

Dynamic Bi-Objective Routing of Multiple Vehicles

Jakob Bossek
School of Computer Science
The University of Adelaide, Australia

Christian Grimme
Information Systems and Statistics
University of Münster, Germany

Heike Trautmann
Information Systems and Statistics
University of Münster, Germany

ABSTRACT

In practice, e.g. in delivery and service scenarios, Vehicle-Routing-Problems (VRPs) often imply repeated decision making on dynamic customer requests. As in classical VRPs, tours have to be planned short while the number of serviced customers has to be maximized at the same time resulting in a multi-objective problem. Beyond that, however, dynamic requests lead to the need for re-planning of not yet realized tour parts, while already realized tour parts are irreversible. In this paper we study this type of bi-objective dynamic VRP including sequential decision making and concurrent realization of decisions. We adopt a recently proposed Dynamic Evolutionary Multi-Objective Algorithm (DEMOA) for a related VRP problem and extend it to the more realistic (here considered) scenario of multiple vehicles. We empirically show that our DEMOA is competitive with a multi-vehicle offline and clairvoyant variant of the proposed DEMOA as well as with the dynamic single-vehicle approach proposed earlier.

CCS CONCEPTS

• **Applied computing** → **Multi-criterion optimization and decision-making**; **Transportation**.

KEYWORDS

Vehicle routing, decision making, multi-objective optimization, dynamic optimization, evolutionary algorithms

ACM Reference Format:

Jakob Bossek, Christian Grimme, and Heike Trautmann. 2020. Dynamic Bi-Objective Routing of Multiple Vehicles. In *Genetic and Evolutionary Computation Conference (GECCO '20)*, July 8–12, 2020, Cancún, Mexico. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3377930.3390146>

1 INTRODUCTION

Routing of multiple vehicles is an important and difficult problem with applications in the logistic domain [34], especially in the area of customer servicing [13]. In postal services, after-sales services, and in business to business delivery or pick up services one or more vehicles have to be efficiently routed towards customers. If customers can request services over time, the problem becomes dynamic: besides a set of fixed customers, new requests can appear at any point in time. Of course, it is desirable that as many customers

as possible are serviced while the tour of any vehicle is kept short. However, it is usually infeasible (due to human resources, labor regulations, or other constraints) to service all customer requests. And clearly, the less customers are left unserved, the longer the tours become. Thus, the problem is inherently multi-objective. Any efficient solution (smallest maximum tour across all vehicles) is a compromise between the desire to service as many customers as possible (e.g. maximize revenue) and the necessity to keep vehicle routes short (minimize costs). At the same time, the dynamic appearance of new customer requests may significantly change the scenario over and over again: new but ignored requests negatively contribute to the objective of visiting as many customers as possible, while the inclusion of new customers (usually) increases tour length and thus changes the compromises on which a selection of a route was originally made by a decision maker (DM).

This dynamic problem has been studied by Bossek et al. [5] for the special case of a single vehicle which answers all requests and travels (in an open tour) from a start to an end depot. The authors devised a dynamic evolutionary approach based on an interactive algorithmic framework that incorporated an evolutionary multi-objective optimization algorithm (EMOA) applied in eras and repeated decision making. However, the applied EMOA is rather unrealistically based on the assumption, that only one vehicle is available.

In this work, we will reuse the framework proposed by Bossek et al. [5] but replace the internal EMOA [4] by an adapted algorithm that is capable of considering multiple vehicles. The inclusion of multiple vehicles changes the problem (and thus the algorithm) considerably: Instead of a single tour (single open TSP), as many tours as considered vehicles have to be optimized simultaneously. This implies changes in problem encoding, in information transfer between generations, and in variation operators.

At the same time, the number of vehicles is explicitly not considered as additional objective. To keep the scenario realistic, the number of vehicles can neither be changed during the optimization process nor in each era. A dynamic change in the number of used vehicles during the process would require most flexible (and thus costly) human resources and is therefore usually infeasible for a company.¹ Additionally, a third objective would turn the originally bi-objective problem in a more complex-to-handle decision scenario for the DM.

The goal and contribution of this work is twofold:

- (1) The dynamic multi-objective vehicle routing problem (MO-VRP) and the dynamic solution approach are extended towards a more realistic scenario by including multiple vehicles. We introduce a significantly changed algorithm within the interactive framework proposed in [5].

Permission to make digital or hard copies of all or part 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 components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
GECCO '20, July 8–12, 2020, Cancún, Mexico

© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-7128-5/20/07...\$15.00
<https://doi.org/10.1145/3377930.3390146>

¹Visits at single or few customers (including direct travel from and return to a depot) would immediately contribute high costs to the total tour costs.

- (2) We analyze the benefit of multiple vehicles in dynamic vehicle routing and compare our approach to an (extended) version of an a-posteriori evolutionary solution approach for this problem [17]. This approach unrealistically knows of all service requests in advance (clairvoyant) and thus needs no dynamic decision making during optimization. Further, we compare the multi-vehicle approach to the dynamic approach by Bossek et al. [5] and investigate the activities of vehicles. The individual activities of vehicles provide information on whether each vehicle contributes to the solution or whether some vehicles stay idle. This evaluation can eventually justify our decision to not include the number of vehicles as third objective.

The work is structured as follows: the next section briefly reflects the related work, while Section 3 formally introduces the dynamic multi-objective problem as described by Bossek et al. [5]. Section 4 then details the algorithmic extensions. The experimental setup as well as empirical results are described and discussed in Sections 5 and 6. Section 7 finally concludes the work and highlights perspectives for future research.

2 RELATED WORK

As the traveling salesperson problem (TSP) is a major sub-problem of the here considered dynamic and multi-objective vehicle routing problem, this paper is naturally related to work on special TSPs, where not all customer locations (or cities) have to be visited. In research, these problems are sometimes referred to as orienteering problems [15], selective TSP [14, 24], or as TSP with profits [11]. However, most of these problems are discussed as single-objective problems [10], although some early work already recognized the (at least) bi-objective character of these problems [22]. Only later work started to solve the orienteering problem in a bi-objective way using an ϵ -constraint approach [2] or approximation schemes [12] that produce Pareto- ϵ -approximations of the efficient solution set. While the both before mentioned approaches are based on repeated single-objective optimization, some authors [20, 21, 30, 35, 37] explicitly solve the bi-objective variants of the orienteering problem using an evolutionary algorithm, however, excluding service requests over time or considering the problem as a-posteriori (non-dynamic). Many of these approaches [21, 30, 35] introduce the number of vehicles as an objective to be minimized while simultaneously minimizing the tour length.

A related a-posteriori variant of the here considered dynamic problem is described by Grimme et al. [17], who propose an NSGA-II-based EMOA. This work has been extended later on by the integration of local search mechanisms [27] and the analysis of local search effects [4]. These works only allow one vehicle (like also described in [15, 36]) but include the number of visited customers (revenue) as second objective besides tour length (costs).

