

TOWARD UNDERSTANDING THE IMPACT OF USER PARTICIPATION IN AUTONOMOUS RIDESHARING SYSTEMS

Wen Shen, Rohan Achar, and Cristina V. Lopes

University of California, Irvine

Irvine, CA 92697, USA

{wen.shen, rachar, lopes}@uci.edu

ABSTRACT

Autonomous ridesharing systems (ARS) promise many societal and environmental benefits, including decreased accident rates, reduced energy consumption and pollutant emissions, and diminished land use for parking. To unleash ARS' potential, stakeholders must understand how the degree of passenger participation influences the ridesharing systems' efficiency. To date, however, a careful study that quantifies the impact of user participation on ARS' performance is missing. Here, we present the first simulation analysis to investigate how and to what extent user participation affects the efficiency of ARS. We demonstrate how specific configurations (e.g., fleet size, vehicle capacity, and the maximum waiting time) of a system can be identified to counter the performance loss due to users' uncoordinated behavior on ridesharing participation. Our results indicate that stakeholders of ARS should base decisions regarding system configurations on insights from data-driven simulations and make tradeoffs between system efficiency and price of anarchy for desired outcomes.

1 INTRODUCTION

Real-time ridesharing (or dynamic carpooling) services have emerged as a convenient and cost-efficient option for commuters to mobilize themselves in cities (Furuhata et al. 2013). These ridesharing systems are transforming urban transportation in an economically feasible, environment-friendly, and socially beneficial way (Furuhata et al. 2013; Shen et al. 2016). Recent advances in autonomous driving technology accelerate this transformation because they have the potential to evolve traditional carpooling services into *autonomous ridesharing systems (ARS)* (Santi et al. 2014; Shen and Lopes 2015; Shen et al. 2016; Alonso-Mora et al. 2017). Through system-wide coordination, ARS promise a host of social and environmental benefits, including alleviated traffic congestion, reduced energy consumption and greenhouse gas emissions, less need for land use, as well as increased safety (Shen and Lopes 2015; Shen et al. 2016).

Prior studies show that three key factors influence the performance of autonomous ridesharing systems (Santi et al. 2014; Gargiulo et al. 2015; Shen et al. 2016; Alonso-Mora et al. 2017). The three factors are: (1) the trip-vehicle assignment (i.e., scheduling, routing, and pricing) algorithms used to match the requests to the vehicles and to compute the fares for the trips (Cordeau 2006; Santi et al. 2014; Shen et al. 2016; Alonso-Mora et al. 2017); (2) the physical structure of the transportation network, such as road networks, the fleet size, the vehicle capacity and the request distribution (Alonso-Mora et al. 2017); and (3) user participation (Shen and Lopes 2015; Shen et al. 2016; Gargiulo et al. 2015). To unleash ARS' potential, it is important for the stakeholders to understand how these factors influence the systems' performance (Shen and Lopes 2015; Gargiulo et al. 2015; Alonso-Mora et al. 2017).

Despite much research on ridesharing, most of them focus on the dynamic assignment problem and the network infrastructure (Cordeau 2006; Santos and Xavier 2013; Shen et al. 2016; Santi et al. 2014; Alonso-Mora et al. 2017). On the one hand, stakeholders of ARS might not be able to implement efficient policies to regulate the systems for societal benefits without knowing how users' individual behavior influences the systems' efficiency (Shen et al. 2017). On the other hand, little attention has been paid to

understanding how passengers' uncoordinated behavior on deciding whether to opt in ridesharing or not affects the systems' performance (e.g., total operational cost).

To fill the gap, we present the first simulation analysis to quantify the impact of user participation on the efficiency of autonomous ridesharing systems with different configurations. Our work advances the state of the art in the following ways:

- We describe a modular, agent-based platform, called *SpaceTime for Autonomous Ridesharing Systems (STARS)*, for simulating large-scale autonomous ridesharing systems.
- We conducted extensive experiments on the STARS platform with 96 different sets of system configurations using the data extracted from the New York City Taxi trip public data. Experimental results demonstrate that tradeoffs between system efficiency and the price of anarchy are necessary and are typically viable in order to achieve desired outcomes.

2 SIMULATING AUTONOMOUS RIDESHARING SYSTEMS

This section discusses methods for solving the trip-vehicle assignment problem, and approaches for modeling users' choice of ridesharing participation.

2.1 Trip-Vehicle Assignment

Traditional algorithms (e.g., the branch-and-cut algorithm (Cordeau 2006), and the annealing meta-heuristic algorithm (Braekers et al. 2014)) for the trip-vehicle assignment problem usually suffer scalability problems and are not applicable to large ridesharing systems that have thousands of vehicles, requests, and streets. This is primarily due to the large search space and the spatial-temporal constraints in the dynamic ridesharing problem (Furuhata et al. 2013; Alonso-Mora et al. 2017).

Recently, Alonso-Mora et al. (2017) extended the static shareability network model by Santi et al. (2014) to dynamic ridesharing with high-capacity vehicles (serving up to ten passengers at the same time). They introduced an anytime algorithm that can compute optimal solutions to the trip-vehicle assignment for large-scale ridesharing systems in real time (Alonso-Mora et al. 2017). The method consists of the following three steps: generating the *request-vehicle (RV)* graph, generating the *request-trip-vehicle (RTV)* graph, and computing the optimal assignment. The anytime approach also runs a load rebalancing process to redistribute the idle vehicles to pick up unassigned requests.

The STARS platform in our work utilizes the anytime algorithm to allow simulating real-time operations of city-scale ridesharing systems.

2.2 Modeling User Participation

In our work, we used *Price of Anarchy (PoA)* (Roughgarden 2005) as the indicator to quantify the extent to which passengers' uncoordinated behavior of ridesharing participation diminishes ARS' performance. This is because price of anarchy has been used as a measure to quantify how the efficiency of a system degrades due to uncoordinated behavior of its agents in many areas, including transportation systems (Youn et al. 2008), smart buildings and power grids (Shen 2013; Shen et al. 2017).

We take a data-driven approach to compute the price of anarchy. Specifically, we select different portions of the passengers as participants, and the rest of the population as non-participating riders. We assume that each of the selected choice profiles follows a stochastic user equilibrium (Daganzo and Sheffi 1977): no user believes she can improve her utility by unilaterally changing her decision on whether to participate or not. This assumption allows us to model passengers' choice behavior based on real-world data instances rather than synthesized samples. Similar assumptions are also made in literature (Youn et al. 2008).

3 SIMULATION PLATFORM

In this section, we describe the *SpaceTime for Autonomous Ridesharing Systems (STARS)* simulation platform. We first introduce the basic design of STARS and then discuss details of its implementation.

