



HAL
open science

Traffic Uncertainty in On-Demand High-Capacity Ride-Pooling

Sebastian Hörl, Felix Zwick

► **To cite this version:**

Sebastian Hörl, Felix Zwick. Traffic Uncertainty in On-Demand High-Capacity Ride-Pooling. 101st Annual Meeting of the Transportation Research Board (TRB), Jan 2022, Washington D.C., United States. hal-03405574

HAL Id: hal-03405574

<https://hal.science/hal-03405574>

Submitted on 27 Oct 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

1 **Traffic Uncertainty in On-Demand High-Capacity Ride-Pooling**

2

3 **Sebastian Hörl**

4 Institut de Recherche Technologique SystemX

5 9 avenue de la Vauve, 91120 Palaiseau, France

6 ORCID number: 0000-0002-9018-432X

7 `sebastian.horl@irt-systemx.fr`

8

9 **Felix Zwick**

10 MOIA GmbH &

11 Institute for Transport Planning and Systems

12 ETH Zürich

13 Stefano-Francini-Platz 5, 8093 Zürich (Switzerland)

14 ORCID number: 0000-0002-4067-4486

15 `felix.zwick@ivt.baug.ethz.ch`

16

17

18 Word Count: 6274 words + 0 table(s) \times 250 = 6274 words

19

20 Submission Date: October 27, 2021

1 ABSTRACT

2 On-demand ride-pooling has been studied frequently in recent literature. However, real-world
3 applications of large-scale ride-pooling systems are rare and the breakthrough of comprehensive
4 systems is yet to come. In this study, we simulate a ride-pooling service based on real-world de-
5 mand data from Hamburg, Germany and evaluate the impact of traffic uncertainty on the system
6 performance using a state-of-the-art ride-pooling algorithm. We propose strategies to remedy cus-
7 tomer constraint violations on wait and travel times due to wrongly estimated travel times. We
8 show that uncertainty leads to a trade-off between violating customer constraints and rejecting re-
9 quests after prior acceptance.

10

11 *Keywords:* ride-sharing, on-demand mobility, MATSim, simulation, uncertainty

1 INTRODUCTION

2 On-demand mobility services have been established for many years and play an essential role in
3 urban mobility. Transport Network Companies (TNCs) such as Uber, Lyft or Didi offer convenient
4 and personalized transportation and match drivers and riders via smartphone app.

5 In recent years, ride-pooling has gained attention in science and practice as it promises
6 to increase system efficiency and vehicle occupancy and to relieve urban traffic. By transporting
7 multiple riders with a similar route in one vehicle, the number of vehicles and vehicle kilometers
8 traveled (VKT) can be drastically reduced (1, 2, 3) compared to on-demand ride-hailing where
9 only one passenger is transported at a time. This way, negative externalities of traffic such as air
10 pollutants, greenhouse gas emissions (4) or noise (1) can be reduced.

11 While on-demand services are already established today, they are assumed to gain even
12 more attraction with automated vehicles in place, as operational costs are expected to drop sub-
13 stantially (5, 6). As the use of on-demand mobility is expected to grow, so will its impact and
14 pooling potential. Therefore, it is critical to design an efficient system to get the most positive im-
15 pact on the transportation system. Multiple pooling algorithms have been developed that dispatch
16 passenger requests and vehicles and find matching trips (7, 8). A frequently used algorithm has
17 been developed by Alonso-Mora et al. (3), which we use and adapt in this study.

18 We implement the pooling algorithm into the multi-agent transport simulation MATSim
19 (9) and use historical demand from the ride-pooling operator MOIA in Hamburg as input for a
20 trip-based simulation. We show the negative impact of traffic uncertainty on the on-demand ride-
21 pooling system in terms of service quality and reliability, which are crucial for customer acceptance
22 in real-world applications.

23 BACKGROUND

24 Ride-pooling, often also named *on-demand ride-sharing*, has caught enormous attention in recent
25 years. Multiple pooling algorithms have been developed and applied in simulation studies to show
26 its impact on the transport system (10, 11, 12, 13). Transport simulations are a common tool to
27 mimic real-world mobility systems and evaluate the impact of changes in the system, such as new
28 transport policies or the introduction of a new mobility mode.

29 Overall, it has been shown that ride-pooling on a large scale has great potential to decrease
30 fleet sizes and VKT in urban areas compared to single-occupied mobility modes such as car or taxi
31 (2, 3, 1). This makes the system so interesting for future mobility applications and policymakers.
32 In reality, however, large-volume ride-sharing services still rarely exist and are mostly limited
33 to small test fleets. Foljanty (14) analyzed the world-wide introduction of new on-demand ride-
34 pooling services and determined an average fleet size of only five vehicles.

35 There are multiple reasons why the promising scientific results could not yet be transferred
36 to reality on a large scale. One major reason is the lack of automated vehicles, which makes the
37 operation of large ride-pooling fleets costly. Bösch et al. (5) identified the costs of (pooled) taxi
38 operations that drastically decrease with automated vehicles, as the main costs of current taxi fleets
39 are determined by the drivers' salaries. Another challenge of transferring the solutions presented
40 in simulation studies to reality is that simulations always simplify reality and ignore certain aspects
41 such as unexpected user behavior, operational challenges and travel time uncertainty, which highly
42 influence the system. The latter is further explored here.

43 We focus on the dispatching algorithm developed by Alonso-Mora et al. (3), which has
44 been applied and further developed in a wide range of publications. Ruch et al. (15) compared

1 the performance of four existing dispatching algorithms in two scenarios and showed that the one
2 developed by Alonso-Mora et al. (3) overall develops best in terms of customer matching rate
3 and VKT reduction. To further improve efficiency of the algorithm, Fielbaum et al. (16) consider
4 dynamic pick-up locations while taking into account the time it takes for customers to walk to the
5 meeting points. Engelhardt et al. (17) explore heuristic approaches for speeding up the customer-
6 vehicle matching process by early filtering of feasible but improbable combinations.

