

# Evolutionary Minimization of Traffic Congestion

Maximilian Böther  
Hasso Plattner Institute  
University of Potsdam

maximilian.boether@student.hpi.de

Louise Molitor  
Hasso Plattner Institute  
University of Potsdam  
louise.molitor@hpi.de

Leon Schiller  
Hasso Plattner Institute  
University of Potsdam

leon.schiller@student.hpi.de

Martin S. Krejca\*  
Sorbonne University, CNRS, LIP6  
Paris, France  
martin.krejca@lip6.fr

Philipp Fischbeck  
Hasso Plattner Institute  
University of Potsdam  
philipp.fischbeck@hpi.de

Tobias Friedrich  
Hasso Plattner Institute  
University of Potsdam  
tobias.friedrich@hpi.de

## ABSTRACT

Traffic congestion is a major issue that can be solved by suggesting drivers alternative routes they are willing to take. This concept has been formalized as a strategic routing problem in which a single alternative route is suggested to an existing one. We extend this formalization and introduce the MULTIPLE-ROUTES problem, which is given a start and destination and aims at finding up to  $n$  different routes that the drivers strategically disperse over, minimizing the overall travel time of the system.

Due to the NP-hard nature of the problem, we introduce the MULTIPLE-ROUTES evolutionary algorithm (MREA) as a heuristic solver. We study several mutation and crossover operators and evaluate them on real-world data of Berlin, Germany. We find that a combination of all operators yields the best result, improving the overall travel time by a factor between 1.8 and 3, in the median, compared to all drivers taking the fastest route. For the base case  $n = 2$ , we compare our MREA to the highly tailored optimal solver by Bläsius et al. [ATMOS 2020] and show that, in the median, our approach finds solutions of quality at least 99.69% of an optimal solution while only requiring 40% of the time.

## CCS CONCEPTS

• **Computing methodologies** → **Heuristic function construction**; *Randomized search.*

## KEYWORDS

Strategic routing, traffic congestion, optimization, evolutionary algorithm

### ACM Reference Format:

Maximilian Böther, Leon Schiller, Philipp Fischbeck, Louise Molitor, Martin S. Krejca, and Tobias Friedrich. 2021. Evolutionary Minimization of Traffic Congestion. In *2021 Genetic and Evolutionary Computation Conference (GECCO '21)*, July 10–14, 2021, Lille, France. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3449639.3459307>

\*This work was supported by the Paris Île-de-France Region.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

GECCO '21, July 10–14, 2021, Lille, France

© 2021 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-8350-9/21/07.

<https://doi.org/10.1145/3449639.3459307>

## 1 INTRODUCTION

Traffic congestion is an increasing problem for urban areas across the world [8]. A solution is to route drivers by proposing them routes that reduce the overall travel time of the system, e.g., by navigation systems. Generally, proposing the *same* route to all drivers is infeasible, as this rather *causes* traffic congestion if the number of drivers is too high. Instead, drivers need to disperse over *different* routes, with some of them taking sub-optimal options into consideration [29] – a cost that some drivers are willing to take [19]. We refer to this setting as *strategic* routing.

A well-studied domain that meets some of these requirements is route planning [3]. Most results consider a time component of each route, e.g., by considering flow over time [20] or predicted congestion [13–15, 22], or they consider multiple routes, where the alternative route needs to be substantially different [1, 24]. However, none of these results take the overall travel time of the system or psychological factors of the drivers into account.

A problem that does consider road capacities and psychological models for route choices by drivers is the recently introduced SINGLE-ALTERNATIVE-PATH (SAP) problem [7], a strategic-routing problem that aims to find an optimal alternative route to a given route for a group of drivers. Still, the SAP problem is restricted to a *single* alternative route and requires one route to be given as an input. In this paper, we naturally extend the SAP problem to the more general MULTIPLE-ROUTES (MR) problem, which aims to minimize the overall travel time of all drivers in a system by proposing a set of routes to them, with the number of routes being controlled by a parameter. In order to account for bounded rationality and differing preferences by the drivers [32], we assume they form a user equilibrium on the given routes, i.e., a state in which no single driver can improve their travel time by choosing a different route.

Since the MR problem is NP-hard, we introduce the MULTIPLE-ROUTES evolutionary algorithm (MREA) to heuristically solve it. The MREA belongs to the class of evolutionary algorithms – nature-inspired metaheuristics that have been applied to great success to hard problems in various domains [12, 27], including non-strategic routing problems, e.g., the VEHICLE ROUTING PROBLEM [5, 25]. The MREA has a population size of  $\mu$ , uses four different mutation operators (changing a single solution), and employs crossover (combining different solutions) to find good solutions to the MR problem.

Using real-world data for the city of Berlin, Germany, provided by TomTom Germany, we evaluate all operators of the MREA for different route scenarios and compare them to the naive solution

of all drivers taking the fastest route. Our results (Table 1) show that using more mutation operators and a larger population size yields better solutions. All three crossover operators that we suggest perform almost equally well, such that one can choose the fastest. Depending on the route scenario, a best configuration of the MREA improves the overall travel time of the system by factors between 1.8 and 3, in the median. Even using a single mutation operator (a population size of 1 as well as no crossover) already improves the solution by factors between 1.5 and 2.8. Further, we adapt the MREA to the SAP problem and compare its solution quality to the deterministic, highly problem-specific exact solver of Bläsius et al. [7]. We find (see Figure 5) that the best configuration of the MREA, in the median, achieves a solution quality of at least 99.69% but only requires 40% of the run time. Overall, our results suggest that the MREA is a heuristic well-suited for solving the MR problem and thus reducing traffic congestion in strategic scenarios.

The remainder of this paper is organized as follows. In Section 2, we formalize the MR problem, and in Section 3, we introduce the MREA. We then analyze the performance of the MREA and the effect of its operators and population size in Section 4. In Section 5, we apply the MREA to the SAP problem and compare it against the algorithm of Bläsius et al. [7]. Finally, we conclude our work in Section 6. For supplementary material, we refer to our repository [6].

## 2 THE MULTIPLE-ROUTES PROBLEM

Given a route network graph  $G = (V, E)$  and a continuous flow of  $k \in \mathbb{R}_{\geq 0}$  drivers per unit of time between an origin  $s \in V$  and a destination  $t \in V$ , we consider routing this flow among  $n \in \mathbb{N}^+$  routes such that no driver can choose a quicker route as long as no other driver cooperatively changes their route as well and the total number of different routes exceeds  $n$ . We call such a state an *n-restricted user equilibrium* (n-UE). The MULTIPLE-ROUTES (MR) problem aims to find an optimal set of  $n$  routes such that the overall travel time of drivers in an n-UE is minimized.

