



# MIT Open Access Articles

## *Toward a Systematic Approach to the Design and Evaluation of Automated Mobility-on-Demand Systems: A Case Study in Singapore*

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

<b>Citation</b>	Spieser, Kevin, Kyle Ballantyne, Treleaven, Rick Zhang, Emilio Frazzoli, Daniel Morton, Marco Pavone. "Toward a Systematic Approach to the Design and Evaluation of Automated Mobility-on-Demand Systems A Case Study in Singapore." in Gereon Meyer, Sven Beiker (editors). Road Vehicle Automation, (Lecture Notes in Mobility), Springer, 2014.
<b>As Published</b>	
<b>Publisher</b>	Springer
<b>Version</b>	Original manuscript
<b>Accessed</b>	Tue Oct 17 07:33:15 EDT 2017
<b>Citable Link</b>	<a href="http://hdl.handle.net/1721.1/82904">http://hdl.handle.net/1721.1/82904</a>
<b>Terms of Use</b>	Creative Commons Attribution-Noncommercial-Share Alike 3.0
<b>Detailed Terms</b>	<a href="http://creativecommons.org/licenses/by-nc-sa/3.0/">http://creativecommons.org/licenses/by-nc-sa/3.0/</a>

# Toward a Systematic Approach to the Design and Evaluation of Automated Mobility-on-Demand Systems

## A Case Study in Singapore

Kevin Spieser, Kyle Treleaven, Rick Zhang, Emilio Frazzoli, Daniel Morton, and Marco Pavone

**Abstract** The objective of this work is to provide analytical guidelines and financial justification for the design of shared-vehicle mobility-on-demand systems. Specifically, we consider the fundamental issue of determining the appropriate number of vehicles to field in the fleet, and estimate the financial benefits of several models of car sharing. As a case study, we consider replacing all modes of personal transportation in a city such as Singapore with a fleet of shared automated vehicles, able to drive themselves, e.g., to move to a customer's location. Using actual transportation data, our analysis suggests a shared-vehicle mobility solution can meet the personal mobility needs of the entire population with a fleet whose size is approximately 1/3 of the total number of passenger vehicles currently in operation.

## 1 Introduction

In light of ongoing urbanization trends, cities face the challenge of maintaining the services and infrastructure necessary to keep pace with the transportation demands of a growing population. When the returns from investment in existing technologies, e.g., road expansion, added bus service, new subway lines, etc., begin to diminish, it is appropriate, perhaps even necessary, to consider new and potentially transformative transportation solutions. A responsible approach to address the merits of a proposed solution is to conduct a systematic analysis of its key operational compo-

---

Kevin Spieser, Kyle Treleaven, and Emilio Frazzoli  
Massachusetts Institute of Technology, Cambridge, MA, e-mail: {kspieser, ktreleav, frazzoli}@mit.edu

Rick Zhang and Marco Pavone  
Stanford University, Palo Alto, CA, e-mail: {rickz, pavone}@stanford.edu

Daniel Morton  
Singapore-MIT Alliance for Research and Technology, Singapore, e-mail: daniel.morton@smart.mit.edu

nents, thereby providing an informed foundation from which to gauge feasibility. It is in this spirit that this paper examines a new solution to personal mobility; namely, that of replacing all modes of personal transport in a city with a fleet of shared autonomous vehicles, i.e., vehicle that are able to drive themselves in traffic, to safely and reliably pick up passengers and deliver them at their intended destination.

Research on autonomous vehicles is currently very active [1, 2]. Proponents of this technology typically point out as the main benefits (i) increased safety, as the automation reduces the effects of human errors, well known to be the leading cause of traffic accidents, (ii) increased convenience and productivity, as humans are absolved from the more burdensome aspects of driving, (iii) increased traffic efficiency and lower congestion, as automated vehicles can precisely monitor one another's position and coordinate their motion to an extent impossible for human drivers, and (iv) reduced environmental impact, as velocity profiles can be carefully tuned to minimize emissions and noise.

For the sake of this article, though, we will concentrate on yet another potential major benefit, i.e., (v) autonomous vehicles as an enabling technology for widespread car sharing. It is well known that most private cars are used less than 10% of the time [3], so car sharing is a clear path towards sustainability—especially if cars do not need a driver to move. Car-sharing services are growing worldwide, but typically do not offer one-way rental options, or if they do they often suffer from limited car availability. If shared cars were able to return to a parking or charging station, or drive to pick up the next customer by themselves, sharing would indeed provide a similar level of convenience as private cars, while providing the sustainability of public transport. Financially, car sharing distributes the cost of purchasing, maintaining, and insuring vehicles across a large user-base, leveraging economies of scale to reduce the cost of personal mobility.

While automated vehicle technology continues to surge forward, less attention has been devoted to the logistics of effectively managing a fleet of potentially thousands of such vehicles. Among those works that do exist, many are of a conceptual nature. Although they raise a number of interesting ideas and suggest novel operational paradigms, they frequently lack the rigor necessary to justify the feasibility of their claims. Those works that do take a design-oriented approach frequently rely heavily on simulations and unrealistic transportation demand models. Consequently, these techniques prove difficult to generalize and neglect salient features that have a fundamental impact on key performance metrics. Recognizing these shortcomings, this work provides some preliminary insights toward a systematic approach to size a fleet of shared vehicles given actual mobility patterns. These results are then applied, using actual transportation and traffic statistics, to the problem of fleet sizing for a shared-mobility system in Singapore. Financial estimates for a variety of car-sharing systems are provided to assess their financial feasibility.

While the main motivation of this paper is provided by automated shared-vehicle systems, the results are applicable to more general cases, including, e.g., a fleet of shared vehicles, each with a human driver, coordinating with other drivers in such a way to maximize the quality of service provided to the customers (as opposed to their own interests, as is the case with current models of taxi services).

