. Even if we did make them finite by putting a cap on the queue lengths and discretizing the prices, curse of dimensionality renders the problem intractable to solve with classical exact dynamic programming algorithms. As such, we resort to approximate dynamic programming methods. Specifically, we define the policy via a deep neural network that takes the full state information of the network as input and outputs the best action⁵. Subsequently, we apply a model-free reinforcement learning algorithm to train the neural network in order to improve the performance of the policy. Since it is model-free, it does not require a modeling of the network (hence, it does not require knowledge of the network parameters), which resolves the third issue associated with the static policy.

We adopted a practical policy gradient method, called Proximal Policy Optimization (PPO), developed in (Schulman et al., 2017), which is effective for optimizing large nonlinear policies such as neural networks. We chose PPO mainly because it supports continuous state-action spaces and guarantees monotonic improvement.⁶

We note that it is possible to apply reinforcement learning to learn a policy in any environment, real or artificial, as long as there is data available. In this work we use our theoretical model described in Section 2 to create the environment and generate data, mainly because there is no electric AMoD microsimulation environment available and also to verify our findings about the static policy. Developing a microsimulator for electric AMoD (like SUMO (Lopez et al., 2018)) and integrating it with a deep reinforcement learning library to create a framework for real traffic experiments remains a future work. To ensure that our numerical experiments are reproducible, in the next subsection, we describe the Markov Decision Process (MDP) that governs this dynamic environment, which is a direct extension of our static model. It is also possible to enrich the environment and the MDP to reflect real life constraints more accurately such as road capacity and charging station constraints. Since the approach we adopt to develop the real-time policy is model-free, it can be applied identically.

In Section 5 we present numerical results on real-time policies developed through reinforcement learning based on dynamic environments generated through our theoretical model. The goal of the experiments is to primarily answer the following questions:

- 1. Can we develop a real-time control and pricing policy for AMoD using reinforcement learning and what are its potential benefits over the static policy?
- 2. How does the policy trained for a specific network perform, if the network parameters change?
- 3. Can we develop a global policy that can be utilized in any network with moderate fine tuning?

The reader may skip reading Section 4.1 if they are not interested in the details of the MDP model used in our numerical experiment.

4.1. The Real-Time Problem as MDP

We define the MDP by the tuple $(\mathcal{S}, \mathcal{A}, \mathcal{T}, r)$, where \mathcal{S} is the state space, \mathcal{A} is the action space, \mathcal{T} is the state transition operator and r is the reward function. We describe these elements as follows:

1. \mathscr{S} : The state space consists of prices of electricity at each node, the queue lengths for each origin–destination pair, and the number of vehicles at each node and each energy level. However, since travelling from node i to node j takes τ_{ij} periods of time, we need to define intermediate nodes. As such, we define $\tau_{ij}-1$ number of intermediate nodes between each origin and destination pair, for each battery energy level ν . Hence, the state space consists of $s_d = m^2 + (\nu_{\text{max}} + 1)((\sum_{i=1}^m \sum_{j=1}^m \tau_{ij}) - m^2 + 2m)$ dimensional vectors in $\mathbb{R}^{s_d}_{>0}$ (We include all the non-negative valued vectors, however, only $m^2 - m$ entries can grow to infinity because they are queue

⁵ In general, the policy is a stochastic policy and determines the probabilities of taking the actions rather than deterministically producing an action.

⁶ Although the policy outputs a continuous set of actions, integer actions can be generated by randomizing. This is done during both training and testing, therefore the RL agent observes the integer state transitions and learns as if the policy outputs integer actions. We discuss how to generate integer actions in more detail in Section 4.1.

lengths, and the rest are always upper bounded by fleet size or maximum price of electricity). As such, we define the elements of the state vector at time t as $s(t) = [p(t) \ q(t) \ s_{veh}(t)]$, where $p(t) = [p_i(t)]_{i \in \mathcal{M}}$ is the electricity prices state vector, $q(t) = [q_{ij}(t)]_{ij \in \mathcal{M}; i \neq j}$ is the queue lengths state vector, and $s_{veh}(t) = [s_{ijk}^v(t)]_{\forall i,j,k,\nu}$ is the vehicle state vector, where $s_{ijk}^v(t)$ is the number of vehicles at vehicle state (i,j,k,ν) . The vehicle state (i,j,k,ν) specifies the location of a vehicle that is travelling between OD pair (i,j) as the k'th intermediate node between nodes i and j, and specifies the battery energy level of a vehicle as ν (The states of the vehicles at the nodes $i \in \mathcal{M}$ with energy level ν is denoted by $(i,i,0,\nu)$).

- 2. \mathscr{N} : The action space consists of prices for rides at each origin–destination pair and routing/charging decisions for vehicles at nodes $i \in \mathscr{M}$ at each energy level v. The price actions are continuous in range $[0,\ell_{\max}]$. Each vehicle at state (i,i,0,v) ($\forall i \in \mathscr{M}, \forall v \in \mathscr{V}$) can either charge, stay idle or travel to one of the remaining m-1 nodes. To allow for different transitions for vehicles at the same state (some might charge, some might travel to another node), we define the action taken at time t for vehicles at state (i,i,0,v) as an m+1 dimensional probability vector with entries in [0,1] that sum up to 1: $\alpha_i^v(t) = [\alpha_{i1}^v(t)...\alpha_{im}^v(t) \ \alpha_{ic}^v(t)]$, where $\alpha_{ic}^{v_{\max}}(t) = 0$ and $\alpha_{ij}^v(t) = 0$ if $v < v_{ij}$. The action space is then all the vectors a of dimension $a_d = m^2 m + (v_{\max} + 1)(m^2 + m)$, whose first $m^2 m$ entries are the prices and the rest are the probability vectors satisfying the aforementioned properties. As such, we define the elements of the action vector at time t as $a(t) = [\ell(t) \ \alpha(t)]$, where $\ell'(t) = [\ell_{ij}]_{i,j \in \mathscr{M}, i \neq j}$ is the vector of prices and $\alpha(t) = [\alpha_i^v(t)]_{\forall i,v}$ is the vector of routing/charging actions.
- 3. \mathcal{F} : The transition operator is defined as $\mathcal{F}_{ijk} = Pr(s(t+1) = j | s(t) = i, a(t) = k)$. We can define the transition probabilities for electricity prices p(t+1), queue lengths q(t+1), and vehicle states $s_{veh}(t+1)$ as follows:

Electricity Price Transitions: Since we assume that the dynamics of prices of electricity are exogenous to our AMoD system, $Pr(p(t+1) = p_2|p(t) = p_1, a(t)) = Pr(p(t+1) = p_2|p(t) = p_1)$, i.e., the dynamics of the price are independent of the action taken. Depending on the setting, new prices might either be deterministic or distributed according to some probability density function at time $t : p(t)^{\sim} \mathcal{P}(t)$, which is determined by the electricity provider.

Vehicle Transitions: For each vehicle at node i and energy level v, the transition probability is defined by the action probability vector $\alpha_i^v(t)$. Each vehicle transitions into state $(i,j,1,v-v_{ij})$ with probability $\alpha_{ij}^v(t)$, stays idle in state (i,i,0,v) with probability $\alpha_{il}^v(t)$ or charges and transitions into state (i,i,0,v+1) with probability $\alpha_{ic}^v(t)$. The vehicles at intermediate states (i,j,k,v) transition into state (i,j,k+1,v) if $k < \tau_{ij} - 1$ or (j,j,0,v) if $k = \tau_{ij} - 1$ with probability 1. The total transition probability to the vehicle states $s_{veh}(t+1)$ given $s_{veh}(t)$ and $\alpha(t)$ is the sum of all the probabilities of the feasible transitions from $s_{veh}(t)$ to $s_{veh}(t+1)$ under $\alpha(t)$, where the probabilities are independent). Note that instead of gradually dissipating the energy of the vehicles on their route, we immediately discharge the required energy for the trip from their batteries and keep them constant during the trip. This ensures that the vehicles have enough battery to complete the ride and does not violate the model, because the vehicles arrive to their destinations with true value of energy and a new action will only be taken when they reach the destination.

