# Day-to-day market evaluation of modular autonomous vehicle fleet operations with en-route transfers

Nicholas S. Caros and Joseph Y.J. Chow*

C2SMART University Transportation Center, Department of Civil & Urban Engineering, New York University, Brooklyn, NY 11201

January 2020

**Abstract**

This study extends the two-sided day-to-day learning framework to simulate the performance of a mobility service using modular autonomous vehicles (MAVs) capable of en-route passenger transfers. An insertion heuristic is used to assign trips to a fleet of vehicles and to determine whether engaging in an en-route transfer is advantageous. The operator acts as an endogenous decision maker, updating the relative weight of the operator cost and user cost within the routing algorithm after each simulation day to optimize profit. Real transit ridership data from the United Arab Emirates are used for an empirical study of three operating strategies: door-to-door service within an urban core, commuter first/last mile service and a hub-and-spoke service. Results are compared with and without en-route transfers to quantify the advantage of the en-route transfer capability for each strategy.
*Keywords*: Day-to-Day Adjustment, Last Mile Problem, Flexible Transit, Ridesharing

## 1 Introduction

Recent advances in sensor, automation and vehicle technology have enabled the development of modular autonomous vehicles (MAVs) such as the NEXT Future Transportation NX1 prototype, the CRRC ART and the Ohmio LIFT. These vehicles are designed to physically connect to one another while in motion to form platoons of two or more vehicles, permitting passengers and freight to move freely between them ("en-route transfers") as illustrated in Figure 1. Prototypes of modular vehicles with en-route transfer capability are currently undergoing testing but have not yet been approved for use on public roads (Fahmy, 2018). While the freight applications of MAVs are evident, the passenger transport benefits of MAVs over traditional connected autonomous vehicles remain unclear. The potential impacts of this technology, which is currently unavailable, have not been widely evaluated in existing literature.

A mobility service with a fleet of MAVs could address what is referred to as the "last mile" transit problem (Guo et al., 2017). This refers to the challenge of providing convenient trips to and from fixed-route transit hubs such as bus depots and train stations. The difficulty of access and egress from these hubs has been found to decrease the use of public transportation (Krygsman et al.,

*Corresponding Author
*Email address:* joseph.chow@nyu.edu

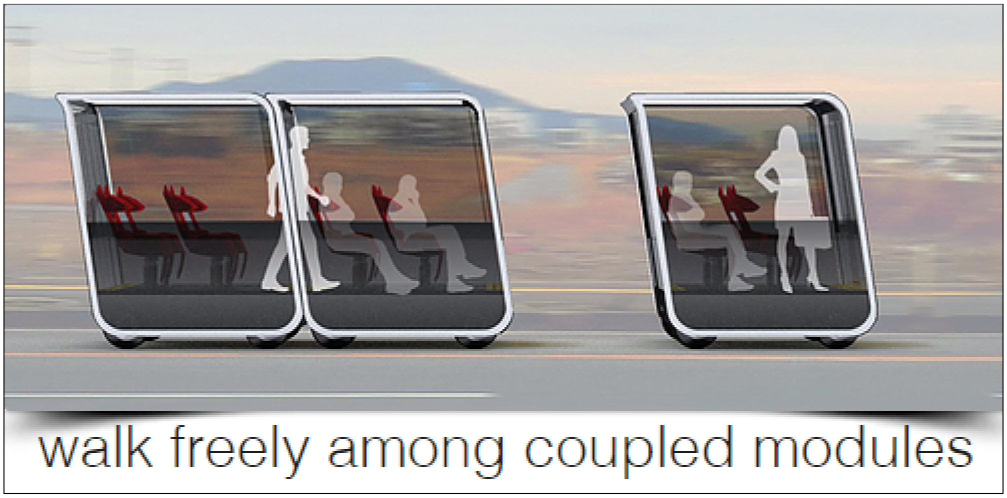


Figure 1: Illustration of modular autonomous vehicles
*Source:* Next Future Transportation (2020)

2004). A feasible solution to the last mile problem could increase transit mode share and reduce reliance on personal vehicle trips, thereby lowering congestion on the road network and improving local air quality (Lesh, 2013).

Attempts to solve the last mile problem with flexible transit service have experienced mixed results. Flexible transit systems designed to serve last mile trips such as Kutsuplus in Helsinki, FLEX in Santa Clara County and Bridj in several American cities did not achieve financial sustainability (Sulopuisto, 2016; Marshall, 2017; Sisson, 2018). Some limited pilot programs including GO Connect in Milton, Ontario have been more successful and many other public pilots are planned or ongoing (Hennessey, 2016; Jiao et al., 2017). One of the primary obstacles that last mile solutions must overcome is passengers' perceived disutility of transfer or interchange between two different travel modes or transit lines. Studies have observed a range of values for the perceived cost of a transfer or interchange between bus and rail. Wardman et al. (2001) revealed that a bus interchange has disutility equivalent to 4.5 minutes of in-vehicle time (IVT), and a rail interchange equivalent to 8 minutes of IVT, in addition to the time needed to complete the interchange. Similarly, Balcombe et al. (2004) found interchange penalties ranging from 5 to 10 minutes of IVT. A more recent study found that travelers perceive modal transfers to be equivalent to 15.2 minutes of IVT for the first transfer in a trip and 17.7 minutes for the second transfer in a trip (Garcia-Martinez et al., 2018).

MAVs could potentially improve outcomes for both the passengers and operators of last mile transit services (Guo et al., 2017; Chen et al., 2019a,b). The capacity to exchange passengers between vehicles provides mobility service operators with an additional dimension of flexibility when assigning users to vehicle routes. In the simple case shown in Figure 2, an en-route transfer reduces average travel times for the passengers while simultaneously reducing the total vehicle distance travelled. Vehicle distance travelled is reduced from 48 km in Scenario 1 (8km + 6km + 10km return trip for each vehicle) to 36 km (5km + 5km + 8km return trip for each vehicle) in Scenario 2, assuming that vehicles must return to their origin location after service is complete. Furthermore, the average passenger travel time (assuming a travel speed of 60 km/hr) is reduced from 11 minutes to 10 minutes.

This study seeks to understand the impact of en-route transfers on users and operators under

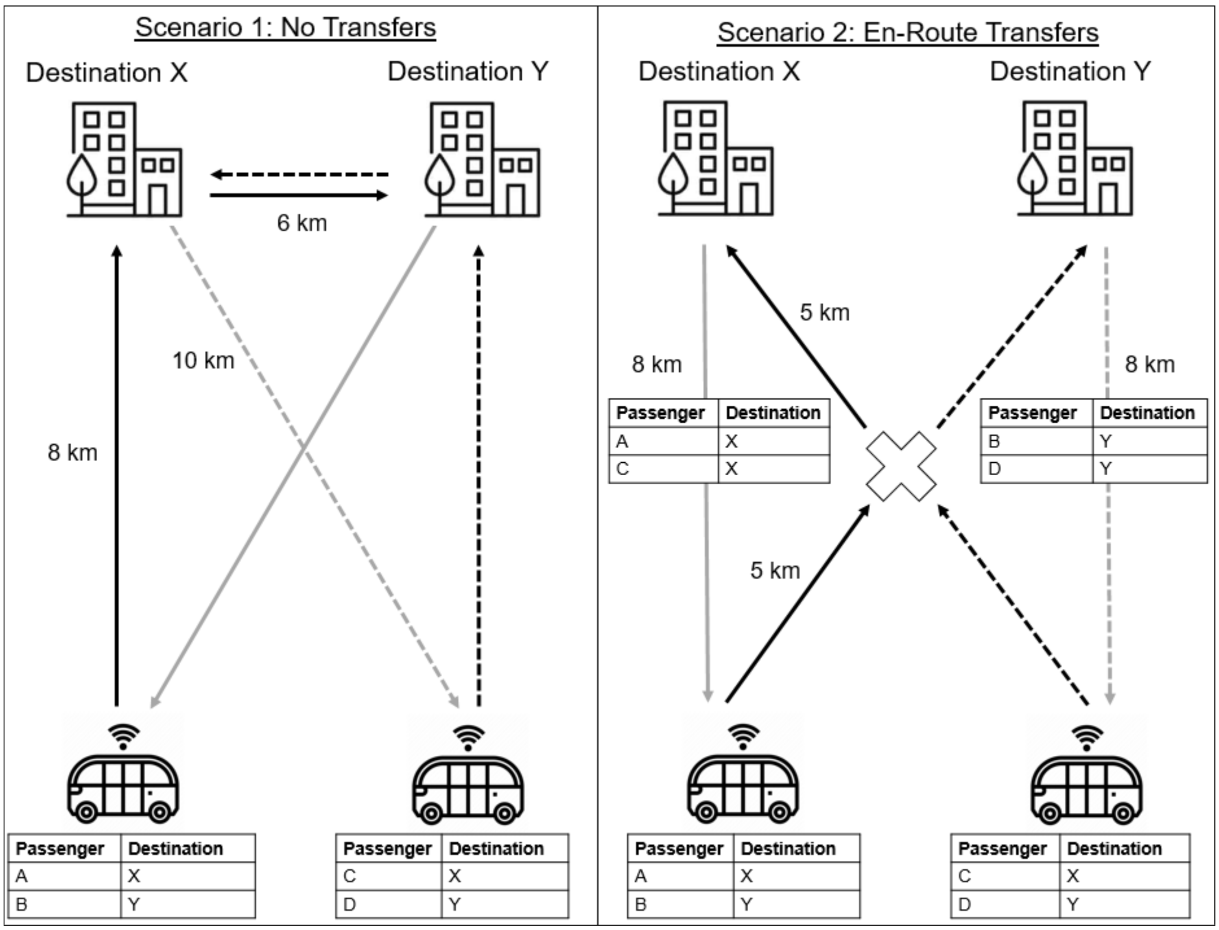


Figure 2: Comparison of vehicle routes with and without en-route transfers

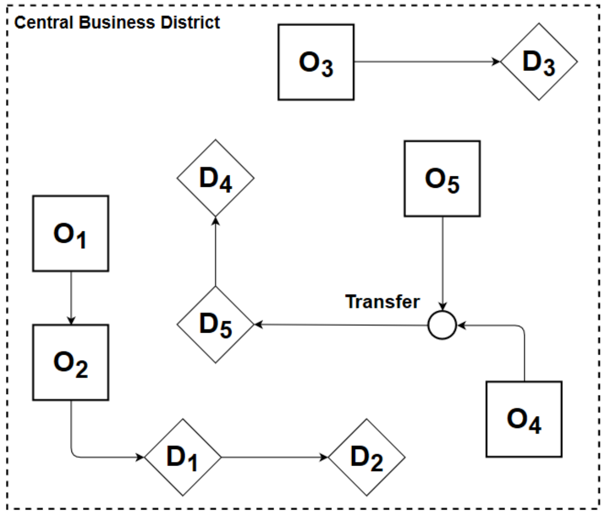


Figure 3(a): Urban door-to-door MAV service with en-route transfers

different last mile service designs. Three types of service designs are considered.

The first service design involves serving relatively short-distance trips where trip origins and destinations are distributed within a central business district, as shown in Figure 3(a). This service design is similar to that of ride hailing and taxi services operating in urban centers. En-route transfers are used in cases where they reduce travel time or operating cost, but they are used less frequently than under the next two service designs. This is because the demand is distributed across a region with no alignment of origins or destinations.

A second service design is to serve passengers for typical last mile trips to and from a fixed route transit station, as shown in Figure 3(b). This type of service would make fixed route transit more convenient for users by eliminating the need to walk, drive or cycle to a station from their trip origin. Using this service design, two MAVs that are collecting passengers from their trip origins could take advantage of the en-route transfer capability by combining passengers into one vehicle, thus freeing the second vehicle to make other pickups. Transfers are more likely to be advantageous in this service design because the origin or destination of each trip will be a transit station, making it more likely that vehicles will be traveling near one another.

The third service design involves stationary, hub-based transfers as a logical initial step in the launch of a full passenger service. Stationary transfers allow the safety and user comfort of the transfer mechanism to be tested before full-speed coupling and decoupling is undertaken. In this service design, MAVs serve long-distance commuters in a hub-and-spoke type of service, replacing both the fixed-route transit and last mile components of a trip. This service design is shown in Figure 3(c). Passengers are first delivered from their origin to a hub location. A two-vehicle convoy then travels to the regional employment center, where the convoy separates to deliver passengers to their destinations. This is similar to a park-and-ride system or multi-modal commuter service; however, it is considerably more convenient for travelers. Only one continuous mode is required to travel from the home to the workplace, and a “transfer” would only require moving seats within a convoy of vehicles.

