A general framework for modeling shared autonomous vehicles

Michael W. Levin (corresponding author)
Doctoral Candidate
Department of Civil, Architectural and Environmental Engineering The University of Texas at Austin
301 E. Dean Keeton St. Stop C1761
Austin, TX 78712-1172
michaellevin@utexas.edu
Ph: 512-471-3548, FAX: 512-475-8744

Tianxin Li
Graduate Research Assistant
Department of Civil, Architectural and Environmental Engineering The University of Texas at Austin

Stephen D. Boyles
Assistant Professor
Department of Civil, Architectural and Environmental Engineering The University of Texas at Austin

Kara M. Kockelman
Professor
Department of Civil, Architectural and Environmental Engineering The University of Texas at Austin

Published in *Computers, Environment and Urban Systems* 64: 373-383 (2017)

Abstract

Shared autonomous vehicles (SAVs) could provide low-cost service to travelers and possibly replace the need for personal vehicles. Previous studies found that each SAV could service multiple travelers, thereby eliminating several personal vehicle trips. However, a major limitation of previous studies is the reliance on custom software packages with unrealistic congestion models, network structures, or travel demand. For effective comparisons with personal vehicle scenarios, a common traffic flow simulator is necessary.

This paper presents an event-based framework for implementing SAV behavior in existing traffic simulation models. We demonstrate this framework in a cell transmission model-based dynamic network loading simulator. We also study a heuristic approach for dynamic ride-sharing. We compared personal vehicles and SAV scenarios on the downtown Austin city network. Without dynamic ride-sharing, the additional empty repositioning trips made by SAVs increased congestion and travel times. However, dynamic ride-sharing resulted in

travel times comparable to those of personal vehicles because ride-sharing reduced vehicular demand. Overall, the results show that realistic traffic flow models should be used for studying SAVs, but with well-chosen SAV fleets and routing algorithms, SAVs could provide acceptable service to travelers.

1 Introduction

Autonomous vehicles could revolutionize transportation. Adaptive cruise control could increase road capacity [16, 26] and reservation-based intersection control [7, 8] could do the same for intersections [13, 22]. The focus of this paper is on integrating models of these traffic flow improvements with shared autonomous vehicle (SAV) behavior. SAVs are a fleet of autonomous SAVs that provide low-cost service to travelers, possibly replacing the need for personal vehicles. Previous studies [1, 10] assuming that all travelers used SAVs found that each SAV could service multiple travelers, reducing the number of vehicles needed in the SAV fleet. Although 100% SAV use is unlikely to occur in the near future, previous results suggest great potential benefits when 100% SAVs becomes viable. Strategies such as preemptive relocation of SAVs for expected demand [10] or dynamic ride-sharing [11] are additional options for improving service.

However, a major limitation of previous studies is that many relied on custom software packages with unspecified or unrealistic congestion models [1,10,11,27] and/or grid networks [10,11]. Although these were important studies for technology demonstration purposes, they lacked realistic flow models. Many studies even assumed that link travel times were constant. This limitation prevents prevent accurate predictions of the benefits of SAVs.

It is clear from a review of previous work that a method of integrating SAVs with realistic congestion models is a common issue without an obvious solution. Moreover, because researchers and practitioners use a variety of traffic models, it is desirable for SAVs to be able to be integrated within their preferred flow model. We address this problem by developing an event-based framework for adding SAVs to a general class of existing traffic simulators. To further justify this framework, we also present results from a calibrated city network demonstrating that not using realistic congestion models can greatly exaggerate the potential benefits of SAVs.

This framework admits a dynamic network loading model of SAVs using the well-established cell transmission model (CTM) [5,6]. We compare SAVs using heuristics for vehicle routing and dynamic ride-sharing based on previous work [10, 11] against personal vehicle scenarios. (Heuristics are used because the vehicle-routing problem is NP-hard [28].) The framework allows us to study SAV behaviors using a more realistic congestion model.

The contributions of this paper are as follows:

1. We propose an event-based framework for implementing SAVs in existing traffic models. This can be adapted for macro-, meso-, or micro-scopic flow models. Our results show that SAVs can cause significant congestion, so using realistic traffic flow models is necessary for accurate estimations of SAV level of service. Therefore, future work on

SAVs should consider using this framework or others to incorporate realistic network models.

2. We demonstrate this framework by studying congestion when SAVs are used to service all travelers, using CTM to propagate flow. We also describe and study a heuristic for dynamic ride-sharing on the downtown Austin city network and compare it with personal vehicle results from dynamic traffic assignment (DTA).

3. We compare SAV scenarios (including dynamic ride-sharing), with personal vehicle scenarios on the calibrated downtown Austin city network. Overall, results show that a smaller SAV fleet can service all travel demand in the AM peak. However, some SAV scenarios also increased congestion because of the additional trips made to reach travelers' origins. Therefore, it is important to model congestion when studying SAVs to attain realistic estimates of quality of service. Furthermore, SAVs may be less effective than previously predicted for peak hour scenarios. Nevertheless, SAVs with dynamic ride-sharing provided service comparable to personal vehicles.

The remainder of this paper is organized as follows: Section 2 discusses recent developments in AV traffic flow and SAV modeling. Section 3 describes a general framework for SAVs. In Section 4, we describe specific behaviors used in our case study. We present experimental results for SAVs and compare with personal vehicle scenarios in Section 5. Section 6 presents our conclusions.

2 Literature review

SAVs differ from personal vehicles as follows:

- With personal vehicles, each traveler drives a vehicle from the origin to the destination, then is assumed to park at the destination. Travelers choose routes to minimize their own travel time, resulting in a dynamic user equilibrium (DUE) in which no vehicle can improve travel cost by changing routes.

- With SAVs, all travelers are serviced by SAVs, and no personal vehicles are used. When travel demand is ready to depart, an SAV drives to the origin, takes the traveler to the destination, and then becomes available to service other demand. This may result in some empty repositioning trips to reach travel demand, but the total number of vehicles on the road may be reduced.

Naturally, SAV behavior raises cost and security issues. SAVs are essentially a fleet of driverless taxis, and replacing personal vehicles with taxis is not cost-effective for most travelers. However, because SAVs are driverless, the cost of travel is much less and is more similar to the costs of vehicle ownership [12]. SAVs may also raise security concerns due to their vulnerability to hacking. However, security issues with SAV implementation are outside the scope of this paper. Complete replacement of personal vehicles by SAVs has been studied

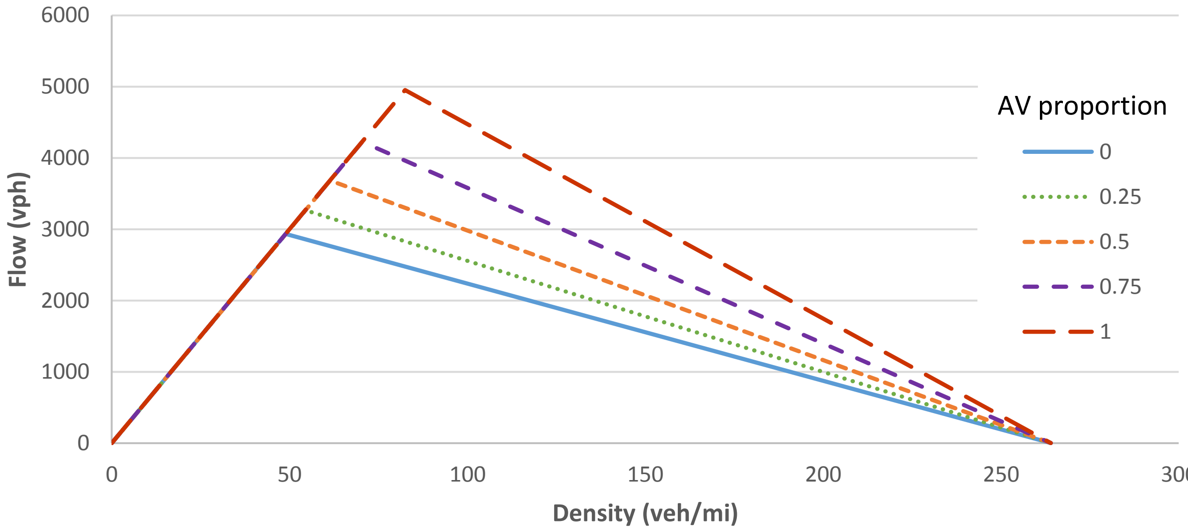

Figure 1: Flow-density relationship as a function of AV proportion for a free flow speed of 60 mph [19]

by previous work [11, 12], and the purpose of this paper is to improve the accuracy of such models.

