A general framework for modeling shared autonomous vehicles

Michael W. Levin (corresponding author)
Graduate Research Assistant
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

Abstract

Shared autonomous vehicles 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, and strategies such as preemptive relocation or dynamic ride-sharing could further improve service. 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 develops an event-based framework for implementing SAV behavior in existing traffic simulation models. We demonstrate this framework by implementing SAV behavior in a cell transmission model-based dynamic traffic assignment simulator. We also study heuristic approaches for preemptive relocation and dynamic ride-sharing. We compared replacing all vehicles with SAVs to the results of dynamic traffic assignment on the downtown

Austin city network. In most scenarios, the additional empty repositioning trips made by SAVs increased congestion and travel times. Preemptive relocation further increased the number of empty trips, making the situation worse in most cases. However, dynamic ride-sharing was competitive against personal vehicles when the number of SAVs was small, and preemptive relocation made slight improvements on dynamic ride-sharing in those scenarios. Overall, the results demonstrate the necessity of using realistic traffic flow models when studying SAVs and comparing SAVs to personal vehicle travel in the same traffic flow model.

1 Introduction

Autonomous vehicles could revolutionize transportation. Adaptive cruise control could increase road capacity [1,2] and reservation-based intersection control [3,4] could do the same for intersections [5,6]. 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 [7,8] 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 [7] or dynamic ride-sharing [9] 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 [7–10] and/or grid networks [7,9]. Although these were important studies for technology demonstration purposes, for accurate comparisons with personal vehicle scenarios a common traffic flow model is necessary. The purpose of this paper is to develop a framework compatible with existing traffic simulation models. This framework allows practitioners to integrate SAVs into their current traffic models to evaluate whether to fund public fleets of SAVs.

This framework admits a dynamic network loading model of SAVs based on a dynamic traffic assignment (DTA) simulator using the cell transmission model (CTM) [11,12]. We compare SAVs using heuristics for preemptive relocation and dynamic ride-sharing, based on previous work [7,9], against personal vehicle scenarios. Heuristics are used because exact algorithms for dynamic ride-sharing (a vehicle routing problem [13]) are NP-hard, and exact algorithms for preemptive relocation may be tractable but require further study. The framework allows us to study preemptive relocation and dynamic ride-sharing using a more realistic congestion model.

The contributions of this paper are as follows: we present an event-based framework for implementing SAVs in existing traffic models. This can be adapted for macro-, meso-, or micro-scopic flow models. We demonstrate this framework by studying congestion when SAVs are used to service all travelers, using CTM to propagate flow. We describe and study heuristics for preemptive relocation and dynamic ride-sharing on the downtown Austin city network and compare them with personal vehicle results from DTA. Overall, results show that a smaller SAV fleet can service all travel demand in the AM peak. However, all SAV scenarios 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 high demand scenarios.

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, with experimental results presented in Section 5. We present our conclusions in Section 6.

2 Literature review

The model developed in this paper builds on previous work on AV traffic flow models (Section 2.1 and SAV models (Section 2.2).

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 [1, 2] and traffic flow stability [14, 15]. However, Levin & Boyles [16] showed that increased road capacity may be offset by greater travel demand, particularly for empty repositioning trips. The reservation-based intersection protocol developed by Dresner & Stone [3,4] makes greater use of intersection capacity, allowing reductions in delays beyond optimized traffic signals in some scenarios [5, 6]. However, due to the computational complexity of the reservation protocol, many previous studies have been limited to small networks [17] or used simplified intersection models that reduced the traffic efficiency [18,19]. The conflict region model of the reservation protocol developed by Levin & Boyles [20] is tractable for large-scale DTA, and is therefore used in this paper.

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 [21]. Fagnant & Kockelman [7] 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 [9] 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. [8] 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. [10] 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, we would like to address some traffic modeling limitations of previous studies. All of them used custom simulation-based models, with many relying on grid-based networks. Many of the traffic congestion models were unrealistic; Fagnant & Kockelman [22] used MATSim [23], but many other studies did not specify or used fixed travel times. As we will demonstrate in Section 5, these models significantly underestimated the congestion resulting from SAVs. Accurate congestion modeling is necessary to evaluate whether replacing personal vehicles with SAVs improves traffic. Furthermore, the custom simulations would be difficult for practitioners to integrate into their existing traffic models. To address these limitations, we present an event-based framework that may be implemented on top of many simulation-based traffic models. We demonstrate this framework by implementing it in a DTA simulator and comparing SAV results with those from DTA.

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 we will demonstrate in our case study, 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.

In this section, we discuss 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 we do not require specific dispatcher behaviors. Section 4 discusses the dispatcher logic used in our case study, including dynamic ride-sharing and preemptive relocation.

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 1, 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.