We evaluate a MAV service empirically under each of the three service designs using demand profiles for commuters traveling between Dubai and Sharjah in UAE. As the system operates under a stochastic dynamic operating policy, a day-to-day adjustment process (Djavadian and Chow,

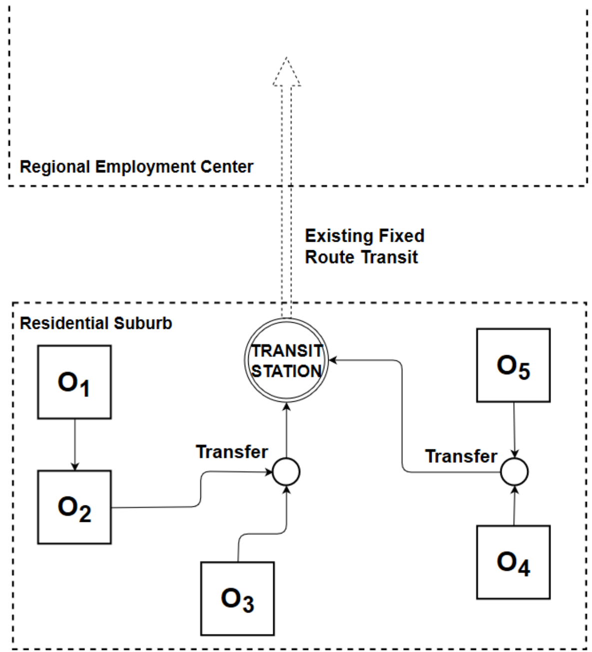


Figure 3(b): Last mile MAV commuter service with en-route transfers

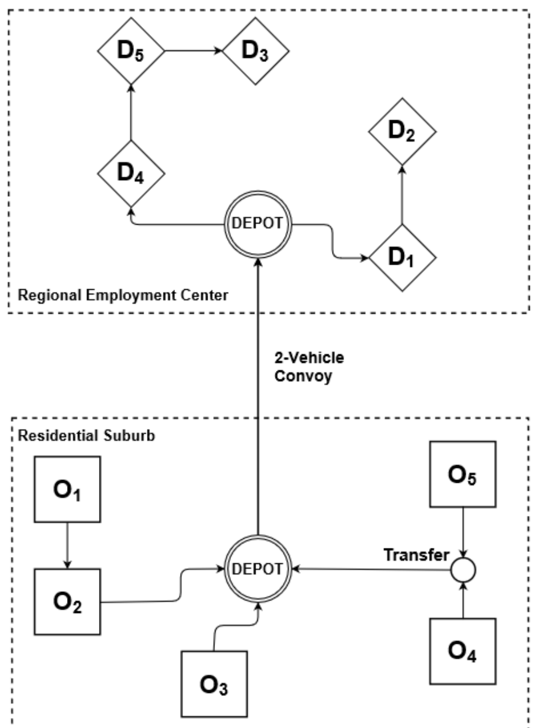


Figure 3(c): Hub-and-spoke MAV service with en-route transfers

2017b) is needed to reach market equilibrium. The demand for the MAV service varies each day based on users' previous experience with the service. Evaluations of transportation systems with fixed demand do not account for competition between modes and services. By creating a simulation with variable demand we are able to evaluate the static effects of a MAV service on individual trips and how those effects impact the equilibrium of a market for transportation services with competing travel modes.

Variable demand also allows for analysis of the long-term effects of the system operator's policies. The service operator interacts with travelers through the assignment of travelers to vehicles. The operator can update their assignment algorithm each day, affecting passenger outcomes and operator profit. The operator's profit function is controlled with a parameter representing the user/operator cost ratio (e.g. see Sayarshad and Chow, 2015). This ratio is the operator's perceived elasticity of user demand to their operating costs and is typically assumed known for operators. However, for new services where no demand information is known, the ratio needs to be learned and adjusted based on observed user demand elasticities. The operator does not know the travelers' choice behavior, so the equilibration of the cost ratio is a method of calibrating operational policy to account for the latent demand function. We propose extending the user-operator day-to-day adjustment process from Djavadian and Chow (2017a) to endogenize the operator's learning of the perceived elasticity in order to maximize profit. While this day-to-day adjustment process is particularly well suited for the MAV application because demand is difficult to estimate for an entirely new transportation mode, it could be applied to any on-demand transportation service.

To the best of our knowledge, this is the first such study to make two important contributions:

1. We extend the day-to-day adjustment method from Djavadian and Chow (2017a) to adjustment of the mobility service operator's routing policy for maximum profit. This methodological contribution endogenizes the operator's routing policy by varying the level of service to learn the unknown latent demand elasticity. Our framework provides mobility service operators with a strategy to adapt to new markets or new services with uncertain demand.

2. We evaluate the benefits of dynamic en-route passenger transfer between vehicles on three different last mile service designs. A case study of commuters between Dubai and Sharjah, UAE drawn from real data is used in the evaluation. Using profit and welfare measures, we derive insights from the computational experiments that can inform service design for future MAV operators. This is the first study to evaluate last-mile operating strategies for MAV passenger transportation.

# 2 Dynamic routing with en-route transfers

## 2.1 Existing literature

Mobility services designed to solve the last mile problem have been studied extensively (Chang and Schonfeld, 1991a,b; Cortés et al., 2005; Mulley and Nelson, 2009; Chong et al., 2013; Wang and Odoni, 2014; Guo et al., 2017). In conventional last mile transit operations, either flexible on-demand service or feeder bus routes are used to connect passengers to a train station (Wang, 2017; Guo et al., 2017; Djavadian and Chow, 2017a,b). Yap et al. (2016) conclude that autonomous vehicles have some potential to solve the last mile problem as passengers become more comfortable with the technology; however, existing research has rarely considered passenger transfers when designing or evaluating mobility services.

General vehicle routing problems, including the dynamic Dial-a-Ride Problem (DDARP) with immediate requests and many-to-many pickup and delivery, has been the subject of a wide body of research (Berbeglia et al., 2010; Irnich et al., 2014). Braekers et al. (2016) provide a comprehensive overview of modern vehicle routing problems and solution methods. Emerging studies of routing problems consider the possibility of transferring passengers or freight at a fixed location. For example, routing algorithms developed in Cortés et al. (2010), Masson et al. (2013), and Rais et al. (2014) involve the transfer of goods at pre-defined transshipment points. This routing problem can be formulated as a mixed-integer linear program and solved using exact methods (Cortés et al., 2010; Rais et al., 2014). Deleplanque and Quilliot (2013) present an insertion heuristic algorithm for the static Dial-a-Ride Problem with transfers at any location. In their research, the freight load or passenger to be transferred is dropped off at the transfer node and picked up at some later time. Bouros et al. (2011) provide a solution to a similar dynamic problem based on solving a shortest path problem and compare it to a two-phase heuristic that combines insertion with tabu search. The latter heuristic is found to have lower cost solutions with higher computational cost (due to the tabu search). One recent study finds the optimal transfer point given a set of passengers using geometric constraints to reduce the search space (Melnick, 2018). An excellent survey of literature in vehicle routing with transshipment can be found in Guastaroba et al. (2016).

This study adapts an insertion heuristic, first developed by Jaw et al. (1986), for MAV operations and uses the proposed day-to-day adjustment model to evaluate the operations with the insertion heuristic. Subsequent research into insertion heuristics is summarized by Cordeau et al. (2005). Many different heuristics with added functionality have been proposed, including demand responsiveness and anticipation (Horn, 2002; van Engelen et al., 2018) and heterogeneous capacity (Madsen et al., 1995; Wong and Bell, 2006). Since Bouros et al. (2011) showed that insertion heuristics work well for routing freight with transfers, their algorithm is adopted (without the tabu search component to reduce computational cost) for the dynamic within-day routing of the MAV system so that the market equilibrium can be assessed for different last mile service designs.

## 2.2 En-route transfer algorithm

Consider a simulation with a set of travel requests, $R$, a set of vehicles available to service those requests, $V$, and a set of possible origin, destination and transfer nodes, $N$. All notation defined for the routing algorithm is included for reference in Table 1 below. Each travel request $r \in R$ has a request time $t_r^a \in \mathbb{Z}^+$, origin node $\mathbf{O}_r \in N$ and destination node $\mathbf{D}_r \in N$. The number of requests for travel is denoted by $|R|$.

The insertion algorithm is implemented in two steps. First, when a new request for travel $r$ is made at time $t_r^a$, the cost function $C(p_v)$ is calculated as shown in Eq. (1) for every feasible insertion position of the new origin node $\mathbf{O}_r \in N$ and destination node $\mathbf{D}_r \in N$ into the route of any vehicle in the fleet. The origin and destination of request $r$ are inserted at positions $k_r^O$ and $k_r^D$ into the route $p_v$ of a vehicle $v \in V$ such that the increase in cost is minimized.

The cost function has three weighted components: passenger travel time $\tau_r$, passenger wait time $\theta_r$ and vehicle distance travelled $d_v$. These terms represent the costs to both the user and the operator for inserting the request into a vehicle route. The arrival time at the origin node $t_r^O$ is determined by summing the time required to travel to all destinations that precede $\mathbf{O}_r$ in route $p_v$. The arrival time at the destination node $t_r^D$ is determined similarly. The travel time for request $r$ is the difference between $t_r^D$ and $t_r^O$, and the wait time for request $r$ is the difference between $t_r^O$ and $t_r^a$. Coefficients $\lambda$, $\rho$ and $\gamma$ govern the relative weight of the objective function components and determine whether user costs or operator costs are prioritized in route insertion decisions. If $\gamma$ is much larger than $\lambda$ and $\rho$, the result will favor fewer active vehicles with high average occupancy. In contrast, a high level of service can be achieved by setting the coefficients such that the costs of additional travel time and wait time dominate the objective function.

A set of insertion positions is considered feasible if the origin is inserted before the destination and the occupancy $u_{vt}$ of vehicle $v$ does not exceed the maximum vehicle capacity $\pi$ at any time $t$. Note that this structure assumes the operator serves all requests; no rejections are permitted. Other insertion heuristics may allow for rejected requests which are heavily penalized.

$$\min C = \sum_{r \in R} (\lambda \tau_r + \rho \theta_r) + \sum_{v \in V} \gamma d_v \tag{1}$$

In a typical dynamic dial-a-ride service with shared trips, this algorithm is repeated for each new request until a time limit is reached. When en-route transfers are considered, the algorithm requires a second step. Following the assignment of the new request to vehicle $v$, a list of other vehicles, $Q_i \subset V$, that might benefit from an en-route transfer with $v$ is generated. A vehicle is added to $Q_i$ if it has a route which will cause it to pass within a distance threshold of vehicle $v$ at some point during its route. For each candidate vehicle in $Q_i$, the cost function in Eq. (1) is calculated for a transfer of any feasible combination of passengers at any possible point in each vehicle's planned route. This includes any added travel time or wait time due to deviation from the initial planned route to the transfer point. If a feasible transfer insertion is found to reduce the objective function compared to the no-transfer set of routes, then the transfer is inserted into the route of both vehicles.

A transfer is considered feasible if both vehicles meet at the same location at the same time. The transfer of a passenger $r$ cannot occur before the passenger has been picked up or after they have been dropped off. A maximum of 2 transfers are permitted per passenger for passenger comfort. Once a transfer is inserted into the route of a vehicle, no changes to that vehicle's route prior to the transfer node are permitted. The two-step trip insertion algorithm is presented in Figure 4.

## 2.3 Illustration of the en-route transfer algorithm

To demonstrate the two steps of the algorithm, consider the simple example shown in Figure 5 with two vehicles $V = \{v_1, v_2\}$ serving four requests $R = \{r_1, r_2, r_3, r_4\}, \quad |R| = 4$ . Initially no requests have arrived to the system. For simplicity, we set $\lambda = \rho = 1$/minute and $\gamma = 1$/km. If we assume vehicles travel at 1 km per minute then one additional kilometer traveled has the same cost as one additional minute of travel or wait time. Requests $r_1$ and $r_2$ arrive at time $t = 1$, therefore $t_1^a = t_2^a = 1$. Inserting the origin and destination of request $r_1$ ($\mathbf{O}_1$ and $\mathbf{D}_1$) into $p_1$, the route of vehicle $v_1$, has $\tau_1 = \sqrt{8}$, $\theta_1 = 2$ and $d_1 = 2 + \sqrt{8}$. The cost is then $C(p_{1,1}) = 1(\sqrt{8}) + 1(2) + 1(2 + \sqrt{8}) = 9.7$. Alternatively, inserting $r_1$ into $p_2$ results in $C(p_{2,1}) = 1(\sqrt{8}) + 1(\sqrt{8}) + 1(\sqrt{8} + \sqrt{8}) = 11.3$. To minimize $C$, $r_1$ is assigned to $p_{1,1}$. By symmetry, $r_2$ is assigned to $p_{2,1}$ for a total system cost of 19.4. An en-route transfer of the passengers would not reduce the cost any further, so no transfers are inserted into the routes. Table 2 shows vehicle route assignments at each decision point.