## 2 Shared-mobility systems

The efficiency gains that shared-vehicle systems can, in theory, offer to both the individual user and society as a whole have been well documented [3]. For select cities, including Singapore, lists of these advantages have even been specifically compiled [4]. To date, the majority of car-sharing programs feature a roundtrip vehicle rental model. In these systems, vehicles must be returned to the same station they were rented from. Zipcar's current rental service, for example, is based on this approach. Eager to capitalize on emerging markets and better serve existing ones, considerable effort has been devoted to characterizing the demand for shared-vehicle mobility in different markets [5]. Naturally, demographic factors, e.g., [5] and [6], and geographic considerations, e.g., [7] and [8], affect demand. However, researchers have also been quick to note the important role quality of service plays in establishing a clientele [9]. For example, by fielding larger vehicle fleets, companies make it easier for patrons to rent a vehicle from a nearby station. This, in turn, draws new members to the program [10]. As this effect takes hold, yet more vehicles are required to maintain the high level of service that initially attracted users.

Noting the limitations in the roundtrip rental model, one-way car-sharing services, such as car2go, have emerged. These services offer the added convenience of being able to return a vehicle to any one of multiple stations throughout the city [9]. However, left unchecked, asymmetries in travel patterns would, in general, create a surplus of vehicles at select stations, while leaving other stations underserved. Rebalancing mechanisms are therefore required to realign the supply of vehicles with the demand. Moreover, how effectively vehicles are shuffled between stations strongly affects vehicle availability, which, in turn, impacts demand for the service.

Simulation-based approaches have been used to infer the viability of various rebalancing schemes and, in turn, gauge consumer demand for one-way car-sharing, e.g., [11–15]. Initial findings suggest that one-way services are ideally suited for densely-populated urban centers. Unfortunately, the lack of insight in the presence of a large number of relevant but uncertain parameters has been noted as a limitation of predominantly simulation-driven methods [9]. A more theoretical direction has been pursued in a number of recent works aimed at understanding optimal rebalancing strategies, and the fundamental limits of stability and performance in car sharing systems, either considering automated or human-driven vehicles [16–18].

Each of the works referenced thus far focus on one or more important aspects of car-sharing systems. However, none of them tackle the more holistic problem of rigorously sizing the fleet for an Automated Mobility-on-Demand (AMoD) service to meet the transportation demand of an actual city. In this regard, [19], similar to this work, is noteworthy; it takes a design-oriented approach to fleet-sizing for hypothetical shared-vehicle systems at three sites across the United States (US). The approach used therein is heavily simulation-based and the spatial component of the demand model was not derived from real-world data. In contrast, we provide guidelines to size an AMoD system for Singapore based on measurable travel characteristics and, to the extent possible given the current technical literature, rigorous theoretical arguments.

### 3 Fleet Sizing for Automated Mobility-on-Demand Systems

In this section, we present the two key technical problems considered in the paper. The problem statements are straightforward, but they are posed in terms of formal mathematical models of the physical environment and transportation demand.

#### 3.1 Problem Formulation

Consider a compact planar region (environment)  $Q \subset \mathbb{R}^2$ . The  $i$ -th transportation demand in a random sequence  $\{(t_i, o_i, d_i)\}_{i \in \mathbb{N}}$  poses the requirement to travel from an origin point  $o_i \in Q$  to a destination point  $d_i \in Q$ . The demand, however, is only *revealed* after time  $t_i$ . Trip requests are to be serviced by vehicles that may transport at most one demand at a time. The average speed of the vehicles  $v$  is assumed to be periodically time-varying. A significant challenge to trip scheduling stems from uncertainty in the travel demand, which we will model probabilistically. Transportation demands arrive according to a non-stationary, (separably) spatio-temporal Poisson process  $(\lambda, f)$ , where  $\lambda$  is the arrival rate function, and  $f$  is a probability distribution function called the *demand distribution*; both are periodically time-varying. The individual trip data are all statistically independent, and the  $i$ -th O-D pair  $(o_i, d_i)$  is conditionally distributed according to  $f(\cdot; t_i)$ . The expected number of demands revealed within a time interval  $[t_1, t_2]$ , and with  $o_i \in Q_1$  and  $d_i \in Q_2$ , for any time  $t$  and regions  $Q_1, Q_2 \subset Q$ , is  $\int_{t_1}^{t_2} \int_{Q_2} \int_{Q_1} \lambda(t) f(p, q, t) dp dq dt$ . The problems of interest are the following:

1. **Minimum fleet sizing:** What is the minimum number of vehicles,  $m_{min}$ , necessary to keep the number of outstanding demands uniformly bounded?
2. **Performance-driven fleet sizing:** How many vehicles,  $m_{per}$ , should be used to ensure that the quality of the service provided to the customer (e.g., vehicle availability, or waiting time) is no less than a given threshold?

The second problem acknowledges an intuitive trade-off between the fleet size and the user experience (beyond the bare minimum).

In the following, we provide techniques to address the fleet sizing problems in the case of Singapore. Interested readers may refer to, e.g., [17] and [18] for similar problems defined with regard to a system of stations embedded in a road network, instead of a compact region in the Euclidean plane.

#### 3.2 Minimum fleet sizing

As mentioned in Section 3.1, the problems of interest pertain to fleet sizing for a pickup-and-delivery system. A detailed theoretical treatment stressing the stochastic and queue-theoretic nature of the problem can be found in [16].

A fleet of  $m$  vehicles is said to *stabilize* the workload if there exists a service (routing) policy  $\pi(m)$  that ensures the expected number of outstanding demands is uniformly bounded. Stability therefore implies that the fleet, as a whole, must be able to cover distance at least as quickly on average as the rate at which service distance accumulates.