Queue Transitions: The queue lengths transition according to the prices and the vehicle routing decisions. For prices $\ell_{ij}(t)$ and induced arrival rate $\Lambda_{ij}(t)$, the probability that $A_{ij}(t)$ new customers arrive in the queue (i,j) is:

$$Pr\bigg(A_{ij}\bigg(t\bigg)\bigg) = rac{e^{-\Lambda_{ij}ig(t)}\Lambda_{ij}(t)^{A_{ij}ig(t)}}{ig(A_{ij}ig(tig)!}$$

Let us denote the total number of vehicles routed from node i to j at time t as $x_{ij}(t)$, which is given by:

$$x_{ij}\left(t\right) = \sum_{\nu = \nu_{ii}}^{\nu_{\max}} x_{ij}^{\nu}\left(t\right) = \sum_{\nu = \nu_{ii}}^{\nu_{\max}} s_{ij1}^{\nu - \nu_{ij}}\left(t + 1\right). \tag{5}$$

Given $s_{veh}(t+1)$ and $x_{ij}(t)$, the probability that the queue length $q_{ij}(t+1) = q$ is:

$$Pr(q_{ij}(t+1) = q|s(t), a(t), s_{veh}(t+1)) = Pr(A_{ij}(t) = q - q_{ij}(t) + x_{ij}(t)),$$

if q > 0, and $Pr(A_{ij}(t) \le -q_{ij}(t) + x_{ij}(t))$ if q = 0. Since the arrivals are independent, the total probability that the queue vector q(t+1) = q is:

$$Pr(\boldsymbol{q}(t+1) = \boldsymbol{q}|\boldsymbol{s}(t), \boldsymbol{a}(t), \boldsymbol{s_{veh}}(t+1)) = \prod_{i \in \mathcal{M}} \prod_{i \neq j \atop i \neq i} Pr\big(q_{ij}(t+1) \big| \boldsymbol{s}(t), \boldsymbol{a}(t), \boldsymbol{s_{veh}}(t+1)\big).$$

Hence, the transition probability is defined as:

$$Pr(s(t+1)|s(t),\boldsymbol{a}(t)) = Pr(\boldsymbol{p}(t+1)|\boldsymbol{p}(t)) \times Pr(s_{veh}(t+1)|s(t),\boldsymbol{\alpha}(t)) \times Pr(q(t+1)|s(t),\boldsymbol{\alpha}(t),s_{veh}(t+1))$$

$$\tag{6}$$

We illustrate how the vehicles and queues transition into new states consequent to an action in Fig. 3.

4. r: The reward function r(t) is a function of state-action pairs at time t : r(t) = r(a(t), s(t)). Let $x_{ic}^{\nu}(t)$ denote the number of vehicles charging at node i starting with energy level ν at time period t. The reward function r(t) is defined as:

$$r(t) = \sum_{i \in \mathscr{M}} \sum_{j \in \mathscr{M}} \ell_{ij}(t) A_{ij}(t) - w \sum_{i \in \mathscr{M}} \sum_{j \in \mathscr{M}} q_{ij}(t) - \sum_{i \in \mathscr{M}} \sum_{v=0}^{v_{\text{max}}-1} (\beta + p_i) x_{ic}^v(t)$$

$$j \neq i \qquad j \neq i \qquad j \neq i$$

$$-\beta \sum_{i \in \mathscr{M}} \sum_{j \in \mathscr{M}} x_{ij}(t) - \beta \sum_{i \in \mathscr{M}} \sum_{j \in \mathscr{M}} \sum_{k=1}^{\tau_{ij}-1} \sum_{v=0}^{v_{\text{max}}-1} s_{ijk}^v(t)$$

$$j \neq i \qquad j \neq i$$

$$(7)$$

The first term corresponds to the revenue generated by the passengers that request a ride for a price $\ell_{ij}(t)$, the second term is the queue cost of the passengers that have not yet been served, the third term is the charging and operational costs of the charging vehicles and the last two terms are the operational costs of the vehicles making trips. Note that revenue generated is immediately added to the reward function when the passengers enter the network instead of after the passengers are served. Since the reinforcement learning approach is based on maximizing the cumulative reward gained, all the passengers eventually have to be served in order to prevent queues from blowing up and hence it does not violate the model to add the revenues immediately.

Using the definitions of the tuple $(\mathcal{S}, \mathcal{A}, \mathcal{T}, r)$, we model the dynamic problem as an MDP. Given large-dimensional state and action spaces with infinite cardinality, we can not solve the MDP using exact dynamic programming methods. As a solution, we characterize the real-time policy via a deep neural network and execute reinforcement learning in order to develop a real-time policy.

4.2. Reinforcement learning method

In this subsection, we go through the preliminaries of reinforcement learning and briefly explain the idea of the algorithm we adopted.

4.2.1. Preliminaries

The real-time policy associated with the MDP is defined as a function parameterized by θ :

$$\pi_{\theta}(\boldsymbol{a}|\boldsymbol{s}) = \pi : \mathscr{S} \times \mathscr{A} \to [0,1],$$

i.e., a probability distribution in the state-action space. Given a state s, the policy returns the probability for taking the action a (for all actions), and samples an action according to the probability distribution. The goal is to derive the optimal policy π^* , which maximizes the discounted cumulative expected rewards J_{π} :

$$J_{\pi^*} = \max_{\pi} J_{\pi} = \max_{\pi} \mathbb{E}_{\pi} \left[\sum_{t=0}^{\infty} \gamma^t r \left(t \right) \right],$$

$$\pi^* = \arg\max_{\pi} \mathbb{E}_{\pi} \left[\sum_{t=0}^{\infty} \gamma^t r \left(t \right) \right],$$

where $\gamma \in (0,1]$ is the discount factor. The value of taking an action a in state s, and following the policy π afterwards is characterized by the value function $Q_{\pi}(s,a)$:

$$Q_{\pi}\left(s,a\right) = \mathbb{E}_{\pi}\left[\sum_{t=0}^{\infty} \gamma^{t} r\left(t\right) \middle| s\left(0\right) = s, a\left(0\right) = a\right].$$

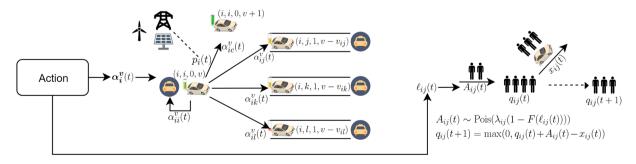


Fig. 3. The schematic diagram representing the state transition of our MDP. Upon taking an action, a vehicle at state $(i, i, 0, \nu)$ charges for a price of $p_i(t)$ and transitions into state $(i, i, 0, \nu+1)$ with probability $\alpha_{ic}^{\nu}(t)$, stays idle at state $(i, i, 0, \nu)$ with probability $\alpha_{il}^{\nu}(t)$, or starts traveling to another node j and transitions into state $(i, j, 1, \nu - \nu_{ij})$ with probability $\alpha_{ij}^{\nu}(t)$. Furthermore, $A_{ij}(t)$ new customers arrive to the queue (i, j) depending on the price $\ell_{ij}(t)$. After the routing and charging decisions are executed for all the EVs in the fleet, the queues are modified.

The value of being in state s is formalized by the value function $V_{\pi}(s)$:

$$V_{\pi}igg(sigg) = \mathbb{E}_{a(0),\pi}igg[\sum_{t=0}^{\infty} \gamma^t rigg(tigg)igg|sigg(0igg) = sigg],$$

and the advantage of taking the action a in state s and following the policy π thereafter is defined as the advantage function $A_{\pi}(s,a)$:

$$A_{\pi}(s, \boldsymbol{a}) = Q_{\pi}(s, \boldsymbol{a}) - V_{\pi}(s).$$

The methods used by reinforcement learning algorithms can be divided into three main groups: 1) critic-only methods, 2) actor-only methods, and 3) actor-critic methods, where the word critic refers to the value function and the word actor refers to the policy (Grondman et al., 2012). Critic-only (or value-function based) methods (such as Q-learning (Watkins and Dayan, 1992) and SARSA (Rummery and Niranjan, 1994)) improve a deterministic policy using the value function by iterating:

$$\boldsymbol{a}^* = \arg\max_{\boldsymbol{a}} Q_{\pi}\left(\boldsymbol{s}, \boldsymbol{a}\right), \pi(\boldsymbol{a}^*|\boldsymbol{s}) \leftarrow 1.$$

Actor-only methods (or policy gradient methods), such as Williams' REINFORCE algorithm (Williams, 1992), improve the policy by updating the parameter θ by gradient ascent, without using any form of a stored value function:

$$\theta\bigg(t+1\bigg) = \theta\bigg(t\bigg) + \alpha \nabla_{\theta} \mathbb{E}_{\pi_{\theta(t)}} \bigg[\sum_{\tau} \gamma^{\tau} r\bigg(\tau\bigg)\bigg].$$

The advantage of policy gradient methods is their ability to generate actions from a continuous action space by utilizing a parameterized policy.

Finally, actor-critic methods (Barto et al., 1983; Witten, 1977) make use of both the value functions and policy gradients:

$$\theta(t+1) = \theta(t) + \alpha \nabla_{\theta} \mathbb{E}_{\pi_{\theta(t)}} [Q_{\pi_{\theta(t)}}(s, a)].$$

Actor-critic methods are able to produce actions in a continuous action space, while reducing the high variance of the policy gradients by adding a critic (value function).

