# A SIMULATION-BASED APPROACH TO COMPARE POLICIES AND STAKEHOLDERS' BEHAVIORS FOR THE RIDE-HAILING ASSIGNMENT PROBLEM

Ignacio Erazo

School of Industrial and Systems Engineering Georgia Institute of Technology 755 Ferst Dr NW Atlanta, GA 30332, USA Rodrigo De la Fuente

Industrial Engineering Department University of Concepción Edmundo Larenas 219 Concepción, Biobío 4030000, CHILE

## ABSTRACT

This study focused on the ride-hailing assignment problem, aiming to optimize drivers' behaviors with respect to simultaneous objectives such as maximizing service level, minimizing  $CO_2$  emissions, and minimizing riders' waiting times. Four different policies were proposed and tested with a real-world case study. With respect to current literature, we present a more realistic simulation model, capturing all characteristics of a ride-hailing system and using road networks to approximate real-time road conditions. Furthermore, it is the first work that tests the effects of different passengers' arrival conditions and analyzes the multiple objectives for different zones of a large city. Results suggest that different passengers' arrival conditions affect the four proposed policies nearly identically. Finally, the policy of drivers remaining static instead of driving while searching for passengers had the highest service level and lowest average distance per ride.

# **1 INTRODUCTION**

Transportation has always been important for the economy, as it increases productivity, connects different countries, and offers jobs, among others. In the past few years, technological advancements greatly changed the way people participate in the economy, becoming the breaking ground for a new type of "on-demand" economy. For transportation, it is translated to new companies like Uber, Didi Chuxing, or Lyft. They offer ride-hailing, ride-sharing, and even food delivery services.

These new companies have had a worldwide impact, propelling studies such as i) how to implement efficient routes for ride-sharing, ii) how to assign drivers and passengers in a ride-hailing system, or iii) how to compute fair prices that help to get a better balance of demand and supply. We focus on ride-hailing as it has a higher global impact; in particular, this work studies the effect that different policies and drivers' behaviors have in how to match cars and riders efficiently. The goal is to maximize the overall service level while minimizing the amount of miles driven per successful ride. Four different policies for the drivers' behaviors are proposed, described, and tested in a real-world case study that considered different passengers' arrival conditions.

With respect to the current literature, this study is the first to optimize drivers' behaviors over two different objectives and to examine the effect of different passengers' arrival conditions. Moreover, this paper also contributes to the field by proposing a realistic simulation model of a ride-hailing system. The model allows for the use of different road network structures, and computes statistics on a micro level, which helps to present service level and waiting time results for different areas of the road network, something that has not been done in previous works.

This paper is organized as follows: Section 2 summarizes the related work. Afterwards, Section 3 explains the methodology used, and Section 4 presents results of the case study. Finally, Section 5 presents the discussion, and in Section 6 the main conclusions are given.

## 2 RELATED WORK

Over the years, several problems related to ride-hailing services have been studied because of the paramount impact they have in modern societies. On a macro level, companies in that sector must 1) forecast demand and balance it with supply, 2) assign cars to pick up every passenger and 3) test and adapt strategies to increase the service level. As a consequence, researchers have focused on these areas using different methods and approaching the problems from several viewpoints.

## 2.1 Forecast Demand and Balance with Supply

With the growing presence of wireless connections, big data and tracking devices, ride-hailing companies have vast amounts of data available. These organizations use that information to improve the spatio-temporal forecast of demand, suggest optimal pick-up and drop-off points, and even identify suspicious accounts (in their mobile applications). Academia also focused on this topic, Moreira-Matias et al. (2013) forecasted demand using historical data, Miao et al. (2016) added spatio-temporally correlated demand and a supply model. Xu et al. (2020) studied the importance of pricing to balance supply and demand.

# 2.2 Assign which Car Will Pick up which Passenger

The ride-hailing problem has been widely studied under several perspectives. Souza et al. (2016) solved it as an assignment problem considering different strategies. Lowalekar et al. (2016) and Maciejewski et al. (2016) proposed two-stage stochastic optimization formulations to solve the general spatio-temporal matching problem, focusing on matching cars and passengers in an online setting. In addition, Alshamsi et al. (2009) and Glaschenko et al. (2009) proposed multi-agent simulation to solve this problem.