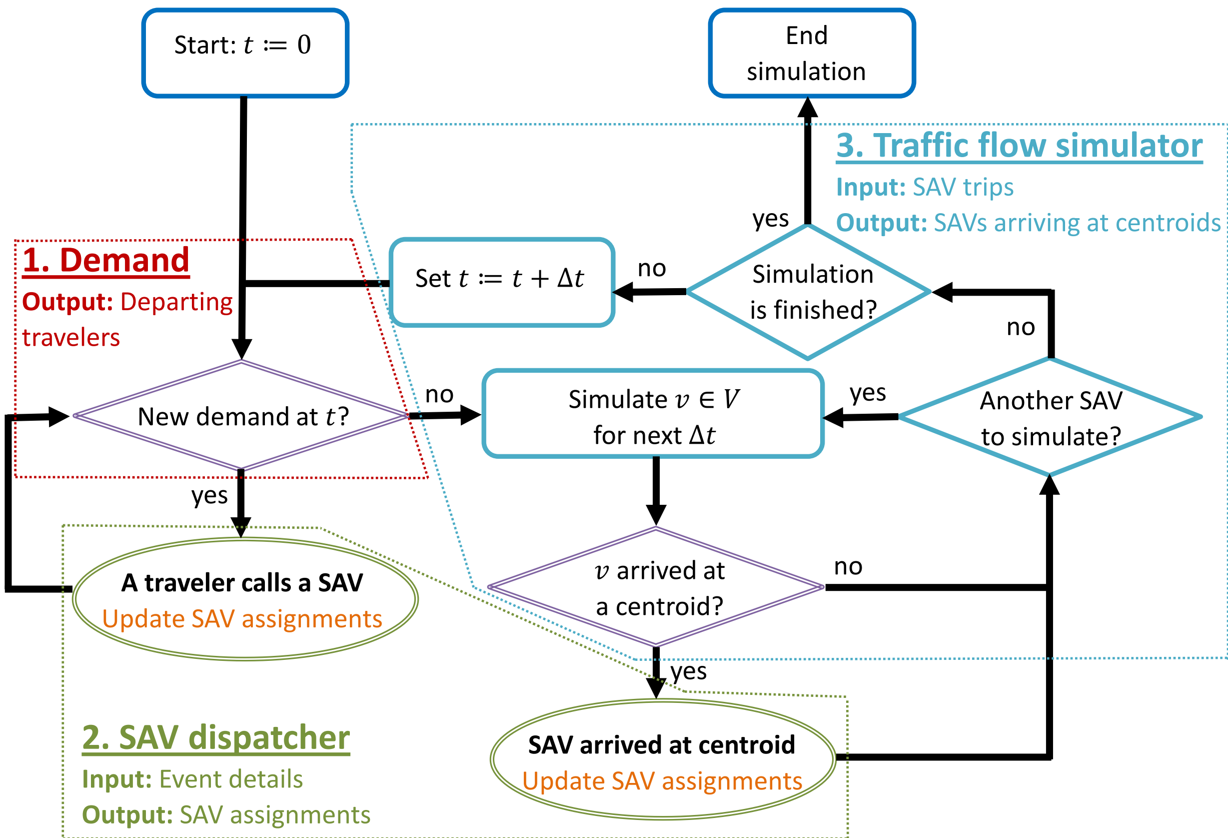

Figure 1: Event-based framework integrated into traffic simulator

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 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 in this paper we primarily refer 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

For this framework, we assume the existence of an 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 [7] 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. If i does not have available parking, v cannot wait at i and must travel elsewhere.

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 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 we will study heuristic rules in later sections.

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. Alternatively, SAV movement may be modeled through micro- or meso-simulation. Any uncertainty in the model is compatible with this framework; the SAV triggers the event only when 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 [20].

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 and preemptive relocation. Results using this implementation are presented in Section 5.

In this case study we assume 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

For this case study we converted personal vehicle trip tables from the morning peak into SAV traveler trip tables. For each vehicle trip, we created 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 origin 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) [11, 12], which is a space and time discretization of the hydrodynamic theory of traffic flow [24, 25]. CTM has been used in, and allows direct comparisons with, large-scale DTA simulators [26]. Because all vehicles are SAVs, we assume that intersections were controlled using the reservation-based protocol of Dresner & Stone [3] for AVs. For computational tractability, we use the conflict region model of reservation-based intersection control proposed by Levin & Boyles [20].

DTA models typically assume that route choice is based on driver experience. Each vehicle individually seeks its shortest route, resulting in a dynamic user equilibrium (DUE) in which no vehicle can improve travel cost by changing routes. Although this concept is based on the analytical static traffic assignment models, it requires further study to be formulated for SAV behavior because SAV trips may depend on stochastic demand. Therefore, we use a dynamic network loading-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.

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 $i \in Z$, we first check 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, we allow that SAV to service d. This is to reduce congestion resulting from sending more SAVs. (As we will demonstrate in Section 5, 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, we assume that travelers get SAVs 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 [7] 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.