7 Several studies have looked at dynamic aspects of the algorithm. Dandl et al. (18) study
8 an approach for deciding early whether to accept or reject requests (rather than waiting until a
9 maximum wait time is reached). Fielbaum and Alonso-Mora (19) study in detail the effect of
10 re-assigning customers to new vehicles once the system has accepted their ride, and relate local
11 delays to the average wait time of the system. Furthermore, they emphasize the link between higher
12 rejection rates and increased reliability on planned pick-up and drop-off times if customer-vehicle
13 assignments are binding. The paper lists a number of other sources of uncertainty, one of them
14 uncertainty about travel times, which are examined in the present paper.

15 The contributions of this paper are the following:

- 16 • We implement the algorithm by Alonso-Mora et al. (3) as part of the transport simulation
17 framework MATSim.
- 18 • We formally analyze the algorithm in presence of travel time uncertainty and propose strate-
19 gies to operate it under such conditions.
- 20 • We underline the relation of customer constraint violations and rejections under traffic un-
21 certainty, which is highly relevant for real-world applications.

22 **METHODOLOGY**

23 In the following section, we will give an overview of the algorithm introduced by Alonso-Mora
24 et al. (3). While the specifics of the algorithm are presented in the given reference, we will focus on
25 how varying travel times affect the functioning of the approach. Afterwards, we provide mitigation
26 strategies for the problems that arise when considering travel time uncertainty.

27 **Algorithm overview**

28 The algorithm presented by Alonso-Mora et al. (3) traverses multiple steps to perform “any-time
29 optimal” assignments between customers of an on-demand service and its vehicles. The assign-
30 ments are optimal at a given time, but not necessarily when looked at over a longer time horizon.
31 In every decision epoch (usually every 30 seconds), a request graph is constructed between all
32 requests that have been submitted since the last epoch or which are currently assigned to but not
33 yet picked up by a vehicle. To do so, each request is compared to each other in terms of their latest
34 pick-up time and their latest drop-off time. These are calculated based on a maximum wait time
35 and an allowed detour compared to the unshared ride time, which are configuration parameters of
36 the algorithm. An edge is added to the request graph if two requests can be served by an imaginary
37 vehicle that would be available at the present time, either at the origin of the first request, or at
38 the origin of the second request, while adhering to the pick-up and drop-off constraints of both
39 requests.

40 The second step of the algorithm makes heavy use of a *routing function*. Given the current
41 location of a vehicle, its on-board requests, and a list of requests that may be served by the vehicle,
42 the *routing function* will return a sequence of pick-up and drop-off tasks for the vehicle which
43 minimizes the travel delay of all involved requests. The travel delay is defined as the difference

1 between the time at which the customer would have arrived on an unshared trip and the drop-off
 2 time that is predicted on the proposed route of the vehicle. In some cases, the routing function will
 3 indicate that no feasible route for the given list of requests can be found.

4 We consider a number of requests referenced by index k with their latest pick-up and drop-
 5 off times $t_{k,pickup}^{latest}$ and $t_{k,dropoff}^{latest}$. Furthermore, for each request, the unshared (direct) arrival time
 6 is known as $t_{k,direct}$. Another important concept is the “active” pick-up time. Initially, after sub-
 7 mission of the request, we have $t_{k,pickup}^{active} = t_{k,pickup}^{latest}$. However, upon first assignment of the request,
 8 the pick-up constraint will be set to the time that is expected based on the assigned route. This
 9 resembles that the operator communicates the expected pick-up time to the customer and needs to
 10 stick to the offer in subsequent decision epochs.

11 Given a vehicle and a set of requests, whenever the *routing function* is called, it will try
 12 to find a sequence of pick-up and drop-off tasks for the vehicle which adheres to a number of
 13 constraints. While many approaches are possible to construct these sequences, without loss of
 14 generality we follow the approach of Alonso-Mora et al. (3) by performing an exhaustive search
 15 for the optimal sequence for vehicles that have up to four passengers on-board, and we perform an
 16 insertive search (keeping the order of on-board request tasks invariant) for higher occupancy. Given
 17 the network travel times, we can calculate expected pick-up and drop-off times for all involved
 18 requests as $\hat{t}_{k,pickup}$ and $\hat{t}_{k,dropoff}$. For every tested chain, apart from the capacity constraint, the
 19 following conditions need to be met for the sequence to be feasible:

$$\hat{t}_{k,pickup} \leq t_{k,pickup}^{active} \quad (1)$$

$$\hat{t}_{k,dropoff} \leq t_{k,dropoff}^{latest} \quad (2)$$

20

21 Note in order for a sequence to be feasible, these constraints need to be fulfilled for both requests
 22 that are to be assigned, and those which have already been picked up by the vehicle. Otherwise,
 23 their arrival time would be freely moved while integrating new requests.

24 If the sequence is feasible, the cost is calculated as:

$$C = \sum_k \max\{ 0, \hat{t}_{k,dropoff} - t_{k,direct} \} \quad (3)$$

25

26 Using the routing function, a list of groups, each consisting of a vehicle, a list of assignable re-
 27 quests, and an associated cost (cumulative delay) is constructed. To do so, first, direct vehicle-
 28 request matches are evaluated. The routing function is evaluated for each vehicle and each request
 29 contained in the request graph. In case a feasible route can be found, the match between the re-
 30 spective vehicle and request is added to the list. Afterwards, all matches between one vehicle and
 31 two requests are evaluated. For that, all existing vehicle-request matches are considered, and for
 32 each group the request graph is used to find all other requests that may be combined with the al-
 33 ready matched request. For each potential match, the routing function is evaluated again to check
 34 if a feasible solution exists. If this is the case, the match between a vehicle and two requests is
 35 added to the list with the respective cost. The algorithm continues, for each vehicle, by examining
 36 the existing matches and by finding feasible combinations of them to construct routes with three,
 37 four, and more assignable requests. Note that there is a limit to the size of the groups as longer
 38 sequences mean longer travel times which, eventually, will not allow adding more requests due
 39 to their constraints, but also due to the capacity of the vehicle which is considered in the *routing*

1 *function*. At the end, a list of potential groups with their costs is available.

