

# Robust Routing in Urban Public Transportation: Evaluating Strategies that Learn From the Past\*

Kateřina Böhmová<sup>1</sup>, Matúš Mihalák<sup>2</sup>, Peggy Neubert<sup>3</sup>,  
Tobias Pröger<sup>1</sup>, and Peter Widmayer<sup>1</sup>

- 1 Department of Computer Science, ETH Zürich  
Universitätstrasse 6, 8092 Zürich, Switzerland  
{katerina.boehmova,tobias.proeger,widmayer}@inf.ethz.ch
- 2 Department of Knowledge Engineering, Maastricht University  
Postbus 616, 6200MD Maastricht, The Netherlands  
matus.mihalak@maastrichtuniversity.nl
- 3 Verkehrsbetriebe Zürich  
Luggwegstrasse 65, 8048 Zürich, Switzerland  
Peggy.Neubert@vbz.ch

---

## Abstract

Given an urban public transportation network and historic delay information, we consider the problem of computing reliable journeys. We propose new algorithms based on our recently presented solution concept (Böhmová et al., ATMOS 2013), and perform an experimental evaluation using real-world delay data from Zürich, Switzerland. We compare these methods to natural approaches as well as to our recently proposed method which can also be used to measure typicality of past observations. Moreover, we demonstrate how this measure relates to the predictive quality of the individual methods. In particular, if the past observations are typical, then the learning-based methods are able to produce solutions that perform well on typical days, even in the presence of large delays.

**1998 ACM Subject Classification** F.2 Analysis of Algorithms and Problem Complexity, F.2.2 Nonnumerical Algorithms and Problems, I.2.6 Learning

**Keywords and phrases** public transportation, route planning, robustness, optimization, experiments

**Digital Object Identifier** 10.4230/OASICS.ATMOS.2015.68

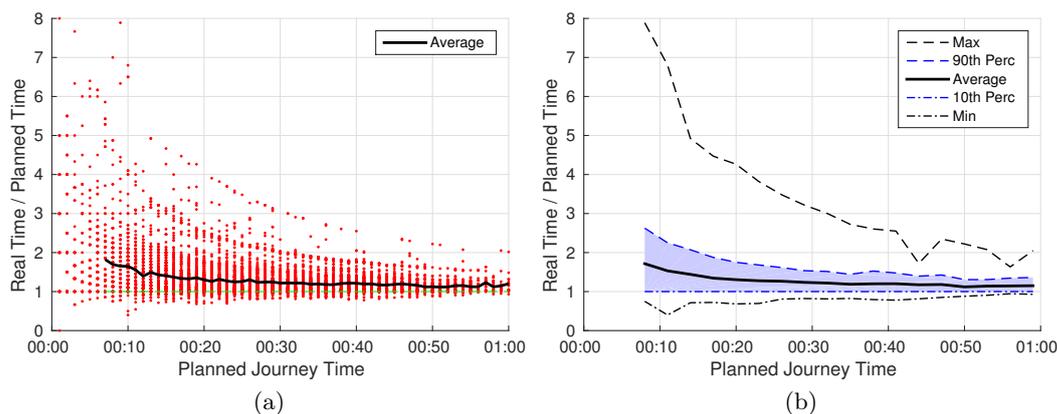
## 1 Introduction

**Motivation.** When using public transportation to travel from a stop  $s$  to a stop  $t$ , we may want to arrive at  $t$  no later than at time  $t_A$ . Determining the right moment to leave  $s$  is nontrivial: We want to reach  $t$  at time  $t_A$  at the latest, but we don't want to leave  $s$  much too early. In an ideal situation, every bus and every tram is on time, and it is sufficient to compute a journey that is planned to leave  $s$  as late as possible but still reaches  $t$  at the latest at  $t_A$ . However, in reality, traffic can be congested and we should expect delays. Thus, we are looking for a *robust* journey from  $s$  to  $t$  that arrives before time  $t_A$ , but still leaves  $s$

---

\* This work has been partially supported by the Swiss National Science Foundation (SNF) under the grant number 200021\_138117/1, and by the EU FP7/2007-2013 (DG CONNECT.H5-Smart Cities and Sustainability), under grant agreement no. 288094 (project eCOMPASS). Kateřina Böhmová is a recipient of a Google Europe Fellowship in Optimization Algorithms, and this research is supported in part by this Google Fellowship.





■ **Figure 1** Distribution of the error coefficients grouped by planned travel time (a), and the average, minimum, maximum, 10th and 90th percentiles of the distribution of the error coefficients in each time slot (b). Times were measured in minutes. For the sake of clarity, (a) does not show journeys with an error coefficient greater than 8, because there are only few ( $< 0.1\%$ ). For the same reason journeys with more than one hour planned travel time are not shown in the figures.

at a “reasonable” time. In real applications one may have additional preferences, such as low travel costs, which we don’t consider for the sake of simplicity.

Firmani et al. [9] observed in an experimental study on the transportation network of Rome that the timetable information and the real movement of the vehicles (based on GPS data) are only mildly correlated. They conclude that an “important issue to investigate is how to compute robust routes” that are “less vulnerable to unexpected events”. Our goal is provide methods for finding such robust routes. Since we only have historic delay data from the public transportation network of Zürich, as a preliminary step we investigated whether our network exhibits a behaviour similar to the Rome network. To make our results comparable to the results of Firmani et al., the methodology and notation of our preliminary study are similar to the methodology and notation in their original article [9].

We selected 10,000 departure and target stops  $s, t$  uniformly at random, set the latest allowed arrival time  $t_A$  to 8:30, and computed the  $st$ -journeys  $j$  that are optimal according to the planned timetable. For each journey  $j$ , we measured the planned travel times  $t_p(j)$  as well as the actual travel times  $t_a(j)$  (on 23 May 2013), and computed the error coefficient  $t_a(j)/t_p(j)$ . Figure 1(a) shows the distribution of the error coefficients grouped by the planned travel times  $t_p(j)$ . High error coefficients occur easily if the planned travel time is small and the vehicle of the planned journey leaves  $s$  a bit too early so that one has to wait for the next vehicle (which may, depending on the line, take up to half an hour in Zürich).

As in [9], we grouped the journeys into 3-minute time slots such that the  $k$ -th slot contains all journeys  $j$  with  $t_p(j) \in (3(k-1), 3k]$ . Figure 1(b) shows the average, minimum, maximum as well as the 10th and the 90th percentile of the distribution of the error coefficients of the journeys in each time slot. Since short journeys sometimes have high error coefficients, for simplicity Figure 1(b) does not incorporate the first two slots. The average error coefficient of the journeys in the remaining slots lies between 1.12 and 1.71 which means that in average a journey may take up to 71% longer than planned. Also, observe that the 90th percentile of the error coefficients of the journeys with 15 minutes travel time is roughly 2. Thus, 10% of the 15-minute journeys take in reality at least twice as long as planned. In overall, we observed that the behaviour in Zürich is comparable to the one in Rome.