Requests $r_3$ and $r_4$ arrive at $t_3^a = t_4^a = 2$, at which time the vehicles have progressed halfway to $\mathbf{O}_1$ and $\mathbf{O}_2$. Step 1 of the routing algorithm is executed, comparing the minimum cost of inserting $r_3$ into $p_1$ and $p_2$. The new request origins $\mathbf{O}_3$ and $\mathbf{O}_4$ can be inserted before or after the other destinations that have already been added to $p_{1,1}$ and $p_{2,1}$. By inspection, the route cost of inserting $r_3$ into $p_{1,2}$ will be minimal if $\mathbf{D}_3$ is served after $\mathbf{D}_1$. The additional cost of this insertion

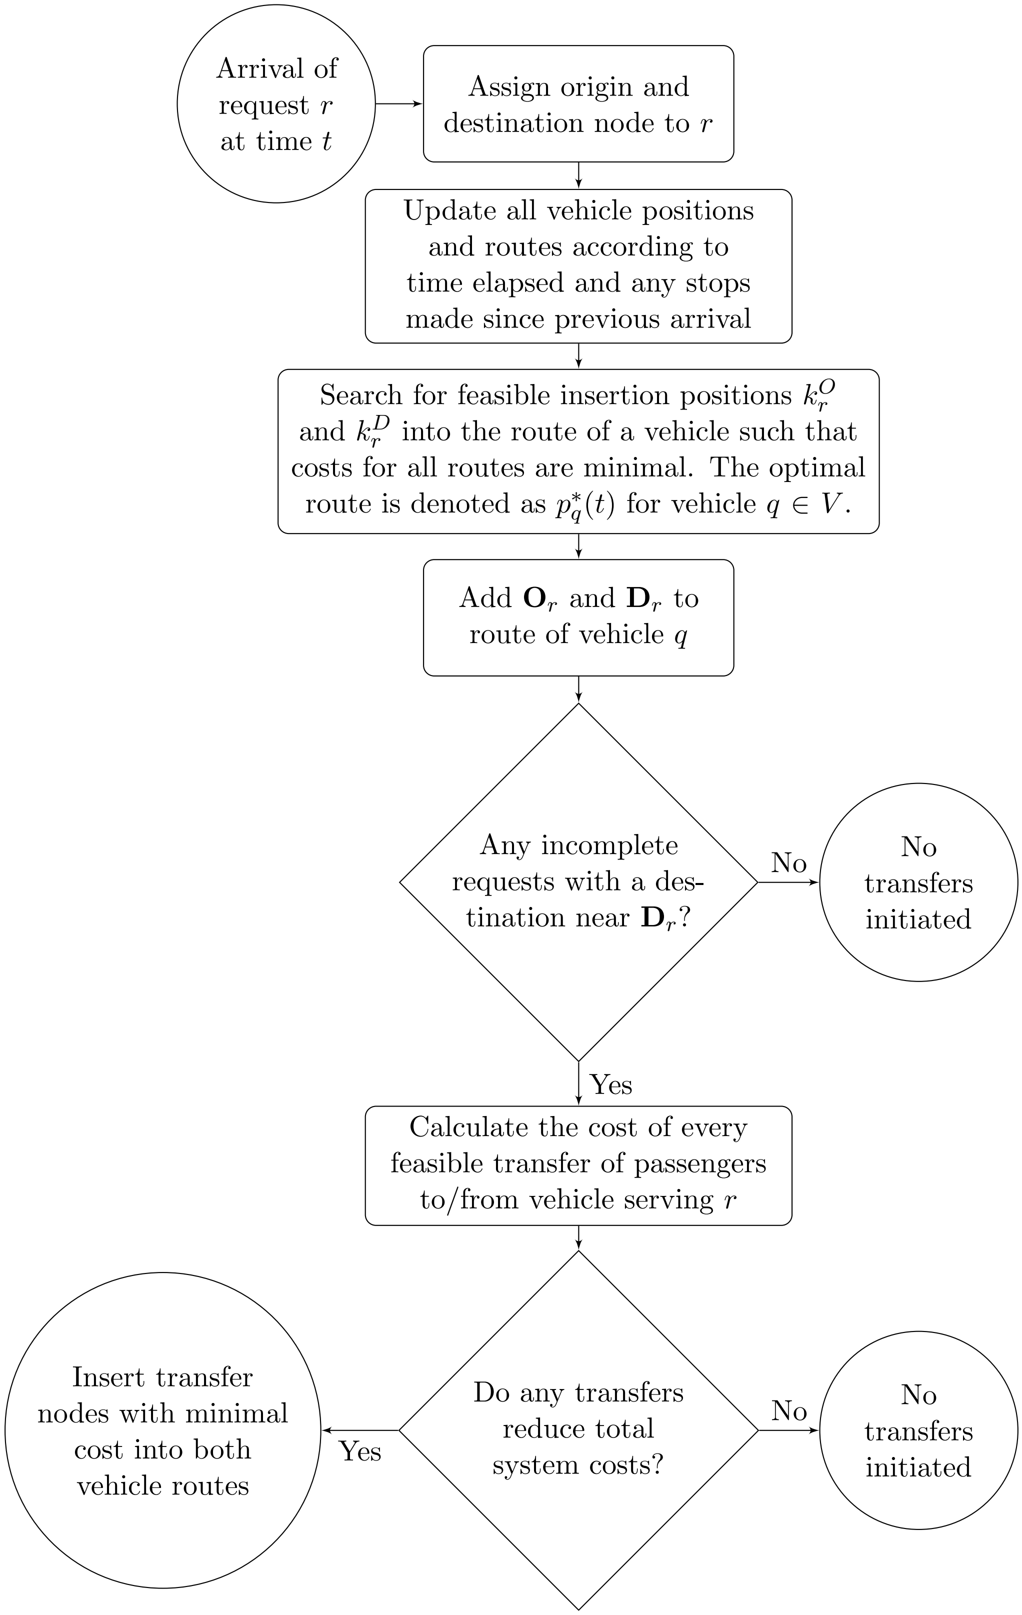


Figure 4: En-route transfer algorithm flowchart

Table 1: Algorithm Notation

| | |
|---|---|
| $N$ | The full set of nodes in the simulation network |
| $n$ | The index corresponding to a specific node where $n \in N$ |
| $R$ | The full set of requests for pickup and delivery |
| $r$ | The index corresponding to an individual request where $r \in R$ |
| $\|R\|$ | The total number of possible requests for travel, where $\|R\| \in \mathbb{Z}^+$ |
| $\mathbf{O}_r$ | A vector representing the Euclidian coordinates for the origin of request $r$ where $\mathbf{O}_r \in N$ |
| $\mathbf{D}_r$ | A vector representing the Euclidian coordinates for the destination of request $r$ where $\mathbf{D}_r \in N$ |
| $t_r^a$ | Time at which request $r$ is made |
| $t_r^O$ | Time at which request $r$ is picked up from the requested trip origin |
| $t_r^D$ | Time at which request $r$ arrives at the requested trip destination |
| $V$ | The fixed full set (or fleet) of vehicles in the simulation |
| $v$ | The index corresponding to an individual vehicle where $v \in V$ |
| $x_{vt}, y_{vt}$ | The coordinates defining the position of vehicle $v$ at time $t$ |
| $p_{vt}$ | The route for vehicle $v$ consisting of a set of origin, destination and transfer nodes $n$ |
| $k_r^O, k_r^D$ | The insertion positions of the origin and destination of request $r$ in a route |
| $u_{vt}$ | The occupancy of vehicle $v$ at time $t$ |
| $\pi$ | The maximum passenger capacity of a vehicle |
| $I$ | The full set of en-route passenger transfers |
| $i$ | The index corresponding to a specific passenger transfer between two or more vehicles where $i \in I$ |
| $\mathbf{s}_i, \mathbf{f}_i$ | Vectors representing the Euclidian coordinates for the start and finish nodes of passenger transfer $i$ where $\mathbf{s}_i, \mathbf{f}_i \in N$ |
| $G_i$ | The subset of nodes which are considered for $\mathbf{s}_i, \mathbf{f}_i$, which are located between the current positions of vehicles participating in transfer $i$ and the destinations of the passengers being considered for transfer where $G_i \subseteq N$ |
| $Q_i$ | The subset of vehicles participating in transfer $i$ where $Q_i \subseteq V$ |
| $\theta_r$ | Wait time of request $r$ |
| $\tau_r$ | Travel time of request $r$ |
| $d_v$ | Distance travelled by vehicle $v$ |
| $\rho$ | Passenger wait time cost coefficient |
| $\lambda$ | Passenger travel time cost coefficient |
| $\gamma$ | Vehicle travel distance cost coefficient |
| $C(p_v)$ | Cost function: the sum of operator and traveller costs for a given route or set of routes |

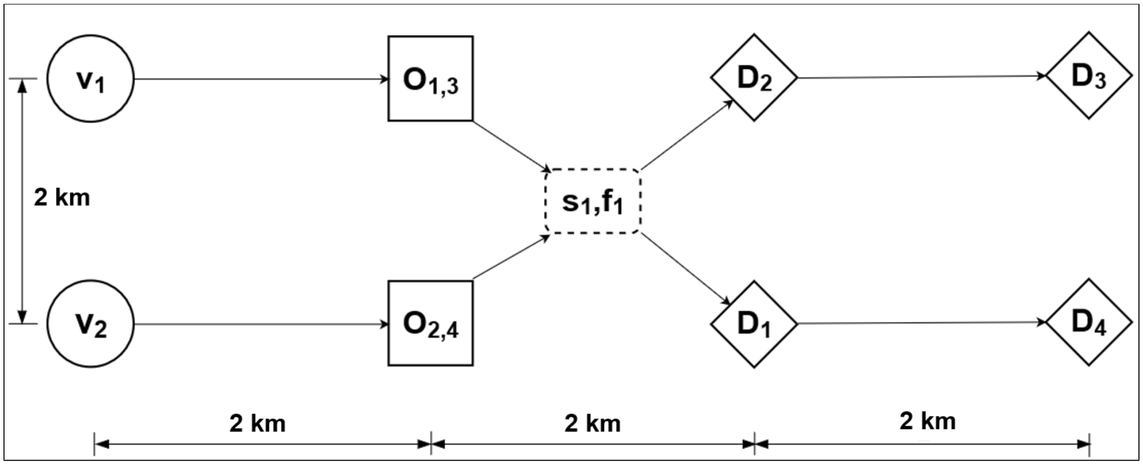


Figure 5: Simple routing problem illustration

Table 2: Vehicle routes in routing problem with transfers

| Time $t$ | $p_{1,t}$ | $p_{2,t}$ |
|---|---|---|
| $t = 0$ | - | - |
| $t = 1$ | $[\mathbf{O}_1, \mathbf{D}_1]$ | $[\mathbf{O}_2, \mathbf{D}_2]$ |
| $t = 2$ (no transfer) | $[\mathbf{O}_1, \mathbf{O}_3, \mathbf{D}_1, \mathbf{D}_3]$ | $[\mathbf{O}_2, \mathbf{O}_4, \mathbf{D}_2, \mathbf{D}_4]$ |
| $t = 2$ (with transfer) | $[\mathbf{O}_1, \mathbf{O}_3, \mathbf{s}_1, \mathbf{f}_1, \mathbf{D}_1, \mathbf{D}_3]$ | $[\mathbf{O}_2, \mathbf{O}_4, \mathbf{s}_1, \mathbf{f}_1, \mathbf{D}_2, \mathbf{D}_4]$ |

is $C(p_{1,2}) = 1(2\sqrt{8}) + 1(1) + 1(\sqrt{8}) = 9.5$. The least additional cost of inserting $r_3$ into route $p_{2,2}$ is 10.2, therefore $r_3$ is assigned to $p_{1,2}$ and similarly $r_4$ is assigned to $p_{2,2}$. The total system cost is now $2(9.7 + 9.5) = 38.4$.

The algorithm then searches for en-route passenger transfers that reduce system cost. In this instance, the vehicles considered for passenger transfer $i$ is $Q_i = \{v_1, v_2\}$. The optimal location for an en-route transfer between every point in $p_{1,2}$ and $p_{2,2}$ is determined. The cost of inserting an en-route transfer at that location is then calculated for every combination of passengers. It is found that if $r_3$ and $r_4$ are transferred immediately after each request is picked up, the resulting system cost is decreased to 35.0 from 38.4. The transfer $i$ is then added to both vehicle routes at $\mathbf{s}_i$, $\mathbf{f}_i$. This example demonstrates how the ability to use passenger transfers reduces overall system cost.