While considering only one vehicle seems to be unrealistic, the inclusion of the number of vehicles as objective is only feasible in the a-posteriori and non-dynamic case. When problem instances change constantly due to customer requests (related examples from logistics and other domains can be found here: [16, 28, 32]) decision making is also a repetitive process over time. However, over time, decisions are constantly renewed building on past decisions which

of course cannot be changed. In vehicle routing, one or more vehicles start at a depot and travel initially decided tours. Later on, new decisions have to take into account the current location as well as newly received or not yet serviced customer requests [5]. Rewinding of previous decisions (i.e. visits of customers) is impossible. As such, the initial decision for a fixed number of vehicles could only be changed by sending vehicles home or activating additional ones. This however, causes additional traveling costs and contradicts (in the real world) human resources' availability or labor regulations. Thus, it is most realistic not to consider the number of vehicles as additional objective in the dynamic case.

In general, dynamic vehicle routing is usually addressed by designing online decision rules, see e.g. [26, 31]. According to Braekers et al. [7] only little work is available on dynamic multi-objective problems. In their survey they mention authors who consider dynamics in service time windows and changing structures of the network [1, 18, 23, 25, 38].

3 PROBLEM NOTATION

The here considered dynamic multi-objective VRP can be denoted as follows: we consider a set of customer locations, which can be partitioned into three disjoint sets, $C = M \cup D \cup \{N - 1, N\}$. The subset M contains all mandatory customers locations that are initially known and have to be visited, while the subset D contains all dynamic customers that appear over time and are not known to the algorithm beforehand. The third subset $\{N - 1, N\}$ denotes the locations of the start and the end depot. Note that we consider the more general case here, in which start and end depot can be different. The more common special case of start and end depot being at the same location is of course included.

We consider two objectives in a minimization problem. The first objective aims for the minimization of the maximum tour length for all vehicles. Let n_v be the number of vehicles and x a solution to the problem, then we denote the tour length for each vehicle $i \in \{1, \dots, n_v\}$ as $L_i(x)$ and determine $f_1(x) = \max_i L_i(x)$. By using the maximum tour length as objective, we expect a balanced usage of vehicles in any solution. The second objective f_2 minimizes the number of unserved dynamic customers².

Clearly, the objectives are in conflict and we need to adopt the notion of Pareto-optimality and dominance to describe compromise solutions for the resulting multi-objective optimization problem. For two solutions x and y , we denote $x < y$ (x dominates y), if x is not worse in any objective and better in at least one objective than y . The set of all non-dominated solutions in search space is called *Pareto set*; its image in objective space is called *Pareto front* [8].

As we consider a dynamic problem, anytime a dynamic customer requests service, the Pareto-set would have to be recomputed for the still unvisited mandatory and dynamic customers and a desired solution needs to be selected. As this is usually infeasible in practice, we discretize time and define n_t intervals of length $\Delta \in \mathbb{R}_{\geq 0}$ called *eras* that partition dynamic requests and subsequent decision making into phases [33].

²Note that we consider the number of unserved dynamic customers to realize a minimization of all objectives. Clearly, the second objective is equivalent to the maximization of served dynamic customers.

At the onset $t = (j - 1) \cdot \Delta$ of each era j , new dynamic customer requests may have appeared. Based on this set and the remaining (not yet visited) mandatory customers in M , we can consider the problem as a static multi-objective optimization problem (MOP) and apply an EMOA to approximate the current Pareto-set. Then, a decision maker (DM) is provided with the compromises and allowed to pick a solution (i.e. a set of n_v tours) which will be realized until the onset of the next era.

Note that in each era $j > 0$ the vehicle has started to realize a tour and possibly has already visited mandatory and/or dynamic customers. Naturally, already realized parts of earlier picked solutions are not reversible anymore. Thus, decisions made in earlier eras may have significant influence on later solutions. We address this challenge by introducing an automated decision making process as proposed in [5] and evaluating different configurations and decision chains, later on.

4 A DYNAMIC MULTI-OBJECTIVE EVOLUTIONARY ALGORITHM

Next we dive into the working principles and algorithmic details of the proposed DEMOA. The algorithm is a natural extension of the DEMOA proposed in Bossek et al. [6] for the single-vehicle version of the considered bi-objective problem. The algorithmic steps are outlined in Alg. 1. The algorithm requires the following parameters: the problem instance comprising of the subsets M of mandatory and D of dynamic customers. Further parameters control the number and the length of eras (n_t and Δ), the number of vehicles n_v and EA-specific arguments like the population size μ and parameters controlling for the strength of mutation (details discussed later). We now describe the DEMOA procedure in general and discuss implementation details (initialization and variation) subsequently.

Algorithm 1 DEMOA

Require: **a)** Customer sets M, D , **b)** nr. of eras n_t ,
c) era length Δ , **d)** nr. of vehicles n_v , **e)** population size μ ,
f) prob. to swap p_{swap} , **g)** nr. of swaps n_{swap}

- 1: $t \leftarrow 0$ \triangleright current time
- 2: $P \leftarrow \emptyset$ \triangleright population (initialized below)
- 3: $T \leftarrow$ list of tours \triangleright empty at the beginning of 1st era
- 4: **for** $i \leftarrow 1$ to n_t **do** \triangleright era loop
- 5: $T^{\leq t} \leftarrow$ list of n_v partial tours already driven by vehicles at time t extracted from list T \triangleright empty in 1st era
- 6: $P \leftarrow \text{INITIALIZE}(\mu, T^{\leq t}, t, P)$ \triangleright see Alg. 2; pass last population of previous era as template
- 7: **while** stopping condition not met **do** \triangleright EMOA loop
- 8: $Q \leftarrow \{\text{MUTATE}(x, T^{\leq t}, p_{\text{swap}}, n_{\text{swap}}) \mid x \in P\}$ \triangleright Alg. 3
- 9: $Q \leftarrow \{\text{LOCALSEARCH}(x) \mid x \in P\}$
- 10: $P \leftarrow \text{SELECT}(Q \cup P)$ \triangleright NSGA-II survival-selection
- 11: $T \leftarrow \text{CHOOSE}(P)$ \triangleright DM makes choice \rightsquigarrow list of n_v tours
- 12: $t \leftarrow t + \Delta$

4.1 General (D)EMOA

Initialization steps (Alg. 1, lines 1-3) consist of declaring a population P and a list T where T_v contains the tour of the corresponding

Algorithm 2 INITIALIZE

Require: **a)** pop. size μ , **b)** initial tours $T^{\leq t}$, **c)** time t ,
d) template population P