### 2.1 Problem Modeling

We follow the formalization by Roughgarden and Tardos [26] but add the constraint of  $n$  routes. Let  $G = (V, E)$  be a directed graph,  $s \in V$ ,  $t \in V$ ,  $k \in \mathbb{R}_{\geq 0}$ , and  $n \in \mathbb{N}^+$ . Further, let  $\mathcal{P}_{s,t}$  denote the set of all routes from  $s$  to  $t$ . A traffic flow  $f: \mathcal{P}_{s,t} \rightarrow \mathbb{R}_{\geq 0}$  is a mapping that assigns to each  $P \in \mathcal{P}_{s,t}$  a value representing the amount of drivers on each edge of  $P$  per unit of time. Note that this value may not be integer. We call a traffic flow *valid* if and only if  $|\{P \in \mathcal{P}_{s,t} \mid f(P) > 0\}| \leq n$  and if  $\sum_{P \in \mathcal{P}_{s,t}} f(P) = k$ . Further, if and only if  $f$  is an n-UE (see Section 2.2), we call the traffic flow *stable*.

The travel time of drivers on an edge  $e \in E$  is determined by a latency function  $\tau_e: \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ . That is, for all  $x \in \mathbb{R}_{\geq 0}$ ,  $\tau_e(x)$  defines the time a single driver needs to travel along  $e$  assuming there are  $x$  agents entering  $e$  per unit of time. We assume  $\tau_e$  to be monotonically increasing and continuous. For a traffic flow  $f$ , the flow  $f_e$  over  $e$  is then  $\sum_{P \in \mathcal{P}_{s,t}: e \in P} f(P)$ , and the overall travel time of drivers on route  $P \in \mathcal{P}_{s,t}$  is  $\tau_P(f) = \sum_{e \in P} \tau_e(f_e)$ .

Last, for each traffic flow  $f$ , we associate a cost  $C(f)$  that denotes the overall travel time of all drivers. Formally,

$$C(f) = \sum_{P \in \mathcal{P}_{s,t}} f(P) \cdot \tau_P(f). \quad (1)$$

The MR problem aims to find a valid and stable traffic flow with minimum cost among all valid and stable traffic flows. We show in the supplementary material [6] that the MR problem is NP-hard.

### 2.2 The User Equilibrium

In routing games, a *user equilibrium* (UE), also known as *Wardrop equilibrium* [30], is a game state where no player has anything to gain by changing only their own strategy [23]. In the MR problem, we consider n-UEs, where no driver can improve their travel time by unilaterally changing their route while the traffic flow stays valid.

Given  $(G, s, t, k)$ , a UE always exists [4, 10, 26]. However, note that in contrast to UEs, an n-UE is a traffic flow with a route set of maximum size  $n$  where drivers are not allowed to choose a new route if this exceeds the number of  $n$  different routes in total. In particular, this means that an n-UE does not have to be unique, as it highly depends on  $n$ . Nonetheless, each valid UE is also an n-UE.

We approximate an n-UE by computing a UE under the constraint of using at most  $n$  routes. To this end, we model the UE as a convex problem [4], which we approximately solve with the FRANK-WOLFE algorithm [17], adjusted such that it makes sure to satisfy the constraint of at most  $n$  routes. The supplementary material [6] contains a detailed description of the algorithm.

## 3 THE MULTIPLE-ROUTES EA

The MREA (Algorithm 1) is an elitist EA for optimizing the MR problem. Given an MR instance  $(G, s, t, k, n)$ , it maintains a population of  $\mu$  route sets (the *individuals*), each of which consists of exactly  $n$  (not necessarily different) routes from  $s$  to  $t$ , and each of which is evaluated as described in Section 2.2. The MREA generates offspring in two different (and exclusive) ways: by (1) expanding the current population via a crossover operation provided as an input and (2) by employing a random number of mutation operators to a copy of each individual. Then, the MREA reduces the population size to  $\mu$  via truncation selection, breaking ties uniformly at random. Note that to avoid a single good individual being copied via crossover and then taking over the entire population, the offspring generated by crossover is only considered for selection if it contains a strict improvement over the parent population. The algorithm stops after a user-defined termination criterion. Although the MREA operates on *sets* of routes, many operators also perform changes to *single* routes. To this end, the subroutine RANDDIJKSTRA is used, which finds a shortest path on  $G$  with randomly perturbed edge weights.

### 3.1 RANDDIJKSTRA

The RANDDIJKSTRA (RD) is a randomized variant of Dijkstra's shortest-path algorithm [16]. Given two nodes  $s$  and  $t$ , it returns a random, yet still short route from  $s$  to  $t$ . RD works like Dijkstra's algorithm, but whenever relaxing an edge  $e$ , its weight  $w$  is perturbed such that  $w \sim N(\tau_e(x), 0.8 \cdot \tau_e(x))$ , where  $\tau_e$  is the latency of  $e$  and where  $x$  is the traffic flow routed from  $s$  to  $t$ . Due to its extensive use, RD contributes the most to the run time of the MREA.

### 3.2 Mutation Operators

In total, the MREA has four different mutation operators: *NewRoute*, *RandomP*, *LinkWP*, and *ExSegment*, each with its own weight. When

---

**Algorithm 1:** The MULTIPLE-ROUTES EA. Note that we use set notation, even though the sets are multisets.

---

**Input:** MR instance  $(G, s, t, k, n)$ , population size  $\mu$ , crossover strategy  $cStra$ , termination criterion

**Output:** Set of  $n$  routes from  $s$  to  $t$

```

1  $P \leftarrow \emptyset$ ;
2 repeat  $\mu$  times
3    $ind \leftarrow$  new individual;
4   repeat  $n$  times  $ind.addRoute(RANDDIJKSTRA(s, t, k))$ ;
5    $P.add(ind)$ ;
6 while termination criterion not met do
7    $C \leftarrow \emptyset$ ;
8   repeat  $\sqrt{\mu^2 - \mu/2}$  times
9      $ind_1, ind_2 \leftarrow$  two uniformly at random chosen
        individuals in  $P$ ;
10     $C.add(cStra(ind_1, ind_2))$ ;
11   $P' \leftarrow$  copy of  $P$ ;
12  for every individual  $ind$  in  $P'$  do
13     $mutations \leftarrow \max(1, \text{Pois}(1.5))$ ;
14     $ops \leftarrow \emptyset$ ;
15    repeat  $mutations$  times
16       $ops.add(\text{randomly weighted selected operator in}$ 
         $\{\text{NewRoute, RandomP, LinkWP, ExSegment}\})$ ;
17    if  $ops$  contains ExSegment then
18       $ops \leftarrow \{\text{ExSegment}\}$ ;
19    apply operators in  $ops$  to  $ind$ ;
20  if no individual in  $C$  is better than the best in  $P$  then
21     $C \leftarrow \emptyset$ ;
22   $P \leftarrow$  the  $\mu$  best individuals in  $C \cup P' \cup P$ ;
23 return the best individual in  $P$ ;

