Inter-Depot Moves and Dynamic-Radius Search for Multi-Depot Vehicle Routing Problems

Jean Bertrand Gauthier^{*,a}, Stefan Irnich^a

^aChair of Logistics Management, Gutenberg School of Management and Economics, Johannes Gutenberg University Mainz, Jakob-Welder-Weg 9, D-55128 Mainz, Germany.

Abstract

Dynamic-radius search, formerly known as sequential search, is an effective neighborhood exploration technique for standard edge-exchange neighborhoods such as 2-opt, 2-opt^{*}, swap, relocation, Or-opt, string exchange, etc. Up to now, it has only been used for vehicle routing problems with a homogeneous fleet and in the single-depot context. In this work, we extend dynamic-radius search to the multi-depot vehicle routing problem, in which 2-opt and 2-opt^{*} moves may involve routes from different depots. To this end, we equip dynamic-radius search with a modified pruning criterion that still guarantees identifying a best-improving move, either intra-depot or inter-depot, with little additional computational effort. We experimentally confirm that substantial speedups of factors of 100 and more are achieved compared to an also optimized implementation of lexicographic search, another effective neighborhood exploration technique using a feasibility-based pruning criterion. As one would expect, better local optima are found on average when allowing inter-depot moves in radius search (positive result). Against intuition, we do however not end up with a better ILS metaheuristic regarding best found solutions, i.e., better average results do not translate into better overall results. We can at least partly explain the latter negative result, which might be useful for other researchers and their attempt to algorithmically optimize their neighborhood exploration procedures.

Key words: Routing, Local search, Sequential search, Dynamic-radius search, Inter-depot

1. Introduction

This paper extends the realm of application of dynamic-radius search, an effective neighborhood exploration technique, to the *multi-depot vehicle routing problem* (MDVRP). Escobar et al. (2014) note that it is rather straightforward to extend most classical neighborhoods to the multi-depot environment except 2-opt and 2-opt^{*}. The description is however quite limited so we formalize the various inter-depot cases that arise in these two neighborhoods and sketch a few others that work naturally with multi-depot. This analysis also serves to equip dynamic-radius search with a modified pruning criterion so that best-improving moves which allow inter-depot edges are found with little additional computational effort. The idea behind this work revolves around the hypothesis that being systematically able to find moves with better gains leads to a local search that can find better local optima. Our analysis leads to positive and negative results. The positive result is that allowing inter-depot moves helps finding better local optima on average. This unfortunately does not translate to the meta-strategy where ultimately only the overall best local optimum is important. While one can compare this best gain idea to Dantzig's pivot rule in the primal simplex algorithm for which it is known that alternative less myopic rules perform better, it is not immediately clear to us that the behavior of a heuristic would be the same as that of an exact method.

^{*}Corresponding author.

Email addresses: jgauthie@uni-mainz.de (Jean Bertrand Gauthier), irnich@uni-mainz.de (Stefan Irnich)

Dynamic-radius search builds on several classical works on the symmetric traveling salesman problem (TSP). In the context of the TSP, Hoos and Stützle (2005) use the expression fixed-radius search to collectively describe the idea of Steiglitz and Weiner (1968) and numerous extensions such as Bentley (1992); Martin et al. (1992); Reinelt (1994); Johnson and McGeoch (1997). For each vertex i, the predecessor p_i and successor s_i in the current TSP tour must be known. Then, for finding improving 2-opt and 3-opt moves, the neighborhood exploration procedures first loop over all vertices i to determine a first deleted edge (p_i, i) or (i, s_i) . The first inserted edge e = (i, j), replacing the deleted edge, must now be shorter, i.e., $c_{ij} < c_{p_i,i}$ or $c_{ij} < c_{i,s_i}$, respectively. This can be interpreted as choosing j as a neighbor of i within the known radius given by $c_{p_i,i}$ or c_{i,s_i} . For 3-opt, the second decision about the second edge to insert must again produce a positive accumulated gain. This selection criterion for inserted edges is know as the gain criterion. Lin and Kernighan (1973) always choose edges according to the gain criterion in their famous Lin-Kernighan variable-depth neighborhood. Irnich et al. (2006) have generalized this idea of using the gain criterion for the (capacitated) vehicle routing problem ((C)VRP) and to VRP-specific moves such as node/vertex relocation, Or-opt, string exchange, etc. Moreover, Irnich (2008b) has shown that these techniques can be further generalized to handle more constrained versions of the VRP such as the VRP with time windows, the periodic VRP, pickup-and-delivery problems, and others.

Irnich et al. (2006) have introduced dynamic-radius search for VRPs under the name sequential search. This naming was inspired by the fact that moves are decomposed into partial moves and complete moves result from sequentially selecting the constituting partial moves. Note that cyclic independent move decomposition as coined in Irnich et al. (2006) relies on the fact that all 2-opt and 3-opt moves can be characterized by a single alternating cycle of deleted and added edges (see Funke et al. 2005; for a deeper analysis of single alternating cycle TSP neighborhoods). However, the term sequential search might nowadays be confused with single-threaded algorithms opposed to parallel/multi-threaded algorithms. Moreover, we think that the term radius better captures the idea of pruning the search tree using the gain criterion whereas dynamic follows in the footstep of the aforementioned fixed version. Certainly not taken lightly, we finally opt for this name change decision.