2 The list of groups is then passed on to the selection stage, in which an Integer Linear
3 Program (see exact definition in Alonso-Mora et al. (3)) is used to select a set of routes to assign
4 to the vehicles in the current time steps. The problem ensures that:

- 5 • Each pending or previously assigned request is present in the selected list of routes *at most*
6 once (otherwise, it would be assigned to multiple vehicles).
- 7 • Each vehicle is presented in the selected list of routes *at most* once (otherwise, there are
8 multiple groups assigned to a vehicle).

9 The objective is to minimize the cumulative cost (delay) of the selected routes. On top,
10 the objective contains a penalty term for every pending request that is not assigned, and a term for
11 every currently assigned request that is not assigned again. These penalties are usually very high to
12 avoid rejecting requests in the first place (*early reject*), but especially after they have been accepted
13 already by the operator (*late reject*):

$$J = \sum_s x_s \cdot C_s + \xi_{\text{early/late}} \cdot (1 - y_k) \quad (4)$$

14 Here, x_s is a binary variable indicating whether group s is selected, and y_k tells whether request k
15 has been assigned at all.

16 After solving the ILP, the selected routes are assigned to the vehicles. Vehicles, which are
17 not assigned a route (anymore) are treated differently. As they may still have on-board requests,
18 which need to be dropped off, the *routing function* is used to calculate a new minimum-delay route
19 for the vehicle to drop off the remaining requests. In case no requests are assigned to a vehicle at
20 all, no special steps need to be taken.

21 After the assignment process, Alonso-Mora et al. (3) propose to consider the requests that
22 could not be assigned in the current epoch and to relocate idling vehicles towards those locations.
23

24 **Varying travel times**

25 The study presented by Alonso-Mora et al. (3) makes a strong assumption on travel times, i.e. they
26 are not varying over the day. Hence, given an origin location and a destination location, the travel
27 time is assumed to stay constant throughout the simulations. This has a couple of implications:

- 28 1. As long as two requests are not rejected or their constraints are violated by the bare progres-
29 sion of time, they remain in the request graph.
- 30 2. As long as requests remain in the vehicle graph, they remain in the group list for a vehicle
31 that does not move.
- 32 3. A group of requests and its corresponding route that has been assigned to a vehicle will
33 always remain feasible in the following decision epoch, and its cost will stay the same.

34 We now introduce varying travel times. In this context, the travel time between two distinct
35 points in the network can change between two decision epochs. Such changes may be caused
36 by targeted traffic control or congestion. All three corollaries stated above do not hold in such a
37 setting:

- 38 1. If travel times increase between two epochs, two requests that may otherwise have remained
39 as a poolable pair in the request graph, can disappear.
- 40 2. If the relevant requests were involved in the assignable group list, a range of groups may
41 disappear in consequence. Note that this poses a filter even before a feasible route has been
42 found for a group.

1 3. For the remaining groups, the routing function may now find that formerly feasible sequences
2 are now infeasible and hence minimum-delay routes may change.

3 Note that the first two effects (1 and 2) strongly affect the rejection behavior of the algo-
4 rithm. While in the context of invariable travel times, rejections happen before assignments, there
5 is now a much higher chance that requests fall out of the request graph and the group list after as-
6 signment. This forces the algorithm to perform *late rejections*, which are highly undesirable from
7 the customers' perspective.

8 The third effect, however, is more severe as it leads the algorithm into an undefined state.
9 As mentioned before, the routing function only tags a route as feasible if the constraints for all
10 assigned and on-board requests of a vehicle are met. While requests that would be picked up late
11 can easily be late-rejected, on-board requests for which their drop-off constraints are violated may
12 lead the algorithm to not finding *any* feasible sequence of drop-off tasks of the remaining requests.
13 This is an undefined condition, as vehicles should not go into idle mode as long as they carry
14 passengers.

15 The two major challenges when using varying travel times with the algorithm from Alonso-
16 Mora et al. (3) are hence, (1) undefined system states when dropping off customers, and, (2) late
17 rejections.

18 **Solution strategy**

19 To mitigate the problems arising with varying travel times, some exploration work has been per-
20 formed, which has led to the following proposals.

21 To overcome the issue of undefined system states, a simple policy is proposed: (A) For the
22 common assignment process, no adjustment is made. Based on the default logic, a vehicle, which
23 cannot satisfy the drop-off constraints of its on-board requests will not appear in any assignable
24 group in the selection stage, and hence, no assignable requests will be attached to it. However,
25 when the "unassigned" route for the vehicle is to be calculated, drop-off constraints are now ig-
26 nored. Note that the *routing function* still selects the route with the shortest cumulative delay.
27 Hence, the vehicle will aim to drop off the on-board customers as quickly as possible and not en-
28 gage in picking up any new passengers until they are all dropped off or their constraints can be
29 fulfilled again.

