# An Adaptive Iterated Local Search for the Mixed Capacitated General Routing Problem

Mauro Dell'Amico(1), José Carlos Díaz Díaz(1), Geir Hasle(2), Manuel Iori(1)

(1) Department of Sciences and Methods for Engineering,

University of Modena and Reggio Emilia,

Via Amendola 2, 42122 Reggio Emilia, Italy

{mauro.dellamico; jose.diaz; manuel.iori}@unimore.it

(2) Department of Applied Mathematics,

SINTEF ICT

P.O. Box 124 Blindern, NO-0314 Oslo, Norway

geir.hasle@sintef.no

**Abstract**

We study the Mixed Capacitated General Routing Problem (MCGRP) in which a fleet of capacitated vehicles has to serve a set of requests by traversing a mixed weighted graph. The requests may be located on nodes, edges, and arcs. The problem has theoretical interest because it is a generalization of the Capacitated Vehicle Routing Problem (CVRP), the Capacitated Arc Routing Problem (CARP), and the General Routing Problem (GRP). It is also of great practical interest since it is often a more accurate model for real world cases than its widely studied specializations, particularly for so-called street routing applications. Examples are urban-waste collection, snow removal, and newspaper delivery. We propose a new Iterated Local Search metaheuristic for the problem that also includes vital mechanisms from Adaptive Large Neighborhood Search combined with further intensification through local search. The method utilizes selected, tailored, and novel local search and large neighborhood search operators, as well as a new local search strategy. Computational experiments show that the proposed metaheuristic

is highly effective on five published benchmarks for the MCGRP. The metaheuristic yields excellent results also on seven standard CARP datasets, and good results on four well-known CVRP benchmarks including improvement of the best known upper bound for one instance.

KEYWORDS: Vehicle Routing; Arc Routing; Mixed Capacitated General Routing Problem; Node, Edge, and Arc Routing Problem; Metaheuristics

## Introduction

Two of the most important optimization problems in freight transportation and logistics are the *Capacitated Vehicle Routing Problem* (CVRP) and the *Capacitated Arc Routing Problem* (CARP).

In the CVRP, a set of customers with known demands must be served by a fleet of identical capacitated vehicles stationed at a central depot. Requests with given demand size are located on the vertices of a graph, and the aim is to route the vehicles along the graph to satisfy all requests with minimum routing cost, obeying vehicle capacity. The graph may be either directed or undirected, and the costs are assigned to links (edges/arcs). The problem has been widely studied (especially in its undirected version) because of the large number of real-world applications it models, including distribution of gasoline, beverage, and food, and collection of solid waste. We refer the interested reader to the books by Golden et al. (2008), and Toth and Vigo (2014).

In the CARP we are still given a weighted undirected/directed graph, but in this case, requests of given size are located on a subset of links. A fleet of identical vehicles, all based at a central depot and having a fixed capacity, is available for serving the requests. The problem is to route vehicles along the graph in a capacity feasible way to serve all requests with minimum routing cost. Also the CARP has been widely studied, because it captures the essence of a wide range of real-world applications, including street sweeping, winter gritting, and snow clearing. In many cases the CARP is also a good model for postal delivery, newspaper delivery, and household waste collection. We refer the interested reader to the survey by Wøhlk (2008), the annotated bibliography by Corberán and Prins (2010), and the book edited by Corberán and Laporte (2014).

In the literature, there has been a tendency to categorize applications as being either a case of node routing, or a case of arc routing. There are, however, important real-world problems whose essential char-

acteristics cannot be captured neither by the CVRP nor by the CARP, as there is a mixture of requests located on nodes and requests located on links. Prins and Bouchenoua (2005) argue that in certain cases of urban-waste collection, most requests may be adequately modeled as located on streets, but some large point-based demands, located for example at schools or hospitals, are better modeled by the use of vertices. In subscription newspaper delivery, requests are basically located in points, but in dense urban or suburban residential areas a CARP model based on the street network may be a good abstraction. In general, qualified abstraction and problem reduction for a CVRP instance through aggregation, for instance with heuristics based on the road network, will create an instance with requests located on nodes, edges, and arcs, see, e.g., Hasle (2014).

To answer the challenges that are induced by these complex problems, several researchers have recently focused their attention on the so-called Mixed Capacitated General Routing Problem (MCGRP). In the MCGRP, requests are located on a subset of vertices, edges, and arcs of a given weighted mixed graph, and a fleet of identical capacitated vehicles based at a central depot is used to satisfy requests with minimum routing cost while adhering to capacity constraints. The MCGRP is able to model a continuum of mixed node and arc routing problems, and hence removes the sharp boundary that is often seen in the literature. As alluded to above, the problem has large practical interest, particularly for so-called street routing applications, see Bodin et al. (2003). The MCGRP is also of interest in combinatorial optimization, because it generalizes both the CVRP, the CARP and many other routing problems, as described in Section 2. Its resulting combinatorial complexity is, however, very high, and solving it to optimality is a difficult task even for moderate-size instances, see Bach et al. (2013) and Bosco et al. (2013).

In this paper, we propose a novel, hybrid metaheuristic, called Adaptive Iterated Local Search (AILS) for easy reference, to solve large-size instances of the MCGRP. It utilizes vital mechanisms from two classical trajectory-based metaheuristics: Iterated Local Search (ILS), see Lourenço et al. (2010), and Adaptive Large Neighborhood Search (ALNS), see Pisinger and Røpke (2007). We have combined these mechanisms in a new way, and introduced several new elements. Novel local search and large neighborhood search operators have been designed, and well-known operators have been tailored to the problem at hand. When ALNS finds solutions with a certain quality, further intensification is performed by local search (LS).

We have designed a new, aggressive LS strategy. In each iteration we explore a large neighborhood consisting of the union of moves from five operators, and investigate all moves with positive savings. The

effect is that we execute all independent moves before generating a new neighborhood.

Our experimental study shows that the resulting algorithm is highly effective. For five MCGRP benchmarks consisting of 409 instances in total, AILS produces 381 best known solutions, 108 of which are new. For three instances, AILS closes the gap for the first time. This brings the number of proven optimal solutions up to 234. AILS fails to find only ten of these. Notably, the AILS also achieves high quality computational results for heavily investigated special cases of the MCGRP, viz. four standard benchmarks for the CVRP, and seven standard benchmarks for the CARP.

The remainder of the paper is organized as follows. In Section 1 we formally describe the MCGRP. In Section 2 we give a survey of the most relevant results in the related literature. In Section 3 we propose our AILS metaheuristic for the MCGRP, and describe the key elements that make it computationally effective. In Section 4 we configure and evaluate the algorithm by means of extensive computational tests, and in Section 5 we draw conclusions.

# 1 Problem Description

The MCGRP is defined on a weighted mixed graph $G = (N, E, A)$, where $N = \{1, 2, \ldots, n\}$ is the set of nodes, $E$ the set of edges and $A$ the set of arcs. Let $c_{ij}$ denote the non-negative *traversal cost* associated with any link $(i, j) \in E \cup A$. The traversal cost, also known as deadheading cost, denotes the cost for traversing the link without servicing it. The traversal cost is 0 for all nodes.

Three subsets $N_r \subseteq N$, $E_r \subseteq E$ and $A_r \subseteq A$ define the *requests*, or *tasks*, i.e., the subsets of, respectively, nodes, edges, and arcs that have demand and must be serviced. Each request has a non-negative *service cost*, $s_i$ for $i \in N_r$ and $s_{ij}$ for $(i, j) \in E_r \cup A_r$, and a non-negative *demand*, $q_i$ for $i \in N_r$ and $q_{ij}$ for $(i, j) \in E_r \cup A_r$. Let $\tau = |N_r| + |E_r| + |A_r|$ be the total number of requests.

A fleet of identical vehicles, all having capacity $Q$, is used to service the requests. The fleet is located in a special node, called the *depot*. Each vehicle performs at most one *route*, that is, it starts from the depot, services a number of requests, and then returns to the depot. Deadheading via non-required links is usually necessary to reach the required ones. A route is *feasible* if the sum of serviced demands does not exceed the vehicle capacity.

The aim of the MCGRP is to define a set $x$ of feasible routes for which every task $t \in N_r \cup E_r \cup A_r$ is

serviced exactly once, and the total cost $z(x)$ is a minimum. Note that the total service cost is constant over all feasible solutions, hence it is sufficient to minimize the sum of traversal costs.

An example of a MCGRP instance is given in Figure 1. Each node is depicted by a circle, drawn with a solid line if the node is a request, by a dashed line otherwise. Node 7 is the depot and is depicted by a square. Similarly, required links are drawn with a solid line, non-required links with a dashed line. The traversal costs are indicated. In this particular instance the traversal costs are symmetric, hence we give only one cost for each pair of arcs connecting the same two vertices. The vehicle capacity is 1437.

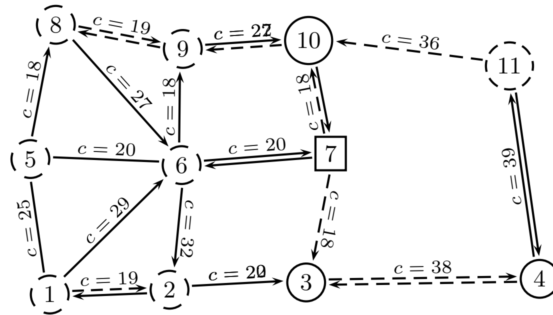


Figure 1: A MCGRP example: Instance `CBMix23`.

A solution for the instance in Figure 1 is illustrated in Figure 2. It consists of four routes, each presented in a separate sub-figure for the sake of clarity. In each sub-figure, the links with solid lines are serviced by the route, links with dashed lines indicate deadheading. We also indicate the demand for each serviced task. Note that Route 1 starts from the depot and visits, in sequence, nodes 10, 6, 5, and 8, then visits 6 again and returns to the depot. It services five tasks, namely 10, (10,6), (5,8), (8,6) and (6,7), with total demand (denoted *load* for short in the figure) equal to 1024. Note also that Route 2 travels three times through node 6, and Route 3 is forced to travel three times between nodes 4 and 11 to perform the two requests (4,11) and (11,4). The resulting solution value is 780, and its optimality was proven by Bosco et al. (2013).