- 1: $Q \leftarrow \emptyset$
- 2: **for** $j \leftarrow 1$ to μ **do**
- 3: **if** P is empty **then** \triangleright 1st era; no template given
- 4: $x.v_i \leftarrow$ random vehicle from $\{1, \dots, n_v\}$ for all $i \in C$
- 5: $x.a_i \leftarrow 1$ for all $i \in M$
- 6: $x.a_i \leftarrow 0$ for all $i \in D$
- 7: $x.p \leftarrow$ random permutation of $C = M \cup D$
- 8: **else** \triangleright repair template
- 9: $x \leftarrow P_j \in P$
- 10: $x.a_i \leftarrow 1$ for all $i \in T^{\leq t}$
- 11: $x.v_i \leftarrow$ vehicle nr. assigned to i in $T^{\leq t}$
- 12: In $x.p$ move sub-sequence of driven tour in $T^{\leq t}$ at the beginning for each vehicle.
- 13: $x.a_i \leftarrow 1$ for each i in random subset of D^{new}
- 14: $Q \leftarrow Q \cup \{x\}$

return Q

vehicle $v \in \{1, \dots, n_v\}$, i.e., T stores the solution the decision maker picked at the end of the previous era. Before the first era begins these tours are naturally empty since no planning was conducted at all. Line 4 iterates over the eras. Here, the actual optimization process starts. The first essential step in each era is – given the passed time t – to determine for each vehicle $v \in \{1, \dots, n_v\}$ the initial tour already realized by vehicle v . This information is extracted from the list of tours T and stored in the list $T^{\leq t}$. Note that again, in the first era the initial tours are empty, as is T , since the vehicles are all located at the start depot. Next, the population is initialized in line 6 and a static EMOA does his job in lines 7 – 10. Here, offspring solutions are generated by mutation followed by a sophisticated genetic local search procedure with the aim to reduce the tour lengths of solutions. Finally, following a $(\mu + \lambda)$ -strategy the population is updated. Therein, the algorithm relies on the survival selection mechanism of the NSGA-II algorithm [9]. Once a stopping condition is met, e.g., a maximum number of generations is reached, the era ends and the final solution set is presented to a decision maker who needs to choose exactly one solution (Alg. 1, line 11). This choice (the list T of n_v tours) determines the (further) order of customers to be visited by the respective service vehicles.

Algorithm 3 MUTATE

Require: **a)** individual x , **b)** initial tours $T^{\leq t}$, **c)** swap prob. p_{swap} ,
d) n_{swap}

- 1: $D_{av} \leftarrow$ dyn. customers available at time t and **not in** $T^{\leq t}$
- 2: $C_{av} \leftarrow$ all customers available at time t and **not in** $T^{\leq t}$
- 3: $p_a \leftarrow 1/|D_{av}|$
- 4: $p_v \leftarrow 1/|C_{av}|$
- 5: Flip $x.a_i$ with prob. p_a for all $i \in D_{av}$
- 6: Change vehicle $x.v_i$ with prob. p_v for all $i \in C_{av}$
- 7: **if** random number in $(0, 1) \leq p_{\text{swap}}$ **then**
- 8: Exchange n_{swap} times each two nodes from C_{av} in $x.t$

return x

4.2 Initialization

We strongly advise the reader to consult Alg. 2 and in particular Fig. 1 in the course of reading the following explanations for visual support. Each individual is built of three vectors $x.v, x.a$ and $x.p$ of length $N - 2$ each of which stores information on the vehicles assigned, the activation status of each customer and a permutation of all customers. In the first era there are no already visited customers, i.e. both P and $T^{\leq t}$ are empty, and the algorithm does not need to take these into account. Hence, in lines 4-7 each individual $x \in P$ is created from scratch as follows: each customer $i \in (M \cup D)$ is assigned a vehicle $v \in \{1, \dots, n_v\}$ uniformly at random. This information is stored in the vector $x.v$. Next, all mandatory customers $i \in M$ are activated by setting the value of a binary string $x.a \in \{0, 1\}^{N-2}$ to 1 which means “active”. In contrast, all dynamic customers $i \in D$ are deactivated ($x.a_i = 0$) since they did not ask for service so far. The final step is to store a random permutation of all customers in the permutation vector $x.p$. Note that during fitness function evaluation in order to calculate the individual tour lengths $L_v(x)$ for each vehicle v only the sub-sequence of positions $i \in \{1, \dots, N - 2\}$ in $x.p$ is considered for which $x.a_i = 1$ and $x.v_i = v$ holds.

Initialization in later eras is different and more complex. Now, the parameter P passed to Alg. 2 – the final population of the previous era – is non-empty and its solutions serve as templates for the new population. We aim to transfer as much information as possible. However, usually the majority of individuals $x \in P$ is in need of repair. This is because as time advances (note that additional Δ time units have passed) further customers $i \in C$ may already have been visited by the vehicle fleet, but it is possible that some of these are inactive in x (i.e., $x.a_i = 0$). Hence, (a) all customers which have already been visited, stored in the list $T^{\leq t}$, are activated and assigned to the responsible vehicle (here we use an overloaded element-of-relation on lists in the pseudo-codes for convenience) and (b) furthermore their order in the permutation string $x.p$ is repaired. The latter step is achieved by moving the sub-sequence of visited customers before the remaining customers assigned to the corresponding vehicle in the permutation string. This step completes the repair procedure and the resulting individual is guaranteed to be feasible. Subsequent steps involve randomly activating customers that asked for service within the last Δ time units.

4.3 Offspring generation

The mutation operator (see Alg. 3) is designed to address all three combinatorial aspects of the underlying problem, i.e., vehicle re-assignment, customer (de)activation and tour permutation. Here, special attention has to be paid to not produce infeasible individuals. Therefore, mutation operates on the subset of customers which have asked for service until now and have not yet been visited. More precisely, each dynamic customer $i \in D_{av} = (D^{\leq t} \setminus T^{\leq t})$ which has not yet been visited is (de)activated with a small probability p_a (note that we treat the list $T^{\leq t}$ as a set here for convenience). Likewise, each of the customers $i \in C^{\leq t} = (C^{\leq t} \setminus T^{\leq t})$ is assigned another vehicle independently with equal probability p_v . The mutation probabilities p_a and p_v are set dynamically such that in expectation only one (de)activation or (re)assignment happens; small changes

Customer sets

$M = \{1, 5, 6\}, D = \{2, 3, 4\}$
 $D^{\leq t} = \{3, 4\}$

(Partial) tours

$T_1 = (1, 5, 3), T_2 = (6)$
 $T_1^{\leq t} = (1, 5), T_2^{\leq t} = (6)$

Encoding of solution x

i	1	2	3	4	5	6
$x.v$	1	2	1	1	1	2
$x.a$	1	0	1	0	1	1
$x.p$	4	1	6	5	3	2




Figure 1: Illustration of the encoding of an individual x . Here, customers $i \in \{1, 3, 5, 6\}$ are active ($x.a_i = 1$) while customers $i \in \{2, 4\}$ are inactive ($x.a_i = 0$); customer 4 however asked for service already since $4 \in D^{\leq t}$. In contrast, customer 2 did not ask for service so far (illustrated by reduced opacity in the plot). The vehicles already visited a subset of customers (illustrated with thick edges): vehicle one serviced customers 1 and 5 (thus $T_1^{\leq t} = (1, 5)$) while customer 6 was visited by vehicle 2 (thus $T_2^{\leq t} = (6)$).