30 While such a strategy will make the algorithm runnable, it can lead to a high number of late
31 rejections. The second strategy (B) is less restrictive. At the beginning of each decision epoch, we
32 examine the currently assigned routes of all vehicles. Although pick-up and drop-off times may
33 have been planned differently in the previous epoch, we can now calculate an updated expectation
34 of the future drop-off times by traversing the routes and summing up travel times. This allows us
35 to calculate an *updated* drop-off time $t_{k,\text{dropoff}}^{\text{updated}}$ for each request, which is used instead of the latest
36 drop-off time for on-board requests that would otherwise violate their constraints. Effectively,
37 this means that the algorithm is allowed to construct routes in which new pick-up tasks are added
38 to a vehicle as long as the on-board requests are not dropped off later than what is the absolute
39 minimum, given current traffic conditions. To favor feasible solutions, a penalty is added for every
40 second of violating a drop-off constraint.

41 Finally, (C) the concept is applied analogously to requests that have been assigned to a
42 vehicle and are currently waiting to be picked up. In such a case, we can calculate $t_{k,\text{pickup}}^{\text{updated}}$ based
43 on the current traffic conditions and used it instead of the active pick-up time of the request for
44 all routes which re-assign the request to the vehicle it is currently assigned to. Less formally, this

1 means that increased traffic will not be a reason for performing a late rejection on the request.
 2 However, if another vehicle is able to pick up the request based on its original active pick-up time,
 3 the reassignment should happen. For that reason, again, a penalty is added to the cost of every
 4 route with violations.

5 Formally, we introduce two new quantities to the algorithm, $t_{k,\text{pickup}}^{\text{required}}$, and $t_{k,\text{dropoff}}^{\text{required}}$ which are
 6 calculated depending on the strategy. By default, we have:

$$t_{k,\text{pickup}}^{\text{required}} = t_{k,\text{pickup}}^{\text{active}} \quad (5)$$

$$t_{k,\text{dropoff}}^{\text{required}} = t_{k,\text{dropoff}}^{\text{latest}} \quad (6)$$

7
 8 For strategy (B), for all requests that are assigned to the currently evaluated vehicle:

$$t_{k,\text{dropoff}}^{\text{required}} = \max \left\{ t_{k,\text{dropoff}}^{\text{updated}}, t_{k,\text{dropoff}}^{\text{latest}} \right\} \quad (7)$$

9

10 Additionally, for (C), for all requests that are assigned to the currently evaluated vehicle:

$$t_{k,\text{pickup}}^{\text{required}} = \max \left\{ t_{k,\text{pickup}}^{\text{updated}}, t_{k,\text{pickup}}^{\text{active}} \right\} \quad (8)$$

11

12 The constraints that are evaluated in the *routing function* are then:

$$\hat{t}_{k,\text{pickup}} \leq t_{k,\text{pickup}}^{\text{required}} \quad (9)$$

$$\hat{t}_{k,\text{dropoff}} \leq t_{k,\text{dropoff}}^{\text{required}} \quad (10)$$

13

14 The objective is modified such that it contains a penalty term for constraint violations:

$$C' = C + \alpha \cdot \left(\sum_k \max \{ 0, \hat{t}_{k,\text{pickup}} - t_{k,\text{pickup}}^{\text{active}} \} + \sum_k \max \{ 0, \hat{t}_{k,\text{dropoff}} - t_{k,\text{dropoff}}^{\text{latest}} \} \right) + \beta \quad (11)$$

15

16 While factor α penalizes the time of violation, β is a constant that is added *if* there have been
 17 *any* violations. The latter parameter should be put to a value which clearly distinguishes feasible
 18 sequences from violating sequences, such that when the routing function evaluates all possible
 19 configurations, it will favor those where no violations take place. Both values should be chosen
 20 in a way such that they generally do not exceed the penalties on unassigned requests in the ILP
 21 objective described in Equation 4.

22

23 Additionally, for (A), we modify the *routing function* such that it ignores drop-off violations
 when examining a potential route that *only* contains on-board requests.

24

25 Finally, note that requests may still be rejected late due to issue (1) noted above. To avoid
 26 that requests are rejected due to them disappearing from the request graph, we forcefully construct
 27 a potential route based on the requests that are currently assigned to a vehicle and add it, if not
 already present, to the list of potential routes with its respective cost.

1 SIMULATION SETUP

2 The following sections describe a simulation environment in which the strategies mentioned above
3 will be tested. First, detailed on the simulation framework based on MATSim are given, second, the
4 relevant data sources of the use case are introduced, and third, the experimental design is covered.

5 MATSim implementation

6 The algorithm by Alonso-Mora et al. (3) has been implemented as a new extension to the multi-
7 agent transport simulation framework MATSim (9)¹. The framework allows to flexibly combine
8 numerous components around the simulation and analysis of transportation systems. Recently, the
9 framework has been used to study the efficiency of on-demand systems (20, 21, 22).

10 The algorithm by Alonso-Mora et al. (3) is implemented directly as a replacement for
11 the fleet optimizer in the DRT (Demand Response Transit) package, which is commonly used to
12 simulate on-demand vehicle services in MATSim. Being part of MATSim leads to a range of
13 aspects under which the algorithm has not been operated so far:

- 14 • Travel times can vary over the day. Such varying travel times can either be imposed directly
15 on the network links, or they may emerge dynamically from the interactions of all travelling
16 agents.
- 17 • Stop times are finite as it takes time for customers to enter or leave the vehicle. Although
18 configurable, by default, a stop time of 60 seconds is used, which is in line with observations
19 from MOIA.
- 20 • Vehicles are bound by physical constraints. If a vehicle is currently traversing a network
21 link, it cannot simply perform a U-turn or change route at any time, but only once it arrives
22 at the next network node.

23 While the first point has been discussed in detail above, the two other points mainly affect
24 the calculation of stop departure and arrival times. Generally, passengers leave the vehicle at the
25 beginning of a stop (i.e. directly after the vehicle has arrived at the destination), while passen-
26 gers enter the vehicle after the finite stop duration. This duration hence needs to be taken into
27 account when examining pick-up constraints and when calculating times for subsequent stops in a
28 sequence.