## 2 Prior Work in the Area

The MCGRP has also been called the Node, Edge, and Arc Routing Problem (NEARP) in the literature. As far as we know, Pandi and Muralidharan (1995) is the first paper that investigates the MCGRP. The authors studied a generalization with a heterogeneous fixed fleet, and a maximum route duration constraint. The

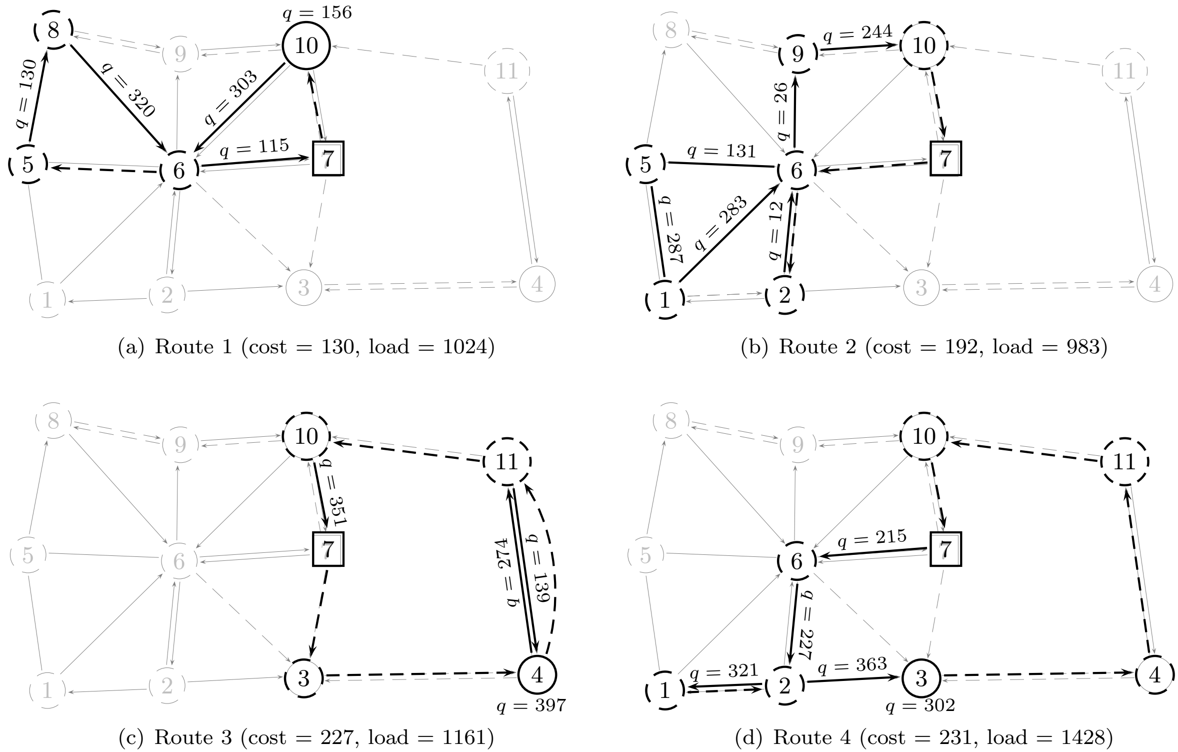


(a) Route 1 (cost = 130, load = 1024)

(b) Route 2 (cost = 192, load = 983)

(c) Route 3 (cost = 227, load = 1161)

(d) Route 4 (cost = 231, load = 1428)

Figure 2: A four-route (optimal) solution for `CBMix23`.

resulting problem, denoted the *Capacitated General Routing Problem* (CGRP), was solved with a route-first-cluster-second heuristic. The algorithm was tested on random test instances inspired from curb-side waste collection in residential areas, and on random instances from the Capacitated Chinese Postman Problem literature.

A few years later, Gutiérrez et al. (2002) studied the homogeneous fixed fleet version of the CGRP, and called it the *Capacitated General Routing Problem on Mixed Graphs* (CGRP-m). In other words, the CGRP-m is a MCGRP with a limited number of vehicles. They proposed an $O(n^3)$ heuristic and compared it with the heuristic by Pandi and Muralidharan (1995), obtaining favorable computational results on a benchmark of 28 instances with the number of vehicles between 2 and 4, and the number of required tasks between 6 and 21.

Prins and Bouchenoua (2005) introduced the Node, Edge, and Arc Routing Problem (NEARP) name for the problem and solved it by means of a memetic algorithm in which a population of solutions is evolved

through a genetic process, and each new solution is post-optimized using five local search operators. The resulting algorithm was tested on benchmark instances from the CVRP and CARP literature, and on a new set of MCGRP instances denoted the CBMix benchmark. Kokubugata et al. (2007) developed a simulated annealing algorithm that makes use of three local search operators. They tested their algorithm on the CBMix instances and provided several new best known solutions. Recently, upper bounding procedures were discussed by Hasle et al. (2012), who obtained interesting computational results by running the commercial VRP solver Spider.

The first lower bounding procedure for the MCGRP was proposed by Bach et al. (2013). It was obtained by adapting a procedure originally developed for the CARP by Wøhlk (2006), based on the solution of a matching problem. The lower bound was tested on the CBMix benchmark, and on two new sets of MCGRP instances proposed by the authors: the BHW benchmark based on well known instances from the CARP literature, and the DI-NEARP benchmark taken from real-world newspaper distribution cases in Norway.

Bosco et al. (2013) proposed the first integer programming formulation for the MCGRP, using three-index variables for nodes and two-index variables for edges and arcs. They extended some valid inequalities originally introduced for the CARP to the MCGRP, and embedded them into a branch-and-cut algorithm. This algorithm was tested on two new benchmarks called mggdb and mgval, each consisting of six sets, and totaling 342 instances. The mggdb benchmark, with 138 instances, was derived from the gdb undirected CARP test set. The mval mixed CARP dataset is the origin of the mgval benchmark that has 204 instances. The authors considered only the 264 instances involving at most seven vehicles in their experiments. For these, the method provided 154 proven optimal solutions. The authors also tested their algorithm on the four smallest-size CBMix instances, providing two optimal solutions.

Bach et al. (2014) proposed a branch-and-cut-and-price method for the MCGRP and investigated its performance on mggdb, one subset of mgval, CBMix, and BHW. At the time, they proved optimality for 31 new mggdb instances. On CBMix, they improved the best known upper bound for two instances and the best known lower bound in 21 cases. Optimality was proven for two new BHW instances.

With the work reported in Bosco et al. (2013) as basis, Irnich et al. (2015) presented a new two-index mathematical model for the MCGRP, and a two-phase branch-and-cut algorithm that utilizes an aggregate formulation to develop an effective lower bounding procedure. They provided numerical results for three of the six parameter subsets of the mggdb and mgval benchmarks. 124 of the 171 instances investigated are

solved to optimality.

A matheuristic for the MCGRP was proposed by Bosco et al. (2014). The authors provided numerical results for the `mggdb`, `mgval`, and `CBMix` benchmarks. A bi-criteria extension of the MCGRP, where the second criterion is route balance, was studied for the first time by Mandal et al. (2015). A memetic algorithm is proposed, and numerical results are reported for the `CBMix` benchmark.

In Section 4, we provide a comprehensive survey of numerical results for the `CBMix`, `BHW`, `DI-NEARP`, `mggdb`, and `mgval` MCGRP benchmarks. We also refer to the Transportation Optimization Portal (TOP) web site SINTEF that attempts to maintain an updated survey of the best known numerical results for all MCGRP benchmarks.

As mentioned in the introduction, the MCGRP generalizes a large number of optimization problems arising in transportation and logistics. A problem classification is presented in Figure 3. The classification is incomplete, because of the large number of variants addressed in the scientific literature. As depicted by the figure, the MCGRP directly generalizes:

- the *Capacitated Vehicle Routing Problem* (CVRP): $N_r = N$, $E_r = \emptyset$ and $A = \emptyset$;
- the *Capacitated Arc Routing Problem* (CARP): $N_r = \emptyset$ and $A = \emptyset$; and
- the *General Routing Problem* (GRP): one vehicle, $Q = +\infty$ and $A = \emptyset$.

The CVRP is one of the most widely studied problems in the combinatorial optimization literature. Recently, exact algorithms based on branch-and-cut-and-price techniques have been proposed by Baldacci et al. (2008) and Baldacci and Mingozzi (2009). Good-quality heuristic solutions have been obtained in the last years by, among others, Gröer et al. (2011) with local search and integer programming embedded into a parallel algorithm, and Vidal et al. (2012) with a hybrid genetic algorithm that can also take into consideration multiple depots or multiple periods. We also mention that there are works aimed at solving the CVRP on an asymmetric cost matrix. The literature on the problem, known as the *Asymmetric CVRP* (ACVRP), is described, e.g., in Toth and Vigo (2002). Note that, since the CVRP is strongly $\mathcal{NP}$-hard, so is the MCGRP.

The CARP has also been widely studied in the literature. Recently, branch-and-cut-and-price algorithms have been presented by Bartolini et al. (2011) and Martinelli et al. (2011). Good-quality heuristic solutions have been obtained via Ant Colony Optimization by Santos et al. (2010). Despite the use of the term "arc"

in its name, the CARP has been originally defined on an undirected graph. Works aimed at solving arc routing problems on directed graphs, and on more general mixed graphs, are described, e.g., in Corberán et al. (2006).

The GRP was introduced by Orloff (1974), to model the problem of collecting requests on nodes and edges of an undirected graph with a single uncapacitated vehicle. A good cutting plane algorithm was proposed by Corberán et al. (2001). Similar to the CVRP and the CARP, also the GRP has been extended to directed and mixed graphs, see, e.g., Corberán et al. (2005). Notably, the GRP generalizes other combinatorial optimization problems, namely:

- the *Rural Postman Problem* (RPP): one uncapacitated vehicle, $A = \emptyset$, $N_r = \emptyset$;
- the *Chinese Postman Problem* (CPP): one uncapacitated vehicle, $A = \emptyset$, $N_r = \emptyset$, $E_r = E$;
- the *Steiner Graphical Travelling Salesman Problem* (SGTSP): one uncapacitated vehicle, $A = \emptyset$, $E_r = \emptyset$;
- the *Graphical Travelling Salesman Problem* (GTSP): one uncapacitated vehicle, $A = \emptyset$, $E_r = \emptyset$, $N_r = N$; and
- the *Travelling Salesman Problem* (TSP): one uncapacitated vehicle, $A = \emptyset$, $E_r = \emptyset$, $N_r = N$.

For the literature on the RPP, CPP, SGTSP, GTSP, TSP and their extensions to directed and mixed graphs, we refer the reader to Corberán et al. (2001), Gutin and Punnen (2002), Corberán et al. (2007), Corberán and Laporte (2014), and references therein.