One way out might be to integrate real-time information into the computation of routes. However, we believe that this is not enough, especially if the journey is planned some time

in advance. For example, a trip to the airport is usually planned a few hours earlier, and the right departure time needs to be computed before the start of the journey. Moreover in reality it often happens that delays occur suddenly and cannot be foreseen in advance, especially not at the time when the journey is planned. For example, consider an *st*-journey that consists of two lines  $l_1$  and  $l_2$ , and imagine that the transfer time between the lines is 2 minutes. Even if  $l_1$  leaves  $s$  on time, every upcoming delay of more than 1 minute (which might always occur) leads to a late arrival at  $t$ .

**Our Contribution.** In [3], we introduced a novel approach for finding robust journeys that uses recorded observations from the past as input—we look for journeys that performed well in the given past observations. Since this approach requires journeys to be comparable in different past days, classical solutions concepts, such as a path in the time-expanded or the time-dependent graph, are not suitable.

In the present paper, we first shortly describe our solution concept and the above-mentioned approach for finding robust routes. Since this approach was originally restricted to learn from historic data of only two different days, we show how it can be generalized to consider historic data from multiple days. We also describe how a stochastic method by Lim et al. [13] for private transportation can be adapted to compute robust journeys in public transportation. After that, we perform an extensive experimental study to evaluate these methods and to investigate different aspects related to robust routing.

**Related Work.** Many approaches to find a fastest journey in a given public transportation network were considered in the literature, see, e.g., a recent survey by Bast et al. [1]. One approach to account for delays is using stochastic methods—the delays are typically modeled as random variables on the edges of the network [4, 10, 15], or on each vehicle [6, 7]. For a given fixed timetable, Disser et al. [8] extended Dijkstra’s algorithm for computing pareto-optimal multi-criteria journeys. Müller-Hannemann and Schnee [14] used a *dependency graph* to predict secondary delays caused by some current primary delays and gave a routing strategy with respect to these delays. Bast et al. [2] studied the robustness of transfer patterns in the presence of delays. They argue that even when delays occur, a reasonably good path is still included in the pattern. Dibbelt et al. [7] modeled the delays using stochasticity and computed a *decision graph* with all the possibly relevant nodes and vehicles instead of a single path. Goerigk et al. [12] assumed that a set of delay *scenarios* is provided, and showed how to compute a journey that arrives on time in every scenario (strict robustness) or a journey with fewest number of unreliable transfers having an almost optimal travel time (light robustness). Goerigk et al. [11] considered journeys, within the setting of delay scenarios, that can be updated if delays occur (recoverable robustness).

## 2 Model

**Network Design.** Let  $\mathcal{S}$  be a set of stops. A *line* is an ordered sequence  $\langle v_1, \dots, v_k \rangle$  of stops from  $\mathcal{S}$ , where  $v_i$  is visited directly before  $v_{i+1}$ . We explicitly distinguish two lines with the same stops but opposite directions. Given a departure stop  $s \in \mathcal{S}$  and a target stop  $t \in \mathcal{S}$ , a sequence of lines  $\langle l_1, \dots, l_{\beta+1} \rangle$  with  $l_i \neq l_{i+1}$  is called an *st-route* if there exist  $\beta + 2$  stops  $v_0 := s, v_1, \dots, v_\beta, v_{\beta+1} := t$  where both  $v_{i-1}$  and  $v_i$  are stops on the line  $l_i$ , and the line  $l_i$  visits  $v_{i-1}$  (not necessarily directly) before  $v_i$ . We say that a *transfer* between the lines  $l_i$  and  $l_{i+1}$  occurs at  $v_i$ . Notice that there might be more than one possible transfer

between two lines. For  $s, t \in \mathcal{S}$  and an integer  $\beta \in \mathbb{N}_0$ , let  $\mathcal{R}_{st}^\beta$  denote the set of all  $st$ -routes with at most  $\beta$  transfers.

A *journey* consists of a departure time  $t_D$ , a route  $\langle l_1, \dots, l_{\alpha+1} \rangle \in \mathcal{R}_{st}^\alpha$  with  $\alpha \leq \beta$ , and a sequence of transfer stops  $\langle v_1, \dots, v_\alpha \rangle$ . Its intuitive interpretation is to leave the stop  $s$  at time  $t_D$ , take the first arriving (trip of) line  $l_1$ , and for every  $i \in \{1, \dots, \alpha\}$ , leave  $l_i$  at stop  $v_i$  and immediately take the next arriving trip of line  $l_{i+1}$ .

**Trips and Timetables.** While the only information associated with a line itself are its consecutive stops, it usually is operated multiple times per day. Each of these concrete realizations is called a *trip*. A *timetable* stores for every stop  $v \in \mathcal{S}$  the arrival and departure times of every trip over a day. We have

1. a *planned* timetable  $T_{plan}$  which we assume to be periodic, i.e., every line realized by some trip  $\tau$  will be realized by a later trip  $\tau'$  again (not necessarily on the same day).
2. a set  $\mathcal{T}$  of *recorded* timetables  $T_i$  that describe how various lines were operated during a given time period (e.g., on a concrete day). These recorded timetables are concrete executions of the planned timetable.

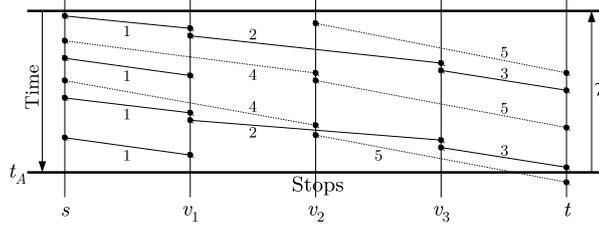
In the following, *timetable* refers both to the planned as well as to a recorded timetable. We assume that timetables respect the FIFO property, i.e. two buses or trams *of the same line* do not overtake each other.

**Goal.** Let  $s, t \in \mathcal{S}$  be the departure and the target stop, and let  $t_A$  be the latest allowed arrival time. Our goal is to use the planned timetable and the recorded timetables in  $\mathcal{T}$  to compute a recommendation in form of one or more (robust) journeys from  $s$  to  $t$  that will likely arrive on time (i.e., at time  $t_A$  or earlier) on a day for which the concrete travel times are not known yet.

