Transportation Research Part B 163 (2022) 258–280

Contents lists available at ScienceDirect

Transportation Research Part B

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

A general maximum-stability dispatch policy for shared autonomous vehicle dispatch with an analytical characterization of the maximum throughput

Michael W. Levin

University of Minnesota, Minneapolis, United States

ARTICLE INFO

Keywords:
Shared autonomous vehicles
Stability
Minimum dispatch plus penalty
Replacement ratio

ABSTRACT

Shared autonomous vehicles (SAVs) have been studied through analytical dispatch methods and simulation. A common question of interest is how many customers can be served per SAV, which necessarily depends on the network characteristics, travel demand, and dispatch policy. We identify equations that describe the maximum set of demands that could be served if an appropriate dispatch policy were chosen. We then provide a dispatch policy that achieves the predicted level of passenger throughput. This is achieved for a general class of SAV behaviors which may include ridesharing, electric SAV recharging, integration with public transit, or combinations thereof. We accomplish this by defining a Markov chain queueing model which admits general SAV behaviors. We say the network is *stable* if the head-of-line waiting times remain bounded, which is equivalent to serving all customers at the same rate at which they request service. We give equations characterizing the stable region $\boldsymbol{\Lambda}$ — the set of demands that could be served by any dispatch policy. We prove that any demand outside $\boldsymbol{\Lambda}$ cannot be completely served. We further prove that our dispatch policy stabilizes the network for any demand in the stable region using Lyapunov drift, establishing $\boldsymbol{\Lambda}$ as the maximum set of demand that can be served. Numerical results validate our calculations using simulation, and we present initial results on calculating $\boldsymbol{\Lambda}$ for a large city network.

1. Introduction

Shared autonomous vehicles (SAVs) are a fleet of automated vehicles that provide mobility-on-demand services to travelers. Currently, Waymo is offering driverless SAV services in limited locations (Gibbs, 2017). Due to the high costs of automated vehicle technology, SAVs may be one of the first widespread uses of automated vehicles (Fagnant and Kockelman, 2015). Furthermore, SAVs may achieve travel costs that are comparable to private vehicle ownership, making them a viable mode for daily commuting trips. Eventual large-scale use of SAVs is likely, which motivates study of how to efficiently dispatch SAVs to passengers and also how to predict the level-of-service that travelers can expect from widespread and frequent SAV use. These two issues are related. A sufficient fleet size is necessary for a good level-of-service, and prior work (Spieser et al., 2014; Fagnant and Kockelman, 2014; Fagnant et al., 2015; Boesch et al., 2016) has attempted to find the replacement ratio of how many travelers can effectively be served by a single SAV. However, the efficiency of the dispatch intuitively affects the level-of-service that can be obtained from a fixed SAV fleet. The goal of this paper is to characterize the set of travel demand that can be served by a given SAV fleet, and in doing so provide a dispatch policy that will serve that demand.

E-mail address: mlevin@umn.edu.

https://doi.org/10.1016/j.trb.2022.07.003
Received 28 September 2021; Received in revised form 25 May 2022; Accepted 4 July 2022
0191-2615/© 2022 Published by Elsevier Ltd.

Modeling SAV performance is made more complex by several technological extensions that affect the behavior of individual vehicles. At the basic level, an SAV travels empty from its current location to a passenger's origin, carries the passenger to their destination, then continues with the next passenger. This empty travel limits the efficiency of SAV systems. Some empty travel is necessary to reach the next passenger's origin, but time spent traveling empty is not directly spent serving passengers. Dynamic ridesharing (Fagnant and Kockelman, 2018; Alonso-Mora et al., 2017a), in which SAVs can carry multiple passengers with different origins and/or destinations simultaneously, can reduce empty travel at the cost of higher in-vehicle travel time for some passengers. SAVs can also be integrated to provide first-mile/last-mile transportation for public transit (Shen et al., 2018; Wen et al., 2018; Gurumurthy et al., 2020a), which reduces the SAV time spent per passenger by taking a passenger to/from public transit stops instead of directly from their origin to their destination. Other papers expect SAVs to be electric (Chen et al., 2016; Chen and Kockelman, 2016; Jäger et al., 2017; Loeb et al., 2018; Loeb and Kockelman, 2019; Zhang and Chen, 2020), which requires SAVs to spend significant time on recharging. All of these extensions to the basic SAV problem affect optimal dispatch behavior and models of their fleet-level performance.

1.1. Prior work on shared autonomous vehicles

SAV modeling (Narayanan et al., 2020) can mostly be characterized as optimization-based methods for SAV dispatch and agent-based simulations of SAV performance. SAVs are often assumed to be centrally dispatched, meaning that the actions of individual vehicles are coordinated at the city level to improve service. This coordination includes matching vehicles to passengers and recharging behavior. Optimizing this coordination is a major topic in the literature. Fundamentally, SAV dispatch is a dial-a-ride problem, a type of vehicle routing problem (Laporte, 1992; Toth and Vigo, 2002; Eksioglu et al., 2009), which is already known to be NP-hard (Cordeau and Laporte, 2007) and has been studied extensively in operations research. Some of the major differences between SAV research and existing dial-a-ride problem algorithms are the problem size (thousands of SAVs vs tens of dial-a-ride vehicles) and stochastic demand for SAVs. Some studies build on network flow concepts (Spieser et al., 2014; Zhang et al., 2015; Rossi et al., 2018) and network queueing models (Zhang and Pavone, 2016; Iglesias et al., 2019) in their optimization design. Others have used model predictive control due to the presence of stochastic and unknown demand (Zhang et al., 2016; Tsao et al., 2018, 2019). Optimizing the empty rebalancing of SAVs is an important sub-problem by itself (Pavone et al., 2012; Hörl et al., 2019; de Souza et al., 2020; Skordilis et al., 2021), and may be necessary due to limited parking availability (Winter et al., 2021b,a). It has been optimized separately or integrated with vehicle matching (Guo et al., 2021).

Optimal dispatch studies tend to be analytically-based, but correspondingly ignore some realistic facets of SAV behavior to reduce the problem complexity. Although large-scale SAV use is likely to affect traffic congestion, only a few studies explicitly consider congestion in their optimization (Levin, 2017; Levin et al., 2019; Salazar et al., 2019; Wollenstein-Betech et al., 2020; Li et al., 2021a), but the computational requirements of these models prevent them from being used on realistic problem sizes. A few others include congestion in their simulations but not in their optimization (Iglesias et al., 2019; Guériau et al., 2020). Chen and Levin (2019) proposed a predictive dynamic user equilibrium model of SAV trip patterns and their corresponding congestion. User equilibrium behavioral responses were later used in a bi-level problem to optimize SAV matching and routing (Ge et al., 2021) and SAV locations and fleet size (Li and Liao, 2020). SAV extensions add further complexity. Without ridesharing, vehicle miles traveled are expected to increase (Moreno et al., 2018; Tirachini and Gomez-Lobo, 2020). However, ridesharing greatly increases the complexity because each vehicle trip serves multiple passenger origins and destinations with a corresponding ordering problem (Alonso-Mora et al., 2017a,b; Cáp and Alonso Mora, 2018; Tsao et al., 2019). Electrification is beneficial for sustainability like reducing greenhouse gas emissions (Greenblatt and Shaheen, 2015; Bauer et al., 2018; Jones and Leibowicz, 2019), but adds further complexity due to SAV battery levels, recharging times, and limited recharging locations (Iacobucci et al., 2019; Boewing et al., 2020).

Other studies have used agent-based simulation to evaluate the fleet-level performance. Agent-based simulation consists of creating "agents" that represent individual vehicles and passengers. These agents exist within a simulation environment that admits certain forms of interaction (e.g. an SAV agent can pick up a passenger agent at the same location). Using agent-based simulation, several studies found that 3–10 personal vehicles could be replaced by 1 SAV (Spieser et al., 2014; Fagnant and Kockelman, 2014; Fagnant et al., 2015; Boesch et al., 2016). Although initial studies used simple simulation environments, e.g. dividing up the simulated region geographically into squares (Fagnant and Kockelman, 2014; Fagnant et al., 2015; Fagnant and Kockelman, 2018), later studies simulate SAVs traveling through traffic networks, sometimes with dynamic congestion (Levin et al., 2017; Gurumurthy et al., 2020b). Although including ridesharing in mathematical programs for SAV dispatch is challenging, it can easily be included in agent-based simulation (Fagnant and Kockelman, 2018; Gurumurthy and Kockelman, 2018; Lokhandwala and Cai, 2018; Vosooghi et al., 2019). The complexities of electric SAV recharging behavior (Chen et al., 2016; Chen and Kockelman, 2016; Jäger et al., 2017; Loeb et al., 2018; Loeb and Kockelman, 2019; Zhang and Chen, 2020) and integration with public transit (Shen et al., 2018; Wen et al., 2018; Huang et al., 2022) are also easy to include in agent-based simulation.

However, agent-based simulations lack an analytical model of fleet-level performance, and the performance depends on factors including SAV dispatch behavior, passenger waiting behavior, network topology, and the passenger trip matrix. This paper aims to address this gap by analytically characterizing the set of demand that can be served by different SAV systems. We aim for a general characterization of SAV service with specific applications to the extensions of ridesharing, electric vehicles, and integration with public transit. This characterization might be useful for SAV operators in their fleet decisions or for metropolitan planning organizations attempting to predict SAV travel behaviors.

1.2. Maximum stability control for shared autonomous vehicles

Characterizing the set of demand that can be served is only convincing if a dispatch policy that provides such service is also given. To achieve both of these goals, we approach the problem using the techniques of max-pressure control, which was originally developed for communications networks by Tassiulas and Ephremides (1992) and later extended to traffic signal timing by Varaiya (2013). Max-pressure control is focused on guaranteeing stability. A network is *stable* if the number of customers queued within remains bounded over time. Intuitively, stability requires that the service rate is equal to the arrival rate of new customers. For SAV dispatch, stability can be defined in terms of the number of customers waiting for travel via SAV. As more customers enter the system, the number of waiting customers remains bounded only if customers exit the system (through pick-up by SAV) at the same rate at which they enter. Although stability is trivially guaranteed in reality because passengers will not wait indefinitely for pick-up, it is still an useful modeling concept because it requires that passengers are served quickly. Most work on max-pressure control models the network as a Markov decision process, where the demand causes stochasticity. When the control policy is given, the model reduces to a Markov chain, and Lyapunov drift techniques are often employed to prove long-run stability.

Applying stability analyses to SAV systems is a recent approach in the literature. It was first introduced by Kang and Levin (2021), who modeled queues of SAV customers and identified a dispatch policy to guarantee stability of those queues for basic SAV systems (without ridesharing, electrification, or public transit integration). However, their dispatch policy was cumbersome as it required a model predictive control approach. Longer planning horizons increased stability but also increased computation time.

Li et al. (2021b) used a different Markov chain approach in which the state is the waiting time instead of the queue length at each customer node. They defined stability in terms of waiting times being bounded and identified a maximum-stable dispatch policy for electric SAVs including recharging. Their dispatch policy is also easier to compute than the policy by Kang and Levin (2021). They used Lyapunov stability techniques to prove that their dispatch policy achieves maximum stability, and equivalently maximum throughput due to Little's Law. However, their scope was limited in two major ways, and the purpose of this paper is to demonstrate the extendability of their analytical approach. First, their model is limited to electric SAVs without ridesharing, but we construct a generalized form that retains the stability properties yet is applicable to a wide range of SAV behaviors. Second, they did not analytically characterize the set of demands that can be served. As previously discussed, predicting the demand served per SAV has been repeatedly studied in the literature using simulations but without analytical structure. As a corollary to finding a general maximum-stable dispatch policy, we also describe the set of demands that could be served by a given fleet of SAVs for a wide range of service behaviors (e.g. ridesharing, integration with public transit).

1.3. Contributions

The contributions of this paper are as follows. We extend the minimum-drift-plus-penalty (MDPP) dispatch policy of Li et al. (2021b) to apply to general SAV systems. We discuss how this extension is applicable to electric SAVs, dynamic ridesharing, and integration with public transit. This extension could also be applied to combinations of those SAV systems or other SAV behaviors. We achieve this generalization by admitting general SAV path assignments that may serve multiple customers and/or include recharging or alternative destinations (public transit stops). We also include preemptive rebalancing in our stability analysis. We provide a general analytical characterization of the stable region which can be used to determine whether a demand rate can be served by a given fleet size. We also prove that any demand rate within the stable region can be served by the MDPP dispatch policy, and any demand rate outside of it cannot be stabilized, which proves that the generalized MDPP policy is maximum-stable. To demonstrate the utility of this general dispatch policy, we show how to apply it and its stable region to the specific SAV extensions of electric SAVs, ridesharing, and public transit integration. This is the first paper to construct a maximum-stable dispatch policy for SAVs with ridesharing or public transit integration, and the first paper to analytically characterize the stable region for electric SAVs. Numerical results are presented to validate the analytical work.

The remainder of this paper is organized as follows. Section 2 introduces the network model and generalized SAV dispatch policy, along with examples of how to adapt it to specific types of SAV systems. Section 3 discusses the stability properties of the dispatch policy and introduces the stable region, which is expanded in Section 4. Numerical results for validation are presented in Section 5, and we conclude in Section 6.

2. Network model and SAV dispatch policy

For this SAV system, we assume that customers enter the network at random times. Upon entering the customer immediately requests SAV service, and waits for the system to dispatch an SAV. An SAV will be assigned to pick up the customer and drive to an acceptable drop-off location. (When integrating with public transit, multiple public transit stops may be acceptable.) Once an SAV assignment is made, it is not changed. Similarly, once an SAV is dispatched to a customer, the customer can be removed from the waiting queue because the SAV is assumed to complete the trip. As part of the dispatch, the SAV may recharge prior to picking up a customer, or might rideshare to reduce travel times. We do not explicitly limit customer waiting times or delays due to ridesharing. Although this is unrealistic, customers are assumed to continue waiting until an SAV is dispatched to them, as opposed to exiting the system after some amount of waiting time has elapsed. This assumption is made because the definition of stability is based on customer waiting times (or equivalently the number of waiting customers). If customers exit the system without service, then the SAV dispatch would not be required to stabilize the system.

Table 1
List of notation.