3.1 System Design

STARS is a multi-agent simulation platform on top of a simulation framework called *Spacetime* (Valadares et al. 2016; Lopes et al. 2017). Spacetime follows a proof-of-concept paradigm called *Predicate Collection Classes (PCCs)* (Lopes et al. 2017) to increase modularity. The design of STARS consists of five aspects: *frame*, *time flow*, *data flow control*, *data model*, and *data store* (See Figure 1a for an abstract view of STARS). We discuss each in turn as follows:

- **Frame** allows system designers to attach applications, fetch objects of desired types, add new objects and delete objects every time step.
- **Data flow control** declares each application’s interaction with its PCCs. There are five main declarations:
 - *Producer*: the application creates objects of specified types;
 - *Getter*: the application fetches objects of specified types;
 - *Setter*: the application updates objects of specified types;
 - *GetterSetter*: the application fetches and updates objects of specified types;
 - *Deleter*: the application deletes object of specified types.
- **Time flow** defines the execution frequency of loop operations (i.e., *pull*, *update* and *push*). The *pull* operation retrieves from the data store to get current information of the objects whose types are included in *Getter* and *GetterSetter*. The *update* operation calls the update function during which the *Frame* collects all the modifications of the setter objects whose types are included in *GetterSetter*. In the *push* operation, the *Frame* send all the modifications collected in the *update* operation to the data store. The loop operations are performed every fixed interval of clock time.
- **Data model** defines the objects of the simulation in terms of classes that are both types and sets. These classes are called *Predicate Collection Classes (PCCs)* (Lopes et al. 2017).
- **Data store** keeps track of the PCC objects. It assumes a distributed simulation environment using pass-by-copy to fulfill data requests, and handles the operation of pushing state updates back.

In the original design of Spacetime, each application is encapsulated and is independent of each other (Valadares et al. 2016; Lopes et al. 2017). However, multi-agent simulations typically require interactions between different applications (agents) (Macal and North 2010). Different from Spacetime, STARS enables communications between different applications on top of each individual application agent. The communications between the interacting applications also follow a chronological order (i.e., arranged in order by time). With a dedicated modular design, the STARS platform supports complex, scalable simulations in a distributed, or multicore environment, making it a good fit for autonomous ridesharing simulations.

3.2 Platform Implementation

We next report the implementation details of the STARS platform. STARS comprises five simulation agents: the *request generating agent* that reads requests from raw data, the *request updating agent* that refreshes the status of requests, the *vehicle updating agent* that initializes vehicles and updates vehicle status, the *data logging agent* that records the state of the ridesharing system, the *dispatching agent* that computes optimal assignments, and the *load rebalancing agent* that redistributes vehicles for picking up unassigned requests. Among them, the *dispatching agent* is composed of three working agents: the RV-graph generator, the RTV-graph generator and the assignment optimizer. The RV-graph generator computes which requests can be paired and which vehicles can service which requests individually. The RTV-graph generator calculates

feasible trips that can be combined and picked up by a vehicle. The assignment optimizer computes the optimal assignment using the anytime algorithm described in the work by Alonso-Mora et al. (2017).

In STARS, each agent independently fulfills its role in the simulation within a frame that has predefined computational space and time flow. In what follows we discuss the role of each agent in turn.

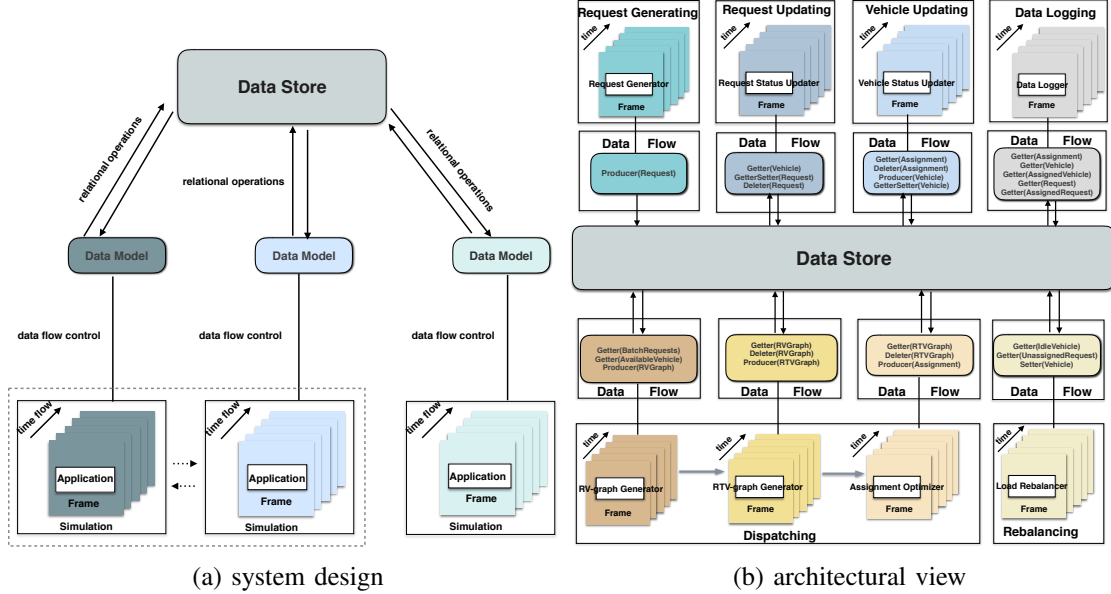


Figure 1: The STARS simulation platform.

- **Request Generating:** The request generating agent produces request objects and pushes them to the data store every time interval θ ($\theta = 1s$ in our experiments). A request is defined as a tuple of origin, destination, the time of the request, the latest acceptable pick-up time, the pick-up time, the expected drop off time, and the earliest possible time for reaching the destination.
- **Request Updating:** Given the locations of the vehicles and the vehicles' passengers (the requests that are being serviced by the vehicles), the request updating agent refreshes the status of the requests every time interval θ ($\theta = 0.5s$ in our experiments). If the destination of a request has been reached, the request will be deleted.
- **Data Logging:** The agent records the state of the assignment, vehicles, and the requests every time interval θ ($\theta = 0.5s$ in our experiments).
- **Dispatching:** The dispatching agent performs three steps using the algorithm described in (Alonso-Mora et al. 2017): Computing the pairwise request-vehicle RV graph, generating the request-trip-vehicle RTV graph and calculating the optimal assignment. We elaborate each step in turn as follows:
 - *Computing RV graph:* Given available vehicles and a batch of requests, this step is to compute the requests that can be pairwise combined, and to search the vehicles that can serve which requests individually ($\theta = 30s$ in our experiments).
 - *Generating RTV graph:* Given the RV graph, this step is to find feasible trips that can be combined and picked up by a vehicle subject to all the constraints. It is done by computing complete subgraphs or cliques of the RV graph ($\theta = 30s$ in our experiments).
 - *Calculating the optimal assignment:* Given the RTV graph, this step is to calculate the optimal assignment of vehicles to trips with one trip per vehicle at most. It is solved incrementally by formalizing the problem into an Integer Linear Program ($\theta = 30s$ in our experiments).

- **Vehicle Updating:** Given the assignment computed by the dispatching agent, the vehicle updating agent adds the requests in the assigned trip to the respective vehicles. Once the an assignment has been processed, it will be deleted from the data store. The agent updates the location of each vehicle according to the designated route associated with it ($\theta = 0.5s$ in our experiments). The vehicles are initialized at the beginning of the simulation.
- **Load Rebalancing:** Given the idle vehicles, and the unassigned requests, the rebalancing agent redistributes vehicles to the locations of unassigned requests ($\theta = 30s$ in our experiments). Specifically, after each assignment, if there are some unassigned requests and some available vehicles, the rebalancing process is executed.

STARS adopts a modular design, making it easily expandable and configurable. By using the state-of-the-art dispatching algorithm (Alonso-Mora et al. 2017), STARS is capable of simulating real-time operations of autonomous ridesharing systems that have thousands of requests and vehicles at the same time in large road networks. Thus, it is feasible to use STARS to conduct extensive simulations with real-world data in ridesharing to gain insights on how user participation influences the efficiency of autonomous ridesharing systems.