## 2.3 Test and Adapt Strategies to Increase the Service Level

Several researchers have considered taxi dispatch systems for idle drivers, with goals like increasing the profit of drivers or the probability of finding passengers. Zou et al. (2013) proposed a system creating routes with the goal of reducing the expected distance to pick up a passenger. Yamamoto et al. (2008) defined two strategies for cars without passengers and studied them. Finally, Lee et al. (2008) and Yuan et al. (2011) created recommendation systems for drivers to tell them where to wait for clients.

Another popular area of study is related to the different "behaviors" a driver can have, and how their actions can be optimized to improve the system's service as a whole. Hoque et al. (2012) used data collected in San Francisco to compute average taxi speeds, frequency of trips' durations, and their distribution in the city with the goal of helping drivers search for passengers. Li et al. (2011) concluded that "hunting" passengers by driving tends to be a better strategy than waiting for them. Henao and Marshall (2019) studied the profits received by drivers and concluded they have better incomes by parking off while waiting instead of driving, unless the waiting time for a new ride is considerably reduced (over 30%) by driving.

# **3** METHODOLOGY

## 3.1 Model Components

A ride-hailing system has three main actors: 1) drivers, 2) passengers, and 3) the transportation company. These actors are linked by their actions in the system, which can be considered as events. Drivers and passengers were included in the model as classes, with their respective attributes and procedures. The following events were considered: drivers entering and leaving the system, passengers entering the system, assignment of drivers to passengers, pick-up, and drop-off of the passengers. Besides that, a *Roaming* event was created in order to instantiate the different scenarios that are going to be compared. This event is used to represent the movements of a driver within the road network when he or she is not assigned to passengers (i.e., the driver is idle). It must be noted that the event *Passenger leaves the system* is not

Erazo and De la Fuente



Figure 1: Relationship between different instances.

implemented in the simulation as an event *per se*, because it is possible to perform those actions without having to execute an event. The relationships between different instances of the *Passenger*, *Driver*, and all the *Event* subclasses are shown in Figure 1.

The main task of matching drivers to passengers was executed by means of a helper function invoked at every arrival of passengers into the system. To do the assignment as in the real systems, information about drivers was stored in a list that had all available drivers, a dictionary having the information about the road network, and by using instance variables of the created objects (*Driver* and *Passenger* instances). Our matching algorithm finds the drivers (meeting requirements) that are closer to the rider and then asks sequentially if they accept the ride. If none accepts, the passenger leaves the system; otherwise, the rider decides if he accepts the ride or not based on the driver's estimated time arrival (ETA).

Moreover, several assumptions have been made: i) Every passenger only enters the system once. There is no loss of generality because every arrival in the system is treated independently. ii) The pricing and payment aspects are out of the scope of this work. It is assumed that the passenger entering the system will pay for the ride if a proper matching is found. iii) The company will always seek to match the passenger with the nearest driver that meets the requirements (time-wise). The matching algorithm is built on that premise. iv) Once a pick-up is scheduled, that driver is no longer available, so he can not be considered for subsequent arrivals until he drops off his assigned passenger. Clearly, ride-sharing is not allowed. v) If no driver that meets requirements is available at the arrival of a passenger, the passenger simply leaves the system. vi) A driver will leave the system only after his last ride or roaming event has finished. If a driver has a tentative leaving time at 7:00 PM, but he is scheduled to serve a ride that will last until 7:17 PM, he will leave the system after the ride.

## 3.2 Model Inputs

In order to use the model, it is necessary to provide some extra information about the different classes, events, and the system being simulated. The system inputs are the road network, the matching algorithm as well as the starting and ending time of the simulation. Driver inputs are their passenger capacity (assigned with a discrete distribution), their "luxury" attribute, the driver criteria to accept or reject rides, the arrival schedule of all drivers in the system, their tentative leaving schedule (a function of the number of successful rides), and their usual arrival locations. Finally, passenger inputs are their capacity and "luxury" requirements (both obtained via a discrete probability distribution), their arrival time, location and destination, and the passenger's criteria to accept or reject the ride after being matched to a specific driver.

## 3.3 Proposed Scenarios

As the main goal of this work is to test how different drivers' behaviors may affect the performance of a ride-hailing system, we considered the following scenarios:

• No-movement scenario: When idle, drivers wait patiently in their location, without moving. That applies to when they enter the system or after dropping off a passenger.