```

---

mutating an individual, the MREA first decides how many mutations to execute consecutively. This number is determined by a Poisson distribution with an expected value of 1.5, but at least one mutation is performed. Afterward, for each mutation to apply, a mutation operator is chosen randomly proportionally to its weight. If ExSegment is chosen, then all other operators are discarded for this mutation. Last, all chosen operators are applied to the individual.

**3.2.1 NewRoute.** The operator chooses a single route randomly proportionally to its inverse traffic flow and replaces this chosen route with one computed by RD. The weight of NewRoute is determined dynamically. In order to have a good exploration–exploitation tradeoff, it is 30 for the first 10 iterations, and then lowered linearly such that it reaches 1 in iteration 200.

**3.2.2 RandomP.** The operator replaces subsegments of a randomly selected subset of routes via RD. The routes to be modified are chosen randomly proportionally to their inverse traffic flow. For each such route, it chooses a start node uniformly at random and a destination node by advancing a number of steps according to the Gaussian distribution  $N(0.25r, 0.5r)$ , where  $r$  is the length of the route. Then, RD is applied to replace the route segment between

these two nodes. As the operator should find a different subsegment between these two nodes, RD increases the costs of the edges of the current route. Last, all cycles that may occur in the route after the replacement are deleted. RandomP has a constant weight of 60.

**3.2.3 LinkWP.** The operator is identical to RandomP except for the choice of delimiting nodes of the subsegment to replace, which are chosen more specifically to the MR instance. For each node  $v$  on a route in which a subsegment should be replaced, RandomP calculates a metric that describes how likely it is for a meaningful subroute to occur at  $v$ . The metric is defined as the sum of the capacities of all outgoing edges of  $v$  except for the edge currently used in the route. The start node is chosen randomly proportionally to the metric. Then, the destination node is chosen randomly by selecting one of the nodes on the original route that comes after the start node, weighted using the same metric.

LinkWP has a constant weight of 30. Note that this weight is lower than the one of RandomP in order to not introduce a too heavy problem-specific bias into the mutation step.

**3.2.4 ExSegment.** The operator swaps subsegments between two routes of the same individual. First, it chooses a pair of different routes from the route set uniformly at random. Then, all cycles are removed from both routes. Afterward, the nodes occurring in both routes are determined, which we call *shared points*. Among the shared points, let the *divergence points* be the nodes who have a different successor in both routes. Further, let the *goto points* be those shared points who have a different predecessor in both routes. ExSegment chooses one divergence point  $v_s$  uniformly at random. Then, it chooses a node  $v_t$  uniformly at random from the set of all goto points that appear after  $v_s$ . The route segments between  $v_s$  and  $v_t$  from both chosen routes are then swapped.

The weight of ExSegment is determined dynamically. If ExSegment was applied within the last 6 iterations, its weight is 0, as this operator is expensive and a too rapid succession of uses is unlikely to change much. If ExSegment was applied more than 6 iterations ago, its weight is determined as follows. It starts at 15 and is increased linearly to 30, depending on the iterations without improvement. The point in time when it reaches exactly 30 depends on the used convergence criterion (see Section 4 for more details).

### 3.3 Crossover Operators

We present three different binary crossover operators. Note, in contrast to the mutation operators, the MREA only uses a single crossover operator, specified by the input. This is due to there being a large tradeoff between run time cost and improvement in solution quality when considering different operators and due to the operators all being versions of the same idea.

Regardless of the operator chosen, the MREA creates  $\sqrt{\mu^2 - \mu/2}$  offspring in each iteration. Note that this number is the square root of all possible  $\binom{\mu}{2}$  2-combinations of  $\mu$  individuals. By the Birthday Paradox, the possibility of a combination of two individuals being chosen at least twice becomes over 50 % once in the order of this value. Thus, when creating  $\sqrt{\mu^2 - \mu/2}$  offspring, we aim to create as many individuals as possible without getting many doubles.

All of our proposed operators consider a score  $D$  that reflects how similar the routes of an individual are. The assumption is that

**Table 1: The median best fitness (lower is better) of the 75 runs (see Section 4.1) for each of the settings from Sections 4.2 to 4.4 for all 11 scenarios. The column  $k$ -Dijkstra states the best possible fitness if all drivers choose the fastest routes, accounting for delays caused by all drivers using the same street, using Dijkstra’s algorithm. This fitness is beaten by any of the MREA configurations. Already the fitness values of the weakest configuration ronly are between 35 % and 63 % of those of  $k$ -Dijkstra. In general, the fitness improves with better configurations (i.e., with entries further to the right, for the columns Mutation Operators and Population Size). The column Crossover Operators shows that the choice of the crossover operator has almost no impact regarding the median. Note that the column wexseg is the same as  $\mu = 1$  and that  $\mu = 4$  is the same as no\_hour due to how we conduct the experiments. Bold numbers indicate a significant change to the previous column (ignoring the deterministic  $k$ -Dijkstra), using the Mann–Whitney  $U$  test [21] with a  $p$ -value of 0.05. We refer to the respective sections for more information.**