This paper builds on previous work on AV traffic flow and intersection control models (Section 2.1) and SAVs (Section 2.2) to model SAV behavior.

2.1 Traffic models of autonomous vehicles

After years of development culminating in AV testing on public roads, the literature has begun to focus on modeling new traffic behaviors available to AVs. Adaptive cruise control could increase capacity [16,26] and traffic flow stability [21,25]. However, Levin & Boyles [17] showed that increased road capacity may be offset by greater travel demand, particularly for empty repositioning trips. Therefore, the flow-density relationship is likely to change in space and time with the proportion of AVs. Levin & Boyles [19], developed a multiclass hydrodynamic theory with varying flow-density relationship, and solved it using a multiclass extension of the cell transmission model [5, 6]. Furthermore, they proposed a first-order car-following model to predict the flow-density relationship as a function of the proportion of AVs, with an example shown in Figure 1.

Dresner & Stone [7, 8] developed *reservation-based intersection control*: vehicles communicate wirelessly with an intersection manager to reserve a space-time path through the intersection. The intersection manager simulates the path on a grid of tiles and accepts the request only if it does not conflict with the reservations of other vehicles. Reservations make greater use of intersection capacity, allowing reductions in delays beyond optimized traffic signals in some scenarios [13, 22]. However, due to the computational complexity of the

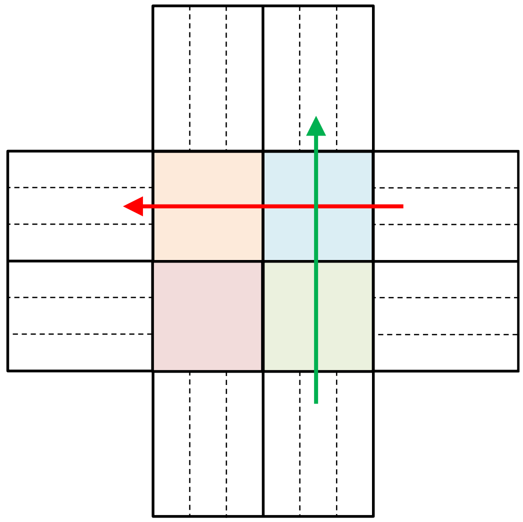

Figure 2: Conflict region model

reservation protocol, many previous studies have been limited to small networks [15] or used simplified intersection models that reduced the traffic efficiency [2, 3]. Levin & Boyles [18] developed the conflict region model of reservations, which is tractable for large-scale DTA, and is therefore used in the simulations in this paper. Instead of simulating vehicle paths along a fine grid of tiles, the conflict region model aggregates tiles into larger *conflict regions*. The conflict regions for a four-way intersection are illustrated in Figure 2. Vehicle turning movements are limited by the capacity of all conflict regions the vehicle must pass through. Different turning movements pass through different sets of conflict regions; for example, left-turning traffic passes through more conflict regions than right-turning traffic.

2.2 Shared autonomous vehicles

Multiple studies have investigated the possibility of using a fleet of SAVs to reduce reliance on personal vehicles and improve mobility and safety [9]. Fagnant & Kockelman [10] estimated that one SAV could provide service to around eleven travelers on a grid network approximation of Austin, Texas with most travelers waiting at most 5 minutes for pick-up, although vehicle travel time increased. Fagnant & Kockelman [11] incorporated dynamic ride-sharing, and found that it could offset the additional vehicle travel time. However, only 10% of personal trips of Austin were included. Further studies on different cities have supported indications that a smaller fleet of SAVs could provide service to all travelers. Burns et al. [1] studied a centrally dispatched SAV system in three different urban and suburban environments. Their findings indicated that a much smaller fleet of SAVs could provide service to all residents with acceptable waiting times. Also, a slightly reduced fleet of taxicabs could improve on wait times and vehicle utilization in Manhattan, New York. Spieser et al. [27] found that a SAV fleet one-third the size of the personal vehicle fleet was sufficient for providing service to Singapore travelers.

Although the results of previous studies are encouraging, they relied on unrealistic traffic congestion models, such as using fixed link travel times [10, 12, 27]. In addition, several studies used grid-based network approximations of cities [10]. SAVs could actually increase the number of trips, as well as vehicle miles traveled, by making repositioning trips to reach new travelers. These increases in demand could result in significantly higher congestion in saturated urban cities. Unfortunately, due to the lack of realistic congestion modeling, the traffic congestion and traveler service benefits of SAVs reported by previous studies may be greatly exaggerated. The lack of realistic congestion models across most previous studies indicates that the problem of integrating SAVs with established traffic flow models does not have an obvious solution. Therefore, this paper presents an event-based framework to build an SAV simulation on top of a general class of existing traffic simulators. We hope this will encourage future studies on SAVs to use more realistic congestion models to obtain more accurate predictions.

3 Shared autonomous vehicle framework

This section presents a general framework for dynamic simulation of SAVs to admit the latest developments in traffic flow modeling and SAV behavior. The framework is built on two events that can be integrated into most existing simulation-based traffic models. The purpose of this framework is to encourage future studies on SAVs to make use of existing traffic models for effective comparisons with current traffic conditions. As the case study will demonstrate, replacing personal vehicles with SAVs for the same number of travelers could increase congestion. To determine whether SAVs are beneficial, it is therefore necessary to compare SAV and personal vehicle scenarios in the same traffic model.

This section discusses the key events defining this framework and the types of responses they warrant. However, the specific responses depend on the dispatcher logic, and for generality this framework does not require specific dispatcher behaviors. Section 4 discusses the dispatcher logic used in our case study, including dynamic ride-sharing.

This framework is based on a traffic simulator operating on a *network* $G = (N, A, Z, V, D)$, where N is the set of nodes, A is the set of links, and $Z \subset N$ is the set of centroids. The network has a set of SAVs V that provide service to the demand D. Note that D is in terms of person trips, not vehicle trips, since travelers will be serviced by SAVs. The integration of the framework with the traffic simulator is illustrated through the simulator logic in Figure 3, with simulator time t and time step Δt. Events and responses are indicated with double lines; the remainder is the standard traffic simulator. The simulation steps are grouped into three modules: 1) demand; 2) SAV dispatcher; and 3) traffic flow simulator. The remainder of this section discusses these modules in greater detail.

3.1 Demand

The demand module introduces demand into the simulation. At each time t, the demand module outputs the set of travelers that request a SAV at t. (This does not include waiting