- Single-movement scenario: After entering the system or dropping off a passenger, the drivers start moving towards a specific location decided by them, using their experience and knowledge about the road network and the passengers' arrival behavior. Once they arrive at that destination, they wait patiently for their next ride, or until the moment they leave the system.
- Nearest-hotspot-movement scenario: The ride-hailing company uses its historical information and predictions to provide a list of the recommended locations where to wait for passengers. Once drivers arrive to the system or drop off a passenger, they go towards the nearest location in that list and then wait there.
- Coordinated-hotspot-movement scenario: The ride-hailing company uses its own information and the driver's location to route the drivers to specific waiting points. For this scenario, it is necessary to give the routing algorithm as input.

Finally, the routing algorithm for the coordinated hotspot scenario may be any kind of repositioning algorithm. For this study, we routed drivers depending on a coefficient based on the priority of each hotspot (a parameter) and the distances (in time units) between the driver and the hotspots. A threshold is defined using the parameter *MinimumCarsRouted* such that the hotspots having less than those drivers routed there are always served first.

## 3.4 Case Study

San Francisco is recognized as one of the most-innovative and progressive cities in the USA and it is also one of the five most-populous primary statistical areas in the country. In fact, *Uber* was launched with the name of "UberCabs" in June 2010, and by the end of 2016, more than 5,700 drivers work on peak hours for ride-hailing companies in San Francisco. They, in combination, drive more than 570,000 miles in the city every weekday, and finish more than 170,000 drives, which represents 15% of all intra-San Francisco vehicle trips (SFCTA 2017). Furthermore, reports have concluded that during the 2010–2016 period, ride-hailing companies have caused 47% of the increase of traveled miles and 55% of the average speed decline in the city (Marshall 2018).

### 3.4.1 Road Network

A buffer around the San Francisco city limits was utilized to delimit the extension of the road network required for this study. The network was extracted using QGIS software and properly modified to create a valid directed graph that consists of 11,372 nodes and 31,428 edges. Each edge was given its respective real-world distance in the horizontal plane by using a 70 x 70 meters *digital elevation model*. With regards to computing traversing times, the distance of each edge was adjusted to include the vertical (elevation) component by means of the following equation:  $Distance = \sqrt{HorizontalDistance^2 + VerticalDistance^2}$ ; and then the slope of the edge was computed as  $Slope = arctg(\frac{VerticalDistance}{HorizontalDistance})$ . The average speed (km/h) for each edge was assigned according to its classification, being the most

The average speed (km/h) for each edge was assigned according to its classification, being the most relevant: Motorway=65, Primary=45, Secondary=35, Tertiary=35, and Residential=25. Moreover, it has been shown that elevation change affects the travel speed because of the additional difficulty of moving uphill or downhill. Thus, each edge had its speed adjusted such that *Ad justedSpeed* = *AverageSpeed* ×  $e^{-0.016 \times Slope}$  (Verma et al. 2017). Finally, the time needed to traverse the edge was computed by computing the ratio between the distance and the adjusted speed.

#### 3.4.2 Passengers

Passengers entering the system required their respective arrival time, arrival node, and destination node considering the spatio-temporal behavior of their requests. For that, real data from SFCTA (2017) were used. Combining pick-up and drop-off data obtained from *Uber* and *Lyft* in 2016, the San Francisco area was divided into 976 medium-sized zones called *travel analysis zones (TAZ)*. Each TAZ was properly

Erazo and De la Fuente



Figure 2: Spatio-temporal comparison of arrivals in San Francisco's ride-hailing systems.

geo-referenced, and the average hourly number of pick-ups and drop-offs in those zones was computed. The information was averaged by hour and day of the week to generate a representative week. A small sample of pick-ups corresponding to these hourly averages is shown in Figure 2, where there are clear differences in the spatio-temporal distribution of passengers' arrivals.

First, for the passengers' temporal distribution, the information provided in SFCTA (2017) was used, as it detailed the hourly city-wide number of pick-ups for each day of the week. Arrivals during the week may be described as a non-homogeneous Poisson process. However, in real systems, correlation between arrivals may exist, and in general, they are not independent, e.g., they can be influenced by sports events, holidays, music festivals, to name a few. Thus, the method described by Nelson and Gerhardt (2011) was used to generate passengers requests. This method uses a stochastic process known as Markov-MECO, which can capture the first three moments for a distribution of interarrival times (Johnson and Taaffe 1989). Moreover, by using this method, it is possible to model processes with a higher or lower coefficient of variation than a regular Poisson process, and the correlation between arrivals can be considered. Finally, to simulate the operations of only one ride-hailing company, *Lyft* was used as a reference, and 36 % of the market share was considered.