All of these methods aim to update the parameters θ (or directly update the policy π for critic-only methods) to improve the policy. In deep reinforcement learning, the policy π is defined by a deep neural network, whose weights constitute the parameter θ . To develop a real-time policy for our MDP, we adopt a practical policy gradient method called Proximal Policy Optimization (PPO).

4.2.2. Proximal Policy Optimization

PPO is a practical policy gradient method developed in (Schulman et al., 2017), and is effective for optimizing large non-linear policies such as deep neural networks. It preserves some of the benefits of trust region policy optimization (TRPO) (Schulman et al., 2015) such as monotonic improvement, but is much simpler to implement because it can be optimized by a first-order optimizer, and is empirically shown to have better sample complexity.

In TRPO, an objective function (the "surrogate" objective) is maximized subject to a constraint on the size of the policy update so that the new policy is not too far from the old policy:

$$\underset{\theta}{\text{maximize}} \ \widehat{\mathbb{E}}_t \left[\frac{\pi_{\theta}(\boldsymbol{a}_t | \mathbf{s}_t)}{\pi_{\theta_{out}}(\boldsymbol{a}_t | \mathbf{s}_t)} \widehat{A}_t \right]$$
(8a)

subject to
$$\widehat{\mathbb{E}}_{t}\left[\mathrm{KL}\left[\pi_{\theta_{old}}(\cdot|s_{t}),\pi_{\theta}(\cdot|s_{t})\right]\right]\leqslant\delta$$
, (8b)

where π_{θ} is a stochastic policy and \widehat{A}_t is an estimator of the advantage function at timestep t. The expectation $\widehat{\mathbb{E}}_t[...]$ indicates the empirical average over a finite batch of samples and $\mathrm{KL}\left[\pi_{\theta_{old}}(\cdot|s_t),\pi_{\theta}(\cdot|s_t)\right]$ denotes the Kullback–Leibler divergence between $\pi_{\theta_{old}}$ and π . Although TRPO solves the above constrained maximization problem using conjugate gradient, the theory justifying TRPO actually suggests using a penalty instead of a constraint, i.e., solving the unconstrained optimization problem

$$\underset{\theta}{\text{maximize}} \ \widehat{\mathbb{E}}_{t} \left[\frac{\pi_{\theta}(\boldsymbol{a}_{t}|\boldsymbol{s}_{t})}{\pi_{\theta_{old}}(\boldsymbol{a}_{t}|\boldsymbol{s}_{t})} \widehat{A}_{t} - \beta \text{KL} \left[\pi_{\theta_{old}}(\cdot|\boldsymbol{s}_{t}), \pi_{\theta}(\cdot|\boldsymbol{s}_{t}) \right] \right], \tag{9}$$

for some penalty coefficient β . TRPO uses a hard constraint rather than a penalty because it is hard to choose a single value of β that performs well. To overcome this issue and develop a first-order algorithm that emulates the monotonic improvement of TRPO (without solving the constrained optimization problem), two PPO algorithms are constructed by: 1) clipping the surrogate objective and 2) using adaptive KL penalty coefficient (Schulman et al., 2017).

1. Clipped Surrogate Objective: Let $r_t(\theta)$ denote the probability ratio $r_t(\theta) = \frac{\pi_{\theta}(a_t|s_t)}{\pi_{\theta,let}(a_t|s_t)}$, so $r(\theta_{old}) = 1$. TRPO maximizes

$$L(\theta) = \widehat{\mathbb{E}}_t \left[\frac{\pi_{\theta}(\boldsymbol{a}_t | \boldsymbol{s}_t)}{\pi_{\theta_{obst}}(\boldsymbol{a}_t | \boldsymbol{s}_t)} \widehat{A}_t \right] = \widehat{\mathbb{E}}_t \left[r_t(\theta) \widehat{A}_t \right]. \tag{10}$$

subject to the KL divergence constraint. Without a constraint however this would lead to a large policy update. To prevent this, PPO modifies the surrogate objective to penalize changes to the policy that move $r_t(\theta)$ away from 1:

$$L^{CLIP}(\theta) = \widehat{\mathbb{E}}_t \Big[\min \Big(r_t(\theta) \widehat{A}_t, \operatorname{clip} \Big(r_t(\theta), 1 - \epsilon, 1 + \epsilon \Big) \widehat{A}_t \Big) \Big], \tag{11}$$

where ϵ is a hyperparameter, usually 0.1 or 0.2. The term $\operatorname{clip}(r_t(\theta), 1 - \epsilon, 1 + \epsilon) \widehat{A}_t$) modifies the surrogate objective by clipping the probability ratio, which removes the incentive for moving r_t outside of the interval $[1 - \epsilon, 1 + \epsilon]$. By taking the minimum of the clipped and the unclipped objective, the final objective becomes a lower bound on the unclipped objective.

- 2. Adaptive KL Penalty Coefficient: Another approach is to use a penalty on KL divergence and to adapt the penalty coefficient so that some target value of the KL divergence d_{targ} is achieved at each policy update. In each policy update, the following steps are performed:
 - Using several epochs of minibatch SGD, optimize the KL-penalized objective

$$L^{KLPEN}\left(\theta\right) = \widehat{\mathbb{E}}_{t} \left[\frac{\pi_{\theta}(\boldsymbol{a}_{t}|\boldsymbol{s}_{t})}{\pi_{\theta_{old}}(\boldsymbol{a}_{t}|\boldsymbol{s}_{t})} \widehat{A}_{t} - \beta KL \left[\pi_{\theta_{old}}(\cdot|\boldsymbol{s}_{t}), \pi_{\theta}(\cdot|\boldsymbol{s}_{t}) \right] \right]$$

$$(12)$$

- Compute $d = \widehat{\mathbb{E}}_t [\text{KL} [\pi_{\theta_{old}}(\cdot|s_t), \pi_{\theta}(\cdot|s_t)]]$
 - If $d < d_{\text{targ}}/1.5, \beta \leftarrow \beta/2$
 - If $d > d_{\text{targ}} \times 1.5, \beta \leftarrow \beta \times 2$.

The updated β is then used for the next policy update. This scheme allows β to adjust if KL divergence is significantly different than d_{targ} so that the desired KL divergence between the old and the updated policy is attained.

A PPO algorithm using fixed-length trajectory segments is summarized in Algorithm 1. Each iteration, each of N (parallel) actors collect T timesteps of data. Then the surrogate loss on these NT timesteps of data is constructed and optimized with minibatch SGD for K epochs.

Algorithm 1. PPO, Actor-Critic Style

```
\label{eq:for_interaction} \begin{aligned} & \textbf{for} \ \text{iteration} = 0, 1, 2, \dots \, \textbf{do} \\ & | \ \textbf{for} \ \text{actor} = 1, 2, \dots, N \ \textbf{do} \\ & | \ \text{Run policy} \ \pi_{\theta_{old}} \ \text{in environment for} \ T \ \text{timesteps.} \\ & | \ \text{Compute advantage estimates} \ \hat{A}_1, \dots, \hat{A}_T \\ & \ \textbf{end} \\ & \ \text{Optimize surrogate} \ L^{CLIP} \ \text{or} \ L^{KLPEN} \ \text{w.r.t.} \ \theta, \ \text{with} \ K \ \text{epochs and minibatch size} \ M \leq NT. \\ & \ \theta_{old} \leftarrow \theta \end{aligned}
```

In this work, we used the PPO algorithm with the clipped surrogate objective, because experimentally it it shown to have better performance than the PPO algorithm with adaptive KL penalty coefficient (Schulman et al., 2017). We refer the reader to (Schulman et al., 2017) for a comprehensive study on PPO algorithms.

In the next section, we present our numerical studies demonstrating the performance of the RL policy.

5. Numerical Study

In this section, we discuss the numerical experiments and results for the performance of reinforcement learning approach to the dynamic problem and compare with the performance of several static policies, including the optimal static policy outlined in Section 3. We solved for the optimal static policy using CVX, a package for specifying and solving convex programs (Grant and Boyd, 2014). To implement the dynamic environment compatible with reinforcement learning algorithms, we used Gym toolkit (Brockman et al., 2016) developed by OpenAI to create an environment. For the implementation of the PPO algorithm, we used Stable Baselines toolkit (Hill et al., 2018).

We chose an operational cost of $\beta = \$0.1$ (by normalizing the average price of an electric car over 5 years (The average electric car in the US is getting cheaper)) and maximum willingness to pay $\ell_{\text{max}} = \$30$. For prices of electricity $p_i(t)$, we generated random prices for different locations and different times using the statistics of locational marginal prices in (http://oasis.caiso.com). We chose a maximum battery capacity of 20kWh. We discretrized the battery energy into 5 units, where one unit of battery energy is 4kWh. The time it takes to deliver one unit of charge is taken as one time epoch, which is equal to 5 min in our setup. The waiting time cost for one period is w = \$2 (average hourly wage is around \\$24 in the United States (United States Average Hourly Wages)).