4 DATA AND METHODS

This section describes the experimental settings, including the dataset used in the simulations, and the procedures followed to investigate the impact of user participation in autonomous ridesharing systems.

4.1 Dataset

We used the New York City Taxi trip Dataset (Donovan and Work 2016) in the experiments. We selected a typical week (i.e., a week without major holidays included) in the year of 2011, from Monday, June 6th 00:00:00, to Sunday, June 12th 23:59:59, as the time duration for simulating the operation of the autonomous ridesharing systems. We extracted the whole week’s trips of which the origins and the destinations were within Manhattan from the raw dataset. The resulting dataset contains 3,014,628 trips ranging from 391,246 (Saturday) to 465,331 (Monday) per day. The trips were originally serviced by a fleet of 13,586 taxis. Each trip in the raw data has the following attributes: pickup datetime, dropoff datetime, passenger count, pickup longitude, pickup latitude, dropoff longitude, and dropoff latitude.

In this work, we defined a request r as a tuple $(o_r, d_r, tr_r, tlp_r, tp_r, td_r, te_r, m_r)$, where o_r denotes the origin, d_r is the destination, tr_r represents the time of the request, tlp_r is the latest acceptable pickup time, tp_r denotes the pickup time, td_r is the expected dropoff time, te_r denotes the earliest dropoff time possible, and m_r indicates user participation – user r ’s decision on whether to participate in a shared ride (i.e., $m_r = 1$) or to opt for a private one (i.e., $m_r = 0$). Since the raw data does not include all the information as required for autonomous ridesharing simulations, we processed the data with the following procedures: we used the pickup datetime as the request time. We calculated the latest acceptable pickup time by $tlp_r = tr_r + \Omega$, where Ω is the maximum waiting time allowed in a ridesharing system. The earliest dropoff time possible te_r is computed by $te_r = tr_r + \tau(o_r, d_r)$, where $\tau(o_r, d_r)$ is the shortest travel time from o_r to d_r . The total delay δ_r due to ridesharing for a serviced request r is calculated by $\delta_r = td_r - te_r$. In our experiments, we required that $\delta_r \leq \Delta$ for all serviced requests, where Δ is the maximum delay allowed in a ridesharing system. That is, for all r , we have $td_r \leq te_r + \Delta$.

We converted the road map of Manhattan into a graph with 4,092 nodes and 9,453 edges. We then estimated the hourly travel time on each edge for each day of the week by using the method described in the work by Santi et al. (2014). With the road network and the travel time estimate, we precomputed the shortest paths and travel time between all the nodes in the network and then stored the calculated results in a lookup table for later use.

4.2 Methods

To study how user participation affects ARS' performance, we selected the following metrics: the service rate (the percentage of requests serviced, the higher the better), the mean waiting time (the smaller the better), the mean in-car delay (the difference between the travel delay and the waiting time, the smaller the better), the percentage of shared rides (the higher the better), the mean travel distance (the smaller the better) and the total cost (the smaller the better). Let R denote the set of requests, P represent the set of profiles that characterize all the passengers' choices on participation, and $p = (m_r)_{r \in R}$ be a generic profile in P . The total cost $\mathcal{C}(p)$ of the system under the profile p is the sum of travel delay δ_r of all serviced requests $r \in R_o \subseteq R$ plus a constant cost c_d ($c_d = 1000$ in our experiments) for each denied request $r \in R_d = R \setminus R_o$. That is,

$$\mathcal{C}(p) = \sum_{r \in R_o} \delta_r + \sum_{r \in R_d} c_d. \quad (1)$$

The price of anarchy due to passengers' uncoordinated behavior of ridesharing participation is defined as

$$PoA = \frac{\max_{p \in P_{sue}} \mathcal{C}(p)}{\min_{p \in P} \mathcal{C}(p)}, \quad (2)$$

where $P_{sue} \subseteq P$ is the set of choice profiles under stochastic user equilibria.

We varied the percentage of user participation from 0% to 100% with an increasing step of 10% to see how different levels of user participation affects the performance of different systems. We used the stratified sampling method (Neyman 1934) to select the participating requests because previous research indicates that differences in both neighborhoods and hour of the day can affect passengers' choices of commute modes (Jackson and Jucker 1982; Schwanen and Mokhtarian 2005). The method first categorized the request data into 888 subgroups according to both spatial differences (37 neighborhoods in Manhattan, see Figure 2b) and temporal variations (24 hours, see Figure 2d). For each subgroup, it then randomly selected corresponding percent of requests as the participating requests ($m_r = 1$) and the rest as the non-participating requests ($m_r = 0$). This produced 11 datasets in total, with each contained the request data for a week. We considered autonomous ridesharing systems with different settings, including vehicle fleet sizes (1000, 2000, 3000, and 4000), vehicle capacities (1, 2, 4, 6, 8, and 10) and maximum waiting time (2, 4, 6, and 8 minutes). The maximum delay time allowed, including both waiting time and in-car delay time, was calculated by doubling the maximum waiting time (i.e., 4, 8, 12, and 16 minutes). This resulted in 96 different ridesharing systems. For each system, we performed a simulation on the STARS platform with each of the 11 generated datasets using the anytime optimal algorithm described in (Alonso-Mora et al. 2017). There were 1,056 simulations in total. For each simulation, we recorded the following information: service rate, average in-car delay ($\delta - \omega$), average waiting time, mean distance traveled by each vehicle during a day, the percentage of shared rides, and the total cost. For each ridesharing system, we then computed the price of anarchy by assuming $P_{sue} \subseteq P = \{p_l\}$, where $l \in \{0\%, 10\%, 20\%, 30\%, \dots, 100\%\}$ is the percentage of user participation, and p_l denotes the profile where there are l percent of users that are willing to take a shared ride. Essentially, a profile is a set of all users' decisions on whether to participate in ridesharing or not.

We performed another two groups of experiments to study if user participation affects the efficiency of the same ridesharing system with different request density and different traffic conditions. In the first group, we varied the request density. The typical travel demand in Manhattan are about 3,000,000 trips per week. We compared three kinds of density: half of the demand (x0.5, by removing the even number of requests), the nominal demand (x1.0), and double of the demand (x2.0, by adding the requests of the subsequent week). In the second group, we used the travel time estimate of three different periods in the day: the estimate at 12:00 (lower speed than the average, see Figure 2c), the estimate at 19:00 (higher speed than the average), and the mean daily travel time estimate (average speed). For each condition in each of the group, we ran a simulation on the same ridesharing system (fleet size = 3000, capacity = 4, maximum waiting time = 6 minutes).