With respect to the passengers' arrival and destination locations, SFCTA (2017) recorded hourly passenger arrival rates for each day of the week in each of the 976 TAZ zones within the city limits. This information was used to compute the probabilities of requesting a ride from a specific "5-zip zone" during a given hour. The centroid of each TAZ polygon from Figure 2 was assigned to its corresponding "5-zip zone". Then, adding up all the TAZ arrivals inside each zone, the hourly probabilities of both passenger pick-up and drop-off zones for each day of the week were obtained. Besides that, a dataset containing over 25,000 GPS traces from 589 taxi rides in San Francisco (Piorkowski et al. 2009) was used to get

the dynamics of the inter-"5-zip" movement that characterized a full trip. Thereby, trips within each zone could be classified by hour and day of the week, and matched to the spatial resolution of the pick-up and drop-off probabilities.

Next, we used previous reports to set the attributes for each passenger. According to data released by *Uber*, around 6% of their drives provide "luxury" services. We assumed that passengers will have a 6% chance of requiring the "luxury" attribute from drivers. Besides that, it has been reported that average and median waiting times for a driver in San Francisco are about three minutes. Thus, to compare scenarios a hard threshold of seven minutes as waiting time was chosen. Once passengers receive the pick-up time assigned to them, if they have to wait more than 7 minutes they will leave the system: otherwise they will accept the ride. Finally, the capacity required, in number of passengers, was set to 0.9 with a probability of being between 1 and 4 (with equal probabilities between those four integers) and 0.1 probability of being either five or six (which means a car with extra room would be required).

## 3.4.3 Drivers

According to SFCTA (2017) there are approximately 45,000 active drivers in the San Francisco area; however, most of them work part-time according to their preferred schedules. Considering the drivers' part-time behavior, a function was created to generate, for each driver, their entering and leaving times in the system, such as to resemble the distributions of the total number of drivers presented in SFCTA (2017).

With respect to to the arrival location of ride-hailing drivers to the system, no information is publicly available. However, as it is possible for them to log in the application at any site, it was decided that each driver would have a random arrival location (among all nodes) where he or she would enter the system. When idle, the driver's movement behavior is ruled by the scenario being tested. However, the *Single random roaming* scenario allows drivers for deciding their destination when performing a roaming trip. To accurately represent the actual decisions made by drivers, the GPS traces from Piorkowski et al. (2009) permitted to characterize the driving patterns followed by idle cabs depending on their location, time of the day, and day of the week.

In order to set the other inputs, previous reports from different sources were used. According to *Uber*, 10% of their drivers offer "luxury" services; thus, drivers were created with a 10% probability of having the "luxury" attribute. Besides that, it is stated that 16.4% of the total number of drivers have cars with extra capacity, so that probability was used for drivers having five or six seats in the simulation, and the rest of the drivers have a capacity of either three or four people. Each driver will accept every trip request with a time distance smaller to their minimum threshold time (a random uniform value between two and four minutes) and will reject every trip with a time distance exceeding their maximum threshold time (random uniform value between 19 and 21 minutes). For all trip requests with a time distance between their minimum and maximum threshold times, the driver would accept the ride with some probability, computed using a decreasing linear function. Finally, we decided to consider 3% of the TAZ areas as *Hotspots*. The priority of each *Hotspot* was just the ratio of the arrivals in that location, divided by the total number of arrivals.

### 4 RESULTS

First, we describe the validation process used to verify that the model was working properly. The riders' spatio-temporal patterns were compared to the real-world data used. Afterwards, the parameters used in the function creating the drivers' arrivals were tuned in order to minimize the difference between the induced drivers' behavior and the data provided by SFCTA (2017), which presented a distribution on the number of drivers in the San Francisco ride-hailing system according to the hour and day.

Next, the case study was performed on the four scenarios described in Section 3.3. The *Single movement* scenario is the closest to the real behavior of drivers, because it is based on GPS traces collected for the city during 2010 from 25,000 rides given by cab drivers in the era before ride-hailing companies. The

Erazo and De la Fuente



Figure 3: Service level per scenario and arrivals' conditions.

passengers' arrival times were generated using the Markov-MECO method described previously. To study the effect that different arrival conditions have on the four proposed scenarios, correlations of 0, 0.35 and 0.7 were used and coefficients of variation of 0.4, 1 and 2 were considered.