### 2.4 Comparison of en-route transfer algorithm to typical ridesharing service

To illustrate the impact of en-route transfers without considering day-to-day adjustment, we created a toy model with a simple deterministic demand profile to simulate a single day of service where the routing algorithm described above is used to assign requests to vehicles and initiate en-route transfers. This was compared to a second scenario with the same demand profile but with en-route transfers forbidden. 200 requests were generated with uniformly-distributed arrival times within a 4-hour period and uniformly-distributed origins and destinations within a simulated 5-mile by 5-mile grid. These simplifications may not accurately represent a typical travel demand profile, but they are suitable for this simple example. A visualization of the simulation in progress is shown in Figure 6. Requests were served by a fleet of 5 vehicles and the user/operator cost ratio was set to

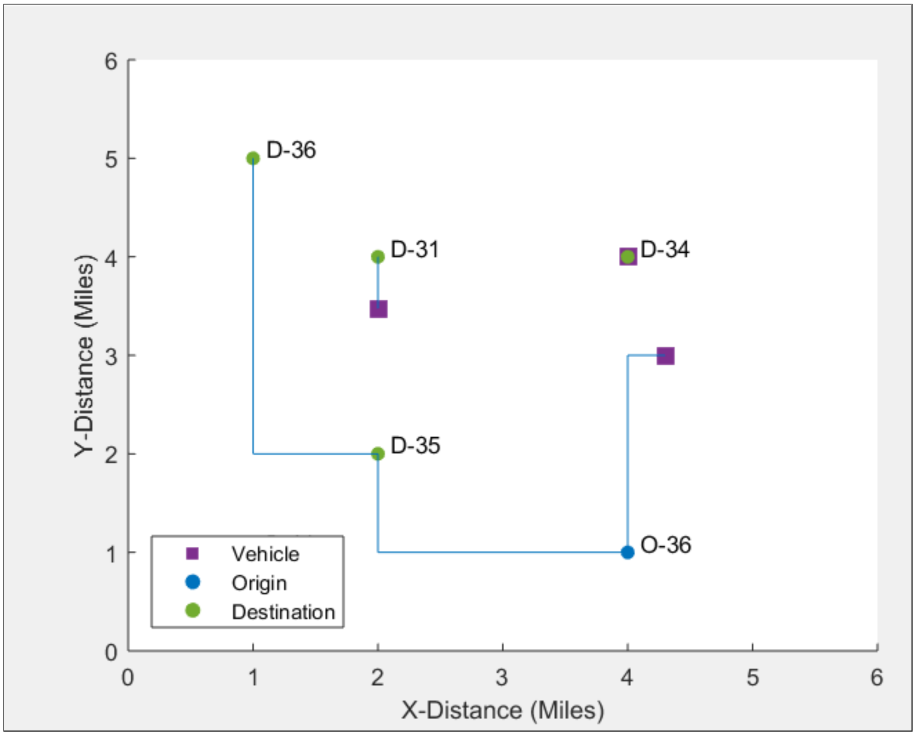


Figure 6: Illustration of simulated vehicle routes with fixed demand

1.0. Average travel times, wait times, and vehicle miles travelled for each scenario are reported in Table 3.

Table 3: Comparison of mobility services with and without en-route transfers

| Service type | Average wait time (min.) | Average travel time (min.) | Total vehicle miles traveled | En-route transfers |
|---|---|---|---|---|
| Without en-route transfers | 2.64 | 10.66 | 720.7 | - |
| With en-route transfers | 2.20 | 10.27 | 741.5 | 13 |

In this illustrative example, users of the service with en-route transfers experience an average of 16.7% less wait time and 4% less travel time when en-route transfers are available. The routing algorithm found 13 routes with transfers that reduced overall system cost, and in each case only one passenger was transferred between vehicles. The improvement in service for the user is offset by an increase of 2.8% in total vehicle miles driven when en-route transfers are used.

The operator may seek to increase the weight of the operator cost term relative to the user cost terms in the routing algorithm to avoid such an increase in travel distance. Similarly, more travelers may choose to use the service due to the level of service improvement, which can impact the performance of the overall system. The "naïve" comparison in this section highlights the lack of learning behavior from the operator and users.

# 3 Day-to-day adjustment

## 3.1 Day-to-day adjustment process overview

Typical day-to-day adjustment processes can be described as "rational behavior adjustment processes" (Yang and Zhang, 2009). Day-to-day adjustment has long been applied to route choice and traffic assignment models (Iida et al., 1992; Hu and Mahmassani, 1997). Djavadian and Chow (2017a) used a day-to-day adjustment process to evaluate flexible transit services that exhibit dynamic within-day operating policies. Jha et al. (1998) and Pacheco et al. (2016) applied the adjustment mechanism to mode choice, using an agent-based day-to-day adjustment approach to simulate the unobservable preferences of a population and determine travel choice utilities for each agent. Their approach is incorporated into this study.

Transportation systems evaluation (Chow, 2018) typically requires survey data to estimate user demand response to different system designs. This is not possible for MAV systems because they are not yet commercially available. As a result, there is no data on user response which also means that an assumed value of certain parameters like user/operator cost ratio can lead to suboptimal operations. One solution is to endogenize the learning behavior of both the operator and the users of the system. The operator's assumptions are updated on a day-to-day basis to reach optimal operations. These decisions include traveler mode choice and departure time choice, and operator choice of parameters used for routing decisions.

Djavadian and Chow (2017b) proposed a two-sided market day-to-day adjustment process to incorporate learning from both operators and users. The two-sided market simulation framework from Djavadian and Chow (2017b) is applied to the MAV system through operational strategies. The framework requires parameterizing the transfer policy such that it can be adjusted from one day to the next based on the information available to the operator: observed performance and demand on previous days. Details of the day-to-day adjustment process used in the simulation model are provided in the following subsections.

The system is characterized by two types of actors: the centralized operator of the MAV service and the agents who choose MAV or an alternative mode to serve their request for travel each day. For each day $n$, an operator/user cost weight $w_n$ is chosen by the operator. The operator/user cost weight $w_n$ is the ratio of the operator cost term $\gamma$ to the travel time and wait time cost terms, $\lambda + \rho$, in the routing algorithm cost function. This decision is based on the operator's profit from the previous day, $P_{n-1}$.

Each simulation day is 24 hours. The simulation continues until a stopping criterion based on sustained day-to-day convergence of the average experienced travel time $\overline{TT_E}$ is reached as determined by Eq. (2). When the parameter $\phi \geq 0.5\%$, convergence is typically reached in under 100 days, although certain parameter settings can lead to oscillation between multiple attractors and no stable solution.

$$\frac{\overline{TT_E}(n-i) - \overline{TT_E}(n-i-1)}{\overline{TT_E}(n-i-1)} \leq \phi \qquad 0 \leq i \leq 2 \tag{2}$$

## 3.2 Day to day demand variation

Each agent has an initial expectation of travel time $\tau_r^0$ based on a linear function of their trip distance. After each day $n$, the agent updates their expected total travel time for the next day $ET_r^{n+1}$ based on their experienced travel time $TT_r^n$ as determined by Eq. (3). Total travel time

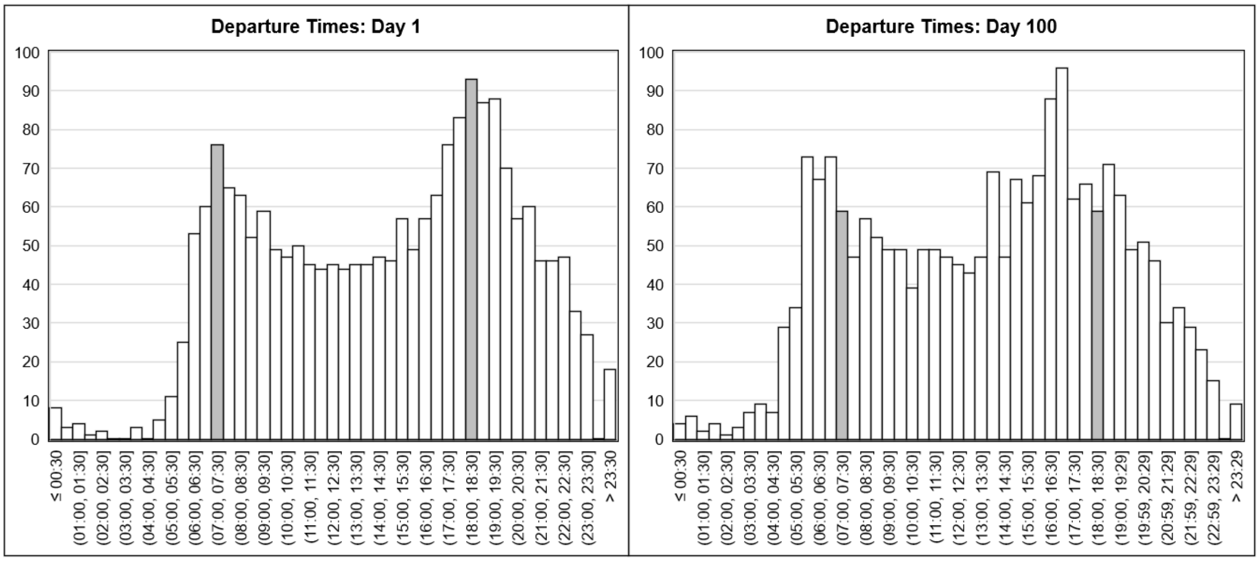


Figure 7: Day-to-day adjustment of simulated agent departure times on Day 1 and Day 100.

includes time spent waiting for pickup. A learning parameter $0 \le \omega \le 1$ is used to weigh the experience of the previous day against all previous experiences. The agent also updates their planned departure time for the next day $t_{d,r}^{n+1}$ by the same amount in order to reach their destination at the desired arrival time (Eq. 4). Figure 7 illustrates how simulated departure times shift to an earlier time period if users experience longer than expected travel times over 100 days.

$$ET_r^{n+1} = ET_r^n + \omega(TT_r^n - ET_r^n) \qquad \forall\ r \in R \tag{3}$$

$$t_{d,r}^{n+1} = t_{d,r}^n + \omega(TT_r^n - ET_r^n) \qquad \forall\ r \in R \tag{4}$$

Agent mode choice is determined for each day using a random utility-based discrete choice model (see Train, 2009) with multiple alternatives. The representative utility $V_{l,r}$ of each alternative $l$ among a set of alternatives $L$ is determined for each agent $r$ based on the mode and agent-specific trip characteristics $x_{l,r,i}$ in Eq. (5) and a coefficient $\beta_i$ weighting the effect of that characteristic on utility. A variate $\varepsilon_{l,r}$ which accounts for taste heterogeneity within the population and the unobserved determinants of utility is added to the representative utility to get the random utility $U_{l,r}$ as shown in Eq. (6).

The probability of an agent choosing mode $m$ among a set of alternative modes $M$ is determined by Eq. (7). A random draw from the probability function generated by the discrete choice model determines the simulated mode choice for each agent. Incorporating this mode choice model adds stochasticity to the simulation and creates day-to-day demand variation based on an agent's experienced level of service.

$$V_{l,r}(x_{l,r}) = \beta_1 x_{l,r,1} + \beta_2 x_{l,r,2} + ... + \beta_i x_{l,r,i} \qquad \forall\ r \in R,\ l \in L \tag{5}$$

$$U_{l,r}(x_{l,r}) = V_{l,r}(x_{l,r}) + \varepsilon_{l,r} \qquad \forall\ r \in R,\ l \in L \tag{6}$$

$$Pr_r(m|M) = \frac{e^{U_{m,r}}}{\sum_{l \in M} e^{U_{l,r}}} \qquad \forall\ r \in R,\ m \in M \tag{7}$$

Some simplifying assumptions are made in determining mode choice. It is assumed that every alternative is available to each agent (i.e. each agent has a car that they are licensed to drive) since car ownership data is not available. To evaluate the mode choice alternatives for each service design described in Section 4, it is necessary to calculate the alternative-specific trip characteristics: travel time and travel cost. Average travel speeds are used to calculate travel time for MAV, public bus, bicycle (Litman, 2009) and walk modes (Knoblauch et al., 1996). These are presented in Table 4. A MAV is assumed to travel at the same speed as a typical car. Ranges of costs per unit distance for motorized modes are summarized in Litman (2017). The average value is used for each mode, converted to $/km. Cycling costs from Litman (2010) are used, and walking cost is assumed to be zero.