ID Scenario	$k$ -Dijkstra	Mutation Operators (see Section 4.2)				Population Size $\mu$ (see Section 4.3)				Crossover Operators (see Section 4.4)			
		ronly	wnewroute	wlinkp	wexseg	$\mu = 1$	$\mu = 2$	$\mu = 4$	$\mu = 8$	no_hour	hour-all	hour-greed	hour-greed-rand
0 Babelsberg – Lichterfelde	153530661	68566307	<b>64240948</b>	64662372	64240948	61980563	<b>59016819</b>	<b>58128546</b>	59016819	58369109	58128546	59016819	
1 Griebnitzsee – Ahrensfelde	186995924	90878267	<b>74610652</b>	74610652	74610652	<b>74602749</b>	<b>74541545</b>	<b>74541545</b>	74541545	74541545	74541545	74541545	
2 KaDeWe – East Side Gallery	28713342	12364817	12292827	12237029	<b>12223856</b>	12223856	<b>12201922</b>	<b>12201922</b>	<b>12201922</b>	12201922	12201922	12201922	
3 Lichterfelde – Prenzlauer Berg	92721673	58477361	<b>50558429</b>	51197745	<b>52541858</b>	52541858	<b>49713176</b>	<b>49071681</b>	<b>48625397</b>	49071681	49071681	49044459	
4 Lichterfelde – Steglitz	86902859	50225754	<b>44211386</b>	<b>44910249</b>	44188815	44188815	<b>43540844</b>	<b>41699314</b>	<b>41699314</b>	41699314	41699314	41699314	
5 Moabit – Birkenwerder	103023491	38395379	<b>37639244</b>	37344703	37387321	37387321	37174180	<b>37174180</b>	<b>37174180</b>	37174180	37174180	37174180	
6 Olympiastadion – Rotes Rathaus	101643446	36504804	<b>36353121</b>	36266833	36347074	36347074	36259504	<b>36259504</b>	36259504	36259504	36259504	36259504	
7 Potsdamer Platz – Pergamonmuseum	40930113	18904046	<b>18764355</b>	18764355	18764355	18764355	<b>18764355</b>	<b>18636335</b>	<b>18558820</b>	18636335	18636335	18716682	
8 Potsdamer Platz – Tempelhofer Feld	12014974	6860945	<b>6407381</b>	6407381	6407381	6407381	6407381	6407381	6407381	6407381	6407381	6407381	
9 Teltow – Hoppegarten	394317060	160818584	<b>141236603</b>	140966107	142441589	142441589	<b>134170060</b>	<b>131325658</b>	<b>130535412</b>	131325658	131157891	131088667	
10 Wannsee – Schönefeld	57461353	24477839	<b>22899049</b>	22876414	22899049	22899049	<b>22876414</b>	<b>22876414</b>	<b>22876414</b>	22876414	22876414	22876414	

a more disjoint route set usually leads to a lower overall travel time due to less congestion on single roads. For an individual  $S$  and an edge  $e \in E$ , let  $c_e^S$  denote the count how often the edge appears in  $S$ . The score  $D$  of  $S$  is defined such that larger values are worse:

$$D(S) = \frac{\sum_{e \in \{e \in E \mid c_e^S > 1\}} c_e^2}{\max\left(1, \sum_{e \in \{e \in E \mid c_e^S = 1\}} c_e\right)}.$$

**3.3.1 Exhaustive Crossover.** The operator considers all  $\binom{2n}{n}$  route sets possible from the routes of the two parents and returns the combination with the highest diversity score.

**3.3.2 Greedy Crossover.** The operator greedily constructs a new route set, guided by the metric  $D$ . It randomly chooses one of the  $2n$  routes of the parents, proportionally to their inverse traffic flow. Each of the following  $n - 1$  routes is chosen such that it optimizes the diversity metric of the thereby created partial route set.

**3.3.3 Randomized Greedy Crossover.** This operator takes the same approach as Greedy Crossover, but instead of greedily choosing the route that maximizes the diversity of the route set, it randomly selects one of the  $2n$  routes, proportionally to the inverse of its diversity score. That is, the more diverse the route set with that route is, the more likely the route is to be chosen.

## 4 PARAMETER EVALUATION

We empirically analyze the utility of the operators of the MREA on the street network of Berlin, Germany. For each operator, we investigate how much the solution quality of the MREA changes when it is added to the algorithm. In Section 4.2, we begin by evaluating the mutation operators, excluding crossover. In Section 4.3, we analyze the impact of the population size  $\mu$ . Last, in Section 4.4, we add


crossover to the MREA, and we compare the quality achieved by the three different crossover operators with each other.

Our evaluations show that using more mutation operators, a larger population size, and crossover are all beneficial for improving the best fitness of the MREA. The largest improvement is made by adding the mutation operators RandomP and NewRoute. Nonetheless, adding more operators generally decreases the spread of the results, in addition to improving them. Table 1 shows a summary of the median best fitness of all our parameter settings.

### 4.1 Experimental Setup

We consider MR instances with the graph  $G$  being the street network of Berlin, Germany, provided by TomTom Germany, and with  $k = 3\,000$ , which is a reasonable choice [7]. Further, we choose  $n = 2$  in order to model proposing a driver with a small choice of fast routes. Choosing larger values makes this choice more troublesome for the driver, and it makes it also more unlikely to find that many different and fast routes. We choose the following 11 highly diverse scenarios (see also Table 1):

- (0) outside the inner-city; country road, non-obvious deviation
- (1) very long; fastest route uses Autobahn (AB; express highway), second fastest route goes through the inner-city
- (2) short, inner-city; lots of possible detours
- (3) long, south to north; covering AB and inner-city side streets
- (4) very short, inner-city; direct route is using side streets, but highway and AB are nearby
- (5) long, start in the city center; choice for highway or AB
- (6) long, inner-city; can be driven almost entirely on a highway
- (7) short, inner-city; different highways or side streets that are reasonable, in a Manhattan-like layout



**Figure 1:** Boxplots (see Section 4.1) of the normalized best fitness of the MREA with  $\mu = 1$  after 150 iterations for the 11 scenarios, with 75 runs per scenario. Each of the four colors, from left to right, represents one of the algorithm configurations explained in Section 4.2. Per scenario, the fitness is normalized to the median of ronly. In general, configurations with more mutation operators (more to the right per scenario) result in a better final fitness. Please refer to Section 4.2.1 for more details.

- (8) medium-long, inner-city; bottleneck at a bridge, but opportunity to split up onto two highways
- (9) long, south-west to east; either long detour using AB or a more direct inner-city highway
- (10) route where the  $k$ -Dijkstra shortest route detours to use AB

For the latency functions, we follow the recommendation of the US Bureau of Public Roads [28], that is, we choose  $\tau_e(x) = (\ell_e/s_e) \cdot 1.15(x/c_e)^2$  where  $s_e$ ,  $c_e$ , and  $\ell_e$  denote free-flow speed, capacity, and length of  $e$ , respectively [7].