We assume that users select one of the recommended journeys, and then travel according to it. One may argue that this assumption is rather strict, because when delays come up, users sometimes spontaneously decide to use a different journey instead. This, however, is a different situation that we do not consider in this paper for two reasons. First, one needs to know the network and the possible backup options well, which might not be the case when one is travelling in a foreign city. Second, as mentioned earlier, delays may occur suddenly, and it might be too late to choose a different journey. Consider, for example, the situation when the alternative journeys don't have any stop in common except for the departure and the target stop. In such a case one has to fix the journey already in advance.

### 3 Robustness

**Overview.** In this section we present some approaches for computing robust journeys. For this we assume that the departure stop  $s$ , the target stop  $t$  and the latest allowed arrival time  $t_A$  were specified by the user and that we already computed a reasonable upper bound  $\beta$  on the maximum number of transfers. Hence,  $s$ ,  $t$ ,  $t_A$  and  $\beta$  are fixed when the journey(s) are computed. We note that, given a route  $r \in \mathcal{R}_{st}^\beta$  and a parameter  $\gamma \in \mathbb{N}$ , we can use the planned timetable  $T_{plan}$  to find a journey  $j$  along  $r$  that leaves  $s$  as late as possible, but not later than time  $t_A - \gamma$ . Thus, as soon as an algorithm identifies both a route  $r$  and a parameter  $\gamma$ , it can also reconstruct the corresponding journey in the planned timetable. These planned journeys will then be recommended to the user.



■ **Figure 2** A timetable with five lines  $\{1, \dots, 5\}$  and two routes  $r_1 = \langle 1, 2, 3 \rangle$  (solid) and  $r_2 = \langle 4, 5 \rangle$  (dotted). The  $x$ -axis denotes the stops  $\{s, v_1, v_2, v_3, t\}$ , the  $y$ -axis the time. If a trip leaves a stop  $v_d$  at time  $t_d$  and arrives at a stop  $v_a$  at time  $t_a$ , it is indicated by a line segment from  $(v_d, t_d)$  to  $(v_a, t_a)$ .  $A_\gamma(T)$  contains  $r_1$  three times and  $r_2$  once.

**Transfer Buffers.** An naïve strategy to increase the reliability of a journey is to enforce an additional buffer time at each transfer or at the end of the trip. The **Buffer- $\xi$**  approach uses  $T_{plan}$  to compute a journey that is planned to leave  $s$  as late as possible, arrives at  $t$  not later than at time  $t_A$ , and that has an additional time of at least  $\xi$  at each transfer of the journey. This especially implies that if a line  $l_i$  is planned to arrive at a transfer stop  $v_i$  at time  $t_i$ , then the next line  $l_{i+1}$  of the journey can only be taken at time  $t_i + \xi$  or later. **Buffer-0** corresponds to an optimal journey in the planned timetable, so we refer to it as **Opt-TT**.

**A Similarity-Based Approach.** In [3], we described how a general approach to robust optimization designed by Buhmann et al. [5] can be used to compute robust journeys. We briefly recall our ideas. Let  $T \in \mathcal{T}$  be a timetable and  $\gamma \in \mathbb{N}_0$ . An *approximation set*  $A_\gamma(T)$  contains all routes  $r \in \mathcal{R}_{st}^\beta$  for which  $T$  contains a journey along  $r$  that leaves  $s$  at time  $t_A - \gamma$  or later, and that arrives at  $t$  at time  $t_A$  or earlier. We assume that  $A_\gamma(T)$  is a *multiset*: a route  $r$  is contained as often as it is realized by a journey starting at time  $t_A - \gamma$  or later, and arriving at time  $t_A$  or earlier (see Figure 2 for an example). The parameter  $\gamma$  can be interpreted as the maximal time that we depart before  $t_A$ . In general we have  $A_0(T) = \emptyset$ , and the size of  $A_\gamma(T)$  grows with increasing  $\gamma$ . If we consider the approximation sets  $A_\gamma(T_1), \dots, A_\gamma(T_k)$  for the timetables  $T_1, \dots, T_k \in \mathcal{T}$ , every approximation set contains only routes that are realized (by a journey) in the same time period  $[t_A - \gamma, t_A]$ , and that are therefore comparable among different approximation sets.

The approach in [3, 5] expects that exactly two timetables  $T_1, T_2 \in \mathcal{T}$  are given. To compute a *robust* route when only two timetables are available, we consider  $A_\gamma(T_1) \cap A_\gamma(T_2)$ : the only chance to find a route that is likely to be good in the future is a route that performed well in *both* recorded timetables. The parameter  $\gamma$  determines the size of the intersection: if  $\gamma$  is too small, the intersection will be empty. If  $\gamma$  is too large, the intersection contains many (and maybe all)  $st$ -routes, and not all of them will be a good choice. Assuming that we knew the “optimal” parameter  $\gamma_{OPT}$ , we could pick a route from  $A_{\gamma_{OPT}}(T_1) \cap A_{\gamma_{OPT}}(T_2)$ . Buhmann et al. [5] suggest to set  $\gamma_{OPT}$  to the value  $\gamma$  that maximizes

$$S_\gamma = \frac{|\mathcal{R}_{st}^\beta| |A_\gamma(T_1) \cap A_\gamma(T_2)|}{|A_\gamma(T_1)| |A_\gamma(T_2)|}. \quad (1)$$

The value  $S_{\gamma_{OPT}}$  measures how similar the timetables  $T_1$  and  $T_2$  are, so Buhmann et al. refer to this ratio as the *similarity* of  $T_1$  and  $T_2$ . They showed that it is always at least 1, and the larger it gets, the more similar  $T_1$  and  $T_2$  are. Of course, if one is only interested in computing  $\gamma_{OPT}$  (and not measuring the similarity itself), one can simply omit the term  $|\mathcal{R}_{st}^\beta|$  in equation (1) as we did in our original work [5].

After  $\gamma_{OPT}$  has been computed, there are two possible approaches to pick a route from  $A_{\gamma_{OPT}}(T_1) \cap A_{\gamma_{OPT}}(T_2)$ . The **Similarity-Rand** approach selects a route  $r$  from the intersection uniformly at random, while **Similarity-MRR** selects the most frequent route  $r$  from the intersection. For both approaches we recommend to depart at least  $\gamma_{OPT}$  units of time in advance. More details can be found in [3, 5].