Table 4: Average travel speed and cost by mode

| Parameter | MAV | Personal Car | Public Bus | Bicycle | Walk |
|---|---|---|---|---|---|
| Travel Speed (km/hr) | 45 | 45 | 20 | 20 | 4.5 |
| Travel Cost ($/km) | 0.3 | 0.5 | 0.3 | 0.1 | 0.0 |

Disutility for previous early and late arrivals is included for the MAV mode. Small (1982) found that travelers perceive an arrival time more than 2.5 minutes after the desired arrival time to have disutility equivalent to 5.5 minutes of added travel time, with another 2.4 minutes of disutility for every additional minute thereafter. Early arrivals also have disutility to travelers, although the effect is smaller: 0.61 minutes of travel time for each minute before the desired arrival time. These penalties are added to the experienced travel time for each agent using the MAV mode. It is assumed that the agent has perfect knowledge of the travel time for other modes and is neither late nor early if an alternative mode is chosen.

Disutility is also included in the case that a passenger is required to make a transfer. This disutility arises from the forced change of seat within two or more connected modular autonomous vehicles during an en-route transfer. We expect that the disutility associated with an en-route transfer would be considerably lower than a transfer between travel modes as the distances would be smaller and the passenger would remain indoors. Furthermore, the disutility is expected to vary with the duration of time provided for the transfer; transfers with short windows would have more disutility but would reduce the amount of time that the two vehicles must remain coupled. It is important to include this disutility in the mode choice model for an accurate comparison across modes, but choosing an appropriate value is not trivial. Forced seat changes are not typically required for any existing transportation mode; therefore, no research has been devoted to modelling the associated disutility. We therefore created two separate test cases for the simulation, a high transfer penalty case and a low transfer penalty case. The sensitivity of the simulation results to variation in the transfer disutility term is tested.

### 3.3 Day-to-day adjustment of operator strategy

In this day-to-day adjustment framework, the mobility service provider seeks to maximize profit. An operational strategy defined as a set of rules and parameters associated with those rules is

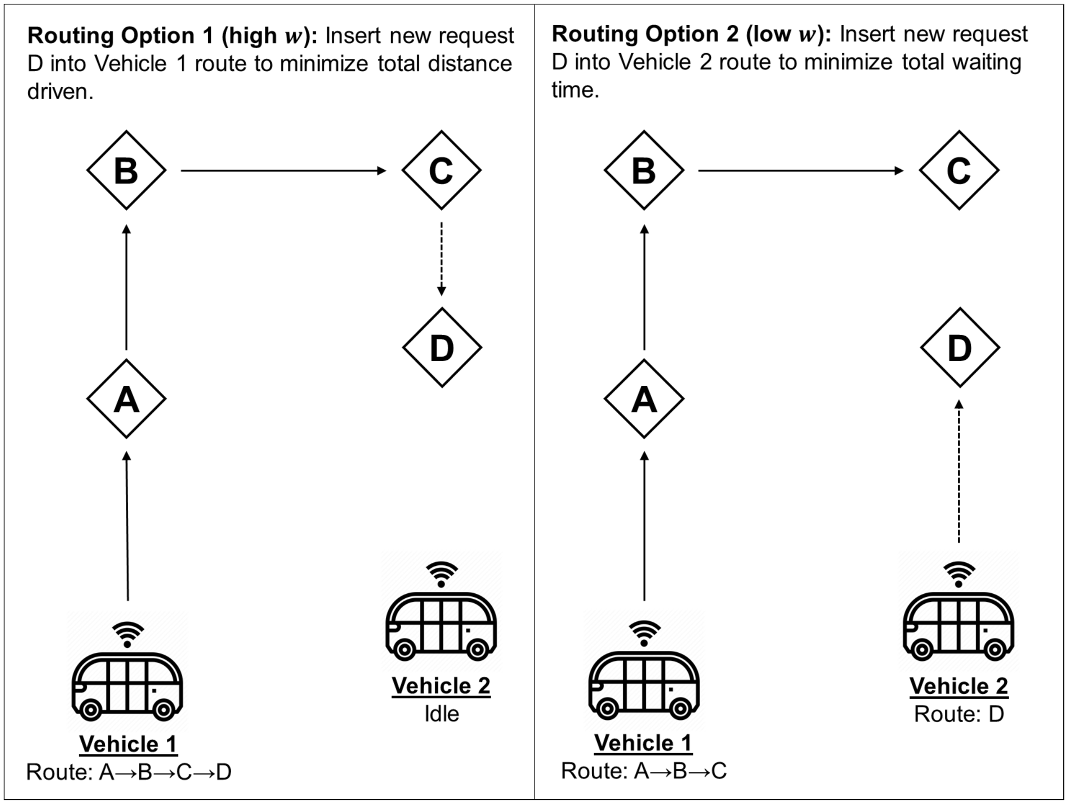


Figure 8: Comparison of routes generated under different cost ratios

chosen for a given day. By varying the operational strategy each day, the operator can learn from and adapt to variable demand conditions. Rather than making capital intensive adjustments such as increasing fleet size, the operator can instead affect the level of service simply by modifying the weight of different terms in the routing algorithm. A novel contribution of this study is the day-to-day endogenization of the relative magnitudes of the terms in the objective function to maximize profit given a set of demand conditions. This framework was developed for the MAV service but could be extended to other on-demand transportation services.

The terms in the objective function are weighted by coefficients: $\lambda$ for travel time, $\rho$ for waiting time and $\gamma$ for vehicle distance travelled. When the operator/user cost ratio $w$ is high, the routing algorithm favors routes which minimize distance travelled over routes which result in less travel time or waiting time for the users. As shown in Figure 8 (Option 1), the algorithm would be more likely to assign a new pickup request to a vehicle whose existing assigned route will cause it to pass near the new request origin, but only after making several other stops, rather than dispatch an idle vehicle that would reach the new request in much less time. Conversely, a low $w$ would cause the operator to select Option 2 where user wait time is minimal.

Operator profit $P_n$ on day $n$ is difference between the sum of revenues from transporting users and the cost to provide each trip. Profit is an analog for operator welfare. Revenues are assumed to be proportional to the distance between origin and destination, and the marginal operating costs to provide service are assumed to be a function of total vehicle distance travelled. Profit is therefore determined by Eq. (8). There is no difference in revenue between shared trips and private trips.

$$P_n = \alpha \sum_{r \in R} ||\mathbf{O}_r - \mathbf{D}_r|| - \mu \sum_{v \in V} d_v \tag{8}$$

Constants $\alpha$ and $\mu$ set the level of revenue and operating cost for each unit of distance. As an illustration of how adjusting $w$ can increase profits, consider $\alpha \gg \mu$ so that each trip served increases profit. Under these conditions, the operator would seek to serve as many trips as possible each day. Given that demand for the MAV mobility service is variable and increases when experienced travel times are low, the operator would be able to increase demand and therefore profit by reducing experienced travel time. To accomplish this, the operator sets a low $w$, forcing the routing algorithm to favor routes with minimal travel time. Other values for $\alpha$ and $\mu$ and variation in other simulation parameters result in differing optimal values for $w$.

Operators may not know the optimal setting for their algorithm given a set of market dynamics, especially when those dynamics are themselves unknown. The simulation assumes an initial $\gamma^0$ which is then adjusted after each day of service based on actual daily profit. The cost coefficient for vehicle miles driven on day $n$ is represented by $\gamma^n$. The ratio of the previous day's profit to the average daily profit over the past $\delta$ days is used to update the magnitude of $\gamma$ in Eq. (9).

$$\gamma^{n+1} = \gamma^n + \frac{\gamma^n - \frac{\sum_{n-\delta}^{n} \gamma^n}{\delta}}{|\gamma^n - \frac{\sum_{n-\delta}^{n} \gamma^n}{\delta}|}(1 - min[\frac{\delta P_n}{\sum_{n-\delta}^{n} P_n}, 0.75]) \tag{9}$$

The day-to-day learning mechanism represented by Eq. (9) for a central operator is designed to converge toward a maximum profit. If the profit is positive, $\gamma$ will continue to increase or decrease depending on the trend across the past $\delta$ days. For example, if $\gamma^n$ is less than the average $\gamma$ over the past $\delta$ days and the profit is positive, then $\gamma$ will be reduced further. If simulation settings are such that nearly every trip served is profitable, the day-to-day operating strategy adjustment forces $\gamma$ to approach zero over time. Similarly, if profitable trips are rare, the weight of the vehicle distance traveled term will rise indefinitely in an attempt to limit operating costs.

The day-to-day change has a lower bound, -25%, to prevent large variation if the stochastic demand results in a day with much lower profits than previous days. The day-to-day operating strategy has been empirically shown to converge from different initial values and when perturbed. Figure 9 shows how the cost weight $\gamma$ converges to a similar range regardless of initial setting from day 1 to 100. Convergence was also tested by simulating positive and negative perturbations every 100 days. The result suggests that there is a $\gamma$ setting that maximizes profit given a set of simulation parameters, and that the day-to-day cost ratio adjustment described in Eq. (9) converges to that setting.

### 3.4 Illustration of cost ratio learning

Consider the following example to demonstrate the day-to-day learning and adjustment of the operator/user cost ratio. A MAV service operator generated a profit $P_n = 1,200$ units on day $n$ during which the operator cost $\gamma^n = 1.0$. Assume $\delta = 5$ days. The average profit $P$ over the past 5 days including day $n$, $\sum_{n-\delta}^{n} P_n/\delta$ is 1,000 units. The average $\gamma$ over the past 5 days is 1.2. The profit has increased by 20% over a period where $\gamma$ was reduced. The operator observes that reducing $\gamma$ leads to larger profits, therefore $\gamma^{n+1}$ will be reduced by 20% compared to $\gamma^n$, or $\gamma^{n+1} = 0.8$. By reducing the cost ratio $w$, the routing algorithm will favor routes with less waiting and travel time, thereby increasing the likelihood of a user choosing the MAV service and boosting operating


Figure 9: Convergence of operating cost weight term after perturbations

revenues. If, after another day of service, $P_{n+1} = 1,100$, the operator would observe that the maximum profit is obtained when $0.8 \leq \gamma \leq 1.0$. The operating cost for the next day is updated using Eq. (9) to get $\gamma^{n+2} = 0.87$. In this case, if no cost ratio learning was implemented, the operator would continue to use a strategy where profits were less than the maximum for the given system parameters. Figure 10 presents the general steps of the day-to-day adjustment process.

# 4 Experiment Design

## 4.1 Simulation overview

The performance of a mobility service capable of en-route transfers is compared against a modern ridesharing service without transfers like UberPool or Lyft Line. Both services are tested with three different service designs: urban door-to-door service, first/last mile service to a transit hub, and hub-and-spoke service. Operations within each day are simulated using the online insertion heuristic described in Section 2 to assign vehicles to new requests. The three last-mile service designs to be compared, with and without en-route transfers, are illustrated in Figure 3.

Each service design uses publicly available transit ridership data including origin and departure time from the Dubai Roads and Transport Authority (Dubai Roads and Transport Authority, 2018) as the demand input. Destination stations were not included in the dataset and were generated based on boarding location distributions. Note that no validation scenario is evaluated because level of service data is not available for this case study, which means predictions cannot be made. The focus of this contribution is on a comparative evaluation of MAV technology to a non-MAV-based on-demand mobility service which is adequate with the service designs defined.

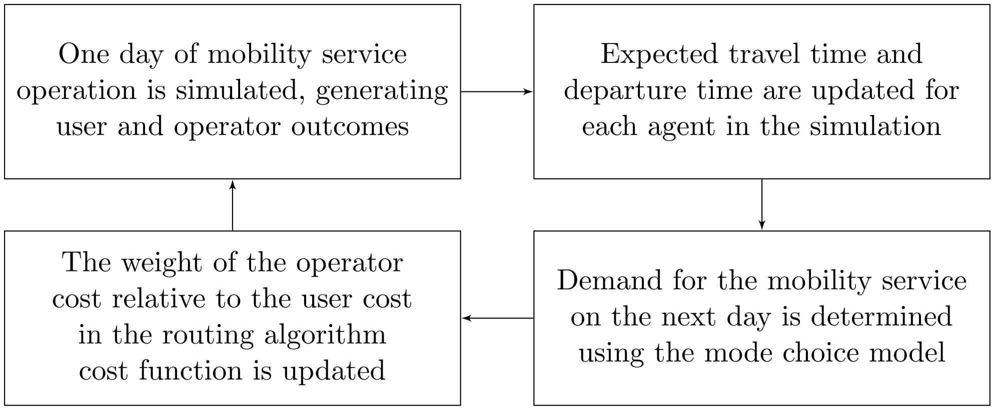


Figure 10: Diagram of the feedback system

These transit demand profiles provide a realistic distribution of demand for transportation in a heavily congested region. Two of the three MAV services are explicitly designed to complement or replace existing public transit, so transit demand provides a reasonable approximation of the future MAV demand. We also anticipate that current transit riders would be likely to adopt the MAV mode share given that it has similar characteristics to public transit: shared vehicles with potential for crowding, low cost, and relatively high travel time.