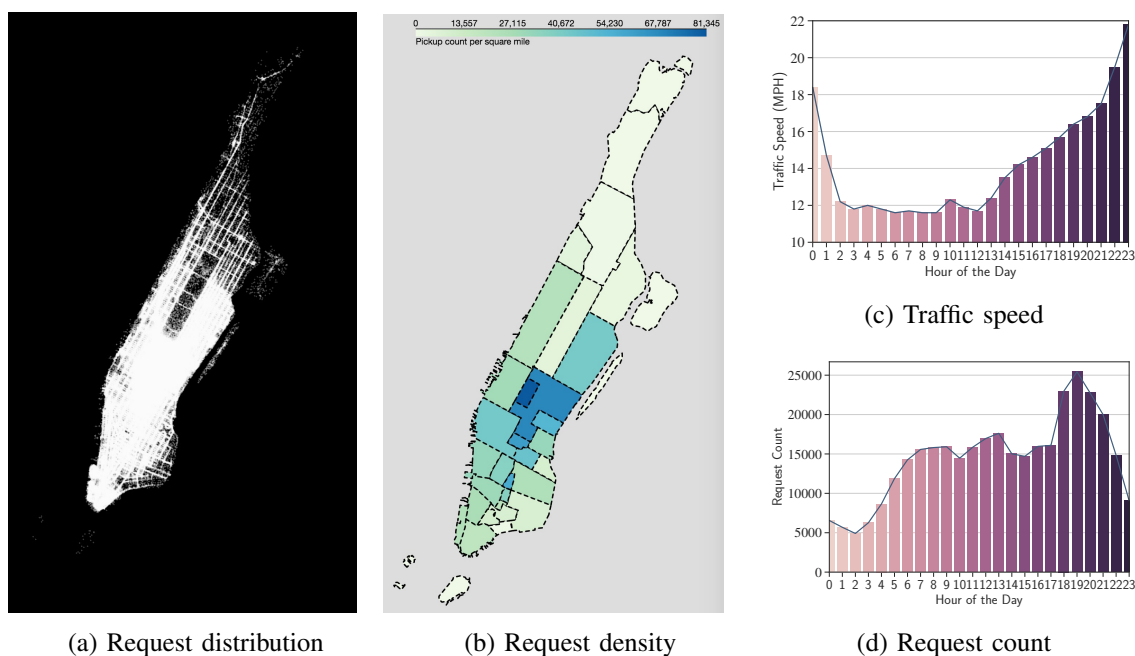


Figure 2: The distribution of request (a), the request count per square mile of each neighborhood (b), the average traffic speed per hour (c), and the request count per hour (d) in Manhattan on Wednesday, June 8th, 2011.

All the simulations were conducted on the same 24-core 3.0GHz Linux machine with 128GB RAM. We used the Gurobi solver (version 7.0.2 with free Academic license) for mixed integer programming optimization. To reduce the total time for simulations, all the time steps for all simulations were multiplied by a discounting factor 0.1. The 1,122 simulations took five weeks.

5 RESULTS

This section describes the experimental results, followed by a discussion of their implications.

5.1 Main Observations

Based on the data points from the numerical simulations, we highlight six main observations. We will next discuss each in turn.

Observation 1. *User participation typically improves the performance of autonomous ridesharing systems, but the degree of improvement generally slows down as user participation increases to a high level.*

This was observed in all the ARS in terms of total cost, service rate, mean waiting time, percentage of shared rides, and average distance traveled by each vehicle during a day. Due to space limitations, we will report the findings mainly based on the total cost.

Figure 3 shows that when the levels of user participation increased, the total cost of all the ridesharing systems ($\text{capacity} \geq 2$) decreased. This was expected since a higher degree of user participation provided the ridesharing systems with more room for system-wide optimization: more requests were paired for shared rides and subsequently were serviced. However, when the level of user participation increased to a high level (i.e., above 80%), the degree of cost reduction slowed down. This was also expected because when the level of user participation reached to certain level, substantial improvement by system optimization no longer sustained due to physical limits of the transportation systems such as spatial and temporal factors (Tsao and Lin 1999).

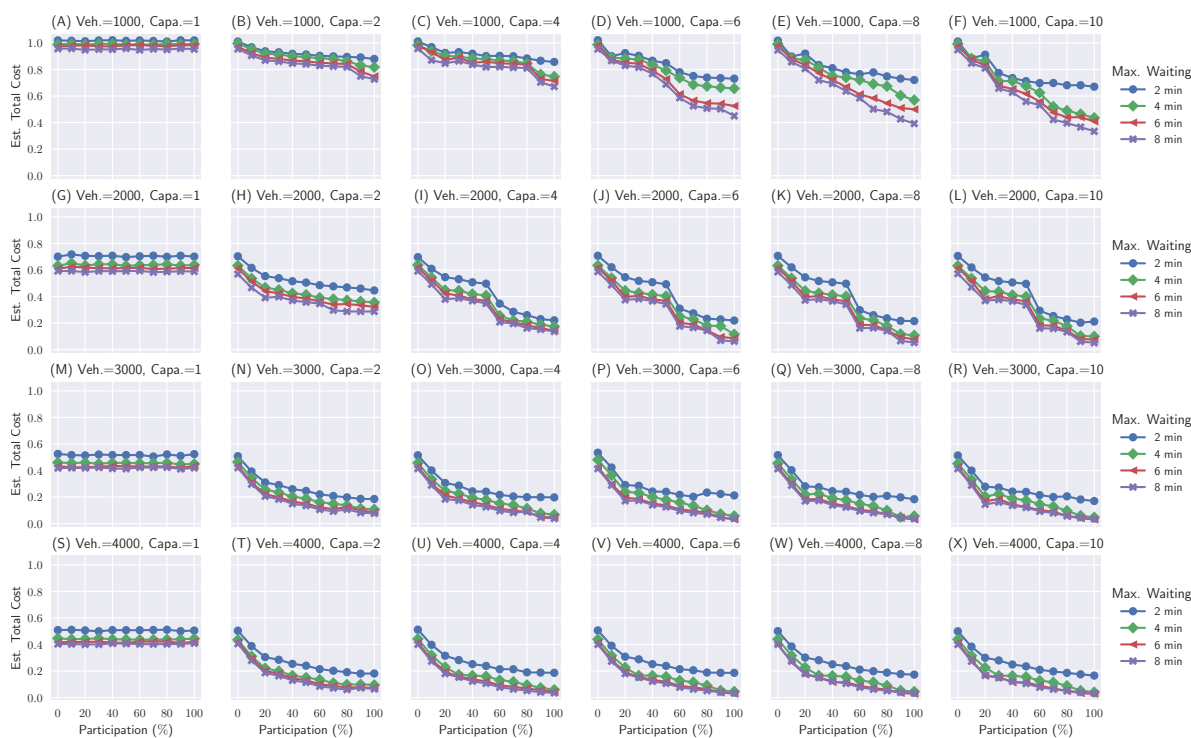


Figure 3: Estimated total cost for ridesharing systems with different configurations.

Note that systems in which the vehicles have the capacity of 1 are excluded since they are not capable of offering carpooling services. In our work, we listed the performance of these systems for comparison purpose only.

Observation 2. *User participation typically has a greater impact on the performance of autonomous ridesharing systems with small fleets of high-capacity vehicles.*

Figure 3 shows that the ARS experienced a larger degree of cost reduction when they were equipped with small fleets of high-capacity vehicles. For instance, higher degree of performance improvement was observed in the ridesharing systems with the following conditions: (1) Vehicle Size = 1000, Capacity = 8 (See Figure 3E); (2) Vehicle Size = 1000, Capacity = 10 (See Figure 3F); (3) Vehicle Size = 2000, Capacity = 8 (See Figure 3K); and (4) Vehicle Size = 2000, Capacity = 10 (See Figure 3L). This trend implies that participation tends to have a greater impact on the performance of ARS with small fleets of high-capacity vehicles than others. An explanation for this trend is that the service rates for the former systems were initially quite low (20% – 40%) due to small fleet sizes and no ridesharing participants. As the level of user participation increases, the large capacity of the vehicles allowed a higher degree of ride sharing, making the degree of performance improvement substantially higher than others.