**Function-Based Approaches.** Let  $T_i \in \mathcal{T}$  be a recorded timetable,  $r = \langle l_1, \dots, l_{\alpha+1} \rangle \in \mathcal{R}_{st}^\alpha$  be a route,  $\tau_1, \dots, \tau_k$  be the trips of line  $l_1$  in  $T_i$  and  $D(\tau_j, s)$  be the departure time of the trip  $\tau_j$  at  $s$ . We define  $\delta_i^r$  as

$$\min_{j \in [1, k]} \left\{ t_A - D(\tau_j, s) \mid \begin{array}{l} \tau_j \text{ can be extended to a journey along } r \text{ that} \\ \text{arrives in } T_i \text{ at stop } t \text{ at time } t_A \text{ or earlier} \end{array} \right\}, \quad (2)$$

which intuitively can be interpreted as follows: to arrive on time using route  $r$  on the day at which  $T_i$  is realized, one has to leave  $s$  at least  $\delta_i^r$  units of time before the latest allowed arrival time  $t_A$ . For a given function  $f : (\mathbb{R}^+)^{|\mathcal{T}|} \rightarrow \mathbb{R}$ , we search for a route  $r \in \mathcal{R}_{st}^\alpha$  that minimizes  $f(\delta_1^r, \dots, \delta_{|\mathcal{T}|}^r)$ . In the following, we describe some possible choices for  $f$ , and we abbreviate  $f(\delta_1^r, \dots, \delta_{|\mathcal{T}|}^r)$  by  $f(r)$ .

For a number  $p \in [1, \infty]$ , the **Norm- $p$**  estimator has the objective function

$$f_{\|\cdot\|}^p(r) = \left\| (\delta_1^r, \dots, \delta_{|\mathcal{T}|}^r) \right\|_p. \quad (3)$$

It is easy to see that  $f_{\|\cdot\|}^1$  selects all routes which in average (w.r.t. the recorded timetables in  $\mathcal{T}$ ) depart as late as possible. Moreover,  $f_{\|\cdot\|}^\infty$  selects all routes minimizing the maximum time between the departure and the latest allowed arrival time  $t_A$ . Such routes can alternatively be seen as routes maximizing the earliest departure time necessary to arrive on time in *all* timetables in  $\mathcal{T}$ . Thus, the **Norm- $\infty$**  estimator is related to the similarity-based approach from the previous paragraph in the following way. Let  $\gamma_{FI} = \min \{ \gamma > 0 \mid \bigcap_{i=1}^{|\mathcal{T}|} A_\gamma(T_i) \neq \emptyset \}$  be the smallest value for  $\gamma$  such that the intersection of all  $\gamma$ -approximation sets is non-empty. One can observe that every route  $r$  contained in  $\bigcap_{i=1}^{|\mathcal{T}|} A_{\gamma_{FI}}(T_i)$  minimizes  $f_{\|\cdot\|}^\infty$  and vice versa. We note that these methods relate to strict robustness [12], but are based on a different solution concept, and learn from past observations given as daily recorded timetables (instead of specifying a set of possible delays).

Now, let  $p \in [1, \infty]$  be arbitrary and let  $r_j^p$  be a route minimizing  $f_{\|\cdot\|}^p$ . To determine how much in advance one has to depart when using  $r_j^p$ , we use our previous observations. For  $p = 1$ , it is reasonable to set  $\gamma_j^p = f^1(r_j^p)/|\mathcal{T}|$  since  $f_{\|\cdot\|}^1$  corresponds to averaging the departure times. For  $p = \infty$ , it is reasonable to set  $\gamma_j^p = f^\infty(r_j^p)$ . For every other  $p \in (1, \infty)$ , we simply scale the time linearly with respect to  $p = 1$  and  $p = \infty$ . More concretely, we set

$$\gamma_j^p = f^\infty(r_j^p) - \left( \frac{f^p(r_j^p) - f^\infty(r_j^p)}{f^1(r_j^p) - f^\infty(r_j^p)} \right) \cdot (f^\infty(r_j^p) - f^1(r_j^p)/|\mathcal{T}|). \quad (4)$$

A different function-based estimator comes from the *mean-risk model* which was just recently used for finding robust routes in private transportation [13]. Let  $c \in \mathbb{R}_0^+$  be the *risk-aversion coefficient*, where  $c = 0$  corresponds to the situation where the risk is being completely ignored. The objective function associated with the **Mean-Risk- $c$**  estimator is

$$f_{MR}^c(r) = \text{Mean}(\delta_1^r, \dots, \delta_{|\mathcal{T}|}^r) + c \cdot \sqrt{\text{Variance}(\delta_1^r, \dots, \delta_{|\mathcal{T}|}^r)}. \quad (5)$$

For a route  $r_j$  minimizing  $f_{MR}^c$ , we simply set  $\gamma_j = f_{MR}^c(r_j)$  as the time one has to depart in advance. Notice that **Mean-Risk-0** is equivalent to **Norm-1**.

## 4 Experimental Results

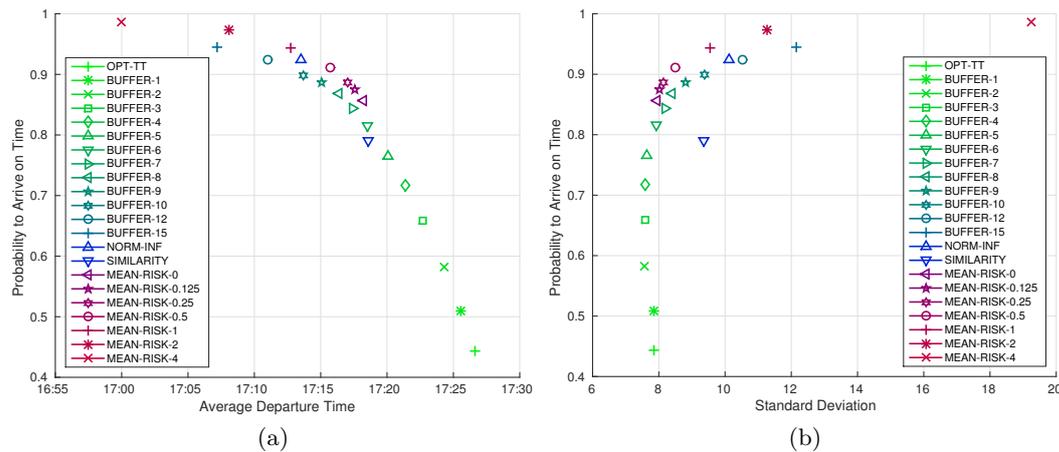
**Experimental Setup.** For an experimental evaluation of the methods proposed in Section 3 we used the tram and bus network of the city of Zürich, Switzerland, which has 401 stops and 292 lines. The recorded timetables  $\mathcal{T} = \{T_1, \dots, T_7\}$  were realized on seven consecutive Thursdays in the period from 4 April to 23 May 2013, ignoring 9 May (which was a public holiday and therefore had different traffic and a different planned timetable).