A formal definition of the MDVRP is in order. It can be defined over a complete undirected graph G = (V, E) where the vertex set V consists of customers N and depots D. Each customer $i \in N$ has a positive demand q_i and we define $q_d = 0$ for every depot $d \in D$. Moreover, a fleet of m homogeneous vehicles characterized by a capacity Q and a maximal tour duration T (one of these limits may be omitted) are stationed at each depot $d \in D$. A global fleet-size limit $\kappa \leq m |D|$ on the number of vehicles may also be given. The edge set E models direct connections between customers as well as depots and customers. For each edge $e \in E$, a routing cost c_e and a travel time t_e are given. A route for depot $d \in D$ is a directed cycle $r = (d = i_0, i_1, i_2, \ldots, i_p, i_{p+1} = d)$ in which all vertices $i_1, i_2, \ldots, i_p \in N$ are different customers. A route r is feasible if it is capacity feasible, i.e., $\sum_{i \in V(r)} q_i \leq Q$, and duration feasible, i.e., $\sum_{e \in E(r)} t_e \leq T$, where V(r) denotes the set of vertices and E(r) the set of edges of r. The cost of a route r is $c_r = \sum_{e \in E(r)} c_e$. Let R_d denote the subset of all feasible, non-empty routes associated with depot $d \in D$. A set R of feasible routes is a solution to the MDVRP if (1) the local and global fleet-size limits are not exceeded, i.e., $|R \cap R_d| \leq m$ for all $d \in D$ and $\sum_{d \in D} |R \cap R_d| = |R| \leq \kappa$, and (2) all customers are serviced exactly once, i.e., $N = \bigcup_{r \in R} (V(r) \cap N)$ and $V(r_1) \cap V(r_2) \cap N = \emptyset$ for all $r_1 \neq r_2 \in R$. Such a solution is optimal if it minimizes the (total) routing costs $c(R) = \sum_{r \in R} c_r$.

Our overall experimental setup uses a rather simple multi-start neighborhood-based local search approach detailed in Section 5. We sketch it now in order to specify precisely the terminology. A neighborhood (e.g., 2-opt) uses moves to map a solution to alternative solutions called *neighbors*. We assume that the available neighborhoods are given by a set Ψ . Hence, for a given neighborhood $\mathcal{N} \in \Psi$ and current solution R, a move $\mu \in \mathcal{N}$ produces from R the neighbor solution $R' = \mu(R)$. A move is improving if its marginal gain defined as $g(\mu, R) = c(R) - c(\mu(R))$ (or simply g when contextually clear) is positive, i.e., g > 0. Moreover, a solution R is a local optimum w.r.t. \mathcal{N} if no improving move $\mu \in \mathcal{N}$ for R exists. Neighborhood exploration is the systematic search for an improving move and associated improving solution. We assume here that the exploration filters out infeasible solutions. Moreover, the pivoting rule (such as first improvement or best improvement) controls whether and (if so) when neighborhood exploration is stopped before all possible moves are considered. Finally, each constructed starting solution R is improved by iterative neighborhood explorations. One of the possible neighborhood exploration technique is dynamic-radius search and we discuss its relationship to other techniques in Section 3.

We combine the neighborhoods $\mathcal{N} \in \Psi$ in a variable neighborhood descent (VND) fashion, i.e., the neighborhoods $\mathcal{N} \in \Psi$ are parameterized and the choice of a next neighborhood to explore is determined according to a priority parameter of each neighborhood. For each starting solution \hat{R} , we finally settle on a heuristic solution \hat{R} which is a local optimum with respect to all $\mathcal{N} \in \Psi$. This solution \hat{R} depends on the starting solution R, the available neighborhoods Ψ , and for each neighborhood $\mathcal{N} \in \Psi$, its prioritization in VND, its pivoting rule, and its neighborhood exploration strategy. We study the trade-off between solution quality and the time it takes to find joint local optima \tilde{R} .

The remainder of this paper is structured as follows: In Section 2, we briefly recall 2-opt and 2-opt^{*} moves for VRPs distinguishing intra/inter-tour moves as well as intra/inter-depot moves. Section 3 discusses neighborhood exploration techniques. The new radius search algorithm for efficiently exploring the extended neighborhoods with inter-depot 2-opt and 2-opt^{*} moves is presented in Section 4. The multi-start neighborhood-based local search approach that we use as a simple metaheuristic is explained in Section 5. Computational results follow in Section 6. Final conclusions are drawn in Section 7.

2. 2-Opt and 2-Opt* Moves in Multi-Depot Vehicle Routing Problems

Notation. A convenient representation of a solution R is obtained by concatenating all routes, in any order and orientation, as one long sequence of vertices. This is known as the *giant route* or *giant tour* representation (Bellmore and Hong 1974). In the following, the resulting sequence, denoted \mathcal{V} , allows us to loop over relevant vertices by writing $i \in \mathcal{V}$. In order to have a unique predecessor p_i and a unique successor s_i for each vertex $i \in \mathcal{V}$, the sequence \mathcal{V} must include different *copies* of the depots, two for each route in order to also distinguish between its source and sink depots (which must represent the identical physical depot). For any route $r \in R$, let the first and last visited customers be denoted by f_r and $l_r \in N$, respectively. Moreover, a reference to the original depot is given by $d_r \in D$. In particular, any two routes r and r' are associated with the same depot if and only if $d_r = d_{r'}$. For any vertex $i \in V$, its associated route is denoted by $r_i \in R$. Finally, we define the corresponding depot d_i of vertex i as d_{r_i} , the first customer f_i as f_{r_i} , and the last customer l_i as l_{r_i} .

Legend. A consequence of having chosen an orientation for each route is that we can depict and write solutions with directed arcs instead of undirected edges. Indicating the direction of traversal makes reading and understanding the following figures easier. However, the underlying MDVRP is still assumed to be symmetric. In the following figures, a snake-shaped link between vertices w and v forms a directed path of arbitrary length (written as $w \rightarrow v$ or $v \stackrel{\leftarrow}{\rightarrow} w$ when it is an inversion of a given path) whereas a straight link indicates a single arc. A solid connection keeps its orientation once the move is completed while a dashed one is inverted as a consequence of the move. An arc is marked for deletion with a loosely dotted pattern while a densely dotted one indicates an insertion. Affected vertices are filled in solid color when they are chosen and a lighter shade when they are implied by a choice. This shading rule also applies when deleted and inserted arcs are implied to repair an otherwise infeasible move. Finally, customers take on circle-shaped vertices whereas different depots are explicitly distinguished by different polygon-shaped vertices (squares and pentagons).