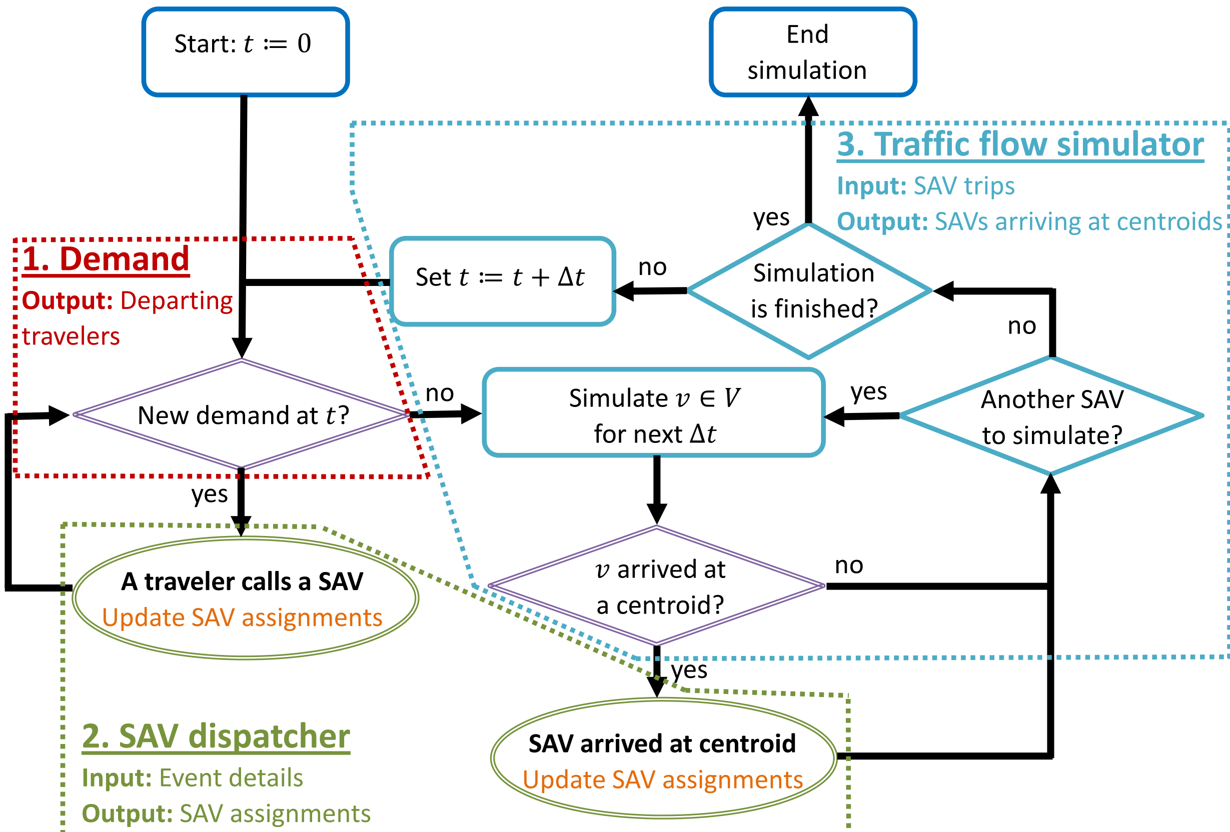

Figure 3: Event-based framework integrated into traffic simulator

travelers.) The demand module of existing traffic simulators may be adapted for this purpose, with the caveat that the demand is in the form of travelers, not personal vehicles. If new demand appears at t, this triggers the corresponding event: a traveler calls a SAV.

Because SAV actions are triggered by a traveler calling a SAV, this framework admits a very general class of demand models. The major requirement is that demand must be separated into packets that spawn at a specific time with a specific origin and destination. Although this paper primarily refers to demand as individual travelers, these packets could also represent a group of people traveling together. Demand cannot be continuous over time because that would trigger a very large number of events. However, in our case study demand and traffic flow are simulated at a timestep of 6 seconds, which is demonstrated to be computationally tractable for city networks.

As a result, this framework can handle both real-time and pre-simulation demand generation. Real-time demand may be randomly generated every simulation step, triggering the event of a traveler calling a SAV when the demand is created. For models with dynamic demand tables, each packet of demand spawns at its departure time and calls a SAV then. In addition, if demand is assumed to be known prior to its departure time, SAVs may choose to preemptively relocate before the traveler appears. However, this requires that travelers plan ahead to schedule a SAV before they depart. A less restrictive assumption is that the productions at each zone are known, and SAVs may preemptively relocate in response to expected travelers. This requires less specific information about the traveler, and trip productions are usually predicted by metropolitan planning organizations.

3.2 SAV dispatcher

This framework assumes the existence of a central SAV dispatcher that knows the status of all SAVs and can make route and passenger assignments. With the range of wireless communication available today, the existence a central dispatcher is a reasonable assumption for SAVs. However, if desired the dispatcher logic could also be chosen to simulate SAVs making individual decisions on their limited information.

The SAV dispatcher module determines SAV behavior, including trip and route choice, parking, and passenger service assignments. The dispatcher operates as an *event handler* responding to the events of a traveler calling a SAV or a SAV arriving at a centroid, and takes as input the event details. The dispatcher is responsible for ensuring that all active travelers are provided with SAV service.

The output of the dispatcher are the SAV behaviors in response to the event. These include SAV vehicle trips (which are passed to the traffic flow simulator), passenger pick-up and drop-off, and parking SAVs that are not needed. At any given time, each SAV is either parked at a centroid or traveling. If a SAV is parked, its exact location must be known.

This framework is event-based, meaning that SAV actions are assigned when one of the following events occurs:

1. A traveler calls a SAV.

2. A SAV arrives at a centroid.

The first event is triggered in response to demand departing (or requesting to depart), and the second is in response to a SAV completing its assigned trip. These can be implemented in most simulation-based frameworks. Instead of a traveler departing by creating a personal vehicle, the traveler calls a SAV. When a SAV completes travel on a path (which should end in a centroid), this also triggers an event so the simulator can check for arriving or departing passengers at that centroid and assign the SAV on its next trip.

3.2.1 A traveler calls a SAV

When a traveler $d \in D$ calls a SAV, the dispatcher should ensure that the demand will be satisfied by a SAV. This could occur in several ways:

1. If an empty SAV $v \in V$ is parked at d's origin, the dispatcher might assign v to immediately pick up d.

2. If an empty SAV $v \in V$ is parked elsewhere, the dispatcher may assign v to travel to d's origin. In this case, the dispatcher might choose to wait to optimize the movement of SAVs. For instance, Fagnant & Kockelman [10] use a heuristic to move SAVs to a closer waiting traveler rather than the first waiting traveler. The dispatcher might also change the path of a traveling SAV to handle the demand.

3. If a SAV $v \in V$ is inbound to d's location, the dispatcher might assign v to service d if possible. However, the dispatcher should consider v's estimated time of arrival (ETA). If v's ETA results in unacceptable waiting time for d, the dispatcher may also send an empty SAV to d to reduce waiting time.

Regardless of the conditions chosen for each action, the dispatcher must ensure that the demand will be handled.

3.2.2 A SAV arrives at a centroid

When a SAV $v \in V$ arrives at a centroid $i \in Z$, it has finished its assigned trip. This should result in two types of actions. First, if v is carrying any travelers destined for i, they should exit v. Second, the dispatcher should assign v to park at i or depart on another trip. There are several possibilities for this assignment:

1. If v still has passengers, it should continue to the next destination. If ride sharing is allowed and the capacity of v permits it, other passengers at i may wish to take v to reduce their waiting time.

2. If v is empty, and a traveler $d \in D$ is waiting at i for a SAV, it is reasonable to assign v to accept d. v may then proceed directly to d's destination or, if dynamic ride-sharing is allowed, to another centroid to pick up another passenger.

3. If no travelers are waiting at i and v is empty, the dispatcher might assign v to pick up a traveler at a different centroid.

4. The dispatcher could also assign v to wait at i until needed for future demand, contingent on parking availability.

5. Finally, the dispatcher might assign v to preemptively relocate to handle predicted demand.

The conditions given above are reasonable but may not be necessary. Optimizing the assignment of actions for the existing and predicted demand could use the possible actions in different ways. For example, v might be assigned to park at i to wait for the expected demand even if v is already carrying passengers. This optimization problem is similar to the class of vehicle routing problems, which are NP-hard. Therefore, solving this optimization is outside the scope of this paper, but later sections will present a heuristic.

3.3 Traffic flow simulator

The traffic flow simulator takes as input SAV trips and their departure times and determines the arrival times of SAVs at centroids. The primary output of the simulator is to trigger the event that an SAV arrived at a centroid at the appropriate time.