We observed that in reality many of the 292 lines have the same ID (such as, e.g., tram 6, bus 31, etc.). This is consequence of our modeling: not only do we distinguish lines travelling in opposite directions, but there are also special lines coming from or going to the depot, lines whose corresponding vehicle turns around in advance, and lines that do not visit certain stops in the evening. Since these special lines operate only on a low frequency and mostly only early in the morning or late in the evening, we ignored them and focused on the “standard” realizations. Hence we effectively used only 118 of the 292 lines. Although the network is rather small in comparison to the networks of other cities, it is well-suited for an experimental study on robustness for two reasons. First, the network is dense enough to provide many different routes between any two stops  $s$  and  $t$ . Second, our study in Section 1 showed that the network is affected by a considerable amount of delays, especially during the rush hours.

For each of the following experiments, we generated 10,000 (30,000 for the experiments on the number of transfers) departure/target pairs  $(s, t) \in \mathcal{S}^2$  with  $s \neq t$  uniformly at random. For each such pair  $(s, t)$ , we computed the smallest  $\beta \in \mathbb{N}_0$  such that  $\mathcal{R}_{st}^\beta \neq \emptyset$  and used this value for the maximum allowed number of transfers. We explicitly set  $\beta = 1$  if there exists a direct  $st$ -route with no transfers at all. In such a case, one might prefer to take an alternative route with only one transfer, probably leading to a shorter travel time. After computing  $\beta$  and  $\mathcal{R}_{st}^\beta$ , we performed the corresponding experiment. We set the target arrival time  $t_A$  to 18:00 except for the experiments that study how the behavior of the methods changes during the day. Unless otherwise stated, the buffer methods used the planned timetable  $T_{plan}$  as input, the similarity-based methods used  $T_5$  and  $T_6$  (recorded on 2 May and 16 May), and the function-based methods used  $T_1, \dots, T_6$  (recorded between 4 April and 16 May). Timetable  $T_7$  (recorded on 23 May) was used to assess the quality of the proposed journeys.

In our experiments we observed that the performance of **Similarity-Rand** and **Similarity-MRR** is nearly identical, so our figures show only the behavior of the latter variant, and for simplicity we refer to both variants as **Similarity**. Also, **Norm-2** performs similarly to **Norm-Inf**, so our figures mostly omit **Norm-2**. Furthermore we observed that it rarely happened that a journey proposed by **Buffer- $\xi$** , **Similarity**, **Norm-Inf** or **Mean-Risk-1** arrived much too early or much too late in the test instance. In all of these cases this was caused either because of a highly non-typical situation in the input or the test instance (e.g., an accident), or because a line was chosen that was not realized regularly (e.g., less than once per hour). Hence we ignored all pairs  $(s, t)$  for which at least one of the methods above computed a journey arriving more than one hour too early or too late.

Our algorithms were implemented in Java 7, and the experiments were performed on one core of an Intel Core i5-3470 CPU clocked at 3.2 GHz with 4 GB of RAM running Debian Linux 7.8. For enumerating all  $st$ -routes in  $\mathcal{R}_{st}^\beta$ , we used the algorithm proposed in [3] which runs on average 35ms. After computing  $\mathcal{R}_{st}^\beta$ , the buffer strategies have an average running time 1ms or less, the similarity-based methods 8ms, and the function-based approaches 24ms. Notice that these running times are faster than the ones described in [3], because we used a smaller network (without the agglomeration).



■ **Figure 3** Comparison of various methods: arrival rate vs. average departure time (a), and arrival rate vs. standard deviation on the arrival time (b).

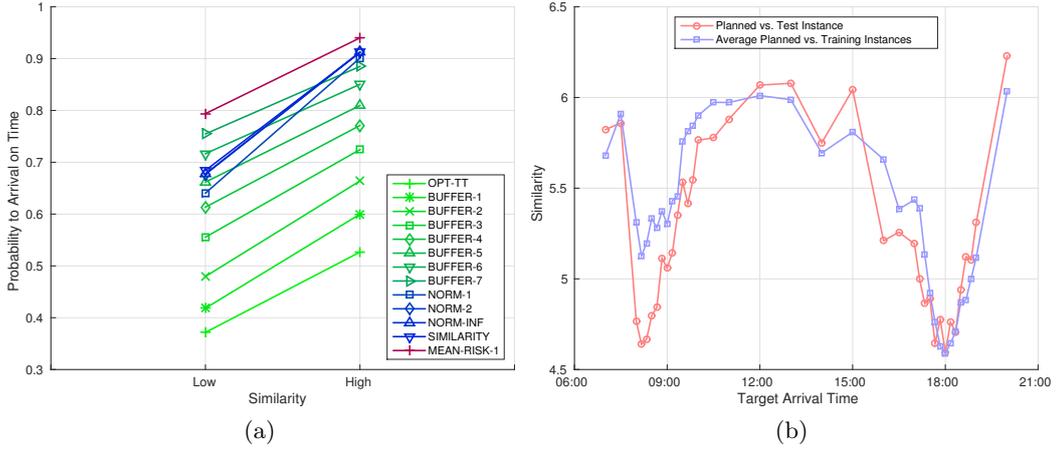
■ **Table 1** Overview of how often the route suggestions of two methods differ.

	Opt-TT	Buffer-3	Buffer-6	Buffer-9	Buffer-12	Norm-1	Norm-Inf	Similarity	Mean-Risk-1
Opt-TT		30, 82%	24, 14%	25, 79%	29, 56%	27, 32%	40, 45%	40, 21%	32, 91%
Buffer-3	30, 82%		31, 60%	25, 97%	24, 89%	30, 31%	40, 05%	40, 72%	32, 54%
Buffer-6	24, 14%	31, 60%		28, 77%	25, 83%	30, 86%	41, 03%	42, 17%	34, 16%
Buffer-9	25, 79%	25, 97%	28, 77%		30, 37%	29, 99%	39, 59%	40, 68%	32, 03%
Buffer-12	29, 56%	24, 89%	25, 83%	30, 37%		31, 77%	40, 83%	42, 48%	33, 99%
Norm-1	27, 32%	30, 31%	30, 86%	29, 99%	31, 77%		27, 48%	31, 30%	14, 33%
Norm-Inf	40, 45%	40, 05%	41, 03%	39, 59%	40, 83%	27, 48%		32, 43%	19, 54%
Similarity	40, 21%	40, 72%	42, 17%	40, 68%	42, 48%	31, 30%	32, 43%		32, 50%
Mean-Risk-1	32, 91%	32, 54%	34, 16%	32, 03%	33, 99%	14, 33%	19, 54%	32, 50%	

**Arrival Rate, Departure Time and Standard Deviation on the Arrival Time.** Intuitively, an earlier departure time leads to a higher probability to arrive on time (i.e., a higher arrival rate), and achieving a higher arrival rate in a network with delays entails a higher standard deviation on the arrival time. Figure 3 compares the proposed methods with respect to these aspects. It shows that, independently of the considered method, there is a clear trade-off between the departure time and the arrival rate (a) as well as between the standard deviation of the arrival time and the arrival rate (b).