29 The assignment based on (3) is executed in fixed intervals, here 30 seconds. Between these
30 decision epochs, all incoming requests are collected in a queue and then processed jointly in the
31 assignment step. Note that, theoretically, when varying travel times are used, we would need to
32 route a large number of movements in the fleet to construct the group lists and estimate pick-up
33 and drop-off times to verify the constraints. This quickly leads to millions of routing tasks per
34 decision epoch and, hence, renders the algorithm computationally heavy. To reduce some of the
35 computational load, we keep a cache of origin-destination relations. Once a travel time between an
36 origin and a destination link is needed, it is first checked whether travel time information for this
37 relation is already in the cache. If so, the respective value is returned. Along with the calculated
38 travel time, we keep the last timestamp in simulation time at which the relation was routed. If the
39 lifetime of a relation has exceeded a specific threshold, it is discarded and re-routed. It is hence
40 possible to either have a rough approximation of the travel time (if the lifetime is high) or very
41 detailed information (if the lifetime is low).

¹The code will be contributed to the open source code base of MATSim.

1 Data preparation

2 For the simulation, we use the stop network and demand data from Europe’s largest ride-pooling
 3 provider MOIA in Hamburg, Germany. MOIA operates with up to 500 vehicles in a 300 km²
 4 service area covering most populated areas of the city shown in Figure 1. Although the input
 5 data reflects the real-world service, it should be noted that the ride-pooling simulation, the used
 6 algorithms and the results only remotely resemble MOIA’s real-world operation.

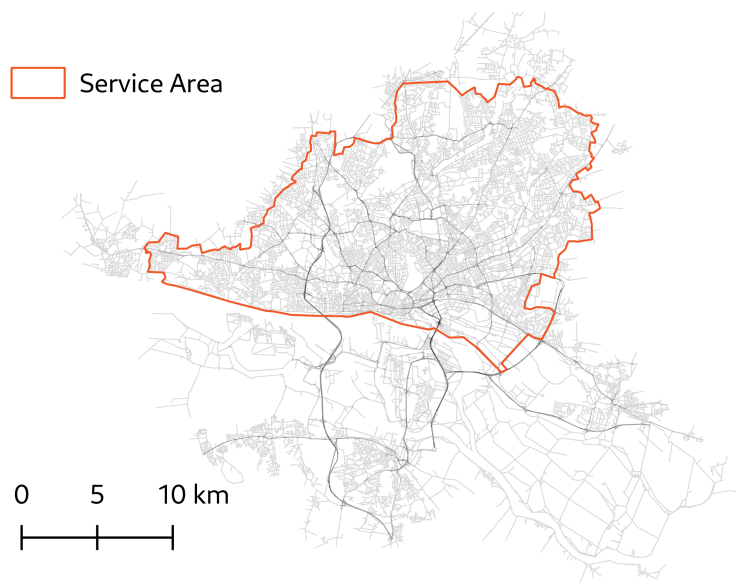


FIGURE 1: Study and service area for Hamburg.

7

8 The street network shown in Figure 1 is based on OpenStreetMap² data. MOIA’s more than
 9 10,000 virtual pick-up and drop-off stops are matched on it. We do not simulate transport modes
 10 other than ride-pooling and therefore do not observe congestion from external traffic in the system.
 11 To obtain realistic travel times that vary throughout the day, we use GPS-based speed data of all
 12 weekdays in November 2019 from TomTom and match it to the MATSim network with the help of
 13 a map-matching algorithm described by Yang and Gidófalvi (23). The network speeds are updated
 14 every 60 minutes.

15 We use 24,032 historical requests from MOIA, which have been collected on 4 typical
 16 weekend days between 19/09/2020 and 10/10/2020. A sample of one fourth of each day’s requests
 17 is used to avoid outlying extreme demand scenarios of a single day. The requests are combined and
 18 assumed to occur on the same simulated day. Requests from the same person within 30 minutes
 19 after the first request are ignored to avoid unrealistically clustered requests.

20 Experimental design

21 To test the strategies presented above in the MATSim environment, two major experiments are
 22 performed.

23 First, we construct a synthetic test case. We impose (uncongested) free flow travel times
 24 that stay constant over the day and use 200 vehicles to serve the demand. In a series of tests, we

