UTTRI

Research Report

Vehicle for Hire Bylaw Review

Report 6: Analysis of Network Impacts: PTC Trip Chaining

Francisco Calderón and Eric J. Miller

May 2019

UTTRI Technical Support for the City of Toronto Vehicle for Hire Bylaw Review

Report No: 6

Analysis of Network Impacts: PTC Trip Chaining

Francisco Calderón, M.Eng.CEM.

Eric J. Miller, Ph.D.

May 18, 2019

UNIVERSITY OF TORONTO
FACULTY OF APPLIED SCIENCE & ENGINEERING
Transportation Research Institute

UTTRI

TABLE OF CONTENTS

List of Tables **Page No.**

List of Figures **Page No.**

1. INTRODUCTION

This technical report presents the work undertaken by the University of Toronto Transportation Research Institute (UTTRI) to model trip chaining of vehicles operating for Private Transportation Companies (PTCs) in support of the City of Toronto's *Vehicle for Hire Bylaw Review*.

The main purpose of this component of the project is to assess network impacts of PTC operations. The data provided to study these impacts consists of the entirety of trips reported by ridehailing service providers from the period of September 1, 2016 to December 31, 2018. Even though these data are comprehensive and extremely useful for analysis, they do not encompass the entirety of ridehailing vehicular movements. Ridehailing vehicles further occupy the network while being en-route to pick-up passengers, as well as when idling or waiting to be paired with the next passenger. Hence, identifying and accounting for these additional vehicle "states" becomes a critical task towards a comprehensive assessment of PTC network impacts.

To align this report´s terminology with its counterpart in the Vehicles-For-Hire Bylaw (City of Toronto, 2016), the mentioned vehicle "states" are defined by three periods. Namely:

- Period 1 (idling): "total time a PTC driver had activated or was logged into a PTC Platform and available to receive or accept requests to provide passenger transportation service".
- Period 2 (en-route): "total time elapsing between the time a passenger request for transportation is accepted by a PTC driver and the arrival of the PTC driver at the passenger's pick up location".
- Period 3 (in-service): "total time elapsing between the time that a PTC driver picks up a passenger(s) until the time the passenger(s) has arrived at their destination(s)".

To quantify the amount of time spent and distance driven by ridehailing vehicles while being in Periods 2 and 3, Project Task 4.2 aims to identify vehicle trip chains, which would provide information related to the periods when vehicles are not transporting passengers. It must be noted that trip chaining is normally not an operational task explicitly performed by service providers, but it rather is the outcome of the dynamics of within-day service provision processes. A brief background on ridehailing operations is presented in Section 2 below to further develop this argument.

As acknowledged in the Project Charter document, undertaking a trip chaining analysis (Task 4.2) involves a high degree of uncertainty and complexity, being contingent upon the availability of various data sets, particularly unique driver identifiers and full path "GPS breadcrumbs" of each PTC VFH trip – the latter categorized by its different periods (1,2, or 3). Given the absence of these data, a prototype model has been developed, built upon existing data (observed demand for ridehailing trips) and an extensive conceptual understanding of ridehailing fundamental operational tasks. Furthermore, it must be acknowledged that the prototype model circumvents the lack of data by endogenizing ridehailing operational processes and vehicles (and its drivers) as agents. In this context, calibration and validation as envisioned in Task 4.3 become practically unfeasible since the prototype model can only provide approximate estimates. Nonetheless, a strategic modelling decision consisted of leaving out wait time variables in the dataset to be used for assessment of the quality of the results product of the modelling efforts undertaken in this report, through comparison of simulated versus observed wait time distributions.

The rest of the report is organized as follows. Section 2 provides background on ridehailing operations. Section 3 describes and analyzes key variables from a modelling perspective. Section 4 documents the development of a prototype model for ridehailing service provision and operations. Section 5 presents and discusses model results. Section 6 concludes the report with discussion on the applicability and limitations of the model and the insights it can provide for the Vehicle for Hire Bylaw Review.

2. Background on Ridehailing Operations

As argued in Calderón and Miller (2019c), the most important operational task performed by ridehailing service providers is arguably to match user trip requests and vehicles. By definition, a ridehailing service can be considered as a platform that facilitates interactions between users and drivers; hence, matching becomes a core component of any ridehailing mobility service. For a detailed discussion on ridehailing mobility services, refer to Calderón and Miller (2019b).

In theory, the matching operational task of ridehailing service providers could be cast as several variants of mathematical optimization problems within the broad class of vehicle routing problems (VRP). For instance, Doerner and Salazar-González (2014) define the conventional Dial-A-Ride-Problem as a VRP with precedence and coupling constraints. It must be noted that these optimization problems fall within the scope of NP-Hard problems, hence these are rarely deployed by service providers for practical applications and real-life operations of ridehailing services (Syed, Irina, & Bogenberger, 2019). Furthermore, a critical challenge arises from the on-demand, responsive nature of ridehailing, which implies that demand for the whole day is not known to service providers and trip requests arrive dynamically throughout the day. Optimization then becomes a moving target, pushing the problem towards dynamic programming approaches, which computational complexity is identified to be $O(n^2 3^n)$ when the problem is cast a single vehicle pick-up and delivery problem (Desrosiers, Dumas, Solomon, & Soumis, 1995). Note that the optimization problem so far is concerned with a single vehicle and single occupant, hence it becomes even more complex if considering multiple vehicles and shared ridehailing operations.

As an alternative, explicitly modelling ridehailing fundamental service operational tasks and driver activity offers a pragmatic and efficient solution to the problem. This can be achieved by following the guidelines established in a conceptual framework for modelling mobility services under the agent-based modelling paradigm, proposed by Calderón and Miller (2019a). To begin with, a ridehailing mobility service involves three actors or agent classes:

- Users that generate trip requests. Note that, an explicit representation of user agents is not necessary for the purpose of this study because trip requests are already observed in the data.
- Vehicles (in this case human-driven) that can fulfill trip requests and make decisions about their activity in the system.

- A service provider that undertakes operational tasks to run the service and acts as a platform to facilitate and manage interactions between users and vehicles.

At a minimum, modelling human-driven vehicle agents should capture their decisions of becoming active/inactive in the system throughout the day. These decisions are arguably dependant on the number of hours worked in a day, cumulative earnings collected, spatial distribution of demand at a given point in time, and the (dynamic) pricing level at a given point in time. More elaborate vehicle activity models can include decisions about where to locate/relocate to provide service, returning to activity or number of working shifts over a day, driver behaviour while waiting for the next customer (idling), etc.

On the other hand, a service provider agent can be modelled by a given combination of sub-models of operational tasks such as: matching, rebalancing and pricing. In a nutshell, matching refers to the process of pairing user trip requests to available vehicles; rebalancing is concerned with maintaining an adequate spatial distribution of the fleet through strategic vehicle relocation; and pricing is typically related to dynamic pricing mechanisms such as surge pricing (Uber, 2019b) or prime time pricing (Lyft, 2018b) – albeit some services operate with fixed pricing schemes[1].

All things considered, this report documents the development of a prototype ridehailing operations model, centred around service provider and driver agents. It should be noted that there is a high degree of uncertainty about how ridehailing providers perform rebalancing tasks or if they even do so; moreover, data limitations imply that modelling rebalancing operations (and a pricing mechanism as well) are inevitably out of reach. Nonetheless, the results of this project suggest that modelling matching can provide a very reasonable approximation of ridehailing service provision.

As a first step, the next section introduces descriptive statistics and focused analyses of key variables from the dataset provided for this study.

3. Key Variables

The data available consists of every trip reported by ridehailing service providers from the period of September 7, 2016 to December 31, 2018. This section explores these data as a first step towards constructing a model of PTC operations. This analysis focusses on the period from September 7, 2016 to December 31, 2016 so that they can be directly linked to the Transportation Tomorrow Survey 2016 (DMG, 2016) and the GTAModel V4.1 (TMG, 2018).

3.1. Trip Status

The dataset includes trips that were initially requested by users but were then cancelled either by users or drivers, as shown in Figure 1. Driver rejection rates are low (5.4%), which are in line with desired operational levels that aim for lower than 10%, usually enforced through driver reward systems (Lyft, 2018a; Uber, 2019a). Conversely, the rate of passenger rejection reaches 17.4%, which is quite high. The reasons for this phenomenon are unknown, and further analyses about

[1] Facedrive (https://www.facedrive.com/); Instaryde (https://www.instaryde.com/).

trip rejection will be conducted in future work, exploring spatiotemporal features and zone-specific average sociodemographic characteristics.

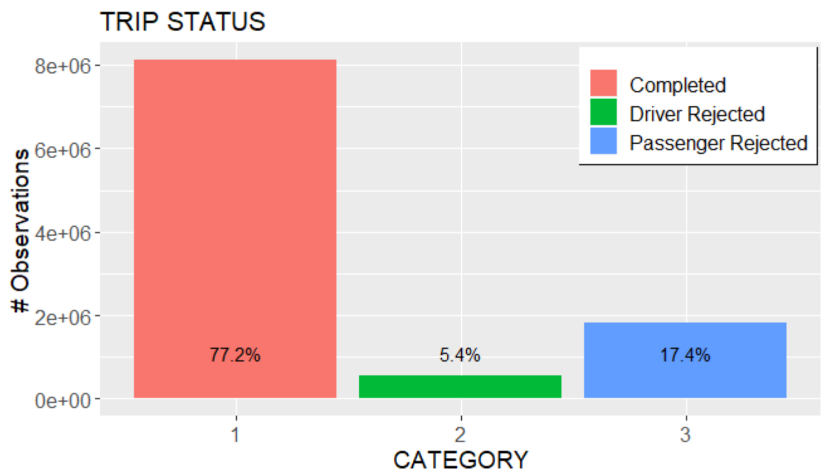

Figure 1: Classification of trips by their status

The remainder of this report considers completed trips only.

3.2. Type of Service Provided and Number of Passengers

Ridehailing service providers offer several service variants to users, yet 96% of all trips correspond to UberX (non-shared rides) and Uberpool, as shown in Figure 2. For model development, all Uber X trips were considered, along with a subset of Uberpool trips consisting only of the first reported Uberpool trip, which is treated as if it was a regular trip starting from its specified origin to its destination. The rationale for this decision is due to the added modelling complexity imposed by shared services, which becomes impractical given the current data limitations.

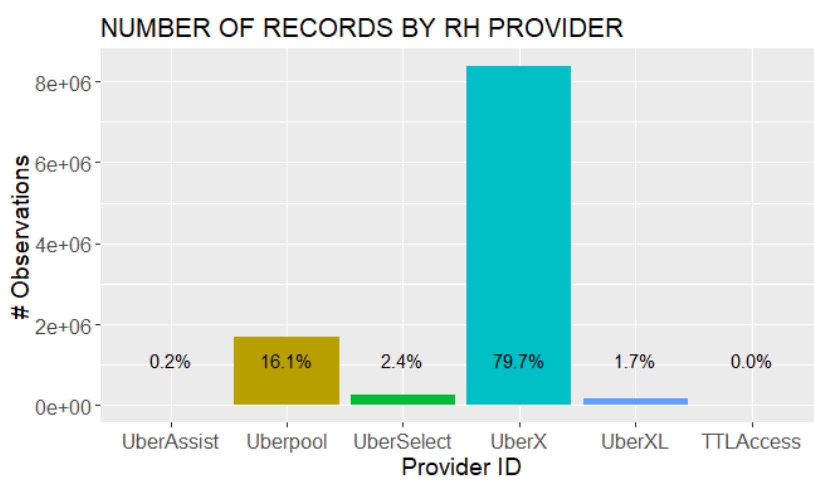

Figure 2: Classification of trips by type of service provided

The distribution of number of passengers per trip is shown in Figure 3, which clearly indicates a predominance of single-occupancy trips. Aside from this, an interesting finding suggests that this distribution is, in principle, inconsistent with the percentage of trips reported as Uberpool. Namely, while 16.1% of trips are reported to be pooled (shared), only 5.3% of trips have more than one passenger. The only feasible explanation that could be formulated is that users might have requested an Uberpool trip, yet no other user joined the trip over the course of it.

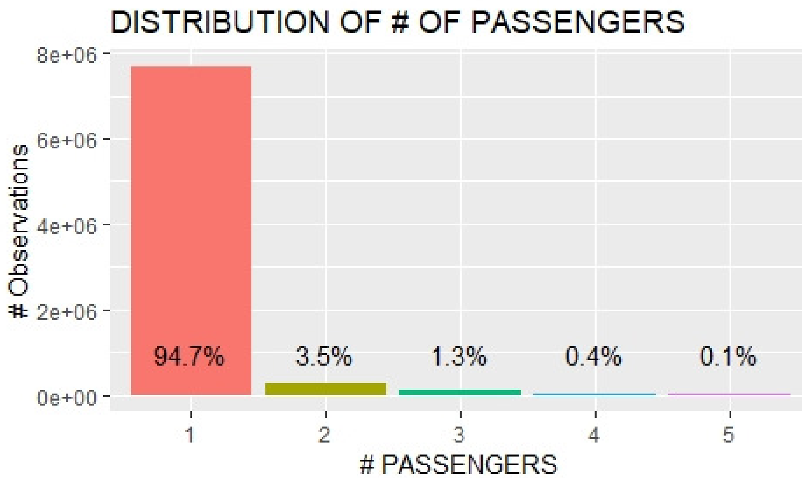

Figure 3: Distribution of number of passengers per trip

Note that the dataset includes records with up to 16 passengers, while records with more than 6 passengers represent only 0.04% of total trips. Thus, these records were excluded from the distribution in Figure 3.

3.3. Trip Request Times

As established in Report 2 of the overall project report series, temporal patterns of trip start times are markedly different among weekdays and weekends, hence, the models developed in this report are tested both for weekdays and weekend days.

A second important consideration regarding trip temporal patterns is related to temporal granularity. Following the rationale proposed in Calderón and Miller (2019a), the most important operational task performed by ridehailing service providers is matching users and vehicles, hence, the selection of time intervals to perform matching is a critical operational and modelling decision. Specifically, longer time steps provide more trip requests to match and hence the opportunity to be more efficient in vehicle allocation, however at the cost of imposing longer wait times for users. On the other hand, shorter time steps provide fewer trip requests that allow for faster computation and short wait times, yet at the expense of efficiency because trade-offs in matching are much more limited. Figure 4 depicts the temporal granularity of trip requests by different time step sizes. In this figure it can be seen that temporal patterns tend to become uniform as time step size decreases.

Considering the patterns shown in Figure 4, a time step of 1 minute was initially chosen for testing, deeming it as a value likely to be used by a service provider that aims to attain a balance among efficiency and service quality. A 1-minute time step is also adopted in similar research efforts on modelling on-demand services (Alonso-mora et al., 2017; Liu, Bansal, Daziano, & Samaranayake, 2018). Real-life operations, however, are usually not aligned with theoretical research experiments, as evidenced in this study, since a 5-minute time interval produces model results that reflect observed data more closely. This is discussed in more detail in Section 5, while the Appendix documents time interval testing in full detail.

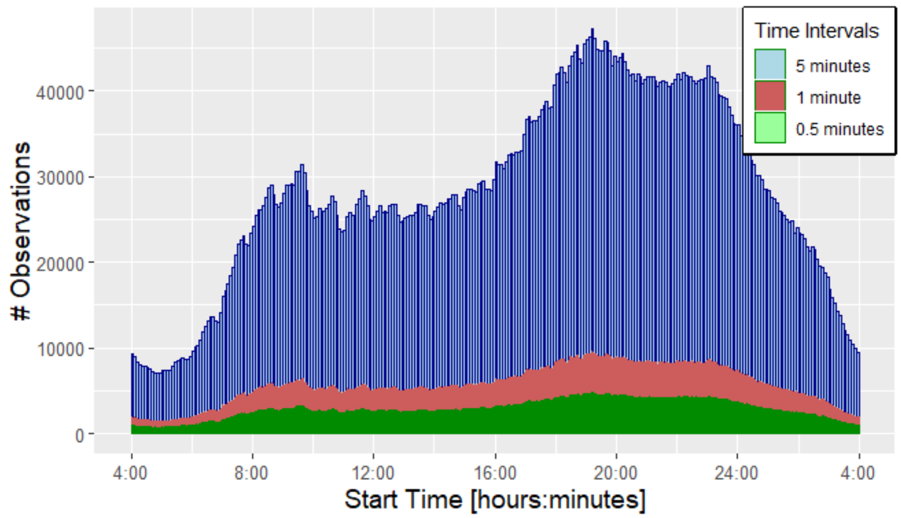

Figure 4: Temporal distribution of trips by different levels of granularity

3.4. Driver Wait Time

This variable represents the time drivers waited to be matched with a passenger. Figure 5a shows that driver wait times are predominantly zero minutes (87.91% of the records), whereas the maximum wait time is 60 minutes. Unfortunately, temporal disaggregation to the minute is too crude to sufficiently understand whether driver wait times are taking place in the range of seconds or are actually zero – which would imply practically instantaneous matching with trip requests.