Both parameter-based methods Buffer- $\xi$  and Mean-Risk- $c$ , form Pareto optimal fronts in both (a) and (b). Clearly, Mean-Risk- $c$  benefits from the additional information from the input instances  $T_1, \dots, T_6$  and it dominates Buffer- $\xi$  in both (a) and (b). The Similarity method needs no parameter adjustment, it is based only on two past timetables, and still proposes solutions with a reasonable arrival rate that do not depart too early. Notice that Norm-Inf (the generalization of Similarity) also benefits from the knowledge of the six past timetables, and without parameter adjustment it produces a solution which gives a very reasonable trade-off between departure time, arrival rate and the standard deviation on the arrival time. Moreover, the solutions proposed by Norm-Inf performed rather well compared to all the competitors (which do require parameter adjustment).

We also investigated whether the arrival rates of different methods differ due to different departure times only, or whether the suggested route(s) also differ. In particular, for any two methods  $M_1$  and  $M_2$ , we studied how often the suggested route(s) of  $M_1$  and  $M_2$  differ. For



■ **Figure 4** Influence of low/high similarity on the arrival rate: comparing various methods (a). Influence of the target arrival time on the similarity of the planned timetable and the test instance  $T_7$ , and on the average similarity between the planned timetable and each of the input instances  $T_1, \dots, T_6$  (b).

some exemplary methods, Table 1 shows that this happens in 14 to 42% of the cases. Notice that in roughly one third of the cases, the routes proposed by *Similarity* differ from the ones proposed by *Norm-Inf* (which can be seen as a generalization of *Similarity*). Also, there is a notable difference between the route suggestions of the different *Buffer* methods. Thus, for enforcing robustness there are better strategies than merely decreasing the departure time.

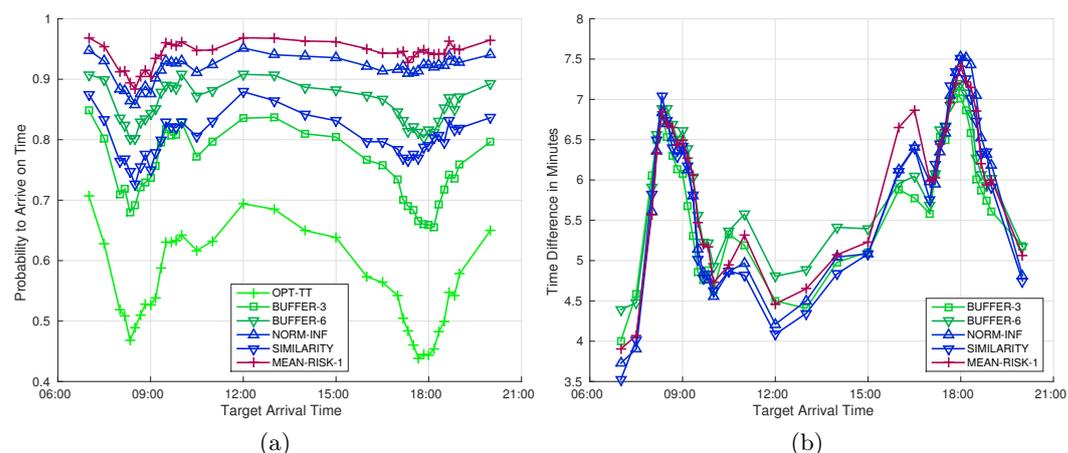
**Influence of the Similarity between Input and Test Instances.** We just saw that journeys proposed by the similarity-based approaches performed rather poorly, with respect to both arrival rate as well as standard deviation on the arrival time. However, we have to take into account that these methods use only two recorded timetables as input: if both differ substantially from the test instance, then in general there is very little one can do. The generic approach by Buhmann et al. [5] works well if both the input and the test instances are typical, i.e., if their mutual similarities is high. Thus we investigate the impact of high and low mutual similarities on the quality of the predictions.

First we note that the similarity  $S_{\gamma_{OPT}}$  does not only depend on the two input instances but also on the origin  $s$  and the destination  $t$ , and on the target arrival time  $t_A$ . Thus, in the following experiments, we do not always use the same timetables  $T_5, T_6$  as input and  $T_7$  for testing, but select for every  $(s, t)$  the timetables whose mutual similarities are as high or as low as possible. Let  $\Upsilon$  be the set of all triples of recorded timetables  $(T_i, T_j, T_k) \in \mathcal{T}^3$  where  $i, j, k$  are mutually different. For a given pair  $(s, t)$  and two timetables  $T_i, T_j \in \mathcal{T}$ , let  $S_{ij}^{st}$  be the similarity of  $T_i$  and  $T_j$  with respect to  $s$  and  $t$ . We selected triples whose minimum (or maximum, respectively) pairwise similarity is as high or as low as possible,

$$(T_1^h, T_2^h, T_3^h) = \arg \max_{(T_i, T_j, T_k) \in \Upsilon} \min \{ S_{ij}^{st}, S_{ik}^{st}, S_{jk}^{st} \} \quad (6)$$

$$(T_1^l, T_2^l, T_3^l) = \arg \min_{(T_i, T_j, T_k) \in \Upsilon} \max \{ S_{ij}^{st}, S_{ik}^{st}, S_{jk}^{st} \} \quad (7)$$

and used  $T_1^h$  and  $T_2^h$  as input and  $T_3^h$  for testing, and for comparison, used  $T_1^l$  and  $T_2^l$  as input and  $T_3^l$  for testing. Even though *Mean-Risk- $c$*  and *Norm- $p$*  could handle more instances, they were given just the two mentioned instances.



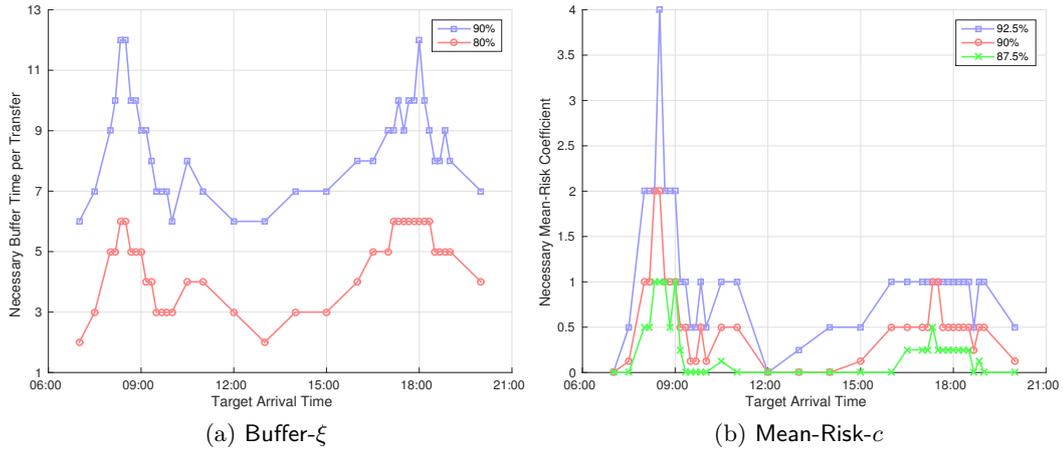
■ **Figure 5** Comparison of various methods: Arrival rate vs. target arrival time (a), and travel time difference to the optimum travel time vs. target arrival time (b).