are preferred. Finally, with probability p_{swap} the permutation vector $x.p$ undergoes n_{swap} sequential exchange/swap operations (limited to customers which are not fixed so far). Occasionally, at certain iterations, a local search (LS) procedure is applied to each individual $x \in P$ (see line 9 in Alg. 1). The LS takes the vehicle mapping $x.v$ and customer activation $x.a$ as fixed and aims to improve the individual path length by means of the sophisticated solver EAX [29] for the Traveling-Salesperson-Problem (TSP). To accomplish this goal, given a solution $x \in P$ and a vehicle $v \in \{1, \dots, n_v\}$ all customers assigned to v in x (appending start and end depot) and their pairwise distances are extracted (nodes 1, 3, 5, 7, 8 for vehicle 1 in Fig. 1). Next, since the optimization of each vehicle tour is a Hamilton-Path-Problem (HPP) on the assigned customers (no round-trip tour), a sequence of distance matrix transformations is necessary such that the TSP solver EAX can be used to find an approximate solution to the HPP (see [19] for details). Again, since partial tours might already have been realized, the HPP-optimization starts at the last node visited by the corresponding vehicle (node 5 in Fig. 1 for vehicle 1, since $T_1^{\leq t} = (1, 5)$ is fixed already and not subject to changes).

5 EXPERIMENTAL METHODOLOGY

Our benchmark set consists of in total 50 instances with each $n = 100$ customers taken from [27]. There are 10 instances with points spread uniformly at random in the Euclidean plane and each 10 clustered instances with the number of clusters $n_c \in \{2, 3, 5, 10\}$. The cluster centers are placed by space-filling Latin-Hypercube-Sampling to ensure good spread. Subsequently, $\lfloor n/n_c \rfloor$ nodes are placed around each cluster center assuring cluster segregation with no overlap. We refer the reader to [27] for more details on the instance generation process. The proportion of dynamic customers

Table 1: DEMOA parameter settings.

Parameter	Value
Nr. of function evaluations per era	65 000
Population size μ	100
Swap probability p_{swap}	0.6
Nr. of swaps n_{swap}	10
Local Search at generations	first, half-time, last




Figure 2: Exemplary Pareto-front approximations for two instances with 75% dynamic requests colored by era and split by problem type (columns) for a single vehicle. Solutions selected by the respective decision maker (0.25-strategy) are highlighted and labeled with the era number. Dashed horizontal lines represent the maximum number of dynamic requests which can remain unserved in the respective era.

is 50% and 75% for each half of the instance set enabling the study of highly dynamic scenarios.

5.1 Dynamic aspects and decision making

We fix $n_t = 7$ eras and set the era-length to $\Delta = \lceil \max_{i \in D} r(i) / n_t \rceil$ where $r(i)$ is the request time of customer i . Δ is consistently ≈ 150 time units across all instances. Naturally, one could start a new epoch once a new customer requests for service. This would result in 50 and 75 eras respectively on our benchmark instances³. However, we argue that in a real-world scenario it is more realistic to make decisions after chunks of requests came in and not every single time. For computational experimentation we automate the decision-making process by considering three different decision maker (DM) strategies. To do so, at the end of each era, we sort the final DEMOA population P in ascending order of the first objective (tour length), i.e. $P_{(1)} \leq P_{(2)} \leq \dots \leq P_{(\mu)}$ where the \leq -relation is with respect to f_1 . Note that in the bi-objective space this sorting results in a descending order with respect to the number of unvisited dynamic customers, our second objective (f_2). The automatic DM now picks the solution $P_{(k)}$ with $k = \lceil d \cdot \mu \rceil$, $d \in [0, 1]$ where increasing d -values correspond to stronger “customer-greediness”, i.e. higher emphasis on keeping the number of unvisited customers low. In our study we cover $d \in \{0.25, 0.5, 0.75\}$ to account for different

³Note that the benchmark sets contains instances with $N = 100$ customers (including two depots) and $\{50\%, 70\%\}$ dynamic customers.

levels of greediness and refer to such a policy as a d -strategy in the following. Certainly, in real-world scenarios, the DM might change his strategy throughout the day reacting to specific circumstances. However, for a systematic evaluation and to keep our study within feasible ranges, we stick to this subset of decision policies.

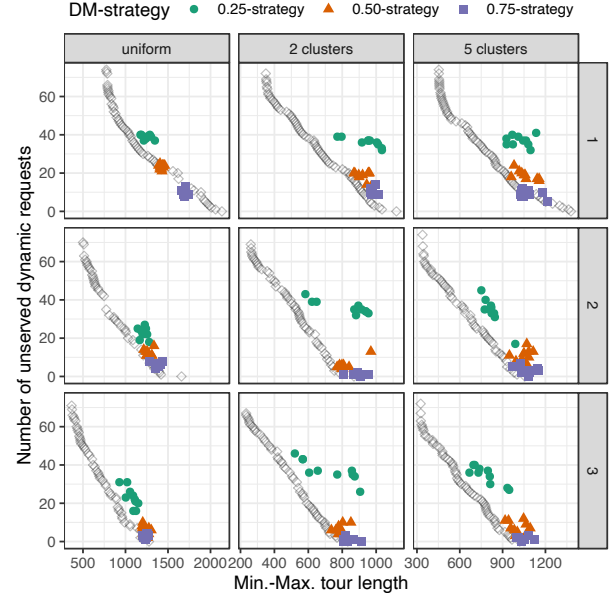


Figure 3: Exemplary visualization of final DM-decisions (i.e., in last era) for three representative instances with 75% dynamic requests colored by decision maker strategy. The data is split by instance (columns) and number of vehicles used (rows).

5.2 Further parameters

The further parameter settings of the DEMOA stem from preliminary experimentation and are gathered in Table 1. For each number of vehicles $n_v \in \{1, 2, 3\}$ and each DM strategy $d \in \{0.25, 0.5, 0.75\}$ we run the DEMOA 30 times on each instance for n_t eras following the d -strategy. Moreover, for a baseline comparison, we run the clairvoyant EMOA multiple times for each n_v on each instance with a stopping condition of 20 000 000 function evaluations. The clairvoyant EMOA works the same way the DEMOA does. However, it has complete knowledge of the request times of dynamic customers a-priori and treats the problem as a static problem.⁴ This idea was originally introduced in [17] for an a-posteriori evaluation of decision making for the single-vehicle variant of the considered problem. We use the Pareto-front approximations of the clairvoyant EMOA as a baseline for performance comparison. We provide the R implementation of our algorithm in an accompanying repository [3].