²www.openstreetmap.org

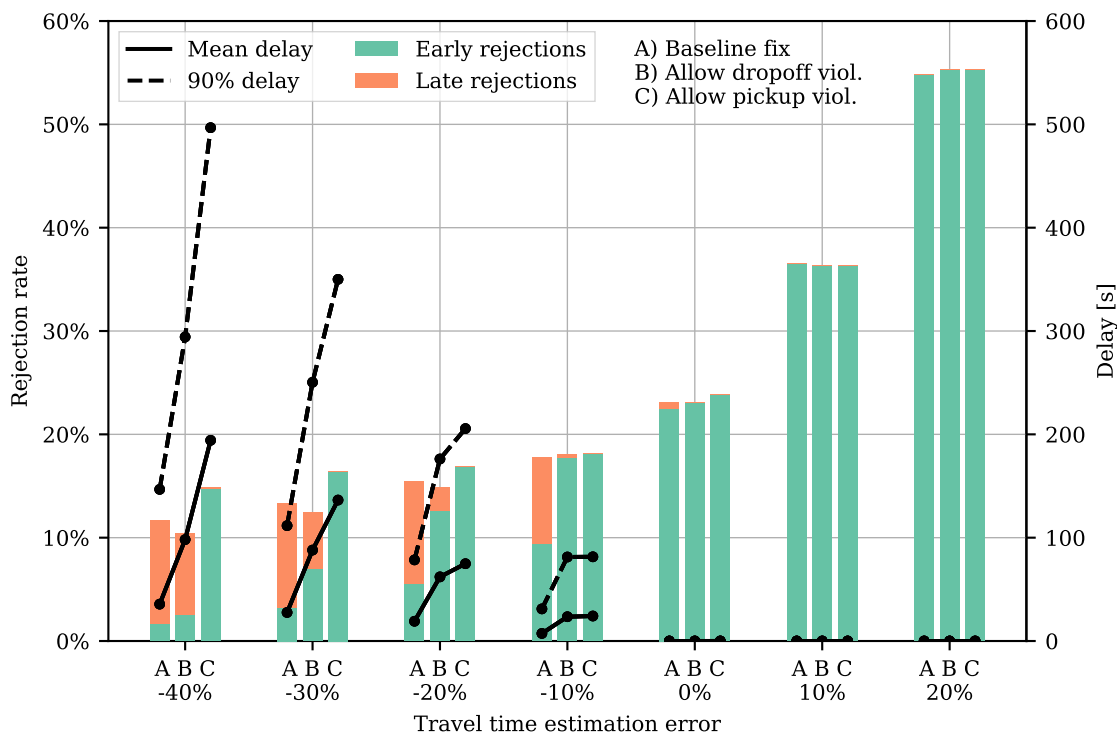


FIGURE 2: Simulation experiments with systematic estimation errors.

1 introduce a systematic bias to the travel times used by the algorithm in contrary to how the vehicle
 2 movements will be simulated. Travel times are scaled from -40% (underestimating arrival times,
 3 overestimating speeds) to 20% (overestimating arrival times, underestimating speeds). The goal
 4 is to see both the impacts of the estimation error for each strategy on the rejection rate (early and
 5 late), and on the additional delays beyond the latest arrival time of the requests.

6 In a second case, realistic travel times are imposed on the network and varied by time of
 7 day, as described above. The demand for Hamburg is served by 300 vehicles. The goals of these
 8 simulations is to see the impact of our proposed algorithm modifications in a realistic setting, based
 9 on real-world data. Uncertainty on travel times in this case stems from the varying travel times and
 10 the fact that a route calculated at one time during the simulation may overestimate or underestimate
 11 the travel time at a later time. Furthermore, we vary the routing interval (or “lifetime” of calculated
 12 routes) to understand its impact on system performance.

13 RESULTS

14 The following sections describe the simulation results for the synthetic and realistic cases.

15 Systematic changes

16 First, the cases with systematic estimation errors of travel time are examined. In Figure 2 we
 17 show a range of estimation errors (from -40% to +20%). For each estimation error, we show the
 18 share of rejections compared to all submitted requests and divide them in early rejections and late
 19 rejections. This is done for the three different strategies described above. On top, the mean and the
 20 90% quantile of the distribution of delays of all served requests are shown. The delays quantify by

1 how much time requests are dropped off after their latest drop-off time constraint.

2 Starting at an exact estimation of travel times, Figure 2 shows that the system has a rejection
3 rate of about 22%. If travel times are overestimated (right side of the plot), the rejection rate
4 increases. This is due to the algorithm being too careful and calculating with too much of slack.
5 Interestingly, an underestimation of travel times (left side) leads to a reduced number of rejections.
6 This indicates that there is a general slack in the system. At most times when estimating travel
7 times exactly, stop sequences for the vehicles have still room for inserting more requests, which
8 are, however, not available due to the density and travel characteristics in the scenario. These
9 capacities, however, are used when travel times are underestimated, because requests need to be
10 reassigned frequently to ensure the promised level of service.

11 However, note that especially for case A, which is the minimum mitigation strategy to run
12 the algorithm with varying travel times, almost all rejections are late rejections. This is neither de-
13 sirable for the operator, nor for the customer. Likewise, as underestimation becomes more severe,
14 the observed delays increase. For -40% estimation error, about 20 seconds delay can be observed
15 on average, with 10% of trips having a delay of more than two minutes.

16 In terms of the three strategies, one can see that they have a strong impact on the system
17 behavior. While, interestingly, the overall rejection rate is barely effected given a certain level of
18 estimation error, the strategies have a clear effect on the ratio between early and late rejections.
19 Allowing violations of drop-off constraints for the assigned vehicle (B) varies in impact depending
20 on the estimation error. For -40% its introduction has only little effect compared to (A), while
21 for lower estimation errors (e.g., -10%) it manages to avoid late rejections almost entirely. For
22 more severe estimation errors, the avoidance of late rejections is only guaranteed by strategy (C)
23 in which once assigned requests will always be served by their initially assigned vehicle. Note that
24 for all mitigation strategies, reducing the number of late rejections is bought by the generation of
25 additional delays, which can reach over 5 minutes if strategy (C) is used.

26 **Realistic system**

27 Clearly, the results from above are based on constructed scenarios to emphasize the effect of the
28 mitigation strategies. It is hence interesting to have a look at a more realistic scenario. Figure 3
29 shows the results for the service with 300 vehicles and realistic travel times. As our aim is to avoid
30 late rejections, only the simple case (A) is shown, along with allowing all violations for assigned
31 vehicles (C). On the bottom axis, we vary between three different routing cache intervals, from
32 rather detailed (15 minutes) to quite sparse (2 hours).

33 The results show that the rejection rate stays rather stable around 32% in all examined
34 cases. However, the routing interval has an impact on the number of late rejections. As routing
35 information becomes more rough, their number is increased. In all cases, introduction of strategy
36 (C) reduces these late rejections. Interestingly, there is only little delay in the realistic scenario,
37 where, even in the worst case, its mean value stays below 1 second on average.

38 **DISCUSSION**

39 From the results above, it becomes evident that one must be careful when applying the algorithm
40 of Alonso-Mora et al. (3) to scenarios with varying travel times. To avoid undefined system states,
41 measures must be taken, but the simplest approach (making sure that vehicles drop off late requests
42 before taking on any additional task) can lead to a high rate of late rejections. The presented
43 advanced strategies are able to mitigate this problem.