Notation	Meaning
$A_c(t) \in \{0,1\}$	Random variable that indicates whether a new customer arrives at node $c \in \mathcal{N}_C$ at time t.
B	Maximum battery capacity.
$C_{v\pi}(t) \in \mathbb{Z}_+$	Cost of assigning SAV $v \in \mathcal{N}_V$ to path π at time t, including empty rebalancing cost.
$\bar{C}_{qrs}$	Cost of travel from q to r to s.
$\bar{C}^b_{qrs}$	Cost of travel from q to r to s with a starting battery level of b.
$\bar{C}_\pi$	Average cost of assigning a SAV to path π. (The cost varies due to the empty rebalancing cost.)
$D^{\not p}(t)$	The total cost associated with dispatch policy $\not p$. $\tilde{D}(t)$ is the total cost of dispatch policy $\tilde{\mathbf{x}}(t)$.
$\delta^\pi_c \in \{0,1\}$	Indicates whether path π serves customer node $c \in \mathcal{N}_C$.
E^b_{qrs}	Energy cost of taking trip from q to r to s with starting battery level of b (this could be negative due to recharging).
$e_c(t) \in \{0,1\}$	Indicates whether node $c \in \mathcal{N}_C$ is empty at time t.
$\epsilon > 0$	A positive constant.
$f_{ij} \in \{0,1\}$	Decision variable indicating whether a trip from i to j is used in path π.
$H_c(t) \in \mathbb{Z}_+$	Waiting time of the head-of-line customer at node $c \in \mathcal{N}_c$ at time t. The vector of waiting times is $\mathbf{H}(t)$.
$K < \infty$	A non-infinite constant.
$\kappa < \infty$	A non-infinite constant.
$\mathcal{L}$	Set of locations, including origins, destinations, and public transit stops.
$L(\mathbf{H}(t))$	The Lyapunov function. $\Delta L(t)$ is the change in the Lyapunov function from time t to time $t+1$
$\ell(\mathbf{y}(t))$	Location of SAVs after dispatch $\mathbf{y}(t)$ is enacted at time t
$\mathbf{\Lambda}$	The stable region of demand.
$\hat{\mathbf{\Lambda}}$	A set of demand defined by constraints (10)–(13). Corollary 1 will prove that $\hat{\mathbf{\Lambda}} = \mathbf{\Lambda}$.
$\mathbf{\Lambda}^0$	The interior of the stable region of demand.
$\bar{\lambda}_c$	Average arrival rate to node $c \in \mathcal{N}_C$. The vector of $\bar{\lambda}_c$ is $\bar{\lambda}$.
M	A large positive constant.
$\mathcal{N} = \mathcal{N}_V \cup \mathcal{N}_C$	Set of all nodes.
$\mathcal{N}_V$	Set of SAV nodes. $\lvert\mathcal{N}_v\rvert$ is the size of the SAV fleet.
$\mathcal{N}^A_V(t) \subseteq \mathcal{N}_V$	Set of SAVs that are available at time t.
$\mathcal{N}_C$	Set of customer nodes. Each customer node $c \in \mathcal{N}_C$ is associated with an origin and destination.
$\eta > 0$	A positive constant.
Π	Set of all possible SAV paths.
π	A path specifying a set of locations to be visited and a set of customer nodes to be served.
σ_i	Arrival time at node i in the ridesharing traveling salesman problem.
$R(\ell(\mathbf{y}(t)))$	Value of having SAVs at locations $\ell(\mathbf{y}(t))$.
S_c	Set of acceptable public transit stations for customer $c \in \mathcal{N}_C$ to connect to public transit.
T	End of the time horizon.
$\tau_c(t) \in \mathbb{Z}_+$	Random variable for the inter-arrival time between the head-of-line customer and the next one.
V	Penalty weight placed on SAV service time.
$W(t)$	Function describing the cost of empty rebalancing in the MDPP policy at time t. $W^{\not p}(t)$ is the cost of empty rebalancing for policy $\not p$.
$x_c(t) \in \{0,1\}$	Indicates whether the head-of-line customer at node $c \in \mathcal{N}_C$ is served at time t. The vector of $x_c(t)$ is $\mathbf{x}(t)$.
$\bar{x}_c \in [0,1]$	Average rate of customer service to node $c \in \mathcal{N}_C$.
$y_{v\pi}(t) \in \{0,1\}$	Whether SAV $v \in \mathcal{N}_V$ is assigned to path π at time t.
$\bar{y}_{v\pi} \in [0,1]$	Average rate of SAV assignment to path π
$\bar{\gamma}_{qrs}$	Average rate of SAV trips from q to r to s.
$\bar{\gamma}^b_{qrs}$	Average rate of SAV trips from q to r to s made with a starting battery level of b.
$\mathcal{Z}_\pi \subseteq \mathcal{N}_C$	The set of customer nodes included in path π.
$\bar{\xi}^v_{c,rs}$	Average number of trips from r to s used to serve customer $c \in \mathcal{N}_C$.

We adopt the queueing model of Li et al. (2021b) but with greater generality to admit ridesharing and public transit integration. We distinguish between the typical traffic network, which is usually associated with physical road infrastructure, and the network we define here. A list of notation is given in Table 1. Consider a network $\mathcal{G} = (\mathcal{N}, \mathcal{A})$ representing customers and vehicles. $\mathcal{N}_C \subsetneq \mathcal{N}$ is the set of customer nodes. Each customer node $c \in \mathcal{N}_C$ has specific pick-up and drop-off locations. When integrating with public transit, c can be associated with multiple acceptable drop-off locations and different public stops. $\mathcal{N}_V \subsetneq \mathcal{N}$ is the set of SAV nodes, and are associated with individual SAVs. The geographical location of SAVs change as they move around the network. Therefore, the cost associated with an SAV v serving a specific customer c varies with time. Together, the sets of customer and SAV nodes comprise all nodes in this network.

Let $\mathcal{L}$ be the set of all locations in the network. Every customer origin and destination are included in $\mathcal{L}$, as well as every public transit stop that a customer could use. $\mathcal{L}$ also includes every intermediate node where SAVs could be located (such as charging stations). We assume that the average travel time between any pair of nodes is known. Although $\mathcal{N}$ is sufficient for defining the maximum stability dispatch policy, $\mathcal{L}$ is needed to explicitly characterize the stable region of demand.

We track the waiting time of the head-of-line (HOL) customer at each customer node. Let $H_c(t)$ be the waiting time of the HOL customer at node c at time t. $H_c(t)$ evolves over time via

$$H_c(t+1) = e_c(t)\left(H_c(t) + 1 - x_c(t)\tau_c(t)\right)^+ + (1 - e_c(t))A_c(t) \tag{1}$$

where $(\cdot)^+ = \max\{\cdot, 0\}$, $e_c(t) \in \{0,1\}$ indicates whether node c is empty at time t, $\tau_c(t)$ is the inter-arrival time between the HOL customer and the next one, and $A_c(t) \in \{0,1\}$ is a random variable that indicates whether a new customer arrives. $x_c(t) \in \{0,1\}$ indicates whether node c is served by an SAV at time t, and is a decision variable for the SAV dispatcher. $x_c(t)$ will be related to SAV assignments later. When $x_c(t) = 1$, the HOL customer is served, and the waiting time at node c transitions to the waiting time of the next customer in line (with headway of $\tau_c(t)$) or 0 if none exists. We assume that $\tau_c(t)$ and $A_c(t)$ are independent identically distributed random variables. Let $\bar{\lambda}_c$ be the average arrival rate to node c, so the probability of an arrival is $1/\bar{\lambda}_c$. Then $\tau_c(t)$ has mean $1/\bar{\lambda}_c$, and $A_c(t)$ has mean $\bar{\lambda}_c$.

SAVs are assigned to paths that serve one or more customers. Each path includes customer pick-ups and drop-offs as well as any repositioning or recharging needed for SAV v prior to picking up the first customer. Essentially, a path π is a set of customer nodes that are served together. When ridesharing, multiple customer nodes may be combined. Because ridesharing can involve multiple orderings of customer pick-ups and drop-offs, including multiple passengers being carried within the SAV simultaneously, we do not impose an ordering on service for the customers included in π. Let $C_{v\pi}(t)$ be the time cost of assigning SAV v to path π and $y_{v\pi}(t) \in \{0,1\}$ indicates whether SAV v is assigned to path π. $C_{v\pi}(t)$ varies with time because it depends on the current location and battery level of v. When π serves a customer c, π must include a visit to c's origin to pick up c followed later by a visit to c's destination to drop off c. It is frequently the case that v will not be located at the start of π at time t. If so, then v will have to travel empty (self-relocate) from its current location to the starting node of π. The cost of that relocation is captured in $C_{v\pi}(t)$, which varies with both π and the location of v at time t.

Let $\delta_c^\pi \in \{0,1\}$ indicate whether path π serves customer c. It is possible for a path to serve multiple customers, such as with ridesharing, and δ_c^π gives a definition for $x_c(t)$:

$$x_c(t) = \sum_{\pi \in \Pi} \sum_{v \in \mathcal{N}_V} y_{v\pi}(t)\delta_c^\pi \tag{2}$$

where Π is the set of all paths. Notice that this definition of π differs from Li et al. (2021b), which only considered paths that serve a single customer and without public transit integration. In contrast, a customer could be served by ridesharing or by transporting that customer to a public transit station. This definition therefore extends the generality of SAV assignments, which will be useful for establishing the stability properties.

Using this definition of π, any SAV can be assigned to any path π, and the associated cost depends on which SAV is being assigned to the path, hence the indices v and π for $C_{v\pi}(t)$. At any given time t, we assume that $C_{v\pi}(t)$ can be calculated for any SAV v and any set of customers π. This calculation involves a combination of shortest path and traveling salesman problems for ridesharing. With a typical vehicle capacity, the number of nodes in the traveling salesman problem is fairly low, and the solutions can be generated offline and stored.

2.1. Minimum-drift-plus-penalty (MDPP) policy

We now define a generalized MDPP policy based on the assignment of vehicles to paths. We assume that only idle SAVs can be dispatched. An SAV is idle if it is not currently completing a trip. This assumption does not reduce the generality because SAVs are assumed to complete their current trip before starting another one anyways. Therefore, we define $\mathcal{N}_V^A(t) \subseteq \mathcal{N}_V$ as the set of available SAVs at time t. The cost of assigning SAV v to path π is already defined as $C_{v\pi}(t)$. We define the cost associated with policy ρ, $D^\rho(t)$, as the sum of the SAV service times after enacting policy ρ at time t:

$$D^\rho(t) = \sum_{v \in \mathcal{N}_V^A(t)} \left(\sum_{\pi \in \Pi} y_{v\pi}^\rho(t) C_{v\pi}(t) \right) \tag{3}$$

where $y_{v\pi}^\rho(t)$ indicates the assignment determined by policy ρ.

Let $V \geq 0$ be the penalty weight placed on vehicle service time. The MDPP policy seeks to minimize the objective

$$\begin{aligned} &\min V D(t) - W(t) - \sum_{c \in \mathcal{N}_C} H_c(t) \sum_{v \in \mathcal{N}_V^A(t)} \sum_{\pi \in \Pi} y_{v\pi}(t)\delta_c^\pi \\ &= V \sum_{v \in \mathcal{N}_V^A(t)} \left(\sum_{\pi \in \Pi} y_{v\pi}(t) C_{v\pi}(t) \right) - W(t) - \sum_{c \in \mathcal{N}_C} H_c(t) \sum_{v \in \mathcal{N}_V^A(t)} \sum_{\pi \in \Pi} y_{v\pi}(t)\delta_c^\pi \end{aligned} \tag{4a}$$

subject to constraints on SAV assignment. Each SAV is assigned to at most one path:

$$\sum_{\pi \in \Pi} y_{v\pi}(t) \leq 1 \qquad \forall v \in \mathcal{N}_V^A(t) \tag{4b}$$

Also, each customer node is served at most once:

$$\sum_{\pi \in \Pi} \sum_{v \in \mathcal{N}_V^A(t)} y_{v\pi}(t)\delta_c^\pi \leq 1 \qquad \forall c \in \mathcal{N}_C \tag{4c}$$

with the requirement that $y_{v\pi}(t)$ is binary:

$$y_{v\pi}(t) \in \{0,1\} \qquad \forall v \in \mathcal{N}_V^A(t), \forall \pi \in \Pi \tag{4d}$$

The term $W(t) < \infty$ in objective (4a) represents the rebalancing value of a path assignment $y_{v\pi}(t)$. Several studies on SAVs have observed improvements in waiting time through preemptive rebalancing of SAVs prior to customer requests (Pavone et al., 2012; Fagnant et al., 2015; Hörl et al., 2019). The $W(t)$ term is used to value such rebalancing assignments in the objective, and we will prove the stability properties of the MDPP policy with rebalancing. A suggested form of $W(t)$ will be discussed later in Section 2.3, but for generality we admit any bounded term.

Objective (4a) includes three components. First, $D(t)$ is the cost associated with vehicle dispatch, given by Eq. (3). We weight this penalty by $V \geq 0$ to decide how much it influences the objective. If V is sufficiently large, trips with a large value of $C_{v\pi}(t)$ will be suboptimal. Since $C_{v\pi}(t)$ includes both the empty rebalancing cost and the path travel time, a larger value of V will reduce the time spent on empty rebalancing by waiting until vehicles near the customer are available. Second, $W(t)$ is the rebalancing value, as discussed above. Third, $\sum_{c \in \mathcal{N}_C} H_c(t) \sum_{v \in \mathcal{N}_V^A(t)} \sum_{\pi \in \Pi} y_{v\pi}(t)\delta_c^\pi$ is the value of serving a customer. $\sum_{v \in \mathcal{N}_V^A(t)} \sum_{\pi \in \Pi} y_{v\pi}(t)\delta_c^\pi$ specifies whether customer c is served, and $H_c(t)$ (the waiting time) is the value of serving customer c. Customers that have been waiting longer are more likely to be served.

We now discuss specific implementations of the MDPP policy for different SAV behaviors.

Ridesharing. With ridesharing, one SAV trip can serve multiple customers. Two or more customers may be combined into one trip, such as by picking up all customers first then dropping them off afterwards. Then some paths π can also serve multiple customers, indicated by the δ_c^π variables. In the MDPP policy, any n customers can be combined into one path, but most ridesharing systems add constraints on the maximum additional travel time resulting from combining customers (e.g. Fagnant and Kockelman, 2018; Alonso-Mora et al., 2017a). Such constraints are fully compatible with the definition of the set of paths Π. Finding the cost $C_{v\pi}(t)$ for SAV v to serve path π involves solving a small traveling salesman problem.

Integration with public transit. With public transit integration, customer c's origin and destination can be relocated to public transit stops. We discuss the case where customers are no longer taken to their final destination but instead dropped off at a public transit stop. Relocating the origin of the SAV trip can be done similarly. Let $\mathcal{S}_c$ be the set of acceptable public transit stations for customer c, and we assume that $\mathcal{S}_c$ includes the final destination s_c. Based on objective (4a) which decreases with smaller service costs $C_{v\pi}(t)$, customer c will always be dropped off at the node in $\mathcal{S}_c$ with the shortest travel time from the origin of c. Although other options are feasible, they would never be optimal for objective (4a). Other options may be optimal if the objective included customer travel times, but the purpose of the MDPP policy is to achieve maximum throughput. By reducing the SAV service time per customer, SAVs can be used to serve more customers overall, which is good for network throughput.

Electric SAVs. Li et al. (2021b) already presented a specific MDPP policy for electric SAVs. In terms of the general formulation (4), SAVs with low battery levels must first recharge before serving a path π. Including recharging results in a larger value of $C_{v\pi}(t)$.

2.2. Solving the generalized MDPP policy

Like Li et al. (2021b), we observe that the solution to the generalized MDPP policy may be simplified by analyzing its structure. Assume for now that $W(t) = 0$. If a path π serves one customer c, then SAV v will not be dispatched on π until

$$H_c(t) \geq V C_{v\pi}(t) \tag{5}$$

However, if π serves multiple customers through ridesharing, then v can be dispatched by the MDPP once

$$\sum_{c \in \mathcal{N}_C} \delta_c^\pi H_c(t) \geq V C_{v\pi}(t) \tag{6}$$

We can then observe that problem (4) can be solved on an event-driven basis, which takes a potentially large problem and simplifies it. The three events triggering a solution to problem (4) are as follows.

A SAV becomes available. When a SAV completes a trip, if there are waiting customers in the system, problem (4) should be evaluated to determine whether that SAV should be dispatched or not.

A customer enters the system. The new customer might have an SAV dispatched immediately if it satisfies inequality (5), or might be combined with another customer for ridesharing if that satisfies inequality (6).

Waiting time elapsed. After a customer c has waited sufficiently long, inequalities (5) or (6) will favor dispatch. For any pairing of SAV v with path π, the time point at which dispatch is optimal can be calculated, and problem (4) can be solved at that time.