Observation 3. *User participation typically has a smaller impact on the performance of autonomous ridesharing systems with a short period of maximum waiting time allowed.*

This was also observed in ARS with all the configurations. For example, Figure 3B demonstrates that systems with maximum waiting time of 2 minutes had a lower degree of cost cutting than the other three systems. The cost of the system with maximum waiting time of 2 minutes dropped from 1.0 to 0.88 as the participation level increased from 0% to 100%, while the system with maximum waiting time of 8 minutes experienced a greater reduction (from 1.0 to 0.72) in the total cost. This reason is that a looser constraint on maximum waiting time enabled the systems to service more requests, which was not possible if given a narrow time frame. Similar patterns were also found when the vehicle capacity varied while other settings kept unchanged: systems with larger capacity tended to have a larger shrinkage in the cost.

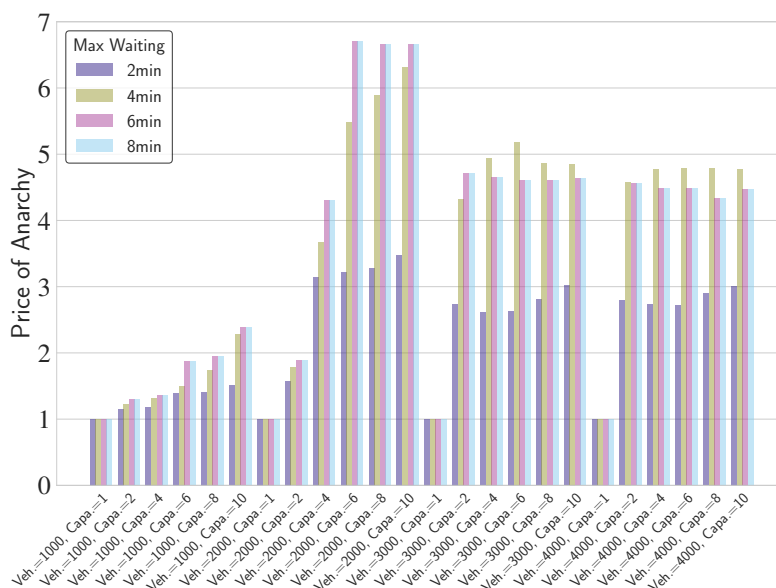


Figure 4: Price of anarchy for ridesharing systems with different configurations.

Observation 4. *As the fleet size increases, the price of anarchy due to users' uncoordinated choices on participation typically first increases and then decreases.*

Figure 4 shows that as the fleet size increased, the price of anarchy first rose steeply (from Vehicle Size = 1000 to 2000) and then dropped to a level with small fluctuations (from Vehicle Size = 2000 to 4000). Specifically, when the fleet size was small (1000), the price of anarchy climbed gradually as the vehicle capacity increased. When the fleet size increased to 2000, the price of anarchy jumped more than doubled when the vehicle capacity was switched from 2 to 4, for all conditions of maximum waiting time. Several factors contributed to this sharp rise: the service rate of the system with vehicle capacity 2 experienced a larger boost than that of the system with vehicle capacity 4, while the changes in the average in-car delay and the mean waiting time kept the same pace.

It is interesting to note that when the fleet size was large (≥ 3000), the price of anarchy fluctuated slightly as the vehicle capacity changed. The price of anarchy for systems with a fleet size of 3000 was almost the same as the systems with a fleet of 4000 vehicles. The reason is that the service rates for both conditions were initially at the same level (about 60%), and systems with both configurations were sufficient for satisfying the travel demand of daily commuters in Manhattan. This implies that a fleet of 3000 vehicles might serve the travel demand of passengers in Manhattan reasonably well. This was consistent with the results by Alonso-Mora et al. (2017). It further indicates that appropriate configurations of the autonomous ridesharing systems can be set to keep the inefficiency caused by passengers' uncoordinated behavior on ridesharing participation minimal.

Observation 5. *The price of anarchy due to users' uncoordinated choices on participation typically keeps steady when the maximum waiting time is sufficiently long.*

When the fleet size was large (≥ 3000), the PoA experienced a significant increase (more than 50%) when switching the maximum waiting time from 2 minutes to 4 minutes. This was expected since the service rates of systems with both were almost the same, while systems with maximum waiting time of 4 minutes had looser constraints and allowed more room the ride matching than the other. As a result, the former experienced a greater improvement of system efficiency.

Interestingly, as the maximum waiting time increased from 6 minutes to 8 minutes, the price of anarchy stayed almost the same despite variations on vehicle capacity or fleet size. The reason is that when the maximum waiting time reaches a sufficient long duration (e.g., 6 minutes), additional waiting time does




Figure 5: A comparison of the price of anarchy and efficiency on a ridesharing system (fleet size =3000, capacity=4, and maximum waiting time = 6 minutes) by varying request density.




Figure 6: A comparison of the price of anarchy and efficiency on a ridesharing system (fleet size =3000, capacity=4, and maximum waiting time = 6 minutes) by varying the time of the day used for travel time estimate.

not improve the performance of ridesharing systems significantly. Other parameters become critical. This implies that it is possible to reduce the influence of user participation on the performance of ridesharing systems by identifying proper maximum waiting time.

Observation 6. Request density or traffic condition typically has little impact on the price of anarchy, although the system efficiency generally boosts when the traffic condition improves or the request density reduces.

Figure 5 demonstrates that user participation had similar impact on ridesharing systems with different request density. Specifically, when the level of user participation increased, the total cost decreased (See Figure 5b). This was expected since higher levels of user participation typically increase the chance of successful ride matching (See Figure 5c) and shared rate (See Figure 5f).

Noticeably, the price of anarchy remained almost unchanged (approximately 5.0) as the density increased (See Figure 5a). The reason is described as follows: from Figure 5c, one can see that the service rates of systems with half of the demand (x0.5), and nominal demand (x1.0) were fairly close, and were higher than the service rate of the system with double demand (x2.0). However, they grew slower than the service rate of the system with double demand. Although the latter system performed worse than the former two in terms of mean waiting time and mean in-car delay, the ratios between the total cost of the worst state (participation level = 0%) and the best state (participation level = 100%) were the same.

Figure 6 illustrates that as the travel condition improved, the system efficiency increased. Similarly, the price of anarchy did not vary according to traffic condition.

Discussion

All the six findings tell a similar story: there are sweet spots (e.g., Vehicle Size = 3000, Capacity = 6, Maximum Waiting Time = 6 minutes) that can be utilized to make tradeoffs to keep the price of anarchy minimal while maintaining a good system efficiency. Stakeholders of ridesharing systems should base decisions regarding system configurations on insights gained from realistic simulations with real-world data when experiments on deployed systems are not possible or too costly.

Our work distinguishes itself from game-theoretic approaches in the sense that we do not assume passengers share the same utility function as many works (Roughgarden 2005; Christodoulou and Koutsoupias 2005; Shen et al. 2016) do. In particular, we do not make the assumption that passengers’ utilities on participation are determined by a unified function or drawn from a probability distribution. This is because passengers’ choices on commuting modes are usually influenced by many factors, such as time variability (Jackson and Jucker 1982), neighborhoods (Schwanen and Mokhtarian 2005), habits and prior experiences (Gargiulo et al. 2015). It is rather difficult or even unrealistic to isolate a utility function or a probability distribution that characterizes all passengers’ decision-making processes.