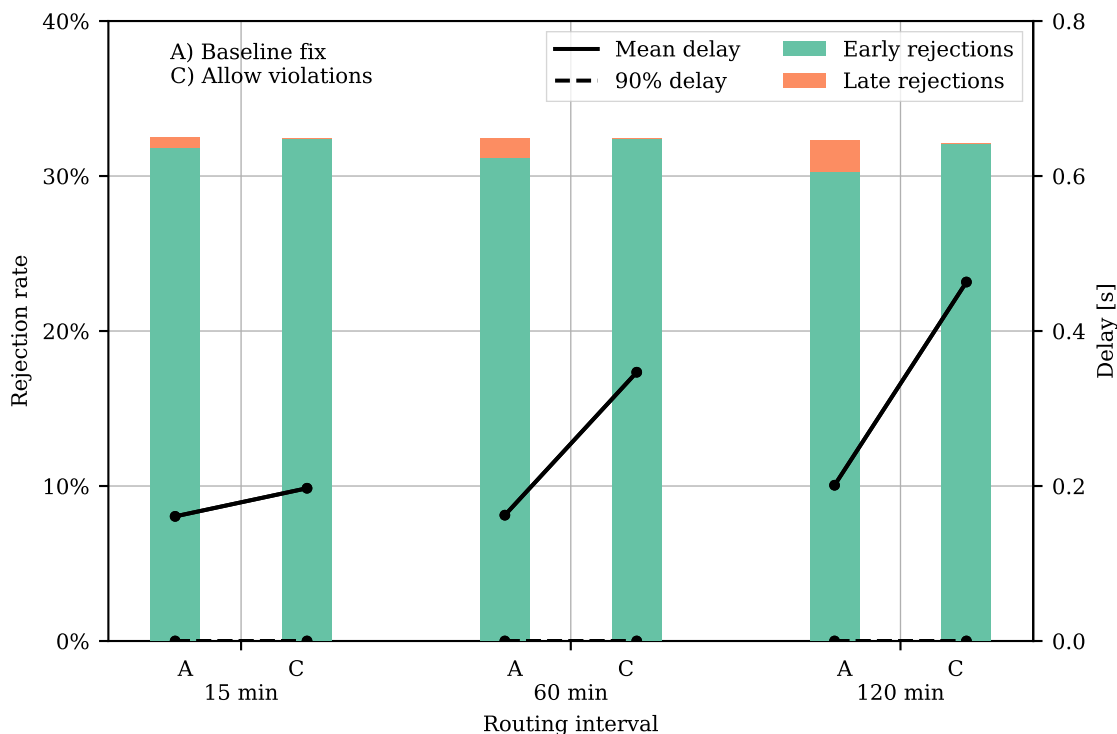


FIGURE 3: Simulation experiments with realistic demand and travel times.

1 Looking at the results from Figure 2, we can state that the strategies “buy” a better per-
 2 formance in terms of late rejections by “paying” with customer delays. And otherwise put, late
 3 rejections or delays are the “price” of having inaccurate estimates of travel time. Note that here we
 4 look at static demand simulations, but in reality excessive delays or late rejections would have a
 5 potentially strong impact on the attraction of the system and customer satisfaction. These findings
 6 are in line with (19) where (in the absence of travel time uncertainty) violations against the first
 7 pick-up time provided to the user, as well as binding customer-vehicle matching, are examined.
 8 Similar to our case, the paper’s authors find that enforcing constraints on the stability of arrival
 9 times towards the user comes with the direct effect of increased rejection rates.

10 However, looking at Figure 2 the impact of delays seems low. First, it should be noted that
 11 the system does not seem to run at its capacity limit and results may differ if a more constrained
 12 system was simulated. This is a task that should be tackled in future research. Second, the simu-
 13 lations in Figure 2 are valuable as they tell us about the robustness of the algorithm towards travel
 14 times in a realistic context. While, here, the routing interval was varied, and we see that with
 15 fewer routing, system performance stays nearly constant, the results open up space for even further
 16 simplifications. For instance, instead of performing detailed routing, zone-based travel time ma-
 17 trices could be used in the future. These insights pose a large potential in speeding up simulation
 18 experiments which make use of the algorithm.

19 The experiments in this paper make use of extreme cases. Either, travel times are affected
 20 strongly by a fixed bias, or they are minor as they are only slightly affected by varying travel
 21 times along the day. Note that we do not cover an entirely realistic case, in which information
 22 on the variability of travel times at any point in time is known. We assume that operators would

1 work with upper bound estimates to provide a robust system, which, however, could be tested
2 in simulation. The case is different for sudden shocks to the system, such as accidents. Both
3 topics relate strongly to the robustness of the ride-pooling system and should be covered in future
4 research. The integration of the algorithm into MATSim makes it possible to test it in a full-scale
5 microscopic traffic simulation in which such phenomena can be replicated in detail.

6 CONCLUSION

7 To conclude, we examine the problem of travel time uncertainty in combination with the algorithm
8 by Alonso-Mora et al. (3) which has been frequently used in literature. However, it has not been
9 used in the context of uncertain travel times, which is a major challenge for real-world ride-pooling
10 systems. We show that the original algorithm is undefined in certain situations when travel times
11 have been estimated too optimistically.

12 Our results show that a trade-off arises necessarily when strategies are proposed to remedy
13 the problem. If travel times are underestimated, the operator either must reject requests after they
14 have been accepted, or allow violations of the latest pick-up and drop-off times that have been
15 communicated to the customer.

However, we show that effects are minor in a scenario with realistic travel times, and we point out that there is a potential in time-slice based aggregation of travel times for speeding up simulation performance.