For this experiment, travel times are assumed to be constant throughout the day; in reality, travel times in Dubai and other urban areas vary throughout the day due to the effects of congestion.

Figure 11 displays typical peak hour traffic between Dubai and the neighboring city Sharjah, where travel time for a 31 km trip can exceed 2 hours. A solution to congestion in this region is desperately needed: 80% of commuters between these two cities experience congestion during their trip and more than half report negative physical and psychological impacts due to the time spent waiting in traffic (Hussein, 2017). The within-day service and day-to-day adjustments are simulated using MATLAB 2017a.

## 4.2 Service Design 1: Urban door-to-door service

Service design 1 uses the real demand distribution from the 3 busiest bus routes within the central area of Dubai, serving approximately 32,000 daily trips between 71 stations. This demand profile tests the MAV service without en-route transfers when serving trips from origin to destination within a dense urban area, similar to a modern taxi or ride-hailing service. Travelers have the option to use the MAV service, drive a personal vehicle, use public transit, walk or bike from their trip origin to destination. Since no disaggregate demand data is available for the travelers, coefficients for the discrete choice model used to determine agent mode choices are adopted from Bhat and Sardesai (2006) and shown in Table 5. The MAV mode is assumed to have model coefficients similar to the shared-ride alternative.

## 4.3 Service Design 2: First/last mile service

The second service design seeks to test the performance of the MAV system in a last mile service where passengers have the same origin or destination (a fixed-route transit hub). In this case, the demand profile for bus service to Dubai from Sharjah is used. These buses primarily serve

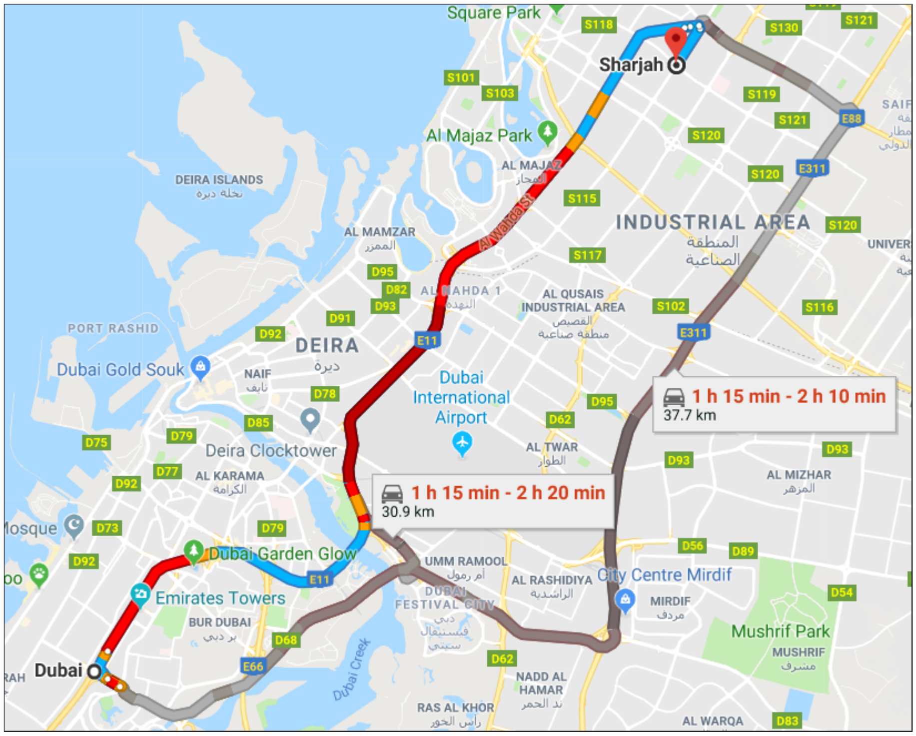


Figure 11: Peak hour travel time from Sharjah to Dubai
*Source:* Google Maps (2018)

Table 5: Mode choice model coefficients (Service design 1)

| Mode | Alternative constant coefficient | Travel time coefficient (utility/min) | Travel cost coefficient (utility/\$) |
|---|---|---|---|
| Car | 0.000 | -0.053 | -0.262 |
| Public transit | 0.000 | -0.053 | -0.262 |
| Bicycle | 0.402 | -0.018 | 0.000 |
| Walk | 0.402 | -0.079 | 0.000 |
| MAV | -2.214 | -0.053 | -0.262 |

commuters who travel from Sharjah to Dubai for work in the morning and then return to Sharjah in the evening.

To simulate first/last mile or feeder service, each traveller from Sharjah is assumed to be travelling to the main bus terminal in Sharjah where they would currently catch an express bus to Dubai. Return trips would originate from the Sharjah bus terminal, terminating at the user's final destination. The demand input contains 20,500 daily trips with 18 stations in Sharjah. In this service design there are four mode choices: drive personal vehicle to the train station, MAV service to train, bicycle to train and bus to train. Logit model coefficients for this set of alternatives were estimated in a stated preference survey by Yap et al. (2016) and are used for the mode choice model in this service design. The coefficients are given in Table 6.

Table 6: Mode choice model coefficients (Service Designs 2 and 3)

| Mode | Alternative constant coefficient | Travel time coefficient (utility/min) | Travel cost coefficient (utility/$) |
|---|---|---|---|
| Car | 0.000 | -0.031 | -0.200 |
| Bus to train | -1.510 | -0.040 | -0.570 |
| Bicycle to train | -1.330 | -0.080 | -0.460 |
| MAV to train | -1.840 | -0.084 | -0.410 |

### 4.4 Service Design 3: Hub-and-spoke service

The third service design tests the performance of the MAV service as a hub-and-spoke service. Passengers are picked up at their trip origin and taken to a hub to be consolidated into a convoy traveling to their destination. The convoy splits once it reaches the opposite hub to deliver travelers to their final destinations. The same demand profile as Service Design 2 is used, with the exception that all trips originating in one city are assigned destinations in the opposite city. Destination stations are assigned according to the distribution of boardings at each station as determined from the bus ridership data. The mode choice model for Service Design 3 is the same as that for Service Design 2 given that both represent commuter travel. The MAV service with en-route transfers is compared to a service without transfers, which does not stop at any hub as no consolidation or convoys can occur without en-route transfers.

### 4.5 Simulation parameters

The simulation parameters are described below in addition to those mentioned previously. Two cases are run for each service design to test traveler sensitivity to the en-route transfer disutility penalty. The high penalty term is similar to the values reported for a mode change transfer, which we considered to be the upper bound for en-route transfer disutility.

- Vehicle capacity: 10 passengers
- Fleet size:
  - Service Design 1: 100 vehicles
  - Service Design 2: 20 vehicles
  - Service Design 3: 30 vehicles
- Boarding and alighting time: 45 seconds
- Learning rate $\omega$: 0.2 (Djavadian and Chow, 2017b)
- En-route transfer penalty
  - Case 1: 1 minute of travel time
  - Case 2: 5 minutes of travel time
- Revenue per kilometer of trip length $\alpha$: \$1.6/km

- Operating cost per kilometer driven $\mu$: \$1.0/km
- Stopping criterion parameter $\phi$: 0.5%

Vehicle capacity is based on the specifications for the NEXT NX1 modular autonomous vehicle (The National UAE, 2018). Fleet size is limited such that reasonable but not excessive delays occur during peak hours given the travel demand profile and type of service. The boarding time assumes that passengers are waiting on the street for the MAV to arrive. Revenue and operating cost per mile are set such that the marginal profits per trip are typically small. This forces the simulated operator to vary $\gamma$ to optimize profit in a competitive market where profit margins are limited.

### 4.6 Evaluation metrics

The evaluation metrics for each scenario are listed below. Daily ridership is a measure of the demand for the service. 6 scenarios are tested: 3 service designs with 2 cases for each of the transfer penalty alternatives. The average of 20 simulation runs for each scenario is used due to the stochasticity of the simulation.

- Average passenger wait time
- Average passenger travel time
- Daily ridership
- Days to reach stopping criterion
- Average daily profit
- Daily consumer surplus

Measuring consumer surplus is a method of quantifying a user's benefit arising from a set of choices or observed variables. It is the scalar sum of the expected value of a set of alternatives. Given a discrete choice model estimation, the change in consumer surplus (alternatively referred to as social welfare or measure of accessibility) due to changes in trip characteristics can be determined. The consumer surplus of a model with mode choice set $M_r$ for an individual $r$ in units of trip cost is calculated using Eq. (10). $\overline{\beta}_{cost}$ represents the average of the random coefficient pertaining to trip cost across the population and $V_{m,r}$ is the representative utility of mode $m$ for individual $r$. Taking the sum of the surplus across all agents and dividing by the number of agents $|R|$ gives the average social welfare $\overline{E}_r$ of a set of choices for each trip. Comparing consumer surplus for different models can be used to determine whether the different mode choices available improve utility for a set of travelers, regardless of which mode they choose.

$$\overline{E}_r = -\frac{1}{\overline{\beta}_{cost}|R|} \sum_{r \in R} ln[\sum_{m \in M_r} e^{V_{mr}}] \tag{10}$$

# 5 Results

## 5.1 Day-to-day adjustment algorithm performance

Average time until convergence and the average cost ratio ($w$) reached at convergence over 20 simulation trials are reported for each service design in Table 7. Convergence of $w$ for the different service designs is shown graphically in Figure 12. En-route transfers do not appear to have a significant effect on the convergence time, although convergence does occur more quickly in Service Design 2. The MAV service with en-route transfers has a higher variability in travel times in all cases. The cost ratio $w$ is higher when en-route transfers are used for Service Designs 1 and 2, while the opposite is true for Service Design 3.

Table 7: Comparison of day-to-day adjustment parameters

| Service design | Service type | Days to converge | $w$ (mean) | $w$ (st. dev.) | $TT_E$ (mean) | $TT_E$ (st. dev.) |
|---|---|---|---|---|---|---|
| 1 | Without en-route transfers | 76.0 | 0.863 | 0.079 | 18.10 | 0.49 |
| | With en-route transfers | 75.1 | 0.887 | 0.060 | 16.43 | 0.51 |
| 2 | Without en-route transfers | 66.3 | 0.691 | 0.035 | 34.72 | 2.81 |
| | With en-route transfers | 68.6 | 0.803 | 0.057 | 34.32 | 3.31 |
| 3 | Without en-route transfers | 74.9 | 0.758 | 0.104 | 46.37 | 4.14 |
| | With en-route transfers | 70.4 | 0.565 | 0.069 | 50.54 | 5.54 |

The parameter that is “learned” through day-to-day adjustment, cost ratio $w$, is relatively consistent across scenarios. Profit is increased when the operator cost is weighted less than costs to the agent given the cost and revenue framework used in this simulation. It is interesting that, for Service Designs 1 and 2, $w$ is higher when en-route transfers are used, suggesting that the transfer capability reduces the travel time and wait time while the day-to-day adjustment drives the reduction in distance travelled. This is consistent with the results of the example in Section 2.4, where en-route transfers used without day-to-day adjustment produced lower agent costs but did not lower overall distance travelled. For the third service design, $w$ is lower when en-route transfers are implemented, indicating that the operator learned to favor more passenger-friendly route design to boost demand. This service design is unique in that agents are consolidated into a convoy for much of the trip, so the marginal increase in vehicle miles travelled for each passenger is minimal, making additional trips largely profitable.