Developed with the TSP in mind, Croes (1958) has devised an algorithm which performs so-called *inversions*. These are essentially 2-opt exchanges in the sense Lin (1965) has generally coined λ -opt. The 2-opt neighborhood has been generalized to single-depot vehicle routing problems in a straightforward manner. Potvin and Rousseau (1995) have introduced the 2-opt* neighborhood to tackle the *vehicle routing problem with time windows* (VRPTW, Desaulniers et al. 2014). The leading observation is that a 2-opt inter-tour move induces two segment inversions in the affected routes (see Figure 1b). A 2-opt* move differs from the latter in that it maintains the general ordering of the customers in the current solution and is thus more likely to produce an alternative solution feasible with respect to time windows (see Figure 1c).



Figure 1: Intra-depot 2-opt and 2-opt* moves.

The canonical description of both 2-opt and 2-opt^{*} can be done with two deleted edges (i, s_i) and (j, s_j) as well as two inserted edges. In the 2-opt case, the inserted edges are (i, j) and (s_i, s_j) whereas in the 2-opt^{*} case they are (i, s_j) and (j, s_i) . The gain of these moves can then be computed for 2-opt as $g = c_{i,s_i} + c_{j,s_j} - c_{ij} - c_{s_i,s_j}$ and for 2-opt^{*} as $g = c_{i,s_i} + c_{j,s_j} - c_{j,s_i}$. This description however only holds if we face an *intra-depot* move, that is, a move that affects a single route (*intra-tour*) or two routes (*inter-tour*) related to the same depot. Figure 1, which uses our multi-depot notation, therefore only captures intra-depot 2-opt and 2-opt^{*} moves.

In the next two subsections, we show that the multi-depot case can nevertheless be correctly taken into account during the analysis of an *inter-depot* move. Observe that an inter-depot move must necessarily be inter-tour. In both neighborhoods, 2-opt and 2-opt^{*}, we break down the possible cases that arise and must be covered by an exhaustive neighborhood exploration. In particular, an edge exchange infeasibly connecting two different depots can be repaired in either of two different ways. Our various visual aids give the final interpretation of the traditional neighborhood move together with a *repair operation* that must be performed to ensure each route has matching source and sink depots. Moreover, we show that all final move configurations fall under well-defined cases (standard, exception, and rejection, see below). The gain computation of the final move is of course affected as a byproduct of the repair operation. It is opportune at this point to separate the presentation of 2-opt and 2-opt^{*}.

2.1. Inter-Depot 2-Opt*

An inter-depot 2-opt^{*} move happens when the two deleted arcs (i, s_i) and (j, s_j) belong to two different routes $r_i \neq r_j$ of two different depots $d_i \neq d_j$. As illustrated in Figure 2, eight cases are sufficient to exhaustively cover repair options. Figures 2a and 2b display the *standard case* where both deleted arcs (i, s_i) and (j, s_j) are either not the two first arcs or not the two last arcs in their respective routes r_i and r_j . This offers two possibilities of swapping the depots: swapping at the source (Figure 2a) and at the sink (Figure 2b). Note that both cases allow the deleted arcs (i, s_i) and (j, s_j) to be at the opposite end of the two routes, so that the new tours are then $(d_i, f_j \rightsquigarrow l_j, d_i)$ and $(d_j, f_i \rightsquigarrow l_i, d_j)$. This is the special case of an exchange of two complete routes between two different depots.

The next four exception cases depicted in Figures 2c-2f happen if exactly one of the deleted arcs is the first (last) and the other one is not the first (not the last). In the exception cases, only three arcs instead of four are finally exchanged, one of the arcs inserted in the standard case is absent (arc (i, s_j) or (j, s_i)). Note that the four cases result from the inherent symmetry, on the one hand between the two routes r_i and r_j (swapping the indices i and j), and on the other hand, between source and sink (reversal of the routes' orientation).

The first six cases shown in Figures 2a–2f cover all feasible inter-depot 2-opt^{*} moves. There exist two more cases visualized in Figures 2g and 2h: (left) depot swap of the source depots with $i = d_i$ and $j = d_j$ and (right) depot swap of the sink depots with $s_i = d_i$ and $s_j = d_j$. These cases are however infeasible because the initially chosen arcs for deletion coincide with the two arcs that one wants to delete to perform the source (sink) depot swap.

Ultimately, every 2-opt^{*} inter-route move with different depots $d_i \neq d_j$ is evaluated with two repair operations (one source and one sink) based on the expressed conditions yielding up to two distinct interdepot 2-opt^{*} moves. Note that the conditions on *i* and *j* shown on the left-hand side of Figure 2 are



(a) Standard case: $i \neq d_i, j \neq d_j$.



(c) Exception case: $i = d_i, j \neq d_j$, arc (i, s_j) is absent.



(e) Exception case: $i \neq d_i, j = d_j$, arc (j, s_i) is absent.



(g) Rejection case: $i = d_i, j = d_j$.



(d) Exception case: $s_i \neq d_i, s_j = d_j, \operatorname{arc}(i, s_j)$ is absent.



(f) Exception case: $s_i = d_i, s_j \neq d_j$, arc (j, s_i) is absent.



(h) Rejection case: $s_i = d_i, s_j = d_j$.

Figure 2: Inter-depot 2-opt* moves, swap of source (left) or sink (right) depots.

independent from the conditions on s_i and s_j shown on the right-hand side. For example, if $i = d_i$ and $j \neq d_j$ holds (case 2c), exactly one of the four cases 2b, 2d, 2f, or 2h is true for s_i and s_j .

2.2. Inter-Depot 2-Opt

For the 2-opt neighborhood, we also raise standard and exception cases as presented in Figure 3, but omit rejection cases for the sake of brevity. The repair operation must always swap sink and source depots of the two different routes. As a mnemonic device, depots finally always remain attached to their current route. In the swap cases 3a, 3c, and 3e, the source/sink roles are preserved, whereas the roles are interchanged in the swap cases 3b, 3d, and 3f.