Given a sequence of points  $p_1, \dots, p_n \in Q$ , let  $\mathcal{D}(p_1, \dots, p_n)$  denote the length of the shortest path through each point in the order specified by the sequence. The average distance that a vehicle must travel in service *per demand* is  $d^{\text{trip}} := \limsup_{i \rightarrow +\infty} \mathbb{E}\{\mathcal{D}(d_{\text{pre}(i)}, o_i, d_i)\}$ , where  $\text{pre}(i)$  is the index of the demand served immediately before the demand  $i$ . If the temporal variation of travel demand is discretized into, e.g., hourly temporal bins, the rate at which work enters the system is  $\sum_k (\lambda_k \cdot d_k^{\text{trip}})$ , where the subscript  $k$  indicates the bin index, and  $d_k^{\text{trip}} := \limsup_{i \rightarrow +\infty} \mathbb{E}\{\mathcal{D}(d_{\text{pre}(i)}, o_i, d_i) : t_i \text{ in bin } k\}$ . A fleet of  $m$  vehicles, each capable of traveling at average speed  $v_k$  during the  $k$ -th bin, is able to cover distance at a daily rate of  $m \sum_k v_k$ . Therefore, a necessary condition for system stability is

$$m > \sum_k \left( \lambda_k \cdot d_k^{\text{trip}} \right) / \sum_k v_k. \quad (1)$$

Dropping the subscript  $k$  for simplicity, the average trip length can be decomposed as  $d^{\text{trip}} = d^{\text{OD}} + d^{\text{E}}$ , where  $d^{\text{OD}} = \mathbb{E}[\mathcal{D}(o_i, d_i)]$ , and  $d^{\text{E}}$  depends on the ordering of demands served, and clearly depends on the routing policy  $\pi(m)$ . In [16], rigorous arguments are used to prove that  $d^{\text{E}}$  is bounded below by a computable quantity that depends on the mobility demand, and that the bound is approachable in practice. Letting  $f^o$  and  $f^d$  denote the first and second factors of  $f$ —i.e., the marginal distributions associated with origins and destinations, respectively—then  $d^{\text{E}} \geq \text{EMD}(f^d, f^o)$ , where  $\text{EMD}$  is a function often called the *Earth mover's distance*, and will be defined shortly. Given the above, condition (1) becomes

$$m > \sum_k \left( \lambda_k (d_k^{\text{OD}} + \text{EMD}(f_k^d, f_k^o)) \right) / \sum_k v_k. \quad (2)$$

Formally, the Earth mover's distance  $\text{EMD}(f^o, f^d)$  is a measure of distance between distributions  $f^o$  and  $f^d$ ; in mathematical terms, given the *ground metric*  $(Q, \mathcal{D})$ , the  $\text{EMD}$  is the first Wasserstein distance [20], usually written as

$$\text{EMD}(f_1, f_2) = \inf_{\gamma \in \Gamma(f_1, f_2)} \int_{Q \times Q} \mathcal{D}(p_1, p_2) d\gamma(p_1, p_2), \quad (3)$$

where  $\Gamma(f_1, f_2)$  is the set of all measures with marginals  $f_1$  and  $f_2$  on the first and second factor, respectively. If distributions  $f_1$  and  $f_2$  are imagined as describing two piles each consisting of a unit of “dirt” (i.e., earth), then  $\text{EMD}(f_1, f_2)$  is intuitively the minimum work (dirt  $\times$  distance) required to reshape  $f_1$  into  $f_2$ .

Although some existing works are keen to emphasize the relationship between  $d^{\text{OD}}$  and  $m_{\text{min}}$ , they often fail to recognize the contribution of  $d^{\text{E}} \geq \text{EMD}(f^o, f^d)$ . This is an unfortunate omission as  $\text{EMD}(f^o, f^d)$  represents the minimum distance,

on average, a vehicle must travel to realign itself with an asymmetrical travel demand, and is a fundamental contributor to system workload. It is justifiable to ignore  $\text{EMD}(f^o, f^d)$  only when  $f^o = f^d$ , because in this case  $\text{EMD} = 0$ . However, in most real cases, including Singapore,  $f^o$  and  $f^d$  are different on the time scales over which trips must be completed.

### 3.3 Performance-driven fleet sizing

The analysis in the previous section provides crucial information about absolute minimum fleet sizes to ensure user demand can (in principle) be met. These results help to answer whether a particular fleet size is “large enough.” However, another important question which is left unanswered is how the size of the fleet impacts the user experience, e.g., by decreasing user wait times or by increasing vehicle availability. It may be quite valuable to quantify more accurately such trade-offs.

In this work, we study performance in terms of vehicle availability. To this purpose, we model an autonomous MOD system as a closed queueing network of  $m$  vehicles and  $N$  disjoint regions  $Q_1, Q_2, \dots, Q_N \subset Q$ . Idle vehicles are parked at the median of each region. When a customer arrives in region  $Q_i$ , destined for  $Q_j$ , a free vehicle in  $Q_i$  is sent to pick up and drop off the customer before parking at the median of  $Q_j$ . Customers arrive in each region  $Q_i \subset Q$  according to a Poisson process with rate  $\lambda_i$  and take vehicles to region  $Q_j$  with probability  $p_{ij}$ . To maintain tractability for analytical results, we consider a simple model where if a region is empty of available vehicles, the customer immediately leaves the system (this model is usually referred to as *loss model*, which models well customer impatience). The performance criterion is then the availability of vehicles in each region, or the probability that a customer will be able to book a vehicle in his/her region.