## 3 Adaptive Iterative Local Search

In this section we discuss the novel hybrid metaheuristic that we propose to search for high-quality solutions of MCGRP instances of realistic size in reasonable time. For easy reference, we call the algorithm *Adaptive Iterative Local Search* (AILS). Parts of AILS uses pseudo-random numbers, but we emphasize that it is a deterministic algorithm that will produce the same path in the search space given the same random seed.

First, we give a description of the overall design of AILS. Main goals are to ensure a good balance between intensification and diversification, and to avoid non-productive search efforts. To these ends, we use

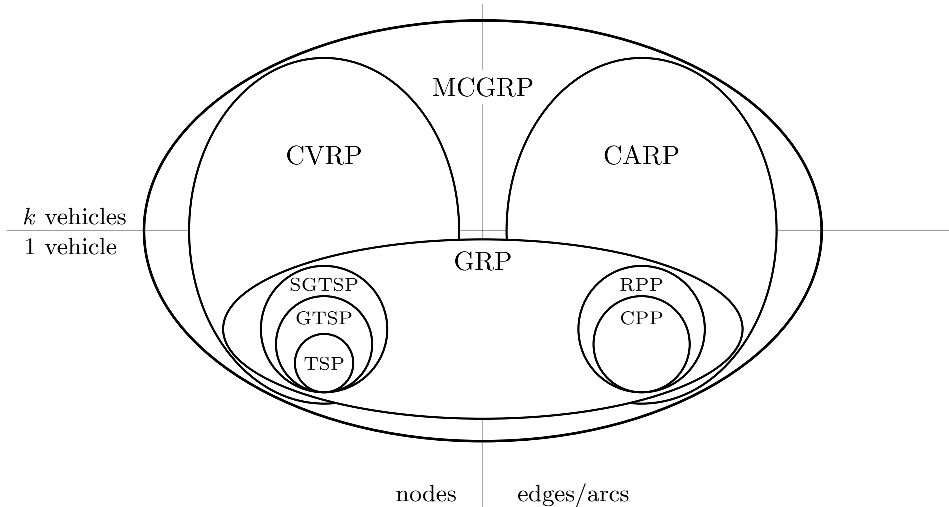


Figure 3: A graphical representation of a problem classification.

the idea of *Iterated Local Search* (ILS) (see, e.g., Lourenço et al. (2010)) that mainly consists of improving a solution through a trajectory based intensification algorithm, and diversification through a perturbation method when intensification stagnates.

AILS combines intensification mechanisms that on their own have proven to be highly effective for a variety of discrete optimization problems, including many variants of the VRP, namely *Adaptive Large Neighborhood Search* (ALNS) proposed by Pisinger and Røpke (2007) and partially based on ideas from Shaw (1997), and deep intensification through *Local Search* (LS). Intensification is performed in *stages*, each consisting of a certain number of iterations.

In one iteration, ALNS destroys and repairs the current solution. The pair of *destructor* and *constructor* operators is probabilistically selected among alternative operators. A reinforcement learning technique is used to update the selection probabilities. Further details on our version of the ALNS are given in Section 3.1. It contains several innovations and non-standard mechanisms. If the solution resulting from a destroy/repair operation has good quality, the search is further intensified through local search with five neighborhood operators, and a new, aggressive move selection strategy. Details are given in Section 3.2. The main diversification mechanism of AILS is a major disruption – a *kick* – applied when a certain search effort has been made without acceptance of a new solution.

Pseudocode for the AILS algorithm, the structure of which is quite simple, is given in Algorithms 1 and

2, with explanation below. Parameters, functions, and procedures used by AILS are briefly described, with final values listed in Table 1. We refer to Section 4.1 for further description of the configuration process.

**Algorithm 1** Adaptive Iterative Local Search

```
 1: function AILS(Instance)
 2:     comment: Initialize global variables
 3:     IterPerStage := N_ITER_PER_STAGE
 4:     k_max := K_MAX(Instance)
 5:     KickCountdown :=ITER_BEFORE_KICK(Instance)
 6:     INITIALIZE_ROULETTE_PROBABILITIES( )
 7:     comment: Construct first solution and take to deep local optimum
 8:     x_init :=CONSTRUCT_INITIAL_SOLUTION(Instance)
 9:     x_incumbent :=LS_FULL(x_init)
10:     x_LocalIncumbent := x_incumbent
11:     comment: Main body: iterative phase
12:     repeat
13:         x_current := x_BestThisStage := x_LocalIncumbent
14:         comment: Execute a batch of intensifying iterations – a stage
15:         IterationCounter := 0;
16:         while IterationCounter < IterPerStage do
17:             IterationCounter := IterationCounter + 1
18:             NewStage := COMBINED_ALNS_AND_LS( )
19:             if NewStage then
20:                 IterPerStage := N_ITER_PER_STAGE - 1
21:                 KickCountdown := iter_BeforeKick
22:                 IterationCounter = IterPerStage
23:             end if
24:         end while
25:         comment: Increase number of iterations
26:         IterPerStage := IterPerStage + 1
27:     until TIMEOUT( )
28:     return x_incumbent
29: end function
```

**Algorithm 2** Combined ALNS and LS

```
1:  function COMBINED_ALNS_AND_LS( )
2:      comment: Returns TRUE if a new stage is required, FALSE otherwise
3:      comment: Check for stagnation
4:      if KickCountdown > 0 then
5:          k :=RANDOM(1, k_max)
6:          x_current :=ROULETTE_DESTROY_AND_REPAIR(k)
7:          if z(x_current) = z(x_BestThisStage) return FALSE
8:          comment: Intensify with LS_FULL, LS1, or LS2 based on random choice and instance size
9:          if RANDOM(0,1)< p_LS_Full then
10:             LS_FULL(x_current)
11:         else
12:             if τ < 200 then LS1(x_current) else LS2(x_current)
13:         end if
14:         KickCountdown := KickCountdown − 1
15:         if z(x_current) < z(x_BestThisStage) then
16:             x_BestThisStage := x_current
17:             comment: Give higher probability to selected Destructor/Constructor pair
18:             UPDATE_ROULETTE_PROBABILITIES( )
19:             UPDATE_INCUMBENTS(x_current)
20:             return TRUE
21:         end if
22:         return FALSE
23:     else
24:         comment: Nothing has happened for a while, make a major, random destroy and repair
25:         k :=RANDOM(τ/2, τ)
26:         x_LocalIncumbent :=RANDOM_DESTROY_AND_REPAIR(k)
27:         x_current := x_LocalIncumbent :=LS_FULL(x_LocalIncumbent)
28:         UPDATE_INCUMBENTS(x_current)
29:         return TRUE
30:     end if
31: end function
```

AILS finds a first feasible solution with simple, fast heuristics. The initial solution is taken to a deep local optimum through an aggressive local search procedure. A main **repeat-until** loop performs Iterated Local Search until timeout. An *intensification stage*, implemented by the inner **for** loop, performs a given number of iterations. Each iteration consists of a destroy and repair cycle, possibly followed by further intensification with local search. A local incumbent for the stage is maintained. When a stage is finished, a new one is started from the local incumbent of the previous one. A kick is performed whenever *stagnation* occurs, i.e., no new solution has been accepted for a certain number of iterations that depends on the computational complexity of the instance. We refer to Table 1 and Section 4.1 for details. The kick utilizes the random destructor and random constructor operators from the ALNS, see Section 3.1. The resulting solution is taken

to a deep local optimum before a new intensification stage is started.

Before AILS starts, the minimum traversal (deadheading) costs $c'_{ij}$ connecting any pair of vertices $i$ and $j$ are computed with the standard Dijkstra algorithm. Recall that $\tau = |N_r| + |E_r| + |A_r|$ is the total number of requests. First, the AILS computes an initial solution $x_{init}$ with the function CONSTRUCT_INITIAL_SOLUTION (Algorithm 1, line 8). Experiments showed that the quality of the initial solution had no significant effect on the final result after a reasonable computing time. Hence, we selected a computationally cheap construction procedure, the *Augment-Merge* heuristic proposed by Golden and Wong (1981). For instances with a given upper bound on the number of vehicles (as the `mggdb` and `mgval` benchmarks) we used a modified version of the Augment-Merge procedure that continues to merge routes, also with negative savings, if the number of routes in the current solution is larger than the upper bound. If this simple procedure fails in finding a feasible solution, we solve a *bin packing problem* where the task demands are objects that have to be packed into bins of capacity equal to the vehicle capacity. We used the powerful variable neighborhood search procedure developed in Dell'Amico et al. (2012), modified so that it stops as soon as the number of bins is not larger than the upper bound. The tasks of each bin are than sequenced with a simple nearest insertion procedure. The initial solution is taken to a deep local optimum, the first version of the incumbent $x_{incumbent}$, by the most powerful LS, called LS_FULL.

After initialization, the AILS enters a main loop that is executed until timeout. Within this loop, combined ALNS and LS is executed for intensification. As acceptance criterion for a new solution, we use simple improvement. We note that the number of iterations per stage is initialized to the parameter N_ITER_PER_STAGE, then increased by one for each stage, as initial experiments indicated that a progressive number of iterations was needed to reinforce the intensification.

Table 1: Final parameter setting

| Parameter/function | Value |
|---|---|
| N_ITER_PER_STAGE | 10 |
| $k_{MAX}(\mathit{Instance})$ | $min(\tau - 2, 50)$ |
| $\beta$ | 0.75 |
| $\gamma$ | 0.1 |
| $p_{LS_{Full}}$ | 0.15 |
| ITER_BEFORE_KICK(*Instance*) | 20.000 * $\max\{1, 20.000/(\tau^2 + \lvert A\rvert + 2\lvert E\rvert + n^2/5)\}$ |

Below follows a description of the functions and procedures used in Algorithms 1 and 2.

TIMEOUT: A function that returns `TRUE` when the given CPU time limit is reached, `FALSE` otherwise.

INITIALIZE_ROULETTE_PROBABILITIES: sets all entries in the scores table $\pi$ for roulette wheel selection to 1. This procedure is invoked in the initialization, line 6.

The following procedures and functions are used in in COMBINED_ALNS_AND_LS.

RANDOM_DESTROY_AND_REPAIR: line 26, is called to make a *kick* after a certain amount of search effort has been performed without acceptance of a new current solution. It calls the Random-Destructor and then the Random-Constructor for a number of tasks randomly drawn from the interval $[\tau/2, \tau)$.

ROULETTE_DESTROY_AND_REPAIR calls the normal, roulette wheel based selection of Destructor and Constructor pair, and the execution of these operators with a randomly drawn $k$ value, see line 6.