3. Neighborhood Exploration Techniques

In the following, we describe and compare the three fundamental neighborhood exploration techniques in local search: lexicographic search, radius search, and granular search. The organization follows this order and we complete these descriptions in Section 3.4 with a comparison and summarizing remarks.

Regardless of the way a neighborhood is explored, the two bottleneck operations of testing moves are: the gain computation and the feasibility check. For the gain computation, it can be straightforward as is the



(a) Standard case: $s_i \neq d_i, j \neq d_j$.



(c) Exception case: $s_i = d_i, j \neq d_j$, arc (s_i, s_j) is absent.





(b) Standard case: $i \neq d_i, s_j \neq d_j$.



(d) Exception case: $i = d_i, s_j \neq d_j$, arc (i, j) is absent.



(e) Exception case: $s_i \neq d_i, j = d_j$, arc (i, j) is absent. (f) Exception case: $i \neq d_i, s_j = d_j$, arc (s_i, s_j) is absent.

Figure 3: Inter-depot **2-opt** moves, source/sink depot roles are preserved (left) or interchanged (right).

case in the studied variant or involved, e.g., needed when computing route durations in the presence of timedependent travel times as proposed by Visser and Spliet (2020). For the feasibility check, it is almost never a trivial task as per the inclusion of non-additive resource constraints. In accordance with Savelsbergh (1990); Irnich (2008b); Vidal et al. (2014), several resource constraints can nevertheless be tested in *constant time* by doing precomputations on the current solution and some segments of arcs (consecutive arcs in the current solution). For the latter computations, segments have to be built-up in a vertex-by-vertex fashion, leading to the lexicographic search paradigm. Constant time tests are well established for capacity, time-window, and pickup-and-delivery constraints. In the following, we use the 2-opt intra-tour move (see Figure 1a) to explain the different search paradigms. A synopsis of the three search principles is shown in Figure 4.

3.1. Lexicographic Search

Lexicographic search as presented by Savelsbergh (1990) consists of an elegant and systematic way to explore a neighborhood using the customer order observed in the current solution. It is especially intuitive for k-edge exchange moves. For exploring the 2-opt neighborhood, lexicographic search explores the vertices i and j using the order given by \mathcal{V} in two nested loops, see Algorithm 1. The first loop iterates over $i \in \mathcal{V}$ whereas the second loop steps over $j > i \in \mathcal{V}$. Hence, the inner loop iterator is always greater than the preceding outer loop while still covering all possibilities for the 2-opt moves.

The key observation, see also Figure 1a, is that in the inner loop the new route of vertex *i* must contain the path $P = (d_i \rightsquigarrow i, j \nleftrightarrow s_i)$, where the orientation of original path $(s_i \rightsquigarrow j)$ has been inverted. Since path *P* grows by one vertex in every inner-loop iteration, the idea of Savelsbergh (1990) is to prune the search based on the *local infeasibility* of *P* (see first if-condition in Algorithm 1), that is, *break* the inner loop if it can be shown that any further vertex in *P* always leads to a resource-infeasible move. Since this is a necessary but not sufficient condition for the feasibility of the overall move, a global feasibility check



Figure 4: Synopsis of the three neighborhood exploration methods for the intra-depot 2-opt neighborhood. It is assumed that both the feasibility and gain of a move are evaluated directly before the if-condition "infeasible or $g \leq \gamma$ " is reached.

(see second if-condition in Algorithm 1) must be performed. Note that in an inter-tour move, the inner loop would *skip* the remaining vertices s_j, \ldots, d_j of route r_j rather than *break*.

For capacity constraints, the local feasibility of P and any path constructed in later iterations of the inner loop amounts to checking $\sum_{i \in V(P)} q_i \leq Q$ (always fulfilled in the intra-tour case). Assuming that the triangle inequality holds for travel times, route duration constraints necessarily require $\sum_{e \in E(P)} t_e \leq T$. For the MDVRP, all this can be tested in $\mathcal{O}(1)$ by summing up demand and travel times for the segment $(d_i \rightsquigarrow i)$ (outer loop over i) and the segment $(j \overleftarrow{\sim} s_i)$ (inner loop over j).

More intricate feasibility conditions such as time windows, pairing, precedence constraints, and many more can be checked in $\mathcal{O}(1)$ as well. Recall that the 2-opt intra-tour move results in the new route $(d_i \rightsquigarrow i, j \stackrel{\leftarrow}{\rightsquigarrow} s_i, s_j \rightsquigarrow d_i) = (P, s_j \rightsquigarrow d_i)$. As a preparatory step, one must compute an upper bound on the resource consumption when arriving at the last segment $(s_j \rightsquigarrow d_i)$. This requires a $\mathcal{O}(n)$ preparation before the exploration is started. Then, the outer loop computes the resource consumption at the end of the first segment $(d_i \rightsquigarrow i)$, while the inner loop computes the resource extension function (Irnich 2008a) for the second segment $(j \stackrel{\leftarrow}{\rightsquigarrow} s_i)$ so that both the resource values are then propagated along the arc (s_i, s_j) and compared against the respective resource upper bounds that were computed in the preparatory step.

Summarizing, the effectiveness of lexicographic search stems from its feasibility-based pruning. It is particularly well suited for a VRP with intricate or very constraining feasibility constraints. If both checking its resource consumption and propagating resource levels over entire segments can be done in constant time, there is no extra effort in the worst-case time complexity when exploring a neighborhood. For 2-opt, the result is a $\mathcal{O}(n^2)$ neighborhood exploration.