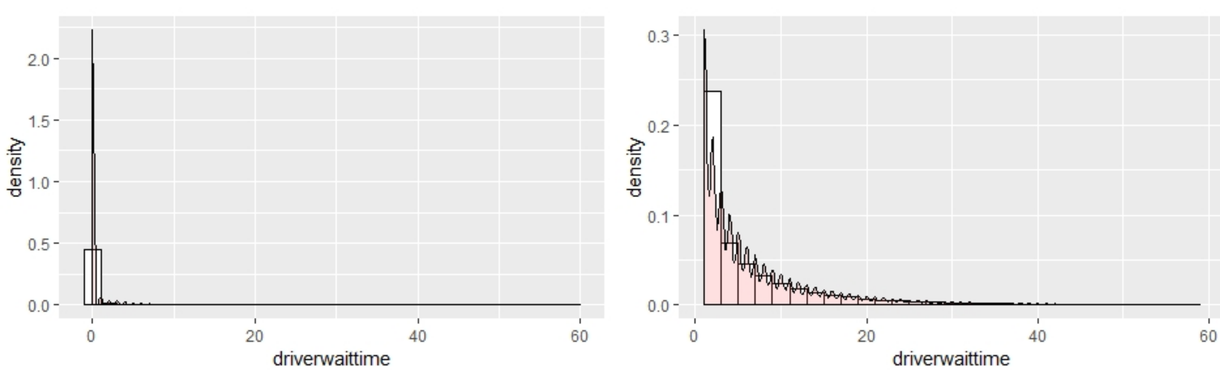

Figure 5: a) Density function of driver wait times [left]; b) Density function of driver wait times greater than zero [right]

To better visualize this variable, Figure 5b only depicts wait times that are greater than zero. Note that these wait times are highly likely to take place in oversupply conditions, while vehicles are idling or cruising and waiting to be allocated trip requests.

3.5. Elapsed Time

This variable represents user wait times, defined as the time elapsed since users request a trip and a vehicle picks them up. Figure 6a shows the distribution of user wait times, which ranges from 0 to 200 minutes, whereas Figure 6b shows the distribution of wait times less than 25 minutes (99.86% of the total records).

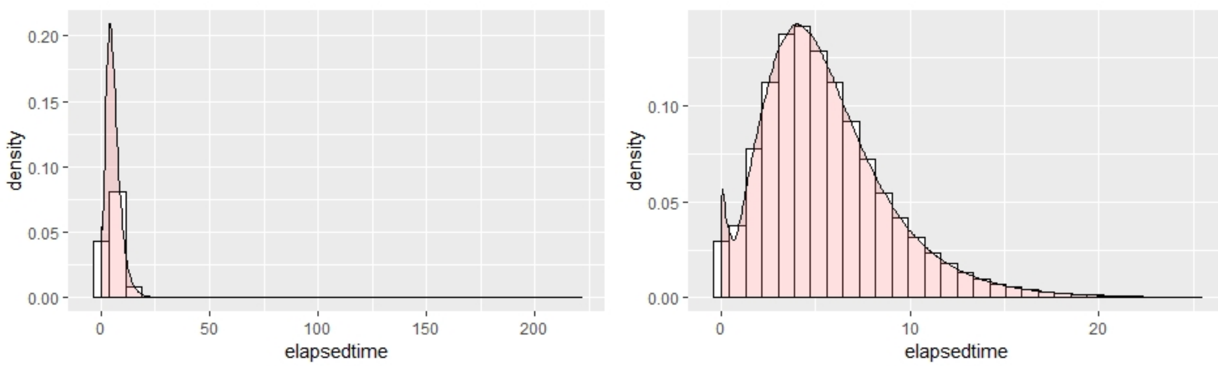

Figure 6: a) Density function of user wait times [left]; b) Density function of user wait times less than 25 minutes [right].

Elapsed time is arguably the most powerful variable of the dataset since it is representative of both service operations and service quality. Despite the appeal to use wait times as an endogenous variable, a strategic modelling decision consists on keeping it as an external control variable, while relying on the model to generate wait times endogenously. After deploying the model, a first-order assessment can be performed by comparing simulated and observed wait time distributions.

3.6. Request acceptance

Request acceptance refers to the time it took for the service provider to accept user requests, which distribution indicates that request acceptance times are extremely short, as shown in Figure 7a. Even though the maximum value observed is 4 minutes, 97.41% of total records are under 0.5 minutes, as depicted in Figure 7b. Furthermore, given that this variable is recorded as fraction of minutes, it could be evidenced that even after filtering, request acceptance times are considerably low, with 73% of all trips were accepted instantaneously (exactly zero values).

The distribution of this variable provides a very valuable insight, since it suggests that service providers accept requests before even finding a match. The supporting argument of this claim is that driver wait times are orders of magnitude longer than request acceptance times.

3.7. Ride duration

This variable represents in-vehicle travel time for users, or equivalently, in-service time for drivers. Figure 8a shows the unfiltered distribution of ride durations, whereas Figure 8b shows 99.97% of the records, corresponding to ride durations less than 100 minutes.

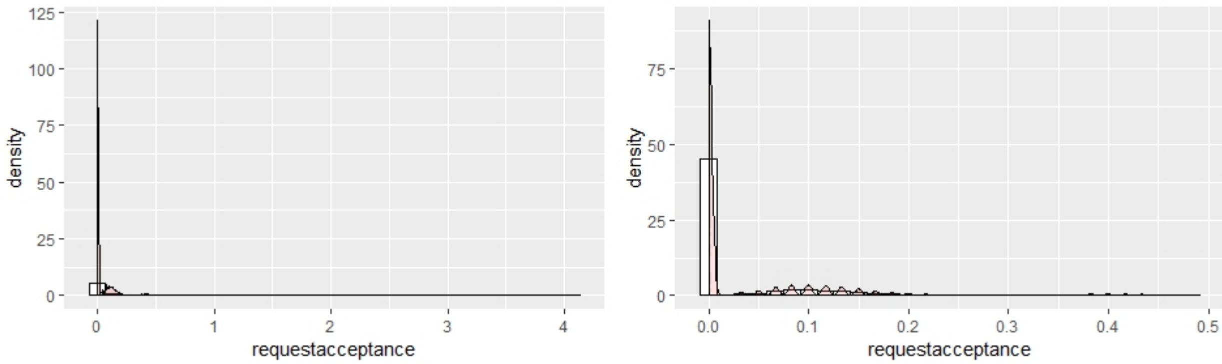

Figure 7: a) Density function of request acceptance times [left]; b) Density function of request acceptance times less than 0.5 minutes [right]

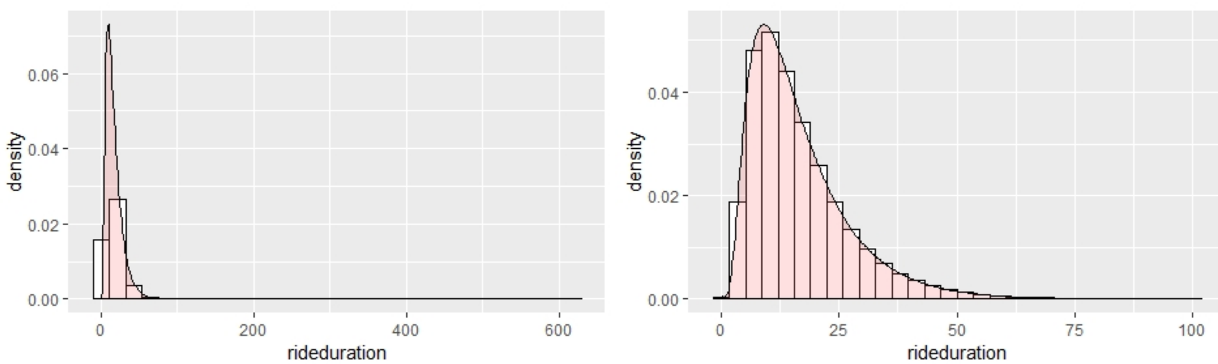

Figure 8: a) Density function of ride duration [left]; b) Density function of ride duration less than 100 minutes [right]

3.8. Spatial variables

Trip locations for both origins and destinations are provided in several formats, the most relevant considered for the models developed here being centreline intersection IDs and municipality IDs. The limitation of centreline intersection IDs is that these only cover the City of Toronto, but not the other municipalities in the region; hence, municipality IDs are used for trips that fall outside of the boundary of the City of Toronto. Naturally, centreline intersection IDs offer a highly granular mapping since these are very densely distributed over the City of Toronto. In contrast, municipalities are geometric area entities, hence offering a much coarser mapping.

It must be emphasized that mapping the observed ridehailing trips in the dataset to a spatial zoning system is an essential task to enable modelling. This is critical not only for spatial visualization of modelling results, but more importantly because vehicle movements (and travel times, by extension) in Periods 1 and 2 are not observed in the data. To fill this data gap, equilibrium congested travel times by relevant time periods are obtained from the GTAModel.

A considerable amount of time was invested into reconciling and linking the mentioned two types of spatial location attributes in the dataset, and the different spatial reference systems, including the one used in the GTAModel. Further details about this task are described below.

To begin with, the dataset contains trips "outside" of the boundary of the GTAModel, with origins/destinations in very distant municipalities (e.g. Kingston). Zones outside of this boundary correspond to the green-coloured zones in Figure 9, whereas zones within the GTAModel boundary are represented by both the blue and red-coloured zones. Moreover, trips outside of the GTAModel boundary cannot be modelled since travel times from/to them are clearly not available.

The next important differentiation is between "internal" (red) and "external" (blue) zones in the GTAModel, where the latter are aggregated into larger zones than regular municipality zones.

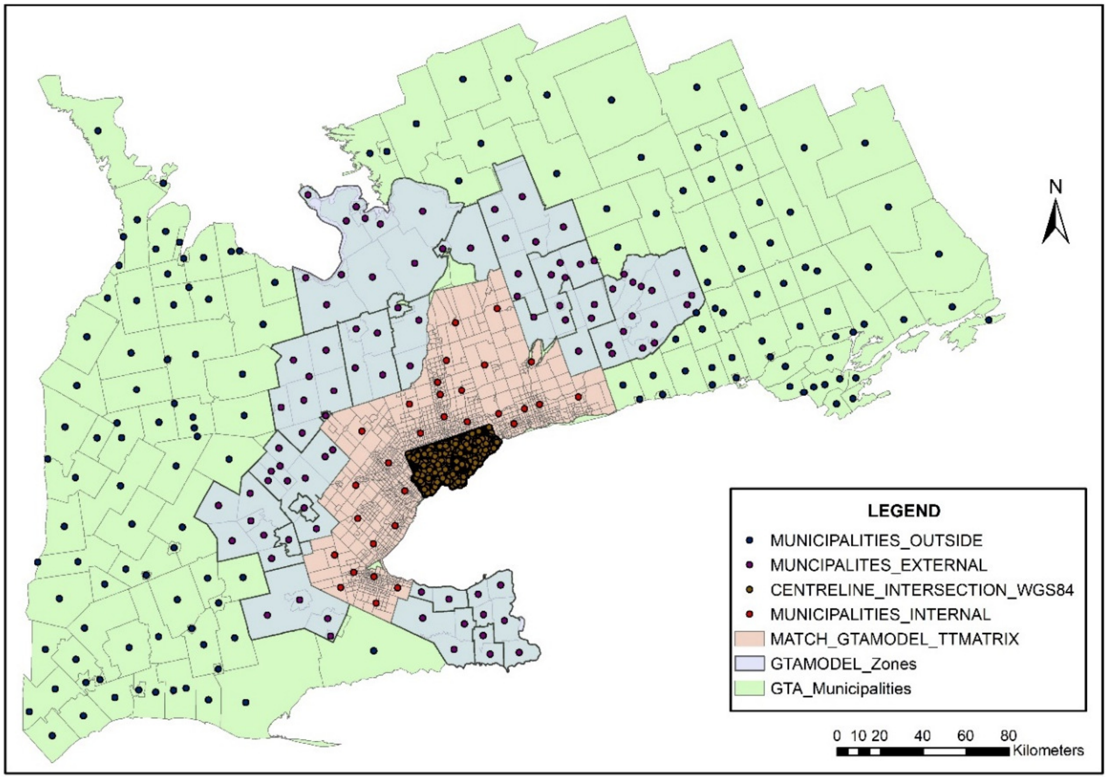

Figure 9: Various spatial references utilized in the study.

Having established all spatial sources, the different mapping combinations that arise are:

- Trip origins/destinations within the City of Toronto boundaries always correspond to internal zones in the GTAModel: trips mapped to zones through centreline IDs.
- Trip origins/destinations outside the City of Toronto boundaries, but still within internal zones in the GTAModel: trips mapped to zones through municipality IDs.
- Trip origins/destinations outside the City of Toronto boundaries, within external zones in the GTAModel: trips mapped to *aggregated* zones through municipality IDs. It should be noted that trips with either origins or destinations corresponding to these external zones are being filtered out due to current issues in the GTAModel; however, there were never more than 20 records filtered out for any of the days modelled in this report.
- Trip origins/destinations outside the City of Toronto boundaries, and outside the boundaries of the GTAModel: not possible to map.

After the mapping process is performed, every origin/destination of a mapped trip is coded as a particular GTAModel zone. As mentioned, this will allow for OD travel time allocation when endogenously generating en-route trips. However, given the macroscopic, zone-based nature of the network assignment component of the GTAModel (EMME[2]), intrazonal travel times are not generated. To overcome this limitation, intrazonal travel times are calculated based on geometric features of each zone, as follows:

$$tt_{ii} = \frac{\sqrt{Area_i}}{6} \Big/ Average\ speed$$

Where the numerator approximates a representative distance considering the area of zone i, and the average speed considered for the calculation is 40 Km/h.

4. Ridehailing Service Provision Prototype Model

Fundamental characteristics of ridehailing mobility services consists of the relationship and interdependency among demand for the service (trip requests), supply of the service (vehicles and its drivers), and an (often dynamic) pricing mechanism that drives the system and brings all pieces together. Given the unavailability of data regarding vehicles (and their drivers) and fares paid by users that would allow for a full-blown operational model, a prototype model has been developed for this project. The model's main premise derives from the fact that observed demand provides the minimum number of vehicles active in the system at any given time. Hence, the first step taken is to instantiate vehicles in the system at "time zero", to match observed demand at the same time. Afterwards, vehicles are generated whenever the active fleet size is lower than observed demand at any given point in time. An implication of this approach is that reactions of drivers to price levels cannot be accounted for, hence, undersupply conditions do not occur, and oversupply conditions will be largely underestimated, limited only to a second-order outcome of vehicles remaining in activity after serving trip requests.

A second key component of the model consists modelling driver activity, which accounts for driver decisions of becoming active/inactive in the system. By assembling this modelling structure, the system can keep track of vehicle agents throughout the day and ultimately generate service performance metrics and vehicle-level metrics such as trip chains.

Another key feature of the model is that it is time-step driven, since it steps through time by discrete time intervals. As established in Section 3.3, a time step of 1 minute was deemed as an adequate initial value. For the interested reader, a discussion on the adequacy of time-step based temporal modelling, as well as an overall modelling framework for mobility services, can be found in recent and on-going research conducted by the authors (Calderón & Miller, 2019b, 2019c, 2019a).

Random days are selected to be run through the model because behavioural and procedural features of within-day ridehailing operations do not vary from day to day. Fridays and weekends are only

[2] EMME is a commercial software specialized in network assignment. For further information: https://www.inrosoftware.com/en/products/emme/

different in their temporal distribution of demand, yet not in how the service operates. Further, it is important to consider that, from an activity perspective, a day starts at 4am in the morning and finishes at 4am in the next day. Both users and drivers follow these "normal" activity patterns. This consideration is particularly important for a logically consistent generation of trip chains, since at 3 am in the morning a driver is likely finishing their day rather than starting a work shift.

Note that, as opposed to more descriptive types of analysis, the computational cost of running the prototype model developed in this study is considerably higher, with processing times for modelling a given single day being in the range of 5-15 minutes (predominantly dependent on the length of the time interval) in an Intel Core i5 8th generation laptop computer. In this sense, random days were selected from the dataset to model their respective within-day operations. While running the model for the whole set of observed days (850) is conceivable, it is deemed as unnecessary for the scope and purpose of this report.

The prototype model is documented next in terms of its required data structures and its high-level components and structure.

4.1. Data Structures

The data structures used in this model are based on the ones proposed in Calderón and Miller (2019a), which are list-based in order to exploit the agent-oriented nature of the model. Namely:

- A trip list (t-list), which contains the full set of trip requests from a given day. Deriving from it, a demand partition (t-listd) is retrieved by the model every time interval and made available to the service provider. This enforces the restriction that demand for the whole day is not known to providers.
- A fleet list (f-list), which contains the total number of vehicles available for operations (absolute fleet size).
- A vehicle list (v-list), which consists of a "snapshot" of the f-list at a time interval, hence it contains the vehicles that are available to the service provider at a given point in time. This list is essential for handling operational within-day dynamics.

These list-based data structures effectively allow to have trip requests and vehicles as the fundamental units of analysis in the system, and more importantly, allow to keep track of vehicles throughout the day.

4.2. High-Level Model Structure and Components

A high-level representation of the model structure is depicted in the high-level flow diagram shown in Figure 10. The algorithmic steps are adapted from Calderón and Miller (2019a), and describe key processes that drive the model, each one of them are described in further detail next.

The first step consists of retrieving from t-list all trips whose request times fall within the current time interval and storing them in t-listd. In terms of information flows, this subset of trip requests is made available to the service provider, and thus, subject for processing within the time interval.