Using an event-driven solution to problem (4) (instead of a simulation based on time steps) could reduce the number of customer and/or SAV nodes present in problem (4), thereby reducing the number of variables.

We further discuss the scenario for which SAVs are integrated with public transit and customers can be dropped off at one of several public transit stops. If we assume that $W(t) = 0$, then the optimal dispatch is always to serve customers in the least time possible, i.e. take them to the nearest acceptable public transit stop. The inclusion of $W(t)$ may alter this behavior: it may become optimal to take the customer to an alternative stop because of its effect on the SAV location balance.

2.3. Preemptive rebalancing in the MDPP policy

Preemptive rebalancing has been studied extensively in previous work (Pavone et al., 2012; Fagnant et al., 2015; Hörl et al., 2019). The goal of this paper is not to extend those studies, but rather to discuss how preemptive rebalancing can be included in stability analyses of SAV systems. We emphasize the "preemptive" part of rebalancing here. All SAV systems require some empty travel between the drop-off location of one customer and the pick-up location of the next customer served. Preemptive rebalancing refers to moving SAVs before assigning them to another customer, and often before those future customers have entered the system. Preemptive rebalancing could also take the form of preemptively recharging SAVs in anticipation of future use. Although Li et al. (2021b) assumed that SAVs would recharge if located at a charging station, they did not assume that SAVs would travel to a charging station unless it was part of a customer-serving trip. Preemptive rebalancing is not necessary for stability, but prior work has shown that it can be effective at reducing average waiting times (Hörl et al., 2019).

Our goal is to give a general definition of $W(t)$ that can encapsulate most strategies for preemptive rebalancing to discuss their stability properties. We first define the possibility of paths that do not serve customers. Such a path π has $\delta_c^\pi = 0$ for all customers c, but still includes SAV travel, such as rebalancing to a new location or traveling to a recharging station and recharging. The cost of rebalancing path π can still be expressed as $C_{v\pi}(t)$, and the assignment of vehicles to rebalancing paths can still be indicated using the $y_{v\pi}(t)$ variables. Define $\boldsymbol{\ell}(t)$ to be the location vector of SAVs at time t, indexed by SAV node $v \in \mathcal{N}_V$. $\ell_v(t)$ is the current location of v, if idle, or the destination of the trip it is currently completing. $\ell_v(t)$ could also contain information about the remaining duration of v's current trip. We define $\boldsymbol{\ell}(\mathbf{y}(t))$ to be the location vector of SAVs after dispatch $\mathbf{y}(t)$ is enacted. $\boldsymbol{\ell}(\mathbf{y}(t))$ differs from $\boldsymbol{\ell}(t)$ when $\mathbf{y}(t)$ assigns idle SAVs to trips; then their location value in $\boldsymbol{\ell}(\mathbf{y}(t))$ changes from their current idle location to the destination of their future trip.

Let $R(\boldsymbol{\ell}(t))$ be a function that indicates the value of having SAVs at locations given by $\boldsymbol{\ell}(t)$. Larger values of $R(\boldsymbol{\ell}(t))$ indicate SAV locations that are more "balanced" or ideal for future operations. We assume that there exists some location vector(s) $\boldsymbol{\ell}(t)$ that maximize $R(\boldsymbol{\ell}(t))$, and other location vector(s) that minimize $R(\boldsymbol{\ell}(t))$, so that $R(\boldsymbol{\ell}(t))$ is bounded. For instance, the state where all SAVs are idle and distributed near customer nodes according to historical demand might achieve the maximum value of $R(\boldsymbol{\ell}(t))$. Then we define $W(t)$ as

$$W(t) = R(\boldsymbol{\ell}(\mathbf{y}(t))) - R(\boldsymbol{\ell}(t)) \tag{7}$$

which is the improvement (or lack thereof) in R due to assigning SAVs to trips $\mathbf{y}(t)$.

When the waiting times $H_c(t)$ are sufficiently small, it is possible that for objective (4a), $W(t) > \sum_{c\in\mathcal{N}_C} H_c(t) \sum_{v\in\mathcal{N}_V^A(t)} \sum_{\pi\in\Pi} y_{v\pi}(t)\delta_c^\pi$, meaning that preemptive rebalancing is preferred to customer service. Sometimes this is intentional. For example, dispatching an SAV to serve a customer that is far away may be less optimal than waiting for a closer SAV to become available. In that case, rebalancing the idle SAV may be preferable to keeping it idle.

3. Network stability

The main goal of the MDPP policy is to achieve maximum stability, i.e. it serves all customers if at all possible. To prove the maximum stability property, we must first formally define stability. We define the network to be strongly stable if there exists some $\kappa < \infty$ such that

$$\limsup_{T\to\infty} \frac{1}{T} \sum_{t=1}^{T} \sum_{c\in\mathcal{N}_C} \mathbb{E}\left[H_c(t)\right] \le \kappa \tag{8}$$

This definition of stability requires that waiting times for all customer nodes are bounded on average. If customer waiting times are unbounded, then customers are waiting for increasingly long times without being served. Although $H_c(t)$ refers to only the HOL customer, customer arrivals are assumed to be independent of the HOL time and a long HOL time $H_c(t)$ indicates a correspondingly large queue of customers at c. Definition (8) of stability is therefore satisfied if and only if the dispatch policy has service rate equal to the arrival rate of customers. In reality, $H_c(t)$ will not increase to infinity because travelers will not wait infinitely long for service. Definition (8) for stability is therefore primarily useful for studying the throughput properties of SAV dispatch rather than as a realistic model of traveler behavior.

Let $\boldsymbol{\Lambda}$ be the stable region, i.e. the set of arrival rates $\bar{\lambda}$ such that there exists some stabilizing dispatch policy. Let $\boldsymbol{\Lambda}^0$ be the interior of the stable region. For any arrival rate $\bar{\lambda} \in \boldsymbol{\Lambda}^0$, there exists a dispatch policy $\tilde{\mathbf{x}}(t)$ and an $\epsilon > 0$ such that

$$\mathbb{E}\left[\tilde{x}_c(t)|\mathbf{H}(t)\right] \ge \bar{\lambda}_c + \epsilon \qquad \forall c \in \mathcal{N}_C \tag{9}$$

Eq. (9) holds because $\bar{\lambda} \in \boldsymbol{\Lambda}^0$, so $\bar{\lambda}_c + \epsilon$ must be within the set $\boldsymbol{\Lambda}$ (possibly on its boundary). Without loss of generality, we assume that $\tilde{W}(t) = 0$. One such policy is the S-only algorithm of Neely (2006) as used by Li et al. (2021b).

3.1. Characterizing the stable region

Li et al. (2021b) was able to prove stability of their MDPP policy without an explicit characterization of $\boldsymbol{\Lambda}$. However, the characterization of $\boldsymbol{\Lambda}$ is particularly useful for determining the fleet size needed to serve a given demand. We therefore develop a more formal characterization in this paper. Conceptually, stability requires that the service rate to customer nodes satisfies their arrival rates:

$$\sum_{\pi \in \Pi} \sum_{v \in \mathcal{N}_V} \bar{y}_{v\pi} \delta_c^{\pi} \geq \bar{\lambda}_c \qquad \forall c \in \mathcal{N}_C \tag{10}$$

where $\bar{y}_{v\pi}$ is the average rate of assignments of vehicle v to path π, defined as

$$\bar{y}_{v\pi} = \lim_{T \to \infty} \frac{1}{T} \sum_{t=1}^{T} y_{v\pi}(t) \tag{11}$$

Eq. (11) gives the average rate of customer service $\bar{x}_c$:

$$\bar{x}_c = \sum_{v \in \mathcal{N}_V} \sum_{\pi \in \Pi} \delta_c^{\pi} \bar{y}_{v\pi} \tag{12}$$

Also, each vehicle cannot be overscheduled:

$$\sum_{\pi \in \Pi} \bar{C}_{\pi} \bar{y}_{v\pi} \leq 1 \qquad \forall v \in \mathcal{N}_V \tag{13}$$

Note that the average cost of traveling path π, $\bar{C}_{\pi}$, includes any rebalancing or recharging travel time required to reach the starting location of π. Therefore, rebalancing is not an explicit part of the characterization of the stable region because rebalancing time is included in constraint (13). These costs will be explored further in Section 4.

Let $\hat{\boldsymbol{\Lambda}}$ be the set of $\bar{\lambda}$ such that there exists a $\bar{\mathbf{y}}$ satisfying constraints (10)–(13). We want to show that $\hat{\boldsymbol{\Lambda}} = \boldsymbol{\Lambda}$. We first prove that if $\bar{\lambda} \notin \hat{\boldsymbol{\Lambda}}$, then the network is unstable (Proposition 1). We also show that if $\bar{\lambda} \in \hat{\boldsymbol{\Lambda}}^0$ (where $\hat{\boldsymbol{\Lambda}}^0$ is the interior of $\hat{\boldsymbol{\Lambda}}$), then inequality (9) holds (Proposition 2). After proving that the MDPP policy is stable if inequality (9) holds (Proposition 3), we will have shown that $\hat{\boldsymbol{\Lambda}} = \boldsymbol{\Lambda}$.

Proposition 1. *If $\bar{\lambda} \notin \hat{\boldsymbol{\Lambda}}$, then the network is unstable.*

Proof. If constraint (13) is violated, then the vehicle assignments are impossible to fulfill. Consider any vehicle assignments $\bar{\mathbf{y}}$ satisfying constraint (13). Since $\bar{\lambda} \notin \hat{\boldsymbol{\Lambda}}$, there always exists at least one customer c such that

$$\bar{\lambda}_c - \sum_{\pi \in \Pi} \sum_{v \in \mathcal{N}_V} \bar{y}_{v\pi} \delta_c^{\pi} > 0 \tag{14}$$

From Eq. (14),

$$\mathbb{E}\left[x_c(t)\right] = \sum_{\pi \in \Pi} \sum_{v \in \mathcal{N}_V} \bar{y}_{v\pi} \delta_c^{\pi} < \bar{\lambda}_c = \frac{1}{\tau_c} \tag{15}$$

Incorporating Eq. (15) into (1), when $e_c(t) = 1$ (i.e. node c is not empty) we have

$$\mathbb{E}\left[H_c(t+1) - H_c(t)\right] = \mathbb{E}\left[\left(1 - x_c(t)\tau_c\right)^+\right] \geq \eta \tag{16}$$

for some $\eta > 0$. Expanding Eq. (16),

$$\sum_{t'=1}^{t} \mathbb{E}\left[H_c(t'+1) - H_c(t')\right] = \mathbb{E}\left[H_c(t) - H_c(0)\right] \geq \eta t \tag{17}$$

Then

$$\frac{1}{T} \sum_{t=1}^{T} \mathbb{E}\left[H_c(t)\right] \geq \frac{1}{T} \sum_{t=1}^{T} \eta t \tag{18}$$

which violates Eq. (8) for stability. □

Proposition 1 derives an upper bound on the stable region of demand by showing that Eqs. (10)–(13) are the boundary of the stable region. Since the stable region describes the set of demand that can be served by any dispatch policy, Proposition 1 establishes that $\hat{\boldsymbol{\Lambda}}$ is the maximum set of demand that can be served by any dispatch policy, including others that have previously been published in the literature. This helps us compare the throughput of the MDPP policy to all other policies by establishing an upper bound on throughput that the MDPP policy will seek to achieve.

Proposition 2. *If $\bar{\lambda} \in \hat{\boldsymbol{\Lambda}}^0$, then Eq. (9) holds.*

Proof. Since $\bar{\lambda} \in \hat{\boldsymbol{\Lambda}}^0$, there exists a $\bar{\mathbf{y}}$ such that for all customer nodes c

$$\sum_{\pi \in \Pi} \sum_{v \in \mathcal{N}_V} \bar{y}_{v\pi} \delta_c^{\pi} \geq \bar{\lambda}_c + \epsilon \tag{19}$$

Then by Eq. (11) there exists a dispatch policy $\mathbf{x}(t)$ such that

$$\mathbb{E}\left[x_c(t)\right] \geq \bar{\lambda}_c + \epsilon \tag{20}$$

which is equivalent to Eq. (9). □

The $\epsilon > 0$ term in Eq. (9) is essential for proving stability. Proposition 2 establishes that for any demand in the interior of the stable region $\hat{\boldsymbol{\Lambda}}$, Eq. (9) will always hold. This prepares us for using that ϵ term to establish that the MDPP policy stabilizes any demand in $\hat{\boldsymbol{\Lambda}}$. Next, we prove the maximum stability property of the general MDPP policy. In Section 4, we further develop the equations defining $\hat{\boldsymbol{\Lambda}}$ for specific types of SAV systems (electric SAVs, ridesharing, and integration with public transit).

3.2. Stability properties

We first define the Lyapunov function $L(\mathbf{H}(t))$ as

$$L(\mathbf{H}(t)) = \frac{1}{2} \sum_{c \in \mathcal{N}_C} \bar{\lambda}_c (H_c(t))^2 \tag{21}$$

We next define $\Delta L(t)$ as the expected difference in the Lyapunov function from t to $t+1$:

$$\Delta L(t) = \mathbb{E}\left[L(\mathbf{H}(t+1)) - L(\mathbf{H}(t)) | \mathbf{H}(t)\right] \tag{22}$$

As we use the same definitions as Li et al. (2021b) for the Lyapunov function, their Lemma 1 can be used directly here:

Lemma 1 (*of Li et al., 2021b*)**.** *There exists a constant* $0 < K < \infty$ *such that*

$$\Delta L(t) \leq K - \sum_{c \in \mathcal{N}_C} H_c(t) \mathbb{E}\left[x_c(t) - \bar{\lambda}_c | \mathbf{H}(t)\right] \tag{23}$$

holds for all $t > 0$.

Lemma 1 is used to simplify the remainder of the proof of the stability properties. Proposition 3 states the maximum stability properties of the generalized MDPP policy. We follow the form of Li et al. (2021b) but generalize the proof to be valid for more complex SAV assignments. Inequality (24a) provides a bound on the customer waiting times, satisfying definition (8) of stability. Inequality (24b) provides a further bound on the dispatch cost of the MDPP policy.