UPDATE_ROULETTE_PROBABILITIES: line 18, increases the probability of selecting a successful Destructor/Constructor pair.

UPDATE_INCUMBENTS: line 19, checks whether the current solution is better than the best solution found, and in case, updates the corresponding variable. Similarly, there is a test whether the current solution improves the global incumbent. The function returns `TRUE` if any of the variables were updated, `FALSE` otherwise.

## 3.1 The Adaptive Large Neighborhood Search Component

The Adaptive Large Neighborhood Search mechanism in AILS is a modified and simplified version of the one proposed by Pisinger and Røpke (2007). To perform well, ALNS must utilize a varied repertoire of destructor and constructor operators, and a qualified mechanism for selecting the operators to employ in a given cycle. Our ALNS design introduces a novel *tree-destructor* that utilizes the graph structure of the instance at hand. Experiments (see Section 4.2) show that it is effective.

Self-adaptation in ALNS is typically achieved through a reinforcement learning mechanism that rewards operators that have been successful in past iterations. The reinforcement learning mechanism in AILS is based on operator pairs rather than on single operators. Any time the destroy and repair mechanism is invoked, a destructor/constructor operator pair is randomly drawn, using roulette wheel selection. These operators are then used to remove and re-insert a randomly drawn number of tasks. A scores matrix $\pi$

contains a measure for the effectiveness of each pair of Destructor and Constructor operators. The score element $\pi_{ij}$ is associated with the $i$-th Destructor and the $j$-th Constructor. AILS uses a very simple score update mechanism. The initial value for all elements is 1. The update procedure increments the value $\pi_{ij}$ by 1 unit, whenever the $i$-th Destructor and the $j-$th Constructor have led to a new current solution. The scores matrix $\pi$ is never reset in AILS, as focused experiments with different reset models never gave significant improvement.

In each iteration, a number $k$ is used as parameter to the randomly selected pair of operators. As is common in the literature, $k$ is randomly drawn from a uniform distribution over an interval $[k_{min}, k_{max}]$. Early experiments revealed that it is beneficial to have a finite probability of selecting very small $k$ values, so we set $k_{min}$ to 1. Common ALNS insight indicates that $k_{max}$ should be an increasing function of instance size, but with an upper limit to avoid excessive computational burden. Focused experiments confirmed that this model is effective, and we found the best upper limit value for $k_{max}$ to be 50. The final form of the function to determine $k_{max}$ is found in Table 1.

Another important simplification relative to the "canonical" ALNS is that the AILS design uses a simple improvement criterion for acceptance of new solutions, rather than a more complex criterion that is common in the literature.

AILS draws upon a repertoire of seven destructor operators, all parameterized by the number of tasks to remove:

1. **Random-Destructor**: $k$ random tasks are selected and removed from the solution;

2. **Task-Destructors**: these are our extensions of the analogous operators used for CVRP and CARP.

   2.a **Node-Destructor**: if $k \leq |N_r|$, $k$ random node tasks are removed from the solution, otherwise the **Random-Destructor** is used;

   2.b **Edge-Destructor**: if $k \leq |E_r|$, $k$ random edge tasks are removed from the solution, otherwise the **Random-Destructor** is used;

   2.c **Arc-Destructor**: if $k \leq |A_r|$, $k$ random arc tasks are removed from the solution, otherwise the **Random-Destructor** is used;

3. **Worst-Destructor**: we define the cost of removing a task $t$ from the current solution $x$ as $\Gamma(t, x) =$

$z(x) - z_{-t}(x)$, where $z_{-t}(x)$ is the cost of the solution without task $t$. The operator removes the $k$ tasks having the highest values of $\Gamma(t, x)$;

4. **Related-Destructor**: this operator was proposed by Shaw (1997, 1998). Its aim is to remove tasks that are somehow close to one another. For the MCGRP, extending the original idea, we define the contiguity of two tasks $r$ and $t$ as:

$$\rho(r,t) = \beta \frac{c'_{rt}}{\max_{su} c'_{su}} + \gamma \frac{|q(r) - q(t)|}{\max_s q(s)} + \delta(r,t), \tag{1}$$

   where $\beta = 0.75$, $\gamma = 0.1$ as recommended in the literature, $c'_{rt}$ is the minimum traversal cost between $r$ and $t$, $q(t)$ is the demand of task $t$, and $\delta(r,t)$ takes value 1 if $r$ and $t$ are in the same route in the current solution, 0 otherwise. We start by selecting a random task, then we add $k-1$ tasks in a greedy way by selecting the next task as the one with minimum distance from the last added task.

5. **Tree-Destructor**: this is a new operator which is particularly effective for MCGRP instances that contain all three task types. It utilizes the instance graph, as illustrated in Figure 1. First, it randomly selects a task as a root node, and then grows a tree in the instance graph $G$, from this root, by using a breadth-first strategy. The growth is halted as soon as $k$ tasks (of any kind) are encountered.

Three constructors are used in AILS, as extensions of operators from the literature. They re-insert $k$ removed tasks in the current solution, one at a time, according to a certain criterion. They iterate until all tasks have been re-inserted. The resulting solution is always feasible, although it may contain additional routes. When the instances have a fixed number of vehicles, as those in Bosco et al. (2013), we manage this particular case by adding a penalty for each route, so that minimizing the objective function also reduces the number of routes. The feasibility check is then modified by including a control on the number of vehicles used.

- **Random-Constructor**: Insert each task, one at a time, according to the order in which they have been removed from the solution by the Destructor, in a random position in the current set of routes. If no feasible position exists, create a new route with only this task;

- **Greedy-Constructor**: In each iteration, the task with the minimal best insertion cost is inserted in its best position;

- **Regret-Constructor**: Compute for each task $t$ its cheapest insertion cost, and its second cheapest insertion cost, and define its regret $r(t)$ as the difference between the two costs. Insert the task having maximum regret in its best position, and then re-iterate, by re-computing regrets, until all tasks have been re-inserted.

The Regret-Constructor has been used, among others, by Røpke and Pisinger (2006), to overcome the myopic behavior of greedy repair.

### 3.2 The Local Search Component

Local search in AILS is based on five operators from the node routing and arc routing literature that will be described in detail below. These operators have been extended to accommodate the MCGRP model. In total, 26 move subtypes have been implemented. However, we have designed a new (as far as we know), more aggressive neighborhood evaluation strategy, as follows. In each iteration, the union of neighborhoods resulting from the five operators applied to the current solution is fully explored. All moves with positive savings are considered, in the order of decreasing savings. All independent moves that led to feasible solutions are executed, before local search proceeds with the next neighborhood exploration from the new current solution.

As is seen in Algorithms 1 and 2, AILS performs intensification through local search in three situations:

- after construction of the initial solution,
- when a solution with sufficient quality has been produced by ALNS,
- after a kick.

The three situations call for different LS variants. After initial construction, and after a kick, the goal is to find a deep local optimum fast. Therefore, we utilize the most powerful local search variant called LS_Full that includes all move types. During the ALNS destroy and repair phase, we have seen through experiments that it becomes too expensive to use LS_Full all the time. Therefore, we designed two reduced local search variants: LS1 and LS2. The details of these are given below. LS1 is used for small instances ($\tau < 200$), for larger instances, we use LS2. However, with a small probability $p_{LS_{Full}}$, see Table 1, we invoke the full local search, regardless of instance size. These choices were made after extensive experiments on standard MCGRP, CARP, and CVRP benchmarks.

AILS utilizes the following set of local search operators:

- **Swap**: exchanges the position of two tasks (both intra and inter-route);
- **Or-opt**: relocates a segment of tasks, either within a route, or from one route to another. The length of the segment to be relocated is limited to $l$ tasks;
- **2-opt**: breaks a route cycle in two segments and re-connects the segments in the only possible way. We adapted to the MCGRP also the seven additional operator subtypes originally proposed by Santos et al. (2010) for the CARP. They result from breaking two different routes in one point each, and reconnecting the segments in all possible ways when reversals are also considered;
- **3-opt**: breaks a route cycle in three segments $A$, $B$ and $C$, and reconnects them in all possible ways, also allowing reversals. There are seven combinations: $AC\bar{B}$, $A\bar{C}B$, $A\bar{C}\bar{B}$, $\bar{A}CB$, $\bar{A}C\bar{B}$, $\bar{A}\bar{C}B$ and $\bar{A}\bar{C}\bar{B}$, where $\bar{X}$ denotes the reversal of $X$. There are six types of intra-route moves (case $A\bar{C}\bar{B}$ is equal to intra-route 2-opt), and seven types of inter-route moves;
- **Flip**: reverses the direction of all the edge tasks of a route.

The Flip operator was proposed for the MCGRP by Prins (2009).

Hence, AILS uses five operators with a total of 26 subtypes: 13 types of 3-opt, 8 types of 2-opt, 2 types of Or-opt, 2 Swap types, and Flip. LS_Full employs all these operators and subtypes. The segment length limit $l$ for Or-opt is 3, and for 3-opt, $|B| \leq 3$. The computationally cheaper LS1 consists of the following operators: Swap, 2-opt, 3-opt with $|B| \leq 1$, and Flip. This limits the search to 18 operator subtypes. LS2 is our cheapest local search variant. It consists of Or-opt with $l = 2$, Swap, and 2-opt, covering 12 operator subtypes.

# 4 Computational Experiments

We coded our algorithm in C++ and ran it on an Intel Xeon E5530 at 2.4 GHz with 23.5 GB of memory, under the Linux Ubuntu 12.04.1 LTS 64-bit operating system. In Section 4.1, we describe AILS parameter configuration, before we move on to a quantitative investigation of main AILS mechanisms in 4.2. The section finishes with a description of extensive computational experiments on standard benchmarks from the literature, for MCGRP, as well as the CARP and CVRP special cases.

### 4.1 Configuration of the Algorithm

AILS has six main parameters. Through a combination of analysis, insights from the literature, and extensive computational experiments on `CBMix`, `BHW`, and `DI-NEARP` instances (see Section 2), we determined a good setting. We refer to Table 1 and explanations in Section 3. Further details of the parameter tuning are not given here, except for an account of how we determined the final form of the iteration limit for the kick.