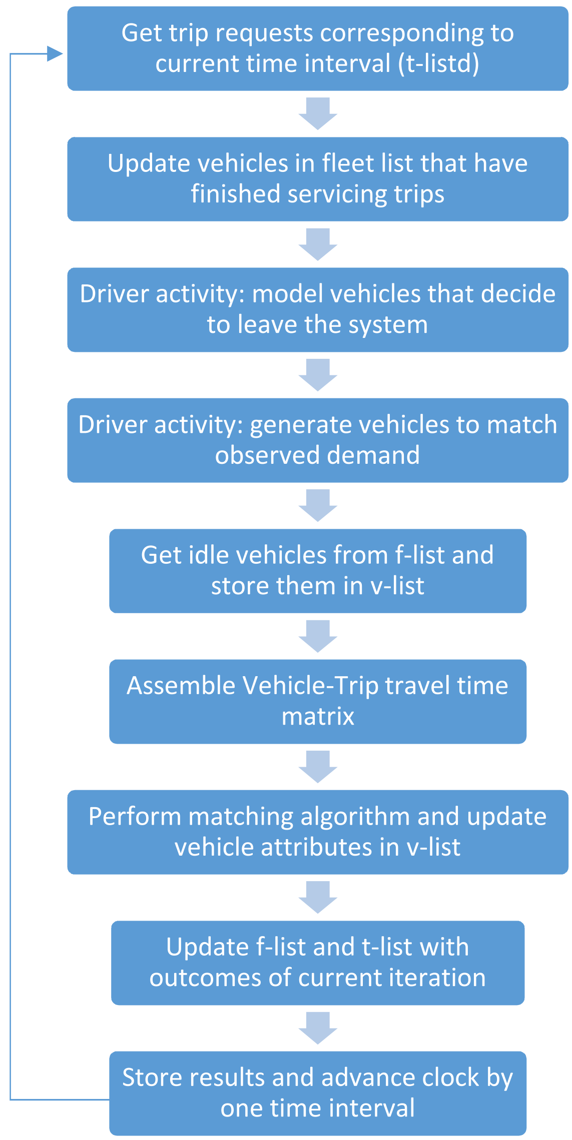

Figure 10: High-level flow diagram of model structure

The second step consists of querying f-list for vehicles that finished servicing before the start of the current time interval and updating their state to "Idle". Note that the location of vehicles that finished servicing a trip becomes the destination of the user trip, implying that the model assumes that drivers stay at user destination after servicing a trip and while waiting for the next trip.

The third step is the first element of modelling driver activity and is concerned with driver decisions to leave the system (stop working). To achieve this, a probabilistic modelling technique is utilized in the following manner:

- Given the lack of driver activity data, relatively ad-hoc assumptions are adopted based on the scarce evidence found about driver activity. Namely, drivers working patterns are sparse and an average driver works 15 hours a week (The Star, 2018). Hence, a normal probability density function (PDF) is assumed to represent the probability of a driver working a certain number of hours in a given shift, as shown in Figure 11. This pdf is assumed to have a mean of 3 hours and a standard deviation of 1 hour, following the logic that a driver might work 5 days per week and thus 3 hours per day on average.

Figure 11: Distribution of number of driving hours

- To generate a probability of leaving the system, the corresponding cumulative density function (CDF) is utilized, since it represents the probability of a driver leaving the system having worked up to a given number of hours. As depicted in Figure 12, this value becomes the probability of success of a Bernoulli trial that ultimately yields a 1/0 binary outcome.

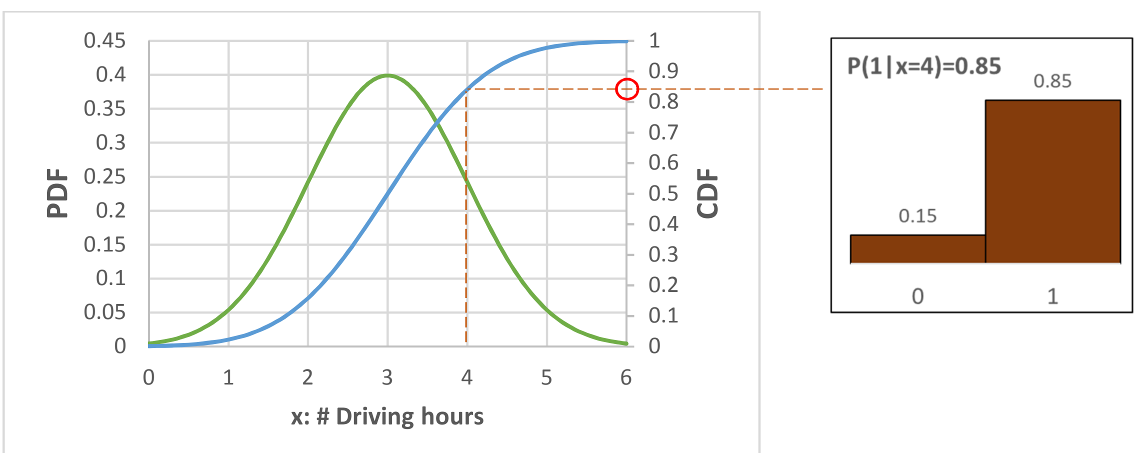

Figure 12: Generation of "active/inactive" value from normal cumulative density function

- This probabilistic model is applied to every active driver in the system, and the state of all drivers that decided to leave the system is updated to "Inactive".

The next step in the algorithm is also related with driver activity, as it enforces the model constraint that at any given time, there must be at least as many drivers as observed trips. Hence, eligible drivers (state = "Inactive") are selected from the f-list to become active. Currently, for drivers that enter the system for the first time, the model instantiates them in random locations throughout the region, but this location generation process can be refined to account for driver residential locations or demand spatiotemporal patterns.

Note that there is a possibility of activating drivers that just became inactive in the previous step, thus, to control for this, drivers that have previously been active in the system must wait a user-defined period before re-entering the system. In addition to this, it is desirable from a behavioural perspective to account for multiple working shifts for drivers as this is the case in real life conditions. This "return to activity" period is currently set to 5 hours, but it could conceivably be also modelled probabilistically for each vehicle agent. The latter option was not implemented to avoid introducing further uncertainty in the model given the lack of data that provides a ground truth on driver activity.

Having modelled relevant driver activity processes, a subset of all Idle vehicles in the f-list is retrieved and stored in the v-list, which completes the required information for the service provider to perform operational tasks.

For a given time interval, all the combinations of location of vehicles available for matching (v-list) and the origins of trip requests (t-listd) are assembled into a Vehicle-Trip Matrix (VTM), in which: rows represent the location of vehicles, columns represent the origin of trips, and cells represent the corresponding time-of-day-dependent travel times obtained from the GTAModel. This matrix is a core instrument of the model since it provides necessary information to perform matching. Furthermore, it must be noted that the travel times stored in this matrix represent both simulated en-route times for vehicles and wait times for users.

The matching algorithm is then deployed until all trip requests are served. Thanks to modular algorithmic design, any desired matching algorithm can be implemented in the model, such as purely random matching, greedy, centralized greedy, Hungarian, etc. For a very good discussion on matching algorithms, the reader is referred to Hanna, Albert, Chen, and Stone (2016). After testing and experimentation (see the Appendix), the current implementation consists of a greedy matching algorithm. This is aligned with the statement made by Hanna et.al. that ridehailing service providers such as Uber deploy simple First-Come-Fist-Serve, nearest-vehicle matching strategies to cope with the real-time dynamics of the service. A brief description of the implemented greedy matching algorithm is provided below.

- Trip requests of a given time interval (t-listd) are not time-ordered, but instead are randomly stored by default (following the also non-time-ordered characteristic of the raw data). Thus, going through t-listd in order can actually be considered as random.
- Each iteration of the algorithm allocates one vehicle to a trip.
- A subset vector from VTM is retrieved, which contains travel time values from the location of all available vehicles to the origin of the trip being processed.
- The minimum travel time from the filtered VTM is found, and its corresponding vehicle and trip request are matched.

- The travel time values corresponding to the matched vehicle (row) and trip request (column) in the original VTM are set to infinity to rule them out from subsequent iterations.
- Update relevant attributes in v-list and t-listd.
- Repeat process until all trips are served.

The final steps consist of updating the f-list with the outcomes recorded in v-list, storing relevant results for post-process visualization and analysis, and increasing clock time by one unit (the size of the time interval) to continue advancing through the day.

5. MODEL RESULTS AND DISCUSSION

Given the agent-oriented nature of the model, a wide variety of system-level and agent-level metrics can be generated and kept tracked. Before diving into results, it is worth highlighting that the core components towards which the model is sensitive, and thus the ones that provide a modelling leverage, are the matching algorithm, the time interval length, and the vehicle generation processes – the latter especially location-wise. This was clearly seen throughout model development as well as in sensitivity tests, which are not included in the main body of this report but can be found in the Appendix. The results presented below consist of the best combination of the abovementioned core components, as assessed against observed wait time distributions. The outputs that are currently generated by the model are presented next for Thursday, April 30, 2017, which was chosen on the grounds that is the day with the highest demand for ridehailing trips observed in the dataset containing detailed minute-by-minute trip records.[3] Alternative days (including Fridays) were also tested, as reported in the Appendix.

General model metrics are presented in Table 1 below, followed by graphical representations of system-wide and agent-level metrics.

Table 1: Model outputs for Thursday 30/03/2017

Metric	Value	Unit
Total number of trips to be served	86018	Trips
Absolute fleet size deployed to provide service	10064	Vehicles
Maximum en-route time (user wait time) observed in the simulation	1.15	Hours
Average number of idling vehicles (over 5-minute intervals)	5.62	Veh/5min
Percentage of time there were 0 vehicles idling in the system	82.99	Percent
Total time that the fleet is active, yet not in-service (Periods 1 & 2)	5149.5	Hours
Percentage of time that fleet is active, yet not in-service (Periods 1 & 2), with respect to overall active time	18.79	Percent
Total VKT when the fleet is en-route (Period 2)	302296	Kilometers
Percentage of VKT when fleet is en-route (Period 2), with respect to overall VKT	29.01	Percent
Average simulated wait time (**true observed value=4.98 minutes**)	6.01	Minutes
Trip connections that would be unfeasible with respect to observed data	14	Percent
Model runtime	5	Minutes

[3] After April 30, 2017, trip records were only stored aggregated to one-hour time intervals, thereby significantly reducing the temporal granularity of any analysis based on these data.

An important consideration to bear in mind is that the amount of idling is inevitably underestimated because of the nature of the model itself. Supply (vehicles) is "catching up" with demand (trip requests) over the day, hence oversupply conditions (and idling by extension) arise in a very limited manner, mostly triggered by rapid demand drops after demand surges (refer to Figures 13 and 18 below).

The most relevant results for the purposes of this study then consist of the total time spent and Vehicle Kilometers Travelled (VKT) by vehicles in the network without a passenger (Periods 1 and 2). Results indicate that for this particular day, en-route and idling together amounted to a total of 5149.5 hours (23% of overall fleet active time) and 302296 kilometers (37.95% of overall fleet VKT). Note that, with respect to distance, cruising behaviour of drivers while idling is not being modelled but rather vehicles are assumed to remain at their last passenger drop-off. Consequently, distance estimates while idling are not considered in the total VKT estimates and thus not reflected in the percentage calculated.

In terms of assessing model performance, the average simulated wait time is very close to the average observed wait time. Further, when comparing the distribution of these variables (refer to Figure 19), the model is performing quite successfully in replicating observed data. Room for improvement certainly exists though, particularly considering the thick tail in the distribution of simulated times. The most likely explanation that has been hypothesized for this result is related with the vehicle generation process of the model, which currently randomly chooses zones within the study area as the initial vehicle locations for instantiated vehicles. It is quite possible that this is unrealistic, specially considering the variety of factors that can influence these processes. A new data source regarding driver residential locations has been recently provided by the City of Toronto to the study team, which can allow for testing alternative location allocation approaches within vehicle generation, but this analysis has not yet been undertaken at the time of this report's preparation.

The model also outputs some control variables such as the absolute fleet size deployed to provide service, which is below the reported 39118 registered drivers by March 2017 and implies that around 25% of registered drivers worked on the modelled day.

It is also worth mentioning that other new data sources regarding aggregate fares and number of active drivers in a day have been recently made available to the authors, which will enable further testing and model development. Again, however, incorporation of this new information into the analysis has not been undertaken.

In terms of graphical outputs, temporal patterns of demand (#trips), supply (# available vehicles), and idling vehicles throughout the day can be clearly observed in Figure 13. Moreover, the secondary vertical axis of the plot refers to the number of dormant vehicles in the system, which are vehicles that have never entered the system throughout the day. Note that the total fleet size is a user-defined parameter in the model that is only used for reference, yet the model is agnostic to fleet size, only throwing an exception when fleet size is not enough to satisfy observed demand.

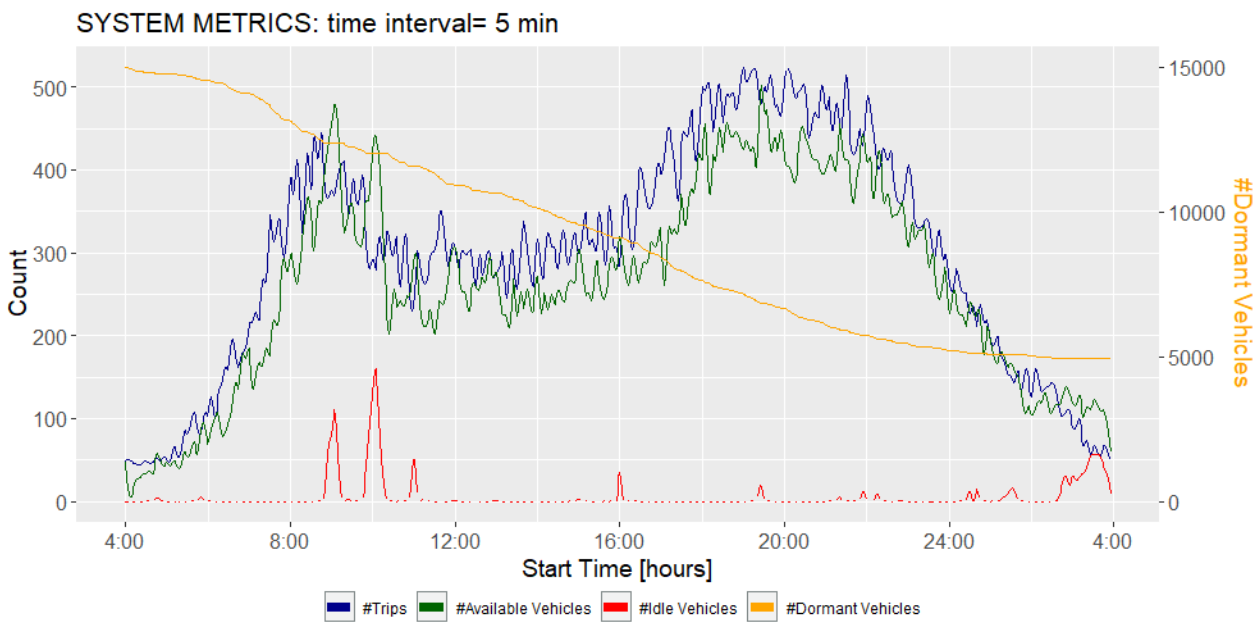

Figure 13: Temporal distribution of main system metrics for Thursday, 30/03/2017.

Figure 14 depicts the temporal variation of the total time spent en-route by the whole fleet, at any given 5-minute interval within the day. As expected, this figure has a shape that is consistent with demand temporal patterns. Further, Figure 15 presents en-route times as a per-vehicle average, meaning that at any given 5-minute interval within the day, the total fleet en-route time is divided by the total number of vehicles matched to trips. Interestingly, the temporal patterns in Figure 15 are roughly inverted with respect to Figure 14, which reflects on the fact that at early and late hours of the day, there are fewer vehicles in the system and trips are more spread out spatially, resulting in longer en-route times from the perspective of the vehicle agents.

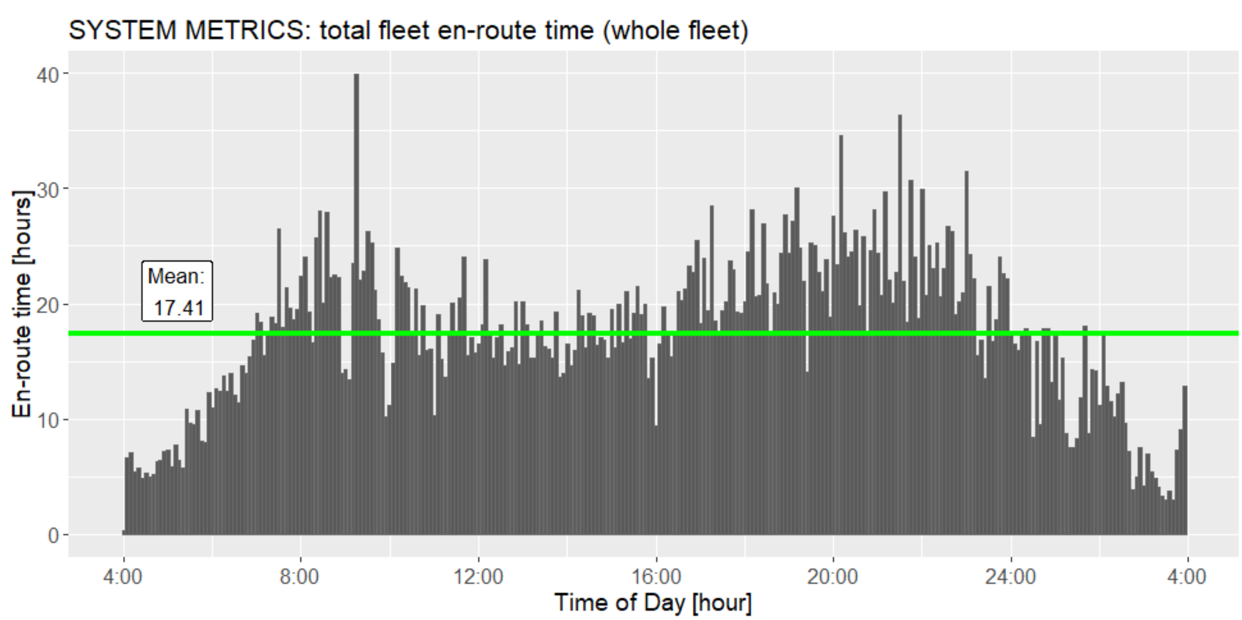

Figure 14: Temporal distribution of fleet total en-route time for Thursday, 30/03/2017.