3.2. Radius Search

The idea of accelerating the neighborhood exploration based on the length of the inserted arc can done with a priori computed bounded candidate lists. This idea can be combined with pruning using the gain criterion. Finally, Irnich et al. (2006) sharpened the gain criterion by incorporating the quality of already detected improving solutions. Accordingly, we present candidate-lists based search (Section 3.2.1), fixedradius search (Section 3.2.2), and dynamic-radius search (Section 3.2.3). In all variants, neighbor lists sorted according to arc costs are created once in initialization. For the special case of costs defined by Euclidean distances, one can imagine the search being conducted starting from the closest neighbor and spiraling outwards until some *break* condition is met. (This break condition remains correct in the non-Euclidean case.) Note that the break condition is identical for 2-opt intra- and inter-tour moves so that implementations of radius search naturally consider both types of moves together.

3.2.1. Bounded Candidate-Lists based Search

Bounded candidate-lists based search follows the idea that good arcs to be inserted should have a small cost. It initially builds, for each vertex i, a bounded length candidate list $\mathbf{N}(i)$ of neighbor vertices j in close proximity of i. The neighbor $j \in \mathbf{N}(i)$ represents the arc (i, j), and only moves inserting the arc (i, j) with $j \in \mathbf{N}(i)$ are considered as possible moves. In a naive implementation, a fixed size σ for the neighborhoods is chosen first (e.g., 50 neighbors) and $\mathbf{N}(i)$ is then filled with the σ closest vertices.

In Algorithm 2, the first if-condition for breaking the inner loop is never fulfilled for a non-constraining radius such as $\rho = \infty$. A speedup solely results from the *bounded* candidate lists, because it reduces the number of arcs (i, j) to be tested. Indeed, for a fixed size σ the considered search space becomes linear in $|\mathcal{V}|$. Obviously, it is however not guaranteed that a true local optimum is found as long as $\sigma = |\mathbf{N}(i)| < |\delta(i)|$.

Another way to interpret bounded candidate-lists based search is to see it as a radius search, as depicted in Algorithm 2, where the radius ρ is not computed inside the outer loop. Instead, ρ is *a priori* chosen (for each *i* possibly in a different way) so that the inner for loop and if-condition can be implemented by filling $\mathbf{N}(i)$ appropriately.

For the TSP, bounded candidate lists have also been constructed on the basis of other criteria. Helsgaun (2000), for example, uses a modified edge weight/cost obtained from an approximation of the Held-Karp lower bound. For the tested instances, all edges of an optimal TSP solution were shown to be contained in candidate lists of smaller size σ . No general guarantee can be given for arbitrary instances. In contrast, the fixed-radius search presented next follows a different line of thought in order to achieve provably local optimal solutions.

3.2.2. Fixed-Radius Search

In routing problems, a move μ comprises the deletion of some arcs and the insertion of the same number of different arcs. It can therefore be decomposed into a number k of partial moves p_1, \ldots, p_k , i.e., $\mu = p_1 \circ p_2 \circ \cdots \circ p_k$, each of which deletes and inserts some of the arcs.

Definition 1. (Irnich et al. 2006) A move μ is cyclic-independent if $\mu = p_{\pi(1)} \circ p_{\pi(2)} \circ \cdots \circ p_{\pi(k)}$ for all cyclic permutations of $\{1, 2, \ldots, k\}$ and it is and cost-independent if $g(\mu) = \sum_{i=1}^{k} g(p_i)$, where $g(p_i)$ is a partial gain associated with p_i .

A neighborhood with cyclic-independent and cost-independent moves can be searched with the gain criterion in light of the following theorem:

Theorem 1. (Lin and Kernighan 1973) If a sequence of k numbers $(g_i)_{i=\{1,...,k\}}$ has a positive sum, i.e., $\sum_{i=1}^{k} g_i > 0$, then there exists a cyclic permutation π of these numbers such that every partial sum is positive, i.e., $\sum_{i=1}^{\ell} g_{\pi(i)} > 0$ for all $1 \le \ell \le k$.

Note that neither Definition 1 nor Theorem 1 claim that for a given neighborhood the move decomposition is unique. Let us exemplify the gain criterion for intra-depot 2-opt moves as depicted in Figures 1a and 1b. First, an improving 2-opt move μ has a gain $g = g(\mu) = c_{i,s_i} - c_{ij} + c_{j,s_j} - c_{s_i,s_j} > 0$. Second, to satisfy Definition 1, one can decompose it into $\mu = p_1 \circ p_2 = p_2 \circ p_1$, where one partial move p_1 is deleting arc (i, s_i) plus inserting arc (i, j) and another partial move p_2 is deleting arc (j, s_j) plus inserting arc (s_i, s_j) . From Theorem 1, we obtain that $g = g(p_1) + g(p_2) > 0$ is improving only if

$$g(p_1) = c_{i,s_i} - c_{ij} > 0$$
 or $g(p_2) = c_{j,s_j} - c_{s_i,s_j} > 0.$ (1)

This is a *necessary* condition for move μ to be improving. By handling both options of the compound condition (1), we obtain independent *radius conditions* based on the cost of two different deleted arcs (i, s_i) or (j, s_j) . This can be exploited algorithmically as done with the break condition in Algorithm 2 which makes fixed-radius search effective: given a vertex i, only those neighbors $j \in \mathbf{N}(i)$ which fulfill the first radius condition $c_{ij} < \rho$ with $\rho = c_{i,s_i}$ need to be inspected.

We have to underline that one must ensure that a move rejected from one partial move can be recovered when the other is analyzed. This can be easily overlooked when simplifying the loop design. Indeed, when testing the deleted arc (i, s_i) with $j \in \mathbf{N}(i)$ and later interchanging the roles of the vertices as (j, s_j) with $i \in \mathbf{N}(j)$, we can observe two facts: we may evaluate the same move twice, and we have never evaluated the second radius condition. The former point is a nuisance, but the latter point implies that this is an incomplete examination of the compound condition (1). Additionally testing over neighbors $s_j \in \mathbf{N}(s_i)$ such that $c_{s_j,s_i} < \rho$ however makes the search exhaustive. Indeed, one can verify that once again interchanging the roles of i and j yields a deleted arc (j, s_j) with $s_i \in \mathbf{N}(s_j)$. For the reader keeping count, we are evaluating in the worst case the same move four times. Among these overlapping move evaluations, two are mandatory to ascertain a complete examination of the neighborhood space and two are redundant by cost symmetry. Since it is impossible to know in advance which of the two partial gains, if any, could fulfill its radius condition, the redundancy is in principle unavoidable by the conservative nature of the gain criterion. As supported by our computational results, we however claim that this redundancy is in practice more limited than what transpires by this quadruple factor (see also the example presented in the next section).