For the experiments, we consider various *settings*. For each, the termination criterion of the MREA is to stop after 150 iterations. The weight of ExSegment (Section 3.2.4) is chosen such that it reaches a value of 30 if there was no improvement in the last  $20\% \cdot 150 = 30$  iterations. Last, we start 75 independent runs of the MREA on all 11 scenarios per setting. We refer to our repository [6] for details.

*Boxplots.* When visualizing our results with boxplots, the box denotes the mid-50% of the 75 runs, and the whiskers denote the mid-90%. All remaining data points are depicted as diamonds.

*Solution space size.* Our results indicate that many runs with different settings have equal fitness. This suggests that the solution space is small, highlighting the impact of adding a new operator.

## 4.2 Analysis of the Mutation Operators

We analyze the utility of the MREA’s four mutation operators (Section 3.2) by considering how well each operator improves the fitness of the overall best solution (Section 4.2.1) and how quickly the algorithm finds such a solution (Section 4.2.2). To this end, we do not employ crossover, and we choose a population size of  $\mu = 1$  in order to see how much a single solution can be improved by mutation only. Further, we consider four different algorithm configurations, starting with a single operator and then adding more operators:

- (1) the MREA has only access to RandomP (ronly),
- (2) ronly but adding NewRoute (wnewroute),
- (3) wnewroute but adding LinkWP (wlinkp),
- (4) using all four operators (wexseg).

*4.2.1 Best overall fitness.* We study the impact of each configuration on the best fitness achieved after our termination criterion of 150 iterations. Our results are depicted in Figure 1.

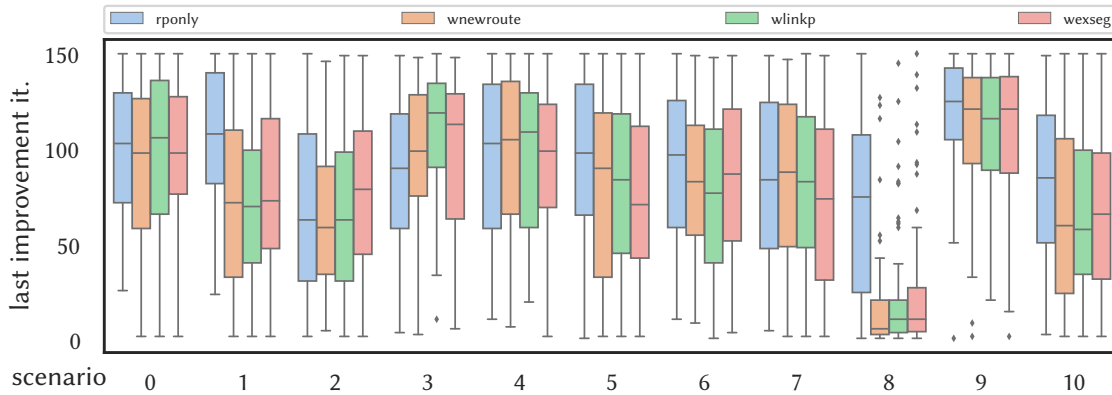
Adding NewRoute yields the largest improvement, with a statistical significance for all scenarios, except for scenario 2. This could be due to it being very short. Thus, RandomP and NewRoute become very similar operations. Averaged over all 11 scenarios, 84% of the wnewroute runs are better and 87% are better or equal to the median of the ronly runs. For scenarios 4 and 8, all runs of wnewroute are better than the median of ronly. This is likely a result of scenarios 4 and 8 requiring two almost disjoint routes, which are more easily found by NewRoute, whereas other scenarios require two nearly identical routes.

Interestingly, in scenarios 3 and 5, LinkWP as well as ExSegment increase the median best fitness. For LinkWP, recall that it prefers edges with a high capacity. If the best routes do not use such edges, LinkWP has no benefit. Nonetheless, averaged over all scenarios, 84% of the wlinkp runs are better and 88% are better or equal to the median of the ronly runs.


For ExSegment, recall that it swaps segments locally optimally, with respect to the segments randomly chosen. Escaping from such a local optimum can prove hard in certain scenarios, especially since we only consider a population size of 1. Still, on average, 38% of the wexseg runs are better and 56% are better or equal to the median of the wlinkp runs, showing a general benefit of ExSegment.

*4.2.2 Speed of convergence.* In order to analyze how quickly the MREA reaches a local optimum from which it cannot escape within its budget of 150 iterations, we consider the last iteration in that the MREA changed the fitness of its best individual. The results are depicted in Figure 2. Note that this analysis does not consider the fitness of each run, only whether it changed in subsequent iterations or not. For a more complete picture, please also refer to the results from Section 4.2.1, which show that, on average, configurations with more operators have a better median performance.

The speed of convergence is dependent on the scenario, and there is no clear trend among the four configurations. The mid-90%



**Figure 2:** Boxplots (see Section 4.1) of the last of, in total, 150 iterations in which the MREA with  $\mu = 1$  improved its best fitness, for the 11 scenarios. Each of the four colors, from left to right, represents one of the algorithm configurations explained in Section 4.2, and each configuration was run 75 times per scenario. Regardless of the scenario, there is a large spread between runs that get stuck quickly and runs that do not converge within 150 iterations. Please refer to Section 4.2.2 for more details.



**Figure 3:** Boxplots (see Section 4.1) of the normalized best fitness of the MREA after 150 iterations for the 11 scenarios, with 75 runs per scenario. Each of the four colors, from left to right, represents a different population size  $\mu$ . Per scenario, the fitness is normalized to the median of  $\mu = 1$ . In general, a higher population size seems more beneficial, but the gain is diminishing. Please refer to Section 4.3 for more details.

are generally close to the extreme values of 0 and 150. Runs close to 0 show that the scenarios are hard, as the MREA gets stuck very quickly. In contrast, runs close to 150 show that the budget of 150 iterations was insufficient for the MREA to converge.


**4.2.3 Conclusion.** Averaged over all scenarios, more mutation operators lead to a better performance. However, this effect is not very well pronounced for the addition of LinkWP, indicating that it should possibly be merged with the similar operator RandomP. Still, using both operators is overall better than just using RandomP.

The large spread in the speed of convergence among all configurations and scenarios suggests that the initialization has a large impact on how easy it is to find improvements, more or less regardless of what configuration is run. This indicates that a larger population size may be beneficial, as it increases the initial diversity.