This model of a MOD system can be analyzed as a closed Jackson network with respect to the vehicles. Jackson networks are a special case of a class of queueing networks known as BCMP networks [21], which, remarkably, admit stationary probability distributions in product form. In our case, regions are mapped into single-server nodes, while routes between each pair of regions are mapped into infinite-server nodes (note that we are not modeling congestion effects). Thanks to the product form of the probability distributions, the availability of vehicles in each region can be efficiently computed using mean value analysis (see, e.g., [22]). Previous work such as [22] used this model to generate guidelines for system design and perform profit-based fleet sizing. However, their analysis shows that without rebalancing the majority of regions can only achieve an availability strictly less than one *even* if  $m$  tends to infinity. To take into account the possibility of vehicle rebalancing offered by autonomous vehicles, we model the rebalancing process as an arrival process of “virtual passengers” with Poisson rate  $\psi_i$  and routing probability  $\alpha_{ij}$ , independent of the real passenger arrival process. As with real passengers, the virtual passengers are lost if no vehicles are available in the booking region upon arrival. We can then optimize the availability with respect to  $\psi_i$  and  $\alpha_{ij}$  to achieve a

balanced system so that availability approaches one as  $m$  tends to infinity, for all regions. The approach is similar to that developed in [17], and allows us to determine performance curves in terms of  $m$  and vehicle availability. Also, it provides *baseline* policies that would guide the development of real-time closed-loop vehicle routing policies as in [17].

## 4 Data sources

In order to apply our analytical results to estimate fleet size for a real-world scenario, we chose to consider Singapore as a case study. Singapore is a fitting venue for at least two reasons: First, we have access to a rich collection of data pertaining to the country, from which to gather the statistics that drive our analytics. Second, despite Singapore's sophisticated and well-subsidized public transportation system, the rate of private vehicle ownership, and correspondingly traffic congestion, continues to increase. Given the island's diminutive size and high population density, officials are limited in the extent to which traditional measures, e.g., roadway expansion, can alleviate rising congestion. In this regard, Singapore is a promising candidate for replacing existing modes of land transport with shared vehicles for personal mobility, i.e., an AMoD system. To support the analysis of such scenario, three complementary data sources are used, as described below.

**The Household Interview Travel Survey**— The Household Interview Travel Survey, or simply HITS, is a comprehensive survey conducted periodically by the Land Transport Authority (LTA) for the purpose of gathering an overview of high-level transportation patterns within Singapore [23]. This work employed the 2008 HITS survey in which 10,840 of the then 1,144,400 households in Singapore were selected to participate in the survey. The HITS database, which summarizes the survey, is structured as follows. For each household surveyed, each resident reported specific details of each trip taken on a recent weekday of interest. In general, each trip is comprised of several stages with a new stage introduced each time the participant switched their mode of transport, e.g., transferred from the subway to bus as part of the same trip. For each trip, the resident reported the trip's origin point, destination point, start time, end time, and the mode of transport, e.g., car, bus, subway, etc., used in each substage.

**Singapore Taxi Data**—To gather ground truth traffic characteristics, we rely on a database of taxi records collected over the course of a week in Singapore in 2012. The data chronicles the movement and activities of approximately 60% of all active taxis by recording each vehicle's GPS coordinates, speed, and passenger status, e.g., "passenger-on-board," "vacant," "responding to call," etc. Owing to the high rate at which recordings are taken, approximately every 30 seconds to 1 minute per vehicle, and the large number of taxis contributing to the database (approximately 15,000), the fleet, collectively, serves as a distributed, mobile, embedded traffic sensor which may be queried to provide an estimate of traffic conditions throughout the city.



**Singapore Road Network**— A graph-based representation of Singapore’s road network is used to determine the most efficient routes automated vehicles should take from point to point in Singapore (whether carrying a passenger or moving to fetch one). When the analysis method required simpler distance evaluations, the average ratio of trip length over Euclidean distance was estimated from the taxi data as a factor  $\beta = 1.38$ .

## 5 Sizing an AMoD solution for Singapore

Having acquired both the necessary analytical tools, and the transportation data, we are now able to compute estimates for the AMoD fleet sizing problem in Singapore.

### 5.1 Minimum fleet sizing

This section describes the methodology used to estimate the quantities appearing in (2). Results are summarized in Figure 1.

**Arrival Rate ( $\lambda$ )**— Let  $\lambda_k^{\text{HITS}}$  represent the average rate at which trips in hour  $k \in \{0, 1, \dots, 23\}$  arrive based solely on the HITS survey. The overall arrival rate, in hour  $k$ , is evaluated as  $\lambda_k = \alpha \lambda_k^{\text{HITS}}$ , where  $\alpha = 1,144,400/10,840 \approx 105.57$  is the scaling factor that, inversely, reflects the fraction of the households that took part in the HITS survey. From the HITS data, 56,839 trips were extracted. After eliminating trips for which the GPS coordinates of  $o$ ,  $d$ , or both were unavailable, 56,673 trips remained.

**Average O-D Distance ( $d^{\text{OD}}$ )**— For each O-D pair in the HITS database, we assume the trip takes place on the shortest path (as measured by distance) connecting  $o$  and  $d$ . Shortest path algorithms, e.g., Dijkstra’s algorithm, are computationally efficient, allowing calculations to be run on a detailed roadmap of Singapore. On an hourly basis,  $d_k^{\text{OD}}$  ranges from a minimum of 6.47km to a maximum of 13.31km.

**Mobility Demand Distribution ( $f$ )**— The road network of Singapore was divided into road segments, each of length no greater than 6km. Each pair of such segments was treated as a bin, and a trip was assigned to bin  $(a, b)$  if its origin was on segment  $a$  and its destination was on segment  $b$ . The demand distribution estimate  $f$  is taken as the distribution whose sampling procedure is: (i) choose a bin  $(a, b)$  with probability proportional to the number of trips, then (ii) produce O-D pair  $(o, d)$ , with  $o$  and  $d$  independent and uniformly distributed, over  $a$  and  $b$ , respectively.