Finally, two equivalent nested loop constructions are possible: an outer loop $i \in \mathcal{V}$ followed by two inner loops $j \in \mathbf{N}(i)$ and $s_j \in \mathbf{N}(s_i)$, or alternatively an outer loop $i \in \mathcal{V}$ for which both deleted arcs (p_i, i) and (i, s_i) are tested followed by a single inner loop $j \in \mathbf{N}(i)$. We retain the latter presentation in Algorithm 2 for aesthetics reasons but advice the former for low-level efficiency. Indeed, looking at the finer details of the implementation, measurable speedup comes from tailoring neighbor lists to the specific inner loops $j \in \mathbf{N}(i)$ and $s_j \in \mathbf{N}(s_i)$ to include or exclude depot arcs, i.e., arcs having a depot copy as an end vertex. Moreover, for 2-opt we see in Section 4.2 that the evaluation of the so-called multi-depot threshold can be simplified by using the same base term which only depends on the deleted arc (i, s_i) .

At this point, we would also like to mention that the 2-opt move can also be decomposed in a different way. If the one partial move is deleting (i, s_i) and inserting (s_i, s_j) , the other partial move is deleting (j, s_j) and inserting (s_j, i) . Anyway, this decomposition is also asymmetric and therefore also requires the distinction of two cases.

Finally, fixed-radius search can use the complete candidate list $\mathbf{N}(i)$ for every vertex *i*, i.e., $\sigma = |V| - 1$ as long as computer memory permits the storage of $\mathcal{O}(n^2)$ elements. Complete candidate lists require a $\mathcal{O}(n^2 \log n)$ preprocessing, where each candidate list is sorted. Using complete candidate lists ensures that fixed-radius search terminates in a local optimum.

3.2.3. Dynamic-Radius Search

As first discussed by Irnich et al. (2006), the gain criterion can be sharpened if a lower bound γ on the best gain g^* is known, e.g., because an improving move has already been found. In this case, the search effort for further improving moves can be potentially lightened by reducing the search radius ρ . The theoretical foundation is the following corollary of Theorem 1.

Corollary 1. (Irnich et al. 2006; p. 2411) If a sequence of k numbers $(g_i)_{i=\{1,...,k\}}$ has a sum greater than g, i.e., $\sum_{i=1}^{k} g_i > g$, then there exists a cyclic permutation π of these numbers for which every partial sum fulfills $\sum_{i=1}^{\ell} g_{\pi(i)} > \ell/k g$ for all $1 \le \ell \le k$.

For a move μ with k partial moves, it means that at the first level the radius can be reduced from ρ^1 to $\rho^1 - \gamma/k$, and at the second level from ρ^2 to $\rho^2 - 2\gamma/k$, etc. For the 2-opt move and its decomposition discussed above, the gain criterion improving condition (1) becomes

$$c_{i,s_i} - c_{ij} > \gamma/2 \quad \text{or} \quad c_{j,s_j} - c_{s_i,s_j} > \gamma/2 \tag{2}$$

with the corresponding radius conditions

$$c_{ij} < \rho \quad \text{with} \quad \rho = c_{i,s_i} - \gamma/2$$
(3a)

and
$$c_{s_i,s_j} < \rho$$
 with $\rho = c_{j,s_j} - \gamma/2.$ (3b)

Comparing (1) and (2) for the deleted arc (i, s_i) , the initial radius, i.e., the one used in fixed-radius search, is $\rho^0 = c_{i,s_i}$ (the superscript 0 corresponds to $\gamma = 0$). The radius that results when the lower bound is exact, i.e., $\gamma = g^*$, is $\rho^* = c_{i,s_i} - g^*/2$. Depending on the previously found improving moves, the radius ρ

actually used in the inner loop of Algorithm 2 is between ρ^* and ρ^0 . Therefore, it is always at least as sharp as fixed-radius search. In the following, we call a radius search that is based on the sharpened gain criterion a dynamic-radius search.



Figure 5: Dynamic-radius search for **2-opt** where the first deleted arc is (4, 15), neighbors in $\mathbf{N}(4)$ and $\mathbf{N}(15)$ inspected with respect to threshold ρ^0 are highlighted. The selected move is an intra-depot (because $d_4 = d_{12} = d$) and inter-tour move. It deletes (12, 44) and inserts (4, 12) and (15, 44). At any time of the search, we have $\rho^* \leq \rho \leq \rho^0$ anywhere within the double arrow of width $g^*/2$. Some specific numbers are $c_{i,s_i} = c_{4,15} = 18.87$, $g^* = 17.28$, and $\rho^* = 18.87 - 17.28/2 = 10.23$ such that customer 40 is never tested whereas customer 13 is not if we already have identified a large enough gain.

Figure 5 depicts part of an MDVRP solution in which we find an improving intra-depot 2-opt inter-tour move as seen in Figure 1b. It deletes arcs (4, 15) and (12, 44), inserts arcs (4, 12) and (15, 44), and inverts the segments (15, 37, 17, d) and (d, 12). Let us assume the deleted arc being tested is $(i, s_i) = (4, 15)$. The concentric circles indicate which of the neighbors $\mathbf{N}(4)$ and $\mathbf{N}(15)$ must be respectively evaluated for the radii ρ^0 and ρ^* . The neighbors are color shaded in one or even two colors if they appear in the area of the circles with radius ρ^0 . As $12 \in \mathbf{N}(4)$ and $44 \in \mathbf{N}(15)$ within radius ρ^0 , the depicted 2-opt move is found twice in fixed-radius search, once with the inserted arc (4, 12) as a neighbor of i = 4 and once more with the inserted arc (15, 44) as a neighbor of $s_i = 15$. In fact, we evaluate it another two times when looking at the deleted arc (12, 44) and its cost $c_{12,44}$ used for the threshold. It is however likely that the two different deleted arcs have quite different costs which already gives a first explanation for a smaller observed redundancy than the aforementioned factor of four.