### 4.3 Analysis of the Population Size

We analyze to what extent the MREA benefits from having a population size larger than 1. Since Section 4.2 suggests that local optima pose a problem for the MREA, a larger population size may help to have alternative solutions to those stuck in local optima. We do not employ crossover but use all four mutation operators, that is, we use the wexseg configuration. Our results are depicted in Figure 3. Note that a higher population size also means more fitness evaluations, as we let each configuration run for 150 iterations. This likely explains the high significances between different configurations.

A larger number of individuals improves the median best fitness and reduces the spread. Although the impact of a larger population size is different among the scenarios, our results suggest that the improvement for  $\mu = 2$  and  $\mu = 4$  provide a large improvement over



**Figure 4:** Boxplots (see Section 4.1) of the normalized best fitness of the MREA after 150 iterations for the 11 scenarios, with 75 runs per scenario. Each of the four colors, from left to right, represents one different crossover operator (including no crossover). Per scenario, the fitness is normalized to the median of no\_heur. The configuration heur-all performs best, but only slightly. There is no clear difference between heur-greed and heur-greed-rand. In general, using crossover reduces the spread of the results. Please refer to Section 4.4 for more details.

$\mu = 1$ . For  $\mu = 8$ , the improvement in comparison to  $\mu = 1$  in median and spread is somewhat smaller. Throughout all scenarios, 61% of the runs with  $\mu = 2$  are better and 77% are better or equal to the median of  $\mu = 1$ ; for the runs with  $\mu = 4$ , these numbers increase to 80% and 93%, respectively. For  $\mu = 8$ , the increase from  $\mu = 4$  is smaller, reaching 88% better runs and 98% better or equal runs.

When comparing to the configuration with  $\mu = 2$ , 40% of the runs with  $\mu = 4$  are better than the median and 85% of the runs are better or equal; for  $\mu = 8$ , these numbers increase to 49% and 96%.

**4.3.1 Conclusion.** Using a larger population size improves the quality of the best fitness and also decreases the spread among the different runs per scenario. However, the computational effort increases with the number of individuals, and the quality gain in fitness from using more individuals varies among the different configurations. Our experiments suggest a sweet spot at 4 individuals.

#### 4.4 Analysis of the Crossover Operators

We analyze the utility of the MREA’s three crossover operators (Section 3.3), measuring the overall best fitness for each operator. To this end, we use all mutation operators, choose  $\mu = 4$ , and consider the following configurations:

- (1) using no crossover (no\_heur),
- (2) using Exhaustive Crossover (heur-all),
- (3) using Greedy Crossover (heur-greed), and
- (4) using Randomized Greedy Crossover (heur-greed-rand).

Our results are depicted in Figure 4. The advantage of crossover strongly depends on the scenario and none are significant. However, averaged over all 11 scenarios, the median best fitness as well as the spread is always reduced when using a crossover operator in comparison to using no crossover. This is also true for the minimum and maximum normalized fitness, highlighting the reduction of outliers. Interestingly, heur-all does not have a large benefit over the two greedy operators. Considering the two greedy strategies, on

average, 19% of the heur-greed runs are better and 72% are better or equal to the median of no\_heur; for heur-greed-rand, we get 18% and 74%, respectively. There is no clear tendency whether heur-greed or heur-greed-rand performs better.

**4.4.1 Conclusion.** In general, crossover improves the result quality of the MREA and also reduces its spread. Among the different crossover operators, Exhaustive Crossover performs best but only slightly. Considering its high computational cost compared to the other two operators, it should not be chosen. Greedy Crossover and Randomized Greedy Crossover provide very good alternatives, each of which performs roughly equally well.

## 5 APPLICATION TO THE SAP PROBLEM

We apply the MREA to the SINGLE-ALTERNATIVE-PATH (SAP) problem [7] and empirically investigate its performance in terms of solution quality and run time. The SAP problem is a special case of the MR problem that fixes a route between  $s$  and  $t$  and aims to find a *single* alternative route such that the overall travel time is minimized. Although the problem remains NP-hard, Bläsius et al. [7] propose a highly specialized algorithm that solves it optimally – the SAP baseline (SAP-B) –, which we compare the MREA against.

Since the SAP problem is a special case of the MR problem, the complexity of the MREA reduces in certain aspects. In addition, we use the insights gained from Section 4 to further improve the MREA. We call the resulting algorithm the SAP-EA.

### 5.1 The SAP-EA

The SAP-EA is a specialization of the MREA for the SAP problem with some modifications to its mutation operators. Since the SAP problem aims to find a single alternative route, an individual in the SAP-EA corresponds to a single route. Further, since determining a user equilibrium for the SAP problem simplifies to equalizing the cost functions of the given and the alternative route, which

results in solving a quadratic equation, the SAP-EA does not use the FRANK-WOLFE algorithm for fitness evaluation.

Regarding the operators proposed in Sections 3.2 and 3.3, the SAP-EA does not employ crossover, as these operators exchange existing routes, which is pointless for a single route. For the same reason, ExSegment is not used. Out of the remaining operators, NewRoute is used unmodified, and RandomP and LinkWP are combined into the new operator RandomPwD. This is due to our results from Section 4.2 showing that LinkWP only provides a small benefit when added but still has its merits for certain scenarios. Last, the SAP-EA always performs exactly one mutation on each individual, using a parameter  $p \in (0, 1)$  instead of operator weights. With probability  $p$ , NewRoute is performed, otherwise RandomPwD.

**5.1.1 RandomPwD.** Similar to RandomP, given a route  $R$  of length  $m$ , RandomPwD replaces a segment of  $R$  between two nodes  $a$  and  $b$  that are  $k$  apart via RandDijkstra. To this end, RandomPwD uses a parameter  $\delta \in [0, 1]$ , which it automatically adjusts. It determines  $k \sim N(\delta \cdot m, (0.05 \cdot m)^2)$ , rounding to the closest whole number, chooses  $a$  uniformly at random, and chooses  $b$  such that it is  $k$  nodes after  $a$ . If there are fewer than  $k$  nodes after  $a$ , then  $b = t$ .

RandomPwD adjusts  $\delta$  according to two parameters  $\alpha \in [0, 1]$  and  $\beta \in \mathbb{N}_{>0}$  in the following way: whenever the SAP-EA does not improve for  $\beta$  iterations, we update  $\delta \leftarrow \alpha \cdot \delta$ .