4.5 Preemptive relocation

Preemptive relocation can reduce waiting times by starting to move SAVs to travelers' locations before they depart. Fagnant & Kockelman [7] studied four strategies for preemptive relocation and found that the best performing heuristic distributed SAVs to each centroid according to the proportion of productions. Since productions are typically determined by a survey of land use, the total expected trip productions at any centroid is likely to be known even if specific traveler departure times are not. Formally, let $\mathcal{P}_i$ be the productions and V_i the set of SAVs parked at $i \in Z$. The number of SAVs to be moved to i is

$$\Delta V_i = \frac{|V_i|}{|V|} - \frac{\mathcal{P}_i}{\sum\limits_{i' \in Z} \mathcal{P}_{i'}}$$

If $\Delta V_i > 0$, ΔV_i SAVs are moved from i; if $\Delta V_i < 0$, $-\Delta V_i$ SAVs are moved to i. Let $\mathcal{Z}^+ = \{i \in Z | \Delta V_i > 0\}$ and $\mathcal{Z}^- = \{i \in Z | \Delta V_i < 0\}$. $\mathcal{Z}^+$ is sorted in decreasing order. For each $i \in \mathcal{Z}^+$, ΔV_i SAVs from i are distributed to the nearest centroids (by travel time) in $\mathcal{Z}^-$. This attempts to minimize the congestion caused by relocation.

5 Case study: experimental results

We performed several sets of experiments to study how SAVs perform relative to personal vehicles, and how the preemptive relocation and dynamic ride-sharing heuristics affect performance. Our experiments were primarily performed on the downtown Austin network, shown in Figure 2. It consists of a downtown grid with freeway and arterial corridors. It has 171 zones, 546 intersections, 1247 links, and 62836 trips over 2 hours in the AM peak. The centroids are significantly disaggregated for this downtown region. 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 (with the SAV framework and cell transmission model 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 data was provided by the Capital Area Metropolitan Planning Organization.

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 about relocating SAVs *during* the AM peak — while travelers are requesting SAVs.) Fagnant & Kockelman [7] 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 4000 and 40,000 SAVs, with increments of 500. For some scenarios, the range was reduced to numbers of SAVs that could provide

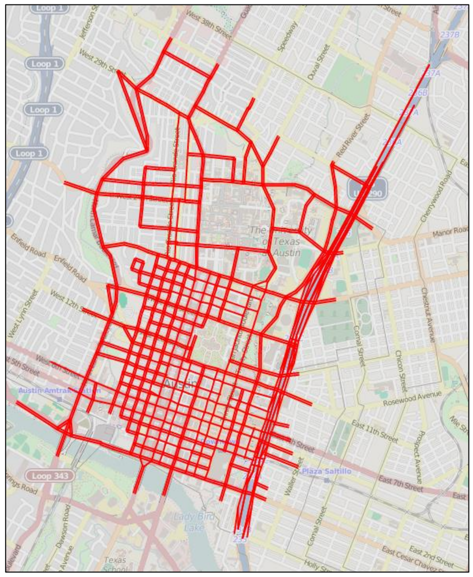

Figure 2: Downtown Austin network

service to all travelers within 6 hours because service was limited by having too few SAVs or too much congestion.

5.1 Personal vehicles

First, to create a base scenario, we solved DTA on downtown Austin, assuming that all travelers use privately owned AVs for their trips. Although SAVs use a dynamic network loading-style route choice, the DTA model assumed drivers based their routes on past experience to find a dynamic user equilibrium. Therefore, the routing strategy in DTA is likely more efficient than the routing strategy for SAVs. 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 7.24 minutes. Since the adoption of reservation controls may be difficult or inefficient if a significant proportion of personal vehicles are not autonomous, both DTA 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 preemptive relocation or dynamic ride-sharing. Figure 3 shows travel time results with 28,500 to 40,000 total SAVs available. (Fewer numbers of SAVs were found to be insufficient to service all travelers after 6 hours.) As the number of SAVs increased, waiting time decreased linearly. 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 3). This indicates that the difference was primarily due to less repositioning trips to pick up the next traveler. It is intuitive that as the number of SAVs increased, the average distance between a waiting traveler and the closest available SAV would decrease.

Overall travel times in this base SAV scenario were much higher than with personal vehicles. In-vehicle travel time, interestingly, decreased for around 31,000 to 32,000 SAVs, then remained nearly constant thereafter. This may be due to a reduction in congestion when SAVs were traveling less for repositioning trips. 33–35 minute in-vehicle travel times, though, are double that of DTA with signals, and five times that of DTA with AVs. Previous studies predicted that each SAV can service multiple travelers with acceptable waiting times — that is still true in these experiments, but the experienced travel times are more similar to those of public transit. Travelers may be unwilling to use SAVs if the travel times are this high.

The difference in travel time is most likely due to additional congestion from empty repositioning trips made to pick up the next traveler. The downtown Austin network is already fairly congested during the AM peak, and the addition of repositioning trips makes matters worse. This is an important result, though, because it demonstrates the value in using a realistic traffic flow model for analyzing congestion. For less congested networks, SAVs might cause only modest increases in congestion. However, for a high traffic city in the AM peak, these results are not encouraging for a switch to SAVs.

5.3 Preemptive relocation

Next, we studied the effects of preemptively relocating SAVs to match the proportion of productions of each centroid, described by Section 4.5. This resulted in very high waiting times with few AVs available. This is likely due to the fairness of assigning SAVs: travelers are prioritized by the time spent waiting. Unless a traveler was waiting at the destination of the relocating SAV, it would be re-assigned to service a different traveler, which is likely why the waiting time was so high when few SAVs were available. Although this is a reasonable policy, alternatives such as that of Fagnant and Kockelman [7], in which travelers are prioritized according to distance from the available SAV, could improve waiting time.