The ITER_BEFORE_KICK function determines the number of iterations without improvement before the kick is invoked. Initially, we performed a set of experiments with a large sample of MCGRP instances to determine the best constant value for ITER_BEFORE_KICK. We tried the following values: 5000, 10.000, 20.000, 30.000, 50.000, and 100.000. Results showed that a good overall choice that balances the potential for further intensification with the potential benefits of diversification is 20.000 iterations. Disabling the kick gave considerably worse average results. For very large instances and reasonable timeout values, the kick will never be invoked within 20.000 iterations, which seems reasonable as the AILS intensification mechanism will not have enough time to stagnate. However, we observed that for small instances, a larger number of iterations before the kick gave better results, indicating that with a 20.000 iteration limit, intensification was interrupted prematurely in many cases.

This observation led us to the conclusion that the iteration limit for the kick should be determined according to the computational complexity of the instance. We identified $\tau$, $|A|$, $|E|$ and $n$ as the main instance metrics of computational complexity. Using a sample of 17 MCGRP benchmark instances (see Section 4.2 for details), we determined a function of these metrics that has a similar behavior to the CPU time required to complete one iteration of AILS. In Figure 4 the solid line reports on the observed CPU time required to execute one iteration (right vertical axis), while the dashed line refers to our function $f_{approx} = \tau^2 + |A| + 2|E| + \frac{1}{5}n^2$ (left vertical axis). We observe that our function gives a reasonable estimate over the sample. We use the reciprocal of $f_{approx}$ in our final ITER_BEFORE_KICK function (see again Table 1): it is designed such that the minimum number of iterations before kick is 20.000, but the value will be higher for instances with low computational complexity.

Further computational experiments, where small variations of the six main AILS parameters were investigated, revealed only minor changes in performance. For a discussion of AILS robustness with respect to choice of random seed, we refer to Section 4.2.

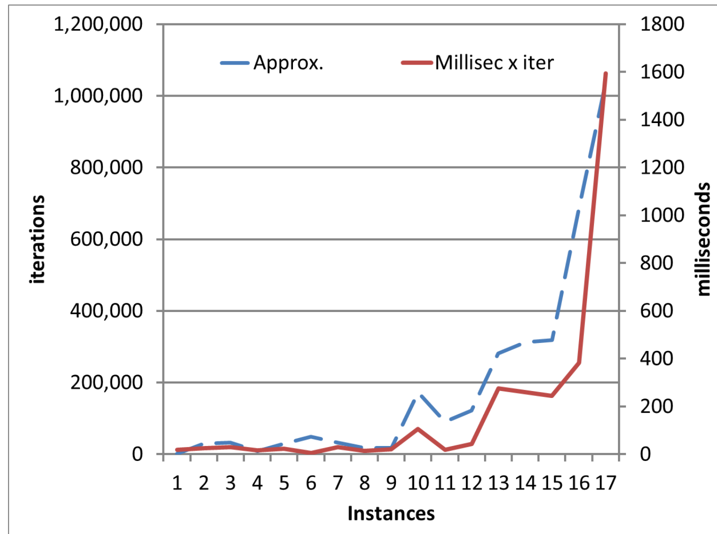


Figure 4: Comparison between observed time for one iteration and approximating function.

## 4.2 Investigation of Main AILS Mechanisms

The special combination of ALNS with conditional LS for further intensification is a main feature of AILS. Further, the Tree-Destructor is an innovation, a new operator in the ALNS component of AILS. This subsection reports from a quantitative assessment of the merit of these two major mechanisms.

For a comprehensive assessment, we selected a sample of 17 MCGRP instances from the `CBMix`, `BHW`, and `DI-NEARP` benchmarks. The sample contains instances of different size and structure, with $\tau$ varying from 91 to 833. In general, they are instances that seemed hard from early experiments. The names of the sample instances and their $\tau$ values are found in the result tables below. To have a larger sample, also investigating the robustness of AILS with respect to random seed, we ran all experiments with ten different seed values. Thus, the total sample size is 170. For all experiments, a CPU time limit of 3600 seconds was imposed.

### 4.2.1 Removing the Local Search Component

We disabled the LS component and ran experiments on the 17 sample MCGRP instances. To remove the unpredictable effect of the kick mechanism that would make comparison more difficult, we disabled it and ran two experiments:

- AILS without the kick (*Basic Configuration*)

- AILS without the kick and with LS disabled

The results are found in Table 2. The first two columns give the instance names and their $\tau$ values. The following five columns give the results from AILS without the kick, here called the "Basic Configuration". For each instance, the minimum, maximum, and average cost over the ten random seeds for the best solution found within the timeout are given. As usual, the constant sum of service costs is not included. The column marked "RSD%" gives the relative standard deviation for the ten cost values, in percent. The average CPU time (seconds) to find the solutions is given in the following column marked $sec_{inc}$. The next five columns, marked "No Local Search", give the corresponding results for the Basic Configuration with LS disabled. The rightmost column marked "$\Delta$%" gives the percentage difference of the average results relative to the Basic Configuration.

We make two important observations from these results. First, the Local Search Component contributes substantially to the performance of AILS. The average deterioration for the sample when the LS component is disabled is more than 5%. There is a consistent deterioration for all instances, and the maximum deterioration is almost 11%.

Second, the RSD values are low for all instances of the sample. For the Basic Configuration, the average RSD is 0.31%, with a maximum of 0.62%. Without the LS component, the average variation is slightly higher at 0.48% with a maximum of 0.94%, which is still low. This is a clear indication that the AILS is robust with respect to selection of random seed.

#### 4.2.2 Removing the Tree-Destructor

We investigated the merit of the novel Tree-Destructor operator in a similar way as for the LS component. Table 3, which has the same layout as Table 2, shows a comparison between the Basic Configuration as defined in Section 4.2.1 above, and the Basic Configuration with the Tree-Destructor disabled. Again, the results show a small variation over ten different seed values. Although the version without the Tree-Destructor only shows a small average deterioration relative to the Basic Configuration, a deterioration is observed for 16 out of 17 instances of the sample. Moreover, when we compare the best cost values found over the ten seeds (in bold), we observe that the Basic Configuration is slightly better.

Based on the results and observations presented we decided to use the Basic Configuration, but with the kick enabled, for our final computational experiments. For all experiments presented below, the timeout

Table 2: Assessment of the Local Search Component of AILS

| Instance | $\tau$ | Basic Configuration | | | | | No Local Search | | | | | $\Delta\%$ |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | | Min cost | Max cost | Avg. cost | RSD% | $sec_{inc}$ | Min cost | Max cost | Avg. cost | RSD% | $sec_{inc}$ | |
| CBMix4 | 98 | 7450 | 7569 | 7507.0 | 0.50 | 2366.11 | 7942 | 8104 | 8041.1 | 0.60 | 263.46 | 7.11 |
| CBMix7 | 168 | 9398 | 9601 | 9491.9 | 0.55 | 2005.13 | 10072 | 10240 | 10184.0 | 0.52 | 890.30 | 7.29 |
| CBMix8 | 177 | 10305 | 10474 | 10370.6 | 0.49 | 2782.04 | 11100 | 11383 | 11206.4 | 0.72 | 693.50 | 8.06 |
| CBMix15 | 91 | 8214 | 8271 | 8240.5 | 0.24 | 1567.61 | 8796 | 9031 | 8945.7 | 0.74 | 1058.23 | 8.56 |
| CBMix16 | 169 | 8714 | 8743 | 8734.4 | 0.15 | 1726.36 | 9374 | 9699 | 9532.9 | 0.94 | 1725.96 | 9.14 |
| CBMix19 | 212 | 16159 | 16374 | 16253.1 | 0.45 | 2530.33 | 17686 | 18022 | 17833.3 | 0.52 | 747.66 | 9.72 |
| BHW9 | 178 | 875 | 889 | 879.8 | 0.43 | 2339.45 | 971 | 981 | 976.4 | 0.33 | 1351.76 | 10.98 |
| BHW12 | 115 | 10873 | 10932 | 10893.9 | 0.14 | 1742.89 | 11264 | 11492 | 11387.2 | 0.58 | 2041.11 | 4.53 |
| BHW15 | 128 | 15371 | 15446 | 15403.7 | 0.14 | 2006.11 | 15804 | 15950 | 15867.5 | 0.29 | 20.70 | 3.01 |
| BHW16 | 410 | 43837 | 44118 | 43982.8 | 0.24 | 3343.97 | 45738 | 46397 | 46023.6 | 0.39 | 69.87 | 4.64 |
| BHW20 | 293 | 16215 | 16385 | 16279.7 | 0.36 | 2566.36 | 17320 | 17828 | 17603.4 | 0.82 | 1334.55 | 8.13 |
| DI-NEARP-n240-Q4k | 240 | 18181 | 18266 | 18202.6 | 0.15 | 2678.55 | 18858 | 19091 | 18947.6 | 0.32 | 29.70 | 4.09 |
| DI-NEARP-n422-Q8k | 422 | 14442 | 14442 | 14442.0 | 0.00 | 621.47 | 14653 | 14705 | 14677.0 | 0.10 | 12.76 | 1.63 |
| DI-NEARP-n442-Q8k | 442 | 43249 | 43360 | 43270.9 | 0.07 | 1495.41 | 44033 | 44294 | 44105.2 | 0.21 | 70.60 | 1.93 |
| DI-NEARP-n477-Q2k | 477 | 22880 | 23104 | 22973.7 | 0.28 | 2069.61 | 23724 | 24030 | 23945.6 | 0.39 | 81.17 | 4.23 |
| DI-NEARP-n699-Q4k | 699 | 39761 | 40505 | 40185.7 | 0.62 | 2936.83 | 41618 | 42326 | 42118.2 | 0.46 | 182.31 | 4.81 |
| DI-NEARP-n833-Q16k | 833 | 32353 | 32852 | 32517.9 | 0.49 | 3191.94 | 34567 | 34862 | 34754.5 | 0.26 | 342.72 | 6.88 |
| Total/average | | 318277 | 321331 | 319630.2 | 0.31 | 2233.54 | 333520 | 338435 | 336149.6 | 0.48 | 642.14 | 5.17 |