Note that the dimension of the state space grows significantly with battery capacity v_{max} , because it expands the states each vehicle can have by v_{max} . Therefore, for computational purposes, we conducted two case studies: 1) Non-electric AMoD case study with a larger network in Manhattan, 2) Electric AMoD case study with a smaller network in San Francisco. We picked two different real world networks in order to demonstrate the universality of reinforcement learning method in establishing a real-time policy. In particular, our intention is to support the claim that the success of the reinforcement learning method is not restricted to a single network, but generalizes to multiple real world networks. Both experiments were performed on a laptop computer with Intel® CoreTM i7-8750H CPU (6 × 2.20 GHz) and 16 GB DDR4 2666 MHz RAM.

5.1. Case Study in Manhattan

In a non-electric AMoD network, the energy dimension ν vanishes. Because there is no charging action⁷, we can perform coarser discretizations of time. Specifically, we can allow each discrete time epoch to cover $5 \times \min_{i,j|i\neq j} \tau_{ij}$ minutes, and normalize the travel times τ_{ij} and w accordingly (For EV's, because charging takes a non-negligible but shorter time than travelling, in general we have $\tau_{ij} > 1$, and larger number of states). The static profit maximization problem in (1) for AMoD with non-electric vehicles can be rewritten as:

$$\max_{x_{ij},\ell_{ij}} \sum_{i \in \mathscr{M}_{j} \in \mathscr{M}} \lambda_{ij} \ell_{ij} \left(1 - F(\ell_{ij}) \right) - \beta_{g} \sum_{i \in \mathscr{M}_{j} \in \mathscr{M}} x_{ij} \tau_{ij}$$
subject to
$$\lambda_{ij} \left(1 - F(\ell_{ij}) \right) \leq x_{ij} \quad \forall i, j \in \mathscr{M},$$

$$\sum_{j \in \mathscr{M}} x_{ij} = \sum_{j \in \mathscr{M}} x_{ji} \quad \forall i \in \mathscr{M},$$

$$x_{ii} \geq 0 \, \forall i, j \in \mathscr{M}.$$
(13)

The operational costs $\beta_g = \$2.5$ (per 10 min, (How much does driving your car cost)) are different than those of electric vehicles. Because there is no "charging" (or refueling action, since it takes negligible time), β_g also includes fuel cost. The optimal static policy is used to compare and highlight the performance of the real-time policy.

We divided Manhattan into 10 regions as in Fig. 4, and using the yellow taxi data from the New York City Taxi and Limousine Commission dataset (https://www1.nyc.gov/site/tlc/about/tlc-trip-record-data.page) for May 04, 2019, Saturday between 18.00–20.00, we extracted the average arrival rates for rides and trip durations between the regions (we exclude the rides occurring in the same region). We trained our model by creating new induced random arrivals with the same average arrival rate using prices determined by our policy. For the fleet size, we used a fleet of 1200 autonomous vehicles (according to the optimal fleet size emerging from the static problem).

For training, we used a neural network with 4 hidden layers and 128 neurons in each hidden layer. The rest of the parameters are left as default as specified by the Stable Baselines toolkit (Hill et al., 2018). In order to get the best policy, we train 3 different models using DDPG (Lillicrap et al., 2015), TRPO (Schulman et al., 2015), and PPO. We trained the models for 10 million iterations, and the performances of the trained models are summarized in Table 1 using average rewards and queue lengths as metrics. Our experiments indicate that the model trained using PPO is performing the best among the three, hence we use that model as our real-time policy.

We compare different policies' performance using the rewards and total queue length as metrics. The results are demonstrated in Fig. 5. In Fig. 5a we compare the rewards generated and the total queue length by applying the static and the real-time policies as defined in Sections 3 and 4. We can observe that while the optimal static policy provides rate stability in a dynamic setting (since the queues do not blow up), it fails to generate profits as it is not able to clear the queues. On the other hand, the real-time policy is able to keep the total length of the queues 100 times shorter than the static policy while generating higher profits.

The optimal static policy fails to generate profits and is not necessarily the best static policy to apply in a dynamic setting. As such, in Fig. 5b we demonstrate the performance of a sub-optimal static policy, where the prices are 5% higher than the optimal static prices to reduce the arrival rates and hence reduce the queue lengths. Observe that the profits generated are higher than the profits generated using optimal static policy for the static planning problem while the total queue length is less. This result indicates that under the stochasticity of the dynamic setting, a sub-optimal static policy can perform better than the optimal static policy. Furthermore, we summarize the performances of other static policies with higher static prices, namely with 5%, 10%, 20%30%, and 40% higher prices than the optimal static prices in Table 2. Among these, an increase of 10% performs the best in terms of rewards. Nevertheless, this

⁷ The vehicles still refuel, however this takes negligible time compared to the trip durations.

⁸ The solution of the static problem yields vehicle flows. In order to make the policy compatible with our environment and to generate integer actions that can be applied in a dynamic setting, we randomized the actions by dividing each flow for OD pair (i,j) (and energy level ν) by the total number of vehicles in i (and energy level ν) and used that fraction as the probability of sending a vehicle from i to j (with energy level ν).



Fig. 4. Manhattan divided into m = 10 regions.

 Table 1

 Performances of RL policies trained with different algorithms.

	Algorithms		
Metrics	DDPG	TRPO	PPO
Average Rewards Average Queue Length	9825.69 431.76	13142.47 87.96	15527.34 68.11

policy does still do worse in terms of rewards and total queue length compared to the real-time policy, which generates around 10% more rewards and results in 70% less queues. Lastly we note that although a 40% increase in prices results in minimum average queue length, this is a result of significantly reduced induced demand and therefore it generates very low rewards.

Next, we showcase that even some heuristic modifications which resemble what is done in practice can do better than the optimal static policy. We utilize the optimal static policy, but additionally utilize a surge-pricing policy. The surge-pricing policy aims to decrease the arrival rates for longer queues so that the queues will stay shorter and the rewards will increase. At each time period, for all OD pairs, the policy is to increase the price by 50% if the queue is longer than 100% of the induced arrival rate. The results are displayed in Fig. 5c. New arrivals bring higher revenue per person and the total queue length is decreased, which stabilizes the network while generating more profits than the optimal static policy. The surge pricing policy results in stable short queues and higher rewards compared to the optimal static policy for the static setting, however, both the real-time policy and the static pricing policy with 10% higher prices are superior. Performances of other surge pricing policies that multiply the prices by 1.25/1.5/2 if the queue is longer than 50%/100%/200% of the induced arrival rates can be found in Table 3. Accordingly, the best surge pricing policy maximizing the rewards is to multiply the prices by 1.25 if the queue is longer than 50% of the induced arrival rate. Yet, our real-time policy still generates around 20% more rewards and results in 32% less queues. We note that a surge pricing policy that multiplies the prices by 2 when the queues are longer than 50% of the induced arrival rates minimizes the queues by decreasing the induced arrival rates significantly, which results in substantially low rewards.

Finally, we test how the static and the real-time policies are robust to variations in input statistics. We compare the rewards generated and the total queue length applying the static and the real-time policies for the arrival rates of May 11, 2019, Saturday between 18.00–20.00. The results are displayed in Fig. 5d. Even though the arrival rates between May 11 and May 4 do not differ much, the static policy is not resilient and fails to stabilize when there is a slight change in the network. The real-time policy, on the other hand, is still able to stabilize the network and generate profits. The neural-network based policy is able to determine the correct pricing and routing decisions by considering the current state of the network, even under different arrival rates.

These experiments show us that we can indeed develop a real-time policy using deep reinforcement learning and this policy is resilient to small changes in the network parameters. The next study investigates the idea of generality, i.e., whether we can develop a global real-time policy and fine-tune it to a specific environment with *few-shots* of training, rather than developing a new policy from scratch.