Because the SAV framework is built on the events of a traveler calling a SAV, and a SAV arriving at a centroid, the framework admits many flow propagation models. The major requirement is that the model be integrated into simulation. After departing, a SAV travels along its assigned path until reaching the destination centroid, at which point it triggers the arrival event. Therefore, the framework must track the SAV travel times to determine arrival times, but its travel time may be evaluated by a variety of flow models. For instance, the travel time could be set as a constant or through link performance functions. SAV movement may also be modeled through micro- or meso-simulation. Any stochasticity in the traffic flow model is compatible with this framework because the SAV triggers the event only after it arrives at its destination. Note that this framework is compatible with other vehicles on the road affecting congestion through link performance functions or simulation-based flow propagation.

Therefore, this SAV framework can be implemented with existing traffic models by modifying them to trigger demand and centroid arrival events. To demonstrate this flexibility, the case study implements this framework on the simulation-based DTA model of Levin & Boyles [19].

4 Case study: framework implementation

This section describes the implementation of the SAV framework on a cell transmission model-based traffic simulator. Although Section 3 discussed how to implement SAVs in existing traffic simulators, the responses of the dispatcher to events were not specified for generality. The purpose of this section is to describe the specific traffic flow simulator and dispatcher logic used in our case study, including the heuristics for dynamic ride-sharing. Results using this implementation are presented in Section 5.

This case study assumes that all vehicles are SAVs: travelers do not have personal vehicles available. This was chosen to study the feasibility of switching to an entirely SAV-based travel model. Furthermore, a mix of SAVs and personal vehicles would complicate the route choice. Finding routes for personal vehicles would require solving DTA, and the many simulations needed to solve DTA would add computation time and complexity to the theoretical model.

4.1 Demand

This case study used personal vehicle trip tables from the morning peak to determine SAV traveler demand. Each vehicle trip was converted into a single traveler trip with the same origin, destination, and departure time. Although some of these vehicle trips may encompass multiple person trips, that information was not available. Furthermore, multiple persons using the same vehicle would likely use the same SAV. Therefore, it would only affect situations in which SAV capacity was a limitation, such as dynamic ride-sharing.

For each trip, the demand module creates a traveler at the appropriate time. Although the demand is fixed, the SAV dispatcher is not programmed to take advantage of demand information. The dispatcher only responds to demand when a traveler was created.

In reality, travelers have more choices available. They could request a SAV in advance, specify time windows for departure or arrival, or change their departure time in response to expected travel times.

4.2 Traffic flow simulator

The traffic flow simulator uses the cell transmission model (CTM) [5,6], which is a Godunov approximation [14] to the hydrodynamic theory of traffic flow [23,24]. CTM discretizes links into *cells* of length $u^{\mathrm{f}}\Delta t$, where u^{f} is the free flow speed and Δt is the simulation time step. Thus, vehicles can traverse at most one cell per time step. Congestion waves from bottlenecks or intersections travel backwards through the cells and reduce vehicle speeds. Since AVs increase capacity [16,26], this simulator use the CTM and flow-density relationship developed by Levin & Boyles [19]. Because all vehicles are SAVs, intersections were controlled using the reservation-based protocol of Dresner & Stone [7,8] for AVs. For computational tractability, the simulator used the conflict region node model of reservation-based intersection control proposed by Levin & Boyles [18].

CTM has been used in, and allows direct comparisons with, large-scale mesoscopic DTA simulators [29]. DTA models [4] typically assume that route choice is based on driver experience. Each vehicle individually seeks its shortest route, resulting in a DUE. DTA algorithms typically consist of three steps, performed iteratively, to find a DUE assignment [20]. First, shortest paths are found for all origin-destination pairs. Then, a fraction of demand is assigned to the new shortest paths. Finally, travel times under the new assignment are evaluated through a mesoscopic flow model such as CTM.

Although DUE is based on the analytical static traffic assignment models, it requires further study to be formulated for SAV behavior due to stochasticity in the SAV trip table. We assume that the SAV dispatcher does not know travel demand or SAV travel times perfectly. Therefore, the list of free SAVs at any given time is stochastic, which results in uncertainty in which SAV will be used to service new demand.

Therefore, we use a dynamic network loading (DNL) -based route assignment. Let π_{rs} be the path stored by the dispatcher for travel from r to s. When a SAV departs to travel from r to s, it is assigned to the stored path π_{rs}. During simulation, when $t \equiv 0 \mod \Delta\mathcal{T}$, where $\Delta\mathcal{T}$ is the update interval, π_{rs} is updated to be the shortest path from r to s based on average link travel times over the interval $[t - \Delta\mathcal{T}, t)$. Our experiments use $\Delta\mathcal{T} = 1$ minute. Note that the path update interval ($\Delta\mathcal{T} = 1$ minute) is different from the traffic flow simulation time step ($\Delta t = 6$ seconds).

4.3 SAV dispatcher

This section describes the specific logic used to assign SAVs in our case study. Although this is only a heuristic for the vehicle routing problem of servicing all travelers, vehicle routing problems in general are NP-hard and solving them in real time is unrealistic. Instead, we describe reasonable behaviors that SAVs could choose.

4.3.1 A traveler calls a SAV

When a traveler $d \in D$ calls a SAV at centroid $i \in Z$, the dispatcher first checks whether there are any SAVs already enroute to i. If a SAV enroute to i is free, or will drop off its last passenger at i, and its ETA at i is less than 10 minutes away, that SAV is assigned to service d. This is to reduce congestion resulting from sending more SAVs. (As Section 5 will demonstrate, moving SAVs more frequently can result in a net travel time increase while decreasing waiting times due to congestion.) If there are multiple travelers waiting at i, travelers are serviced in a first-come-first-serve (FCFS) order — with some exceptions for dynamic ride-sharing. Therefore, we look at the ETA of the SAV that would be assigned to d, if one exists.

Otherwise, we search for the parked SAV that is closest (in travel time) to i. If it could arrive sooner than the ETA of the appropriate enroute SAV, it is assigned to travel to i to provide service to d. This is a FCFS policy: the traveler that requests a SAV first will be the first to get picked up, even if the SAV could sooner reach a traveler departing later. Although Fagnant & Kockelman [10] initially restricted SAV assignments to those within 5 minutes of travel to improve the system efficiency, FCFS is also a reasonable policy for dispatching SAVs. If all SAVs are busy, then d is added to the list of waiting travelers $\mathcal{W}$.

4.3.2 A SAV arrives at a centroid

If a SAV $v \in V$ is free after reaching centroid $i \in Z$ (either because v is empty, or because v drops off all passengers at i), and there are waiting travelers at i, then it is assigned to carry the longest waiting traveler. Note that v may not be the same SAV that was dispatched to that traveler. Due to stochasticity in the flow propagation model, it is possible that the order of arrival of SAVs may differ. However, there is no significant difference between two free SAVs in terms of carrying a single traveler. Therefore, we assign them to travelers in FCFS order.

If v still has passengers after reaching i (which is possible when dynamic ride-sharing is permitted), then v is assigned to travel to the next passenger's destination. However, travelers waiting at i have the option of entering v if it makes sense for their destination. This is discussed further in Section 4.4.

If v is free after reaching i and no demand is waiting at i, then v is dispatched to the longest-waiting traveler in $\mathcal{W}$. If multiple SAVs become free at the same time, the one closest to the longest-waiting traveler in $\mathcal{W}$ will be sent. If $\mathcal{W}$ is empty, then v will park at i until needed. We assume for this study that centroids have infinite parking space, as there are no personal vehicles in this network. However, it would be possible to model limited parking by assigning v to travel somewhere else if parking was not available at i.

4.4 Dynamic ride-sharing

We also consider the possibility of dynamic ride-sharing. Following the principle of FCFS, we give precedence to the longest-waiting traveler. However, we allow other passengers to enter the SAV if they are traveling to the same, or a close destination. Specifically, suppose that the SAV $v \in V$ is initially empty, and the longest-waiting traveler at $i \in Z$ is d_0, traveling from i to $j \in Z$. If there is another traveler d_1 also traveling from i to j, then d_1 may take the same SAV. If there is a traveler d_2 traveling from i to $k \in Z$, and there is room in the SAV, d_2 may also take the same SAV if the additional travel time is sufficiently low. Let t_{ij} be the expected travel time from i to j. Then d_2 will take the SAV if $t_{ij} + t_{jk} \leq (1 + \epsilon)t_{ik}$. Otherwise, d_2 will wait at i. If d_2 decides to take the SAV, then any other waiting travelers at i also traveling from i to k may enter the SAV. Although this violates FCFS, this is permitted because it does not impose any additional travel time on the SAV.