As the number of SAVs increased, waiting time decreased linearly, although it was still much higher than the base scenario. One potential reason is the additional congestion resulting from relocating SAVs. This is illustrated by the much higher empty VMT resulting from relocations, shown in Figure 4. Relocating resulted in around 400,000 vehicle miles of empty travel. This did not decrease as the number of SAVs increased, as it did in the base scenario (Figure 3), which likely contributed to the increasing in-vehicle travel times. The in-vehicle travel time increased linearly with the number of SAVs, which is indicative of those additional SAVs contributing significantly to congestion. In fact, beyond 20,500 SAVs, congestion prevented effective service for all travelers. Although waiting time decreased, the increases in travel time resulted in only small decreases in total travel time.

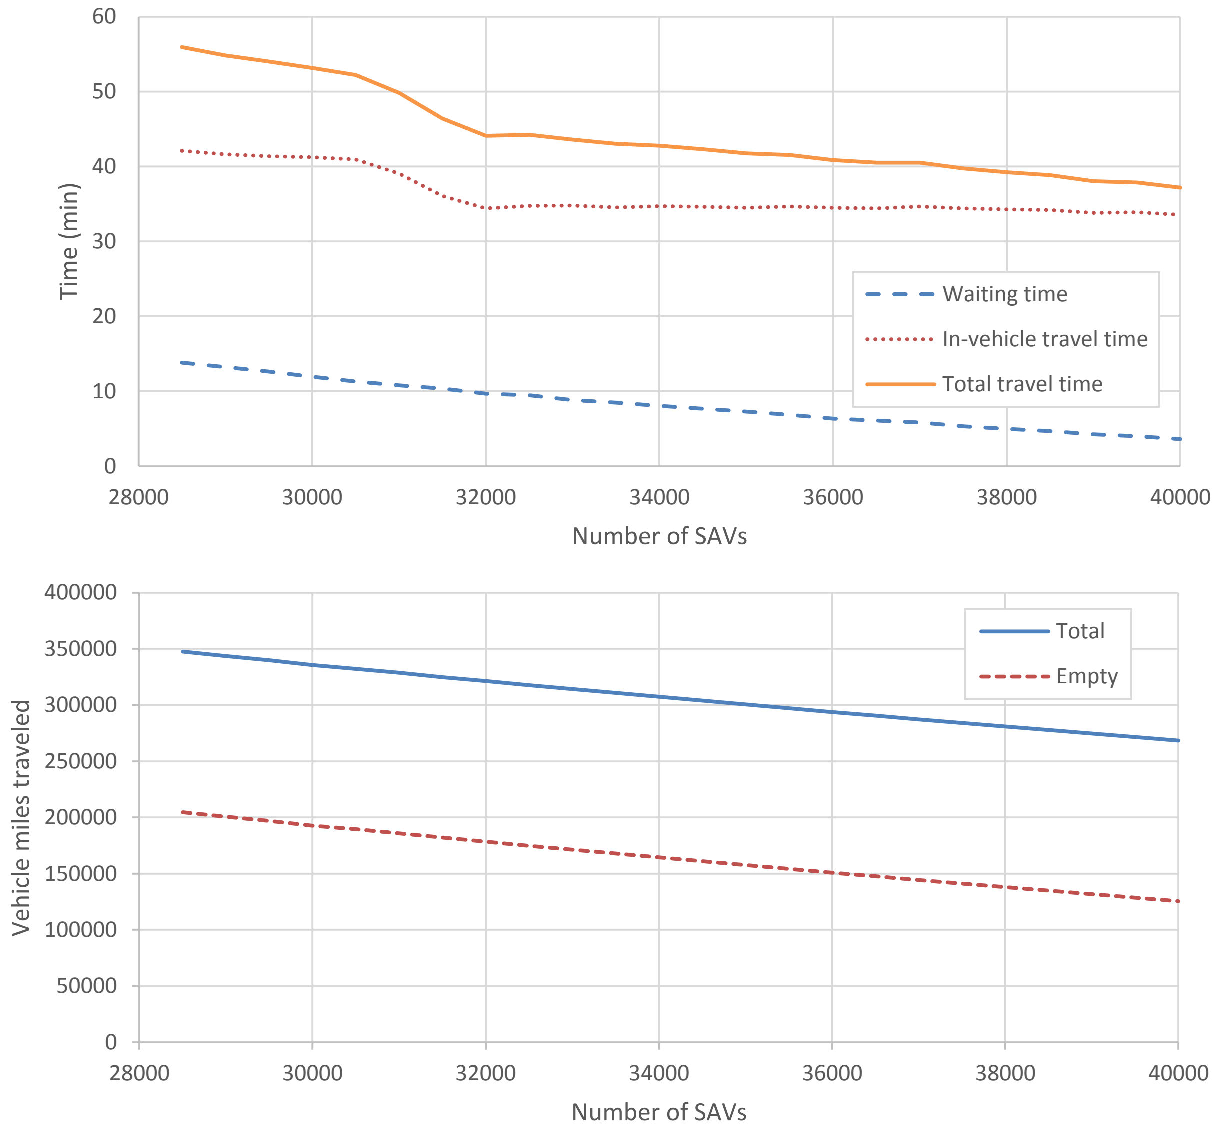

Figure 3: Travel time and VMT for the base SAV scenario

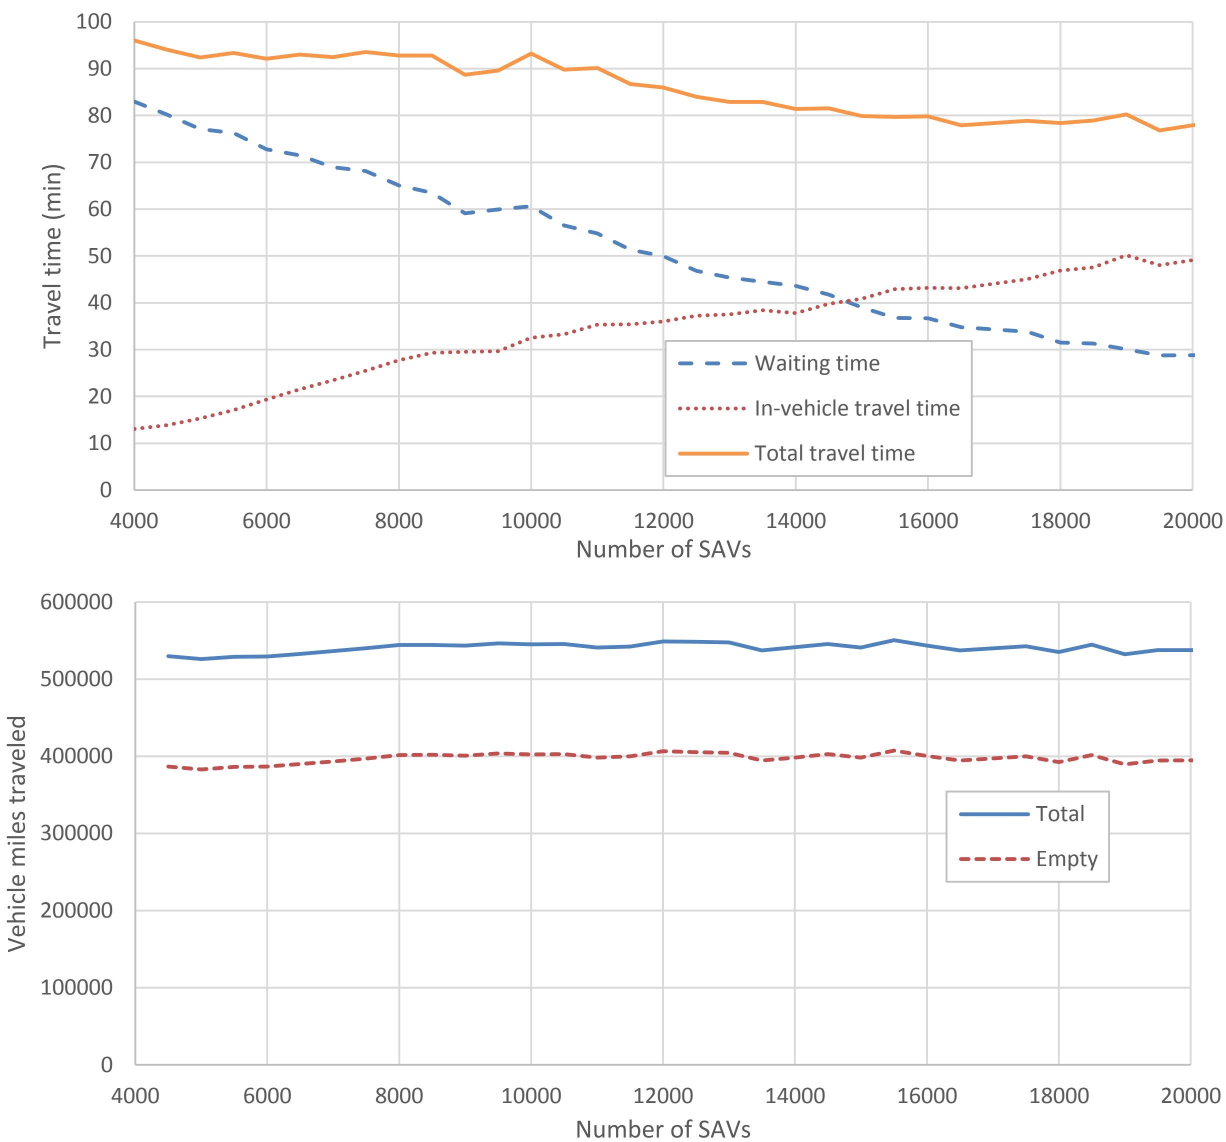

Figure 4: Travel time and VMT for the preemptive relocation scenario

5.4 Dynamic ride-sharing

Compared with the base and pre-emptive relocation SAV scenarios, dynamic ride-sharing allowed SAVs to provide in-vehicle travel times competitive with personal vehicles. SAV capacity was 4 passengers, and ϵ was set at 0.4 [9]. At the minimum scenario of 4000 SAVs, average in-vehicle travel time was 12.4 minutes and the average waiting time was only 5.1 minutes. as shown in Figure 5. For travelers who call a SAV a few minutes before they plan to leave, a 5.1 minute waiting time is easily forgivable. 