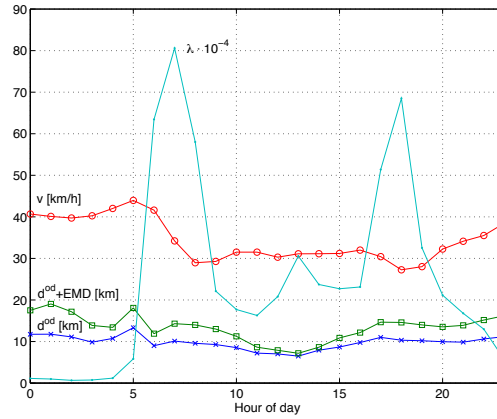
**Earth Mover’s Distance** — To estimate EMD, Singapore was partitioned into regions  $R_1, \dots, R_N$ . Origin and destination points of trips were assigned to the corresponding regions, thus defining pick-up and drop-off bins. The EMD is computed using a linear program that minimizes the amount of work, i.e., the cumulative distance traveled by all points in the pick-up bins, required to transform the distribution

of origin points into the distribution of destination points. The distances involved in this calculation are inherently Euclidean; the previously described scaling factor  $\beta = 1.38$  was used to approximate the distance on the underlying road network. See also [24] for a more accurate method to estimate EMD on road networks.

**Average Velocity ( $v$ )**—Taxi data was used to determine a conservative estimate of the average speed at which occupied taxis move about the city in the *current* traffic conditions. This value is then used as an estimate for  $v$  in (2). Note that this does not take into account potential decreases in congestion due to vehicle sharing. To determine how fast, on average, an individual taxi travels, the total distance traveled by the taxi, with a passenger on board, was divided by the total associated time during each hour of the day over the course of a typical week in Singapore.

**Minimum fleet size**—Given the aforementioned quantities, (2) yields that at least 92,693 automated vehicles are required to ensure the transportation demand remains uniformly bounded. Note however, that this should only be seen as a lower bound on the fleet size, since customer waiting times would be unacceptably high.

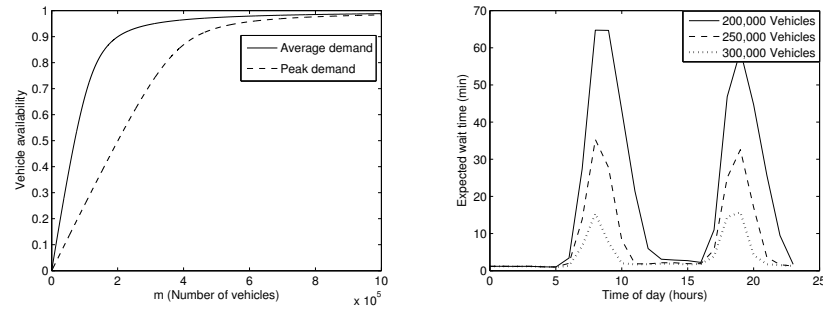
**Fig. 1** Summary of the data necessary for the evaluation of the minimum fleet size. According to Eq. (2) the minimum fleet size to serve all of Singapore’s mobility demand is 92,693 shared vehicles. To gain an appreciation for the level of vehicle sharing possible in an AMoD system of this size, consider that at 1,144,400 households in Singapore, there would be roughly one shared car every 12.3 households.



## 5.2 Performance-driven fleet sizing

Finally, we consider how much the fleet size should be increased in order to reduce the waiting times of customers to acceptable levels. We use the technique described in Section 3.3. To apply the approach to Singapore, the HITS data is first used to partition the city’s road network into  $N = 100$  regions using  $k$ -means clustering. This number of regions corresponds to an average driving time from booking to pickup of 2.3 minutes. The system parameters  $\lambda_i$ ,  $p_{ij}$ , and  $T_{ij}$  are estimated using trip data between regions.

Vehicle availability was analyzed in two representative cases. The first was chosen as the 2-3pm bin, since it is the one that is the closest to the “average” traffic



**Fig. 2** (Left) Performance curve with 100 regions, showing the availability of vehicles vs. the size of the system for both average demand (2-3pm) and peak demand (7-8am). (Right) Average wait times over the course of a day, for systems of different sizes.

condition. The second case considers the 7-8am rush-hour peak. Results are summarized in Figure 2 (left). With about 200,000 vehicles availability is about 90% on average, but drops to about 50% at peak times. With 300,000 vehicles in the fleet, availability is about 95% on average and about 72% at peak times.

In a real MoD system passengers would typically wait for the next available vehicle rather than leave the system immediately if no vehicles are available upon booking. Thus, it is important to characterize how the availability criterion relates to customer waiting times in a practical system. We characterize the customer waiting times through simulation, using a closed-loop rebalancing policy inspired by the loss model, where rebalancing is performed every 30 minutes by minimizing the distance travelled by rebalancing vehicles while evenly distributing the free vehicles across all the stations. For the average-demand case, a fleet of 200,000 vehicles corresponds to expected booking times of less than a minute, to which one must add the pickup driving time, for a total of about 3 minutes between booking and pickup.

Figure 2 (right) shows simulation results of average wait times over the course of a day. For 250,000 vehicles, the maximum wait times during peak hours is around 30 minutes, which is comparable with typical congestion delays during rush hour. With 300,000 vehicles, peak wait times are reduced to less than 15 minutes. To put these numbers into perspective, in 2011 there were 779,890 passenger vehicles operating in Singapore [25].

## 6 Financial Analysis

Other benefits notwithstanding, financial considerations will undoubtedly factor in to if and when cities switch to an AMoD system. To understand the costs associated with such a move, we consider the total mobility cost (TMC) for users in two competing transportation models. In each case, we consider not only the explicit costs to access mobility, but also hidden costs attributed to the time invested in var-

ious mobility-related activities. Within this framework, our analysis indicates that an AMoD system is a financially viable alternative to private vehicle ownership in Singapore. Moreover, to gain an appreciation for the financial benefits of installing AMoD systems in other markets, we provide similar estimates for a typical city in the US. Throughout, all costs are reported in US dollars, with an assumed exchange rate of 1.25 SGD/USD.