For each of the four scenarios and each of the nine different arrival conditions, ten replications of one week of operations of a ride-hailing company, with around 36 % of the market share, in the San Francisco area were simulated. Experiments were run on a *Dell* computer with *Intel core i7-7,700 HQ CPU* and 16 GB of RAM memory, taking around 18 hours for each replication (for the four scenarios), and a proper random seed was given to make the scenarios comparable.

## 4.1 Overall Service Level

Figure 3 presents the average service level for the ten replications of each scenario under the nine different arrival conditions. It is evident that all the tested scenarios have a service level above 95 % under all conditions. However, the *No-movement* and *Single-movement* scenarios have clearly a higher service level than the other two. The advantage appearing in those two scenarios is explained by the fact that demand tends to be concentrated in the downtown area, as was shown in Figure 2, so that hotspots are mostly located there, resulting in drivers ending up far away from non-dense arrival zones. This in turn implies that the *No-movement* and *Single-movement* strategies are better prepared to serve clients from areas outside downtown, as seen in Figure 5. The number of rides rejected by passengers in the hotspot scenarios are nearly triple the value for the first two scenarios. The difference in service level between the *No-movement* and *Single-movement* and the two hotspot policies is always between 2 % and 3 %, which is not insignificant as 2 % represents around 8,400 rides per week. It must also be noted that the *Nearest Hotspot* scenario has a higher service level than the *Coordinated Hotspot* scenario under all arrival conditions. Additionally, for all scenarios, a higher correlation in arrivals (with constant variance coefficient) produces a smaller service level, as it becomes harder to balance supply and demand when all arrivals are occurring in short intervals of time.

Figure 4 shows the average number of missed arrivals to the system (rides) per block of time during the week. Each block consists of three hours and the results for the 56 blocks, and the four policies are presented under the most "chaotic" arrival conditions and under the "Poisson" arrival conditions. All policies show the same trend, with higher values around the peak hours. It is also straightforward to

Erazo and De la Fuente



Figure 4: Missed rides in the city per block of time.

see that the "chaotic" arrival conditions put more pressure on the system, therefore causing more rides to be lost. The *No-movement* and *Single-movement* policies show a very close behavior, with the former outperforming slightly the latter in quiet hours, whereas the opposite is true in peak hours. Furthermore, the *Nearest Hotspot* and *Coordinated Hotspot* also have a very similar performance, and it can be seen that their difference in the quiet hours with respect to the other two models is lower than 200 arrivals per block. However, it can exceed 800 rides in the peak hours.

## 4.2 Overall Distance Driven per Ride

The *No-movement* scenario has the lowest distance driven per passenger under all arrival conditions, whereas the *Single-movement* scenario has the highest distance driven per passenger. In fact, the *No-movement* scenario has around 48 % reduction of distance driven with respect to the actual driver behavior (represented by the *Single-movement* behavior). For all arrival conditions, the *Nearest Hotspot* scenario has a smaller distance driven than the *Coordinated Hotspot* scenario by around 125 meters in average, and both scenarios have savings of 25 % with respect to the *Single-movement* behavior.

A higher correlation in arrivals (with constant variance coefficient) increases the distance driven by the hotspot scenarios. However, the impact in the *No-movement* and *Single-movement* scenarios is not significant. This is probably due to the fact that both hotspot scenarios have more trouble serving demand from areas outside the downtown, and that is further increased when the arrivals' correlation is higher.

### 4.3 Zone Analysis

Figure 5 shows the average service level, average waiting time, and average number of rides lost per TAZ on Fridays, between 6:00 PM and 9:00 PM (the busiest time of the week) for the four proposed policies. The passengers' arrival conditions are given by a correlation of 0.7 and variance coefficient 2, recreating the most "chaotic" environment. It is easy to see that under these conditions that the *No-movement* scenario is the one that provides the highest performance, as it has a higher service level and less waiting time in most of the TAZ. Next, the *Single-movement* policy presents a similar behavior near downtown, but performance declines outside of it. Both hotspot policies provide a similar performance in downtown, but the waiting times are higher and the service level decreases outside of it.



Figure 5: Average service level, average waiting time, and average missed rides per TAZ.

Comparing the *Single-movement* and *Nearest Hotspot* policies, the former outperforms the latter in all the TAZs that are outside downtown. The difference between the service level means is 5.37 %, but the first quartiles differ in 19 % (with respect to total arrivals). With respect to the average waiting times, similar results are obtained, with differences of nearly 40 seconds, whereas the difference of the maximum average waiting time for a TAZ is around 55 seconds.