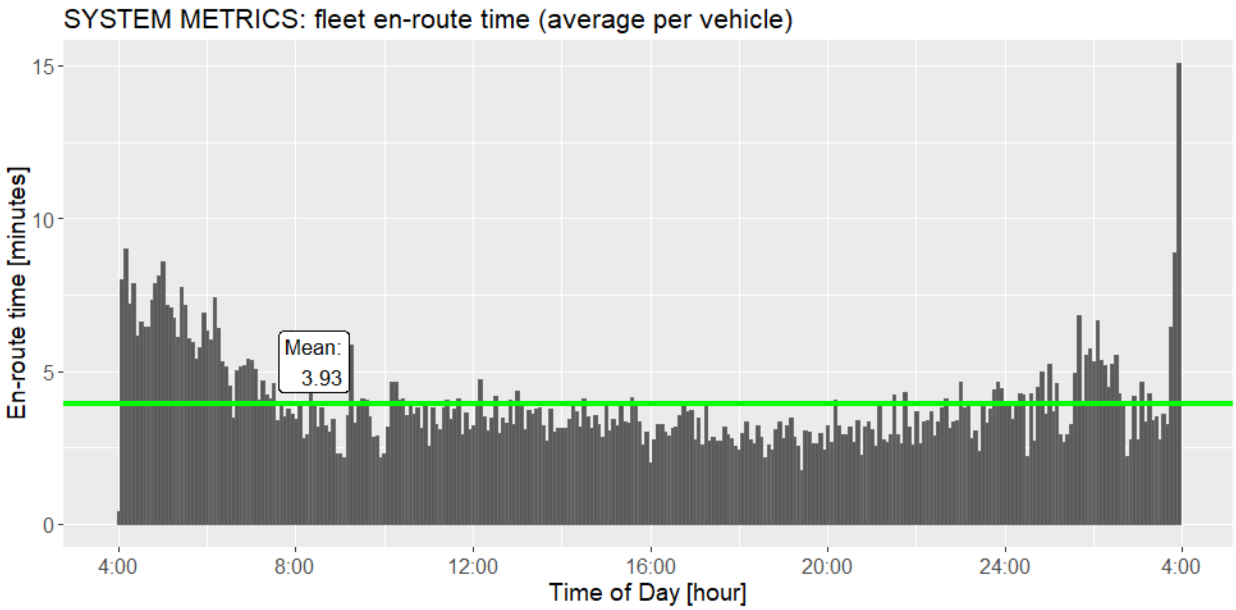

Figure 15: Temporal distribution of average en-route time per vehicle for Thursday, 30/03/2017.

The amount of VKT by the whole active fleet, by time of day, is shown in Figure 16. As expected, it is clearly mirroring the en-route distribution shown in Figure 14, since both reflect vehicles that are en-route, yet each figure plots a different metric. Likewise, Figure 17 depicts the average VKT covered by individual vehicles, and its patterns also mirror Figure 15.

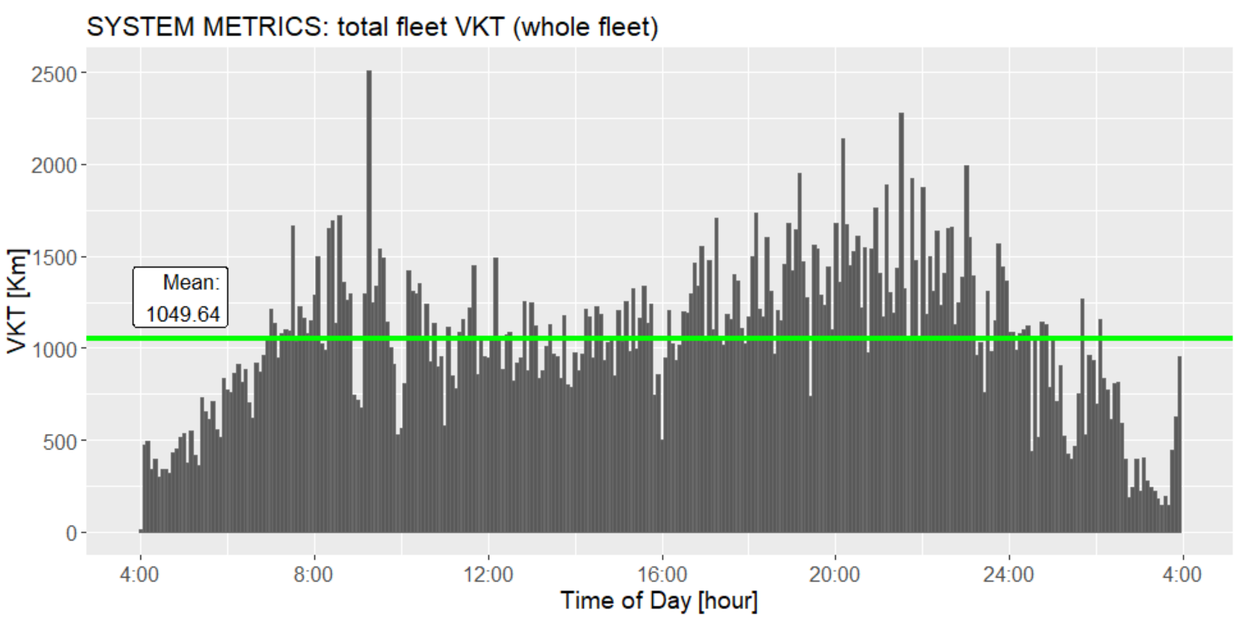

Figure 16: Temporal distribution of total fleet VKT for Thursday, 30/03/2017

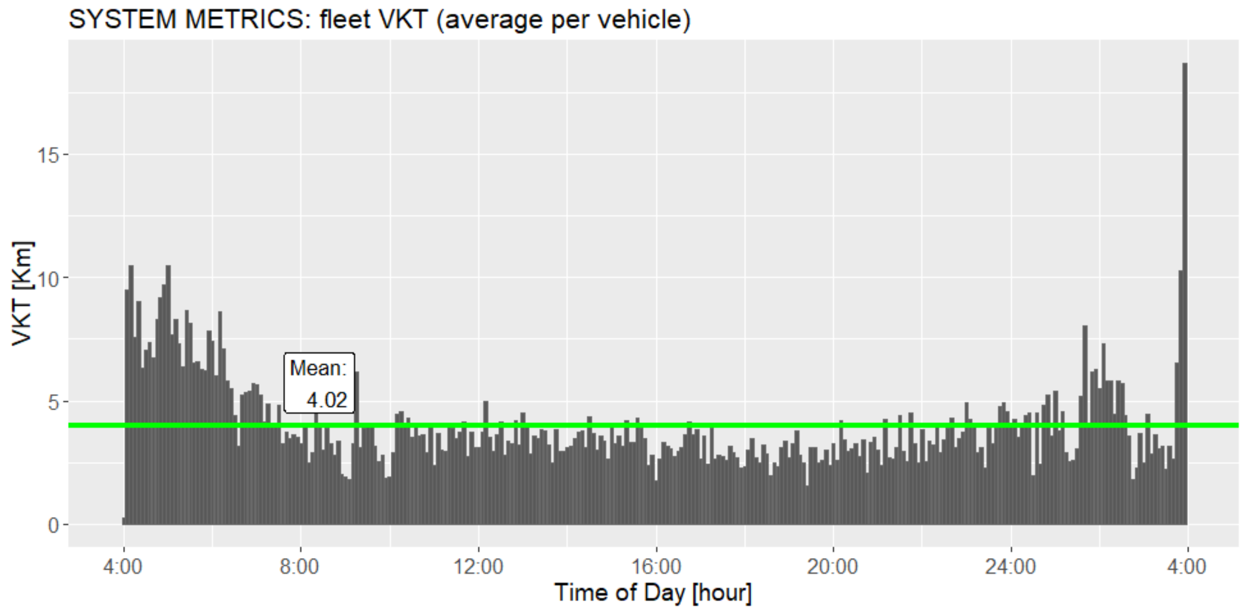

Figure 17: Temporal distribution of average VKT per vehicle for Thursday, 30/03/2017.

As mentioned earlier, idle times are only slightly captured by the model due to its data limitations, hence the low values found in Figure 18 presumably provide a very low estimate on this phenomenon.

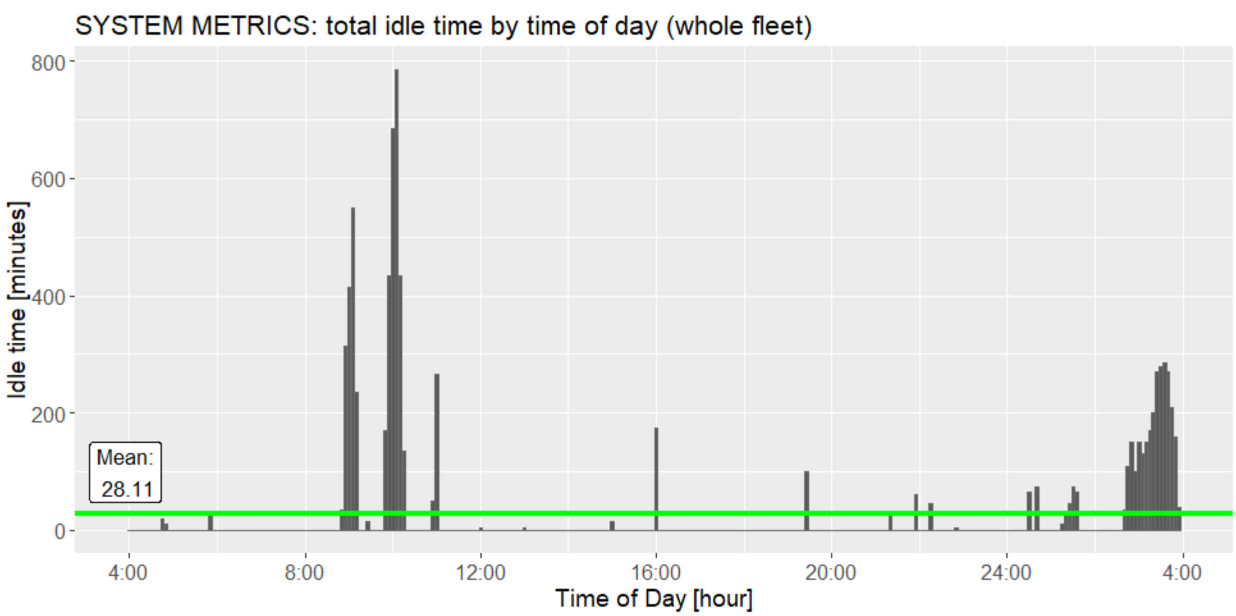

Figure 18: Temporal distribution of total fleet idle time for Thursday, 30/03/2017.

Figure 19 shows a comparison of simulated and observed wait times, which consist of the main control check of the model. An analysis on the influencing factors and possible explanations of the slight differences observed has already been discussed above.

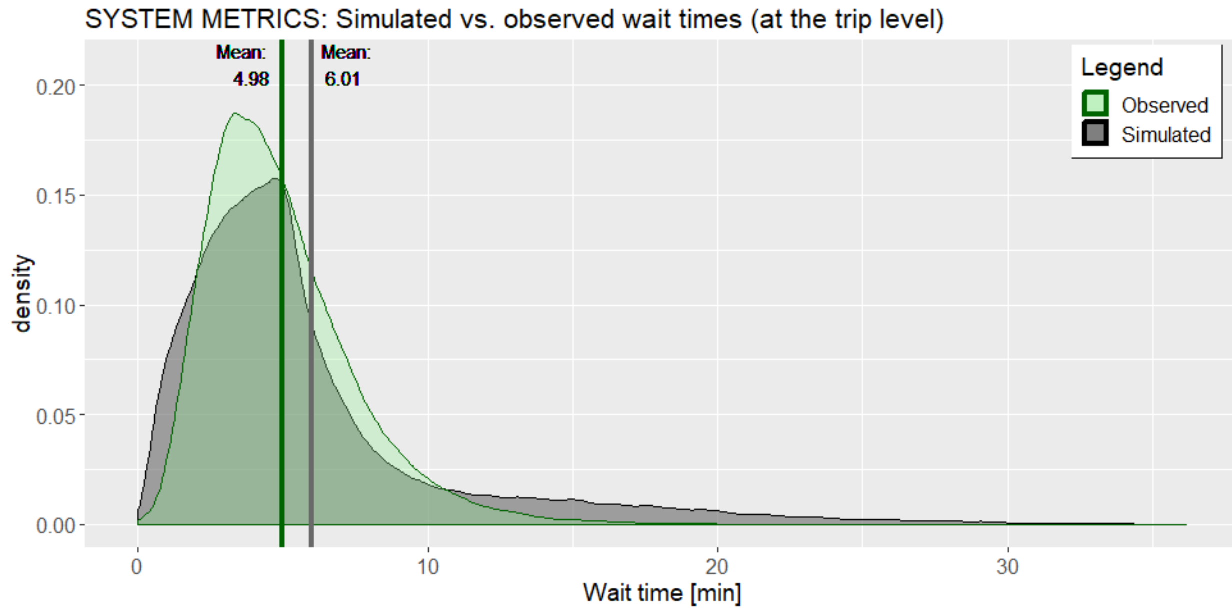

Figure 19: Comparison among distributions of simulated and observed wait times for Thursday, 30/03/2017.

Additionally, by recording the trip requests that each vehicle is assigned to over the course of the day, trip chains can be generated and thus also the number of trips per vehicle, as shown in Figure 20. What also stands out is that the distribution of number of trips per vehicle has a thin, long tail, indicating that a very small subset of vehicles serves up to 28 trips in a day. This could also be an implication of location allocation in the vehicle generation process. In any event, 28 trips in a day is not a mathematical or logical impossibility: it could happen for a driver working 8 hours that has constantly been servicing short trips (likely within the CBD); this translates into servicing one trip every 15 minutes, with 2 minutes in-between consecutive trips.

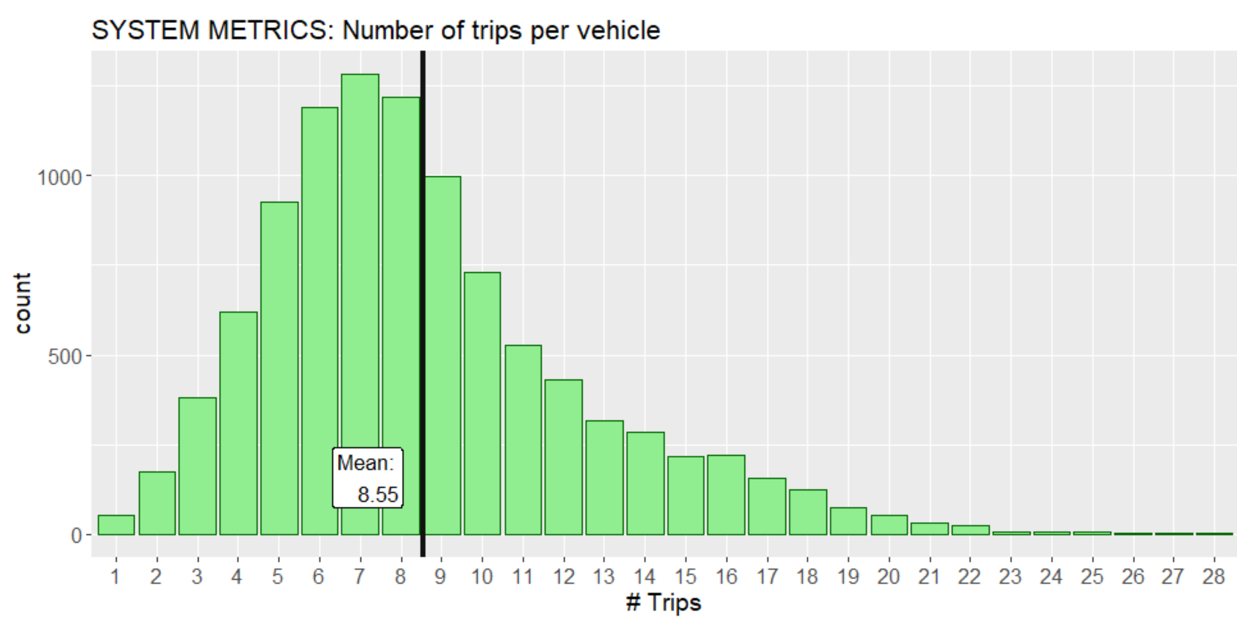

Figure 20: Number of trips per vehicle for Thursday, 30/03/2017.