Figure 4(a) shows that all methods benefit when the similarity of the three instances is high. The arrival rates of both Norm- $p$  and especially Similarity increase significantly. We observed that Similarity outperforms Norm- $p$  when the similarity is low, which is reasonable: for a low similarity, the routes in the first intersection of the approximation sets as well as the route that maximizes the average departure time are too much influenced by the noise in the input instances. However, Similarity can still let the approximation sets grow beyond the first intersection so that more stable solutions are contained (which Norm- $p$  cannot). On the other hand, if the similarity is high, then there is so little noise in the data that  $S_\gamma$  is maximized already at the first  $\gamma$  for which the intersection is non-empty, thus Similarity and Norm- $p$  are nearly identical.

Of course these results cannot directly be used for designing an algorithm, since the test instance is unknown. Nevertheless we believe that the results are interesting because they demonstrate the power of the similarity-based approach.

**Influence of the Target Arrival Time.** Figure 5 shows how the behavior of the methods, in terms of the arrival rate (a) and travel time (b), changes over the day. In particular, we can observe a clear influence of the morning and evening rush hours. Interestingly, the two rush hours affect the arrival rates of different methods differently. Specifically, the timetable-based method Buffer- $\xi$  is greatly affected by both rush hours while the learning strategies are less affected by the evening rush hour.

To understand this behavior, consider Figure 4(b). The red curve shows how the value of the similarity of  $T_{plan}$  and the test instance  $T_7$  changes during the day. In particular, we see a significant drop of the similarity during rush hours. Notably, the two dips corresponding to morning and evening rush hour are of the same height. This suggests that on the day corresponding to  $T_7$ , during the morning rush hour, there was a similar amount of irregularities with respect to  $T_{plan}$ , as during the evening one. The blue curve in Figure 4(b) shows the changes during the day of the averaged value of similarity of  $T_{plan}$  and each of the training instances  $T_1 - T_6$ . Also there the similarity drops during rush hours, but we clearly see that the morning dip is significantly lower than the evening one. This suggests that in the recorded timetables  $T_1 - T_6$  used for learning, the amount of irregularities (with respect to  $T_{plan}$ ) was lower in the morning than in the evening. Thus, when comparing the two



■ **Figure 6** Necessary parameter to achieve a specified arrival rate in  $T_7$  depending on the target arrival time.

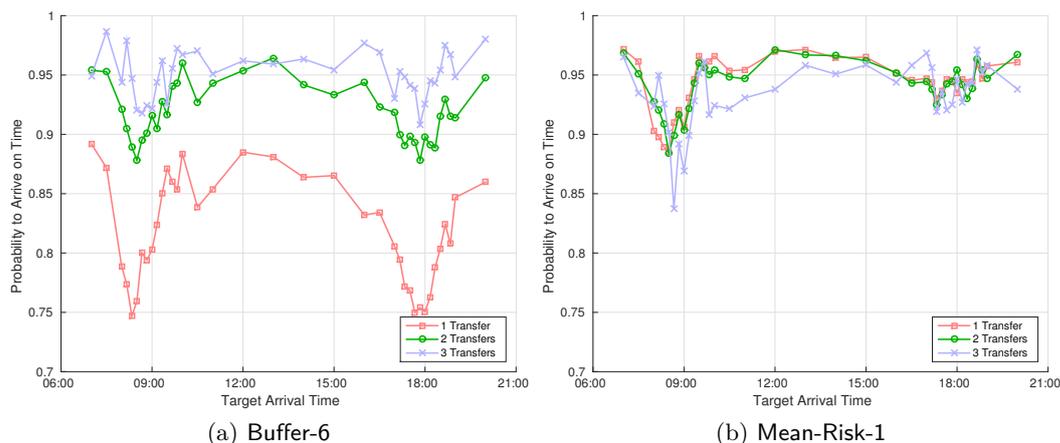
curves, we see a significant gap between them during the morning rush hour, but a relative match during the rest of the day. This suggests that the test instance  $T_7$  contained during the day a similar amount of irregularities as it is expected on a typical day (represented by  $T_1 - T_6$ ), with the only exception of the morning rush hour, where it was less regular.

Let us now relate what we observed in Figures 4(b) and 5(a). Since Buffer- $\xi$  is based solely on  $T_{plan}$ , any irregularities with respect to  $T_{plan}$  occurring in  $T_7$  (captured by the red curve in Figure 4(b)) affect its arrival rate. This explains why the arrival rate of Buffer- $\xi$  drops both in the morning and evening rush hour and exhibits two dips of nearly the same height. On the other hand, the methods that use the information from the past observations (e.g., Mean-Risk- $c$ ) are trained to account for a certain amount of irregularities. Since the situation in  $T_7$  in the evening is typical, the solutions proposed by these methods are prepared for it and their arrival rate is almost not affected by the evening rush hour. In contrast, morning rush hour causes their arrival rate to drop significantly and this maps to the discrepancy of the red and blue curve in Figure 4(b).

In Figure 5(b) we observe that during peak hours, the travel time increases. Interestingly, the required travel time does not depend on the method nor whether it is on time or not. Thus, to achieve higher probability to arrive on time, one has to depart earlier (as seen in Figure 3(a)), but does not need to increase the time spent traveling. We believe that this is the case because the network of Zürich is quite dense, hence there exist different alternative journeys with comparable travel times.

**Choice of the Parameters for Buffer- $\xi$  and Mean-Risk- $c$ .** Figure 6(a) displays the minimum value of the parameter  $\xi$  of Buffer- $\xi$  that would be necessary to achieve arrival rates of 80%, and 90% of the cases in  $T_7$ , and how this value changes over the day. We see that, affected by the daily rush hours, this parameter varies significantly, suggesting that the Buffer- $\xi$  strategy needs a non-trivial amount of parameter adjustment. We observe that the dips corresponding to morning and evening rush hours are of the same height. Again, we can directly relate this behavior with the similarity of  $T_{plan}$  and  $T_7$  (captured by the red curve in Figure 4(b)).