### 5 DISCUSSION

The *No-movement* scenario compared favorably versus all the other scenarios in driving efficiency, and its service level nearly matched the highest one obtained by the *Single-movement* scenario. Differences in distances driven between these two scenarios were relevant, as the static scenario resulted in a decrease of 48 %. In fact, the drivers' real-world behavior is closely represented by the *Single-movement* scenario, as the roaming events followed experienced cab drivers' GPS traces. Assuming that it is possible to change that behavior in the long term to the one proposed in the *No-movement* scenario, a reduction in driving distances of around 1.73 miles per ride would be obtained, while maximizing the service level. This helps to save more than 11 million dollars, and 57 million pounds of  $CO_2$  would not be released to the atmosphere per year (savings computed considering the average annual gasoline price for San Francisco in 2019 of 3.76 dollars per gallon, 7% of cars as electric, a performance of 35 miles per gallon of gasoline for every gasoline-fueled car, and 19.64 pounds of  $CO_2$  being released to the atmosphere after burning one gallon of gasoline).

In an ideal world, it would be advisable for drivers to wait for passengers in their respective locations, as the system and society as a whole would receive the benefits (less congestion, less pollution, etc). However, the solution is not plausible in the short and medium terms because of a lack of parking lots for this use in the city and because of the drivers' desire for a maximization of their revenues instead of a maximization of benefits for the society. Thus, the static scenario should only be used as a tool to compare other scenarios versus an ideal behavior. Considering the other three proposed scenarios, the *Single-movement* scenario has the largest service level, which is explained by the fact that drivers know where to go in order to maximize their rides. However, as drivers in ride-hailing companies usually have less experience, it can be expected that the results for the real ride-hailing system may be worse.

As seen in results, the *Nearest Hotspot* outperforms the *Coordinated Hotspot* strategy, and when compared to the *Single-movement* policy, it would help to save 7.9 million dollars per year and 38 million pounds of *CO*<sub>2</sub> released in the atmosphere. That being said, the *Nearest Hotspot* strategy seems to be the best at balancing efficiency and service level. However, when the demand is harder to satisfy, that policy is prone to do very well in the downtown area, but to fail at covering the needs of riders arriving far away from the hotspots. Besides that, when comparing to the *Single-movement* policy, the riders need to wait approximately 40 seconds more for a driver, which may not seem like a significant difference in absolute terms, but shows the knowledge of taxi drivers, as their movements seem optimized to serve the highest amount of passengers and arrive to them quickly. Finally, during the peak hours, the difference in missed rides between the *Nearest Hotspot* and *Single-movement* policies is amplified. Thus, any strategy seeking to implement the former policy should plan to react accordingly.

The implementation of the *Nearest Hotspot* policy is plausible, and would bring several economical and environmental benefits to the community. However, as it was evidenced by the comparison during the most "chaotic" conditions, it has some negative trade-offs that must be addressed. First, this strategy should be executed with funds to build parking lots for drivers. In peak hours, around 6,500 drivers work in San Francisco; considering a typical parking space of 5 by 2.4 meters. If half of the drivers are idle, a total of 39,000 square meters of parking lots in the hotspot locations would be required. Also, the ride-hailing companies should help using their data and developing procedures to compute the best hotspot locations. Moreover, other means of transportation should be adjusted to provide support in the areas that are not well covered by the ride-hailing system (i.e., outside downtown), or other incentives could be used to bring more drivers to serve those areas in the moments of greater needs (i.e., pricing, benefits, ...).

Finally, while the results obtained are valid for this case study, we assume that some of them can also be generalized to other ride-hailing systems. In particular, the reduction in  $CO_2$  emissions by the "hotspot" strategies and the *No-movement* scenario should be applicable to every ride-hailing system, but also the increasing difficulty on implementing such strategies. Also, the fact that different arrival conditions seem to have a similar effect on every tested scenario is something that should be possible to generalize, because it only represents an increase on the stress the system receives, and that holds across cities. Also, it is clear that some other results depend on the spatio-temporal behavior of drivers and riders, the road network, as well as the matching and repositioning algorithms. But, the presented framework is useful to get results under any change of the aforementioned parameters.