The final output of the model consists of trip chains of every vehicle that entered the system. This is an input that was originally requested by the City Staff to compare against the fleet size minimization approaches tested by them under this project, but it also useful as a comparative output since it allows for comparison between the endogenously generated trip chains of the model and observed pick-up and drop-off variables. As per Table 1, 14% of subsequent trip connections would be unfeasible when compared to observed data, meaning that the observed drop-off time of a trip is later than the observed pick-up time of the next trip in the chain.

In any event, the percentage of infeasible connections is a completely expected one since a considerably large portion of service operations – along with drivers – are endogenously modelled due to lack of observed data. In particular, the actual matching mechanism performed by the service provider is not known to modellers due to proprietary rights, hence the outcomes of matching translate into path-dependant vehicle locations, which are also then linked to user origins by en-route times obtained from an external source (GTAModel). Given these circumstances, an inflated 14% is an acceptable and encouraging outcome, especially considering that average simulated wait times are relatively close to observed wait times and their distributions are very similar in range and shape. To illustrate the outputs generated at the trip chain level, Table 2 presents an example of a trip chain of a randomly chosen vehicle within this simulation.

Table 2: Example of a vehicle trip chain for Thursday, 30/03/2017.

Trip Chain for Vehicle#: 510				
Trip #	**PTCID**	**Observed Request time [hrs]**	**Observed Pick-up time [hrs]**	**Observed Drop-off time [hrs]**
1	32538DE31D564CCA93906CB239EBD0A6	20.1150	20.4064	20.6258
2	8759FBF230DB4086A5E130CFF8596872	20.4047	20.4517	21.1519
3	F6894B5F15514DCB8A43168B57016FF7	21.3258	21.3706	21.4083
4	32354316E07B4368BC50626757A1684F	21.4461	21.4778	21.5894
5	87AAC7E8A0DF407DB734DDEE34CD04A1	21.6303	21.7269	21.7950
6	3893B30A40244AD4B1CD8DF6DAF9D284	21.6919	21.7856	22.0767
7	ADC2DEFA3FF84054A8EFC3F4A2ABC289	22.0575	22.0989	22.3322
8	5C44953CE2D84F66AB5A88B8C64D9D0A	22.3347	22.4139	22.5489
9	DF43E18079CD4E1A9EC5CB16362D5DF7	22.5553	22.6275	22.9156
10	6D4156AED607499CA5A1F908E81412EB	22.8622	22.9239	23.1289

Instances of infeasible connections can be observed in Table 2 when comparing the observed drop-off times of trip 1 and 5 with the observed pickup time of trips 2 and 6, respectively. Evidently, however, full sequential consistency exists between **modelled** pickup and drop-off times.

Model results were highly consistent when testing different days and demand patterns, which suggests that the model is stable and robust with respect to its different exogenous inputs. Relevant tests against day-to-day variations are also documented in the Appendix.

6. Conclusions and Remarks for the VfH By-Law Review

The results presented in this report demonstrate that, when confronted with the lack of observed operational data regarding driver activity and pricing, a solid conceptualization of the operations of a ridehailing mobility service applied to observed demand proved to be successful in replicating wait time distributions. The potential and capabilities of the prototype model developed in this study are promising and could be further enhanced with additional sources of operational data. Results, albeit encouraging, should be handled with caution and should not be used to make definitive claims. Nevertheless, broader and more generic conclusions can be asserted with a higher degree of confidence, such as the percentage of time (out of total active time) that vehicles spend en-route and idling (Periods 1 and 2) is around 19%, and the percentage of VKT that can be attributed to vehicles en-route (Period 2) is around 29%.

The results mentioned provide an "educated guess" and a proxy estimation of potential network impacts of ridehailing services. As mentioned, data limitations constrained the model scope to only be able to enforce the realization of observed demand, via generation of vehicles. This implies that oversupply conditions are largely overlooked, hence estimates of idling are very low by construction. By adding at least one more fundamental component – either driver information or fares paid by time of day, oversupply conditions could be estimated in a much better capacity. More generally, further data provision can significantly improve the accuracy and applicability of the modelling efforts presented here.

In terms of robustness, the model proved able to produce stable and consistent results for different random days, both weekdays and weekend days. Furthermore, the sensitivity tests (see Appendix) with respect to variations in time interval size, matching algorithms, and vehicle generation/instantiation corroborated the authors' hypotheses on them being key leverage variables with respect to ridehailing service operations.

The recommendation that can be offered for the *Vehicle for Hire Bylaw Review* project is to consider the results shown here as an approximate estimate of ridehailing operations. Hopefully, the results shown in this report can motivate and expedite the data provision possibilities initially contemplated under the project, which would allow the authors to considerably relax some of the assumptions made, resulting in much stronger analyses that could well be considered for definitive evidence-based decision making.

Acknowledgements & Disclaimer

This report was prepared as part of the project *UTTRI Technical Support for the City of Toronto Vehicle for Hire Bylaw Review*, funded by the City of Toronto. The findings and opinions expressed in this report are those of the authors alone and do not necessarily reflect those of the City of Toronto.

REFERENCES

Alonso-mora, J., Wallar, A., Frazzoli, E., Rus, D., Alonso-mora, J., Samaranayake, S., … Rus, D. (2017). On-demand high-capacity ride-sharing via dynamic trip-vehicle assignment. *Proceedings of the National Academy of Sciences*, *114*(3), 462–467. https://doi.org/10.1073/pnas.1611675114

Calderón, F., & Miller, E. J. (2019a). A Conceptual Framework for Modelling the Supply Side of Mobility Services within Large-Scale Agent-Based Models. *Submitted to Transportation Research Part A: Policy and Practice*.

Calderón, F., & Miller, E. J. (2019b). A Literature Review of Mobility Services: Definitions , Modelling State- of-the-Art , and Key Considerations for a Conceptual Modelling Framework. *Submitted to Transport Reviews*.

Calderón, F., & Miller, E. J. (2019c). A new outlook on ridehailing: spatiotemporal patterns and commuting analysis from the Greater Toronto and Hamilton Area. In *Procedia Computer Science*. Elsevier B.V.

City of Toronto. LICENSING OF VEHICLES-FOR-HIRE ARTICLE (2016). Toronto.

Desrosiers, J., Dumas, Y., Solomon, M. M., & Soumis, F. (1995). Time Constrained Routing and Scheduling. *Handbooks in Operations Research and Management Science*, *8*(C), 35–139. https://doi.org/10.1016/S0927-0507(05)80106-9

DMG. (2016). Transportation Tomorrow Survey-University of Toronto Data Management Group. Retrieved May 20, 2018, from http://www.tts2016.ca/en/

Doerner, K. F., & Salazar-González, J.-J. (2014). Pickup-and-Delivery Problems for People Transportation. In *VEHICLE ROUTING Problems, Methods, and Applications*.

Hanna, J. P., Albert, M., Chen, D., & Stone, P. (2016). Minimum Cost Matching for Autonomous Carsharing. In *Proceedings of the 9th IFAC Symposium on Intelligent Autonomous Vehicles (IAV 2016)*. Leipzig, Germany. Retrieved from http://www.cs.utexas.edu/~aim/papers/IAV16-Hanna.pdf

Liu, Y., Bansal, P., Daziano, R., & Samaranayake, S. (2018). A framework to integrate mode choice in the design of mobility-on-demand systems. *Transportation Research Part C: Emerging Technologies*. https://doi.org/10.1016/j.trc.2018.09.022

Lyft. (2018a). Acceptance rate – Lyft Help. Retrieved May 11, 2019, from https://help.lyft.com/hc/en-ca/articles/115013077708-Acceptance-rate

Lyft. (2018b). Prime Time for drivers – Lyft Help. Retrieved May 12, 2019, from https://help.lyft.com/hc/en-us/articles/115012926467-Prime-Time-for-drivers

Syed, A. A., Irina, G., & Bogenberger, K. (2019). Neural Network Based Metaheuristic Parameterization, with Application to Vehicle Matching Problem in On Demand Mobility Services. In *The 98th Annual Meeting of the Transportation Research Board*. Washington, DC, January.

The Star. (2018). Uber Canada to force drivers to take a break after working 12-straight hours | The Star. Retrieved May 11, 2019, from https://www.thestar.com/business/2018/02/21/uber-canada-to-force-drivers-to-take-a-break-after-working-12-straight-hours.html

TMG. (2018). Documents-Travel Modelling Group. Retrieved June 1, 2018, from https://tmg.utoronto.ca/documents/#gtamodel-documentation

Uber. (2019a). Legal | Uber. Retrieved May 11, 2019, from https://www.uber.com/legal/rewards-program/uberpro/us-en/

Uber. (2019b). What is surge? | Uber Partner Help. Retrieved February 6, 2019, from https://help.uber.com/partners/article/what-is-surge?nodeId=e9375d5e-917b-4bc5-8142-23b89a440eec

APPENDIX: SENSITIVITY TESTS

Initially, a 1-minute time interval and centralized greedy algorithm were adopted, the latter consists of the following steps:

- Each iteration of the algorithm allocates a vehicle to a trip.
- An initial travel time threshold is defined.
- A subset of the Vehicle-Trip Matrix (VTM, as per Section 4.2) is found, formed by all travel time values below the current threshold. If no values found, then the threshold is increased progressively.
- The minimum travel time from the filtered VTM is found, and its corresponding vehicle and trip request are matched.
- The values in VTM of the corresponding matched row and column are set to infinity to rule them out from consideration in subsequent iterations.
- Update relevant attributes in v-list and t-listd.
- Repeat process until all trips are served.

As Table A1 shows, a centralized greedy outperforms a greedy algorithm in practically all fronts, except for the percentage of trip connections of vehicle trip chains that would be infeasible when compared against observed data. This exception is not one to overlook, since, theoretically, it is expected for a centralized greedy implementation to be superior; however, this does not mean that service providers rely on this approach in real life. In fact, this argument becomes stronger when comparing Figure A15 and Figure A16.

Table A3: Comparison of main model metrics under different matching algorithms: 1-min time interval.

Metric	C-Greedy	Greedy	Unit
Total number of trips to be served	86018	86018	Trips
Absolute fleet size deployed to provide service	9801	10023	Vehicles
Maximum en-route time (user wait time) observed in the simulation	1.33	1.08	Hours
Average number of idling vehicles (over 5-minute intervals)	5.18	4.918	Veh/5min
Percentage of time there were 0 vehicles idling in the system	64.35	65.462	Percent
Total time that the fleet is active, yet not in-service (Periods 1 & 2)	6615.7	7438.9	Hours
Percentage of time that fleet is active, yet not in-service (Periods 1 & 2), with respect to overall active time	22.92	25.05	Percent
Total VKT when the fleet is en-route (Period 2)	403923	450702	Kilometers
Percentage of VKT when fleet is en-route (Period 2), with respect to overall VKT	38.75	43.24	Percent
Average simulated wait time (**true observed value=4.98 minutes**)	5.028	5.61	Minutes
Trip connections that would be unfeasible with respect to observed data	18	15	Percent
Model runtime	15	10	Minutes

All figures presented below correspond to the default testing day of Thursday, March 30, 2017.

Figure A21: Temporal distribution of main system metrics for c-greedy algorithm.

Figure A22: Temporal distribution of main system metrics for greedy algorithm.

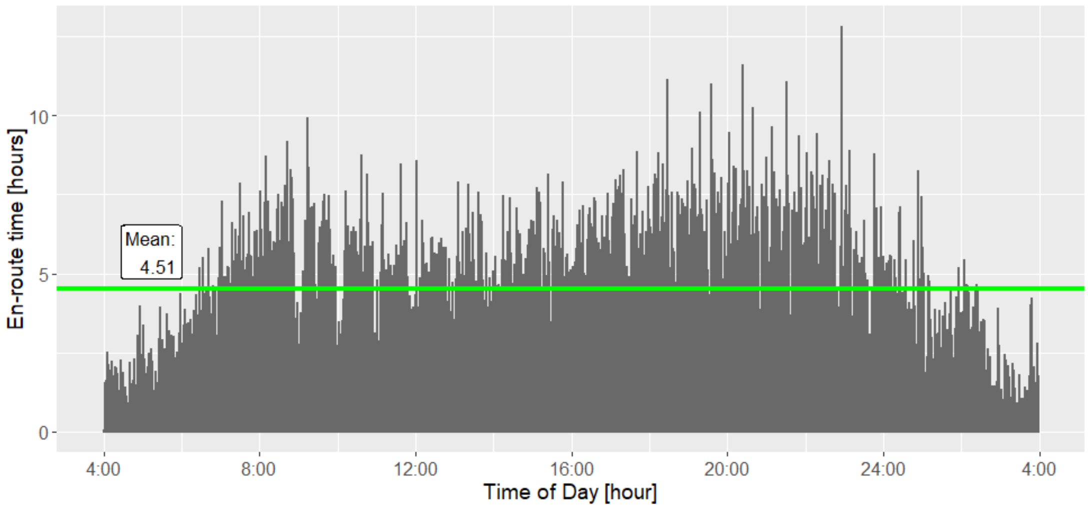

Figure A23: Temporal distribution of fleet total en-route time for c-greedy.

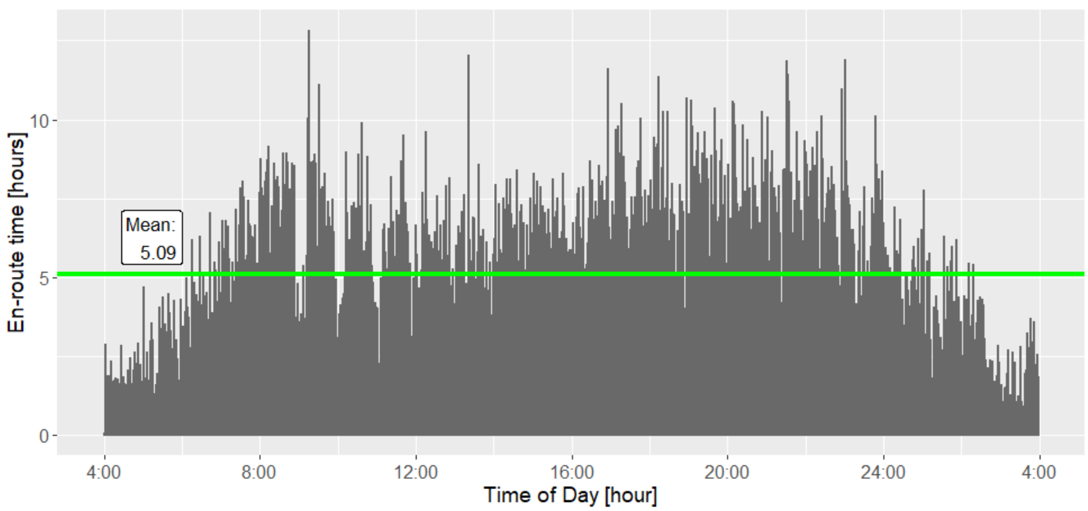

.5

Figure A24: Temporal distribution of fleet total en-route time for greedy.

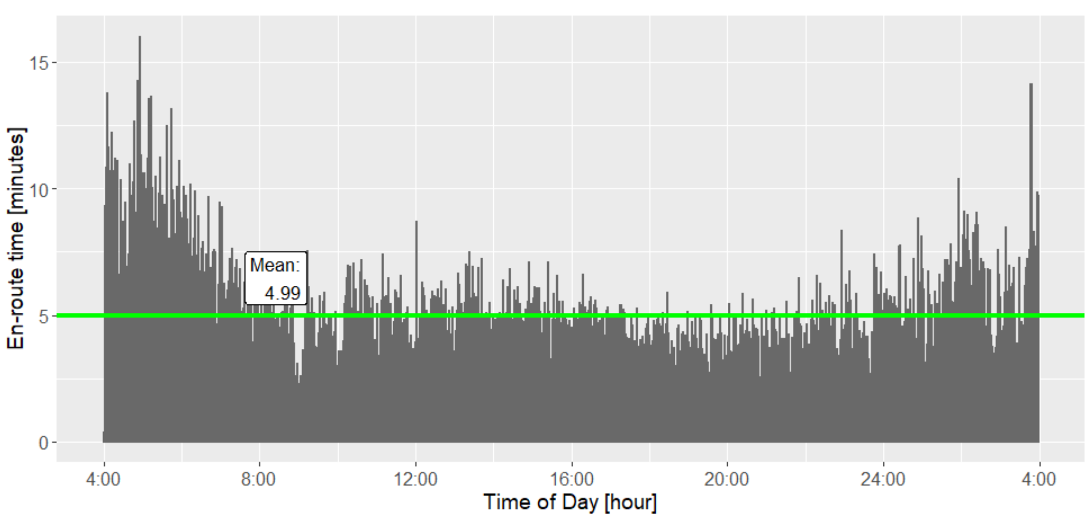

Figure A25: Temporal distribution of average en-route time per vehicle for c-greedy.