REFERENCES

1. Felix Zwick, Nico Kuehnel, Rolf Moeckel, and Kay W. Axhausen. Agent-based simulation of city-wide autonomous ride-pooling and the impact on traffic noise. *Transportation Research Part D: Transport and Environment*, 90:102673, 2021.
2. Luis M. Martinez and José Manuel Viegas. Assessing the impacts of deploying a shared self-driving urban mobility system: An agent-based model applied to the city of Lisbon, Portugal. *International Journal of Transportation Science and Technology*, 6(1):13–27, 2017.
3. Javier Alonso-Mora, Samitha Samaranayake, Alex Wallar, Emilio Frazzoli, and Daniela Rus. On-demand high-capacity ride-sharing via dynamic trip-vehicle assignment. *Proceedings of the National Academy of Sciences of the United States of America*, 114(3):462–467, 2017.
4. Jeffery B. Greenblatt and Samveg Saxena. Autonomous taxis could greatly reduce greenhouse-gas emissions of US light-duty vehicles. *Nature Climate Change*, 5(9):860–863, 2015.
5. Patrick M. Bösch, Felix Becker, Henrik Becker, and Kay W. Axhausen. Cost-based analysis of autonomous mobility services. *Transport Policy*, 64:76–91, 2018.
6. Henrik Becker, Felix Becker, Ryosuke Abe, Shlomo Bekhor, Prawira F. Belgiawan, Junia Compostella, Emilio Frazzoli, Lewis M. Fulton, Davi Guggisberg Bicudo, Krishna Murthy Gurumurthy, David A. Hensher, Johan W. Joubert, Kara M. Kockelman, Lars Kröger, Scott Le Vine, Jai Malik, Katarzyna Marczuk, Reza Ashari Nasution, Jeppe Rich, Andrea Papu Carrone, Danqi Shen, Yoram Shifan, Alejandro Tirachini, Yale Z. Wong, Mengmeng Zhang, Patrick M. Bösch, and Kay W. Axhausen. Impact of vehicle automation and electric propulsion on production costs for mobility services worldwide. *Transportation Research Part A: Policy and Practice*, 138:105–126, 2020.

7. Shuo Ma, Yu Zheng, and Ouri Wolfson. T-share: A large-scale dynamic taxi ridesharing service. In *Proceedings - International Conference on Data Engineering*, pages 410–421, 2013.
8. Joschka Bischoff, Michal Maciejewski, and Kai Nagel. City-wide shared taxis: A simulation study in Berlin. *IEEE Conference on Intelligent Transportation Systems, Proceedings, ITSC*, 2017.
9. Andreas Horni, Kai Nagel, and Kay W Axhausen. *The multi-agent transport simulation MATSim*. Ubiquity Press, London, 2016.
10. S. Hörl, C. Ruch, F. Becker, E. Frazzoli, and K.W. Axhausen. Fleet operational policies for automated mobility: A simulation assessment for Zurich. *Transportation Research Part C: Emerging Technologies*, 102:20–31, 2019.
11. Abood Mourad, Jakob Puchinger, and Chengbin Chu. A survey of models and algorithms for optimizing shared mobility. *Transportation Research Part B: Methodological*, 123: 323–346, 2019.
12. Anna Pernestål and Ida Kristoffersson. Effects of driverless vehicles : Comparing simulations to get a broader picture. *European Journal of Transport and Infrastructure Research*, 1(19):1–23, 2019.
13. Peng Jing, Hanbin Hu, Fengping Zhan, Yuexia Chen, and Yuji Shi. Agent-Based Simulation of Autonomous Vehicles: A Systematic Literature Review. *IEEE Access*, 8:79089–79103, 2020.
14. Lukas Foljanty. On-Demand Ridepooling Market: 2020 Recap. 2020.
15. Claudio Ruch, ChengQi Lu, Lukas Sieber, and Emilio Frazzoli. Quantifying the Efficiency of Ride Sharing. *IEEE Transactions on Intelligent Transportation Systems*, pages 1–6, 2020.
16. Andres Fielbaum, Xiaoshan Bai, and Javier Alonso-Mora. On-demand ridesharing with optimized pick-up and drop-off walking locations. *Transportation Research Part C: Emerging Technologies*, 126:103061, 2021.
17. Roman Engelhardt, Florian Dandl, and Klaus Bogenberger. Speed-up Heuristic for an On-Demand Ride-Pooling Algorithm. *arXiv:2007.14877 [cs, eess, math]*, 2020.
18. Florian Dandl, Roman Engelhardt, and Klaus Bogenberger. On the Dynamism of User Rejections in Mobility-on-Demand Systems. *arXiv:2104.01326 [cs, eess]*, 2021.
19. Andrés Fielbaum and Javier Alonso-Mora. Unreliability in ridesharing systems: Measuring changes in users’ times due to new requests. *Transportation Research Part C: Emerging Technologies*, 121:102831, 2020.
20. Ihab Kaddoura, Joschka Bischoff, and Kai Nagel. Towards welfare optimal operation of innovative mobility concepts: External cost pricing in a world of shared autonomous vehicles. *Transportation Research Part A: Policy and Practice*, 136:48–63, 2020.
21. Reza Vosooghi, Jakob Puchinger, Marija Jankovic, and Anthony Vouillon. Shared autonomous vehicle simulation and service design. *Transportation Research Part C: Emerging Technologies*, 107:15–33, 2019.
22. Sebastian Hörl, Felix Becker, and Kay W. Axhausen. Simulation of price, customer behaviour and system impact for a cost-covering automated taxi system in Zurich. *Transportation Research Part C: Emerging Technologies*, 123:102974, 2021.

23. Can Yang and Győző Gidófalvi. Fast map matching, an algorithm integrating hidden Markov model with precomputation. *International Journal of Geographical Information Science*, 32(3):547–570, 2018.