This offer is extended, in FCFS order, for all travelers waiting at i until v is full. For instance, suppose a passenger d_3 departing after d_2 is traveling from i to $l \in Z$. Because of FCFS, v must service d_2 first, but if $t_{ij} + t_{jk} + t_{kl} \leq (1 + \epsilon)t_{il}$, then d_3 will still take SAV v from i.

The logic is slightly different when v arrives at i already carrying a passenger. In that case, precedence is given to all passengers already in v because they have been traveling. However, travelers in i may enter v — at the back of the queue — if the additional travel time is less than ϵ of the direct travel time.

The problem of dynamic ride-sharing is a vehicle routing problem with all SAVs. In general, vehicle routing problems can admit solutions in which a SAV picks up several passengers before dropping any off. The heuristic in this case study does not do that due to complexity, although that behavior could certainly be implemented within this framework. In practice, due to the necessity of tractability when solving vehicle routing problems in real-time in response to demand, similar simple heuristics are likely to be used. Even with this restricted form of dynamic ride-sharing, the benefits over non-ride-sharing SAVs are significant, as shown in Section 5.

5 Case study: experimental results

We performed several sets of experiments to study how SAVs (Sections 5.2 through 5.3) perform relative to personal vehicles (Section 5.1), and how the dynamic ride-sharing heuristic affects performance. Our experiments were performed on the downtown Austin network, shown in Figure 4. It consists of a downtown grid with freeway and arterial corridors. It has 171 zones, 546 intersections, 1,247 links, and 62,836 trips over 2 hours in the AM peak. The centroids are significantly disaggregated for this downtown region, so we did not include intra-zonal trips in the trip table. The network was calibrated by the Network Modeling Center to match traffic data from the Capital Area Metropolitan Planning Organization.

This is only a subnetwork of the larger Austin region, which has 1.2 million trips. This subnetwork was used because computation times were around 30–40 seconds per scenario on an Intel Xeon running at 3.33 GHz (implemented in Java), allowing many scenarios to be studied. However, many trips bound for the downtown grid originate from outside the subnetwork region. We approximated them as arriving from one of the subnetwork boundaries. The 62,836 trips within the downtown subnetwork is sufficient for a large-scale, realistic study of SAVs.

Initially, SAVs were distributed proportionally to productions: centroid $i \in Z$ started with $|V|\frac{\mathcal{P}_i}{\sum\limits_{i' \in Z} \mathcal{P}_{i'}}$ parked SAVs, which corresponds to $\Delta V_i = 0$. We assumed that all SAVs could be relocated overnight to fulfill these proportions at the start of the AM peak. (Preemptive relocation is a strategy for relocating SAVs *during* the AM peak — while travelers are requesting SAVs.)

Fagnant & Kockelman [10] used a seeding run to determine the minimum number of SAVs necessary to service all travelers. However, a seeding run may have biased the number of SAVs to be lower. Instead of a seeding run, we performed sensitivity analyses to study how increasing numbers of SAVs affected level of service. In some scenarios (such as dynamic ride-sharing) we observed that fewer numbers of SAVs performed better due to lower congestion. In other scenarios, greater numbers of SAVs improved service. The following charts contain experiments using between 1000 and 60,000 SAVs, with increments of 500. For some scenarios, the range was reduced to numbers of SAVs that could provide service to all travelers within 6 hours because service was limited by having too few SAVs or too much congestion.

5.1 Personal vehicles

For comparison, we also considered two personal vehicle scenarios on the downtown Austin network:

1. All travelers drive personal non-autonomous vehicles. This represents current traffic conditions, and shows

2. All travelers use personal AVs, and use AV capacity and intersection improvements. This is an alternative to SAVs in which travelers own the AVs.

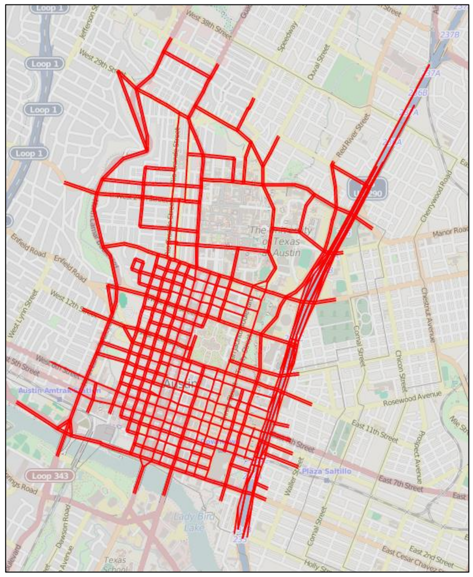

Figure 4: Downtown Austin network

Table 1: Results from personal vehicle scenarios

Scenario	Avg. travel time	Vehicle miles traveled
Personal conventional vehicles	15.24 min	146096 mi
Personal autonomous vehicles	4.12 min	142455 mi

For the private vehicle scenarios, we assumed that travelers chose routes to minimize their own travel time, resulting in a DUE. Therefore, we used DTA to find route choice for personal vehicle scenarios.

One potential issue with comparing these personal vehicle scenarios with SAVs is the different methods used for route choice. For personal vehicles, we assumed DUE behavior, and for SAVs, we assumed DNL behavior determined by the SAV dispatcher. DUE is widely accepted for modeling personal vehicle behavior [4]. DNL was used for SAVs because the SAV dispatcher is modeled to react to travel demand as it appears. Therefore, to handle stochastic demand, the SAV dispatcher should rely on current rather than historical traffic conditions in its route assignments. (Furthermore, a traffic assignment problem has not been formulated for SAVs, and consequently it is not known how to solve DTA for SAVs.)

Results from personal vehicle scenarios are shown in Table 1. Overall, when using personal vehicles with traffic signals, travelers experienced an average travel time of 15.24 minutes. When signals were replaced with reservation controls, average travel times were reduced to 4.12 minutes. Since the adoption of reservation controls may be difficult or inefficient if a significant proportion of personal vehicles are not autonomous, both personal vehicle scenarios may be reasonable for comparison against SAVs. We assume that if SAVs were to replace all personal vehicles, reservation controls would be used.

5.2 Shared autonomous vehicles

The initial SAV scenario did not include dynamic ride-sharing. Figure 5 shows travel time results with 17,500 to 60,000 total SAVs available. Fewer numbers of SAVs were found to be insufficient to service the 2 hours of travel demand after 6 hours. Greater numbers of SAVs reduced both waiting time and in-vehicle travel time. With more SAVs, more vehicles were available near traveler origins, and fewer empty repositioning trips reduced congestion.

As the number of SAVs increased, waiting time decreased consistently, although with diminishing returns. With 39,500 or more SAVs, average waiting times were below 1 minute. Waiting times approached 0 because SAVs were assumed to be initially distributed according to trip productions. Therefore, with 62,836 or more SAVs, waiting times would be 0. Of course, one of the goals of SAVs is to reduce the total number of vehicles in [10].

Because the demand is from the AM peak, much of the waiting time results from SAVs carrying travelers to the downtown region then making an empty repositioning trip to the next traveler's origin. However, waiting times were only 10.3 minutes with 17,500 SAVs. With 25,500 or more SAVs, average waiting times were less than 5 minutes. These average waiting times could be acceptable to travelers.

The average in-vehicle travel time (IVTT) was higher than the personal vehicle scenarios at low numbers of SAVs. This shows that a small SAV fleet requires many empty repositioning trips to service travelers. The empty repositioning trips result in greater demand and therefore congestion. This is particularly relevant for peak hour scenarios, which result in the greatest number of empty repositioning trips because most trips are to or from the central business district. SAV models that do not include realistic travel time predictions would not be able to predict the congestion caused by a small SAV fleet.