Few-shot Learning: A common problem with reinforcement learning approaches is that because the agent is trained for a specific environment, it fails to respond to a slightly changed environment. Hence, one would need to train a different model for different environments (different network configurations, different arrival rates). However, this is not a feasible solution considering that training one model takes millions of iterations. As a more tractable solution, one could train a global model using different environments, and then calibrate it to the desired environment with fewer iterations rather than training a new model from scratch. We tested this phenomenon by training a global model for Manhattan using various arrival rates and network configurations that we extracted from different 2-h intervals (We trained the global model for 10 million iterations). We then trained this model for the network configuration and arrival rates on May 6, 2019, Monday between 15.00–17.00. The results are displayed in Fig. 6. Even with no additional training, the global model performs better than the specific model trained from scratch for 2 million iterations. Furthermore, with only few iterations, it is possible to improve the performance of the global model significantly. This is an anticipated

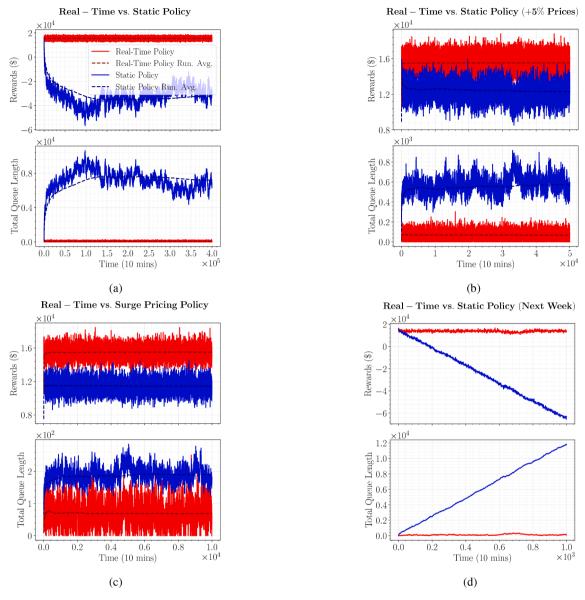


Fig. 5. Comparison of different policies for Manhattan case study. The legends for all figures are the same as the top left figure, where red lines correspond to the real-time policy and blue lines correspond to the static policies (We excluded the running averages for (d), because the static policy diverges). In all scenarios, we use the rewards generated and the total queue length as metrics. In (a), we demonstrate the results from applying the real-time policy and the optimal static policy. In (b), we compare the real-time policy with the static policy that utilizes 5% higher prices than the optimal static policy. In (c), we utilize a surge pricing policy along with the optimal static policy and compare with the real-time policy. In (d), we employ the real-time policy and static policy developed for May 4, 2019, Saturday for the arrivals on May 11, 2019, Saturday.

Table 2Performances of static pricing policies for Manhattan case study.

		% of opt. static prices					
Metrics	105%	110%	120%	130%	140%		
Average Rewards Average Queue Length	12234.13 584.05	14112.77 231.93	13739.35 74.64	12046.91 30.88	9625.82 14.20		

result, because although the network configurations and arrival rates for different 2-h intervals are different, the environments are not fundamentally different (the state transitions are governed by similar random processes) and hence it is possible to generalize a global policy and fine-tune it to the desired environment with fewer number of iterations.

Table 3Performances of surge pricing policies for Manhattan case study.

Surge Multiplier		Queue Threshold					
	Ę	50%		100%		200%	
	Queue	Rewards	Queue	Rewards	Queue	Rewards	
1.25	101.25	13022.83	186.56	12897.30	380.34	12357.33	
1.5	91.89	12602.90	178.22	12589.71	370.18	12233.95	
2	83.15	5272.04	162.99	6224.69	337.01	7485.75	

Table 4Performances of surge pricing policies for San Francisco case study.

		Queue Threshold					
Surge Multiplier	5	50%		100%		200%	
	Queue	Rewards	Queue	Rewards	Queue	Rewards	
1.25	67.62	718.66	75.92	715.02	99.56	687.45	
1.5	25.16	650.90	34.32	687.71	49.94	708.38	
2	14.06	331.21	20.55	455.25	44.44	611.23	

Few – shot Learning

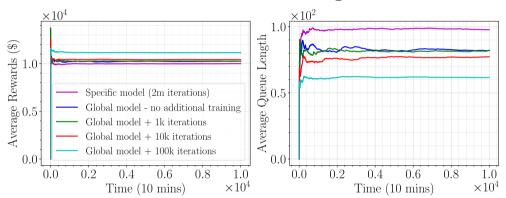


Fig. 6. Performances of the specific model that is trained from scratch and fine-tuned global model (for different amounts of fine-tuning as specified in the legend): rewards (left) and queue lengths (right).

5.2. Case Study in San Francisco

We conducted the case study in San Francisco by utilizing an EV fleet of 420 vehicles. We divided San Francisco into 7 regions as in Fig. 7, and using the traceset of mobility of taxi cabs data from CRAWDAD (Piorkowski et al., 2009), we obtained the average arrival rates and travel times between regions (we exclude the rides occurring in the same region).

In Fig. 8, we compare the charging costs paid under the real-time policy and the static policy. The static policy is generated by using

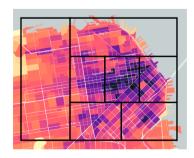


Fig. 7. San Francisco divided into m = 7 regions. Map obtained from the San Francisco County Transportation Authority (http://tncstoday.sfcta.org/).

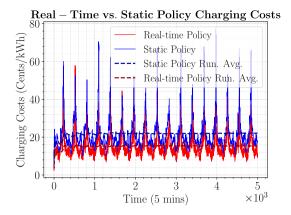


Fig. 8. Charging costs for the optimal static policy and the real-time policy in San Francisco case study.

the average value of the electricity prices, whereas the real-time policy takes into account the current electricity prices before executing an action. Therefore, the real-time policy provides cheaper charging options by utilizing smart charging strategies, decreasing the average charging costs by 25%.

In Fig. 9a, we compare the rewards and the total queue length resulting from the real-time policy and the static policy. In Fig. 9b, we compare the RL policy to the static policy with 5% higher prices than the optimal static policy, and summarize performances of several other static pricing policies in Table 5. In Fig. 9c, we use the static policy but also utilize a surge pricing policy that multiplies the prices by 1.5 if the queues are longer than 100% of the induced arrival rates. The performances of other surge pricing policies are also displayed in Table 4. Similar to the case study in Manhattan, the results demonstrate that the performance of the trained real-time policy is superior to the other policies. In particular, the RL policy is able to generate around 24% more rewards and result in around 75% less queues than the best heuristic policy, which utilizes 30% higher static prices than the optimal static policy.

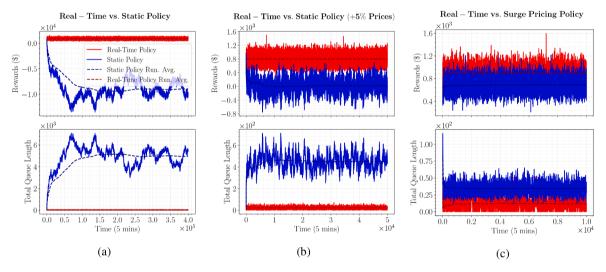


Fig. 9. Comparison of different policies for San Francisco case study. The legends for all figures are the same as the top left figure, where red lines correspond to the real-time policy and blue lines correspond to the static policies. In all scenarios, we use the rewards generated and the total queue length as metrics. In (a), we demonstrate the results from applying the real-time policy and the optimal static policy. In (b), we compare the real-time policy with a sub-optimal static policy, where the prices are 5% higher than the optimal static policy. In (c), we utilize a surge pricing policy along with the optimal static policy and compare with the real-time policy.

Table 5Performances of static pricing policies for San Francisco case study.

		% of opt. static prices				
Metrics	105%	110%	120%	130%	140%	
Average Rewards Average Queue Length	4.98 456.83	485.65 211.04	696.38 87.15	721.89 45.28	682.76 25.66	

6. Conclusion

In this paper, we developed a real-time control policy based on deep reinforcement learning for operating an AMoD fleet of EVs as well as pricing for rides. Our real-time control policy jointly makes decisions for: (1) vehicle routing in order to serve passenger demand and to rebalance the empty vehicles, (2) vehicle charging in order to sustain energy for rides while exploiting geographical and temporal diversity in electricity prices for cheaper charging options, and (3) pricing for rides in order to adjust the potential demand so that the network is stable and the profits are maximized. Furthermore, we formulated the static planning problem associated with the dynamic problem in order to define the optimal static policy for the static planning problem. When implemented correctly, the static policy provides stability of the queues in the dynamic setting, yet it is not optimal regarding the profits and keeping the queues sufficiently low. Finally, we conducted case studies in Manhattan and San Francisco that demonstrate the performance of our developed policy. The two case studies on different networks indicate that reinforcement learning can be a universal method for establishing well performing real-time policies that can be applied to many real world networks. Lastly, by doing the Manhattan study with non-electric vehicles and San Francisco study with electric vehicles, we have also demonstrated that a real-time policy using reinforcement learning can be established for both electric and non-electric AMoD systems.