Proposition 3. *When* $\bar{\lambda} \in \boldsymbol{\Lambda}^0$*, there exist constants* $0 < K < \infty$ *and* $0 < \epsilon < \infty$ *such that*

(a)
$$\limsup_{T \to \infty} \frac{1}{T} \sum_{t=1}^{T} \sum_{c \in \mathcal{N}_C} H_c(t) \leq \frac{1}{T\epsilon} \sum_{t=1}^{T} \left(V \mathbb{E}\left[\tilde{D}(t)\right] - \mathbb{E}\left[V D^{\mathrm{MDPP}}(t)\right] + \mathbb{E}\left[W^{\mathrm{MDPP}}(t)\right] \right) + \frac{K}{\epsilon} \tag{24a}$$

(b)
$$\limsup_{T \to \infty} \frac{1}{T} \sum_{t=1}^{T} \mathbb{E}\left[D^{\mathrm{MDPP}}(t)\right] \leq \frac{1}{T} \sum_{t=1}^{T} \mathbb{E}\left[\tilde{D}(t)\right] + \frac{K}{V} \tag{24b}$$

Proof. Since $\bar{\lambda} \in \boldsymbol{\Lambda}^0$, by assumption there exists a policy $\tilde{\mathbf{x}}(t)$ satisfying Eq. (9). From Eq. (23) of Lemma 1, adding $V\mathbb{E}\left[D^{\mathrm{MDPP}}(t) | \mathbf{H}(t)\right] + W^{\mathrm{MDPP}}(t)$ to both sides of the inequality yields

$$\begin{aligned} &\mathbb{E}\left[V D^{\mathrm{MDPP}}(t) - W^{\mathrm{MDPP}}(t) | \mathbf{H}(t)\right] + \Delta L(t) \\ &\quad \leq \mathbb{E}\left[V D^{\mathrm{MDPP}}(t) - W^{\mathrm{MDPP}}(t) | \mathbf{H}(t)\right] + K - \sum_{c \in \mathcal{N}_C} H_c(t) \mathbb{E}\left[x_c(t) - \bar{\lambda}_c | \mathbf{H}(t)\right] \end{aligned} \tag{25}$$

By the definition of the MDPP policy,

$$\begin{aligned} &\mathbb{E}\left[V D^{\mathrm{MDPP}}(t) - W^{\mathrm{MDPP}}(t) | \mathbf{H}(t)\right] - \sum_{c \in \mathcal{N}_C} \mathbb{E}\left[H_c(t) x_c^{\mathrm{MDPP}}(t) | \mathbf{H}(t)\right] \\ &\quad \leq \mathbb{E}\left[V \tilde{D}(t) - \tilde{W}(t) | \mathbf{H}(t)\right] - \sum_{c \in \mathcal{N}_C} \mathbb{E}\left[H_c(t) \tilde{x}_c(t) | \mathbf{H}(t)\right] \end{aligned} \tag{26}$$

so Eq. (25) becomes

$$\begin{aligned} &V\mathbb{E}\left[D^{\mathrm{MDPP}}(t) - W^{\mathrm{MDPP}}(t) | \mathbf{H}(t)\right] + \Delta L(t) \\ &\quad \leq \mathbb{E}\left[V \tilde{D}(t) - \tilde{W}(t) | \mathbf{H}(t)\right] + K - \sum_{c \in \mathcal{N}_C} H_c(t) \mathbb{E}\left[\tilde{x}_c(t) - \bar{\lambda}_c | \mathbf{H}(t)\right] \end{aligned} \tag{27}$$

Since $\tilde{W}(t) = 0$, we can drop this term from the remaining equations. By Eq. (9),

$$\mathbb{E}\left[V D^{\text{MDPP}}(t) - W^{\text{MDPP}}(t)|\mathbf{H}(t)\right] + \Delta L(t) \leq \mathbb{E}\left[V \tilde{D}(t)|\mathbf{H}(t)\right] + K - \epsilon \sum_{c \in \mathcal{N}_C} H_c(t) \tag{28}$$

Taking expectations and summing from $t = 1$ to T yields

$$\begin{aligned}&\sum_{t=1}^{T} \mathbb{E}\left[V D^{\text{MDPP}}(t) - W^{\text{MDPP}}(t)\right] + \mathbb{E}\left[L(\mathbf{H}(t))\right] - \mathbb{E}\left[L(\mathbf{H}(0))\right] \\ &\quad \leq \sum_{t=1}^{T} V\mathbb{E}\left[\tilde{D}(t)\right] + TK - \epsilon \sum_{t=1}^{T} \sum_{c \in \mathcal{N}_C} H_c(t)\end{aligned} \tag{29}$$

Rearranging the terms in Eq. (29) and dividing by $T\epsilon$ yields

$$\frac{1}{T}\sum_{t=1}^{T}\sum_{c \in \mathcal{N}_C} H_c(t) \leq \frac{1}{T\epsilon}\sum_{t=1}^{T}\left(V\mathbb{E}\left[\tilde{D}(t)\right] - \mathbb{E}\left[V D^{\text{MDPP}}(t) + W^{\text{MDPP}}(t)\right]\right) + \frac{1}{T\epsilon}\mathbb{E}\left[L(\mathbf{H}(0))\right] + \frac{K}{\epsilon} \tag{30}$$

which yields result (24a). Rearranging the terms in Eq. (29) a different way and dividing by VT yields

$$\frac{1}{T}\sum_{t=1}^{T}\mathbb{E}\left[D^{\text{MDPP}}(t)\right] \leq \frac{1}{T}\sum_{t=1}^{T}\mathbb{E}\left[\tilde{D}(t)\right] + \frac{K}{V} \tag{31}$$

which yields result (24b). □

At this point, we can complete our proof that $\hat{\boldsymbol{\Lambda}} = \boldsymbol{\Lambda}$.

Corollary 1. *Eqs. (10)–(13) define a set $\hat{\boldsymbol{\Lambda}}$ such that $\hat{\boldsymbol{\Lambda}} = \boldsymbol{\Lambda}$.*

Proof. By Proposition 1, if $\bar{\lambda} \notin \hat{\boldsymbol{\Lambda}}$, then $\bar{\lambda}$ cannot be stabilized, so $\bar{\lambda} \notin \boldsymbol{\Lambda}$ also. By Proposition 2, if $\bar{\lambda} \in \hat{\boldsymbol{\Lambda}}^0$, then Eq. (9) holds, which means that $\bar{\lambda} \in \boldsymbol{\Lambda}^0$ by Proposition 3. □

3.3. Discussion

Corollary 1 analytically compares the throughput of the MDPP policy to every other policy. Specifically, it proves that the MDPP policy achieves at least as much throughput as all other policies. We accomplish this through the following steps. First, we create a definition of stability in Eq. (8) that is equivalent to serving all demand. If some demand is not served, then the average waiting times will continuously increase with time. Second, we prove that if $\bar{\lambda} \notin \boldsymbol{\Lambda}$, then no dispatch policy can stabilize the network (Proposition 1). This establishes an upper bound on the throughput of every other policy. Finally, we prove that if $\bar{\lambda} \in \boldsymbol{\Lambda}^0$ (the interior of $\boldsymbol{\Lambda}$) then the MDPP policy will stabilize it (Proposition 3). Therefore, the MDPP policy becomes ϵ close to the maximum throughput possible by any policy, where $\epsilon \to 0$.

In the process, we have also derived a general characterization of $\boldsymbol{\Lambda}$, the stable region of demand. Any demand outside of $\boldsymbol{\Lambda}$ cannot be served by any dispatch policy, and any demand in the interior of $\boldsymbol{\Lambda}$ can be served by the MDPP policy. Therefore, the boundary of $\boldsymbol{\Lambda}$ determines whether a demand rate $\bar{\lambda}$ can be served by the given SAV fleet. In other words, given certain origin–destination customer demand rates, we can determine analytically the minimum number of SAVs are needed to serve them. This addresses the replacement ratio problem. The analytical characterization of $\boldsymbol{\Lambda}$ could be used to determine whether a demand rate $\bar{\lambda}$ is contained within the set $\boldsymbol{\Lambda}$ by evaluating whether there exists a $\bar{\mathbf{y}}$ satisfying Eqs. (10) and (13).

Eq. (24a) establishes an upper bound on the expected value of the waiting time. Deriving a hard upper bound on all waiting times is not possible: a large number of customers entering the system simultaneously, which is possible due to stochastic demand, could exceed any hard upper bound on waiting times. However, the expected waiting time is bounded by Eq. (24a). The upper bound in Eq. (24a) depends on ϵ, which is the difference between $\boldsymbol{\Lambda}$ and $\boldsymbol{\Lambda}^0$. In other words, ϵ is the excess capacity available to the system. Increasing ϵ, or equivalently increasing the excess capacity, will reduce the bound on waiting times. ϵ can be increased by reducing the demand for a fixed fleet size. In other words, Proposition 3 can be used to derive an upper bound on the maximum demand that can be served while maintaining a given bound on waiting times.

3.4. Impacts of preemptive rebalancing on performance bounds

Recall that the term $W(t)$ represents the anticipated reduction in waiting time achieved by the MDPP policy (4). $W(t) > 0$ indicates a reduction. However, a positive value of $W(t)$ *increases* the upper bound on headway in inequality (24a), which seems counterintuitive. Mathematically, we can see that $W^{\text{MDPP}}(t)$ switches signs in Eq. (30) to form the upper bound on $H_c(t)$. Eventually, once $H_c(t)$ becomes sufficiently large, $\sum_{c \in \mathcal{N}_C} H_c(t) \sum_{v \in \mathcal{N}_V} \sum_{\pi \in \Pi} y_{v\pi}(t)\delta_c^\pi$ will dominate $W(t)$, and the optimal solution to problem (4) will serve passengers instead of scheduling rebalancing. Until $H_c(t)$ is sufficiently large, vehicles might rebalance instead of serving passengers. Although a good rebalancing strategy will anticipate future demand, future demand is always somewhat uncertain. In the worst case, such rebalancing adds to the waiting time of passengers. Inequality (24a) provides an upper bound on the worst case, not a tight value on the average passenger waiting time.

4. Characterizing the stable region

Previous work (Spieser et al., 2014; Fagnant and Kockelman, 2014; Fagnant et al., 2015; Boesch et al., 2016) has attempted to determine the replacement ratio, or how much demand can be served by a single SAV. In terms of stability, for a given demand we need to find a fleet size such that the demand is within the corresponding stable region $\boldsymbol{\Lambda}$. Of course, the size of $\boldsymbol{\Lambda}$ increases with the number of SAVs. Eqs. (10) and (13) describe $\boldsymbol{\Lambda}$ at a high level. Eq. (10) is fairly explicit, but the term $\bar{C}_\pi$ in Eq. (13) is not easily computable. However, if we know the starting location q of the SAV, then $C_{q\pi}$ can be computed. Many of these equations will be extended using the tuple (q, π) as SAV flow indices instead of π alone.

Our previous results rely on the average SAV flows $\bar{y}_{v\pi}$ on path π, which is indexed by v. In other words, the indices of these variables suggest that we might require different path flows for each individual SAV to achieve maximum stability. Such path flows would be difficult to generalize to arbitrarily-sized fleets of many SAVs. It is helpful to first observe that $\bar{y}_{v\pi}$ is independent of the individual SAV v. To do that, we show that we can choose a value of the SAV-dependent average flows $\bar{y}_{v\pi}$ that is constant per SAV in Proposition 4.

Proposition 4. *Suppose that $\bar{\lambda} \in \boldsymbol{\Lambda}$. Then there exists values of $\bar{y}_{v\pi}$ such that $\bar{y}_{v\pi} = \bar{y}_{v'\pi}$ for all $v, v' \in \mathcal{N}_{\mathrm{V}}$ satisfying Eqs. (10) and (13).*

Proof. Because $\bar{\lambda} \in \boldsymbol{\Lambda}$, there exists some $\bar{y}'_{v\pi}$ satisfying Eqs. (10) and (13). Take

$$\bar{y}_{v\pi} = \frac{1}{|\mathcal{N}_{\mathrm{V}}|} \sum_{v \in \mathcal{N}_{\mathrm{V}}} \bar{y}'_{v\pi} \tag{32}$$

for all π and v. Then

$$\sum_{\pi \in \Pi} \sum_{v \in \mathcal{N}_{\mathrm{V}}} \bar{y}_{v\pi} \delta_c^\pi = \sum_{\pi \in \Pi} \sum_{v \in \mathcal{N}_{\mathrm{V}}} \bar{y}'_{v\pi} \delta_c^\pi \geq \bar{\lambda}_c \qquad \forall c \in \mathcal{N}_{\mathrm{C}} \tag{33}$$

and

$$\sum_{\pi \in \Pi} \bar{C}_\pi \bar{y}_{v\pi} = \frac{1}{|\mathcal{N}_{\mathrm{V}}|} \sum_{v \in \mathcal{N}_{\mathrm{V}}} \sum_{\pi \in \Pi} \bar{C}_\pi \bar{y}'_{v\pi} \leq 1 \tag{34}$$

where the last inequality holds because $\bar{y}'_{v\pi}$ satisfies constraint (13), which verifies that $\bar{y}_{v\pi}$ satisfies Eqs. (10) and (13). □

Since we can find a single value of $\bar{y}_{v\pi}$ that applies to all SAVs, we can now show that we can observe that the demand rates that can be served increase linearly with the fleet size. A similar result was obtained by Kang and Levin (2021) but only for SAVs without ridesharing, electric vehicles, or public transit integration. We can provide a more general result here. When papers (e.g. Fagnant and Kockelman, 2018) discuss a replacement ratio, they implicitly assume that the fleet size required to serve customers increases linearly with the demand via the replacement ratio. It is not immediately obvious that the relationship is linear. Proposition 5 establishes that the relationship is linear, and equivalently establishes the existence of the replacement ratio.

Proposition 5. *If the fleet size increases by a rate α, then the average demand rates that can be served also increase by α.*

Proof. By Proposition 4 there exists a $\bar{y}_{v\pi}$ which is identical for all v satisfying constraints (10) and (13). Then constraint (10) can be rewritten as

$$\bar{\lambda}_c \leq \sum_{\pi \in \Pi} \sum_{v \in \mathcal{N}_{\mathrm{V}}} \bar{y}_{v\pi} \delta_c^\pi = |\mathcal{N}_{\mathrm{V}}| \sum_{\pi \in \Pi} \bar{y}_{v\pi} \delta_c^\pi \qquad \forall c \in \mathcal{N}_{\mathrm{C}} \tag{35}$$

Suppose that the fleet size increases from F to αF. Then inequality (35) becomes

$$\alpha \bar{\lambda}_c \leq \alpha F \sum_{\pi \in \Pi} \bar{y}_{v\pi} \delta_c^\pi \qquad \forall c \in \mathcal{N}_{\mathrm{C}} \tag{36}$$

which indicates that $\alpha \bar{\lambda}_c$ demand can be served. □

The purpose of this section is to expand Eq. (13) for different types of SAV services. We start by considering the simple case of one customer per trip and non-electric SAVs (without recharging delays), then expand to ridesharing, electric SAVs, and integration with public transit. We note that although (Li et al., 2021b) developed a maximum-stable MDPP policy for electric SAVs with recharging, they did not characterize the stable region explicitly or expand their results to other variations of SAV behavior.

The difference between $\boldsymbol{\Lambda}$ and its interior $\boldsymbol{\Lambda}^0$ can be described in terms of inequalities (10) and (13). If both inequalities are strict, then they describe the interior of the stable region $\boldsymbol{\Lambda}^0$. In our stable region analyses, it is therefore sufficient to expand Eqs. (10) and (13) to define $\boldsymbol{\Lambda}$, and then the characterization of $\boldsymbol{\Lambda}^0$ follows immediately.