This AM peak hour scenario required far more SAVs than 1 per 9.3 travelers [12]. 1 SAV could replace at most 3.6 personal vehicles, and total travel time was significantly higher there. SAV fleet size is likely to be determined by peak hour demand because peak hour travel patterns are the most difficult to serve with SAVs.

However, with only 22,000 SAVs, the average IVTT was less than the personal non-AV scenario of 15.24 minutes (Table 1). The average IVTT never decreased below 9.8 minutes — higher than the 4.12 minutes of the personal AV scenario, but small enough to be feasible for travelers. This was probably due to the route choice heuristic used in this scenario. Personal AVs used DUE behavior, whereas SAVs did not. Better heuristics for SAV routing could therefore decrease the IVTT further for SAVs. Still, the average IVTT was not substantially higher than the personal AV scenario.

Vehicle miles traveled (VMT) and empty VMT — miles traveled while not carrying any passengers — decreased at the same rate as the number of SAVs increased (Figure 5). This indicates that the difference was primarily due to less repositioning trips to pick up the next traveler, rather than changes in route choice. It is intuitive that as the number of SAVs increased, the average distance between a waiting traveler and the nearest (in travel time) available SAV would decrease. The average passenger miles traveled was consistently 2.27 miles.

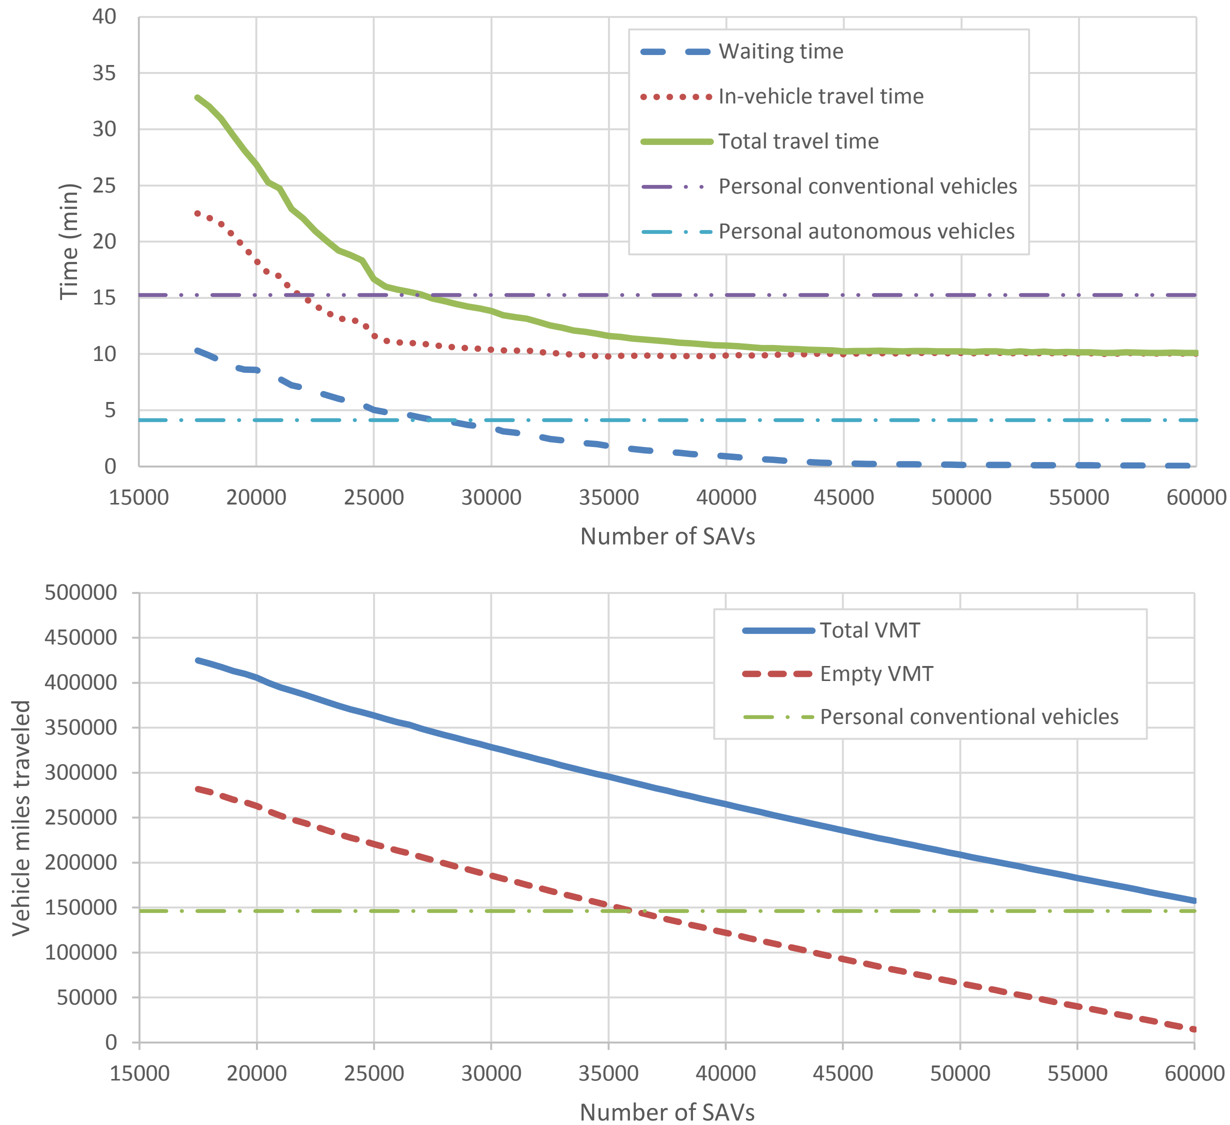

Figure 5: Travel time and VMT for the base SAV scenario

5.3 Dynamic ride-sharing

Dynamic ride-sharing greatly affected level of service for travelers as shown in Figure 6. With dynamic ride-sharing, 1000 SAVs were actually sufficient to service all demand. Each SAV could carry up to 4 passengers, although they would travel with less if no travelers were waiting. However, because most trips were to the central business district, SAVs could easily combine trips because traveler destinations were relatively close. Surprisingly, optimal service was provided with just 2000 SAVs, or a ratio of 1 SAV to 31.4 travelers. This is significantly higher than the 1 SAV to 9.3 travelers [12] although of course here each SAV was probably carrying 3 to 4 passengers.

The least average total travel time was 6.46 minutes with 2000 SAVs, comparable with the 4.12 minutes with the personal AV scenario (Table 1). 5.41 minutes was due to IVTT, with 1.04 minutes due to waiting time. These travel and waiting times might be further reduced with a better heuristic for dynamic ride-sharing. Therefore, with such a low travel time, SAVs with dynamic ride-sharing could be an effective replacement for personal AVs. Furthermore, the size of the SAV fleet used is so small relative to the number of travelers that full replacement might be feasible. The cost per traveler are also likely to be significantly reduced due to car-sharing and the lack of driver. Further study in different demand scenarios and on different networks is needed, but this result suggests that SAVs could be a cost-effective form of paratransit with a high level of service.

Waiting times were consistently low with 2000 or more SAVs. This is probably because most travelers had relatively close destinations, so ride-sharing was frequently used. Strangely, IVTT peaked at 17.54 minutes with 11,000 SAVs. This was likely because SAVs did not wait around for ride-sharing with later-departing travelers. Therefore, the 11,000 SAVs made more trips, carrying fewer travelers per trip, and increased congestion. Figure 7 shows that passenger miles traveled increased as the number of SAVs increased because ride-sharing was used less. With greater than 11,000 SAVs, travel times decreased because less empty repositioning trips were needed, decreasing vehicle demand. VMT, and empty repositioning miles traveled, was highest around 14,500 SAVs (Figure 6). With our heuristic, a fleet of between 5500 and 17,500 SAVs was less efficient than a smaller fleet. Therefore, future work on SAVs should study more effective heuristics for the dynamic ride-sharing problem.