Table 3: Assessment of the AILS Tree-Destructor

| | | Basic Configuration | | | | | No Tree Destructor | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Instance | $\tau$ | Min cost | Max cost | Avg. cost | RSD% | $sec_{inc}$ | Min cost | Max cost | Avg. cost | RSD% | $sec_{inc}$ | $\Delta$% |
| CBMix4 | 98 | **7450** | 7569 | 7507.0 | 0.50 | 2366.11 | 7475 | 7572 | 7529.7 | 0.37 | 2780.88 | 0.30 |
| CBMix7 | 168 | **9398** | 9601 | 9491.9 | 0.55 | 2005.13 | 9436 | 9564 | 9494.1 | 0.45 | 1936.40 | 0.02 |
| CBMix8 | 177 | **10305** | 10474 | 10370.6 | 0.49 | 2782.04 | 10326 | 10482 | 10390.2 | 0.46 | 2255.34 | 0.19 |
| CBMix15 | 91 | **8214** | 8271 | 8240.5 | 0.24 | 1567.61 | 8221 | 8285 | 8263.6 | 0.21 | 1490.54 | 0.28 |
| CBMix16 | 169 | 8714 | 8743 | 8734.4 | 0.15 | 1726.36 | **8710** | 8755 | 8738.4 | 0.20 | 2282.77 | 0.05 |
| CBMix19 | 212 | 16159 | 16374 | 16253.1 | 0.45 | 2530.33 | **16045** | 16390 | 16260.6 | 0.69 | 2761.19 | 0.05 |
| BHW9 | 178 | 875 | 889 | 879.8 | 0.43 | 2339.45 | **873** | 889 | 880.4 | 0.52 | 2153.09 | 0.07 |
| BHW12 | 115 | **10873** | 10932 | 10893.9 | 0.14 | 1742.89 | 10882 | 10932 | 10906.3 | 0.14 | 1817.34 | 0.11 |
| BHW15 | 128 | 15371 | 15446 | 15403.7 | 0.14 | 2006.11 | **15370** | 15488 | 15423.6 | 0.22 | 2086.28 | 0.13 |
| BHW16 | 410 | 43837 | 44118 | 43982.8 | 0.24 | 3343.97 | **43836** | 44286 | 43996.2 | 0.32 | 3181.37 | 0.03 |
| BHW20 | 293 | 16215 | 16385 | 16279.7 | 0.36 | 2566.36 | **16197** | 16378 | 16260.4 | 0.33 | 3154.39 | -0.12 |
| DI-NEARP-n240-Q4k | 240 | **18181** | 18266 | 18202.6 | 0.15 | 2678.55 | **18181** | 18272 | 18205.9 | 0.17 | 2214.66 | 0.02 |
| DI-NEARP-n422-Q8k | 422 | **14442** | 14442 | 14442.0 | 0.00 | 621.47 | **14442** | 14467 | 14447.0 | 0.07 | 1341.39 | 0.03 |
| DI-NEARP-n442-Q8k | 442 | **43249** | 43360 | 43270.9 | 0.07 | 1495.41 | 43264 | 43360 | 43302.7 | 0.11 | 923.48 | 0.07 |
| DI-NEARP-n477-Q2k | 477 | **22880** | 23104 | 22973.7 | 0.28 | 2069.61 | 22896 | 23099 | 23006.9 | 0.30 | 2277.07 | 0.14 |
| DI-NEARP-n699-Q4k | 699 | **39761** | 40505 | 40185.7 | 0.62 | 2936.83 | 39795 | 40592 | 40258.9 | 0.64 | 2735.50 | 0.18 |
| DI-NEARP-n833-Q16k | 833 | **32353** | 32852 | 32517.9 | 0.49 | 3191.94 | 32414 | 32785 | 32580.7 | 0.34 | 3108.61 | 0.19 |
| Total/average | | **318277** | 321331 | 319630.2 | 0.31 | 2233.54 | **318363** | 321596 | 319945.6 | 0.32 | 2264.72 | 0.10 |

$sec_{tot}$ for AILS was set to 3600 seconds. For each instance, we also report the CPU time $sec_{inc}$ to find the solution with the given cost. All detailed experimental results are found in Tables A.1-A.15 in Appendix A (online only).

### 4.3 Results on MCGRP benchmarks

All five MCGRP benchmarks that currently exist in the literature were used for empirical assessment of AILS, namely, `CBMix` proposed by Prins and Bouchenoua (2005), `BHW` and `DI-NEARP` proposed by Bach et al. (2013), and `mggdb` and `mgval` by Bosco et al. (2013). The `CBMix` benchmark consists of 23 randomly generated instances on mixed graphs that imitate real street networks. They contain from 11 to 150 nodes, and from 27 to 332 links. The instances have a number of requests between 20 and 212, located on a combination of nodes, edges, and arcs. On average, the 50% of the nodes, edges, and arcs have to be serviced.

The `BHW` set has 20 test problems generated by modifying well-known instances from the CARP literature, including `gdb` instances (see Golden et al. (1983)), `val` instances (see Benavent et al. (1992)), and `egl` instances (see Li and Eglese (1996)). Instances contain from 11 to 72 nodes, 0 to 51 edges, and 22 to 380 arcs. The number of requests varies from 20 to 410, and on average, about 62% of the nodes, edges, and arcs have requests.

The `DI-NEARP` benchmark with 24 instances originates from six real-life newspaper carrier routing cases in Norway. There are four different variants, corresponding to a reasonable range of capacity values for each case. The instances contain from 563 to 1120 nodes, from 815 to 1450 edges, but no arcs. The number of requests varies from 240 to 833, and roughly 1/3 of the nodes and edges require service.

In contrast with `CBMix`, `BHW`, and `DI-NEARP`, `mggdb` and `mgval` include a fleet size constraint. Both consist of six subsets, each corresponding to a specific value of the parameter $\beta \in \{0.25, 0.30, 0.35, 0.40, 0.45, 0.50\}$ that controls the number of required links in the original CARP instance that have been shifted to adjacent vertices. Each of the six `mggdb` subsets has 23 instances with between 18 and 48 tasks, and between 3 and 10 vehicles. For `mgval`, each of the six subsets has 34 instances. The number of tasks is between 38 and 129, and fleet size varies between 2 and 10.

In Table A.1, we compare upper bounds for the `CBMix` instances yielded by all five MCGRP approximation methods that we have found in the literature, namely:

- MA: the memetic algorithm by Prins and Bouchenoua (2005), run on an Intel Pentium III at 1.0 GHz;

- SA: the simulated annealing algorithm by Kokubugata et al. (2007), run on an Intel Pentium IV at 1.8 GHz;
- Spider: the commercial VRP solver tested in Hasle et al. (2012), run on an Intel Core i7 at 3.07 GHz;
- MH: a matheuristic proposed in Bosco et al. (2014), run on an Intel Xeon Quad at 3.0 GHz;
- AILS: our proposed metaheuristic, run on an Intel Xeon E5530 at 2.4 GHz.

For MA and SA, the termination condition was the number of iterations without new accepted solutions. Spider, MH, and AILS were stopped after a given CPU time limit. MA, Spider, MH, and AILS were run just once on each instance, whereas SA was run ten times by varying the random seed generator. It is worth noting that Spider has been implemented to solve a large variety of routing problems; it is not specifically designed for the MCGRP. In each line of Table A.1, we give the name of the instance addressed, the $\tau$ value, the best known lower bound (column $LB$) from Bach et al. (2013), Bosco et al. (2013), Bach et al. (2014), and Irnich et al. (2015), as well as upper bounds and CPU times yielded by the branch-and-cut-and-price exact method of Bach et al. (2014) (columns marked B&C&P). The latter results were obtained with an Intel Core 2 Duo CPU P8700 at 2.53 GHz, and CPLEX 12.4 to solve the LP-relaxation. The time limit $sec_{tot}$ was six hours. Note that for all exact method results, the reported CPU times are equal to the time limit, unless an optimal solution has been proven, as the time to find the best upper bound is not known. For MA, we give the solution value $z$ it obtains, and the CPU seconds required to run to completion, $sec_{tot}$. For SA, we give the average solution value $avg\ z$ over the 10 attempts and the average CPU time $sec_{tot}$ in seconds to run to completion. For Spider, that was run a single time for two hours on each instance, we provide the solution value $z$ and the CPU time in seconds in which the incumbent solution was found, $sec_{inc}$. For MH, the time limit has not been reported. AILS was given a time limit of one CPU hour. For both, we report the cost and the time $sec_{inc}$ in seconds needed to reach the reported solution value. The bottom three lines give, for each method, the number of optimal upper bounds, the number of best known upper bounds (including optimal values), and the number of instances for which no feasible solution has been found within the time limit. The best objective function values obtained are reported in bold. We observe that the AILS is very effective on the `CBMix` benchmark. It produces all best known solutions but one, 19 of them for the first time. The sum of objective values over all instances is 1.8% lower than for the best competitor MH, even though the average CPU time to find the reported solutions is substantially lower on a similar computer.

In Table A.2 we present the results we obtained on the BHW test set, in a similar way as in Table A.1. Here, the lower bounds reported are the best among the two exact methods that have been tested on this benchmark, and we refer to Bach et al. (2013), and Bach et al. (2014). For BHW, the only competing approximation method is Spider. We observe that the branch-and-cut-and-price of Bach et al. (2014) yields four proven optimal solutions. For these instances, optimal solutions are also found by both approximation methods. In addition, the Bach et al. (2014) method provides three upper bounds that are not competitive, and no upper bound for the remaining 13 instances. Compared to the competition, AILS provides better or equally good solutions on all BHW instances, with 1.8% lower total cost than Spider, and 15 new best solutions. BHW5 was closed for the first time by AILS.

In Table A.3 we compare again the performance of AILS with Spider, this time on the DI-NEARP benchmark. The lower bounds are taken from Bach et al. (2013), as no exact method has been tested so far on this set of 24 large-size instances. Also here, the performance of the AILS is good. Indeed, it yields novel best known solutions to all but two instances, lowering the total cost with 1.4% relative to Spider. For some instances, both algorithms find the incumbent solution close to timeout. This is an indication of the complexity of this test bed, and we believe further improvements are possible for many of these instances.

Tables A.4-A.9 compare the performance of AILS on the 138 mggdb instances with all (to our knowledge) competitors in the literature:

- B&C: the branch-and-cut method by Bosco et al. (2013), run on two Intel Xeon Quad at 3 GHz, using CPLEX 12.2;
- B&C&P: the branch-and-cut-and-price exact method of Bach et al. (2014), run on Intel Core 2 Duo CPU P8700 at 2.53 GHz, and CPLEX 12.4;
- B&C2: the branch-and-cut method by Irnich et al. (2015), with the same hardware and basic software as for B&C;
- MH: a matheuristic proposed in Bosco et al. (2014), run on an Intel Xeon Quad at 3.0 GHz;
- AILS: our proposed metaheuristic, run on an Intel Xeon E5530 at 2.4 GHz.

.