Dynamic-radius search can reduce the observed redundancy even further although it is not as direct. Imagine this particular move yields the best gain g^* , the threshold ρ being used at any given time in the exploration is then anywhere in the interval $[\rho^*, \rho^0]$ depending on the value γ . It is tempting to look at the smaller overlapping area of the circles but this is mostly irrelevant since moves evaluated with vertices in this area are not comparable, e.g., inserted arcs (4, 17) and (15, 17) do not lead to the same move. The reduced observed redundancy simply comes from the smaller radius which potentially contains much less vertices to evaluate. In the most optimistic scenario, testing another deleted arc (i, s_i) with a larger cost $c_{i,s_i} \geq \gamma/2$, the threshold ρ becomes zero or even negative, thus implying that no neighbors must be analyzed at all. The consequence of this is that the reduced observed redundancy is very tangible although unpredictable as it depends on the loop construction and the observed gain. We stress again the conservative nature of criterion (2) and take a look at when redundancy is maximal. For fixed-radius search, it occurs when $c_{i,s_i} = c_{j,s_j}$. Many more conditions need to align for maximal redundancy in dynamic-radius search, that is, it occurs when $g^* = \varepsilon$ is rather small and the arc cost of all four arcs is almost identical, e.g., $c_{i,s_i} = c_{j,s_j} = c_{ij} = c_{s_i,s_j} + \epsilon$ which leads to $\rho^0 = \rho^* + \epsilon/2$. These observations are independent of how one decomposes the move into partial moves or how one implements the inner loops.

Finally, we close with a geometric interpretation of dynamic-radius search and an open question for future research. Every unit reduction of the radius ρ reduces the area of the admissible neighbors quadratically by definition of a circle. Any radius $\rho > c_{15,44}$ would suffice for $44 \in \mathbf{N}(15)$ to qualify for the gain criterion. Note that such a radius is smaller than ρ^* except when the deleted edges are of the same length. The question is therefore: Is there a way to predict the smallest but sufficiently large radius ensuring that a move with maximum gain is identified?

3.3. Granular Search

We briefly review granular search to clarify its relationship to radius search. Granular search is a neighborhood search and exploration technique that has been introduced via granular tabu search implementations for VRPs (Toth and Vigo 2003; Escobar et al. 2014; Schneider et al. 2017). The idea is an extension of bounded candidate-lists, as explained in Section 3.2.1, where arcs to be inserted are still ordered but now stored in a single global list, instead of one candidate list per vertex. This global list, denoted \mathcal{A} , is called the generator-arc list. Granular search explores the given neighborhood by considering only those moves where one specific inserted arc, the generator arc, is in \mathcal{A} . In case of the 2-opt neighborhood, the generator arc can be defined as the arc (i, j). Since (i, j) completely determines the 2-opt move, neighborhood exploration boils down to loop over $(i, j) \in \mathcal{A}$ and to implicitly construct and evaluate the feasibility and gain of the associated neighbor solution (the latter is possible in constant time for the MDVRP). Note that for a 2-opt move and the second inserted arc (s_i, s_j) it is not required that $(s_i, s_j) \in \mathcal{A}$ holds. In the same vein, redundancy occurs if both arcs (i, j) and (s_i, s_j) are present in \mathcal{A} .

Granular search can only guarantee that an improving move is found whenever one exists if the generatorarc list comprises all feasible arcs. Typically, \mathcal{A} is a heavily truncated list so that granular search only explores heuristically. Moreover, in the tabu search context, for which granular search was invented, one is interested in a best but not necessarily improving move. A generator-arc list is well-suited for this task.

The granularity aspect of granular search comes from a partitioning of the generator-arc list, that is, $\mathcal{A} = \bigcup_{u=1}^{U} \mathcal{A}_u$, in which the sorted order of the arcs in \mathcal{A} is also maintained. If no improving move is found with \mathcal{A}_u , the exploration is pursued in \mathcal{A}_{u+1} for all $1 \leq u \leq U - 1$, see Algorithm 3.

The speed of granular search also results from the increased flexibility of maintaining the generator arcs in any order. Indeed, the above mentioned implementations for VRPs exploit that a better selection of generator arcs (i.e., the choice of \mathcal{A}) and ordering of generator arcs (i.e., their assignment to $\mathcal{A}_1, \mathcal{A}_2, \ldots, \mathcal{A}_U$) often results when pseudo-costs are used instead of the given routing costs c_{ij} .

3.4. Comparison and Remarks

The synopsis of lexicographic, radius, and granular search provided by Figure 4 as well as the above discussion highlights the different ideas behind the neighborhood exploration techniques: Lexicographic search primarily prunes the search tree on the basis of local feasibility and is therefore well suited for strongly constrained VRPs. In contrast, radius search primarily prunes on the basis of the gain criterion and can be expected to be less effective for strongly constrained VRPs, because here many infeasible moves look promising from a gain's perspective. Later results will however show that for loosely constrained VRPs like the MDVRP, radius search typically outperforms lexicographic search. Granular search prunes on the basis of a heuristic pre-selection of generator arcs, which can either be based on feasibility or cost criteria or a mix of both.

In all cases, we have the freedom to decide which comes first: the feasibility test or the gain computation. One should decide this by comparing the computational effort and effectiveness of both tasks.

We have presented lexicographic and radius search as *best improvement* exploration strategies. Even though both can be prematurely stopped when any improving and feasible solution has been found (this is

first improvement), the idea of dynamic-radius search is to not stop but to explicitly exploit previously found improving solutions that lead to reduced search radii. By design, granular search follows a best improvement strategy per generator-arc list subset \mathcal{A}_u . First improvement would only really make sense if U = 1 in which case the initial sorting of the arcs is even more crucial.