6 Conclusions

This paper presented an event-based framework for implementing SAV behavior in existing traffic simulation models. The framework relies on two events: travelers calling SAVs, and SAVs arriving at centroids, that are orthogonal to traffic flow models. This allows comparisons with personal vehicle scenarios through solving traffic assignment in the same simulator. We implemented this SAV framework on a cell transmission model-based dynamic traffic assignment simulator as well as a heuristic approach to dynamic ride-sharing. Then, we studied replacing personal vehicles with SAVs in the downtown Austin network with AM

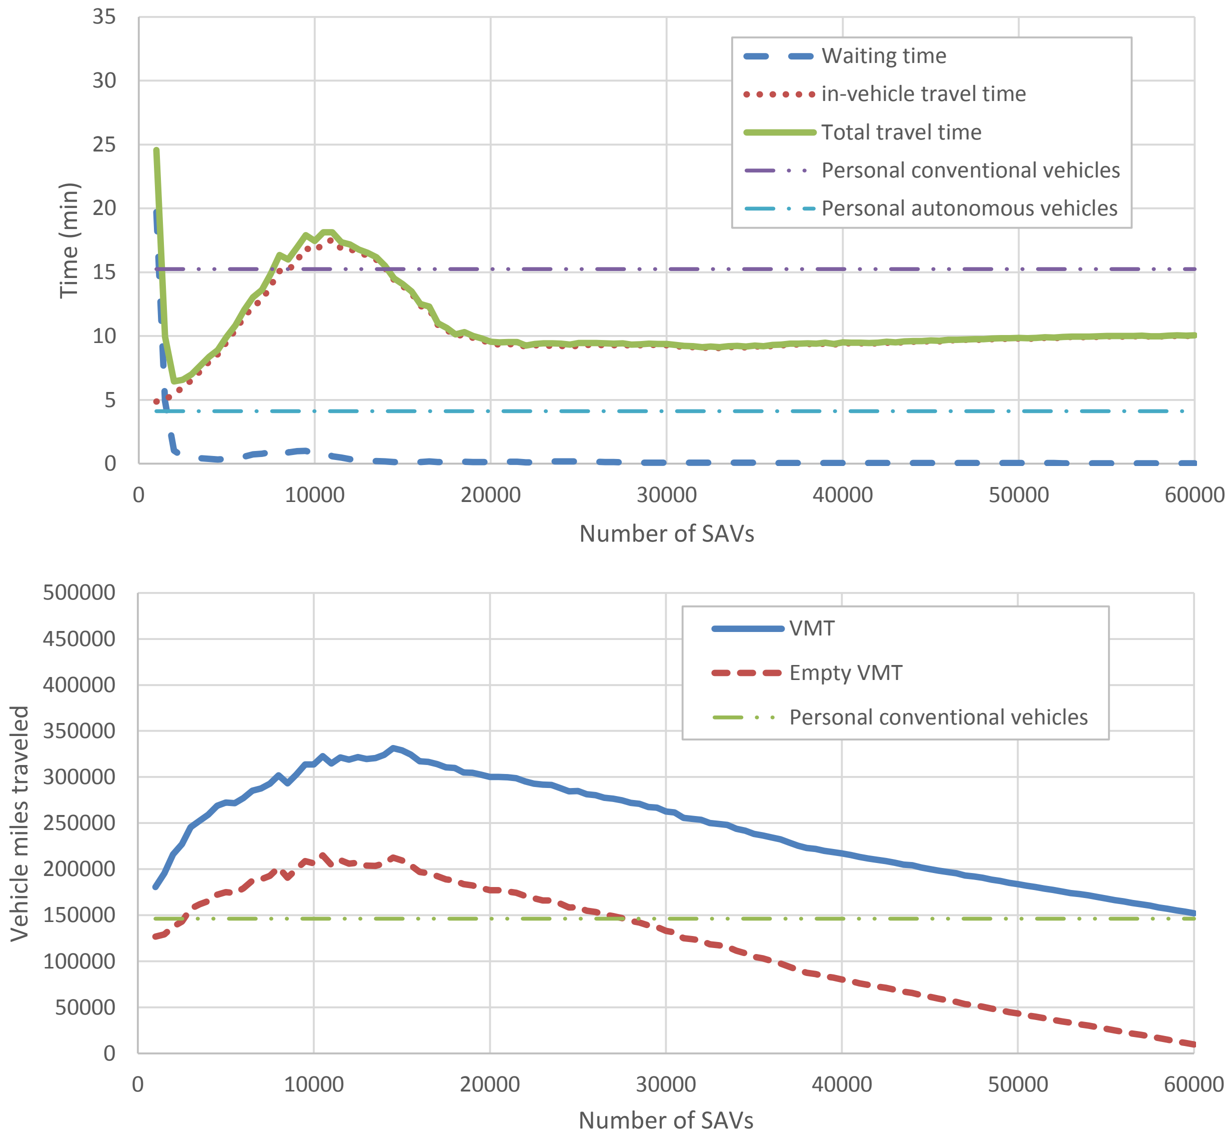

Figure 6: Travel time and VMT for the dynamic ride-sharing scenario

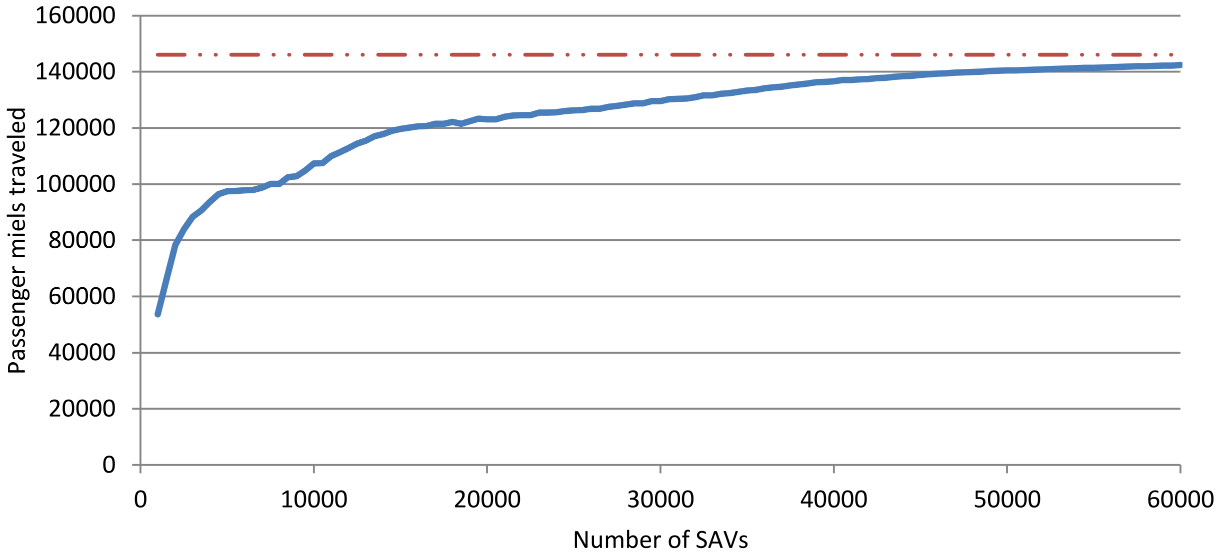

Figure 7: Passenger miles traveled for the dynamic ride-sharing scenario

peak demand. Most SAV scenarios resulted in greater congestion due to empty repositioning trips to reach travelers' origins.

Using SAVs without dynamic ride-sharing resulted in higher travel time than personal AVs. These levels of service appear to be lower than predicted by previous studies. Furthermore, a much larger SAV fleet size was needed for the AM peak. Although this paper used heuristics to solve the vehicle routing problem, finding an optimal solution in real-time in response to demand is impractical because the vehicle routing problem is NP-hard. Furthermore, previous studies also used similar heuristics. Therefore, these results demonstrate the importance of using realistic traffic flow models to study the additional congestion resulting from SAVs, and comparing SAVs with personal vehicles with a common traffic flow model. This paper also provides the framework to integrate SAV behavior into such models.