4.1. One customer per SAV trip

We start by focusing on the simple system where each SAV trip serves one customer (without electric SAVs or public transit integration). Then customer node c can be rewritten as origin–destination pair $(r, s) \in \mathcal{L}^2$. In other words, a customer waiting at

node c is equivalently waiting for travel from r to s. In a slight overload of notation, we define $\bar{C}_{rs}$ as the average travel time from r to s. We assume that travel times may vary randomly due to traffic conditions and therefore use the average time from r to s. However, $\bar{C}_{\pi} \neq \bar{C}_{rs}$ because the SAV may not be located at r, which necessitates additional travel time to reach r. We therefore want to identify the average number of trips spent traveling from q to r (to pickup customer c) then from r to s, which sufficiently describes $\bar{C}_{\pi}$. To that end, we define $\bar{C}_{qrs}$ to be the average travel time from q to r to s. Let $\bar{\gamma}_{qrs}$ be the average flow of SAVs from q to r to s.

$$\sum_{q \in \mathcal{L}} \bar{\gamma}_{qrs} = \sum_{v \in \mathcal{N}_V} \bar{y}_{v\pi} \tag{37}$$

where π is serving the customer from r to s, but includes information about the starting location q also. By Proposition 4, we can choose $\bar{y}_{v\pi} = \frac{1}{|\mathcal{N}_V|} \bar{\gamma}_{qrs}$ to obtain $\bar{y}_{v\pi}$ from $\bar{\gamma}_{qrs}$. By conservation of flow, the average flow into q must also equal the average flow out of q. Therefore,

$$\sum_{(r,s) \in \mathcal{L}^2} \bar{\gamma}_{qrs} = \sum_{(s,r) \in \mathcal{L}^2} \bar{\gamma}_{srq} \qquad \forall q \in \mathcal{L} \tag{38}$$

The cost associated with path assignments of $\bar{\gamma}^{v}_{qrs}$ is

$$\sum_{v \in \mathcal{N}_V} \sum_{\pi \in \Pi} \bar{C}_{\pi} \bar{y}_{v\pi} = \sum_{q \in \mathcal{L}} \sum_{(r,s) \in \mathcal{L}^2} \bar{\gamma}_{qrs} \bar{C}_{qrs} \leq |\mathcal{N}_V| \tag{39}$$

to match Eq. (13). Constraint (10) becomes

$$\sum_{q \in \mathcal{L}} \bar{\gamma}_{qrs} \geq \bar{\lambda}_{rs} \qquad \forall (r,s) \in \mathcal{L}^2 \tag{40}$$

where $\bar{\lambda}_{rs}$ is used to mean $\bar{\lambda}_c$ for customers traveling from r to s. Constraints (38)–(40) characterize the stable region when each vehicle serves 1 customer at a time.

4.2. Ridesharing

We next consider the stable region for ridesharing path assignments. Here π indicates a set of customers to be served. For instance, with two customers, $\pi = \{c_1, c_2\}$ where $c_1 \neq c_2$ because a single customer node is never served twice by one path. We again want to calculate $\bar{C}_{\pi}$, the average travel time to serve customers in π. We define $\bar{C}_{q\pi}$ to be the average travel time to serve customers π with the SAV starting at node q. With multiple travelers to be served, finding $\bar{C}_{q\pi}$ depends on the order in which travelers are served. This is a traveling salesman problem with node ordering. A customer's origin node must be visited before their destination node. Assuming that minimizing the time required for v to serve customers in π is desired, the problem can be formulated as follows. Let $\mathcal{Z}_{\pi} = \{r_c, s_c : c \in \pi\}$ be the set of nodes to visit, which includes the origin r_c and destination s_c for every customer c in π. If multiple customers have the same origin or destination, we create duplicate co-located nodes for clarity in the formulation. We form a network with nodes $\mathcal{L}_{q\pi} = \{q\} \cup \mathcal{Z}_{\pi} \cup \{\diamond\}$ and links between the nodes where $\diamond$ represents the dummy node. The cost of travel between any node and $\diamond$ is 0. Let $f_{ij} \in \{0,1\}$ indicate whether v travels from i to j, and let $\bar{C}_{ij}$ be the average travel time from i to j. Let σ_i indicate the order in which i is visited, which is necessary to ensure that the origin node for a customer is visited before their destination node. The cost of visiting nodes in a specific order given by the f_{ij} variables is defined by objective (41a). The remainder of problem (41) defines the constraints needed to ensure that f_{ij} specifies service for all customers in $\mathcal{Z}_{\pi}$.

$$\text{min} \qquad \bar{C}_{q\pi} = \sum_{(i,j) \in \mathcal{L}_{q\pi}} f_{ij} \bar{C}_{ij} \tag{41a}$$

$$\text{s.t.} \qquad \sum_{j \in \mathcal{L}_{q\pi}} f_{qj} = 1 \tag{41b}$$

$$f_{iq} = 0 \qquad \forall i \in \mathcal{L}_{q\pi} \tag{41c}$$

$$\sum_{i \in \mathcal{L}_{q\pi}} f_{i\diamond} = 1 \tag{41d}$$

$$f_{\diamond j} = 0 \qquad \forall j \in \mathcal{L}_{q\pi} \tag{41e}$$

$$\sum_{i \in \mathcal{L}_{q\pi}} f_{ir_c} = 1 \qquad \forall c \in \mathcal{N}_C \tag{41f}$$

$$\sum_{i \in \mathcal{L}_{q\pi}} f_{is_c} = 1 \qquad \forall c \in \mathcal{N}_C \tag{41g}$$

$$\sum_{j \in \mathcal{L}_{q\pi}} f_{r_c j} = 1 \qquad \forall c \in \mathcal{N}_C \tag{41h}$$

$$\sum_{j \in \mathcal{L}_{q\pi}} f_{s_c j} = 1 \qquad \forall c \in \mathcal{N}_C \tag{41i}$$

$$\sigma_{r_c} \le \sigma_{s_c} + 1 \qquad \forall c \in \mathcal{N}_C \tag{41j}$$

$$\sigma_i \le \sigma_j + 1 + M f_{ij} \qquad \forall (i,j) \in \left(\{q\} \cup \mathcal{Z}_\pi\right)^2 \tag{41k}$$

$$f_{ij} \in \{0,1\} \qquad \forall (i,j) \in \left(\{q\} \cup \mathcal{Z}_\pi\right)^2 \tag{41l}$$

Constraint (41k) ensures that $\sigma_i < \sigma_j$ if link f_{ij} is used, which enables constraint (41j) to check that the origin r_c of customer c is visited before the destination s_c. The dummy node $\diamond$ is used to ensure that constraint (41i) always holds, as $\diamond$ is the last node visited. Like other traveling salesman problems, problem (41) is NP-hard. However, for the typical 2–4 customer capacity of a passenger car, the number of possible orderings is fairly small.

SAVs typically have capacity limits on the number of passengers that can simultaneously be traveling in the same SAV. Since our SAV dispatch is defined in terms of paths, limiting the number of passengers in an SAV is equivalent to limiting the number of customers served by a single path. We solve problem (41) to find a path π that serves all customers in some set $\mathcal{Z}_\pi$. To limit the number of passengers served, we need to limit the size of $\mathcal{Z}_\pi$. For instance, if SAVs can hold at most 4 passengers, then we only consider grouping passengers into sets $\mathcal{Z}_\pi$ satisfying $|\mathcal{Z}_\pi| \le 4$.

In problem (41), we assume that the SAV seeks to serve the passengers in π as quickly as possible. In general, it is possible to create examples such that choosing a suboptimal $\bar{C}_{q\pi}$ with a different end node $s_{q\pi}$ could reduce the empty travel costs for the next passenger being served. Choosing such a path requires some knowledge of the next passenger that would be served by the SAV, and is not necessary for stability by Proposition 3.

Once $\bar{C}_{q\pi}$ is determined by problem (41), we can more explicitly characterize the stable region. In a slight overload of notation, let $s_{q\pi}$ be the last node visited when customers π is served by an SAV starting from node q. Conservation of flow requires that the number of trips leaving node q relates to the number of trips arriving at node q:

$$\sum_{\pi \in \Pi} \bar{\gamma}_{r\pi} = \sum_{q \in \mathcal{L}} \sum_{\pi \,:\, s_{q\pi} = r} \bar{\gamma}_{q\pi} \qquad \forall r \in \mathcal{L}, \forall v \in \mathcal{N}_V \tag{42}$$

Like constraint (39), we can rewrite constraint (13):

$$\sum_{v \in \mathcal{N}_V} \sum_{\pi \in \Pi} \bar{C}_\pi \bar{y}_{v\pi} = \sum_{q \in \mathcal{L}} \sum_{\pi \in \Pi} \bar{\gamma}_{q\pi} \bar{C}_{q\pi} \le |\mathcal{N}_V| \tag{43}$$

Constraints (42) and (43) together with (10) characterize the stable region with ridesharing.

In this ridesharing scenario, it is possible that grouping customers into one set π does not offer any benefits over treating them separately. For instance, suppose that $\pi = \{c_1, c_2\}$ but the optimal solution to problem (41) is to first visit the origin and destination c_1 and only after to visit the origin and destination of c_2. Then we could equivalently form sets $\pi_1 = \{c_1\}$ and $\pi_2 = \{c_2\}$ and observe that $\bar{C}_{q\pi} = \bar{C}_{q\pi_1} + \bar{C}_{s_{q\pi_1}\pi_2}$. In other words, grouping customers c_1 and c_2 results in the same travel cost to v.

However, the main challenge of finding the best average vehicle flows satisfying constraints (10), (42), and (43) is that the number of possible paths greatly increases with the number of customers served per path. With 2 customers served per SAV trip, the number of possible $\bar{\gamma}_{q\pi}$ variables is $O\left(\mathcal{L}^5\right)$ because each customer has an origin and destination location. It is still not obvious how to solve this problem for large networks.

4.3. Integrating SAVs with public transit

We consider again the case where $|\pi| = 1$, but now assume integration with public transit. The SAV no longer needs to carry customer c from their origin to their destination. Instead, they can carry customer c to an acceptable public transit station. Recall that $\mathcal{S}_c$ is the set of acceptable public transit stations for customer c, with $s_c \in \mathcal{S}_c$. Like in Section 4.1, each SAV trip from q serving customer set π can again be described as the tuple $(q, r, s) \in \mathcal{L}^3$, meaning that constraints (38) and (39) describe the set of valid average path assignments.

The challenge in defining public transit service is that each customer has multiple valid destination nodes (the set $\mathcal{S}_c$). These destination nodes can overlap with other customers, meaning that examples exist where a trip from r to s could be a valid service for several different customers. With the assumption of one customer per trip, we need to distinguish which customer is being served. Let $\bar{\xi}_{c,rs}$ be the average number of trips r to s used to serve customer node c. Obviously, $\bar{\xi}_{c,rs} = 0$ if $s \notin \mathcal{S}_c$ or $r \ne r_c$. Then we can disaggregate the number of trips from r to s by which customer is being served:

$$\sum_{c \in \mathcal{N}_C} \bar{\xi}^v_{c,rs} = \sum_{q \in \mathcal{L}} \bar{\gamma}_{qrs} \tag{44}$$

To ensure that each customer is served, stable region constraint (10) can be rewritten as

$$\sum_{(r,s) \in \mathcal{L}^2} \bar{\xi}_{c,rs} \ge \bar{\lambda}_c \qquad \forall c \in \mathcal{N}_C \tag{45}$$

Constraint (39) together with constraint (44) suggests that when the demand rate $\bar{\lambda}$ is small, SAVs can serve customers by taking them directly to their destination while retaining stability. Choosing (r, s) pairs with larger $\bar{C}_{qrs}$ reduces the number of (q, r, s) trips served by constraint (39) but if $\bar{\lambda}$ is small enough, then those (q, r, s) trips are sufficient for stability. On the other hand, when $\bar{\lambda}$ is larger, the cost required to serve each customer can be reduced by choosing different (r, s) pairs. Instead of taking customer c

directly to s_c, the SAV can instead take them to some public transit stop in S_c with a lower $\bar{C}_{qrs}$ cost. That enables more (q, r, s) trips to be made by constraint (39), which increases the number of customers per hour that can be served.

The MDPP policy objective (4a) chooses the minimum cost paths to serve each customer, which of course ensures stability. The above discussion suggests that other maximum-stable policies exist which are more generous to customers when demand is small. Finding such policies and showing their stability properties in an interesting question for future work.

4.4. Electric SAVs with recharging

We again consider one customer per trip, but now assume that SAVs are electric and require recharging. Unlike internal combustion engine vehicles, recharging requires significant time which adds to the time required to pickup a passenger. The stability properties of electric SAVs were previously studied by Li et al. (2021b) using a very similar model, but they did not characterize the stable region. This section therefore describes the first attempt to give equations that identify which demand rates can be served.

Like in Section 4.1, we assume that customers are served one at a time, meaning that every SAV trip can be described by the tuple $(q, r, s) \in \mathcal{L}^3$. However, the cost depends on whether recharging is required, which depends on the SAV battery level. Let C^b_{qrs} be the travel time required for trip (q, r, s) to serve customer (r, s) with an initial battery level of b. Let B be the maximum battery capacity. We discretize B into different levels denoted by b to index the decision variables. More intervals achieves a more precise characterization of the stable region but at the cost of more variables.

If recharging is not required, then the SAV can travel directly from q to r to s. However, if the initial battery level is low and requires charging, then the SAV must first travel from q to some recharging station, spend time recharging, and only afterwards proceed to r. We assume that the battery capacity is sufficient for customer (r, s) to be served for a fully charged SAV departing any charging station in the network. This is likely realistic for most cities and typical electric vehicle battery capacity.

Let E^b_{qrs} be the energy cost of taking that trip, so the final charge upon reaching s would be $b - E^b_{qrs}$. Some trips have multiple recharging options, such as recharging fully or only partially. For simplicity, we will assume that a single recharging option is chosen per (q, r, s) trip with initial battery level b, but this approach can easily be modified to represent multiple recharging options albeit with more decision variables. We admit $E^b_{qrs} < 0$, meaning that the SAV recharged as part of the trip. We further require that the final charge $b - E^b_{qrs}$ is sufficiently high for the SAV to reach another recharging station. In other words, the SAV cannot become stuck due to lack of charge after arriving at s.

Let γ^b_{qrs} be the average number of (q, r, s) trips undertaken with a starting battery level of b. Then the associated cost is

$$\sum_{v \in \mathcal{N}_V} \sum_{\pi \in \Pi} \bar{C}_\pi \bar{y}_{v\pi} = \sum_{q \in \mathcal{L}} \sum_{(r,s) \in \mathcal{L}^2} \sum_{b=0}^{B} \bar{\gamma}^b_{qrs} \bar{C}^b_{qrs} \le |\mathcal{N}| \tag{46}$$

which replaces Eq. (10). Rewriting Eq. (13) requires expanding conservation to include both vehicle locations and their battery levels. Starting with the conservation Eq. (38),

$$\sum_{(q,r) \in \mathcal{L}^2} \sum_{b' : b' + E^{b'}_{qrs} = b} \bar{\gamma}^{b'}_{qrs} = \sum_{(r,q) \in \mathcal{L}^2} \bar{\gamma}^b_{srq} \qquad \forall s \in \mathcal{L}, \forall b \in [0, B] \tag{47}$$

The right hand side of Eq. (47) refers to trips departing s with battery level b. Conservation means that an incoming trip must have arrived at s with battery level b. This is achieved if SAV v starts trip (q, r, s) with battery level b' and $b' + E^{b'}_{qrs} = b$, which is indicated on the left hand side of Eq. (47).

5. Numerical results

Li et al. (2021b) explored the waiting time properties of the MDPP policy for electric SAVs. Rather than repeat their results for other types of SAVs systems, we focus on validating the stable region predicted by Eqs. (10) and (13). These equations can be used to quickly compute the minimum fleet size needed to serve a given demand.

Proposition 3 established that the MDPP policy stabilizes any demand in $\boldsymbol{\Lambda}^0$, and Corollary 1 proved that Eqs. (10)–(13) establish a set $\hat{\boldsymbol{\Lambda}} = \boldsymbol{\Lambda}$. We now aim to explore the relationship between $\hat{\boldsymbol{\Lambda}}$ and $\boldsymbol{\Lambda}$ through simulation. We can only verify stability of demand vectors $\bar{\lambda} \in \boldsymbol{\Lambda}^0$, which requires the existence of some $\epsilon > 0$ difference between $\bar{\lambda}$ and the boundary of $\boldsymbol{\Lambda}$, as given by Eq. (9). When ϵ is very small, then Eq. (24a) predicts the bound on $\sum_{c \in \mathcal{N}_C} H_c(t)$ to be large, which makes stability difficult to detect numerically. Some differences between the simulated stable region and the predicted boundary are to be expected based on the minimum ϵ needed.