12.4 minutes in-vehicle travel times improve over average travel times with personal vehicles and traffic signals, and are only around 5 minutes greater than personal vehicles with reservation controls. As the number of SAVs increased, though, travel times increased until they were comparable with the non-ride-sharing scenario. Waiting times were overall much lower. This was probably because travelers with nearby destinations could share the same SAV, when one arrived. This performed best with the fewest SAVs available: although this increased waiting times, SAV utilization was greater.

Figure 5 shows that VMT peaked with around 23,000 SAVs. With only 4000 SAVs, VMT was low because of the low number of SAVs, but dynamic ride-sharing allowed just 4000 SAVs to service 62836 travelers in the AM peak. Note that the difference between total and empty VMT increases as the number of SAVs increases due to the reduction in average number of passengers carried per SAV. This demonstrates an interesting result: when ride-sharing is possible, having fewer SAVs is sometimes more efficient. It reduces congestion and maximizes the utilization of each SAV because travelers accumulate as they wait for one of the few SAVs to arrive for pick-up.

4000 SAVs corresponds to a 93.6% reduction in the number of vehicles: each SAV services an average of 15.7 travelers. This efficiency is similar that of previous studies, such as one SAV servicing 11 travelers [7]. However, the observed efficiency is at least partially due to the network topology: due to considering only the downtown region, traveler origin/destinations are fairly close together. If a regional network were used, the efficiency would likely decrease.

Preemptive relocation was somewhat detrimental when used with dynamic ride-sharing, as shown in Figure 6. When the number of SAVs was below 10,000, preemptive relocation slightly reduced waiting times. At higher numbers of SAVs, though, relocation still had a waiting time of around 3–4 minutes. This probably resulted from high congestion delaying relocating vehicles from arriving. Beyond 20,000 SAVs, the congestion caused by the additional relocations prevented travelers from reaching their destination. Travel time increased significantly with the number of SAVs, mostly due to increases in in-vehicle travel time from congestion. However, travel time with ride-sharing and relocation increased at a lower rate than travel time with just ride-sharing. In fact, when the number of SAVs was between 4000 or 10,000, preemptive relocation with ride-sharing had slightly lower travel times than ride-sharing alone. However, at higher numbers of SAVs, with ride-sharing available most SAVs were relocating, resulting in high congestion and worse travel times than in the base case. As the number of SAVs increased, the empty VMT increased as well, resulting in around 100,000 additional miles traveled at 20,000 SAVs when relocation and ride-sharing was used compared to ride-sharing alone (Figure 6).