The competing transportation models will be referred to as Systems 1 and 2. In System 1, users access personal-mobility by purchasing (or leasing) a private vehicle. Vehicles in System 1 must be operated by a human driver and are referred to as *human-driven cars* (HDCs). In this way, System 1 represents personal mobility as we know it today. Conversely, in System 2, users access personal mobility by subscribing to a shared AMoD fleet of vehicles. Vehicles in System 2 are referred to as *shared self-driving cars* (SSDCs).

Based on the findings in Section 5.2, estimates for an AMoD fleet size in Singapore correspond to a sharing ratio of approximately 3.5–4.5. For simplicity, we will assume that, on average, 4 people effectively share a single SSDC.

### 6.1 The explicit cost of mobility

For System  $i$ , the cost of service ( $COS_i$ ) is defined to be the sum of all explicit costs associated with accessing mobility. For example, in System 1,  $COS_1$  reflects, among other expenditures, the costs to individually purchase, service, park, insure, and fuel a private vehicle. In Singapore, the estimated annual cost to own a mid-sized car, including parking expenses, is approximately \$18,162/year [26]. In the US, assuming an annual mileage of 21,580 km/year, and factoring in the \$1,992/year spent on parking [27], the equivalent figure is \$11,315/year [28]. The disparity between Singaporean and US numbers is due primarily to hefty ownership taxes and traffic tolls within Singapore.

In System 2, fielding a fleet of SSDCs will, initially, require retrofitting production vehicles with the sensors, actuators, and computational power required for automated driving. While still relatively expensive, it is expected that with technological advancement, the needed components and customizations will become more affordable. Assuming some economies of scale for large fleets, we estimate the necessary retrofit of a mid-sized car can be completed for a one-time fee of \$15,000. Automated capabilities will gradually be incorporated into production cycles, with fully automated vehicles eventually rolling off assembly lines. Cost savings associated with mass production suggest these figures have the potential to be significantly smaller in coming years and as AMoD systems become more prevalent.

From the fleet-sizing arguments of Section 5.2, one SSDC in System 2 can effectively serve the role of 3.5–4.5 HDCs in System 1. However, this reduction in vehicles on the road requires a typical SSDC to drive much farther, per day, than an HDC in System 1. Consequently, a typical SSDC in System 2 will depreciate at a faster rate than an HDC in System 1. Accounting for both usage-driven and age-

related depreciation, we conservatively estimate that a SSDC will have an average lifespan of 2.5 years [29]. Moreover, these high utilization rates and the shared nature of SSDCs will require significant maintenance and cleaning budgets to uphold high levels of customer safety and satisfaction.

The routine of a typical SSDC consists of dropping off one passenger and immediately departing to pick up a new passenger. A positive side-effect of this functionality is a drastic reduction, as compared to System 1, in the demand for parking spaces on high-value land. Moreover, should an overabundance of SSDCs develop, e.g., in the hours after the morning rush to work, surplus SSDCs can park themselves in structures or on lots occupying low-valued land (or even earn extra revenue providing a solution for logistics, e.g., shipping parcels and goods within the city). With respect to fuel usage, the central authority that manages System 2 may negotiate bulk fuel deliveries and benefit from discounted rates. Additionally, SSDCs may be programmed to drive in fuel-efficient ways, e.g., by employing gradual acceleration and proactive braking techniques to realize further savings. Tallying the aforementioned costs on a fleet-wide scale and distributing the sum evenly among the intended user base (i.e., the entire population) gives a  $\text{COS}_2$  of \$12,563/year in Singapore and \$9,728/year in the US.

According to COS values, it is more affordable to access mobility in System 2 than System 1. However, the analysis thus far does not reflect the value of the time saved in System 2 by avoiding the more burdensome mobility-related obligations in System 1. For example, users in System 2 not only avoid paying for parking, as reflected in  $\text{COS}_2$ , they also spare themselves the hassle of searching for parking spaces. As the following discussion attests, accounting for these factors further substantiates the financial advantages of AMoD technology.

## 6.2 *The hidden cost of mobility*

Following an approach first pioneered to explore the hidden costs of owning a personal computer in the 1980s [30], we define the value of time (VOT) to be the monetary valuation of the total time invested in mobility related activities. For example, in System 1,  $\text{VOT}_1$  reflects, among other commitments, the time spent taking a car to get a tuneup, paying (or contesting) traffic tickets, renewing license plates, and driving the car. The total mobility cost of System  $i$  is then given by  $\text{TMC}_i = \text{COS}_i + \text{VOT}_i$ ,  $i = 1, 2$ .

The American National Household Travel Survey estimates that an individual spends 465 hours/year in their car [31]. In addition, drivers begin and end each trip by spending an estimated 4 minutes traveling to or from their parked vehicle [32], or, at an average of 3.8 trips per day [33], 175 hours/year. Factoring in the time required to renew license plates, pay tickets, tow a broken-down vehicle, wait while the vehicle is serviced etc., we estimate the total time spent on vehicle ownership and operation related activities in System 1 to be 885 hours/year in the US. Similar studies are not yet available for Singapore. However, given an average travel

distance of 19,000 km/year [25], and the average driving speed on roads, private vehicle owners in Singapore spend, on average, 458 hours/year driving in their car. Factoring in the time parking and other related activities brings this total to 747 hours/year.