5.1. Theoretical stable region

We assume that customer demand proportions are constant but vary the total number of customers. In other words, $\bar{\lambda}_c$ is constant but we introduce a factor α such that the total demand is $\alpha \sum_{c \in \mathcal{N}_C} \bar{\lambda}_c$. To predict the boundary of the stable region, we solve the problem

$$\max \quad \alpha \tag{48a}$$

$$\text{s.t.} \quad \sum_{\pi \in \Pi} \sum_{v \in \mathcal{N}_V} \bar{y}_{v\pi} \delta^\pi_c \ge \alpha \lambda_c \qquad \forall c \in \mathcal{N}_C \tag{48b}$$

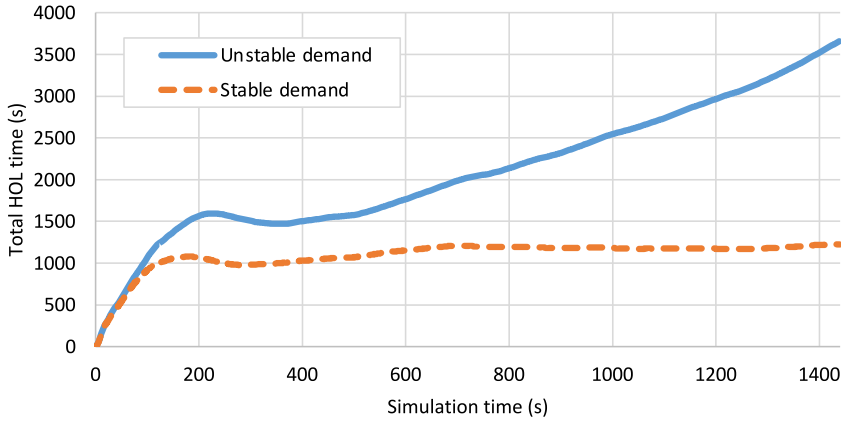

Fig. 1. Illustration of HOL time $\frac{1}{T}\sum_{t=1}^{T}\sum_{c\in\mathcal{N}_C} H_c(t)$.

Eq. (13)

where α is an additional decision variable, $\bar{\lambda}_c$ is constant, Eq. (48b) is based on (10) with the addition of serving $\alpha\bar{\lambda}_c$ demand, and constraint (13) is expanded as discussed in Section 4. We could alternatively use the objective $\max \sum_{c\in\mathcal{N}_C} \bar{\lambda}_c$ with $\bar{\lambda}_c$ as a decision variable. However, certain demand patterns (e.g. symmetric vs. asymmetric demand) can be served with less empty travel, but realistic travel demand does not follow ideal patterns. Therefore, problem (48) assumes that the demand proportions are fixed and seeks to serve as much demand as possible with those proportions. We used IBM CPLEX to solve problem (48)., then calculated the replacement ratio as $\frac{1}{|\mathcal{N}_V|}\sum_{c\in\mathcal{N}_C}\alpha\bar{\lambda}_c$ customers per hour served per SAV.

We can write another program to find the minimum fleet size needed to serve demand rates $\bar{\lambda}$. By Proposition 4, only a single vehicle index v is required.

$$\min \quad F \tag{49a}$$

$$\text{s.t.} \quad F\sum_{\pi\in\Pi}\bar{y}_{v\pi}\delta_c^{\pi} \geq \bar{\lambda}_c \qquad \forall c\in\mathcal{N}_C \tag{49b}$$

Eq. (13)

Since problems (48) and (49) are equivalent, we will focus on demonstrating problem (48) in the numerical results. Problem (49) is more useful for finding the minimum SAV fleet size.

5.2. Simulated stable region

Finding the stable region from simulation is difficult. The left hand side of stability definition (8) can easily be calculated for any given simulation. However, determining whether it is bounded is more difficult for two reasons. The $\mathbb{E}\left[H_c(t)\right]$ term can be approximated through Monte Carlo simulations but not calculated directly. Definition (8) holds as $T\to\infty$, but an infinite time horizon is not practical for simulation either. Fig. 1 illustrates the HOL times from two different demand rates on the Sioux Falls network. Even when the network is stable, and with a 24hr simulation time horizon, fluctuations in the average HOL time are observed which make erroneous misidentification of stability possible. To reduce the impact of fluctuations, we take a 1hr moving average of Eq. (8). Then, we identify the network as stable if the average increase in averaged HOL time for the last 2 h of simulation is less than 0.1s per hour per customer. 10 Monte Carlo simulations are used to approximate $\mathbb{E}\left[H_c(t)\right]$. Theoretically, there should be an α defined in problem (48) at the boundary of the stable region such that any $\alpha' > \alpha$ will result in an unstable network and any $\alpha' < \alpha$ will result in a stable network. Due to stochasticity, this is not guaranteed for any random realization of demand. However, after the use of Monte Carlo simulations, we assume that such properties of α can be identified in simulation too. We use a line search to find that maximum α such that demand of $\alpha\bar{\lambda}$ is stable for a given fleet size $|\mathcal{N}_V|$, and report the value of $\alpha|\bar{\lambda}|$ as the maximum stable demand. Results in Section 5.3 validate that this method of detecting stability achieves results similar to those predicted mathematically.

5.3. Validation of stable region calculations

For the purposes of validation, we use the Sioux Falls network with 24 nodes and 76 links, shown in Fig. 2. This network has been modified from original versions to include bus routes from Levin et al. (2019) and electric SAV charging stations. The bus routes and charging stations were chosen arbitrarily to highlight SAV behaviors and are not intended to represent reality. We use the free flow travel times for all links. Since travel times are assumed to be exogenous constants, the use of these travel times is

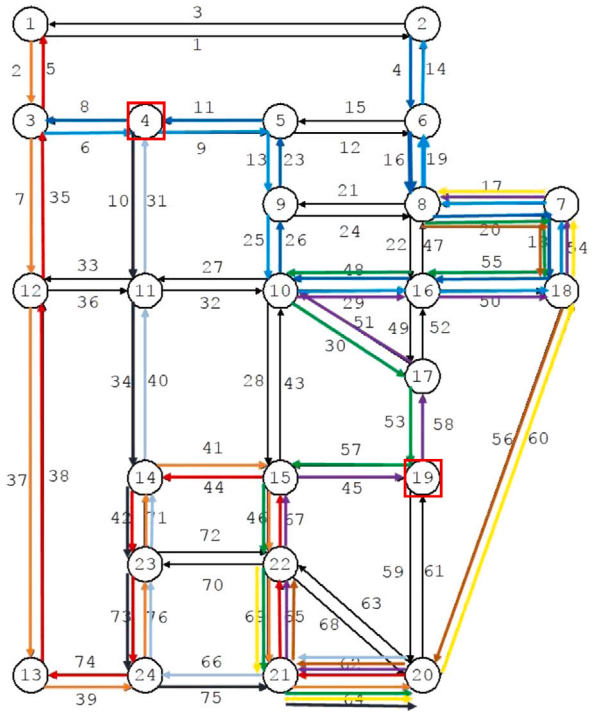

Fig. 2. Sioux Falls network with added bus routes and electric vehicle charging stations. Colored arrows indicate different bus routes, and red squares around nodes indicate charging stations.

not limiting. The Sioux Falls network is chosen for its size because many simulations can be quickly run on this network. The goal of this section is to validate the use of Eqs. (10) and (13) to determine the stable region, which admits the use of those equations for calculating the stable region on other networks. In Section 5.4, we will explore the effects of different SAV behaviors on the maximum stable demand for a larger network.

5.3.1. *One passenger per SAV*

We first focus on the simple case where each SAV carries one passenger at a time without public transit integration or electric vehicle recharging. Fig. 3 shows the simulated and calculated maximum stable demand with fixed origin–destination proportions. Some differences are to be expected due to the challenges in detecting stability discussed in Section 5.2. However, the simulated and calculated lines match up very closely. Surprisingly, the simulated stable demand has a slightly larger slope. This is best explained by a decrease in the average time spent serving a passenger, $\bar{C}$, calculated by:

$$\bar{C} = \lim_{T\to\infty} \frac{1}{T} \sum_{t=1}^{T} \frac{\sum_{\pi\in\Pi} y_{v\pi}(t) C_{v\pi}(t)}{\sum_{\pi\in\Pi} y_{v\pi}(t)} \tag{50}$$

As $\sum_{c\in\mathcal{N}_C} \bar{\lambda}_c \to \infty$, $\bar{C}$ from simulation should approach the value predicted by the stable region Eqs. (10) and (13), calculated as

$$\bar{C} = \frac{\sum_{v\in\mathcal{N}_V} \sum_{\pi\in\Pi} \bar{C}_\pi \bar{y}_{v\pi}}{\sum_{v\in\mathcal{N}_V} \sum_{\pi\in\Pi} \bar{y}_{v\pi}} \tag{51}$$

Note that Eqs. (50) and (51) are equal due to Eq. (11), but Eq. (50) can be approximated by simulation and Eq. (51) is found as a byproduct of solving problem (48).

$\bar{C}$ is shown in Fig. 4 for the maximum stable demand in simulation plotted in Fig. 3. Note that since link travel times are constant, the in-vehicle travel time for this scenario is fixed. Therefore, any increase in $\bar{C}$ is due to empty travel time while SAVs travel to their assigned customer pick-up location. For simulation, $\bar{C}$ is observed to be generally decreasing with the fleet size. The reason for this trend is that as the demand and fleet size increase, it is easier to match passengers with nearby SAVs. In other words, discretization of SAVs and customers creates less overhead for the dispatch as fleet size increases, and the simulated value of $\bar{C}$ approaches the calculated value. The decreasing trend in $\bar{C}$ with respect to the fleet size explains the relative increase in the maximum stable demand estimated from simulation. The simulated stable region still closely follows a linear trend as predicted by Proposition 5.

The stable region detected through simulation also depends on the parameter V in the MDPP. Fig. 5 shows the change in $\bar{C}$ with respect to V. For small values of V, such as around 0.6, SAVs are dispatched more quickly, which reduces the dispatch delay.

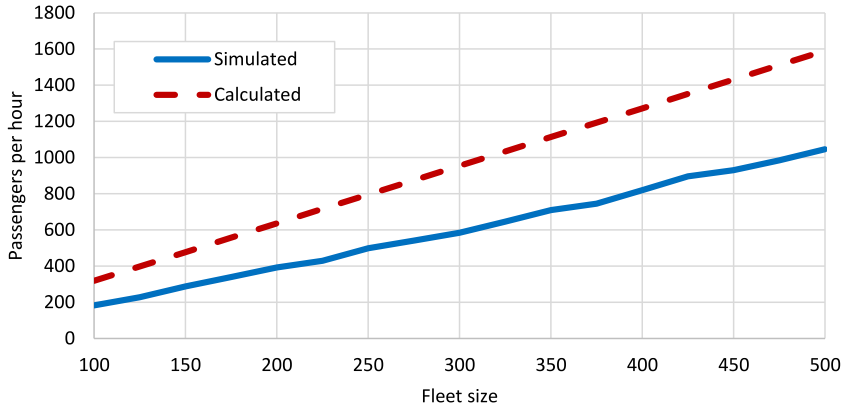

Fig. 3. Simulated and calculated maximum stable demand (using fixed origin–destination proportions) for the Sioux Falls network with one passenger per SAV trip.

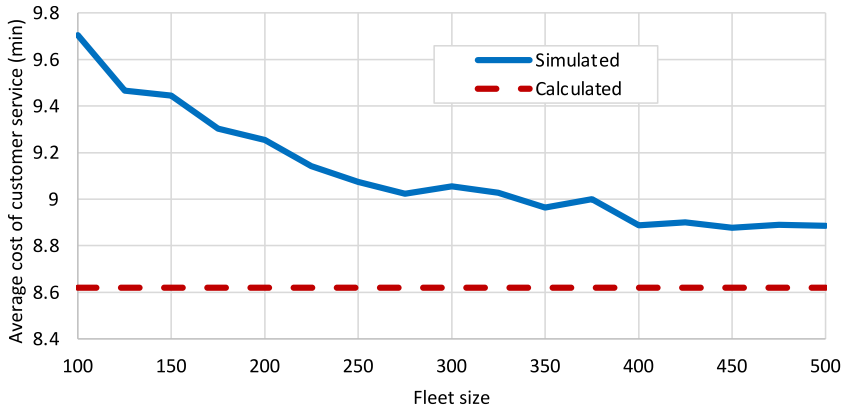

Fig. 4. Simulated and calculated average $\bar{C}$ for the Sioux Falls network with one passenger per SAV trip.

However, lower values of V also cause $\bar{C}$ to be higher because the MDPP dispatches SAVs immediately, meaning that SAV-to-customer assignments may be less efficient. Therefore when $V \leq 0.5$, the dispatch delay is actually quite high because as SAVs are assigned on inefficient paths, they are less available for customers. If the number of waiting customers becomes sufficiently large, the efficiency would return but at the cost of having large average waiting times for dispatch. In contrast, larger values of V reduce $\bar{C}$ but increase the dispatch delay because SAVs are not dispatched immediately when a customer arrives. As V increases, $\bar{C}$ approaches the minimum value calculated by Eq. (51). At some larger values of V, $\bar{C}$ from simulation actually is slightly lower than the value from Eq. (48), but that could be caused by stochasticity in the demand in simulation.

5.3.2. *Integration with public transit*

We adopt the fictional bus routes of Levin et al. (2019) for the Sioux Falls network (shown in Fig. 2) to evaluate the stable region demand. Due to their significant coverage of the network, many customers can be served entirely by bus, and others can be served more quickly by dropping them off at a bus stop. However, bus transfers were not considered: if a transfer was required, SAVs would be used to take a customer directly to a bus route that reaches their destination. The maximum walking distance to and from bus stops was assumed to be 0.25 miles. When determining how to serve a customer, the objective is to maximize the number of customers served. Hence, SAVs will prefer for customers to use a bus route as much as possible to reduce the time that SAVs spent on service.

Fig. 6 compares the simulated and calculate stable demand. Overall, the results are similar to Fig. 3, but slightly more differences are observed between the simulated and calculated values at large fleet sizes. The discrepancy is partly due to the large number of customers served entirely by public transit, which makes detection of a stable demand in simulation more difficult. Public transit does not have capacity limitations, yet customers served mostly or entirely by public transit still affects the denominator when calculating the average increase in average HOL time. We also note that the simulated stable region is slightly higher than the calculated stable region. Again, this is due to the difficulty in determining whether a demand is stable from simulation alone.

Fig. 7 shows the average simulated and calculated $\bar{C}$. Unlike Fig. 4, the simulated $\bar{C}$ does not appear to be converging to the calculated $\bar{C}$. We note that even at $|\mathcal{N}_V| = 100$, the difference between the simulated and calculated values of $\bar{C}$ are quite small

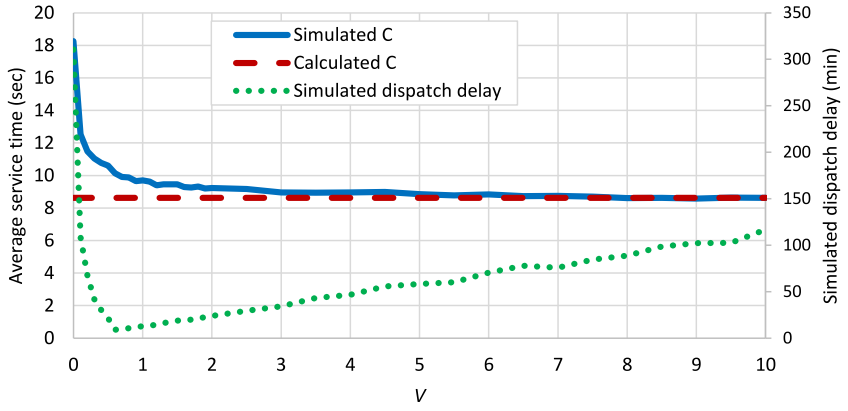

Fig. 5. Change in $\bar{C}$ and average dispatch delay with respect to parameter V in the MDPP.