While illuminating, we must be careful not to overgeneralize these results. Although we conducted 1,122 simulations on ridesharing systems with 96 configurations, our work is an exploratory study rather than an exhaustive one. Similarly, when sampling the ride requests as the ridesharing participants, it was not possible for us to include all the combinations. Our work is intended to spur a pipeline of further studies to deepen our understanding of how human dynamics impact the efficiency of large social-technological systems, as well as to shed light on integrating advances in programming languages into the development of scalable, data-driven, multi-agent based applications. Future work is needed to identify how to best do so.

6 CONCLUSION

In this paper, we describe a modular design for agent-based simulations of ridesharing systems. We introduce a multi-agent simulation platform called STARS, for simulating city-scale, autonomous ridesharing systems. Via extensive experiments on STARS with real-world datasets in ridesharing, we present the first simulation analysis toward quantifying the impact of users’ uncoordinated behavior on participation in autonomous ridesharing systems. The two messages to take away are: (1) specific configurations (e.g., fleet size, vehicle capacity, and the maximum waiting time) of ridesharing systems can be identified to counter the effect of user participation on the system efficiency; and (2) tradeoffs between system efficiency and price of anarchy are needed and are often feasible to achieve desired outcomes.

There are a number of interesting avenues for further research. One of them is to investigate whether our results will generalize to autonomous ridesharing systems with different road networks. Another one is to develop novel approaches (e.g., incentive mechanisms, information structures and regulations) to reduce the price of anarchy in autonomous ridesharing systems. We also find it very rewarding to explore new methods for scalable multi-agent simulations.

ACKNOWLEDGMENTS

This work was supported by National Science Foundation through grant no. CCF-1526593.

REFERENCES

- Alonso-Mora, J., S. Samaranayake, A. Wallar, E. Frazzoli, and D. Rus. 2017. “On-demand high-capacity ride-sharing via dynamic trip-vehicle assignment”. *Proceedings of the National Academy of Sciences* 114(3):462–467.
- Braekers, K., A. Caris, and G. K. Janssens. 2014. “Exact and meta-heuristic approach for a general heterogeneous dial-a-ride problem with multiple depots”. *Transportation Research Part B: Methodological* 67:166–186.
- Christodoulou, G., and E. Koutsoupias. 2005. “The price of anarchy of finite congestion games”. In *Proceedings of the Thirty-Seventh Annual ACM Symposium on Theory of Computing*, 67–73. ACM.
- Cordeau, J.-F. 2006. “A branch-and-cut algorithm for the dial-a-ride problem”. *Operations Research* 54(3):573–586.
- Daganzo, C. F., and Y. Sheffi. 1977. “On Stochastic Models of Traffic Assignment”. *Transportation Science*:253–274.

- Donovan, B., and D. Work. 2016. “New York City Taxi Trip Data (2010-2013)”.
- Furuhata, M., M. Dessouky, F. Ordóñez, M.-E. Brunet, X. Wang, and S. Koenig. 2013. “Ridesharing: The state-of-the-art and future directions”. *Transportation Research Part B: Methodological* 57:28–46.
- Gargiulo, E., R. Giannantonio, E. Guercio, C. Borean, and G. Zenezini. 2015. “Dynamic ride sharing service: are users ready to adopt it?”. *Procedia Manufacturing* 3:777–784.
- Jackson, W. B., and J. V. Jucker. 1982. “An empirical study of travel time variability and travel choice behavior”. *Transportation Science* 16(4):460–475.
- Lopes, C. V., R. Achar, and A. Valadares. 2017. “Predicate Collection Classes”. *Journal of Object Technology* 16(2).
- Macal, C. M., and M. J. North. 2010. “Tutorial on agent-based modelling and simulation”. *Journal of simulation* 4(3):151–162.
- Neyman, J. 1934. “On the two different aspects of the representative method: the method of stratified sampling and the method of purposive selection”. *Journal of the Royal Statistical Society* 97(4):558–625.
- Roughgarden, T. 2005. *Selfish routing and the price of anarchy*, Volume 174. MIT Press Cambridge.
- Santi, P., G. Resta, M. Szell, S. Sobolevsky, S. H. Strogatz, and C. Ratti. 2014. “Quantifying the benefits of vehicle pooling with shareability networks”. *Proceedings of the National Academy of Sciences* 111(37):13290–13294.
- Santos, D. O., and E. C. Xavier. 2013. “Dynamic taxi and ridesharing: a framework and heuristics for the optimization problem”. In *Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence*, 2885–2891. AAAI Press.
- Schwanen, T., and P. L. Mokhtarian. 2005. “What affects commute mode choice: neighborhood physical structure or preferences toward neighborhoods?”. *Journal of Transport Geography* 13(1):83–99.
- Shen, W. 2013. “Regulating Self-Adaptive Multi-Agent Systems with Real-Time Interventions”. *Masdar Institute of Science and Technology*.
- Shen, W., A. A. Khemiri, A. Almehrezi, W. A. Enezi, I. Rahwan, and J. W. Crandall. 2017. “Regulating Highly Automated Robot Ecologies: Insights from Three User Studies”. In *Proceedings of the 5th International Conference on Human Agent Interaction (HAI'17)*, 111–120. ACM.
- Shen, W., and C. Lopes. 2015. “Managing Autonomous Mobility on Demand Systems for Better Passenger Experience”. In *Proceedings of the 18th International Conference on Principles and Practice of Multi-Agent Systems*, Volume 9387, 20–35. Springer 2015.
- Shen, W., C. V. Lopes, and J. W. Crandall. 2016. “An Online Mechanism for Ridesharing in Autonomous Mobility-on-Demand Systems”. In *Proceedings of IJCAI 2016*, 475–481. AAAI Press.
- Tsao, H.-S. J., and D.-J. Lin. 1999. “Spatial and temporal factors in estimating the potential of ride-sharing for demand reduction”. *California Partners for Advanced Transit and Highways (PATH)*.
- Valadares, A., C. V. Lopes, R. Achar, and M. Bowman. 2016. “CADIS: Aspect-Oriented architecture for collaborative modeling and simulation”. In *Winter Simulation Conference, 2016*, 1024–1035. IEEE.
- Youn, H., M. T. Gastner, and H. Jeong. 2008. “Price of anarchy in transportation networks: efficiency and optimality control”. *Physical Review Letters* 101(12):128701.