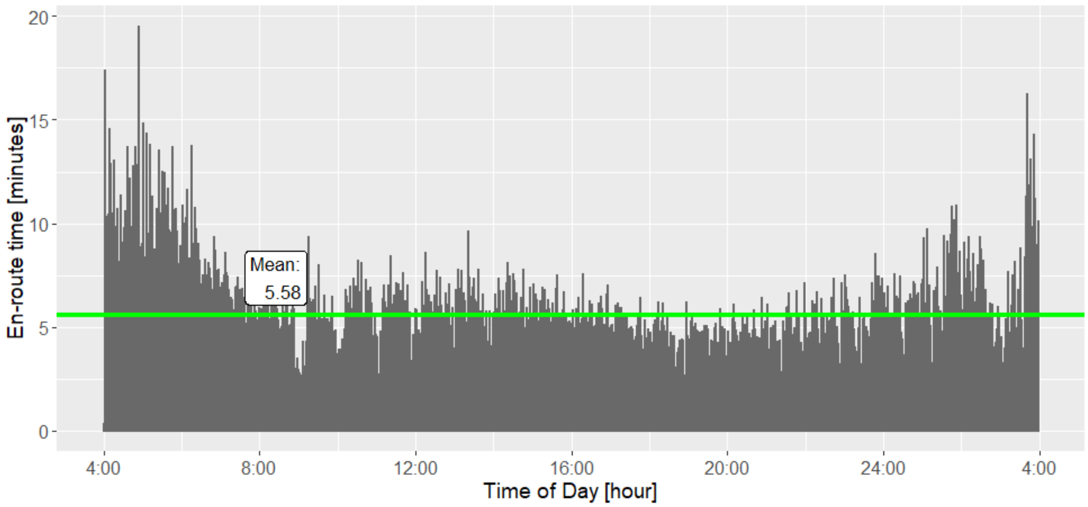

Figure A26: Temporal distribution of average en-route time per vehicle for greedy.

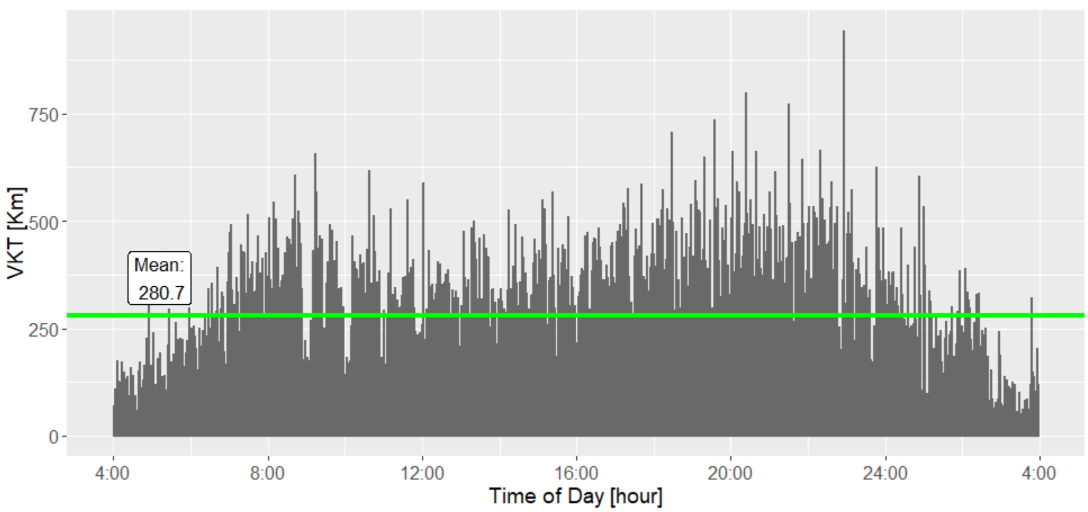

Figure A27: Temporal distribution of total fleet VKT for c-greedy.

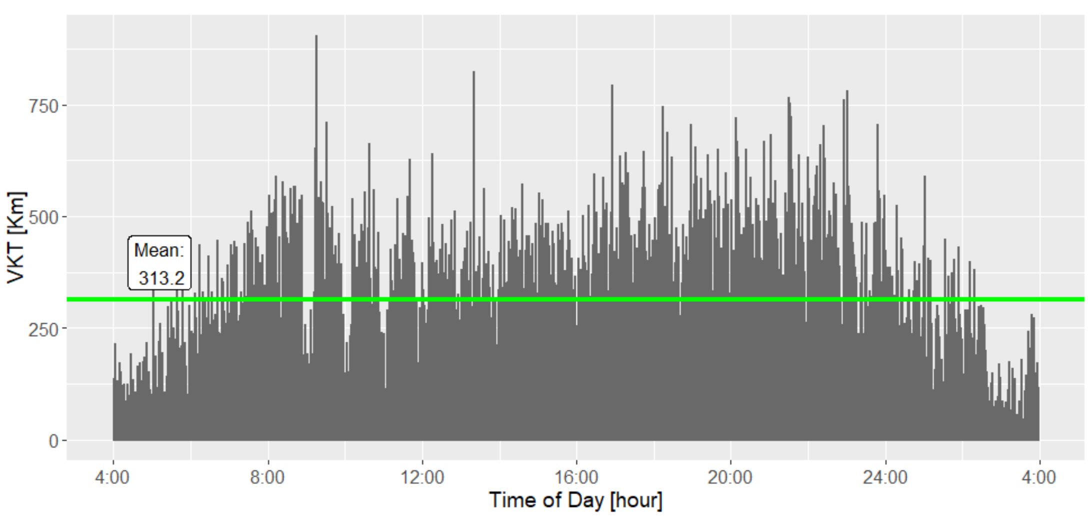

Figure A28: Temporal distribution of total fleet VKT for greedy.

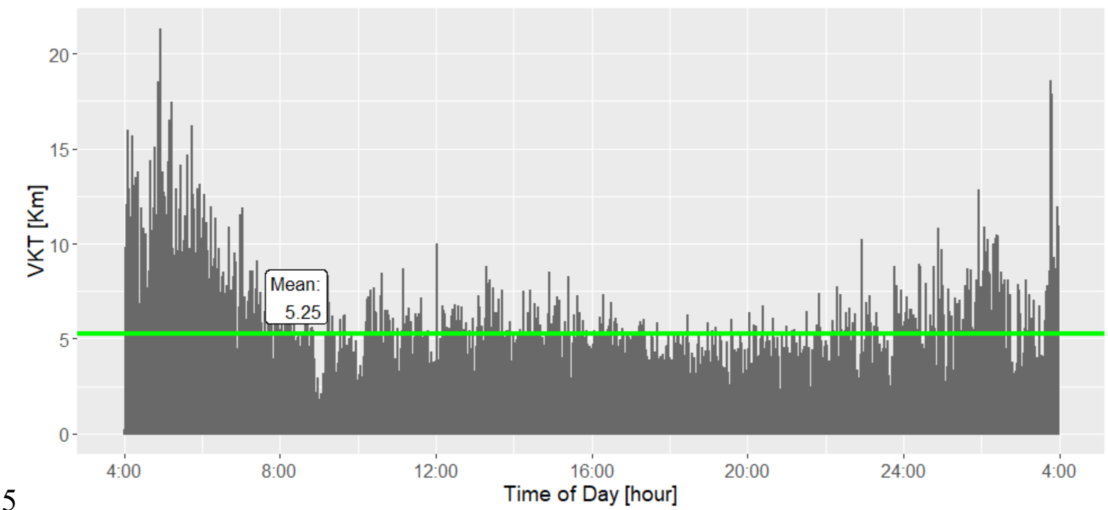

Figure A29: Temporal distribution of average VKT per vehicle for c-greedy.

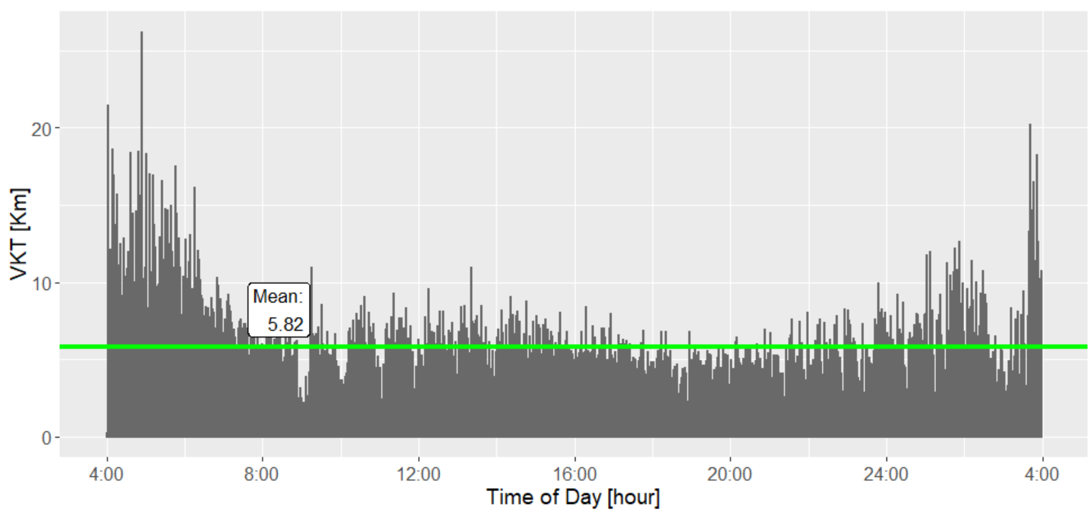

Figure A30: Temporal distribution of average VKT per vehicle for greedy.

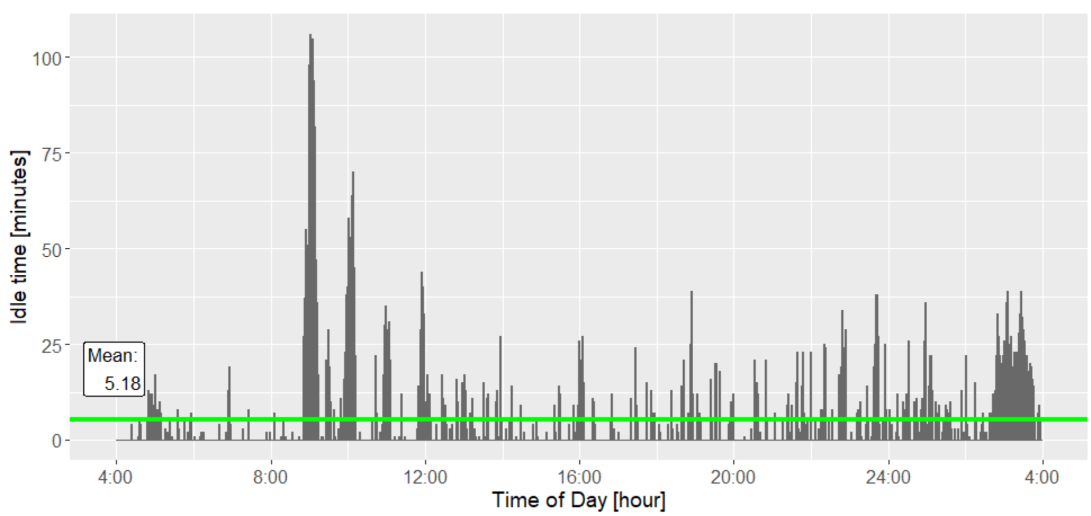

Figure A31: Temporal distribution of total fleet idle time for c-greedy.

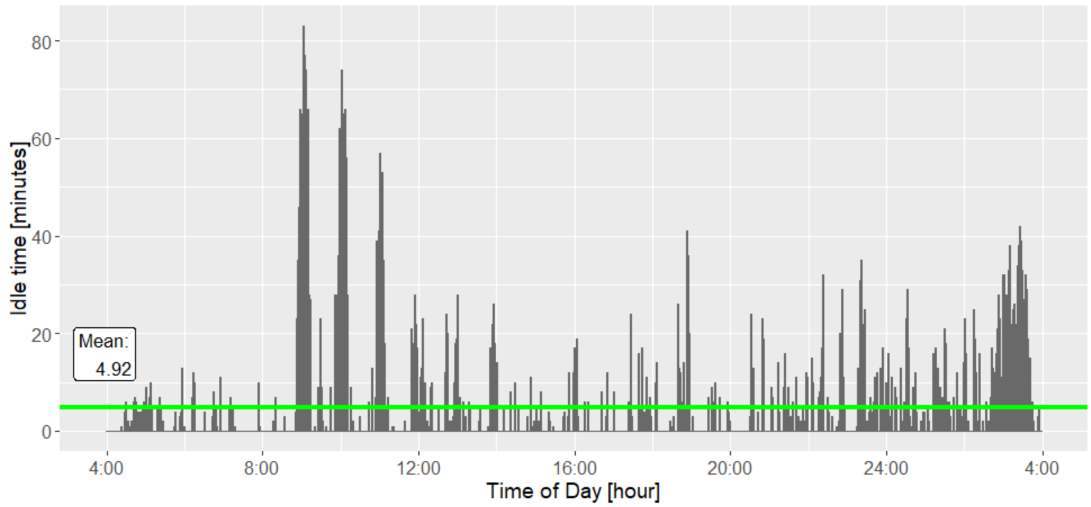

Figure A32: Temporal distribution of total fleet idle time for greedy.

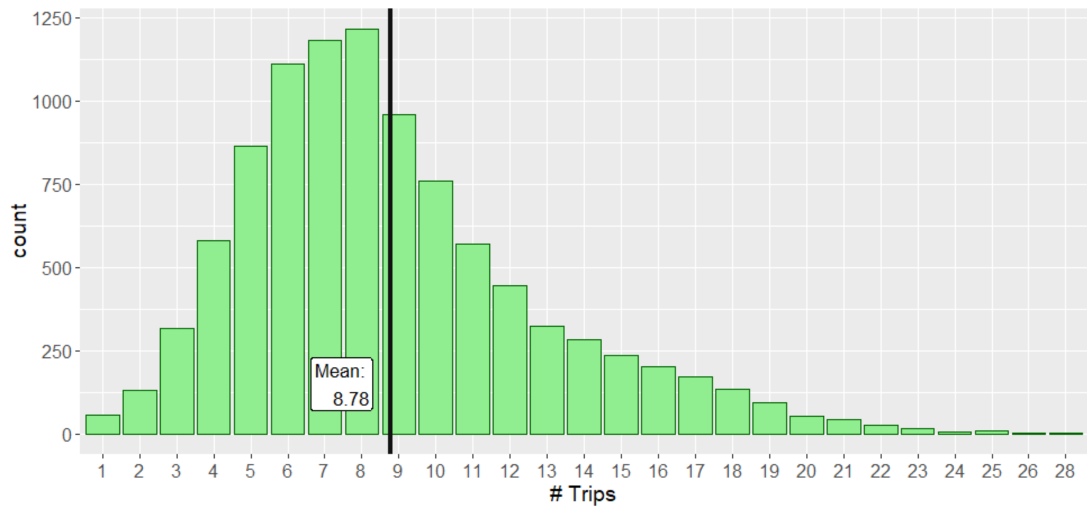

Figure A33: Number of trips per vehicle for c-greedy.

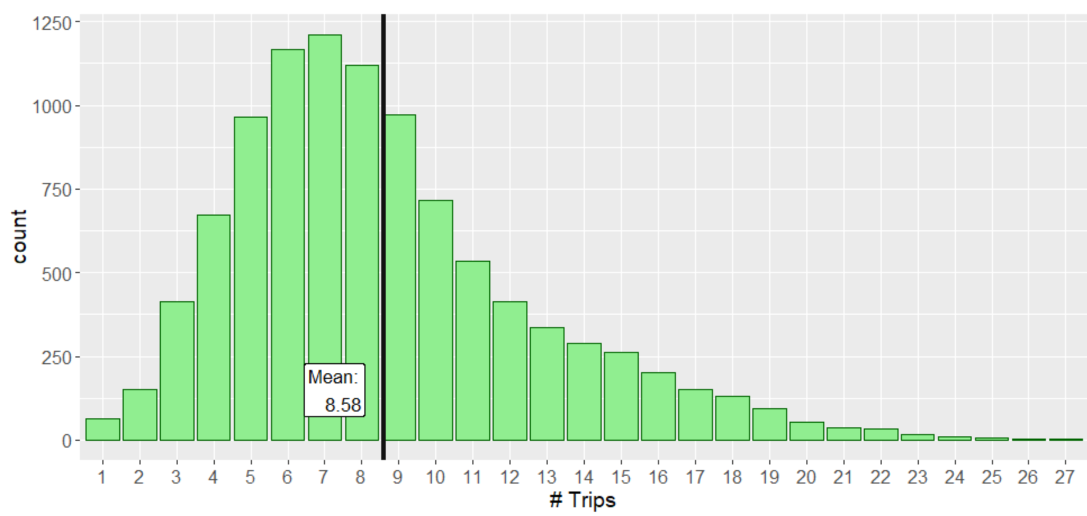

Figure A34: Number of trips per vehicle for greedy.