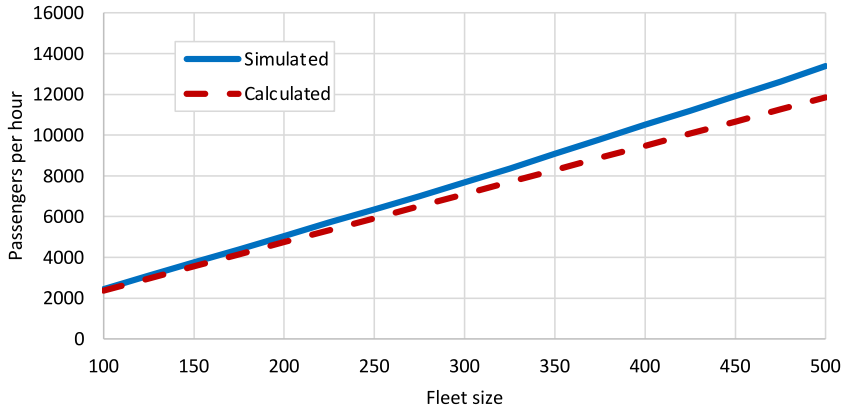

Fig. 6. Simulated and calculated maximum stable demand (using fixed origin–destination proportions) for the Sioux Falls network with public transit integration.

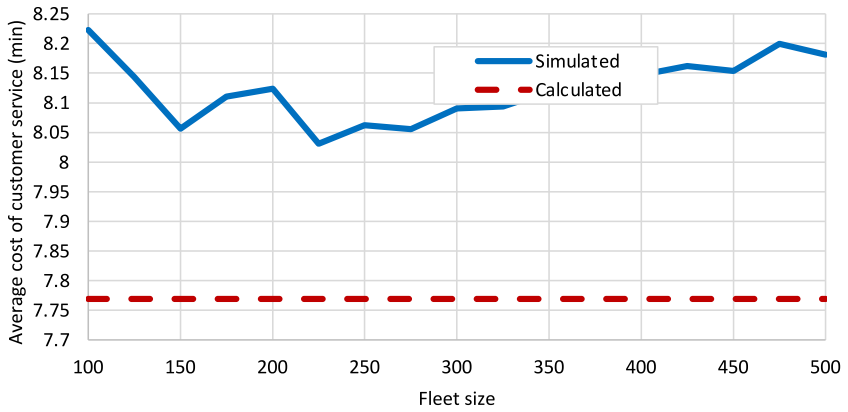

Fig. 7. Simulated and calculated average $\bar{C}$ for the Sioux Falls network with public transit integration.

already. Although the fleet sizes in Fig. 7 are the same as shown in Fig. 4, the number of customers included is much higher because public transit admits many more passengers served. More total customers reduces the error due to assigning discrete SAVs to discrete customers and also reduces the effects of stochastic demand on the simulated $\bar{C}$.

5.3.3. *Electric SAVs*

We next compare the predicted stable region with the stable demand in simulation for electric SAVs. As shown in Fig. 2, recharging stations were placed at nodes 4 and 19, which required travel to and from these nodes for charging. SAV battery capacity

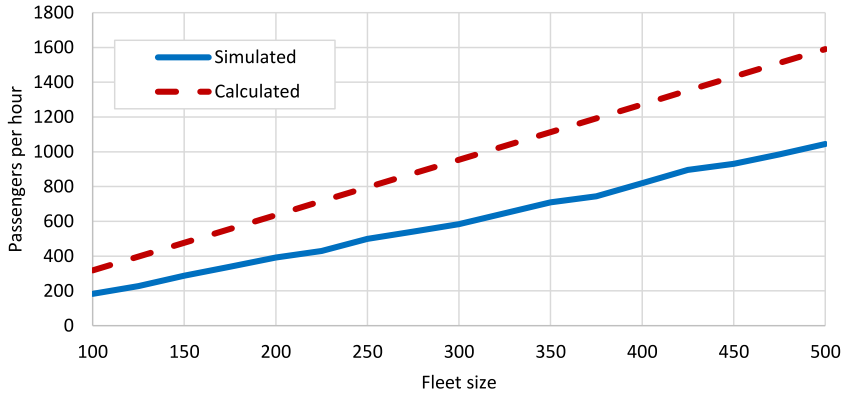

Fig. 8. Simulated and calculated maximum stable demand (using fixed origin–destination proportions) for the Sioux Falls network with electric SAVs.

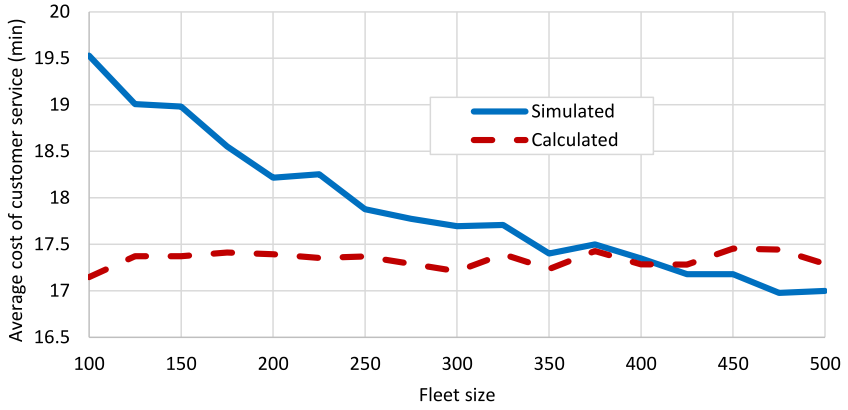

Fig. 9. Simulated and calculated average $\bar{C}$ for the Sioux Falls network with electric SAVs.

was set at 80 miles, with 1 h required to recharge. These values are based on parameters in Chen et al. (2016), but are not intended to provide a prediction of current or future technology. Instead, the goal is to choose values from the literature that provide a significant and easily discernable impact on SAV service.

Fig. 8 presents the maximum stable demand both from problem (48) and estimated through simulation. Unlike Figs. 3 and 6, Fig. 8 exhibits a significant difference between the simulated and calculated maximum stable demand. This difference is due to how the recharging behavior of electric SAVs introduces large variations in service times. When electric SAVs run low on battery, they must recharge as part of their dispatched assignment. Based on the given parameters, recharging itself can require up to 1 h, but is also combined with empty travel to and from the charging station. When calculating the stable region, this recharging behavior is anticipated and expected. However, for the MDPP, such recharging behavior means that low-battery SAVs will not be dispatched until the HOL time has increased to match, i.e. we need $H_c(t) > C_{v\pi}(t)$ where $C_{v\pi}(t)$ includes the recharging. Such recharging reduces the number of SAVs dispatched, thereby increasing the average HOL time and making stability more difficult to detect. Consequently, the simulations shown in Fig. 8 do not detect stability at the same demand levels predicted by problem (48).

Figs. 9 and 10 further illustrate the issue. In Fig. 9, the average $\bar{C}$ is larger than the calculated $\bar{C}$ for small fleet sizes, then actually decreases below the calculated $\bar{C}$ for $|\mathcal{N}_V| \geq 400$. Fig. 10 shows the time spent traveling empty at different fleet sizes, which is always greater than the calculated empty time. The difference is larger than the empty time without recharging shown in Fig. 4; this is because problem (48) can plan recharging behavior, but the MDPP will not be as successful in planning such recharging behavior unless the number of waiting customers is large. Correspondingly, SAVs have to travel more to get to recharging stations than expected by problem (48) unless the simulation runs sufficiently long. The 12 h of simulation used in these results was not sufficiently long to converge to average behavior. When the simulated $\bar{C}$ is smaller than the calculated $\bar{C}$ in Fig. 9, it indicates larger stochasticity in demand and the corresponding recharging requirements or a lack of recharging resulting from low-battery vehicles no longer being assigned to customers by the MDPP due to their high dispatch cost. Either way, these results suggest that achieving the throughput predicted by problem (48) may come with unacceptable level-of-service for customers. Planning for a margin between the expected demand and the maximum throughput may provide smaller waiting times. Furthermore, the dispatch of electric SAVs to recharging stations by the MDPP as formulated in this paper or by Li et al. (2021a) could be further improved.

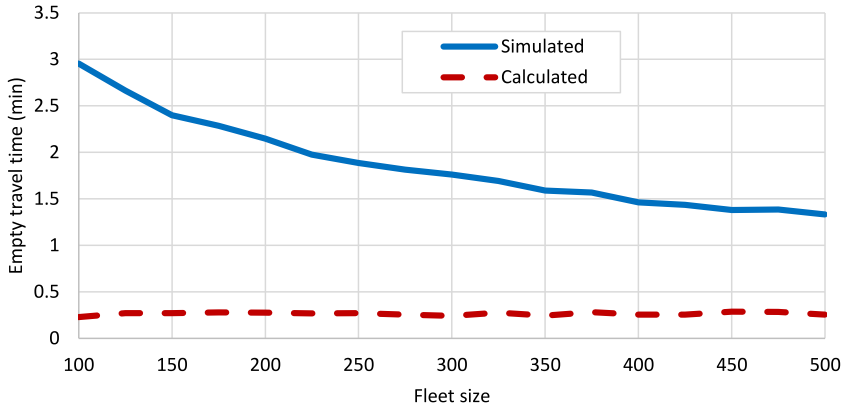

Fig. 10. Simulated and calculated average empty travel time for the Sioux Falls network with electric SAVs.

Table 2
Maximum stable demand on the Sioux Falls network.

SAV system	Maximum stable demand	$\bar{C}$	Empty travel time	Recharging time
Basic SAV	6.95 pass/h/SAV	8.62 min	0.00 min	N/A
Ridesharing	10.64 pass/h/SAV	11.07 min	0.14 min	N/A
Electric SAVs	3.18 pass/h/SAV	17.15 min	0.23 min	8.85 min
SAVs with public transit	23.70 pass/h/SAV	7.77 min	2.64 min	N/A

5.3.4. Summary of Sioux Falls results

Table 2 summarizes the predictions of the stable region calculations. The demand appears to be almost symmetric as the basic SAV system theoretically requires less than 0.01 min of empty time. Electric vehicle recharging results in a massive decrease in the customers per hour served per SAV. That decrease is almost entirely due to the large time spent recharging. Despite only two nodes having charging stations in these experiments, the empty time increases only slightly. This can be accomplished by having SAVs visiting the charging stations while serving nearby customers. The 80 mile battery range is sufficient for multiple customer trips before recharging is required, which also reduces the charging time spent per customer trip. The large recharging time of 21.74 min per trip reflects the 4hr recharging time required to completely charge an empty battery. However, note that these results are heavily dependent on the SAV battery capacity and their recharging rate. As these parameters improve, the impacts of recharging on customer service should decrease.

We attempted to investigate the stable region of ridesharing through simulation. Although 1 simulation of ridesharing can be performed easily, the computation times were prohibitive for a line search with 10 Monte Carlo repetitions. This occurs because of the large number of ridesharing paths; we found that 260,680 paths with 2 customers could be created where no customer exceeds the direct travel time by more than 20%. Consequently, we only report the calculated stable region in Table 2. Ridesharing included at most 2 customers, and each customer could be inconvenienced by at most 20% of their normal time. Nevertheless, ridesharing was effective at increasing the passengers served per hour. The average trip cost increased from 8.62 min to 11.07 min, but the addition of multiple customers per trip for some trips more than made up for that increase.

When public transit is integrated, the number of customers served per SAV increases significantly, but much of this increase is due to customers that can be served by transit alone. We still calculated the customers served per hour using the maximum α from problem (48) although some customer nodes do not require SAV service at all. The time per customer service decreases only slightly, but the SAV trips are no longer symmetric and necessitated greater empty time.

5.4. Estimating the stable region on downtown Austin network

We now use Eqs. (10) and (13) to explore how different SAV behaviors affect the stable region on a realistic city network. We use the downtown Austin network since it includes realistic bus routes. The downtown Austin network, which has 171 zones, 546 intersections, 1247 links, was calibrated to match observed AM peak data in 2011 by the Network Modeling Center of The University of Texas at Austin. Levin et al. (2017) provides more information on this network. However, it was too large to solve problem (48) with ridesharing; the problem size exceeded the available computer memory. Even solving problem (48) with a large number of discrete battery levels was exceeding memory requirements. With only 2 battery levels, CPLEX was unable to solve problem (48) on a Windows computer with 16 GB of memory. These results suggest that solving problem (48) on city-size networks with EVs or ridesharing may require further work to be useful.

Table 3 reports the customers served per hour per SAV for different types of SAV systems. Because of Proposition 5, this number is constant with respect to the fleet size. Table 3 also shows how using SAVs as a first-mile collector service for public transit can increase the demand served by the combined SAV-public transit system. Because the network models the AM peak and most bus

Table 3
Maximum stable demand on the downtown Austin network.

SAV system	Maximum stable demand	$\bar{C}$	Empty time
Basic SAV	5.88 pass/h/SAV	9.75 min	2.49 min
SAVs with public transit	10.02 pass/h/SAV	5.97 min	1.94 min

routes in the downtown Austin network are designed to take travelers downtown, we focused on a first-mile collector service but not a last-mile service. We assumed a 0.25mile walking distance to and from bus stops. For simplicity, we did not consider transfers between bus routes. Instead, travelers transferred from an SAV to a bus route, or traveled directly there on SAV, whichever had less SAV travel time. As expected, integration with public transit increases the maximum stable demand and decreases the average time spent per customer. However, the improvement is less than observed in Sioux Falls. Part of the reason is that our experimental version of Sioux Falls had a very extensive public transit system to demonstrate the benefits of integration. Another possibility is that the larger geographical area is not fully covered with the bus system, so more SAV trips are used.

6. Conclusions

The definition of the stable region $\boldsymbol{\Lambda}$ for general types of SAV systems is particularly useful because it characterizes whether a given demand rate $\bar{\lambda}$ can be served by a fleet size $|\mathcal{N}_V|$. Until now most estimations of $\boldsymbol{\Lambda}$ have been based on simulations and were dependent on the network topology and the specific dispatch policy. Corollary 1 establishes that Eqs. (10)–(13) describe whether a given demand $\bar{\lambda} \in \boldsymbol{\Lambda}$ by showing that these equations identify whether a $\bar{\lambda}$ can be stabilized by the MDPP policy or cannot be stabilized at all. (Proposition 1). We also prove that the MDPP policy serves as many customers as any other dispatch policy. These results hold for a general class of SAV behaviors, including ridesharing, electric SAVs with recharging, and integration with public transit. Numerical results verify the calculations of the maximum stability demand and demonstrate their application on a large city network.