⁴Consider the clairvoyant EMOA as the proposed DEMOA (see Alg. 1) run for one era with $t = 0$ and all customers available from the beginning accounting for request times in the tour length calculations.

6 EXPERIMENTAL RESULTS

In the following, we analyze the proposed dynamic approach (DEMOA) for one to three vehicles with an adapted clairvoyant implementation of the original approach by Bossek et al. [4]. For a fair comparison, that approach has been extended to deal with multiple vehicles and is denoted as EMOA in the following results.

In a first step, we show exemplary results of our dynamic approach to visually introduce the era concept. In continuation of the approach of Bossek et al. [5], we also briefly investigate the influence of decision making to final decision location, when different greediness preferences are considered. Then, we compare the performance of the EMOA and the proposed DEMOA with respect to three different performance measures. In detail, we investigate the overall performance with respect to problem type and fleet size, the performance with respect to different decision strategies, and the overall performance gain induced by using multiple vehicles. Finally, we zoom into dedicated solution instances (and their dynamic evolution process) to learn about the behavior of vehicles on clustered and uniform problem instances.

6.1 Pareto-front approximations and decisions in the dynamic scenario

The dynamic nature of the problem and online decision making imply that we do not have a single point in time at which the algorithm performance can be evaluated. In the beginning only mandatory customers are available and the tour planning task is equivalent to (multi-vehicle and open) TSP solving. However, as dynamic customer requests appear over time, the initially planned tour(s) must be modified to allow compromises between tour length and number of visited customers. From this point on, a Pareto-set of solutions has to be considered. Following the era concept of decision making, at a dedicated point in time, a compromise is chosen for realization by a DM. From that time onward, realization starts and vehicles travel the decided tour. Of course, new dynamic requests appear over time. These are considered at the end of the next era and form a set of new compromises. However, that set of compromises has to consider the already realized partial tour of the vehicles, which cannot be reverted. Consequently, already visited (formerly) dynamic customers reduce the upper bound of unserved customers in compromise sets for future decisions. The effect of repeated decision making and continuous realization of decisions is exemplarily shown in Figure 2. The non-dominated fronts for decisions in all seven eras are shown for a uniform and a clustered instance, respectively. For visual comparison, the clairvoyant EMOA results are also shown. Horizontal dashed lines denote the upper bound of unserved dynamic customers in each stage of decision making. Clearly, in the last era (7, brown points), more customers have been visited by the traveling vehicles than in the first era. Thus, the upper bound has decreased. The upper bound and the range of possible decisions is also depending on the decision strategy. A greedy strategy, which aims to reduce the number of unserved customers will favor solutions with many served customers and thus influence realization of longer tours. Less greedy strategies will favor realizations with shorter tours and less visited customers. This behavior was already observed by Bossek et al. [5] for a single vehicle. In Figure 3, we confirm an analogous behavior

also for multiple vehicles and our algorithmic approach. Therein, visualizations of the results of different strategies (0.25, 0.5, and 0.75 priority of the second objective) and multiple runs for different topologies (uniform and clustered) as well as for different numbers of vehicles are shown. We find, that decision strategies are reflected in the final decision locations. The less greedy strategy produces solutions with more unserved customers than very greedy strategies - independent of the vehicle number.

6.2 Dynamic and clairvoyant performance and the influence of multiple vehicles

In order to evaluate the approximation quality of the DEMOA, the dynamic nature of the problem has to be respected. From the decision maker's perspective, the last era (and thus the last non-dominated front) includes all previous decisions and can be compared with the clairvoyant results. However, due to the continuously decreasing upper bounds of unserved customers during the optimization process (see 6.1), the EMOA approximation covers a wider range of solutions than the approximated Pareto-front of the final era. This is considered by our comparison of equivalent ranges of the DEMOA and EMOA results.

In Figure 4, we compare the Hypervolume indicator [39] of the DEMOA and EMOA by (1) determining instance-wise the minimal upper bound for the number of unserved customers (objective f_2) in the last approximated Pareto-front and for all independent runs of the DEMOA. Then (2), we reduce the solutions of the EMOA to those below the before determined upper bound for f_2 . From the union of the reduced EMOA results and the DEMOA results, we (3) compute a reference point for Hypervolume computation. Here we compare both algorithms for the medium greedy 0.5 strategy. We find, that the DEMOA results outperform the EMOA results for uniform problem instances with both 50% and 75% dynamic customers. For clustered instances, performance is in the same range but seemingly depending on the specific instance topology and dynamic customer ratio and service request times. However, we can conclude that the results of the DEMOA for the more realistic dynamic scenario are not necessarily worse than those of the clairvoyant approach. While the clairvoyant EMOA approach knows about the request times of all dynamic customers at $t = 0$ and considers all potential customers in compromise generation, the DEMOA optimizes tours of era $i + 1$ based on partially realized (and unchangeable) tours from era i . This often reduces the size of the tour planning problem significantly and allows to gain very good solutions - some of those even outperform the EMOA solutions in uniform instances.

Another observation from Figure 4 shows that the aforementioned advantage of problem complexity reduction due to dynamism becomes neglectable, when both algorithms consider multiple vehicles. With increasing number of vehicles, the EMOA tends to outperform the approximation quality of the DEMOA.

While we focused on a single decision strategy in the previous discussion, Figure 5 investigates multiple strategies together with tour length properties. Instead of the hypervolume, we now focus on the tour length objective (f_1) and measure the distance of solutions to the clairvoyant solution. The idea behind this measure is, that (in the range of the common objective space of DEMOA and EMOA) usually any number of unvisited dynamic customers is covered