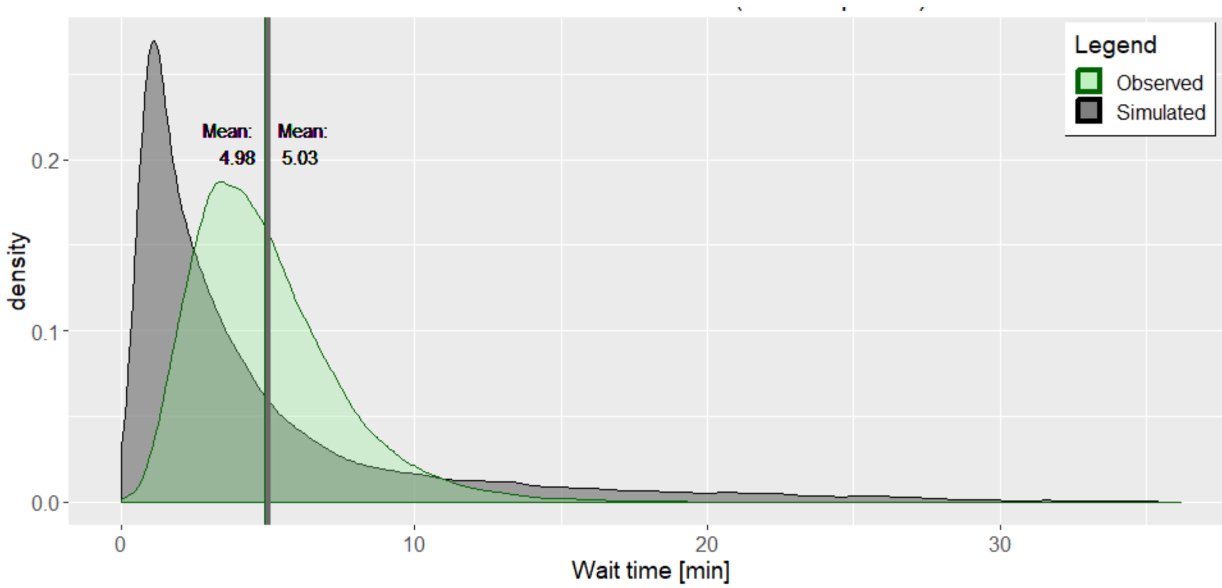

Figure A35: Comparison among distributions of simulated and observed wait times for c-greedy.

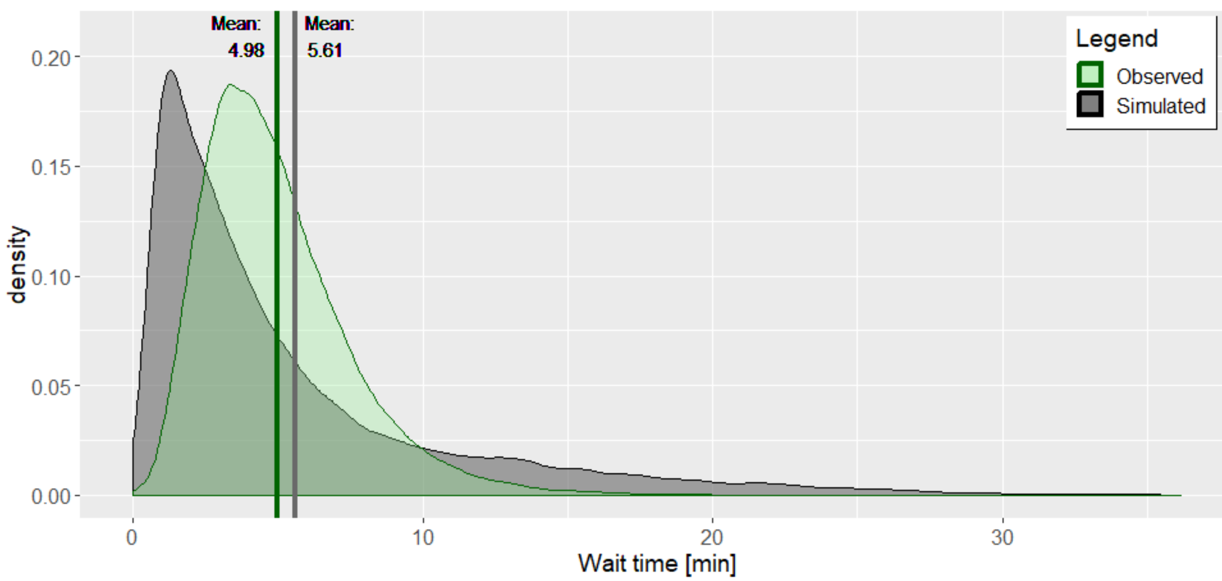

Figure A36: Comparison among distributions of simulated and observed wait times for greedy.

The improvement in terms of replicating observed wait times when deploying a conventional greedy algorithm is remarkable, as Figure A16 shows. However, this implementation of the model (greedy+1-min time intervals) is systematically generating shorter wait times than reality (simulated distribution is more skewed to the right). Moreover, the "tail" of the simulated distribution is wider.

Given the complex nature of the model and the fact that matching and time interval are mutually dependant model components, different time intervals must be tested to assess their effects. Hence, to continue with the analysis, the same matching algorithms are compared, but now with a 5-minute time step. Table A2 summarizes model metrics for this new comparison.

Table A4: Comparison of main model metrics under different matching algorithms: 5-min time interval.

Metric	C-Greedy	Greedy	Unit
Total number of trips to be served	86018	86018	Trips
Absolute fleet size deployed to provide service	9919	10064	Vehicles
Maximum en-route time (user wait time) observed in the simulation	1.68	1.15	Hours
Average number of idling vehicles (over 5-minute intervals)	5.49	5.62	Veh/5min
Percentage of time there were 0 vehicles idling in the system	84.03	82.99	Percent
Total time that the fleet is active, yet not in-service (Periods 1 & 2)	4452.4	5149.5	Hours
Percentage of time that fleet is active, yet not in-service (Periods 1 & 2), with respect to overall active time	16.98	18.79	Percent
Total VKT when the fleet is en-route (Period 2)	272211	302296	Kilometers
Percentage of VKT when fleet is en-route (Period 2), with respect to overall VKT	26.12	29.01	Percent
Average simulated wait time (**true observed value=4.98 minutes**)	5.59	6.01	Minutes
Trip connections that would be unfeasible with respect to observed data	16	14	Percent
Model runtime	8	5	Minutes

A remarkable difference between the metrics of Table A1 and Table A2 consists of a large improvement in metrics related with vehicles in Periods 1 and 2. Albeit not necessarily claiming that a 5-minute time interval is reflecting real-life operations, it certainly improves service provider efficiency. This is quite a predictable outcome since having more trip requests to match provides service operators with the opportunity to be more efficient.

The figures presented next are reduced to wait time distributions because the different time aggregation impose a scaling effect that renders plots incomparable in terms of value magnitudes. Overall patterns are stable though, with the exception of idling, as shown in Figures A17 and A18.

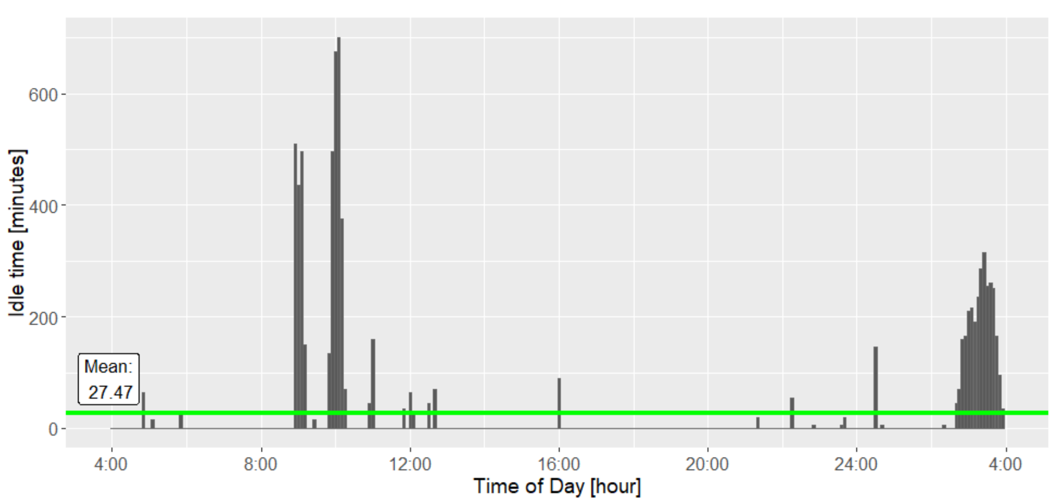

Figure A37: Temporal distribution of total fleet idle time for c-greedy and 5-min interval.

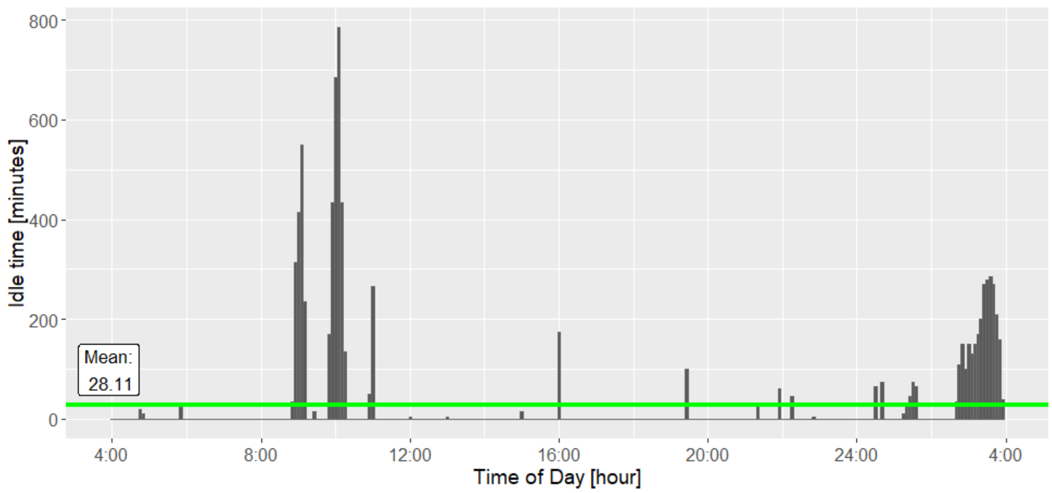

Figure A38: Temporal distribution of main system metrics for greedy algorithm with 5-min intervals.

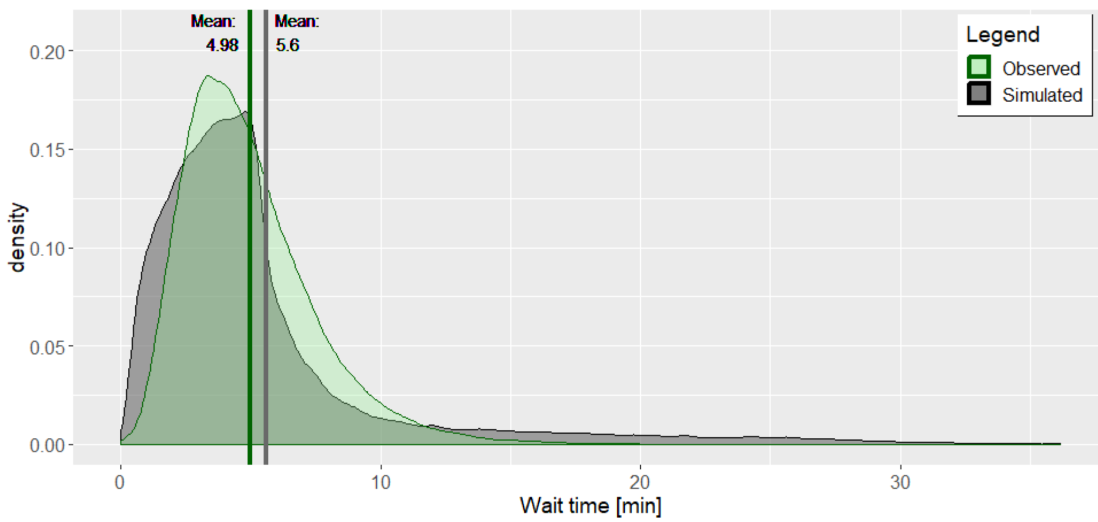

Figure A39: Comparison among distributions of simulated and observed wait times for c-greedy with 5-min intervals.

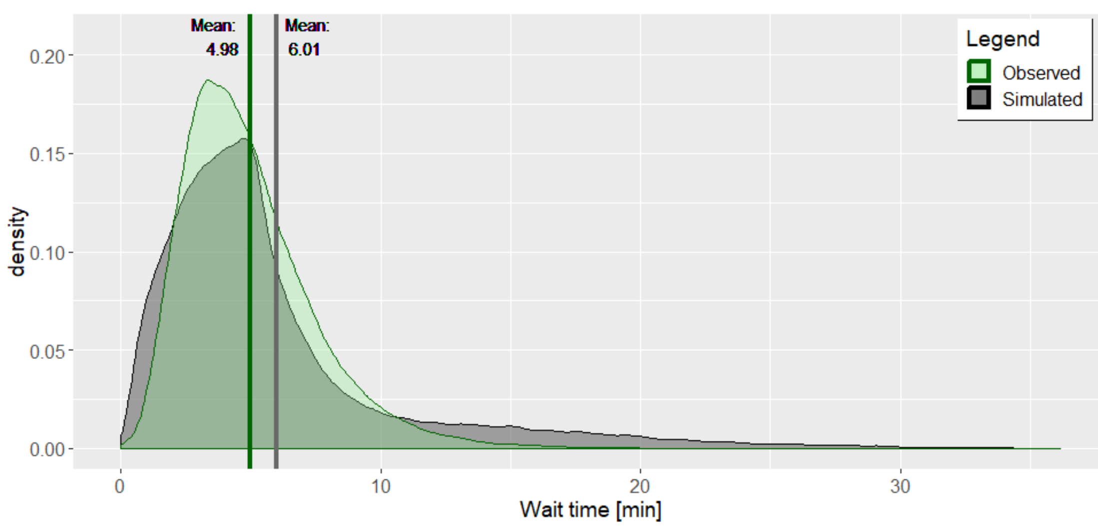

Figure A40: Comparison among distributions of simulated and observed wait times for greedy with 5-min intervals.

Based on Figures A19 and A20, the effect of time step size has been found to be a dominating factor with more influence than the matching algorithm. At this stage, the preferred model setting is a greedy matching coupled with a 5-minute time step. Its only disadvantage is still the thick tail of the distribution; however, it is possible that this is due to the current implementation of vehicle generation processes. As argued in Section 2, there are several factors that arguably influence the location within the service coverage in which drivers decide to start/re-start activity. This is on-going research and it can be enhanced considerably by the availability of additional data sources.

To be comprehensive in testing the model, and by recognizing that time interval is a critical component, the next testing efforts were directed towards different time interval sizes, as summarized in Table A3. Only a regular greedy matching is considered now, following the findings reported so far.

Table A5: Comparison of main model metrics under different time intervals.

Metric	**Greedy 30 sec**	**Greedy 1 min**	**Greedy 5 min**	**Unit**
Total number of trips to be served	86018	86018	86018	Trips
Absolute fleet size deployed to provide service	10220	10032	10023	Vehicles
Maximum en-route time (user wait time) observed in the simulation	1.28	1.22	1.27	Hours
Average number of idling vehicles (over 5-minute intervals)	4.88	4.87	5.03	Veh/time interval
Percentage of time there were 0 vehicles idling in the system	55.66	66.16	82.64	Percent
Total time that the fleet is active, yet not in-service (Periods 1 & 2)	8760.8	7481.0	5168.1	Hours
Percentage of time that fleet is active, yet not in-service (Periods 1 & 2), with respect to overall active time	28.25	24.77	18.85	Percent
Total VKT when the fleet is en-route (Period 2)	538970	453609	302296	Kilometers
Percentage of VKT when fleet is en-route (Period 2), with respect to overall VKT	43.22	39.06	29.01	Percent
Average simulated wait time (**true observed value=4.98 minutes**)	6.28	5.64	6.03	Minutes
Trip connections that would be unfeasible with respect to observed data	14	15	14	Percent
Model runtime	9	6.8	5.2	Minutes

It is useful to consider the results from Table A3 from a service provider efficiency perspective. In general, the longer the time interval, the more time a service provider can aggregate demand, and ultimately, the more efficient its operations can be. This can be observed in the 5-minute metrics where there is comparatively: higher percentage of time with 0 vehicles idling; less amount of time vehicles spent in Periods 1 and 2; and less VKTs while vehicles are in Period 2.

From a user perspective, the conclusions are very different. To begin with, longer time intervals imply that wait times would be (by model construction) half of the length of the time interval on average – assuming requests are uniformly distributed over the interval for the sake of argument. Accordingly, larger wait times are observed when comparing 1-minute to 5-minute time intervals, which reflect the argument above in that despite a more efficient provider, large time interval implies that matching results can be communicated to users only after the 5-minute period anyway.

In the case of shorter time intervals, a less efficient service operator (more myopic) is also detrimental to user wait times. The reason being that with too short time intervals, there are fewer agents in both vehicle and driver "pools", hence not much alternative matching possibilities exist to get the most of matching trade-offs. This in turn translates into users getting lower-quality options and hence subject to be longer in wait times, as shown in Table A3.

Wait time distribution plots are presented next in Figures A21 and 22. By comparing these against Figure A20, it can be argued that the latter is the one that replicates observed patterns more closely.