## 5.2 Results for Service Design 1: Urban door-to-door service

The results for Service Design 1 for both cases are presented in Table 8. In Case 1 (S1C1) an average of 642 daily transfers occurred, with an average of 1.09 agents transferred in each. 5.2%

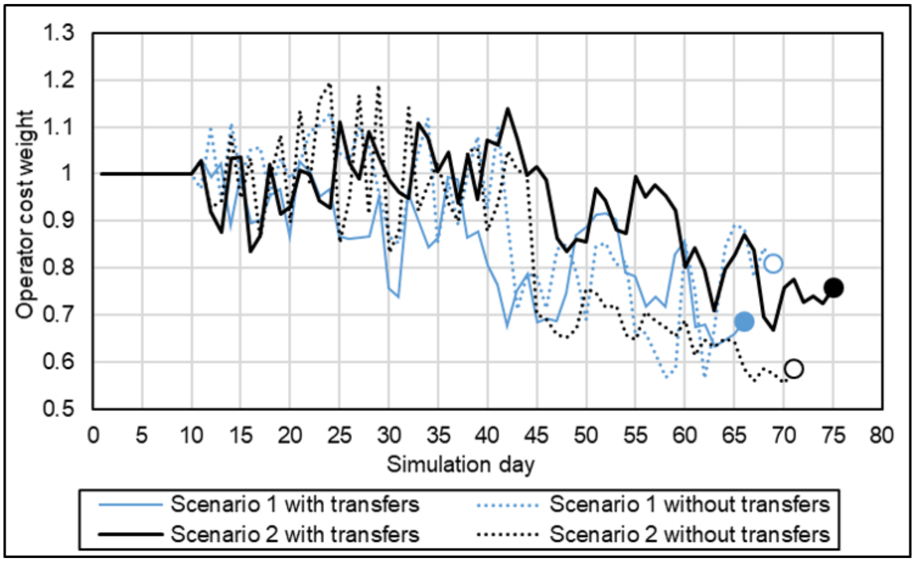


Figure 12: Convergence of cost ratio $w$ under different conditions

of the agents choosing the MAV mode were involved in at least one transfer. The low frequency of transfers can be explained by the scattered nature of the demand profile within the urban area, making it unlikely that two vehicles would be passing near each other with a similar bearing. In Case 2 (S1C2), there were fewer transfers due to a reduction in daily trips compared to Case 1 (S1C1) from 13,344 to 12,096 or -9.4%. Case 2 saw an average of 1.22 agents involved in each transfer.

The results suggest that en-route transfers offer a significant advantage to the mobility operator. Average wait times and travel times are reduced significantly (-10.6% and -9.2%, respectively), and the average profit is increased by more than 20% for Case 1. The improved wait and travel times generated a higher utility (+7.5%) for passengers of the MAV mode, increasing the demand for the service (+2.7%). Each transfer reduced the distance traveled by nearly 6 miles. This is most likely because requests bound for locations on the periphery of the service area can be combined, reducing the number of vehicles that must travel that distance and then travel empty to the next pickup location.

The simulation results are highly sensitive to the changing transfer disutility penalty. The higher penalty in Case 2 causes a significant reduction in average daily users (-6.9%), which in turn lowers the daily profit (from \$161,271 to \$151,911, or -5.8%). This result suggests that MAV research would benefit from further investigation into the perceived disutility of a within-vehicle transfer. A transfer cost could be added to the routing algorithm cost function to avoid inserting en-route transfers with marginal system cost reductions if the transfer disutility is significant.

## 5.3 Results for Service Design 2: First/last mile service

The last mile service (Service Design 2) saw considerably more transfers per user than Service Design 1. In Case 1 (S2C1), an average of 305 daily transfers occurred with an average of 1.95 agents transferred in each. 29% of all MAV passengers were transferred between vehicles. The common origin and destination in the last mile problem makes transfers more advantageous and common. In Case 2 (S2C2), the number of transfers decreased to an average of 283 per day, although

Table 8: Service Design 1 simulation results

(a) Service Design 1, Case 1 (S1C1)

| Service type | Avg. wait time (min) | Avg. travel time (min) | Avg. daily MAV trips | Total VMT | Avg. profit ($) | $\overline{E}_r$ / agent |
|---|---|---|---|---|---|---|
| Without en-route transfers | 5.84 | 18.10 | 12,987 | 28,821 | 134,240 | 3.19 |
| With en-route transfers | 5.22 | 16.43 | 13,344 | 25,003 | 161,271 | 3.43 |
| % change from S1C1 without transfers | -10.6% | -9.2% | +2.7% | -13.2% | +20.6% | +7.5% |

(b) Service Design 1, Case 2 (S1C2)

| Service type | Avg. wait time (min) | Avg. travel time (min) | Avg. daily MAV trips | Total VMT | Avg. profit ($) | $\overline{E}_r$ / agent |
|---|---|---|---|---|---|---|
| With en-route transfers | 4.72 | 15.93 | 12,096 | 24,581 | 151,911 | 3.37 |
| % change from S1C1 without transfers | -19.2% | -12.0% | -6.9% | -14.7% | +13.2% | +5.6% |

the average number of people transferred in each increased to 2.09.

En-route transfers lowered the average wait time (-11.0%) and travel distance (-8.4%) considerably for this service design, as shown in Table 9, although the average travel time is similar with and without transfers (-1.2%). This could be attributed to polar demand profile where each route starts or ends at the transit hub. En-route transfers allow the operator to consolidate passengers into a small number of vehicles travelling to the transit station, freeing up the other vehicles to make additional pickups. The number of daily riders (+0.9%), profit (+10.1%) and consumer welfare per user (+1.8%) were all increased when en-route transfers were used. Each transfer had less of an impact on the profit and the distance traveled in Service Design 2 when compared to Service Design 1. The structure of the demand led to vehicles making trips to and from the transit hub. Transfers often occurred very close to the transit hub at the end of the trip, limiting their benefit.

Much like Service Design 1, the number of daily users was found to be fairly sensitive to changes in the transfer penalty term. Increasing the penalty from 1 minute to 5 minutes reduced the average daily riders from 2,051 (S2C1) to 1,824 (S2C2), a decline of more than 11%, therefore reducing profit from $43,076 to $39,264 (-8.8%). The remaining passengers experienced better service with reduced average wait time (12.11 min to 11.19 min, or -7.6%) and average travel time (34.23 min to 31.30 min, or -8.8%), but the higher disutility for en-route transfers meant users who had been involved in a transfer were much less likely to use the service in the future. Trips declined from 2,051 to 1,824 when the transfer penalty was increased to 5 minutes of travel time for Case 2 (S2C2). The profit is approximately the same as the previous case with no transfers (+0.4%).

Table 9: Service Design 2 simulation results

(a) Service Design 2, Case 1 (S2C1)

| Service type | Avg. wait time (min) | Avg. travel time (min) | Avg. daily MAV trips | Total VMT | Avg. profit ($) | $\overline{E}_r$ / agent |
|---|---|---|---|---|---|---|
| Without en-route transfers | 13.63 | 34.72 | 2,033 | 11,368 | 39,108 | 7.59 |
| With en-route transfers | 12.11 | 34.23 | 2,051 | 10,413 | 43,076 | 7.73 |
| % change from S2C1 without transfers | -11.0% | -1.2% | +0.9% | -8.4% | +10.1% | +1.8% |

(b) Service Design 2, Case 2 (S2C2)

| Service type | Avg. wait time (min) | Avg. travel time (min) | Avg. daily MAV trips | Total VMT | Avg. profit ($) | $\overline{E}_r$ / agent |
|---|---|---|---|---|---|---|
| With en-route transfers | 11.19 | 31.30 | 1,824 | 9,430 | 39,264 | 7.61 |
| % change from S2C1 without transfers | -17.9% | -9.9% | -10.3% | -17.0% | +0.4% | +0.3% |

## 5.4 Results for Service Design 3: Hub-and-spoke service

Table 10 presents the results for Service Design 3. This service design involves a transfer for each agent when they reach the designated depot in their departure city, unless they arrive at the depot when no other vehicles are present. Additional transfers while travelling to and from the depots also occur. In Case 1 (S3C1), an average of 24% of agents were transferred in an additional en-route transfer. An average of 2.12 agents were transferred between vehicles during each transfer. These results are similar to Service Design 2, which is not surprising given that the demand profile before and after the convoy portion of the route is identical to the last-mile service design. In Case 2 (S3C3), the number of transfers decreased to an average of 95 per day, although the average number of people transferred in each increased to 2.31.

The service types with and without en-route transfers have very different results. Because the MAV service consolidates agents into a convoy, the travel distance and wait time are reduced by 17.9% and 19.2%, respectively, when en-route transfers are used. This impact is independent of the additional transfers between the depot and the trip ends which also contribute to the savings. Once again, the increase in the transfer disutility function for Case 2 (S3C2) reduces the average daily riders (from 1,112 to 1,003 or -9.8%), profit (from $41,656 to $37,224 or -10.6%) and consumer welfare (from 10.23 to 9.98 or -2.4%).

Table 10: Service Design 3 simulation results

(a) Service Design 3, Case 1 (S3C1)

| Service type | Avg. wait time (min) | Avg. travel time (min) | Avg. daily MAV trips | Total VMT | Avg. profit ($) | $\overline{E}_r$ / agent |
|---|---|---|---|---|---|---|
| Without en-route transfers | 14.52 | 46.37 | 1,024 | 9,456 | 36,651 | 9.83 |
| With en-route transfers | 11.92 | 50.54 | 1,112 | 7,637 | 41,656 | 10.23 |
| % change from S3C1 without transfers | -17.9% | +9.0% | +8.6% | -19.2% | +13.6% | +4.1% |

(b) Service Design 3, Case 2 (S3C2)

| Service type | Avg. wait time (min) | Avg. travel time (min) | Avg. daily MAV trips | Total VMT | Avg. profit ($) | $\overline{E}_r$ / agent |
|---|---|---|---|---|---|---|
| With en-route transfers | 10.10 | 52.85 | 1,003 | 6,798 | 37,224 | 9.98 |
| % change from S3C1 without transfers | -30.4% | +14.0% | -2.1% | -28.1% | +1.6% | +1.5% |

## 5.5 Discussion

This computational experiment produced several high-level empirical findings which are summarized as follows.

- Operation of Service Design 1 offers the highest profit by far, even without en-route transfers, compared to the other service designs with en-route transfers. This implies that new last mile service deployment, regardless of having MAV technology, stands to benefit the private operator most by operating under Service Design 1.
- On the other hand, Service Design 3 offers the highest consumer welfare compared to the other service designs and cases, even without en-route transfers. New last mile service deployment should consider operating a hub-and-spoke commuter service if they seek to maximize public welfare as a public agency, regardless of the availability of MAV technology.
- If a future MAV provider is looking to "convert" an existing last mile operation, they should investigate locations with a Service Design 1 type of service because that service stands to benefit the most from the implementation of MAVs with en-route transfer technology. En-route transfers provide the greatest increase in profit and consumer welfare within an urban door-to-door service.

The results of this simulation suggest that en-route transfers can improve consumer welfare and operator profits in all three tested service designs. Assuming the MAV technology with en-route transfers reaches maturation, implementation of a commuter MAV service connecting the two cities in the empirical study would have the greatest impact on vehicle miles travelled and thereby congestion. This type of service can be implemented in stages, with transfers initially occurring only at designated hubs during testing of the full-speed transfer mechanism. The results for an urban door-to-door service are also promising, demonstrating that en-route transfers could mitigate the congestion that is typically caused by ride-hailing services (Clewlow and Mishra, 2017). The day-to-day adjustment framework could inform operating strategies for other emerging mobility systems. It is valuable for deployment in new markets where demand elasticity is difficult to estimate, or for emerging transportation modes such as autonomous ridehailing services or urban air mobility where the latent demand function is unknown. By operationalizing the user/operator cost parameter, the operator is able to adapt to revealed demand.

It should be noted that the results of this simulation represent a lower bound due to the sub-optimal insertion heuristic used for routing. An routing algorithm that generates globally optimal routes and transfers could produce further performance improvements for en-route transfers.

# 6 Conclusion

Several contributions are made to the evaluation and operation of future mobility services. A new procedure is developed for mobility operators to learn transportation market dynamics through routing algorithm adjustment. The operator/user cost weighting in a routing algorithm is parameterized and varied endogenously resulting in profit maximization regardless of initial setting. In uncertain markets with limited data, this day-to-day operational policy adjustment allows the operator to learn the latent demand function of the market. This presents a new set of operating service designs for mobility services entering a new region with a new product. MAV operating service designs are evaluated using the day-to-day learning framework to understand the impact on traveler and operator welfare relative to existing on-demand mobility services.

The results of the computational experiments suggest that MAVs with en-route transfer capability could be deployed to improve service and increase profits in a mobility services market. En-route transfers were found to increase the traveler and operator welfare (profit) most in an urban door-to-door service. Congestion would be reduced most in a suburb-city hub-and-spoke service. Transfers are more likely to be used in a last-mile or commuter service than in an urban door-to-door service, although the impact of each transfer is less.

The urban door-to-door service is most similar to existing ridehailing operations. The results presented in this study suggest that both operators and customers would benefit from the development of MAVs with en-route transfer capability. The demonstrated potential for congestion alleviation and opportunity to encourage a transit mode shift by addressing the last-mile challenge should also encourage transportation agencies and transit operators to consider this emerging technology among their long term plans. There is no existing regulatory framework or testing procedure for passenger vehicles that couple and uncouple while in motion to exchange riders. Given the potential societal benefits, policy makers should develop comprehensive testing regimens that can be used to verify the safety of these vehicles before they are permitted on public roads.