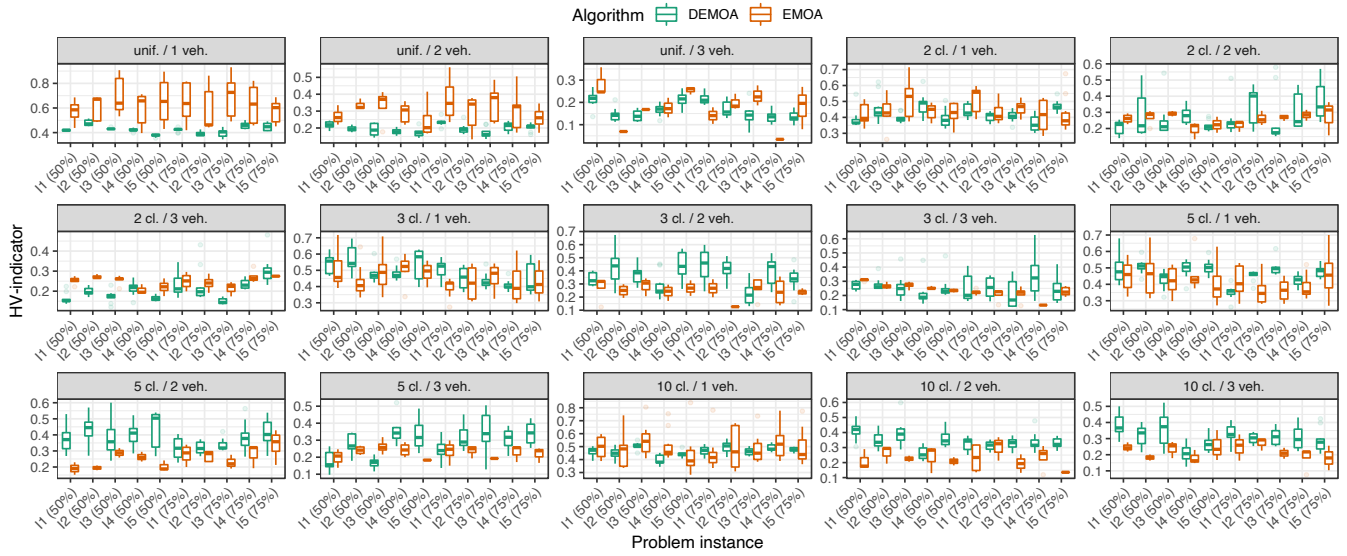


Figure 4: Boxplots of hypervolume-indicator (lower is better) for all instances split by problem type and fleet size. Results are shown for the 0.5-strategy for visual clarity, but the omitted results show the same patterns. Hypervolume values are calculated by (1) instance-wise calculation of the minimal upper bound for the number of unvisited dynamic customers in the last era across all independent runs of the DEMOA, (2) determining the reference point based on the union of all Pareto-front approximations of DEMOA and EMOA chopped at the bound and (3) calculating the HV-value based on this reference set.

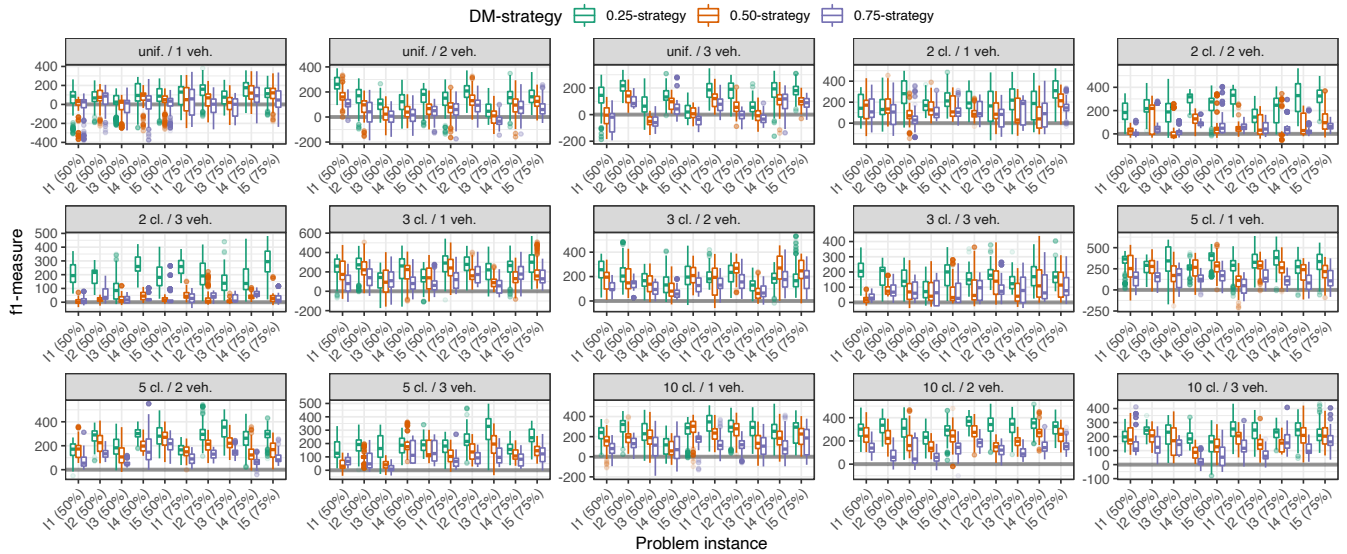


Figure 5: Boxplots of the f_1 -measure (tour length of DEMOA solution minus tour length of best clairvoyant EMOA solution with the same number of unvisited customers) calculated on basis of all solutions in the last era.

by the approximated Pareto-front. The convergence quality of a solution set can thus also be expressed by the difference of DEMOA tour length and EMOA tour length. If DEMOA outperforms EMOA, the respective value is negative. The equilibrium of solution quality is denoted by a gray vertical line at 0 in Figure 5.

While we are able to confirm the general observations from above, we now have detailed insights into the effect of decision

strategies. As a clear trend, we find better tour lengths per instance topology and dynamism, when greediness in decision making w.r.t. reduction of number of unserved customers (f_2) is increased. This holds for the single and multiple vehicle case. The reason for this observation is similar to the argument discussed before. The more customers have to be serviced, the more complex the tour planning problem becomes for the clairvoyant EMOA, while more and more

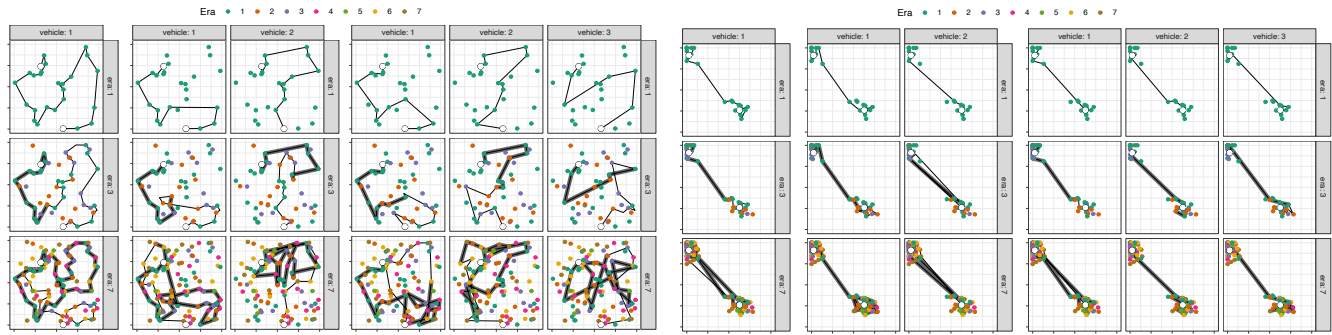


Figure 6: Exemplary tours in eras 1, 3 and 7 (final) with 0.75-strategy for a uniform instance (left) and a clustered instance (right). For each instance we show the tours the DM picked for a scenario of a single vehicle (left-most column), two vehicles (columns 2 and 3) and three columns (remaining three columns). Bold edges represent the irreversible tour parts which already have been realized by the corresponding vehicle (columns) in the respective era (rows).