B&C2 has been run only on the first three of the `mggdb` subsets. We report CPU times for B&C2 for these instances, and for B&C&P for the remaining instances. Again, note that the reported CPU times for exact methods are equal to the time limit, unless an optimal solution has been proven. The lower bounds are the best ones among the three exact methods reported in the MCGRP literature. For details, we refer to Bosco et al. (2013), Bach et al. (2014), and Irnich et al. (2015).

For 121 of the 138 `mggdb` instances, optimal solutions have been proven by the three exact methods. For 114 of these, AILS finds solutions with optimal upper bounds, in less than one second for most of the cases. For 128 of the 138 instances, AILS finds a solution with the best known upper bound. The MH matheuristic finds 65 optimal solutions and 67 best known upper bounds. The average CPU time for finding the reported solutions is 18.5 seconds for MH and 7.5 seconds for AILS. Over the `mggdb` benchmark, AILS provides 1.4% better upper bounds than MH. The average gap between the AILS upper bounds and the best known lower bounds is less than 0.7%.

Similarly, in Tables A.10-A.15, we compare results for a total of 204 `mgval` instances for the same methods as for `mggdb`, and AILS. Again, the B&C2 has been run only on the first three subsets. B&C&P has provided results on the last subset ($\beta = 0.50$) only. For the exact methods, CPU times are reported for B&C2 for the first three subsets, for B&C&P for the last subset, and for B&C for the remaining two subsets. In total, 105 optimal values have been proven for `mgval`. MH finds 84 of these, whereas AILS finds 102. MH produces 103 best known upper bounds versus 189 for AILS. The average gap to the best lower bound is 4.2% for MH and 1.9% for AILS. Note that no non-trivial lower bound is known for 18 of the `mgval-0.40` and `mgval-0.45` instances. The average CPU time for finding the reported solutions is 1099.2 seconds for MH and 93.7 seconds for AILS.

We conclude from the observations presented above that AILS is a highly competitive approximation method for the MCGRP. Over a total of 409 instances for the five MCGRP benchmarks used in the literature, it yields 381 best known upper bounds and 224 out of 234 currently known optimal solution values within a reasonable time limit.

### 4.4 Results on CVRP and CARP Instances

In Tables 4 and 5, we compare AILS with some of the best performing approximation methods for the CARP and the CVRP, according to the recent surveys in Prins (2014) (Section 7.6), and Laporte et al.

(2014) (Section 4.7), respectively. Note that the reported gap values are calculated relative to the best known solutions as of 2012, some of which have been improved since then, but relative performance should be clear. Again, we refer to Prins (2014) and Laporte et al. (2014) for updated details on a subset of the benchmarks.

For the CARP, we tested seven well-known benchmarks, 23 `gdb` instances proposed in Golden et al. (1983), 34 `val` instances proposed in Benavent et al. (1992), 24 `egl` instances proposed in Li and Eglese (1996), and 100 `bmcv` instances proposed in Beullens et al. (2003) in four datasets (`C`, `D`, `E` and `F`). In Table 4, we compare our approach against six highly competitive CARP metaheuristics:

- GLS: proposed by Beullens et al. (2003), based on guided local search (run on a Pentium II at 500 MHz);
- MA-CARP: proposed by Lacomme et al. (2004a), based on genetic algorithms (run on a Pentium III at 1 GHz);
- BACO: proposed by Lacomme et al. (2004b), based on ant colony optimization (run on a Pentium III at 800 MHz);
- VNS: proposed by Polacek et al. (2008), based on variable neighborhood search (run on a Pentium IV at 3 GHz). Polacek et al. (2008) reported two sets of results, here we only report the "3.0 GHz" solutions;
- TSA: proposed by Brandão and Eglese (2008), based on tabu search (run on a Pentium Mobile at 1.4 GHz). Brandão and Eglese (2008) report results of two versions of TSA, here we show the best configuration, i.e., the second one;
- Ant-CARP: proposed by Santos et al. (2010), based on ant colony optimization (run on an Intel Pentium III at 1 GHz). Santos et al. (2010) report results of two versions of the Ant-CARP, the median of the best one is reported here (Ant-CARP_12);

Note that '–' means that the method has not been tested on this benchmark. We also compare with MA, the memetic algorithm for the MCGRP by Prins and Bouchenoua (2005), run on an Intel Pentium III at 1.0 GHz. We observe that AILS is among the very best competitors on the CARP.

For the CVRP, four heavily investigated benchmarks were used: 14 instances proposed in Christofides and Eilon (1969) and Christofides et al. (1979), 13 instances proposed in Taillard (1993), 20 instances from Golden et al. (1998), and 12 instances from Li et al. (2005). The four first lines in Table 5 shows the average gap for four of the best performing metaheuristics for the CVRP. The last two rows show results for two MCGRP metaheuristics, namely the memetic algorithm of Prins and Bouchenoua (2005) and AILS. The CVRP metaheuristics are the following:

- GRASP: proposed by Prins (2009), based on GRASP and evolutionary local search (run on a 2.8 GHz Pentium 4);
- MB: proposed by Mester and Bräysy (2007), based on active-guided evolution strategies (run on a 2.8 GHz Pentium 4);
- MA-CVRP: proposed by Nagata and Bräysy (2009), based on memetic algorithm (run on a 3.2 GHz Xeon);
- PARALLEL: proposed by Gröer et al. (2011), based on parallel algorithm (run on 50 computers, each dual-core 2.3 GHz Xeon);

We see that AILS is highly competitive also for the CVRP, except for the Golden et al. (1998) instances where the quality is still reasonable. AILS finds a new best known solution for the `D151-14c` (`CMT9`) instance proposed in Christofides et al. (1979) in 180.1 CPU seconds. The detailed solution is given in Appendix B (online only).

Table 4: Average percentage above the BKS for top-performing CARP algorithms in the literature.

| | Problem set | | | | | | |
|---|---|---|---|---|---|---|---|
| Algorithm | `gdb` | `val` | `egl` | `C` | `D` | `E` | `F` |
| GLS | 0.000 | 0.032 | – | 0.047 | 0.011 | 0.098 | 0.000 |
| MA-CARP | 0.025 | 0.132 | 0.805 | – | – | – | – |
| BACO | 0.154 | 0.351 | 2.348 | – | – | – | – |
| VNS | – | 0.056 | 0.538 | – | – | – | – |
| TSA | 0.070 | 0.100 | 0.725 | 0.054 | 0.164 | 0.168 | 0.249 |
| Ant-CARP | 0.102 | 0.083 | 0.558 | 0.210 | 0.083 | 0.360 | 0.199 |
| MA | 0.285 | – | – | – | – | – | – |
| AILS | 0.000 | 0.054 | 0.328 | 0.024 | 0.155 | 0.125 | 0.017 |

Table 5: Average percentage above the BKS for top-performing CVRP algorithms in the literature.

| Algorithm | Problem set: Christofides et al. (1969, 1979) | Taillard (1993) | Golden et al. (1998) | Li et al. (2005) |
|---|---|---|---|---|
| GRASP | 0.071 | – | 0.525 | – |
| MB | 0.027 | 0.236 | 0.263 | 0.202 |
| MA-CVRP | 0.030 | 0.096 | 0.210 | – |
| PARALLEL | 0.085 | 0.131 | 0.411 | 0.299 |
| MA | 0.389 | – | – | – |
| AILS | 0.159 | 0.167 | 1.455 | 0.433 |

# 5 Conclusions

The Mixed Capacitated General Routing Problem, also called the Node, Edge, and Arc Routing Problem, provides a capacitated multi-vehicle VRP variant that captures an arbitrary mixture of requests on links and nodes. As far as we know, the problem was first studied by Pandi and Muralidharan (1995). Despite the fact that the MCGRP is scientifically interesting and has considerable practical value, it has received limited attention. Recently, however, several metaheuristics, a lower bound procedure, an ILP formulation, three exact methods, and a matheuristic have been proposed.

In this paper, we report the design and investigation of a new hybrid metaheuristic, called AILS, containing several innovations and non-standard mechanisms, for solving MCGRP instances also of industrial size. Computational experiments on five MCGRP benchmarks show excellent performance, with best known solutions to 64 of 67 instances of the `CBMix`, `BHW`, and `DI-NEARP` benchmarks in reasonable time, 55 of which are new. For the smaller size `mggdb` and `mgval` benchmarks designed to investigate exact methods, AILS finds 317 of 342 best known upper bounds, and 216 out of 226 proven optimal solutions. A comparative assessment of the AILS metaheuristic on 181 CARP and 59 CVRP instances proves that our metaheuristic is also among the best for special cases of the MCGRP, in fact improving the best known solution for the much studied `D151-14c` CVRP instance proposed by Christofides et al. (1979).

## Acknowledgements

This work has been partly financed by the Research Council of Norway, under contracts 176869/V30, 187293/I40, 205298/V30, and 217108/I40. We thank two anonymous referees whose comments have greatly improved the quality of this paper.

## References

L. Bach, G. Hasle, and S. Wøhlk. A lower bound for the node, edge, and arc routing problem. *Computers & Operations Research*, 40(4):943–952, 2013. ISSN 0305-0548.

L. Bach, J. Lysgaard, and S. Wøhlk. A Branch-and-Cut-and-Price Algorithm for the Mixed Capacitated General Routing Problem. In *Lukas Bach: Routing and Scheduling Problems - Optimization using Exact and Heuristic Methods*, Ph.D. Thesis, chapter 2, pages 37–75. Aarhus University, 2014.

R. Baldacci and A. Mingozzi. A unified exact method for solving different classes of vehicle routing problems. *Mathematical Programming*, 120:347–380, 2009.

R. Baldacci, N. Christofides, and A. Mingozzi. An exact algorithm for the vehicle routing problem based on the set partitioning formulation with additional cuts. *Mathematical Programming*, 115:351–385, 2008.

E. Bartolini, J.-F. Cordeau, and G. Laporte. Improved lower bounds and exact algorithm for the capacitated arc routing problem. *Mathematical Programming*, 137(1–2):1–44, 2011.

E. Benavent, V. Campos, A. Corberán, and M. Mota. The Capacitated Arc Routing Problem. Lower Bounds. *Networks*, 22:669–690, 1992.

P. Beullens, L. Muyldermans, D. Cattrysse, and D. Van Oudheusden. A Guided Local Search Heuristic for the Capacitated Arc Routing Problem. *European Journal of Operational Research*, 147(3):629–643, 2003.

L. Bodin, V. Maniezzo, and A. Mingozzi. Street routing and scheduling problems. In Randolph W. Hall and Frederick S. Hillier, editors, *Handbook of Transportation Science*, volume 56 of *International Series in Operations Research & Management Science*, pages 413–449. Springer US, 2003. ISBN 978-0-306-48058-4.