**5.1.2 Comparison to the SAP-B.** Although the SAP-EA is specialized for the SAP problem, it is still a general heuristic applicable to different fitness functions. In contrast, the SAP-B is explicitly tailored to solving the SAP problem with monotone cost functions per edge, such as the flow of traffic, as in our setting. Thus, the SAP-B fails for other costs, for example, when optimizing for overall low CO<sub>2</sub> emissions of strategic drivers in a street network. In such a setting, the SAP-EA is still applicable without change.

## 5.2 Empirical Investigations

We compare the SAP-EA to the SAP-B on the street network of Berlin, Germany, with respect to best fitness as well as run time. Recall that the SAP-B is an optimal algorithm. Thus, the SAP-EA cannot outperform it in regard to best fitness.

**5.2.1 Experimental Setup.** We use the same setup as in Section 4.1, with the following differences. We consider 25 scenarios chosen uniformly at random from the set of cluster centers of  $s$ - $t$  pairs, computed by the BIRCH [31] algorithm. The clustering is based on real-world traffic density data provided by TomTom Germany. Per scenario, we choose  $k \in \{500, 1000, 1500, 2000\}$  and  $p \in \{0.0, 0.01, 0.05, 0.1, 0.2, 0.3, 0.4\}$ , and we perform 20 runs per value of  $k$  and  $p$ . The SAP-EA terminates after 1000 iterations or whenever it does not improve for 100 iterations. Last, we choose  $\mu = 1$ . We used a machine with two Intel Xeon Gold 5118 CPUs and 64 GiB of memory.

**5.2.2 Experimental Evaluation.** Our results are depicted in Figure 5. The maximum of all medians in the score ratio is 1.009, for  $p = 0$ , which is already very close to an optimal fitness. The median decreases up to  $p = 0.2$  and increases afterward. Further, the spread is smallest for  $p = 0.2$ , making this configuration preferable. However, the run time ratio increases for higher values of  $p$  both in median and spread, as RandomPwD is computationally more expensive



**Figure 5: The ratio of the best fitness (left) and the run time (right) of the SAP-EA with  $\mu = 1$ ,  $\alpha = 0.4$ ,  $\beta = 35$ , and different values of  $p$  compared to the SAP-B. The run times of SAP-B range from 0.3 seconds to 30 minutes, with better time ratios for higher SAP-B run times. Each boxplot contains the data of all 20 runs per value of  $k$  and per each of the 25 scenarios, totaling to 2000 points per box. The orange line depicts the median, the box the mid-50% of the data, and the whiskers the mid 95%. A higher value of  $p$ , i.e., an increased use of RandomPwD, yields generally better solutions and a smaller spread but also increases the run time. A sweet spot seems to be around  $p = 0.05$ . Please also refer to Section 5.2.2.**

than NewRoute. Since the configuration with  $p = 0.05$  is very close to the best configuration both in score and time, we deem it the best configuration out of all.

## 6 CONCLUSION

We introduced and empirically analyzed the MULTIPLE-ROUTES EA, an evolutionary algorithm designed to suggest alternative routes for street networks with a high flow of traffic with the aim to reduce the overall travel time of all drivers. To this end, we introduced the NP-hard MULTIPLE-ROUTES problem, allowing for a precise modeling of our setting. For the MREA, we proposed four mutation and three crossover operators, and we showed that using all mutation operators yields the best results. Further, each crossover operator reduces the spread of the results. Last, we applied the MREA to a more specific setting, aiming to find a single alternative route to a given route, and compared it to a highly specialized optimal algorithm. Although the MREA is more general, it is capable of competing with the tailored algorithm while often being faster.

Overall, our results suggest that the MREA is well-suited for the highly complex problem of distributing traffic. For future work, we propose to extend the MREA to island models [9], a parallelization method well suited for EAs [2]. Another direction is to use data sets that measure other criteria, for example, the emission of cars. We believe that the MREA is well suited for such settings.

## ACKNOWLEDGMENTS

We thank TomTom Location Technology Germany GmbH for supplying us with the data necessary for our experiments. Further, we thank Galassi et al. [18] for the GNU Scientific Library and Dagum and Menon [11] for OpenMP.