Appendix A. Proof of Proposition 1

To prove Proposition 1, we first formulate the static optimization problem via a network flow model that characterizes the *capacity region* of the network for a given set of prices $\ell_{ij}(t) = \ell_{ij} \ \forall t$ (Hence, $\Lambda_{ij}(t) = \Lambda_{ij} \ \forall t$). The capacity region is defined as the set of all arrival rates $[\Lambda_{ij}]_{i,j\in\mathscr{M}}$, where there exists a charging and routing policy under which the queueing network of the system is stable. Let x_i^v be the number of vehicles available at node i, α_{ij}^v be the fraction of vehicles at node i with energy level v being routed to node i, and α_{ic}^v be the fraction of vehicles charging at node i starting with energy level v. We say the static vehicle allocation for node i and energy level v is feasible if $\alpha_{ic}^v + \sum_{j\in\mathscr{M}} \alpha_{ij}^v \leqslant 1$.

The optimization problem that characterizes the capacity region of the network ensures that the total number of vehicles routed from i to j is at least as large as the nominal arrival rate to the queue (i,j). Namely, the vehicle allocation problem can be formulated as follows:

$$\min_{x_{i}^{\mu}, \alpha_{i}^{\mu}, \alpha_{i}^{\mu}} \rho \tag{14a}$$

s.t.
$$\Lambda_{ij} \leq \sum_{v=v_{ij}}^{v_{\text{max}}} x_i^v \alpha_{ij}^v \quad \forall i, j \in \mathcal{M},$$
 (14b)

$$\rho \geqslant \alpha_{ic}^{v} + \sum_{j \in \mathcal{M}} \alpha_{ij}^{v} \quad \forall i \in \mathcal{M}, \ \forall v \in \mathcal{V},$$

$$\tag{14c}$$

$$x_i^{\nu} = x_i^{\nu-1} \alpha_{ic}^{\nu-1} + \sum_{i \in \mathcal{M}} x_i^{\nu+\nu_{ji}} \alpha_{ji}^{\nu+\nu_{ji}} \quad \forall i \in \mathcal{M}, \ \forall \nu \in \mathcal{V},$$

$$\tag{14d}$$

$$a_{i_{c}}^{v_{max}}=0 \quad \forall i \in \mathscr{M},$$
 (14e)

$$\alpha_{ij}^v = 0 \quad \forall v < v_{ii}, \ \forall i, j \in \mathcal{M},$$
 (14f)

$$x_i^v \geqslant 0, \ \alpha_{ii}^v \geqslant 0, \ \alpha_{iv}^v \geqslant 0, \ \forall i,j \in \mathcal{M}, \ \forall v \in \mathcal{V},$$
 (14g)

$$x_i^v = a_{iv}^v = a_{ii}^v = 0 \quad \forall v \notin \mathscr{V}, \, \forall i, j \in \mathscr{M}.$$
 (14h)

The constraint (14c) upper bounds the allocation of vehicles for each node i and energy level v. The constraints (14d)–(14f) are similar to those of optimization problem (1) with $x_i^v = x_{ic}^v + \sum_{i \in \mathcal{M}} x_{ij}^v, \alpha_{ic}^v = x_{ic}^v / x_i^v$, and $\alpha_{ij}^v = x_{ij}^v / x_i^v$.

Lemma 1. Let the optimal value of (14) be ρ^* . Then, $\rho^* \leq 1$ is a necessary and sufficient condition of rate stability of the system under some routing and charging policy.

Proof. Consider the fluid scaling of the queueing network, $Q_{ij}^{rt} = \frac{q_{ij}(\lfloor rt \rfloor)}{r}$ (see (Dai, 1995) for more discussion on the stability of fluid models), and let Q_{ij}^{t} be the corresponding fluid limit. The fluid model dynamics is as follows:

$$Q_{ij}^{t} = Q_{ij}^{0} + A_{ij}^{t} - X_{ij}^{t},$$

where A^t_{ij} is the total number of riders from node i to node j that have arrived to the network until time t and X^t_{ij} is the total number of vehicles routed from node i to j up to time t. Suppose that $\rho^* > 1$ and there exists a policy under which for all $t \geqslant 0$ and for all origin–destination pairs $(i,j), Q^t_{ij} = 0$. Pick a point t_1 , where $Q^{t_1}_{ij}$ is differentiable for all (i,j). Then, for all $(i,j), \dot{Q}^{t_1}_{ij} = 0$. Since $\dot{A}^{t_1}_{ij} = \Lambda_{ij}$, this

implies $\dot{X}^{t_1}_{ij} = \Lambda_{ij}$. On the other hand, $\dot{X}^{t_1}_{ij}$ is the total number of vehicles routed from i to j at t_1 . This implies $\Lambda_{ij} = \sum_{\nu=\nu_0}^{\nu_{\max}} x_i^{\nu} \alpha_{ij}^{\nu}$ for all (i,j) and there exists α^{ν}_{ij} and α^{ν}_{ic} at time t_1 such that the flow balance constraints hold and the allocation vector $[\alpha^{\nu}_{ij} \ \alpha^{\nu}_{ic}]$ is feasible, i.e. $\alpha^{\nu}_{ic} + \sum_{j=1}^{m} \alpha^{\nu}_{ij} \leqslant 1$. This contradicts $\rho^* > 1$.

Now suppose $\rho^* \leqslant 1$ and $\alpha^* = [\alpha_{ij}^{v^*} \ \alpha_{ij}^{v^*}]$ is an allocation vector that solves the static problem. The cumulative number of vehicles routed from node i to j up to time t is $S_{ij}^t = \sum_{\nu=0}^{\nu_{\max}} x_i^{\nu} \alpha_{ij}^{\nu} t = \sum_{\nu=0}^{\nu_{\max}} x_i^{\nu} \alpha_{ij}^{\nu} t > \Lambda_{ij} t$. Suppose that for some origin–destination pair (i,j), the queue $Q_{ij}^{t_1} \geqslant \epsilon > 0$ for some positive t_1 and ϵ . By continuity of the fluid limit, there exists $t_0 \in (0,t_1)$ such that $Q_{ij}^{t_0} = \epsilon/2$ and $Q_{ij}^t > 0$ for $t \in [t_0,t_1]$. Then, $\dot{Q}_{ij}^t > 0$ implies $\Lambda_{ij} > \sum_{\nu=\nu_0}^{\nu_{\max}} x_i^{\nu} \alpha_{ij}^{v}$, which is a contradiction. \square

By Lemma 1, the *capacity region* C_{Λ} of the network is the set of all $\Lambda_{ij} \in \mathbb{R}^+$ for which the corresponding optimal solution to the optimization problem (14) satisfies $\rho^* \leq 1$. As long as $\rho^* \leq 1$, there exists a routing and charging policy such that the queues will be bounded away from infinity.

The platform operator's goal is to maximize its profits by setting prices and making routing and charging decisions such that the system remains stable. In its most general form, the problem can be formulated as follows:

$$\max_{\ell_{ij}, x_i^{\nu}, \alpha_{ij}^{\nu}, \alpha_{ic}^{\nu}} \quad U\left(\Lambda_{ij}(\ell_{ij}), x_i^{\nu}, \alpha_{ic}^{\nu}, \alpha_{ic}^{\nu}\right)$$
subject to
$$\left[\Lambda_{ij}(\ell_{ij})\right]_{i,i \in \mathcal{M}} \in C_{\Lambda},$$
(14i)

where $U(\cdot)$ is the utility function that depends on the prices, demand for rides and the vehicle decisions.

Recall that $x_{ic}^{\nu} = x_i^{\nu} a_{ic}^{\nu}$ and $x_{ij}^{\nu} = x_i^{\nu} a_{ij}^{\nu}$. Using these variables and noting that $a_{ic}^{\nu} + \sum_{j \in \mathcal{M}} a_{ij}^{\nu} = 1$ when $\rho^* \leqslant 1$, the platform operator's profit maximization problem can be stated as (1). A feasible solution of (1) guarantees rate stability of the system, since the corresponding vehicle allocation problem (14) has solution $\rho^* \leqslant 1$.