APPENDIX

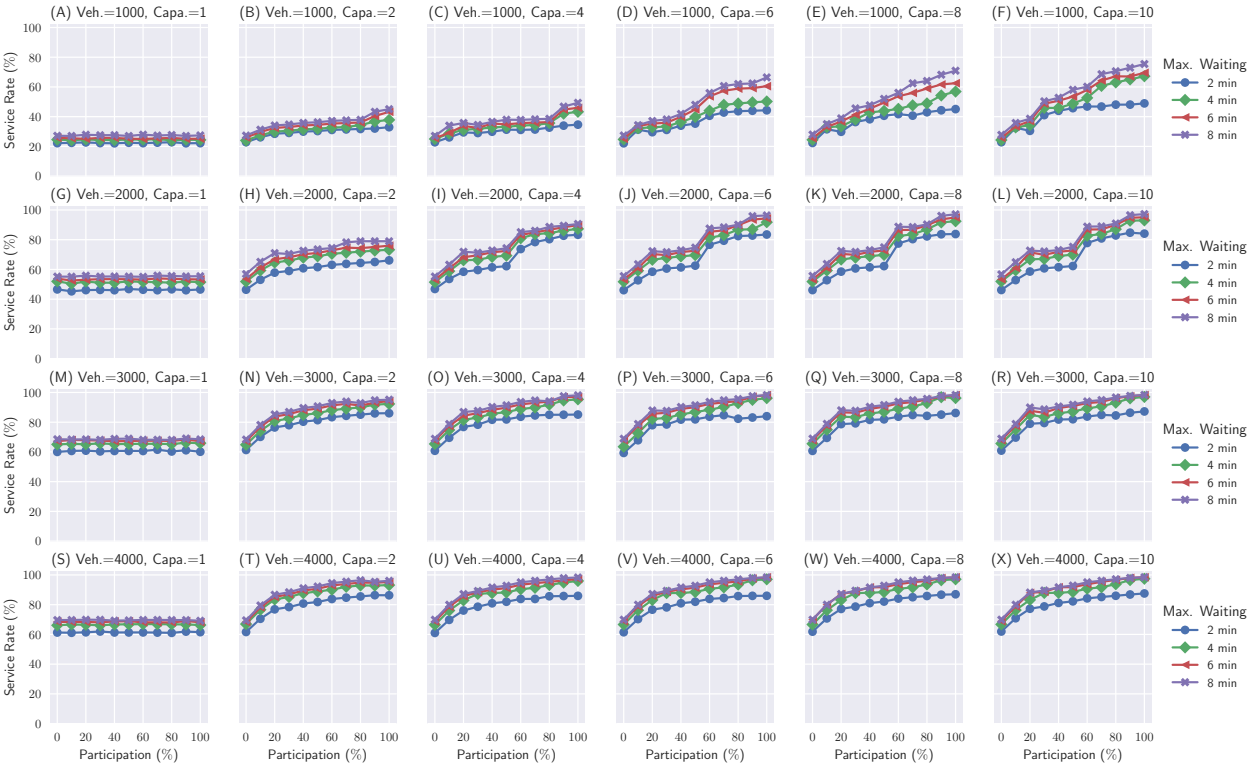


Figure 7: Service rate (percentage of the requests serviced) for autonomous ridesharing systems with different configurations.

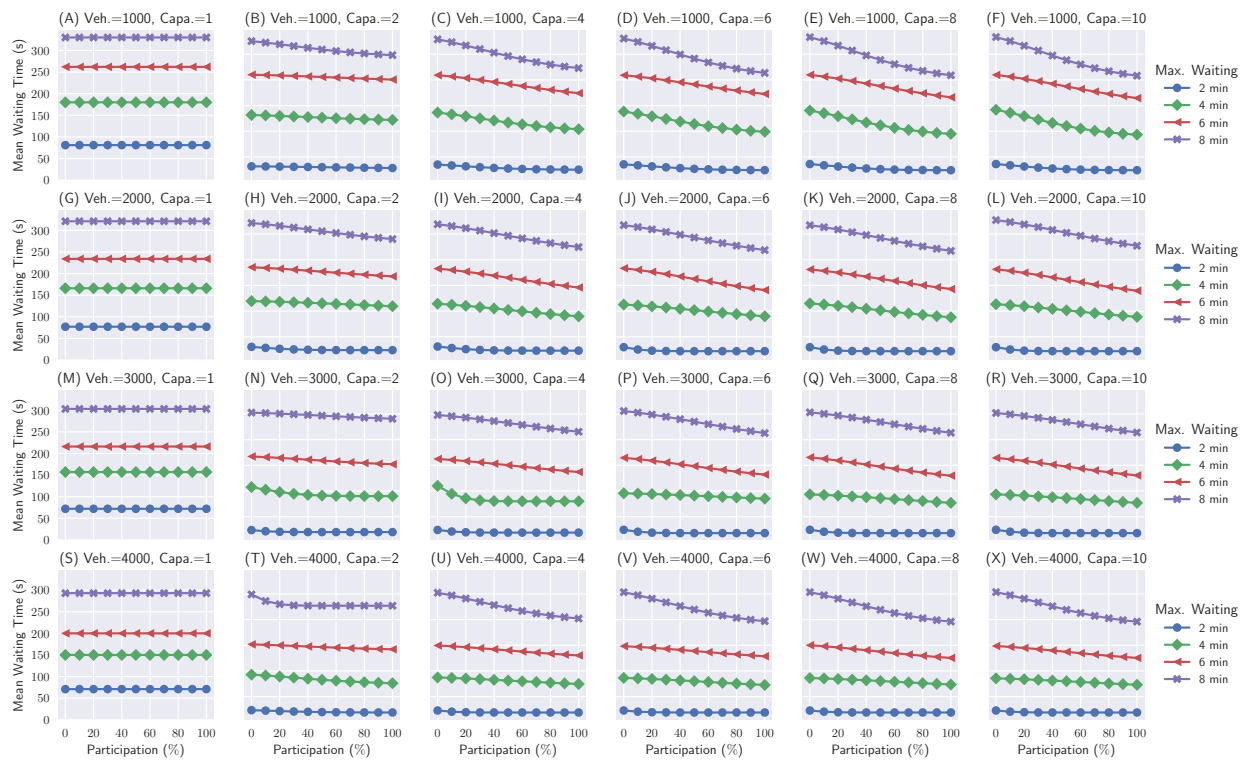


Figure 8: Mean waiting time for autonomous ridesharing systems with different configurations.

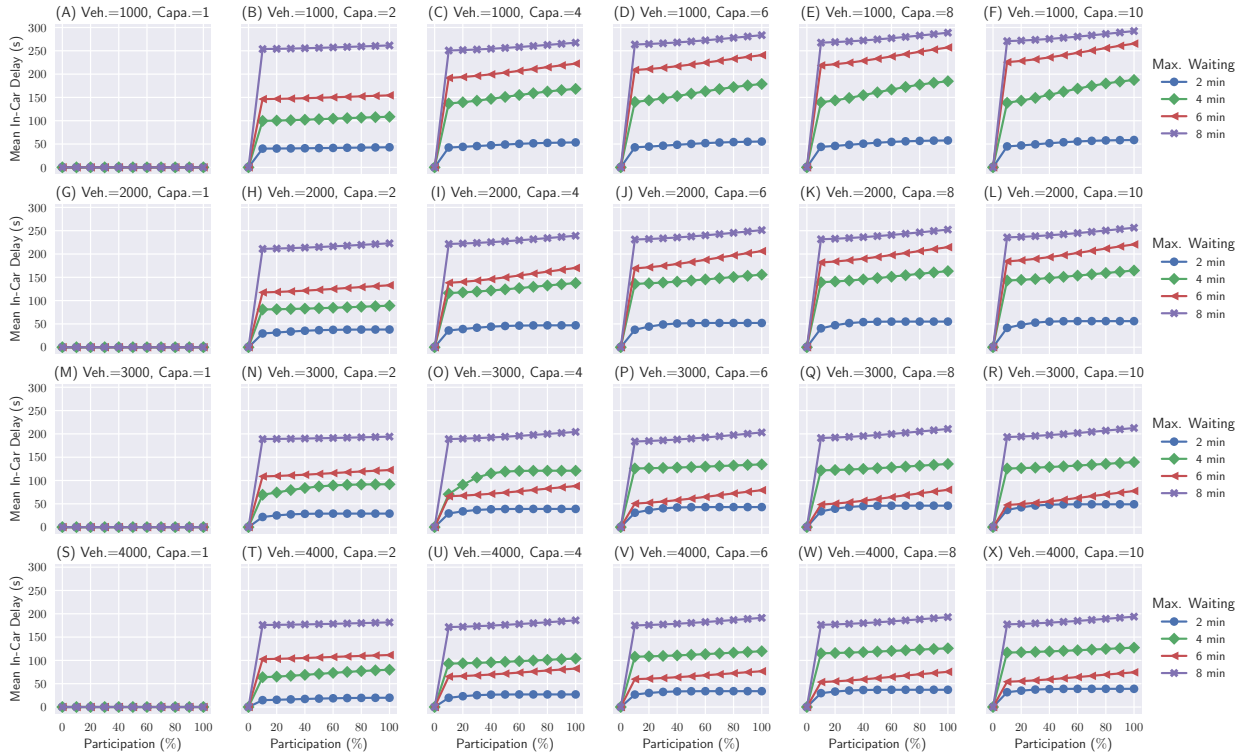


Figure 9: Mean in-car delay time for autonomous ridesharing systems with different configurations.