Similarly, Figure 6(b) displays the value of the coefficient  $c$  of Mean-Risk- $c$  that would be necessary to achieve arrival rates of 78.5%, 90%, and 92.5% in  $T_7$ , and its development



■ **Figure 7** Influence of the number of transfers on the arrival rate.

during the day. We observe that this value is greatly affected by the morning rush hour. On the other hand, the dip corresponding to the evening rush hour is visible, but not too significant. Again, we link this behavior of the value to the observed similarity of the training/test instances with the planned timetable—the two curves captured by Figure 4(b). Recall that in the morning rush hour there is a gap between the two curves in Figure 4(b) indicating that the situation in  $T_7$  was not typical with respect to previous observations. As we see in Figure 6(b), the value of the coefficient  $c$  has to be quite large to compensate for the unexpected irregularities. In contrast, in a situation that is typical (i.e., when the two curves in 4(b) approximately match), the Mean-Risk- $c$  method performs well and fine-tuning of the parameters is not crucial. For instance, a coefficient  $c$  set to 1 leads to reasonably robust solutions.

**Influence of the Number of Transfers.** Figure 7(a) shows that the arrival rate of Buffer- $\xi$  (for  $\xi = 6$ ) is quite sensitive to the number of transfers. This suggests that the number of transfers is another aspect (of possibly many aspects) which has to be taken into account when searching for the best parameter for Buffer- $\xi$ . In contrast, Figure 7(b) shows that the influence of the number of transfers on the arrival rate of Mean-Risk- $c$  (for  $c = 1$ ) is almost negligible. Thus, there is no need to fine-tune the coefficient  $c$  to compensate for this aspect. We remark that we generally observed that the arrival rate of the methods based on the past observations is not very sensitive to the number of transfers.

## 5 Conclusion

We observed a clear trade-off: to achieve a higher probability to arrive on time in a network with delays, one has to depart earlier and expect higher standard deviation on the arrival time. On the other hand, the average travel time itself does not change with robustness or the choice of a routing method.

Methods based solely on the planned timetable, where the robustness is achieved by adding buffer times, need a non-trivial parameter adjustment for which many aspects need to be considered (time of the day, number of transfers, etc.). The methods that learn from past benefit from the additional knowledge: If the test instance is typical with respect to the past observations, these strategies perform well, Mean-Risk- $c$  does not need much fine-tuning, and Norm-Inf without parameter adjustment proposes a highly competitive solution with

reasonable trade-offs. We have seen that *Similarity* gives a good measure of the amount of irregularities in the network and can help to detect typical situations. Notably, it considers complex solutions (journeys), and thus it has a potential to capture behavior that cannot be observed only locally. We believe that this measure is worth further exploring, and by considering various aspects (e.g., how different approaches would benefit if *Similarity* was used to preselect typical instances for training) it can bring us even closer to the goal of robust routing.

The existence of equally good alternative journeys is one of the reasons why we believe that it was reasonable to choose the public transportation network of Zürich for our experiments, although the network is rather small in comparison to the public transportation networks of other cities. An interesting question is whether the algorithms are still sufficiently fast on larger networks. We believe that due to our solution concept (i.e., sequences of lines), the running time depends on the number of *lines* rather than the number of stops. In that respect, the network of Zürich is not exorbitantly small: For example, the public transportation network in Vienna has more than six times as many stops, but only less than two times as many lines. Hence, if the number of feasible *st*-routes (with a bounded number of transfers) is not too large, the algorithms should still work fast. Otherwise one could try to generate meaningful alternative routes in advance. Investigating these aspects and also whether our qualitative results hold for other cities are clearly interesting questions that we plan to investigate further.

**Acknowledgements.** We wish to thank the Verkehrsbetriebe Zürich (VBZ) for providing historic real-world delay data.

---

## References

- 1 Hannah Bast, Daniel Delling, Andrew Goldberg, Matthias Müller-Hannemann, Thomas Pajor, Peter Sanders, Dorothea Wagner, and Renato Werneck. Route planning in transportation networks. Technical Report MSR-TR-2014-4, Microsoft Research, 2014.
- 2 Hannah Bast, Jonas Sternisko, Sabine Storandt, et al. Delay-robustness of transfer patterns in public transportation route planning. In *ATMOS*, pages 42–54, 2013.
- 3 Kateřina Böhmová, Matúš Mihalák, Tobias Pröger, Rastislav Šrámek, and Peter Widmayer. Robust routing in urban public transportation: How to find reliable journeys based on past observations. In *ATMOS*, pages 27–41, 2013.
- 4 Justin Boyan and Michael Mitzenmacher. Improved results for route planning in stochastic transportation. In *SODA*, pages 895–902, 2001.
- 5 Joachim M. Buhmann, Matúš Mihalák, Rastislav Šrámek, and Peter Widmayer. Robust optimization in the presence of uncertainty. In *ITCS*, pages 505–514, 2013.
- 6 Julian Dibbelt, Thomas Pajor, Ben Strasser, and Dorothea Wagner. Intriguingly simple and fast transit routing. In *SEA*, pages 43–54, 2013.
- 7 Julian Dibbelt, Ben Strasser, and Dorothea Wagner. Delay-robust journeys in timetable networks with minimum expected arrival time. In *ATMOS*, pages 1–14, 2014.
- 8 Yann Disser, Matthias Müller-Hannemann, and Mathias Schnee. Multi-criteria shortest paths in time-dependent train networks. In *WEA*, pages 347–361, 2008.
- 9 Donatella Firmani, Giuseppe F. Italiano, Luigi Laura, and Federico Santaroni. Is time-tabling routing always reliable for public transport? In *ATMOS*, pages 15–26, 2013.
- 10 H Frank. Shortest paths in probabilistic graphs. *Operations Research*, 17(4):583–599, 1969.
- 11 Marc Goerigk, Sacha Heße, Matthias Müller-Hannemann, Marie Schmidt, and Anita Schöbel. Recoverable robust timetable information. In *ATMOS*, pages 1–14, 2013.

- 12 Marc Goerigk, Martin Knoth, Matthias Müller-Hannemann, Marie Schmidt, and Anita Schöbel. The price of robustness in timetable information. In *ATMOS*, pages 76–87, 2011.
- 13 Sejoon Lim, Christian Sommer, Evdokia Nikolova, and Daniela Rus. Practical route planning under delay uncertainty: Stochastic shortest path queries. In *Robotics: Science and Systems VIII*, 2012.
- 14 Matthias Müller-Hannemann and Mathias Schnee. Efficient timetable information in the presence of delays. In *Robust and Online Large-Scale Optimization*, pages 249–272. Springer, 2009.
- 15 Evdokia Nikolova, Jonathan A Kelner, Matthew Brand, and Michael Mitzenmacher. Stochastic shortest paths via quasi-convex maximization. In *ESA*, pages 552–563, 2006.