Appendix B. Proof of Proposition 2

For brevity of notation, let $\beta + p_i = P_i$. Let ν_{ij} be the dual variables corresponding to the demand satisfaction constraints and μ_i^v be the dual variables corresponding to the flow balance constraints. Since the optimization problem (1) is a convex quadratic maximization problem (given a with uniform $F(\cdot)$) and Slater's condition is satisfied, strong duality holds. We can write the dual problem as:

$$\underset{\nu_{ij}, \mu_i^*}{\text{minmax}} \sum_{i=1}^{m} \sum_{i=1}^{m} \left(\lambda_{ij} \left(1 - \frac{\ell_{ij}}{\ell_{\text{max}}} \right) (\ell_i - \nu_{ij}) \right) \tag{16a}$$

s.t.
$$\nu_{ij} \geqslant 0$$
, (16b)

$$\nu_{ii} + \mu_i^v - \mu^{v-v_{ij}} - \beta \tau_{ii} \leq 0, \tag{16c}$$

$$\mu_i^v - \mu_i^{v+1} - P_i \leqslant 0 \quad \forall i, j, v.$$
 (16d)

For fixed ν_{ij} and μ_i^{ν} , the inner maximization results in the optimal prices:

$$\ell_{ij}^* = \frac{\ell_{\text{max}} + \nu_{ij}}{2}. \tag{17}$$

By strong duality, the optimal primal solution satisfies the dual solution with optimal dual variables ν_{ij}^* and $\mu_i^{\nu^*}$, which completes the first part of the proposition. The dual problem with optimal prices in (17) can be written as:

$$\min_{\nu_{ij},\mu_{i}^{r}} \sum_{i=1}^{m} \sum_{j=1}^{m} \frac{\lambda_{ij}}{\ell_{\max}} \left(\frac{\ell_{\max} - \nu_{ij}}{2} \right)^{2}$$
(18a)

subject
$$to\nu_{ii} \geqslant 0$$
, (18b)

$$\nu_{ij} + \mu_i^{\nu} - \mu_i^{\nu-\nu_{ij}} - \beta \tau_{ij} \leqslant 0, \tag{18c}$$

$$\mu_i^{\nu} - \mu_i^{\nu+1} - P_i \leqslant 0 \quad \forall i, j, \nu.$$
 (18d)

The objective function in (18a) with optimal dual variables, along with (17) suggests:

$$P = \sum_{i=1}^{m} \sum_{i=1}^{m} \frac{\lambda_{ij}}{\ell_{\text{max}}} \left(\ell_{\text{max}} - \ell_{ij}^{*}\right)^{2},$$

where profits *P* is the value of the objective function of both optimal and dual problems. To get the upper bound on prices, we go through the following algebraic calculations using the constraints. The inequality (18d) gives:

$$\mu_{\nu}^{\nu-\nu\mu} \leqslant \nu_{ii} P_i + \mu_{\nu}^{\nu}, \tag{19}$$

and equivalently:

$$\mu_i^{\nu_i \nu_j} \le \nu_{ii} P_i + \mu_i^{\nu}$$
. (20)

The inequalities (18c) and (18b) yield:

$$\mu_i^{\nu} - \mu_i^{\nu - \nu_{ij}} - \beta \tau_{ij} \leq 0,$$

and equivalently:

$$\mu_i^{\mathrm{v}} - \mu_i^{\mathrm{v} - \mathrm{v}_{ji}} - \beta \tau_{ji} \leqslant 0, \tag{21}$$

Inequalities (19) and (21):

$$\mu_i^v \leqslant \mu_i^v + \beta \tau_{ji} + \nu_{ji} P_i. \tag{22}$$

And finally, the constraint (18c):

$$\nu_{ij} \leq \beta \tau_{ij} + \mu_i^{v-v_{ij}} - \mu_i^{v} \stackrel{20}{\leq} {}^{(20)}\beta \tau_{ij} + v_{ij}P_j + \mu_i^{v} - \mu_i^{v} \stackrel{22}{\leq} {}^{(22)}\beta \tau_{ij} + v_{ij}P_j + \beta \tau_{ji} + v_{ji}P_i.$$

Replacing $P_i = p_i + \beta$ and rearranging the terms:

$$\nu_{ij} \leqslant \beta(\tau_{ij} + \tau_{ii} + \nu_{ij} + \nu_{ij}) + \nu_{ij}p_j + \nu_{ij}p_j. \tag{23}$$

Using the upper bound on the dual variables ν_{ii} and (17), we can upper bound the optimal prices.

References

Barto, A.G., Sutton, R.S., Anderson, C.W., 1983. Neuronlike adaptive elements that can solve difficult learning control problems. IEEE Trans Syst., Man, Cybernet., vol. SMC-13, no. 5, pp. 834–846, Sep. 1983.

Bimpikis, K., Candogan, O., Sabán, D., 2019. Spatial pricing in ride-sharing networks. Oper. Res. 67, 744-769.

Bongiovanni, C., Kaspi, M., Geroliminis, N., 2019. "The electric autonomous dial-a-ride problem." Transport. Res. Part B: Methodol., vol. 122, pp. 436–456. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0191261517309669.

Boyaci, B., Zografos, K.G., Geroliminis, N., 2017. "An integrated optimization-simulation framework for vehicle and personnel relocations of electric carsharing systems with reservations." Transport. Res. Part B: Methodol., vol. 95, pp. 214–237, 2017. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0191261515301119.

Brockman, G., Cheung, V., Pettersson, L., Schneider, J., Schulman, J., Tang, J., Zaremba, W., 2016. Openai gym.

Chen, T.D., Kockelman, K.M., 2016. Management of a shared autonomous electric vehicle fleet: Implications of pricing schemes. Transp. Res. Rec. 2572 (1), 37–46. Chen, T.D., Kockelman, K.M., Hanna, J.P., 2016. Operations of a Shared, Autonomous, Electric Vehicle Fleet: Implications of Vehicle & Charging Infrastructure Decisions. Transport. Res. Part A: Policy Pract. 94, 243–254.

Dai, J.G., 1995. On positive harris recurrence of multiclass queueing networks: A unified approach via fluid limit models. Annals Appl. Probab. 5, 49–77. Grant, M., Boyd, S., 2014. "CVX: Matlab software for disciplined convex programming, version 2.1," http://cvxr.com/cvx, Mar. 2014.

Grondman, I., Busoniu, L., Lopes, G.A.D., Babuska, R., Nov 2012. A survey of actor-critic reinforcement learning: Standard and natural policy gradients. IEEE Trans. Syst., Man, Cybernet. Part C (Appl. Rev.) 42 (6), 1291–1307.

Guan, Y., Annaswamy, A.M., Tseng, H.E., 2019. "Cumulative prospect theory based dynamic pricing for shared mobility on demand services." CoRR, vol. abs/1904.04824, 2019. [Online]. Available: http://arxiv.org/abs/1904.04824.

Guériau, M., Dusparic, I., 2018. "Samod: Shared autonomous mobility-on-demand using decentralized reinforcement learning." In: 2018 21st International Conference on Intelligent Transportation Systems (ITSC), Nov 2018, pp. 1558–1563.

Han, M., Senellart, P., Bressan, S., Wu, H., 2016. Routing an autonomous taxi with reinforcement learning. In: CIKM.

Hill, A., Raffin, A., Ernestus, M., Gleave, A., Traore, R., Dhariwal, P., Hesse, C., Klimov, O., Nichol, A., Plappert, M., Radford, A., Schulman, J., Sidor, S., Wu, Y., 2018. "Stable baselines," https://github.com/hill-a/stable-baselines, 2018.

How much does driving your car cost, per minute? [Online]. Available: https://www.bostonglobe.com/ideas/2014/08/08/how-much-driving-really-costs-per-minute/BqnNd2q7jETedLhxxzY2CI/story.html.

[Online]. Available: http://oasis.caiso.com.

[Online]. Available: http://tncstoday.sfcta.org/.

[Online]. Available: https://www1.nyc.gov/site/tlc/about/tlc-trip-record-data.page.

[Online]. Available: https://www.cbinsights.com/research/autonomous- driverless-vehicles-corporations-list/.

Iglesias, R., Rossi, F., Wang, K., Hallac, D., Leskovec, J., . Pavone, J., 2017. "Data-driven model predictive control of autonomous mobility-on-demand systems." CoRR, vol. abs/1709.07032, 2017. [Online]. Available: http://arxiv.org/abs/1709.07032.

Kaelbling, L.P., Littman, M.L., Moore, A.W., 1996. Reinforcement learning: A survey. J. Artif. Intell. Res. 4, 237–285.

Lazar, D.A., Bıyık, E., Sadigh, D., Pedarsani, R., 2019. "Learning how to dynamically route autonomous vehicles on shared roads," arXiv preprint arXiv:1909.03664, 2019.

Li, L., Lv, Y., Wang, F., 2016. Traffic signal timing via deep reinforcement learning. IEEE/CAA J. Autom. Sin. 3 (3), 247-254.

Lillicrap, T.P., Hunt, J.J., Pritzel, A., Heess, N., Erez, T., Tassa, Y., Silver, D., Wierstra, D., 2015. "Continuous control with deep reinforcement learning," arXiv preprint arXiv:1509.02971, 2015.