A. Bosco, D. Laganà, R. Musmanno, and F. Vocaturo. Modeling and solving the mixed capacitated general routing problem. *Optimization Letters*, 7(7):1451–1469, 2013.

A. Bosco, D. Laganà, R. Musmanno, and F. Vocaturo. A matheuristic algorithm for the mixed capacitated general routing problem. *Networks*, 64(4):262–281, 2014.

J. Brandão and R. Eglese. A deterministic tabu search algorithm for the capacitated arc routing problem. *Computers & Operations Research*, 35(4):1112–1126, 2008.

N. Christofides and S. Eilon. An algorithm for the vehicle-dispatching problem. *Oper. Res. Quart.*, 20(3): 309–318, 1969.

N. Christofides, A. Mingozzi, and P. Toth. The vehicle routing problem. In N. Christofides, A. Mingozzi, P. Toth, and C. Sandi, editors, *Combinatorial Optimization*, pages 315–338. Wiley, Chichester, 1979.

A. Corberán and D. Laporte. *Arc Routing: Problems, Methods, and Applications.* SIAM Monographs on Discrete Mathematics and Applications, Philadelphia, 2014.

A. Corberán and C. Prins. Recent results on arc routing problems: An annotated bibliography. *Networks*, 23:50–69, 2010.

A. Corberán, A.N. Letchford, and J.M. Sanchis. A cutting plane algorithm for the general routing problem. *Mathematical Programming, Series A*, 90:291–316, 2001.

A. Corberán, G. Mejia, and J.M. Sanchis. New results on the mixed general routing problem. *Operations Research*, 53:363–376, 2005.

A. Corberán, E. Mota, and J.M. Sanchis. A comparison of two different formulations for arc routing problems on mixed graphs. *Computers & Operations Research*, 33:3384–3402, 2006.

A. Corberán, I. Plana, and J.M. Sanchis. A branch & cut algorithm for the windy general routing problem and special cases. *Networks*, 49:245–257, 2007.

M. Dell'Amico, J.C. Díaz Díaz, and M. Iori. The bin packing problem with precedence constraints. *Operations Research*, 60:1491–1504, 2012.

B. Golden, S. Raghavan, and E. Wasil (eds.). *The Vehicle Routing Problem: Latest Advances And New Challenges*, volume 43 of *Operations Research/Computer Science Interfaces Series*. Springer, Berlin, 2008.

B.L. Golden and R.T. Wong. Capacitated Arc Routing Problems. *Networks*, 11(3):305–315, 1981.

B.L. Golden, J.S. DeArmon, and E.K. Baker. Computational experiments with algorithms for a class of routing problems. *Computers & Operations Research*, 10:47–59, 1983.

B.L. Golden, E.A. Wasil, J.P. Kelly, and I.M. Chao. Metaheuristics in Vehicle Routing. In T. Crainic and G. Laporte, editors, *Fleet management and logistics*, pages 33–56. Boston, MA: Kluwer, 1998.

C. Gröer, B. Golden, and E. Wasil. A Parallel Algorithm for the Vehicle Routing Problem. *INFORMS Journal on Computing*, 23:315–330, 2011.

J.C.A. Gutiérrez, D. Soler, and A. Hérvas. The capacitated general routing problem on mixed graphs. *Revista Investigacion Operacional*, 23:15–26, 2002.

G. Gutin and A.P. (eds.) Punnen. *The Traveling Salesman and its Variations.* Kluwer, Dordrecht, 2002.

G. Hasle. Arc routing applications in newspaper delivery. In A. Corberán and D. Laporte, editors, *Arc Routing: Problems, Methods, and Applications*, pages 371–395. SIAM Monographs on Discrete Mathematics and Applications, Philadelphia, 2014.

G. Hasle, O. Kloster, M. Smedsrud, and K. Gaze. Experiments on the node, edge, and arc routing problem. Technical Report A23265, ISBN 978-82-14-05288-6, SINTEF, 2012.

S. Irnich, D. Laganà, C. Schlebusch, and F. Vocaturo. Two-phase branch-and-cut for the mixed capacitated general routing problem. *European Journal of Operational Research*, 243(1):17–29, 2015.

H. Kokubugata, A. Moriyama, and H. Kawashima. A practical solution using simulated annealing for general routing problems with nodes, edges, and arcs. In *Engineering Stochastic Local Search Algorithms. Designing, Implementing and Analyzing Effective Heuristics*, volume 4638, pages 136–149. Springer Berlin/Heidelberg, 2007.

P. Lacomme, C. Prins, and W. Ramdane-Chérif. Competitive memetic algorithms for arc routing problems. *Annals of Operations Research*, 131(1):159–185, 2004a.

P. Lacomme, C. Prins, and A. Tanguy. First competitive ant colony scheme for the carp. In M. Dorigo, M. Birattari, C. Blum, L.M. Gambardella, F. Mondada, and T. Stützle, editors, *Ant Colony Optimization and Swarm Intelligence*, volume 3172 of *Lecture Notes in Computer Science*, pages 426–427. Springer Berlin Heidelberg, 2004b. ISBN 978-3-540-22672-7.

G. Laporte, S. Ropke, and T. Vidal. Heuristics for the Vehicle Routing Problem. In P. Toth and D. Vigo, editors, *Vehicle routing: problems, methods, and applications*, chapter 4, pages 87–116. SIAM, 2014. ISBN 978-1-611973-58-7.

F. Li, B. Golden, and E. Wasil. Very large-scale vehicle routing: New test problems, algorithms, and results. *Computers & Operations Research*, 32(5):1165–1179, 2005.

L.Y.O. Li and R.W. Eglese. An Interactive Algorithm for Vehicle Routing for Winter-Gritting. *Journal of the Operational Research Society*, 47:217–228, 1996.

H.R. Lourenço, O.C. Martin, and T. Stützle. Iterated local search: Framework and applications. In M. Gendreau and J.-Y. Potvin, editors, *Handbook of Metaheuristics, second edition*, volume 146 of *International Series in Operations Research & Management Science*, pages 363–398. Springer, Berlin, 2010.

S. K. Mandal, D. Pacciarelli, A. Løkketangen, and G. Hasle. A memetic NSGA-II for the bi-objective mixed capacitated general routing problem. *Journal of Heuristics*, 21(3):359–390, 2015.

R. Martinelli, D. Pecin, M. Poggi, and H. Longo. A branch-cut-and-price algorithm for the capacitated arc routing problem. In M.P. Pardalos and S. Rebennack, editors, *Experimental Algorithms*, volume 6630 of *Lecture Notes in Computer Science*, pages 315–326. Springer Berlin Heidelberg, 2011.

D. Mester and O. Bräysy. Active-guided evolution strategies for large-scale Vehicle Routing Problems. *Computers & Operations Research*, 34:2964–2975, 2007.

Y. Nagata and O. Bräysy. Edge Assembly based Memetic Algorithm for the Capacitated Vehicle Routing Problem. *Networks*, 54:205–215, 2009.

C.S. Orloff. A fundamental problem in vehicle routing. *Networks*, 4(1):35–64, 1974.

R. Pandi and B. Muralidharan. A capacitated general routing problem on mixed networks. *Computers & Operations Research*, 22:465–478, 1995.

D. Pisinger and S. Røpke. A general heuristic for vehicle routing problems. *Computers & Operations Research*, 34(8):2403–2435, 2007.

M. Polacek, K.F. Doerner, R.F. Hartl, and V. Maniezzo. A variable neighborhood search for the capacitated arc routing problem with intermediate facilities. *Journal of Heuristics*, 14(5):405–423, 2008.

C. Prins. A GRASP * Evolutionary Local Search Hybrid for the Vehicle Routing Problem. In F.B. Pereira and J. Tavares, editors, *Bio-inspired Algorithms for the Vehicle Routing Problem*, volume 161 of *Studies in Computational Intelligence*, pages 35–53. Springer, Berlin, 2009.

C. Prins. The Capacitated Arc Routing Problem: Heuristics. In A. Corberán and G. Laporte, editors, *Arc routing: problems, methods, and applications*, chapter 7, pages 131–157. SIAM, 2014. ISBN 978-1-611973-66-2.

C. Prins and S. Bouchenoua. A Memetic Algorithm Solving the VRP, the CARP and General Routing Problems with Nodes, Edges and Arcs. In *Recent Advances in Memetic Algorithms*, volume 166, pages 65–85. Springer Berlin / Heidelberg, 2005.

S. Røpke and D. Pisinger. An Adaptive Large Neighborhood Search Heuristic for the Pickup and Delivery Problem with Time Windows. *Transportation Science*, 40(4):455–472, 2006.

L. Santos, J. Coutinho-Rodrigues, and J. R. Current. An improved ant colony optimization based algorithm for the capacitated arc routing problem. *Transportation Research Part B: Methodological*, 44(2):246–266, 2010.

P. Shaw. A New Local Search Algorithm Providing High Quality Solutions to Vehicle Routing Problems. Technical report, University of Strathclyde, 1997.

P. Shaw. Using Constraint Programming and Local Search Methods to Solve Vehicle Routing Problems. In *Proceedings CP-98 (Fourth International Conference on Principles and Practice of Constraint Programming)*, 1998.

SINTEF. TOP website. URL `http://www.sintef.no/top`.

É.D. Taillard. Parallel iterative search methods for vehicle routing problems. *Networks*, 23:661–673, 1993.

P. Toth and D. Vigo. An overview of vehicle routing problems. In P. Toth and D. Vigo, editors, *The Vehicle Routing Problem*, pages 1–26. SIAM Monographs on Discrete Mathematics and Applications, Philadelphia, 2002.

P. Toth and D. Vigo. *Vehicle Routing: Problems, Methods, and Applications*. SIAM Monographs on Discrete Mathematics and Applications, Philadelphia, second edition, 2014.

T. Vidal, T.G. Crainic, M. Gendreau, N. Lahrichi, and W. Rei. A hybrid genetic algorithm for multidepot and periodic vehicle routing problems. *Operations Research*, 60:611–624, 2012.

S. Wøhlk. New lower bound for the capacitated arc routing problem. *Computers & Operations Research*, 33:3458–3472, 2006.

S. Wøhlk. A decade of capacitated arc routing. In B. Golden, S. Raghavan, and E. Wasil, editors, *The Vehicle Routing Problem: Latest Advances And New Challenges*, volume 43 of *Operations Research/Computer Science Interfaces Series*, pages 29–48. Springer, Berlin, 2008.