The size of the agent pool, vehicle fleet and geographical area used in these experiments is limited by computation time. MAV service operators with en-route transfers would benefit from economies of scale as there are more options for transfers when there are more vehicles and agents. Future studies could operationalize the simulation on a more efficient computational setting such as C++, or an agent-based traffic simulator that includes road capacity effects such as MATSIM. These simulation tools would allow for evaluation of the impact of a MAV service with en-route transfers in an urban area like New York City where ride-hailing operators serve hundreds of thousands of daily trips.

While more advanced heuristics have been developed in recent years (Laporte et al., 2014; Wassan and Nagy, 2014; Vidal et al., 2019), we chose to use a relatively straightforward routing algorithm in order to focus the paper on the main contributions: evaluating operating policies for MAV services and day-to-day adjustment of operating parameters to learn an unknown latent demand function. As a result, our findings may underestimate the potential efficiency of MAV services. Future extensions of this work could include more advanced routing algorithms to provide an estimate for further performance improvements.

## Acknowledgements

The authors were partially supported by the C2SMART University Transportation Center.

## References

Balcombe, R., Mackett, R., Paulley, N., Preston, J., Shires, J., Titheridge, H., Wardman, M., and White, P. (2004). The demand for public transport: a practical guide. *Transportation Research Laboratory*.

Berbeglia, G., Cordeau, J.-F., and Laporte, G. (2010). Dynamic pickup and delivery problems. *European Journal of Operational Research*, 202(1):8–15.

Bhat, C. R. and Sardesai, R. (2006). The impact of stop-making and travel time reliability on commute mode choice. *Transportation Research Part B: Methodological*, 40(9):709–730.

Bouros, P., Sacharidis, D., Dalamagas, T., and Sellis, T. (2011). Dynamic pickup and delivery with transfers. In *International Symposium on Spatial and Temporal Databases*, pages 112–129. Springer.

Braekers, K., Ramaekers, K., and Van Nieuwenhuyse, I. (2016). The vehicle routing problem: State of the art classification and review. *Computers & Industrial Engineering*, 99:300–313.

Chang, S. K. and Schonfeld, P. M. (1991a). Integration of fixed-and flexible-route bus systems. *Transportation Research Record*, 1308(1).

Chang, S. K. and Schonfeld, P. M. (1991b). Optimization models for comparing conventional and subscription bus feeder services. *Transportation Science*, 25(4):281–298.

Chen, Z., Li, X., and Zhou, X. (2019a). Operational design for shuttle systems with modular vehicles under oversaturated traffic: Continuous modeling method. *Transportation Research Part B: Methodological*.

Chen, Z., Li, X., and Zhou, X. (2019b). Operational design for shuttle systems with modular vehicles under oversaturated traffic: Discrete modeling method. *Transportation Research Part B: Methodological*, 122:1 – 19.

Chong, Z., Qin, B., Bandyopadhyay, T., Wongpiromsarn, T., Rebsamen, B., Dai, P., Rankin, E. S., and Ang, M. H. (2013). Autonomy for mobility on demand. In *Intelligent Autonomous Systems 12*, pages 671–682. Springer.

Chow, J. Y. J. (2018). *Informed Urban Transport Systems: Classic and Emerging Mobility Methods toward Smart Cities*. Elsevier.

Clewlow, R. R. and Mishra, G. S. (2017). Disruptive transportation: The adoption, utilization, and impacts of ride-hailing in the united states. Technical report, UC Davis Institute of Transportation Studies.

Cordeau, J.-F., Gendreau, M., Hertz, A., Laporte, G., and Sormany, J.-S. (2005). New heuristics for the vehicle routing problem. In *Logistics Systems: Design and Optimization*, pages 279–297. Springer.

Cortés, C. E., Matamala, M., and Contardo, C. (2010). The pickup and delivery problem with transfers: Formulation and a branch-and-cut solution method. *European Journal of Operational Research*, 200(3):711–724.

Cortés, C. E., Pagès, L., and Jayakrishnan, R. (2005). Microsimulation of flexible transit system designs in realistic urban networks. *Transportation Research Record*, 1923(1):153–163.

Deleplanque, S. and Quilliot, A. (2013). Transfers in the on-demand transportation: the darpt dial-a-ride problem with transfers allowed. In *Multidisciplinary International Scheduling Conference: Theory and Applications (MISTA)*, pages 185–205.

Djavadian, S. and Chow, J. Y. J. (2017a). An agent-based day-to-day adjustment process for modeling ‘mobility as a service’with a two-sided flexible transport market. *Transportation Research Part B: Methodological*, 104:36–57.

Djavadian, S. and Chow, J. Y. J. (2017b). Agent-based day-to-day adjustment process to evaluate dynamic flexible transport service policies. *Transportmetrica B: Transport Dynamics*, 5(3):281–306.

Dubai Roads and Transport Authority (2018). Dubai Pulse, Bus Ridership Data. www.dubaipulse.gov.ae/data/allresource/rta-bus, accessed 2018-04-27.

Fahmy, T. (2018). Dubai tests autonomous pods in drive for smart city. Reuters. https://www.reuters.com/article/us-emirates-transportation-autonomous/dubai-tests-autonomous-pods-in-drive-for-smart-city-idUSKCN1GD5G6, accessed March 1, 2018.

Garcia-Martinez, A., Cascajo, R., Jara-Diaz, S. R., Chowdhury, S., and Monzon, A. (2018). Transfer penalties in multimodal public transport networks. *Transportation Research Part A: Policy and Practice*, 114:52–66.

Google Maps (2018). Sharjah - United Arab Emirates to Dubai - United Arab Emirates - Google Maps. https://maps.google.com/, accessed November 28, 2018.

Guastaroba, G., Speranza, M. G., and Vigo, D. (2016). Intermediate facilities in freight transportation planning: a survey. *Transportation Science*, 50(3):763–789.

Guo, Q.-W., Chow, J. Y. J., and Schonfeld, P. (2017). Stochastic dynamic switching in fixed and flexible transit services as market entry-exit real options. *Transportation Research Procedia*, 23:380–399.

Hennessey, M. (2016). Milton's GO Connect Service pilot a success. *Inside Halton*. https://www.insidehalton.com/news-story/6511299-milton-s-go-connect-service-pilot-a-success/, accessed April 23, 2018.

Horn, M. E. T. (2002). Fleet scheduling and dispatching for demand-responsive passenger services. *Transportation Research Part C: Emerging Technologies*, 10(1):35–63.

Hu, T.-Y. and Mahmassani, H. S. (1997). Day-to-day evolution of network flows under real-time information and reactive signal control. *Transportation Research Part C: Emerging Technologies*, 5(1):51–69.

Hussein, A. (2017). Impact of traffic congestion on drivers commuting between sharjah and dubai emirates. In *Health and Environment Conference*, page 95.

Iida, Y., Akiyama, T., and Uchida, T. (1992). Experimental analysis of dynamic route choice behavior. *Transportation Research Part B: Methodological*, 26(1):17–32.

Irnich, S., Toth, P., and Vigo, D. (2014). Chapter 1: The family of vehicle routing problems. In *Vehicle Routing: Problems, Methods, and Applications, Second Edition*, pages 1–33. SIAM.

Jaw, J.-J., Odoni, A. R., Psaraftis, H. N., and Wilson, N. H. (1986). A heuristic algorithm for the multi-vehicle advance request dial-a-ride problem with time windows. *Transportation Research Part B: Methodological*, 20(3):243–257.

Jha, M., Madanat, S., and Peeta, S. (1998). Perception updating and day-to-day travel choice dynamics in traffic networks with information provision. *Transportation Research Part C: Emerging Technologies*, 6(3):189–212.

Jiao, J., Miro, J., and McGrath, N. (2017). What Public Transit Can Learn From Uber and Lyft. *CityLab*. https://www.citylab.com/transportation/2017/11/what-public-transit-can-learn-from-uber-and-lyft/544637/, accessed April 23, 2018.

Knoblauch, R. L., Pietrucha, M. T., and Nitzburg, M. (1996). Field studies of pedestrian walking speed and start-up time. *Transportation Research Record*, 1538(1):27–38.

Krygsman, S., Dijst, M., and Arentze, T. (2004). Multimodal public transport: an analysis of travel time elements and the interconnectivity ratio. *Transport Policy*, 11(3):265–275.

Laporte, G., Ropke, S., and Vidal, T. (2014). Chapter 4: Heuristics for the vehicle routing problem. In *Vehicle Routing: Problems, Methods, and Applications, Second Edition*, pages 87–116. SIAM.

Lesh, M. C. (2013). Innovative concepts in first-last mile connections to public transportation. In *Urban Public Transportation Systems 2013*, pages 63–74. American Society of Civil Engineers.

Litman, T. (2009). Transportation cost and benefit analysis. *Victoria Transport Policy Institute*, 31.

Litman, T. (2010). *Evaluating public transit benefits and costs: Best Practices Guidebook*. Victoria Transport Policy Institute.

Litman, T. (2017). *Autonomous vehicle implementation predictions*. Victoria Transport Policy Institute.

Madsen, O. B. G., Ravn, H. F., and Rygaard, J. M. (1995). A heuristic algorithm for a dial-a-ride problem with time windows, multiple capacities, and multiple objectives. *Annals of Operations Research*, 60(1):193–208.

Marshall, A. (2017). How a failed experiment could still be the future of public transit. *Wired*. https://www.wired.com/2017/03/failed-experiment-still-future-public-transit/, accessed April 23, 2018.

Masson, R., Lehuédé, F., and Péton, O. (2013). An adaptive large neighborhood search for the pickup and delivery problem with transfers. *Transportation Science*, 47(3):344–355.

Melnick, L. (2018). A dynmaic ridesharing routing algorithm with en-route transfers. Unpublished Master's Thesis.

Mulley, C. and Nelson, J. D. (2009). Flexible transport services: A new market opportunity for public transport. *Research in Transportation Economics*, 25(1):39–45.

Next Future Transportation (2020). Home: Next Future Transportation. https://www.next-future-mobility.com/home, accessed May 16, 2020.

Pacheco, M., Sharif Azadeh, S., and Bierlaire, M. (2016). A new mathematical representation of demand using choice-based optimization method. In *16th Swiss Transport Research Conference*.

Rais, A., Alvelos, F., and Carvalho, M. S. (2014). New mixed integer-programming model for the pickup-and-delivery problem with transshipment. *European Journal of Operational Research*, 235(3):530–539.

Sayarshad, H. R. and Chow, J. Y. J. (2015). A scalable non-myopic dynamic dial-a-ride and pricing problem. *Transportation Research Part B: Methodological*, 81(4):539–554.

Sisson, P. (2018). Microtransit: How cities are, and aren't, adapting transit technology. *Curbed*. https://www.curbed.com/2018/1/9/16871474/microtransit-mass-transit-uber-lyft, accessed April 23, 2018.

Small, K. A. (1982). The scheduling of consumer activities: work trips. *The American Economic Review*, 72(3):467–479.

Sulopuisto, O. (2016). Why Helsinki's innovative on-demand bus service failed. *Citiscope*. http://archive.citiscope.org/story/2016/why-helsinkis-innovative-demand-bus-service-failed, accessed April 23, 2018.

The National UAE (2018). Autonomous pods the future of city driving. https://www.thenational.ae/business/technology/autonomous-pods-the-future-of-city-driving-1.730283, accessed July 16, 2018.

Train, K. E. (2009). *Discrete choice methods with simulation*. Cambridge University Press.

van Engelen, M., Cats, O., Post, H., and Aardal, K. (2018). Enhancing flexible transport services with demand-anticipatory insertion heuristics. *Transportation Research Part E: Logistics and Transportation Review*, 110:110–121.

Vidal, T., Laporte, G., and Matl, P. (2019). A concise guide to existing and emerging vehicle routing problem variants. *European Journal of Operational Research*.

Wang, H. (2017). Routing and scheduling for a last-mile transportation system. *Transportation Science*, 53(1):131–147.

Wang, H. and Odoni, A. (2014). Approximating the performance of a "last mile" transportation system. *Transportation Science*, 50(2):659–675.

Wardman, M., Hine, J., and Stradling, S. (2001). Interchange and travel choice-volumes 1 and 2. *Transport Research Series*.

Wassan, N. A. and Nagy, G. (2014). Vehicle routing problem with deliveries and pickups: modelling issues and meta-heuristics solution approaches. *International Journal of Transportation*, 2(1):95–110.

Wong, K. I. and Bell, M. G. H. (2006). Solution of the dial-a-ride problem with multi-dimensional capacity constraints. *International Transactions in Operational Research*, 13(3):195–208.

Yang, F. and Zhang, D. (2009). Day-to-day stationary link flow pattern. *Transportation Research Part B: Methodological*, 43(1):119–126.

Yap, M. D., Correia, G., and Van Arem, B. (2016). Preferences of travellers for using automated vehicles as last mile public transport of multimodal train trips. *Transportation Research Part A: Policy and Practice*, 94:1–16.