Figure 7: Classical dominated hypervolume distributions (higher is better) of representative instances. The trend is the same for all 50 benchmark instances.

customers are fixed in the dynamic scenario and do not have to be considered for tour planning anymore.

A dedicated view on the effect of using multiple vehicles in the scenario (online and offline) is provided in Figure 7. We find for all investigated cluster configurations and the uniform distribution of customers, that multiple vehicles are advantageous regarding classical hypervolume comparison⁵; the results are statistically highly significant with respect to Wilcoxon-Mann-Whitney tests at significance level $\alpha = 0.001$ in 100% of the cases, i.e. two vehicles are significantly better than one and three better than two. At the same time, we observe that the highest gain in solution quality is associated with the step from one to two vehicles. By including a third vehicle, only little is gained. This is probably rooted in the overhead associated with the distances travelled from the start depot to the first customer and from the last customer to the end depot. These distances occur for each vehicle and have to be travelled no matter how short the remaining tour becomes. Thus, this overhead will naturally bound the amount of reasonably applicable vehicles.

6.3 Exemplary analysis of vehicle tours

In this paragraph, we briefly investigate representative examples of generated solutions. Figure 6 details the evolution of tours for a single vehicle as well as two, and three vehicles for a uniform and

⁵Here the reference point is determined for each problem (independent of the number of vehicles and the strategy). The dominated hypervolume is then calculated for each algorithm, DM-strategy and number of vehicles.

clustered instance in eras one, three, and 7. In the single vehicle case, we can just observe the dynamic adaptation of the planned tour (thin line) towards the realized (and irreversible) tour (bold lines) over time. In the two and three vehicle scenarios, we can nicely observe, how the vehicles automatically partition the customers space. It is obvious, that no vehicle stays idle. Moreover, we find, that each vehicle is assigned similar workload. This behavior is partly rooted in the design of our algorithm. As objective f_1 minimizes the maximum tour length across all vehicles, selection pressure forces similar workload to each vehicle. Future research has to clarify, whether this desirable feature from a real-world application perspective is always advantageous from the optimization perspective.

7 CONCLUSION AND OUTLOOK

In this work, we successfully extended an already high-performing single-vehicle approach to the more realistic dynamic multiple vehicle scenario and we proposed two measures for comparing DEMOA quality to the performance of the related clairvoyant EMOA variant. We find that the algorithmic enhancements ensure a nice distribution of workload (in terms of tour lengths and number of customers served) between the involved vehicles without the necessity for explicitly optimizing for this kind of balance. At the same time and especially on instances with random uniformly distributed locations, the DEMOA can even outperform the offline EMOA variant with full knowledge of request times. Due to concurrent realization of planned tours, Pareto front approximations of the DEMOA's decision eras naturally concentrate on constantly shrinking problem sizes. These reduced problems can then be solved more effectively than the complete (offline) problem. Also, variations of decision makers' preferences and decision chains were investigated. With increasing degree of "greediness", i.e. a stronger focus on minimizing the number of unserved customers, of course overall tour lengths increase, both in the single- as well as in the multiple vehicle scenario. The informative and sophisticated visualization approaches presented here can possibly foster the facilitation of decision processes along the DEMOA run and may offer perspectives for future dynamic, tool-based decision support systems.