#### **6** CONCLUSIONS

One week of operations of a ride-hailing company in the city of San Francisco with 36 % of the market share was simulated. Drivers and passengers were created using historical data considering the spatiotemporal behavior that they exhibit. Besides that, nine different set of parameters were used to produce the passengers' arrivals, which allowed us not only to compare the proposed scenarios but also to understand the effect that the variation coefficient and correlation of passenger arrivals have in the performance of the ride-hailing system. Results suggest that the effects of the passengers' arrival conditions are similar on the performance of the system for all the policies proposed and for all conditions. This makes sense, because the changes on the passengers' arrival conditions increase or diminish the stress received by the ride-hailing system, but as all proposed strategies have a similar (and very high) service level, it is expected that they are similarly equipped to handle the variations in the system's dynamics. Moreover, if drivers remained static while waiting for passengers, instead of driving while searching for them, the service level would exceed 98 % and no other strategy would surpass that service level. Furthermore, when compared to the actual situation, the distances driven by drivers would be reduced by 48 %. More than 11 million dollars would be saved per year in gasoline, and 56 million pounds of  $CO_2$  would not be released into the atmosphere. Nevertheless, it is not realistic to follow that strategy in the short term, because there is no parking space for drivers, particularly in downtown. Other scenarios tested proved to provide a better balance between service level and driving distance. However, switching to more-efficient systems like the Nearest Hotspot scenario requires help from the community and ride-hailing companies, in addition to investments in parking spaces, and some adjustments of the whole transportation system to be prepared for the period with higher demand. Therefore, it should be regarded a medium-term project.

Future work should be performed with the goal of searching for other ways to enhance the performance of the system. This may be done by introducing different variants to match drivers with passengers, for example, cars changing their assigned passengers while going to pick them up or conducting as the matching between drivers and passengers in fixed time intervals to obtain more-optimal pairings, among other strategies. Additionally, other scenarios may be tested, like drivers being assigned to zones or giving them a fixed schedule depending on the projected demand. Finally, pricing aspects could be considered to understand their effects on the ride-hailing system and their effect on the drivers' behaviors as well as to improve the efficiency of the ride-hailing system in less-served areas.

#### REFERENCES

- Alshamsi, A., S. Abdallah, and I. Rahwan. 2009. "Multiagent Self-Organization for a Taxi Dispatch System". In Proceedings of the 8th International Conference on Autonomous Agents and Multiagent Systems, edited by C. Sierra and C. Castelfranchi, 21–28. Richland, South Carolina: International Foundation for Autonomous Agents and Multiagent Systems.
- Glaschenko, A., A. Ivaschenko, G. Rzevski, and P. Skobelev. 2009. "Multi-Agent Real Time Scheduling System for Taxi Companies". In *Proceedings of the 8th International Conference on Autonomous Agents and Multiagent Systems*, edited by C. Sierra and C. Castelfranchi, 29–36. Richland, South Carolina: International Foundation for Autonomous Agents and Multiagent Systems.
- Henao, A., and W. E. Marshall. 2019. "An Analysis of the Individual Economics of Ride-Hailing Drivers". Transportation Research Part A: Policy and Practice 130:440–451.