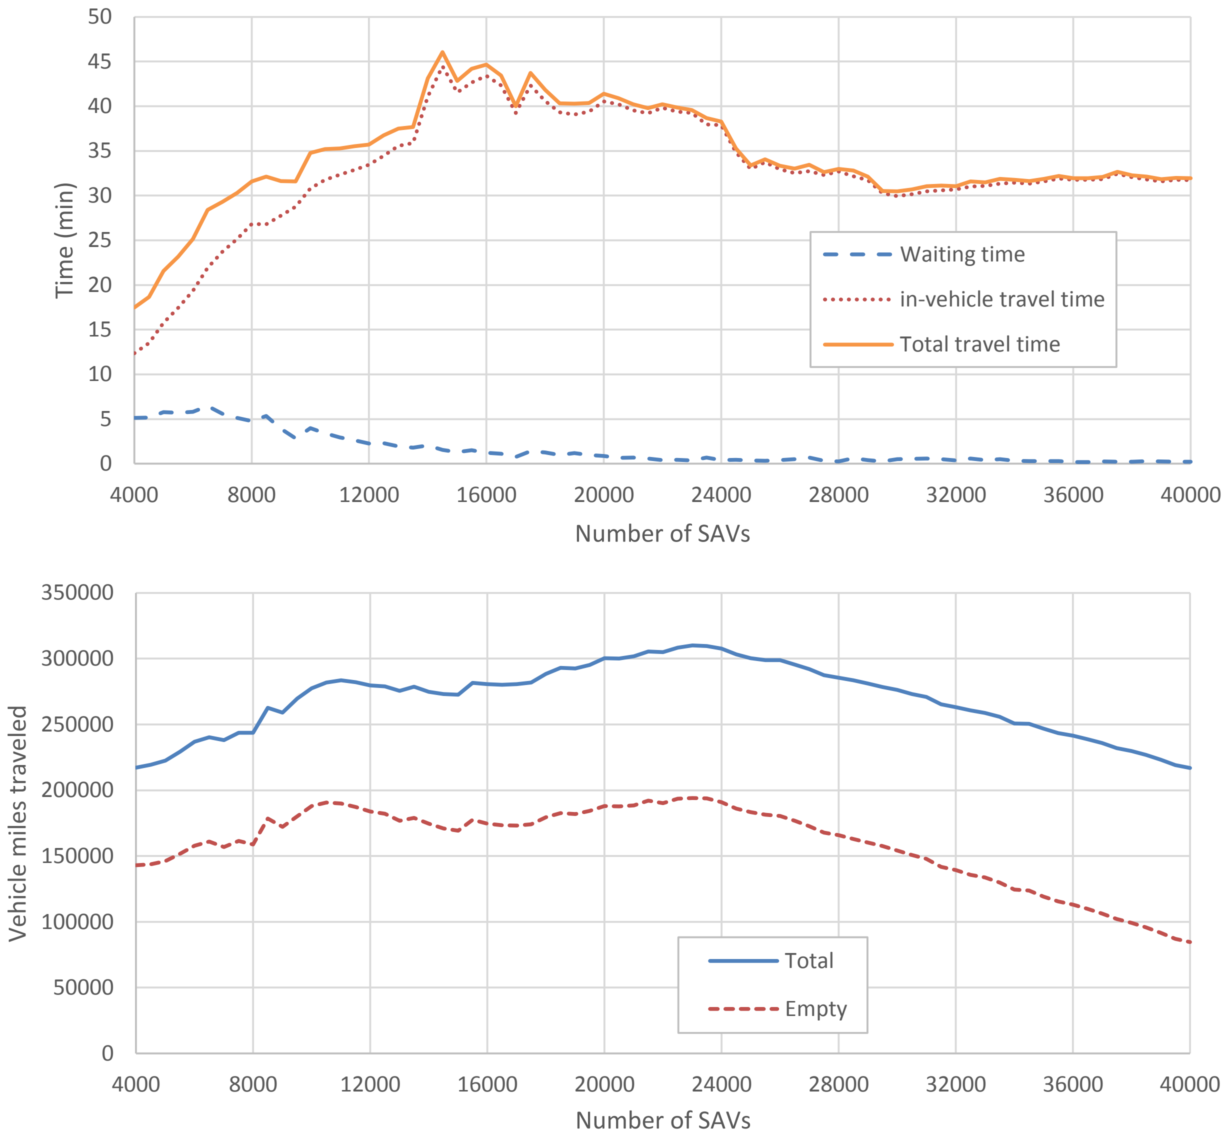

Figure 5: Travel time and VMT for the dynamic ride-sharing scenario

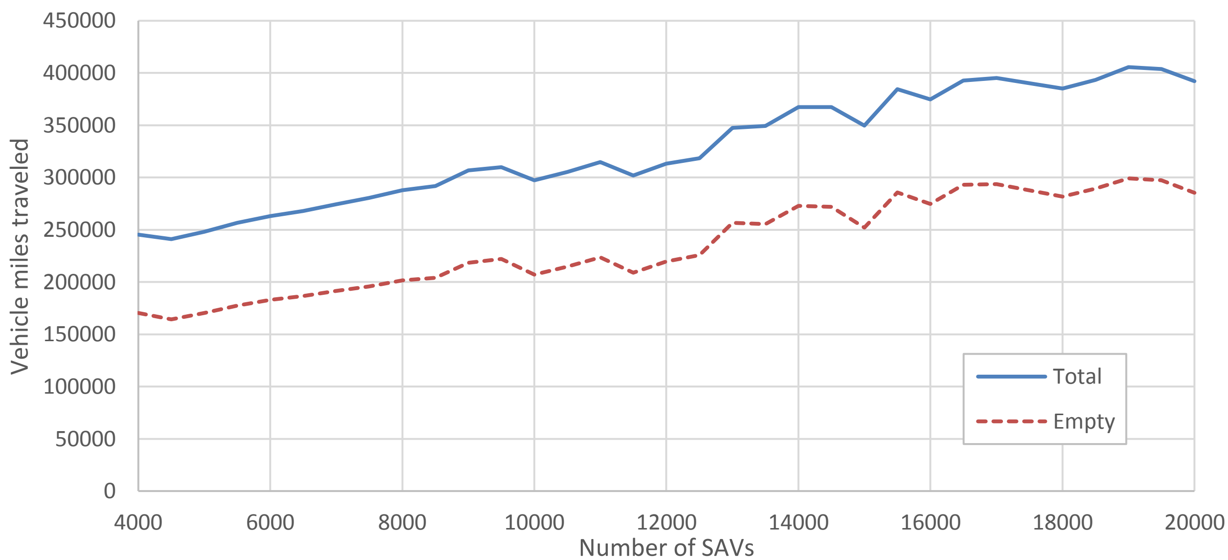

Figure 6: Travel time and VMT for the dynamic ride-sharing and preemptive relocation scenario

5.5 Discussion

The high travel times across most SAV scenarios indicates that greater optimization is necessary before 100% SAVs are competitive with personal vehicles. Agencies would not be willing to invest in a SAV fleet unless they improve traffic. Considering the positive results of previous studies, the results in this paper demonstrate the necessity of using realistic congestion models when studying SAVs.