REFERENCES

- [1] Mohamed Barkaoui and Michel Gendreau. 2013. An Adaptive Evolutionary Approach for Real-time Vehicle Routing and Dispatching. *Comput. Oper. Res.* 40, 7 (2013), 1766–1776. <https://doi.org/10.1016/j.cor.2013.01.022>
- [2] J.-F. Berube, M. Gendreau, and J.-Y. Potvin. 2009. An exact [epsilon]-constraint method for bi-objective combinatorial optimization problems: Application to the Traveling Salesman Problem with Profits. *European Journal of Operational Research* 194, 1 (2009), 39–50.
- [3] Jakob Bossek. 2020. GitHub repository with R implementation. <https://github.com/jakobbossek/dynvrp/>.
- [4] J. Bossek, C. Grimme, S. Meisel, G. Rudolph, and H. Trautmann. 2018. Local Search Effects in Bi-objective Orienteering. In *Proc. of the Genetic and Evolutionary Computation Conference (GECCO '18)*. ACM, New York, NY, USA, 585–592. <https://doi.org/10.1145/3205455.3205548>
- [5] Jakob Bossek, Christian Grimme, Stephan Meisel, Günter Rudolph, and Heike Trautmann. 2019. Bi-Objective Orienteering: Towards a Dynamic Multi-Objective Evolutionary Algorithm. In *Proceedings of the 10th International Conference on Evolutionary Multi-Criterion Optimization (EMO)*. East Lansing, Michigan, USA.
- [6] J. Bossek, C. Grimme, S. Meisel, G. Rudolph, and H. Trautmann. 2019. Bi-Objective Orienteering: Towards a Dynamic Multi-Objective Evolutionary Algorithm. In *Evolutionary Multi-Criterion Optimization (Lecture Notes in Computer Science (LNCS))*, Vol. 11411. Springer, 1–12.
- [7] Kris Braekers, Katrien Ramaekers, and Inneke Van Nieuwenhuysse. 2016. The vehicle routing problem: State of the art classification and review. *Computers & Industrial Engineering* 99 (2016), 300–313. <https://doi.org/10.1016/j.cie.2015.12.007>
- [8] Carlos A Coello Coello, Gary B Lamont, and David A Van Veldhuizen. 2006. *Evolutionary Algorithms for Solving Multi-Objective Problems (Genetic and Evolutionary Computation)*. Springer-Verlag New York, Inc., Secaucus, NJ, USA.
- [9] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan. 2002. A Fast and Elitist Multi-objective Genetic Algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation* 6, 2 (2002), 182–197.
- [10] Mauro Dell'Amico, Francesco Maffioli, and Peter Vaerbrand. 1995. On prize-collecting tours and the asymmetric travelling salesman problem. *International Transactions in Operational Research* 2, 3 (1995), 297–308.
- [11] Dominique Feillet, Pierre Dejax, and Michel Gendreau. 2005. Traveling Salesman Problems with Profits. *Transportation Science* 39, 2 (2005), 188–205.
- [12] C. Filippi and E. Stevanato. 2013. Approximation schemes for bi-objective combinatorial optimization and their application to the TSP with profits. *Computers & Operations Research* 40, 10 (2013), 2418–2428.
- [13] Truls Flatberg, Geir Hasle, Oddvar Kloster, Eivind J. Nilssen, and Atle Riise. 2007. *Dynamic And Stochastic Vehicle Routing In Practice*. Springer US, Boston, MA, 41–63. https://doi.org/10.1007/978-0-387-71722-7_3
- [14] Michel Gendreau, Gilbert Laporte, and Frederic Semet. 1998. A branch-and-cut algorithm for the undirected selective traveling salesman problem. *Networks* 32, 4 (1998), 263–273.
- [15] Bruce L. Golden, Larry Levy, and Rakesh Vohra. 1987. The orienteering problem. *Naval Research Logistics* 34, 3 (1987), 307–318.
- [16] Kannan Govindana, Mohammad Fattahib, and Esmaei Keyvanshokooch. 2017. Supply chain network design under uncertainty: A comprehensive review and future research directions. *European Journal of Operational Research* 263, 1 (2017), 108–141.
- [17] C. Grimme, S. Meisel, H. Trautmann, G. Rudolph, and M. Wölcck. 2015. Multi-Objective Analysis of Approaches to Dynamic Routing of a Vehicle. In *ECIS 2015 Completed Research Papers. Paper 62*. AIS Electronic Library. <https://doi.org/10.18151/7217333>
- [18] Lianxi Hong. 2012. An Improved LNS Algorithm for Real-time Vehicle Routing Problem with Time Windows. *Comput. Oper. Res.* 39, 2 (2012), 151–163. <https://doi.org/10.1016/j.cor.2011.03.006>
- [19] Roy Jonker and Ton Volgenant. 1983. Transforming Asymmetric into Symmetric Traveling Salesman Problems. *Oper. Res. Lett.* 2, 4 (November 1983), 161–163.
- [20] N. Jozefowicz, F. Glover, and M. Laguna. 2008. Multi-objective Meta-heuristics for the Traveling Salesman Problem with profits. *Journal of Mathematical Modelling and Algorithms* 7, 2 (2008), 177–195.
- [21] He-Yau Kang and Amy HI Lee. 2018. An Enhanced Approach for the Multiple Vehicle Routing Problem with Heterogeneous Vehicles and a Soft Time Window. *Symmetry* 10, 11 (2018), 650.
- [22] C. P. Keller and M. Goodchild. 1988. The multiobjective vending problem: A generalization of the traveling salesman problem. *Environ. Planning B: Planning Design* 15, 2 (1988), 447–460.
- [23] Mostepha R. Khouadjia, Briseida Sarasola, Enrique Alba, Laetitia Jourdan, and El-Ghazali Talbi. 2012. A Comparative Study Between Dynamic Adapted PSO and VNS for the Vehicle Routing Problem with Dynamic Requests. *Appl. Soft Comput.* 12, 4 (2012), 1426–1439. <https://doi.org/10.1016/j.asoc.2011.10.023>
- [24] Gilbert Laporte and Silvano Martello. 1990. The selective travelling salesman problem. *Discrete Applied Mathematics* 26, 2-3 (1990), 193–207.
- [25] Sandro Lorini, Jean-Yves Potvin, and Nicolas Zufferey. 2011. Online Vehicle Routing and Scheduling with Dynamic Travel Times. *Comput. Oper. Res.* 38, 7 (2011), 1086–1090. <https://doi.org/10.1016/j.cor.2010.10.019>
- [26] S. Meisel. 2011. *Anticipatory Optimization for Dynamic Decision Making*. Operations Research/Computer Science Interfaces Series, Vol. 51. Springer New York.
- [27] S. Meisel, C. Grimme, J. Bossek, M. Wölcck, G. Rudolph, and H. Trautmann. 2015. Evaluation of a Multi-Objective EA on Benchmark Instances for Dynamic Routing of a Vehicle. In *Proc. of the Genetic and Evolutionary Computation Conference (GECCO '15)*. ACM, New York, NY, USA, 425–432. <https://doi.org/10.1145/2739480.2754705>
- [28] J. Mula, R. Poler, J.P. Garcia Sabater, and F.C. Lario. 2006. Models for production planning under uncertainty: A review. *International Journal of Production Economics* 103, 1 (2006), 271–285.
- [29] Y. Nagata and S. Kobayashi. 2013. A Powerful Genetic Algorithm Using Edge Assembly Crossover for the Traveling Salesman Problem. *INFORMS Journal on Computing* 25, 2 (2013), 346–363.
- [30] Beatrice Ombuki, Brian J. Ross, and Franklin Hanshar. 2006. Multi-Objective Genetic Algorithms for Vehicle Routing Problem with Time Windows. *Applied Intelligence* 24, 1 (Feb 2006), 17–30. <https://doi.org/10.1007/s10489-006-6926-z>
- [31] V. Pillac, M. Gendreau, C. Guéret, and A. L. Medaglia. 2013. A review of dynamic vehicle routing problems. *European Journal of Operational Research* 225, 1 (2013), 1–11.
- [32] M. L. Pinedo. 2012. *Scheduling: Theory, Algorithms, and Systems* (4th ed.). Springer, New York.
- [33] C. Raquel and X. Yao. 2013. Dynamic Multi-objective Optimization: A Survey of the State-of-the-Art. In *Evolutionary Computation for Dynamic Optimization Problems*, S. Yang and X. Yao (Eds.). Springer, Berlin Heidelberg, 85–106.
- [34] Verena Schmid, Karl F Doerner, and Gilbert Laporte. 2013. Rich routing problems arising in supply chain management. *European Journal of Operational Research* 224, 3 (2013), 435–448.
- [35] Kay Chen Tan, Yoong Han Chew, and LH Lee. 2006. A hybrid multiobjective evolutionary algorithm for solving vehicle routing problem with time windows. *Computational Optimization and Applications* 34, 1 (2006), 115.
- [36] Pieter Vansteenwegen, Wouter Souffriau, and Dirk Van Oudheusden. 2011. The orienteering problem: A survey. *European Journal of Operational Research* 209, 1 (2011), 1–10.
- [37] Shijin Wang, Xiaodong Wang, Xin Liu, and Jianbo Yu. 2018. A Bi-Objective Vehicle-Routing Problem with Soft Time Windows and Multiple Depots to Minimize the Total Energy Consumption and Customer Dissatisfaction. *Sustainability* 10, 11 (2018), 4257.
- [38] Min Wen, Jean-François Cordeau, Gilbert Laporte, and Jesper Larsen. 2010. The Dynamic Multi-period Vehicle Routing Problem. *Comput. Oper. Res.* 37, 9 (2010), 1615–1623. <https://doi.org/10.1016/j.cor.2009.12.002>
- [39] Eckart Zitzler, Kalyanmoy Deb, and Lothar Thiele. 2000. Comparison of Multi-objective Evolutionary Algorithms: Empirical Results. *Evol. Comput.* 8, 2 (2000), 173–195. <https://doi.org/10.1162/106365600568202>