To monetize the preceding values, we use the Value of Travel Time Savings (VTTS) numbers laid out by the Department of Transportation (DOT) for performing a Cost Benefit Analysis (CBA) of transportation scenarios in the US [34], [35]. The CBA is used by governments to decide whether or not to proceed with major traffic-related projects, e.g., bridge construction or highway expansion. For various trip scenarios, the VTTS is expressed as a fraction of the median income based on the level of comfort for various in-car trips; less comfortable scenarios incur higher costs. For example, in free-flowing traffic, personal trips on local roads are priced at 50% of the median wage (\$10.80/hour in Singapore and \$12/hour in the US) [36]. Business trips on local roads are valued at 100% of the median wage [37]. Personal trips between cities are considered a greater inconvenience; as such, they are priced slightly higher, at 75% of the median wage. Traveling on heavily congested traffic increases the VTTS to 150% of the median wage [38]. Similar values are available for pricing other-driving related activities, e.g., parking a car.

Applying the appropriate VTTS values based on actual driving patterns gives  $VOT_1 = \$14,460/\text{year}$  in Singapore and  $\$18,295/\text{year}$  in the US. Adding in the associated COS gives an annual  $TMC_1$  of  $\$32,622/\text{year}$  or  $\$1.72/\text{km}$  in Singapore and  $\$29,610/\text{year}$  or  $\$1.37/\text{km}$  in the US. The latter value is significantly higher than the  $\$0.49/\text{km}$  reported by AAA for travel in the US [39]. Furthermore, for all the media attention paid to gasoline prices, fuel costs comprise only six percent of  $TMC_1$ .

To compute  $VOT_2$ , we take a closer look at the activity breakdown associated with taking a trip in an AMoD system. This includes the time spent requesting, waiting for, entering, traveling in, and exiting an SSDC. Given the capabilities of an SSDC, users in System 2 spend no time parking and limited time walking to and from the vehicle. We assume that requesting an SSDC would take no more than 1 minute, and that the fleet is sized such that users wait, on average, no more than 5.5 minutes for a requested vehicle to show up. Given AMoD systems do not yet exist, there are no published VTTS value for the time spent traveling in a SSDC. We price sitting comfortably in an SSDC while being able to work, read, or simply relax at 20% of the median wage. This is significantly lower than the average of 67% of the median wage rate used to compute  $VOT_1$ .

Working from the figures above,  $VOT_2$  is  $\$4,959/\text{year}$  in Singapore and  $\$5,527$  in the US/year, approximately one third of the corresponding  $VOT_1$ . For an individual who is a high wage earner or spends an above average amount of time traveling by car, the gains are even greater.

**Table 1** Summary of the financial analysis of the mobility-related cost for the mobility models discussed in the text. (The average Singaporean drives 18,997 km in a year, the average American drives 21,581 km in a year.)

	Cost [USD/km]						Yearly cost [USD/year]					
	Singapore			United States			Singapore			United States		
	COT	VOT	TMC	COT	VOT	TMC	COT	VOT	TMC	COT	VOT	TMC
HDC	0.96	0.76	<b>1.72</b>	0.52	0.85	<b>1.37</b>	18,162	14,460	<b>32,622</b>	11,315	18,295	<b>29,610</b>
SSDC	0.66	0.26	<b>0.92</b>	0.45	0.26	<b>0.71</b>	12,563	4,959	<b>17,522</b>	9,728	5,527	<b>15,256</b>
SDMC	0.66	0.51	<b>1.17</b>	0.45	0.50	<b>0.95</b>	12,563	9,683	<b>22,246</b>	9,728	10,835	<b>20,563</b>
PSDC	1.09	0.22	<b>1.31</b>	0.62	0.21	<b>0.83</b>	20,712	4,160	<b>24,872</b>	13,408	4,567	<b>17,976</b>
Taxi	1.06	0.26	<b>1.32</b>	—	—	—	20,169	4,959	<b>25,128</b>	—	—	—

### 6.3 Alternate mobility models

To further illustrate how the shared and automated nature of System 2 reduces the TMC for the average user, we briefly consider three additional systems. System 3 consists of Shared Dual-Mode Cars (SDMCs). An SDMC is driven by a human when one or more passengers is onboard, but drives autonomously when vacant. SDMCs therefore have the ability to rebalance themselves in order to meet the travel demand. System 4 is comprised of Personal Self-Driving Cars (PSDCs), each functionally equivalent to an SSDC, but owned and operated by a single individual. Finally, System 5 models a world in which human-driven taxis provide personal mobility (in place of private cars) for the population. For each system, we used similar techniques to estimate the TMC for installations in both Singapore and the US. (The taxi model was only evaluated in Singapore, due to the central role played by taxis in Singapore’s transportation system, and to the numerous different ways taxi services are operated throughout the US.)

### 6.4 Discussion

A summary of the COS, VOT, and TMC of the five systems is provided in Table 1. Remarkably, combining COS and VOT figures, the TMC for SSDCs is roughly half of that for HDCs in both Singapore and the US. To put this into perspective, these savings represent about one third of GDP per capita. On a relative basis, the savings afforded by AMoD technology in Singapore stem largely from the ability to split the hefty cost of car ownership. In the US, the savings are predominantly the result of being able to travel more comfortably and eliminate parking activities.

From the preceding arguments, the true cost to access mobility includes not only an explicit financial investment, but also a significant investment of valuable time. These factors combined, our analysis reveals it is much more affordable to access mobility in an AMoD system compared to traditional mobility models based on private vehicle ownership.

## 7 Conclusions and Future Directions

This paper has provided analytical guidelines for rigorously sizing Automated Mobility-on-Demand (AMoD) systems based on transportation data. Namely, bounds were provided for the minimum number of vehicles required to (i) meet the transportation demand in Singapore and (ii) keep the waiting times below an acceptable threshold. Results suggest that an AMoD solution could meet the personal mobility need of the entire population of Singapore with a fleet whose size is approximately 1/3 of the total number of passenger vehicles currently in operation. Moreover, a financial analysis indicates AMoD systems are a financially viable alternative to more-traditional means of accessing personal mobility.