Congestion could also be reduced by improving the dispatch behavior of SAVs used in this paper. Dispatching SAVs to service travelers, with preemptive relocation and dynamic ride-sharing, is in general a vehicle routing problem. Vehicle routing problems are NP-hard [13], and for SAVs the problem must be solved in real-time in response to demand. Therefore, any practical implementation will necessarily use some type of heuristic for tractability. However, greater study on heuristics could improve efficiency.

On the other hand, these studies assumed that SAVs could park at centroids for arbitrary amounts of time. However, traditional taxis are usually traveling while empty, especially in downtown regions where parking may be limited. This is somewhat modeled by the preemptive relocation scenarios, but taxis may also circle high production centroids while waiting for demand. Such additional travel could increase congestion beyond what is predicted in this paper.

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 heuristic approaches to preemptive relocation and dynamic ride-sharing. Then, we studied replacing personal vehicles with SAVs in the downtown Austin network with AM peak demand. Most SAV scenarios resulted in greater congestion due to empty repositioning trips to reach travelers' origins. The best preemptive relocation strategy of Fagnant & Kockelman [7] further increased congestion due to the additional repositioning trips, and in most scenarios made congestion even worse. Dynamic ride-sharing was effective at reducing congestion by combining traveler trips. Interestingly, ride-sharing had the best travel times when the number of SAVs was small (4000 SAVs providing service to 62,836 travelers), and these travel times improved over personal vehicles and traffic signals, and were competitive with personal vehicles and reservation controls. More SAVs decreased waiting times, but also decreased the number of passengers per SAV and correspondingly increased congestion. When used with dynamic ride-sharing, preemptive relocation slightly reduced both in-vehicle travel time and waiting time for small numbers of SAVs.

Most SAV scenarios resulted in much higher travel time than personal vehicles. These levels of service appear to be significantly lower than predicted by previous studies. 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. Considering that only a small subset of the tested SAV scenarios compared well against personal vehicles, future work should further explore which conditions allow SAVs to be used effectively. Perhaps SAVs are most useful in moderate, but not high, demand, such as during off-peak hours for work-based trips or home-based non-work trips.

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 — although preemptive relocation modeled that to some degree, constantly moving SAVs might increase congestion but decrease wait times. Additionally, better methods for preemptive relocation and dynamic ride-sharing could improve SAV service, although any solution algorithms will have to be tractable for real-time execution in response to dynamic demand.

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

[2] S. Shladover, D. Su, and X.-Y. Lu, "Impacts of cooperative adaptive cruise control on freeway traffic flow," *Transportation Research Record: Journal of the Transportation Research Board*, no. 2324, pp. 63–70, 2012.

[3] K. Dresner and P. Stone, "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*, pp. 530–537, IEEE Computer Society, 2004.

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

[5] D. Fajardo, T.-C. Au, S. Waller, P. Stone, and D. Yang, "Automated intersection control: Performance of future innovation versus current traffic signal control," *Transportation Research Record: Journal of the Transportation Research Board*, no. 2259, pp. 223–232, 2011.

[6] Z. Li, M. Chitturi, D. Zheng, A. Bill, and D. Noyce, "Modeling reservation-based autonomous intersection control in vissim," *Transportation Research Record: Journal of the Transportation Research Board*, no. 2381, pp. 81–90, 2013.

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

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

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

[10] K. Spieser, K. Treleaven, R. Zhang, E. Frazzoli, D. Morton, and M. Pavone, "Toward a systematic approach to the design and evaluation of automated mobility-on-demand systems: A case study in singapore," in *Road Vehicle Automation*, pp. 229–245, Springer, 2014.

[11] C. F. Daganzo, "The cell transmission model: A dynamic representation of highway traffic consistent with the hydrodynamic theory," *Transportation Research Part B: Methodological*, vol. 28, no. 4, pp. 269–287, 1994.

[12] C. F. Daganzo, "The cell transmission model, part ii: network traffic," *Transportation Research Part B: Methodological*, vol. 29, no. 2, pp. 79–93, 1995.

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

[14] P. Y. Li and A. Shrivastava, "Traffic flow stability induced by constant time headway policy for adaptive cruise control vehicles," *Transportation Research Part C: Emerging Technologies*, vol. 10, no. 4, pp. 275–301, 2002.

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

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

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

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

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

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

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

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

[23] M. Balmer, K. Meister, M. Rieser, K. Nagel, K. W. Axhausen, K. W. Axhausen, and K. W. Axhausen, *Agent-based simulation of travel demand: Structure and computational performance of MATSim-T*. ETH, Eidgenössische Technische Hochschule Zürich, IVT Institut für Verkehrsplanung und Transportsysteme, 2008.

[24] M. J. Lighthill and G. B. Whitham, "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*, vol. 229, pp. 317–345, The Royal Society, 1955.

[25] P. I. Richards, "Shock waves on the highway," *Operations research*, vol. 4, no. 1, pp. 42–51, 1956.

[26] A. K. Ziliaskopoulos and S. T. Waller, "An internet-based geographic information system that integrates data, models and users for transportation applications," *Transportation Research Part C: Emerging Technologies*, vol. 8, no. 1, pp. 427–444, 2000.