- Lin, K., Zhao, R., Xu, Z., Zhou, J., 2018. "Efficient large-scale fleet management via multi-agent deep reinforcement learning." In: Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, ser. KDD '18. New York, NY, USA: Association for Computing Machinery, 2018, p. 1774–1783. [Online]. Available: doi: 10.1145/3219819.3219993.
- Lopez, P.A., Behrisch, M., Bieker-Walz, L., Erdmann, J., Flötteröd, Y.-P., Hilbrich, R., Lücken, L., Rummel, J., Wagner, P., Wießner, E., 2018. "Microscopic traffic simulation using sumo." In: The 21st IEEE International Conference on Intelligent Transportation Systems. IEEE, 2018. [Online]. Available: https://elib.dlr.de/124092/
- Mao, C., Shen, C., 2018. "A reinforcement learning framework for the adaptive routing problem in stochastic time-dependent network." Transport. Res. Part C: Emerg. Technol., vol. 93, pp. 179–197, 2018. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0968090X18307617.
- Mao, C., Liu, Y., Shen, Z.-J.M., 2020. "Dispatch of autonomous vehicles for taxi services: A deep reinforcement learning approach." Transport. Res. Part C: Emerg. Technol., vol. 115, p. 102626, 2020. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0968090X19312227.
- Miao, F., Han, S., Lin, S., Stankovic, J.A., Huang, H., Zhang, D., Munir, S., He, T., Pappas, G.J., 2016. "Taxi dispatch with real-time sensing data in metropolitan areas: A receding horizon control approach." CoRR, vol. abs/1603.04418, 2016. [Online]. Available: http://arxiv.org/abs/1603.04418.
- Miao, F., Han, S., Hendawi, A.M., Khalefa, M.E., Stankovic, J.A., Pappas, G.J., 2017. "Data-driven distributionally robust vehicle balancing using dynamic region partitions." In: 2017 ACM/IEEE 8th International Conference on Cyber-Physical Systems (ICCPS), April 2017, pp. 261–272.
- Mukherjee, J.C., Gupta, A., Dec 2015. A review of charge scheduling of electric vehicles in smart grid. IEEE Syst. J. 9 (4), 1541-1553.
- Pavone, M., Smith, S.L., Frazzoli, E., Rus, D., 2012. Robotic load balancing for Mobility-on-Demand systems. Int. J. Robot. Res. 31 (7), 839-854.
- Pedarsani, R., Walrand, J., Zhong, Y., 2017. Robust scheduling for flexible processing networks. Adv. Appl. Probab. 49 (2), 603–628.
- Piorkowski, M., Sarafijanovic-Djukic, N., Grossglauser, M., 2009. "CRAWDAD dataset epfl/mobility (v. 2009–02-24)," Downloaded from https://crawdad.org/epfl/mobility/20090224, Feb. 2009.
- Repoux, M., Kaspi, M., Boyaci, B., Geroliminis, N., 2019. "Dynamic prediction-based relocation policies in one-way station-based carsharing systems with complete journey reservations." Transport. Res. Part B: Methodol., vol. 130, pp. 82–104, 2019. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S019126151930102X.
- Rossi, F., Iglesias, R., Alizadeh, M., Pavone, M., 2018. "On the interaction between autonomous mobility-on-demand systems and the power network: models and coordination algorithms." Robotics: Science and Systems XIV, Jun 2018.
- Rossi, F., Zhang, R., Hindy, Y., Pavone, M., 2018. Routing autonomous vehicles in congested transportation networks: Structural properties and coordination algorithms. Auton. Robots 42 (7), 1427–1442.
- Rummery, G., Niranjan, M., 1994. "On-line q-learning using connectionist systems," Technical Report CUED/F-INFENG/TR 166, 11 1994.
- Schulman, J., Levine, S., Moritz, P., Jordan, M.I., Abbeel, P., 2015. "Trust region policy optimization," CoRR, vol. abs/1502.05477, 2015. [Online]. Available: http://arxiv.org/abs/1502.05477.
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A., Klimov, O., 2017. "Proximal policy optimization algorithms," arXiv preprint arXiv:1707.06347, 2017.
- Sheppard, C.J.R., Bauer, G.S., Gerke, B.F., Greenblatt, J.B., Jenn, A.T., Gopal, A.R., 2019. "Joint optimization scheme for the planning and operations of shared autonomous electric vehicle fleets serving mobility on demand." Transport. Res. Rec., vol. 2673, no. 6, pp. 579–597, 2019. [Online]. Available: doi: 10.1177/0361198119838270.
- Spieser, K., Samaranayake, S., Frazzoli, E., 2016. "Vehicle routing for shared-mobility systems with time-varying demand." In: 2016 American Control Conference (ACC), July 2016, pp. 796–802.
- Su, W., Eichi, H., Zeng, W., Chow, M., 2012. A survey on the electrification of transportation in a smart grid environment. IEEE Trans. Industr. Inf. 8 (1), 1–10. Swaszek, R.M.A., Cassandras, C., 2019. Load balancing in mobility-on-demand systems: Reallocation via parametric control using concurrent estimation. In: 2019 IEEE Intelligent Transportation Systems Conference (ITSC), pp. 2148–2153.
- The average electric car in the US is getting cheaper. [Online]. Available: https://qz.com/1695602/the-average-electric-vehicle-is-getting-cheaper-in-the-us/. Tsao, M., Iglesias, R., Pavone, M., 2018. "Stochastic model predictive control for autonomous mobility on demand," CoRR, vol. abs/1804.11074, 2018. [Online].
- Available: http://arxiv.org/abs/1804.11074.

 Tucker, N., Turan, B., Alizadeh, M., 2019. Online Charge Scheduling for Electric Vehicles in Autonomous Mobility on Demand Fleets. In: Proc. IEEE Int. Conf. on Intelligent Transportation Systems.
- Turan, B., Tucker, N., Alizadeh, N., 2019. "Smart Charging Benefits in Autonomous Mobility on Demand Systems," In Proc. IEEE Int. Conf. on Intelligent Transportation Systems, 2019. [Online]. Available: https://arxiv.org/abs/1907.00106.
- United States Average Hourly Wages. [Online]. Available: https://tradingeconomics.com/united-states/wages.
- Veldman, E., Verzijlbergh, R.A., Jan 2015. Distribution grid impacts of smart electric vehicle charging from different perspectives. IEEE Trans. Smart Grid 6 (1),
- Volkov, M., Aslam, J., Rus, D., 2012. "Markov-based redistribution policy model for future urban mobility networks." In: Conference Record IEEE Conference on Intelligent Transportation Systems, pp. 1906–1911, 09 2012.
- Walraven, E., Spaan, M.T., Bakker, B., 2016. "Traffic flow optimization: A reinforcement learning approach." Engineering Applications of Artificial Intelligence, vol. 52, pp. 203–212, 2016. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0952197616000038.
- Warrington, J., Ruchti, D., 2019. "Two-stage stochastic approximation for dynamic rebalancing of shared mobility systems," Transport. Res. Part C: Emerg. Technol., vol. 104, pp. 110–134, 2019. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0968090X18314104.
- Watkins, C.J.C.H., Dayan, P., 1992. "Q-learning," Machine Learning, vol. 8, no. 3, pp. 279-292, May 1992. [Online]. Available: doi: 10.1007/BF00992698.
- Wei, Q., Rodriguez, J.A., Pedarsani, R., Coogan, S., 2019. "Ride-sharing networks with mixed autonomy," arXiv preprint arXiv:1903.07707, 2019.
- Wen, J., Zhao, J., Jaillet, P., 2017. "Rebalancing shared mobility-on-demand systems: A reinforcement learning approach." In: 2017 IEEE 20th International Conference on Intelligent Transportation Systems (ITSC), Oct 2017, pp. 220–225.
- Williams, R.J., 1992. "Simple statistical gradient-following algorithms for connectionist reinforcement learning." Mach. Learn., vol. 8, no. 3, pp. 229–256, May 1992. [Online]. Available: doi: 10.1007/BF00992696.
- Witten, I.H., 1977. An adaptive optimal controller for discrete-time markov environments. Inf. Control 34, 286-295.
- Zhang, R., Pavone, M., 2016. Control of robotic Mobility-on-Demand systems: A queueing-theoretical perspective. Int. J. Robot. Res. 35 (1-3), 186-203.
- Zhang, R., Rossi, F., Pavone, M., 2016. "Model predictive control of autonomous mobility-on-demand systems." In: 2016 IEEE International Conference on Robotics and Automation (ICRA), May 2016.
- Zhu, F., Ukkusuri, S.V., 2014. "Accounting for dynamic speed limit control in a stochastic traffic environment: A reinforcement learning approach," Transportation Research Part C: Emerging Technologies, vol. 41, pp. 30–47, 2014. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0968090X1400028X.
- Zhu, F., Aziz, H.A., Qian, X., Ukkusuri, S.V., 2015. "A junction-tree based learning algorithm to optimize network wide traffic control: A coordinated multi-agent framework." Transport. Res. Part C: Emerg. Technol., vol. 58, pp. 487–501, 2015, special Issue: Advanced Road Traffic Control. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0968090X14003593.