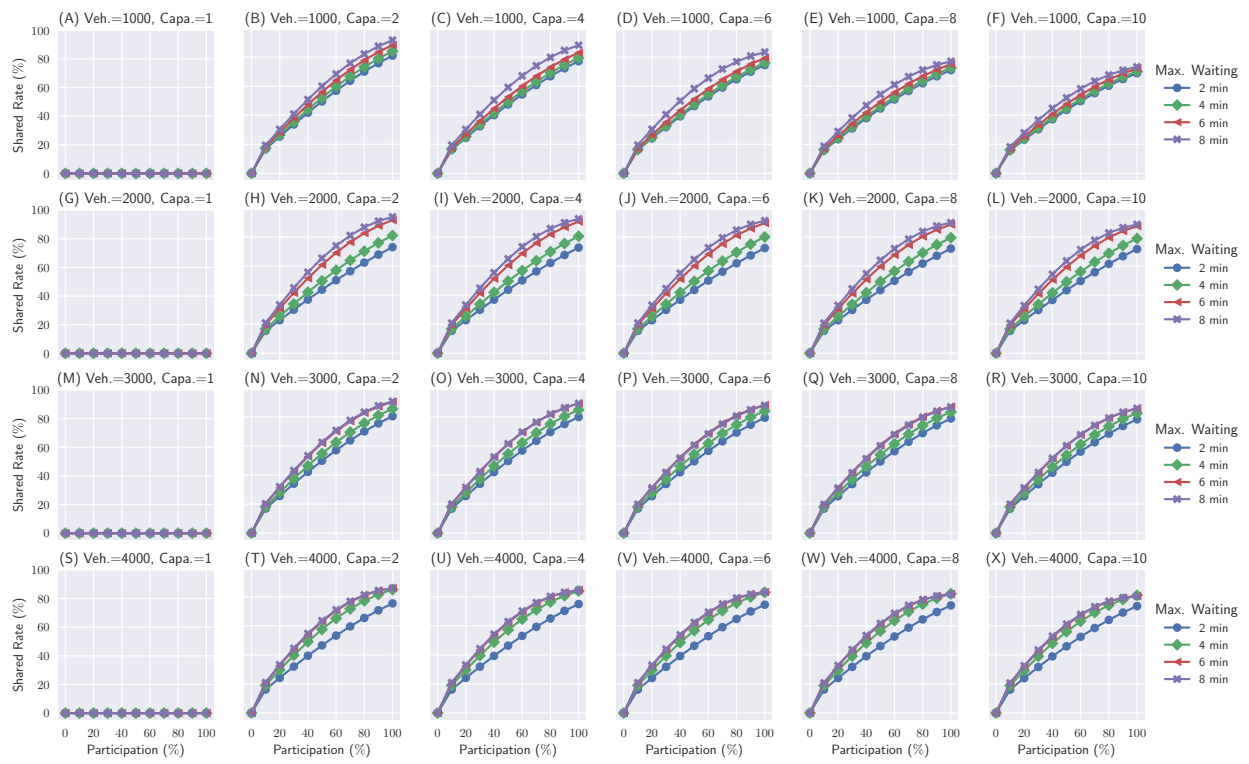


Figure 10: Percentage of shared rides for autonomous ridesharing systems with different configurations.

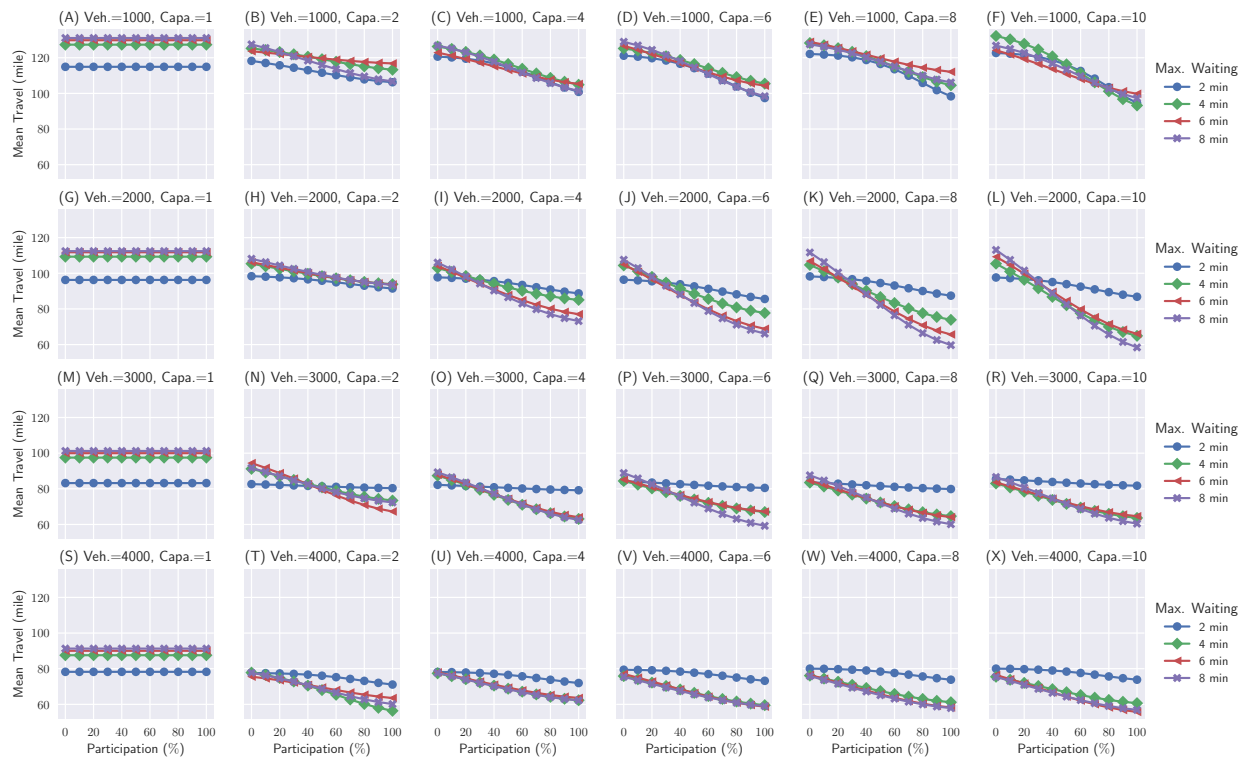


Figure 11: Average distance traveled by each vehicle during a single day for autonomous ridesharing systems with different configurations.