Although the MDPP policy is proven to have maximum stability, it may not minimize waiting times. In particular, SAVs are not dispatched to customers until the weighted cost of dispatch $VC_{v\pi}(t)$ exceeds the customer waiting time $H_c(t)$, which may result in a dispatch delay when $V > 0$. Therefore, it is likely possible to achieve maximum stability with lower customer waiting times. Preemptive rebalancing is also useful for reducing waiting times (Pavone et al., 2012; Fagnant et al., 2015; Hörl et al., 2019), which is included in the stability analysis but not explored directly. The stability properties of other dispatch policies can be evaluated using the concepts developed here, such as evaluating whether stability definition (8) holds analytically or numerically. Fig. 1 demonstrates how to evaluate stability numerically, and Proposition 3 provides one approach to establishing stability properties analytically. Overall, the general work in this paper may lead to new SAV dispatch policies which maximize stability while improving waiting times for customers. However, excessively high waiting times are unrealistic, as customers without an assigned vehicle will eventually exit the system and find alternative transportation. Therefore, future work should also consider stability in models with customer abandonment.

Adding exiting passengers to this methodology requires significant modifications. We can define $z_c(t)$ to be the number of customers abandoning the system, and $\omega_c(t)$ to be the total number of abandoned customers, at node c at time t. Then

$$\omega_c(t+1) = \omega_c(t) + z_c(t) \tag{52}$$

because exited passengers cannot be picked up by an SAV. Then Eq. (1) becomes

$$H_c(t+1) = e_c(t)\left(H_c(t) + 1 - (x_c(t) + z_c(t))\tau_c(t)\right)^+ + (1 - e_c(t))A_c(t) \tag{53}$$

We can redefine stability as there exists a $\kappa < \infty$ such that

$$\limsup_{T\to\infty} \frac{1}{T}\sum_{t=1}^{T}\sum_{c\in\mathcal{N}_C}\left(\mathbb{E}\left[H_c(t)\right] + \mathbb{E}\left[\omega_c(t)\right]\right) \le \kappa \tag{54}$$

Since $\omega_c(t)$ is non-decreasing with time, stability requires that $\mathbb{E}\left[z_c(t)\right] = 0$, i.e. the average number of customers abandoning the system due to lack of service is 0. That definition is consistent with maximum throughput, but it is not clear how to prove that $\mathbb{E}\left[z_c(t)\right] = 0$.

A separate problem exists for calculating $z_c(t)$. We could enforce a maximum waiting time $H_c^{\max}$ by defining $z_c(t)$ as

$$z_c(t) = \begin{cases} 1 & H_c(t) > H_c^{\max} \\ 0 & \text{else} \end{cases} \tag{55}$$

Another possible assumption is that $z_c(t) = A_c(t)$ when the number of waiting customers at c reaches some threshold, which essentially enforces a finite buffer size for customers waiting at c. However, we would need to modify the state to explicitly track the number of waiting customers to determine when the threshold is reached. In summary, it appears that tracking customer abandonment within this model is possible with some modification, but proving stability with customer abandonment is far more difficult. Since customer abandonment is a realistic feature of mobility-on-demand systems, we believe that future work should consider addressing it.

We also noticed another challenge in these analyses: solving linear program (48) can become computationally intensive for more complex SAV systems. In particular, the linear programs for electric SAVs and SAVs with ridesharing exceeded the available computer memory on a typical desktop for the downtown Austin network when using IBM CPLEX. The number of variables increases greatly because of the extra dimensions involved in tracking battery levels or multiple passengers per trip. Further work on simplifying these problems for larger networks could improve the practical utility of this method.

CRediT authorship contribution statement

Michael W. Levin: Conceptualization, Methodology, Software, Analysis, Writing.

Acknowledgment

We gratefully acknowledge the support of the National Science Foundation, United States, Award no. 1935514.

References

Alonso-Mora, J., Samaranayake, S., Wallar, A., Frazzoli, E., Rus, D., 2017a. On-demand high-capacity ride-sharing via dynamic trip-vehicle assignment. Proc. Natl. Acad. Sci. 114 (3), 462–467.
Alonso-Mora, J., Wallar, A., Rus, D., 2017b. Predictive routing for autonomous mobility-on-demand systems with ride-sharing. In: 2017 IEEE/RSJ International Conference on Intelligent Robots and Systems. IROS, IEEE, pp. 3583–3590.
Bauer, G.S., Greenblatt, J.B., Gerke, B.F., 2018. Cost, energy, and environmental impact of automated electric taxi fleets in manhattan. Environ. Sci. Technol. 52 (8), 4920–4928.
Boesch, P.M., Ciari, F., Axhausen, K.W., 2016. Autonomous vehicle fleet sizes required to serve different levels of demand. Transp. Res. Rec. 2542 (1), 111–119.
Boewing, F., Schiffer, M., Salazar, M., Pavone, M., 2020. A vehicle coordination and charge scheduling algorithm for electric autonomous mobility-on-demand systems. In: 2020 American Control Conference. ACC, IEEE, pp. 248–255.
Cáp, M., Alonso Mora, J., 2018. Multi-objective analysis of ridesharing in automated mobility-on-demand. In: Proceedings of RSS 2018: Robotics-Science and Systems XIV.
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. Transp. Res. A Policy Pract. 94, 243–254.
Chen, R., Levin, M.W., 2019. Dynamic user equilibrium of mobility-on-demand system with linear programming rebalancing strategy. Transp. Res. Rec. 2673 (1), 447–459.
Cordeau, J.-F., Laporte, G., 2007. The dial-a-ride problem: models and algorithms. Ann. Oper. Res. 153 (1), 29.
de Souza, F., Gurumurthy, K.M., Auld, J., Kockelman, K.M., 2020. An optimization-based strategy for shared autonomous vehicle fleet repositioning. In: VEHITS. pp. 370–376.
Eksioglu, B., Vural, A.V., Reisman, A., 2009. The vehicle routing problem: A taxonomic review. Comput. Ind. Eng. 57 (4), 1472–1483.
Fagnant, D.J., Kockelman, K.M., 2014. The travel and environmental implications of shared autonomous vehicles, using agent-based model scenarios. Transp. Res. C 40, 1–13.
Fagnant, D.J., Kockelman, K., 2015. Preparing a nation for autonomous vehicles: opportunities, barriers and policy recommendations. Transp. Res. A Policy Pract. 77, 167–181.
Fagnant, D.J., Kockelman, K.M., 2018. Dynamic ride-sharing and fleet sizing for a system of shared autonomous vehicles in Austin, Texas. Transportation 45 (1), 143–158.
Fagnant, D.J., Kockelman, K.M., Bansal, P., 2015. Operations of shared autonomous vehicle fleet for Austin, Texas, market. Transp. Res. Rec. 2563 (1), 98–106.
Ge, Q., Han, K., Liu, X., 2021. Matching and routing for shared autonomous vehicles in congestible network. Transp. Res. E Logist. Transp. Rev. 156, 102513.
Gibbs, S., 2017. Google sibling waymo launches fully autonomous ride-hailing service. Guardian 7.
Greenblatt, J.B., Shaheen, S., 2015. Automated vehicles, on-demand mobility, and environmental impacts. Curr. Sustain. Renew. Energy Rep. 2 (3), 74–81.
Guériau, M., Cugurullo, F., Acheampong, R.A., Dusparic, I., 2020. Shared autonomous mobility on demand: A learning-based approach and its performance in the presence of traffic congestion. IEEE Intell. Transp. Syst. Mag. 12 (4), 208–218.
Guo, X., Caros, N.S., Zhao, J., 2021. Robust matching-integrated vehicle rebalancing in ride-hailing system with uncertain demand. Transp. Res. B 150, 161–189.
Gurumurthy, K.M., Kockelman, K.M., 2018. Analyzing the dynamic ride-sharing potential for shared autonomous vehicle fleets using cellphone data from Orlando, Florida. Comput. Environ. Urban Syst. 71, 177–185.
Gurumurthy, K.M., Kockelman, K.M., Zuniga-Garcia, N., 2020a. First-mile-last-mile collector-distributor system using shared autonomous mobility. Transp. Res. Rec. 2674 (10), 638–647.
Gurumurthy, K.M., de Souza, F., Enam, A., Auld, J., 2020b. Integrating supply and demand perspectives for a large-scale simulation of shared autonomous vehicles. Transp. Res. Rec. 2674 (7), 181–192.
Hörl, S., Ruch, C., Becker, F., Frazzoli, E., Axhausen, K.W., 2019. Fleet operational policies for automated mobility: A simulation assessment for Zurich. Transp. Res. C 102, 20–31.
Huang, Y., Kockelman, K.M., Garikapati, V., 2022. Shared automated vehicle fleet operations for first-mile last-mile transit connections with dynamic pooling. Comput. Environ. Urban Syst. 92, 101730.
Iacobucci, R., McLellan, B., Tezuka, T., 2019. Optimization of shared autonomous electric vehicles operations with charge scheduling and vehicle-to-grid. Transp. Res. C 100, 34–52.
Iglesias, R., Rossi, F., Zhang, R., Pavone, M., 2019. A BCMP network approach to modeling and controlling autonomous mobility-on-demand systems. Int. J. Robot. Res. 38 (2–3), 357–374.
Jäger, B., Agua, F.M.M., Lienkamp, M., 2017. Agent-based simulation of a shared, autonomous and electric on-demand mobility solution. In: 2017 IEEE 20th International Conference on Intelligent Transportation Systems. ITSC, IEEE, pp. 250–255.
Jones, E.C., Leibowicz, B.D., 2019. Contributions of shared autonomous vehicles to climate change mitigation. Transp. Res. D Trans. Environ. 72, 279–298.
Kang, D., Levin, M.W., 2021. Maximum-stability dispatch policy for shared autonomous vehicles. Transp. Res. B 148, 132–151.
Laporte, G., 1992. The vehicle routing problem: An overview of exact and approximate algorithms. European J. Oper. Res. 59 (3), 345–358.
Levin, M.W., 2017. Congestion-aware system optimal route choice for shared autonomous vehicles. Transp. Res. C 82, 229–247.

Levin, M.W., Kockelman, K.M., Boyles, S.D., Li, T., 2017. A general framework for modeling shared autonomous vehicles with dynamic network-loading and dynamic ride-sharing application. Comput. Environ. Urban Syst. 64, 373–383.
Levin, M.W., Odell, M., Samarasena, S., Schwartz, A., 2019. A linear program for optimal integration of shared autonomous vehicles with public transit. Transp. Res. C 109, 267–288.
Li, Q., Liao, F., 2020. Incorporating vehicle self-relocations and traveler activity chains in a bi-level model of optimal deployment of shared autonomous vehicles. Transp. Res. B 140, 151–175.
Li, Y., Long, J., Yu, M., 2021a. A time-dependent shared autonomous vehicle system design problem. Transp. Res. C 124, 102956.
Li, L., Pantelidis, T., Chow, J.Y., Jabari, S.E., 2021b. A real-time dispatching strategy for shared automated electric vehicles with performance guarantees. Transp. Res. E Logist. Transp. Rev. 152, 9292.
Loeb, B., Kockelman, K.M., 2019. Fleet performance and cost evaluation of a shared autonomous electric vehicle (SAEV) fleet: A case study for Austin, Texas. Transp. Res. A Policy Pract. 121, 374–385.
Loeb, B., Kockelman, K.M., Liu, J., 2018. Shared autonomous electric vehicle (SAEV) operations across the Austin, Texas network with charging infrastructure decisions. Transp. Res. C 89, 222–233.
Lokhandwala, M., Cai, H., 2018. Dynamic ride sharing using traditional taxis and shared autonomous taxis: A case study of NYC. Transp. Res. C 97, 45–60.
Moreno, A.T., Michalski, A., Llorca, C., Moeckel, R., 2018. Shared autonomous vehicles effect on vehicle-km traveled and average trip duration. J. Adv. Transp. 2018.
Narayanan, S., Chaniotakis, E., Antoniou, C., 2020. Shared autonomous vehicle services: A comprehensive review. Transp. Res. C 111, 255–293.
Neely, M.J., 2006. Energy optimal control for time-varying wireless networks. IEEE Trans. Inform. Theory 52 (7), 2915–2934.
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.
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.
Salazar, M., Tsao, M., Aguiar, I., Schiffer, M., Pavone, M., 2019. A congestion-aware routing scheme for autonomous mobility-on-demand systems. In: 2019 18th European Control Conference. ECC, IEEE, pp. 3040–3046.
Shen, Y., Zhang, H., Zhao, J., 2018. Integrating shared autonomous vehicle in public transportation system: A supply-side simulation of the first-mile service in Singapore. Transp. Res. A Policy Pract. 113, 125–136.
Skordilis, E., Hou, Y., Tripp, C., Moniot, M., Graf, P., Biagioni, D., 2021. A modular and transferable reinforcement learning framework for the fleet rebalancing problem. IEEE Trans. Intell. Transp. Syst..
Spieser, K., Treleaven, K., Zhang, R., Frazzoli, E., Morton, D., Pavone, M., 2014. Toward a systematic approach to the design and evaluation of automated mobility-on-demand systems: A case study in Singapore. In: Road Vehicle Automation. Springer, pp. 229–245.
Tassiulas, L., Ephremides, A., 1992. Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks. IEEE Trans. Automat. Control 37 (12), 1936–1948.
Tirachini, A., Gomez-Lobo, A., 2020. Does ride-hailing increase or decrease vehicle kilometers traveled (VKT)? A simulation approach for Santiago de Chile. Int. J. Sustain. Transp. 14 (3), 187–204.
Toth, P., Vigo, D., 2002. The Vehicle Routing Problem. SIAM.
Tsao, M., Iglesias, R., Pavone, M., 2018. Stochastic model predictive control for autonomous mobility on demand. In: 2018 21st International Conference on Intelligent Transportation Systems. ITSC, IEEE, pp. 3941–3948.
Tsao, M., Milojevic, D., Ruch, C., Salazar, M., Frazzoli, E., Pavone, M., 2019. Model predictive control of ride-sharing autonomous mobility-on-demand systems. In: 2019 International Conference on Robotics and Automation. ICRA, IEEE, pp. 6665–6671.
Varaiya, P., 2013. Max pressure control of a network of signalized intersections. Transp. Res. C 36, 177–195.
Vosooghi, R., Puchinger, J., Jankovic, M., Vouillon, A., 2019. Shared autonomous vehicle simulation and service design. Transp. Res. C 107, 15–33.
Wen, J., Chen, Y.X., Nassir, N., Zhao, J., 2018. Transit-oriented autonomous vehicle operation with integrated demand-supply interaction. Transp. Res. C 97, 216–234.
Winter, K., Cats, O., Martens, K., van Arem, B., 2021a. Parking space for shared automated vehicles: How less can be more. Transp. Res. A Policy Pract. 143, 61–77.
Winter, K., Cats, O., Martens, K., van Arem, B., 2021b. Relocating shared automated vehicles under parking constraints: assessing the impact of different strategies for on-street parking. Transportation 48 (4), 1931–1965.
Wollenstein-Betech, S., Houshmand, A., Salazar, M., Pavone, M., Cassandras, C.G., Paschalidis, I.C., 2020. Congestion-aware routing and rebalancing of autonomous mobility-on-demand systems in mixed traffic. In: 2020 IEEE 23rd International Conference on Intelligent Transportation Systems. ITSC, IEEE, pp. 1–7.
Zhang, T.Z., Chen, T.D., 2020. Smart charging management for shared autonomous electric vehicle fleets: A puget sound case study. Transp. Res. D Trans. Environ. 78, 102184.
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: Robotics and Automation (ICRA), 2016 IEEE International Conference on. IEEE, pp. 1382–1389.
Zhang, R., Spieser, K., Frazzoli, E., Pavone, M., 2015. Models, algorithms, and evaluation for autonomous mobility-on-demand systems. In: American Control Conference (ACC), 2015. IEEE, pp. 2573–2587.