Given the multifaceted nature of AMoD systems, the results reported herein suggest a number of natural extensions. For example, it remains to thoroughly quantify the extent to which traffic congestion can be alleviated via an automated shared-vehicle solution. In other words, although our solution requires far fewer vehicles than are currently on the road, it mandates the total distance traveled be greater, due to empty vehicles realigning themselves with the demand distribution. Given these competing forces, it remains to certify the net effect is indeed a reduction in travel times and congestion levels.

**Acknowledgements** This work was partially supported by the Singapore National Research Foundation, under the Future Urban Mobility SMART IRG program. The authors are grateful to the Land Transport Authority of Singapore for providing access to the HITS data.

## References

1. M. Buehler, K. Iagnemma, and S. Singh, editors. *The DARPA Urban Challenge: Autonomous Vehicles in City Traffic*. Number 56 in Springer Tracts in Advanced Robotics. Springer, 2009.
2. A. Eskandarian, editor. *Handbook of Intelligent Vehicles*. Springer, 2012.
3. W.J. Mitchell, C.E. Borroni-Bird, and L.D. Burns. Reinventing the automobile: personal urban mobility for the 21st century. *MIT Press*, 2010.
4. B. Edelman. Towards improvement in Singapore's transportation efficiency and environmental impact. *submission to the National Climate Change Secretariat of Singapore*, 2011.
5. C. Costain, C. Ardon, and K. Habib. Synopsis of users behavior of a carsharing program: a case study in Toronto. *Proc. of Transportation Research Board Annual Meeting*, 2012.
6. D. Brook. Carsharing: startup issues and new operational models. *Proc. Transportation Research Board Annual Meeting*, 2004.
7. R. Cervero. City CarShare: First-year travel demand impacts. *Transportation Research Record*, 2003.
8. T. Stillwater, P. Mokhtarian, and S. Shaheen. Carsharing and the built environment: A GIS-based study of one U.S. operator. *Institute of Transportation Studies, University of California*, 2008.
9. D. Jorge and G. Correia. Carsharing systems demand estimation and defined operations: a literature review. *European Journal of Transportation Infrastructure Research*, 2013.
10. A. Lorimier and El-Genaidy. Understanding the factors affecting vehicle usage and availability in carsharing networks: a case study of Communauto carsharing from Montreal, Canada. *International Journal of Sustainable Transportation*, 2011.



11. D. Papanikolaou. A new system dynamics framework for modelling behavior of vehicle sharing systems. *Proc. of Symposium on Simulation for Architecture and urban design*, 2011.
12. J. Barrios. On the performance of flexible carsharing: a simulation-based approach. *J.*, 2012.
13. M. Barth and M. Todd. Simulation model performance analysis of a multiple shared vehicle system. *Transportation Research Record*, 1999.
14. A. Kek, R. Cheu, and M. Chor. Relocation simulation model for multiple-station shared-use vehicle systems. *Transportation Research Record*, 2006.
15. D. Efthymiou, C. Antoniou, and Y. Tyrinopoulos. Spatially aware model for optimal site selection: method and application in a greek mobility center. *Transportation Research Record*, 2012.
16. K. Treleaven, M. Pavone, and E. Frazzoli. Asymptotically optimal algorithms for one-to-one pickup and delivery problems with applications to transportation systems. *IEEE Trans. Automatic Control*, (9):2261–2276.
17. M. Pavone, S. L. Smith, E. Frazzoli, and D. Rus. *Int. Journal of Robotics Research*.
18. S.L. Smith, M. Pavone, M. Schwager, E. Frazzoli, and D. Rus. Rebalancing the rebalancers: Optimally routing vehicles and drivers in mobility-on-demand systems. In *American Control Conference*, 2013.
19. L. D. Burns, W.C. Jordan, and B. A. Scarborough. Transforming personal mobility. *The Earth Institute*, 2013.
20. L. Ruschendorf. The Wasserstein distance and approximation theorems. *Probability Theory and Related Fields*, 70:117–129, 1985.
21. Forest Baskett, K. Mani Chandy, Richard R. Muntz, and Fernando G. Palacios. Open, closed, and mixed networks of queues with different classes of customers. *J. ACM*, 22(2):248–260, April 1975.
22. David K George and Cathy H Xia. Fleet-sizing and service availability for a vehicle rental system via closed queueing networks. *European Journal of Operational Research*, 211(1):198–207, 2011.
23. Singapore Land Transport Authority. 2008 household interview travel survey background information. 2008.
24. K. Treleaven and E. Frazzoli. An Explicit Formulation of the Earth Movers Distance with Continuous Road Map Distances. *ArXiv e-prints*, September 2013.
25. Land Transport Authority. Singapore land transit statistics in brief. 2012.
26. R. Ong. Ride or drive?: The cost of car ownership vs. public transport in Singapore. *MoneySmart.sg*, 2011.
27. Colliers International. 2012 parking rate survey. 2012.
28. American Automobile Association. Your driving costs. 2013.
29. Minnesota Department of Transportation. The per mile costs of operating automobiles and trucks, 2003.
30. L. M. Ellram. Total cost of ownership. In *Handbuch Industrielles Beschaffungsmanagement*, pages 659–671. Springer, 2002.
31. Department of Transportation. Summary of travel trends: 2009 national household travel survey. 2013.
32. TRB Transit Cooperative Highway Research Program. Transit capacity and quality of service manual.
33. U.S. Department of Transportation. Summary of travel trends: 2009 national household travel survey. 2009.
34. U.S. Department of Transportation. Revised departmental guidance on valuation of travel time in economic analysis. 2011.
35. HEATCO. Harmonized European approaches for transport costing and project assessment. 2006.
36. Singapore yearbook of manpower statistics. 2012.
37. U.S. Department of Transportation. Revised departmental guidance valuation of travel time economic analysis. 2011.
38. Victoria Transport Policy Institute. Transportation cost and benefit analysis II - travel time costs. 2013.
39. American Automobile Association. Your driving costs, 2013.