- Hoque, M. A., X. Hong, and B. Dixon. 2012. "Analysis of Mobility Patterns for Urban Taxi Cabs". In 2012 International Conference on Computing, Networking and Communications (ICNC), edited by V. Tarokh and K. Makki, 756–760. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Johnson, M. A., and M. R. Taaffe. 1989. "Matching Moments to Phase Distributions: Mixtures of Erlang Distributions of Common Order". *Journal Communications in Statistics. Stochastic Models* 5:711–743.
- Lee, J., I. Shin, and G.-L. Park. 2008. "Analysis of the Passenger Pick-Up Pattern for Taxi Location Recommendation". In 2008 Fourth International Conference on Networked Computing and Advanced Information Management, edited by D. Delen, 199–204. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Li, B., D. Zhang, L. Sun, C. Chen, S. Li, G. Qi, and Q. Yang. 2011. "Hunting or Waiting? Discovering Passenger-Finding Strategies from a Large-Scale Real-World Taxi Dataset". In 2011 IEEE International Conference on Pervasive Computing and Communications Workshops (PERCOM Workshops), edited by D. J. Cook and J. Liu, 62–68. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Lowalekar, M., P. Varakantham, and P. Jaillet. 2016. "Online Spatio-Temporal Matching in Stochastic and Dynamic Domains". In AAAI'16 Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, edited by D. Leake and J. Lester, 3271–3277. Cambridge, Massachussets: Association for the Advancement of Artificial Intelligence.
- Maciejewski, M., J. Bischoff, and K. Nagel. 2016. "An Assignment-Based Approach to Efficient Real-Time City-Scale Taxi Dispatching". *IEEE Intelligent Systems* 31:68–77.
- Marshall, A. 2018. "Uber and Lyft Made Traffic Worse in San Francisco. But Its Complicated". WIRED Transportation, October 16<sup>th</sup>. https://www.wired.com/story/uber-lyft-san-francisco-traffic-report/, accessed 7<sup>th</sup> April 2021.
- Miao, F., S. Lin, S. Munir, J. A. Stankovic, H. Huang, D. Zhang, T. He, and G. J. Pappas. 2016. "Taxi Dispatch With Real-Time Sensing Data in Metropolitan Areas: A Receding Horizon Control Approach". *IEEE Transactions on Automation Science* and Engineering 13(2):463–478.
- Moreira-Matias, L., J. Gama, M. Ferreira, J. Mendes-Moreira, and L. Damas. 2013. "On Predicting the Taxi-Passenger Demand: A Real-Time Approach". In *EPIA 2013: Progress in Artificial Intelligence*, edited by L. M. Correia, L. P. Reis, and J. M. Cascalho, 54–65. Berlin Heidelberg: Springer.
- Nelson, B., and I. Gerhardt. 2011. "Modelling and Simulating Non-Stationary Arrival Processes to Facilitate Analysis". *Journal* of Simulation 5:3–8.
- Piorkowski, M., N. Sarafijanovic-Djukic, and M. Grossglauser. 2009. "CRAWDAD Dataset epfl/mobility". https://crawdad.org/ epfl/mobility/20090224/cab, accessed 10<sup>th</sup> May 2020.
- SFCTA 2017. San Francisco County Transportation Authority, "TNCs Today", https://www.sfcta.org/projects/tncs-today, accessed 7<sup>th</sup> April 2021.
- Souza, M. P., A. A. M. de Oliveira, M. de Arruda Pereira, F. A. L. Reis, P. E. M. Almeida, E. J. Silva, and D. S. Crepalde. 2016. "Optimization of Taxi Cabs Assignment Using a Geographical Location-Based System in Distinct Offer and Demand Scenarios". *Brazilian Journal of Cartography* 68:1143–1155.
- Verma, P., K. Kurar, and S. Saram. 2017. "Road Network Impedance Factor Modelling Based on Slope and Curvature of the Road.". International Journal of Advanced Remote Sensing and GIS 6:2274–2280.
- Xu, Z., Y. Yin, and J. Ye. 2020. "On the Supply Curve of Ride-hailing Systems". Transportation Research Part B: Methodological 132:29-43.
- Yamamoto, K., K. Uesugi, and T. Watanabe. 2008. "Adaptive Routing of Cruising Taxis by Mutual Exchange of Pathways". In KES 2008: Knowledge-Based Intelligent Information and Engineering Systems, edited by I. Lovrek, R. J. Howlett, and L. C. Jain, 559–566. Berlin Heidelberg: Springer.
- Yuan, J., Y. Zheng, L. Zhang, X. Xie, and G. Sun. 2011. "Where to Find My Next Passenger?". In *Proceedings of the 13th International Conference on Ubiquitous Computing*, edited by J. Landay and Y. Shi, 109–118. New York, NY: Association for Computing Machinery.
- Zou, Q., G. Xue, Y. Luo, J. Yu, and H. Zhu. 2013. "A Novel Taxi Dispatch System for Smart City". In DAPI 2013: Distributed, Ambient, and Pervasive Interactions, edited by N. Streitz and C. Stephanidis, 326–335. Berlin Heidelberg: Springer.

#### **AUTHOR BIOGRAPHIES**

**IGNACIO ERAZO** is currently pursuing his Ph.D. in Operations Research at the Georgia Institute of Technology. His research interests include the development of efficient optimization algorithms and heuristics for applications in logistics and also large-scale simulation-optimization procedures. His email address is iien3@gatech.edu.

**RODRIGO DE LA FUENTE** is an assistant professor at the University of Concepcion, and received his Ph.D. in Industrial Engineering from the North Carolina State University. He is interested in large-scale simulation, stochastic processes, and bridging simulation and machine learning. He email address is radelafu@ncsu.edu.