However, dynamic ride-sharing was highly effective at reducing congestion by combining traveler trips. Interestingly, ride-sharing had the best travel times when the number of SAVs was small (2000 SAVs providing service to 62,836 travelers), and these travel times were comparable or improved over personal vehicle scenarios. This shows that with effective routing heuristics and the right fleet size, SAVs could replace personal vehicles as paratransit or individual taxis.

Future studies should analyze how SAVs perform in a greater variety of scenarios, including varying demand and network topology. The experiments in this paper focused on a downtown grid network; a more suburban area with greater distance trips may be affected differently. This framework could also be used to study replacing traditional taxi service with SAVs. Taxis are typically constantly moving , which might increase congestion but decrease wait times. Additionally, better methods for vehicle routing and dynamic ride-sharing could improve SAV service, although any solution algorithms will have to be tractable for real-time execution in response to stochastic demand.

In addition, models using this framework could be used for travel demand and mode choice analyses. Travelers' trip choices typically depend on travel times, which could be greatly increased from congestion caused by SAVs. Many previous studies have assumed that all personal vehicle travel is replaced by SAVs [10–12]. In reality, SAVs add another mode option to personal vehicles and mass transit, and a fraction of travelers will choose each mode. The utility for each mode depends on travel times, for which congestion is a major factor. In particular, SAV congestion and routing affects both in-vehicle travel times as well as time spent waiting for pickup. With SAVs comprising a large fraction of vehicles on the road, SAVs will also affect the travel times of other modes as well. Of course, the number of travelers choosing the SAV mode will correspondingly affect the congestion caused by SAVs. To find mutually consistent travel demand, mode choice, and traffic congestion solutions, a SAV model with realistic congestion should be integrated into planning models to better predict the impacts of SAVs on city traffic patterns.

Acknowledgements

The authors gratefully acknowledge the support of the Data-Supported Transportation Operations & Planning Center, the National Science Foundation under Grant No. 1254921, and the Texas Department of Transportation.

References

[1] Burns, L. D., Jordan, W. C., and Scarborough, B. A. Transforming personal mobility. *The Earth Institute* (2013).

[2] Carlino, D., Boyles, S. D., and Stone, P. Auction-based autonomous intersection management. In *Intelligent Transportation Systems-(ITSC), 2013 16th International IEEE Conference on* (2013), IEEE, pp. 529–534.

[3] Carlino, D., Depinet, M., Khandelwal, P., and Stone, P. Approximately orchestrated routing and transportation analyzer: Large-scale traffic simulation for autonomous vehicles. In *Intelligent Transportation Systems (ITSC), 2012 15th International IEEE Conference on* (2012), IEEE, pp. 334–339.

[4] Chiu, Y.-C., Bottom, J., Mahut, M., Paz, A., Balakrishna, R., Waller, T., and Hicks, J. Dynamic traffic assignment: A primer. *Transportation Research E-Circular*, E-C153 (2011).

[5] Daganzo, C. F. The cell transmission model: A dynamic representation of highway traffic consistent with the hydrodynamic theory. *Transportation Research Part B: Methodological 28*, 4 (1994), 269–287.

[6] Daganzo, C. F. The cell transmission model, part ii: network traffic. *Transportation Research Part B: Methodological 29*, 2 (1995), 79–93.

[7] Dresner, K., and Stone, P. Multiagent traffic management: A reservation-based intersection control mechanism. In *Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems-Volume 2* (2004), IEEE Computer Society, pp. 530–537.

[8] Dresner, K., and Stone, P. Traffic intersections of the future. In *Proceedings of the National Conference on Artificial Intelligence* (2006), vol. 21, Menlo Park, CA; Cambridge, MA; London; AAAI Press; MIT Press; 1999, p. 1593.

[9] Fagnant, D. J., and Kockelman, K. Preparing a nation for autonomous vehicles: opportunities, barriers and policy recommendations. *Transportation Research Part A: Policy and Practice 77* (2015), 167–181.

[10] Fagnant, D. J., and Kockelman, K. M. The travel and environmental implications of shared autonomous vehicles, using agent-based model scenarios. *Transportation Research Part C: Emerging Technologies 40* (2014), 1–13.

[11] Fagnant, D. J., and Kockelman, K. M. Dynamic ride-sharing and optimal fleet sizing for a system of shared autonomous vehicles. In *Transportation Research Board 94th Annual Meeting* (2015), no. 15-1962.

[12] Fagnant, D. J., Kockelman, K. M., and Bansal, P. Operations of a shared autonomous vehicle fleet for the austin, texas market. *Transportation Research 28* (2015).

[13] Fajardo, D., Au, T.-C., Waller, S., Stone, P., and Yang, D. Automated intersection control: Performance of future innovation versus current traffic signal control. *Transportation Research Record: Journal of the Transportation Research Board*, 2259 (2011), 223–232.

[14] Godunov, S. K. A difference method for numerical calculation of discontinuous solutions of the equations of hydrodynamics. *Matematicheskii Sbornik 89*, 3 (1959), 271–306.

[15] Hausknecht, M., Au, T.-C., and Stone, P. Autonomous intersection management: Multi-intersection optimization. In *Intelligent Robots and Systems (IROS), 2011 IEEE/RSJ International Conference on* (2011), IEEE, pp. 4581–4586.

[16] Kesting, A., Treiber, M., and Helbing, D. Enhanced intelligent driver model to access the impact of driving strategies on traffic capacity. *arXiv preprint arXiv:0912.3613* (2009).

[17] Levin, M. W., and Boyles, S. D. Effects of autonomous vehicle ownership on trip, mode, and route choice. In *Transportation Research Board 94th Annual Meeting* (2015), no. 15-2147.

[18] Levin, M. W., and Boyles, S. D. Intersection auctions and reservation-based control in dynamic traffic assignment. In *Transportation Research Board 94th Annual Meeting* (2015), no. 15-2149.

[19] Levin, M. W., and Boyles, S. D. A multiclass cell transmission model for shared human and autonomous vehicle roads. *Transportation Research Part C: Emerging Technologies* (2015).

[20] Levin, M. W., Pool, M., Owens, T., Juri, N. R., and Waller, S. T. Improving the convergence of simulation-based dynamic traffic assignment methodologies. *Networks and Spatial Economics* (2014), 1–22.

[21] Li, P. Y., and Shrivastava, A. Traffic flow stability induced by constant time headway policy for adaptive cruise control vehicles. *Transportation Research Part C: Emerging Technologies 10*, 4 (2002), 275–301.

[22] Li, Z., Chitturi, M., Zheng, D., Bill, A., and Noyce, D. Modeling reservation-based autonomous intersection control in vissim. *Transportation Research Record: Journal of the Transportation Research Board*, 2381 (2013), 81–90.

[23] Lighthill, M. J., and Whitham, G. B. On kinematic waves. ii. a theory of traffic flow on long crowded roads. In *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences* (1955), vol. 229, The Royal Society, pp. 317–345.

[24] Richards, P. I. Shock waves on the highway. *Operations research 4*, 1 (1956), 42–51.

[25] Schakel, W. J., Van Arem, B., and Netten, B. D. Effects of cooperative adaptive cruise control on traffic flow stability. In *Intelligent Transportation Systems (ITSC), 2010 13th International IEEE Conference on* (2010), IEEE, pp. 759–764.

[26] Shladover, S., Su, D., and Lu, X.-Y. Impacts of cooperative adaptive cruise control on freeway traffic flow. *Transportation Research Record: Journal of the Transportation Research Board*, 2324 (2012), 63–70.

[27] Spieser, K., Treleaven, K., Zhang, R., Frazzoli, E., Morton, D., and Pavone, M. 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, 2014, pp. 229–245.

[28] Toth, P., and Vigo, D. *The vehicle routing problem.* Society for Industrial and Applied Mathematics, 2001.

[29] Ziliaskopoulos, A. K., and Waller, S. T. An internet-based geographic information system that integrates data, models and users for transportation applications. *Transportation Research Part C: Emerging Technologies 8*, 1 (2000), 427–444.