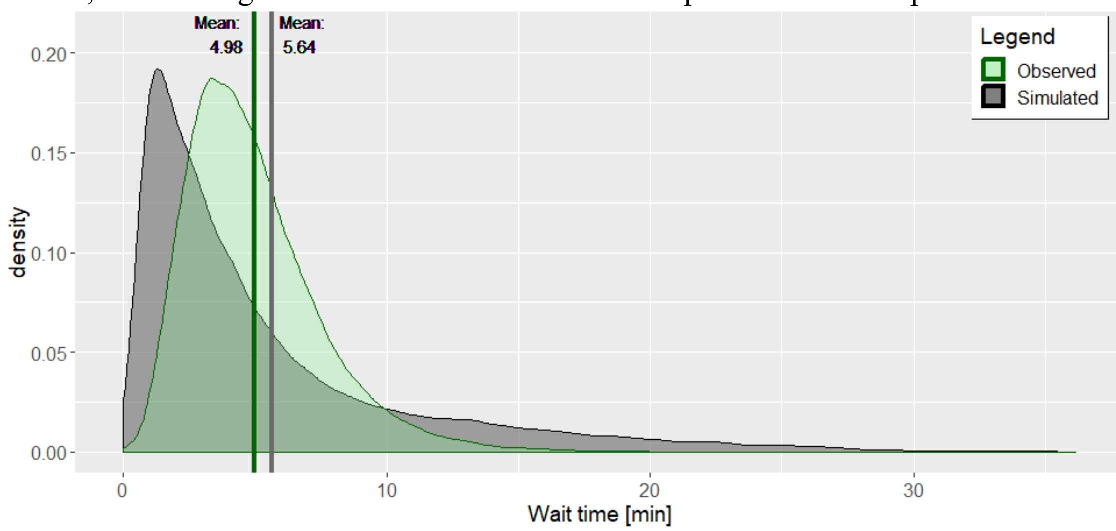

Figure A41: Comparison among distributions of simulated and observed wait times for greedy with 1-min intervals.

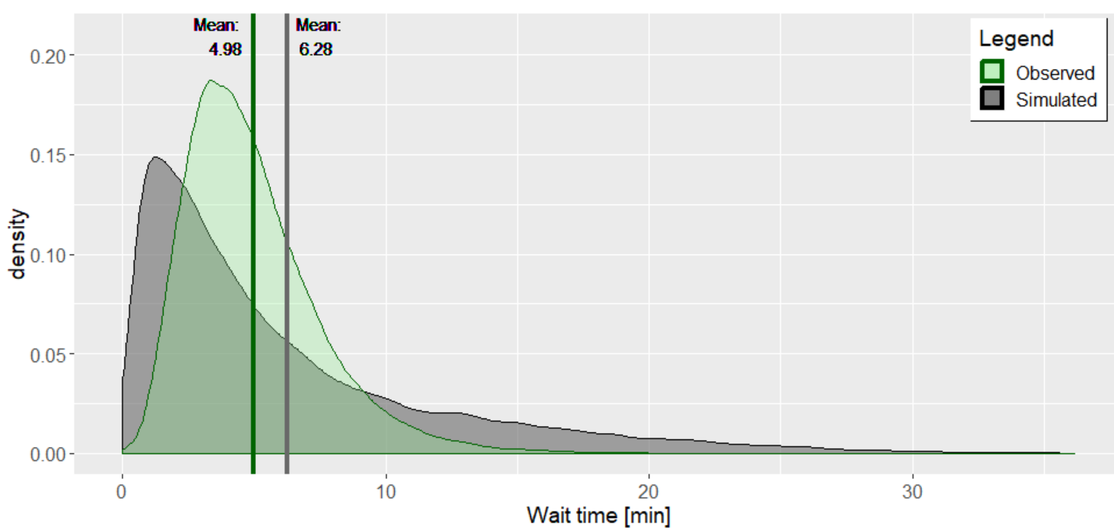

.7

Figure A42: Comparison among distributions of simulated and observed wait times for greedy with 0.5-min intervals.

As mentioned in Section 6 of this report, it should be noted that the tests carried forward for the model were only concerned with testing variations in the combination of different modelling components, as opposed to a calibration exercise. In this sense, the model results reported here are "raw" and have not been calibrated to any extent.

SENSITIVITY TO DIFFERENT DAYS

So far in the report, outputs and results have been shown only for Thursday March 30, 2017. Hence, the last tests consist of running the model for other days randomly chosen from the dataset. Note that a Tuesday, December 13, 2016 was also included for testing, following the findings and suggestions documented in Report 2 of the overall report series. All these days are run with the combination of greedy matching and 5-minute time interval, general metrics are shown on Table A4, and only temporal patterns and wait time distributions are shown.

Table A6: Comparison of main model metrics across different days.

Metric	Tues.	Wed.	Fri.	Unit
Total number of trips to be served	64042	55106	81375	Trips
Absolute fleet size deployed to provide service	8206	7533	10095	Vehicles
Maximum en-route time (user wait time) observed in the simulation	1.15	1.08	1.18	Hours
Average number of idling vehicles (over 5-minute intervals)	5.49	3.6	1.63	Veh/time interval
Percentage of time there were 0 vehicles idling in the system	80.9	81.25	90.63	Percent
Total time that the fleet is active, yet not in-service (Periods 1 & 2)	4163.5	3908	4960.3	Hours
Percentage of time that fleet is active, yet not in-service (Periods 1 & 2), with respect to overall active time	19.31	19.55	17.56	Percent
Total VKT when the fleet is en-route (Period 2)	243989	234340	299200	Kilometers
Percentage of VKT when fleet is en-route (Period 2), with respect to overall VKT	30.47	31.21	28.01	Percent
Average simulated wait time	6.28	6.67	4.43	Minutes
Trip connections that would be unfeasible with respect to observed data	16	15	18	Percent
Model runtime	4.2	3.7	5.2	Minutes

As Figures A23-28 clearly show, the model performs quite well for different days as well. Despite a somewhat lower accuracy in replicating observed wait time distributions, good performance can also be claimed for Fridays, whose demand patterns are clearly very different than other days. In sum, the prototype model has shown high resilience to withstand day-of-week variations.

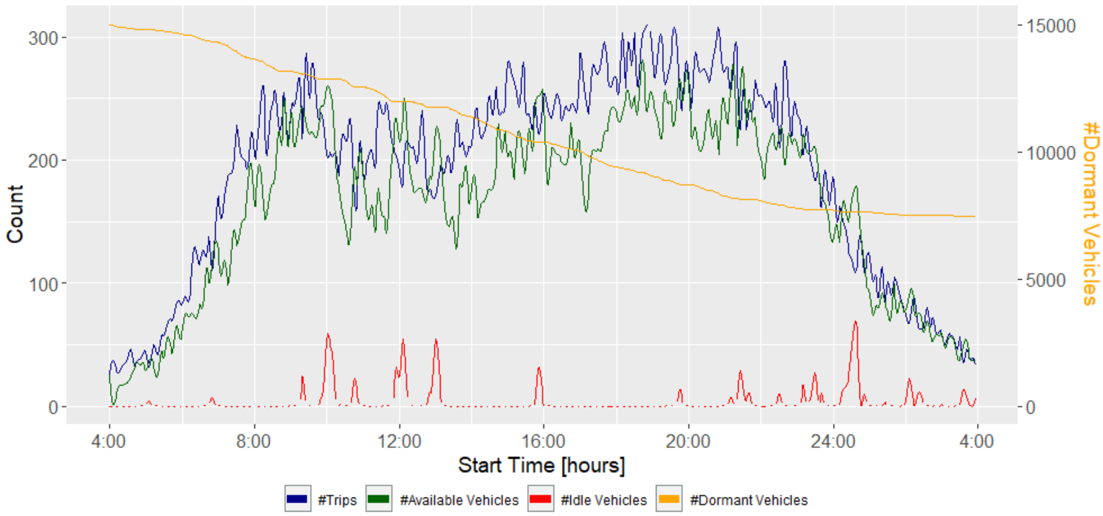

Figure A43: Temporal distribution of main system metrics for Wednesday 19/10/2016.

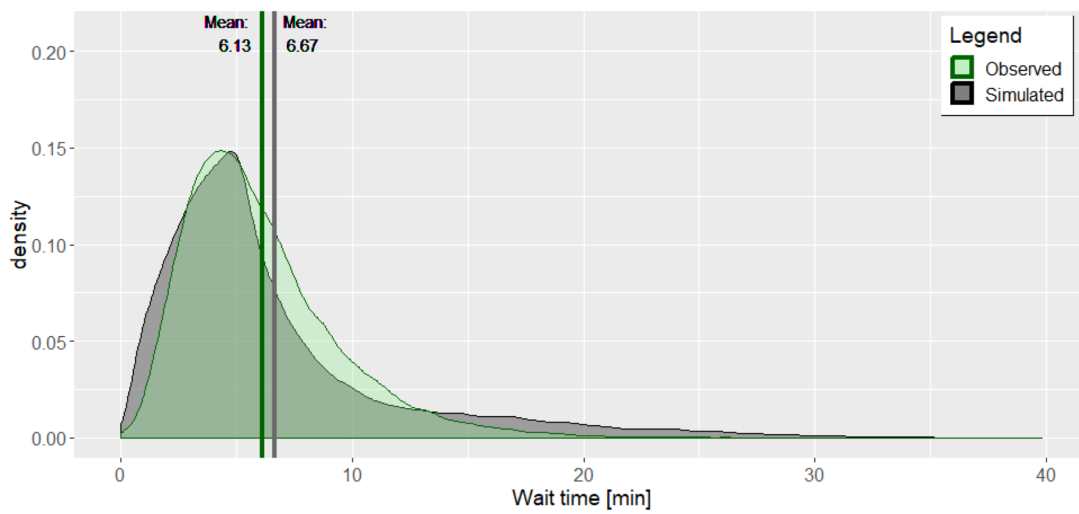

Figure A44: Comparison among distributions of simulated and observed wait times Wednesday 19/10/2016.

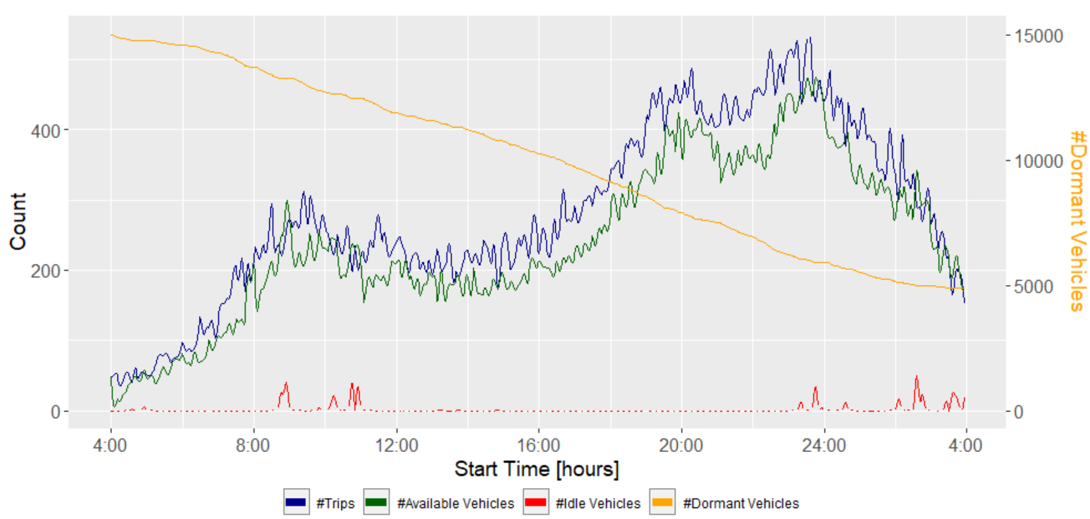

Figure A45: Temporal distribution of main system metrics for Friday 07/10/2016.

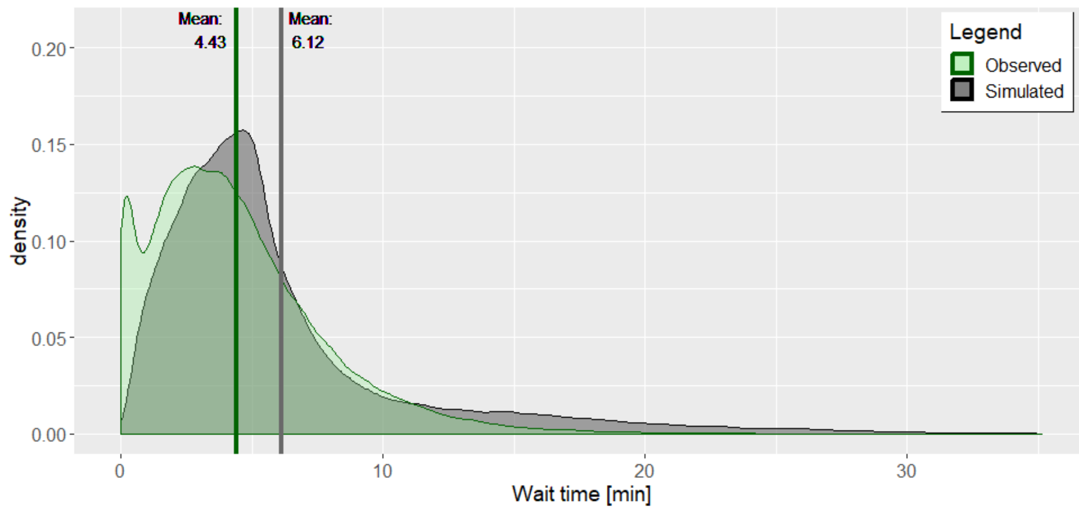

Figure A46: Comparison among distributions of simulated and observed wait times Friday 07/10/2016.

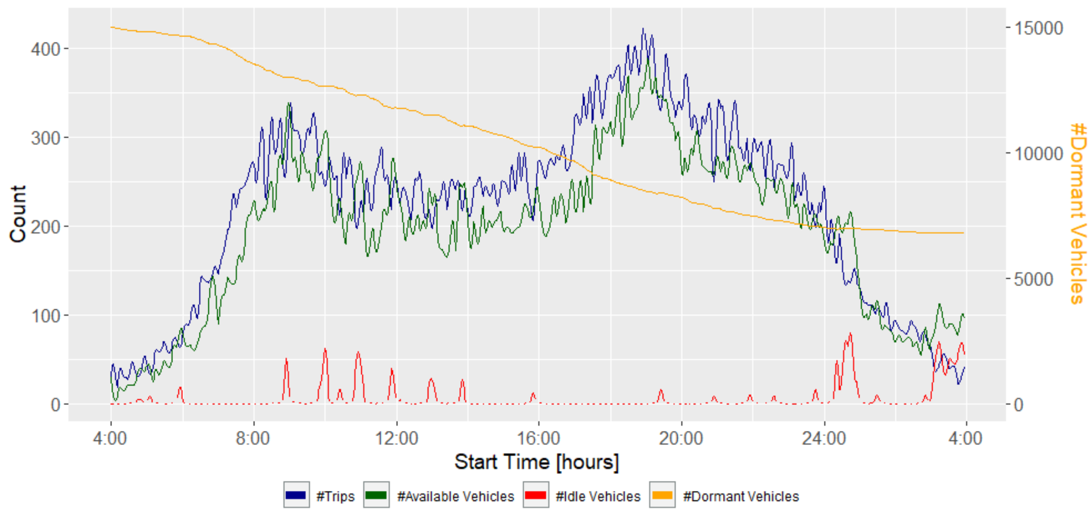

Figure A47: Temporal distribution of main system metrics for Tuesday 13/12/2016.

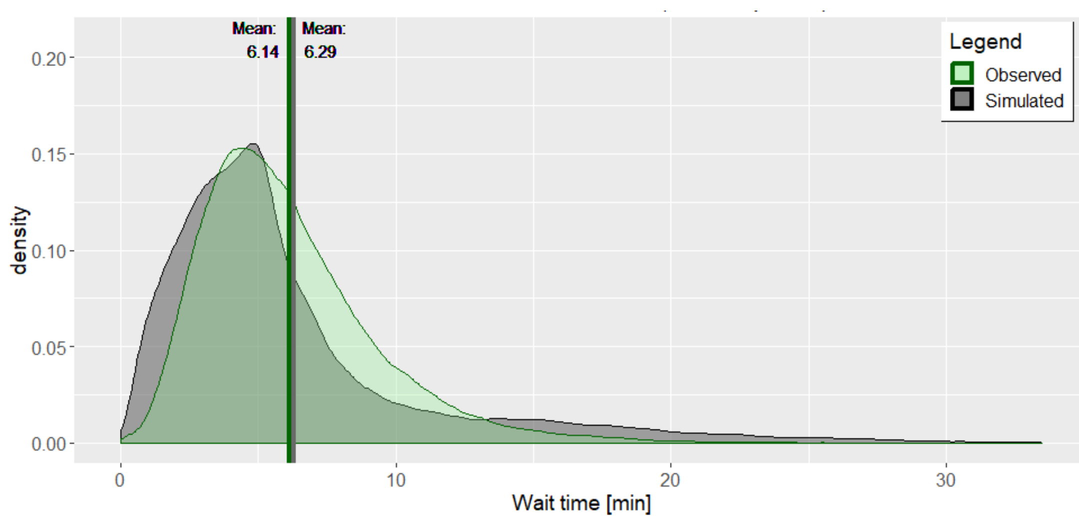

Figure A48: Comparison among distributions of simulated and observed wait times Tuesday 13/12/2016.