## REFERENCES

- [1] Ittai Abraham, Daniel Delling, Andrew V. Goldberg, and Renato F. Werneck. 2013. Alternative Routes in Road Networks. *Journal of Experimental Algorithmics* 18, Article 1.3 (2013), 17 pages. <https://doi.org/10.1145/24444016.24444019>
- [2] Enrique Alba, Gabriel Luque, and Sergio Nesmachnow. 2013. Parallel metaheuristics: recent advances and new trends. *International Transactions on Operational Research* 20, 1 (2013), 1–48. <https://doi.org/10.1111/j.1475-3995.2012.00862.x>
- [3] Hannah Bast, Daniel Delling, Andrew V. Goldberg, Matthias Müller Hannemann, Thomas Pajor, Peter Sanders, Dorothea Wagner, and Renato F. Werneck. 2016. Route Planning in Transportation Networks. In *Algorithm Engineering: Selected Results and Surveys*. Springer International Publishing. [https://doi.org/10.1007/978-3-319-49487-6\\_2](https://doi.org/10.1007/978-3-319-49487-6_2)
- [4] Martin J. Beckmann, Charles B. MacGuire, and Christopher B. Winsten. 1956. *Studies in the Economics of Transportation*. Yale University Press.
- [5] Jean Berger and Mohamed Barkaoui. 2003. A Hybrid Genetic Algorithm for the Capacitated Vehicle Routing Problem. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO) 2003*. 646–656. [https://doi.org/10.1007/3-540-45105-6\\_80](https://doi.org/10.1007/3-540-45105-6_80)
- [6] Thomas Bläsius, Maximilian Böther, Philipp Fischbeck, Tobias Friedrich, Alina Gries, Falk Hüffner, Otto Kißig, Martin S. Krejca, Pascal Lenzner, Louise Molitor, Leon Schiller, Armin Wells, and Simon Wietheger. 2021. *Strategic Routing GitHub Repository*. Retrieved April 03, 2021 from <https://github.com/MaxiBoether/strategic-routing>
- [7] Thomas Bläsius, Maximilian Böther, Philipp Fischbeck, Tobias Friedrich, Alina Gries, Falk Hüffner, Otto Kißig, Pascal Lenzner, Louise Molitor, Leon Schiller, Armin Wells, and Simon Wietheger. 2020. A Strategic Routing Framework and Algorithms for Computing Alternative Paths. In *Proceedings of the 20th Symposium on Algorithmic Approaches for Transportation Modelling, Optimization, and Systems (ATMOS) 2020*. 10:1–10:14. <https://doi.org/10.4230/OASICS.ATMOS.2020.10>
- [8] Nick Cohn. 2019. *The TomTom Traffic Index: An Objective Measure of Urban Traffic Congestion*. Retrieved January 27, 2021 from <https://www.tomtom.com/blog/road-traffic/urban-traffic-congestion/>
- [9] James Patrick Cohoon, Shailesh U. Hegde, Worthy N. Martin, and Dana S. Richards. 1987. Punctuated Equilibria: A Parallel Genetic Algorithm. In *Proceedings of the Second International Conference on Genetic Algorithms and Their Application*. 148–154. <https://doi.org/10.18130/V3219P>
- [10] Stella Dafermos. 1980. Traffic Equilibrium and Variational Inequalities. *Transportation Science* 14, 1 (1980), 42–54. <https://doi.org/10.1287/trsc.14.1.42>
- [11] Leonardo Dagum and Ramesh Menon. 1998. OpenMP: An Industry Standard API for Shared Memory Programming. *IEEE Computational Science and Engineering* 5, 1 (1998), 46–55. <https://doi.org/10.1109/99.660313>
- [12] Kalyanmoy Deb and Christie Myburgh. 2016. Breaking the Billion-Variable Barrier in Real-World Optimization Using a Customized Evolutionary Algorithm. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO) 2016*. 653–660. <https://doi.org/10.1145/2908812.2908952>
- [13] Daniel Delling. 2009. Time-Dependent SHARC-Routing. *Algorithmica* 60, 1 (2009), 60–94. <https://doi.org/10.1007/s00453-009-9341-0>
- [14] Daniel Delling and Dorothea Wagner. 2009. Time-Dependent Route Planning. In *Robust and Online Large-Scale Optimization: Models and Techniques for Transportation Systems*. Springer, 207–230. [https://doi.org/10.1007/978-3-642-05465-5\\_8](https://doi.org/10.1007/978-3-642-05465-5_8)
- [15] Ugur Demiryurek, Farnoush Banaei-Kashani, and Cyrus Shahabi. 2010. A Case for Time-Dependent Shortest Path Computation in Spatial Networks. In *Proceedings of SIGSPATIAL 2010*. ACM, 474–477. <https://doi.org/10.1145/1869790.1869865>
- [16] Edsger Wybe Dijkstra. 1959. A Note on Two Problems in Connexion with Graphs. *Numer. Math.* 1, 1 (1959), 269–271. <https://doi.org/10.1007/BF01386390>
- [17] Marguerite Frank and Philip Wolfe. 1956. An Algorithm for Quadratic Programming. *Naval Research Logistics Quarterly* 3, 1-2 (1956), 95–110. <https://doi.org/10.1002/nav.3800030109>
- [18] Mark Galassi, Jim Davies, James Theiler, Brian Gough, Gerard Jungman, Patrick Alken, Michael Booth, Fabrice Rossi, and Rhys Ulerich. 2019. *GNU Scientific Library Reference Manual*. Network Theory Ltd.
- [19] Alexander Kröller, Falk Hüffner, Łukasz Kosma, Katja Kröller, and Mattia Zeni. 2021. Driver Expectations towards Strategic Routing. *Transportation Research Record* (2021). To appear.
- [20] Ekkehard Köhler, Rolf H. Möhring, and Martin Skutella. 2009. Traffic Networks and Flows over Time. In *Algorithmics of Large and Complex Networks*. Springer, 166–196. [https://doi.org/10.1007/978-3-642-02094-0\\_9](https://doi.org/10.1007/978-3-642-02094-0_9)
- [21] Henry B. Mann and Donald R. Whitney. 1947. On a Test of Whether one of Two Random Variables is Stochastically Larger than the Other. *The Annals of Mathematical Statistics* 18, 1 (1947), 50–60. <https://doi.org/10.1214/aoms/1177730491>
- [22] Giacomo Nannicini, Daniel Delling, Dominik Schultes, and Leo Liberti. 2011. Bidirectional A\* search on time-dependent road networks. *Networks* 59, 2 (2011), 240–251. <https://doi.org/10.1002/net.20438>
- [23] John F. Nash. 1950. Equilibrium Points in n-Person Games. *Proceedings of the National Academy of Sciences* 36, 1 (1950), 48–49. <https://doi.org/10.1073/pnas.36.1.48>
- [24] Andreas Paraskevopoulos and Christos D. Zaroliagis. 2013. Improved Alternative Route Planning. In *Proceedings of ATMOS 2013*. Schloss Dagstuhl, 108–122. <https://doi.org/10.4230/OASICS.ATMOS.2013.108>
- [25] Jean-Yves Potvin. 2009. State-of-the Art Review—Evolutionary Algorithms for Vehicle Routing. *INFORMS Journal on Computing* 21, 4 (2009), 518–548. <https://doi.org/10.1287/ijoc.1080.0312>
- [26] Tim Roughgarden and Eva Tardos. 2002. How Bad is Selfish Routing? *J. ACM* 49, 2 (2002), 236–259. <https://doi.org/10.1145/506147.506153>
- [27] Dan Simon. 2013. *Evolutionary Optimization Algorithms*. Wiley-Blackwell.
- [28] United States Bureau of Public Roads, Office of Planning, Urban Planning Division. 1964. *Traffic Assignment Manual for Application with a Large, High Speed Computer*.
- [29] Mariska Alice van Essen. 2018. *The Potential of Social Routing Advice*. Ph.D. Dissertation. University of Twente. <https://doi.org/10.3990/1.9789055842377>
- [30] John Glen Wardrop. 1952. Some Theoretical Aspects of Road Traffic Research. *Proceedings of the Institution of Civil Engineers* 1, 3 (1952), 325–362. <https://doi.org/10.1680/ipeds.1952.11259>
- [31] Tian Zhang, Raghu Ramakrishnan, and Miron Livny. 1996. BIRCH: An Efficient Data Clustering Method for Very Large Databases. *ACM SIGMOD Record* 25, 2 (1996), 103–114. <https://doi.org/10.1145/235968.233324>
- [32] Shanjiang Zhu and David Levinson. 2015. Do People Use the Shortest Path? An Empirical Test of Wardrop's First Principle. *PLOS ONE* 10, 8 (2015), 1–18. <https://doi.org/10.1371/journal.pone.0134322>