4. Dynamic-Radius Search for Inter-Depot 2-Opt and 2-Opt* Moves

Using dynamic-radius search in the multi-depot environment, one must realize that the radii ρ^0 and ρ^* , as defined in (3) for the intra-depot 2-opt cases in Figures 1a and 1b, do not account for the additional source/sink depot swap cost that may occur. Using these radii, we are no longer guaranteed to find the remaining improving moves, let alone provide the best gain available, unless we find a way to correctly consider this otherwise neglected cost prior to the radius breakpoint, that is, before testing neighbor vertices in an inner loop. For this purpose, we introduce a *correction term* τ on the standard threshold ρ which bounds from above the potential gain any source or sink depot swap can produce with respect to the current solution R and the explored neighborhood.

Let us further clarify the need for a correction term by presenting some facts regarding the inter-depot 2-opt^{*} and 2-opt moves as broken down in the cases of Figures 2a-2f and 3a-3f. Note that we shorthand the latter expression to 'case xy' and even only use the index 'y' in mathematical formulas of a given neighborhood.

Table 1 summarizes all cases by listing their corresponding gains. The headers (standard, source, sink, and exception) reflect the left-/right-side structure and refer, respectively, to move-defining, exchanged arcs of the *standard* cases 1c, 2a and 2b of 2-opt* and *standard* cases 1b, 3a and 3b of 2-opt, the arcs exchanged at the *source* (resp. roles preserved) and the *sink* (resp. roles interchanged) to repair the otherwise infeasible depot assignment, and the added arc of the *exception* cases that differs from the standard cases. It is obvious that the gain computations differ significantly from one another and especially compared to the standard ones.

2-opt^*	Fig.	Gain g from edges			
Move type		standard	source	sink	exception
Intra-depot	1c	$c_{i,s_i} - c_{j,s_i} + c_{j,s_j} - c_{i,s_j}$			
Inter-depot	2a	$c_{i,s_i} - c_{j,s_i} + c_{j,s_j} - c_{i,s_j}$	$+c_{d_i,f_i}-c_{d_i,f_i}+c_{d_i,f_i}-c_{d_i,f_i}$		
	$2\mathrm{b}$	$c_{i,s_i} - c_{j,s_i} + c_{j,s_j} - c_{i,s_j}$		$+c_{l_i,d_i}-c_{l_i,d_j}+c_{l_j,d_j}-c_{l_j,d_i}$	
	2c	$c_{i,s_i} - c_{j,s_i} + c_{j,s_j}$	$+c_{d_j,f_j}-c_{d_i,f_j}$		$-c_{d_j,s_j}$
	2d	$c_{i,s_i} - c_{j,s_i} + c_{j,s_j}$		$+c_{l_i,d_i}-c_{l_i,d_j}$	$-c_{i,d_i}$
	$2\mathrm{e}$	c_{i,s_i} $+c_{j,s_j}-c_{i,s_j}$	$+c_{d_i,f_i}-c_{d_j,f_i}$	-	$-c_{d_i,s_i}$
	2f	$c_{i,s_i} + c_{j,s_j} - c_{i,s_j}$	-	$+c_{l_j,d_j}-c_{l_j,d_i}$	$-c_{j,d_j}$

Table 1: Gains of the **2-opt**^{*} and **2-opt** moves.

(a) The seven different cases of **2-opt***.

2-opt	Fig.	Gain g from edges			
Move type		standard	roles preserved	roles interchanged	exception
Intra-tour	1a	$c_{i,s_i} - c_{ij} + c_{j,s_j} - c_{s_i,s_j}$			
Intra-depot Inter-depot	1b 3a	$\begin{array}{ccc} c_{i,s_i} - c_{ij} + c_{j,s_j} - c_{s_i,s_j} \\ c_{i} - c_{ij} + c_{i} - c_{ij} \end{array}$	$+c_1$, $-c_2$, $+c_3$, c_3 , $-c_4$, c_3		
inter-depot	3b	$\begin{array}{cccccccccccccccccccccccccccccccccccc$	$+c_{l_i,d_i}-c_{d_j,l_i}+c_{d_j,f_j}-c_{f_j,d_i}$	$+c_{d_{i},f_{i}}-c_{d_{i},f_{i}}+c_{l_{i},d_{i}}-c_{l_{i},d_{i}}$	
	3c	$c_{i,s_i} - c_{ij} + c_{j,s_j}$	$+c_{d_j,f_j}-c_{f_j,d_i}$		$-c_{d_j,s_j}$
	3d	$c_{i,s_i} \qquad + c_{j,s_j} - c_{s_i,s_j}$		$+c_{l_j,d_j}\!-\!c_{l_j,d_i}$	$-c_{d_j,j}$
	3e	$c_{i,s_i} \qquad +c_{j,s_j} - c_{s_i,s_j}$	$+c_{l_i,d_i}\!-\!c_{d_j,l_i}$		$-c_{i,d_i}$
	3f	$c_{i,s_i} - c_{ij} + c_{j,s_j}$		$+c_{d_i,f_i}-c_{d_j,f_i}$	$-c_{s_i,d_i}$

(b) The eight different cases of **2-opt**.

In Table 2, the various cases that may arise are conditioned based on the known deleted arc (i, s_i) given by the outer loop and whose cost is central to the threshold computation. For each neighborhood, we have an *explicit* and an *implicit* column. This distinction is explained and utilized in the upcoming multi-depot threshold analysis. In order to correctly apply dynamic-radius search, we must match the multi-depot threshold with inner loops that are based on neighbor lists $\mathbf{N}(i)$ and $\mathbf{N}(s_i)$.

	2-opt* Figure 2		2-opt Figure 3	
Conditions on i and s_i	Explicit	Implicit	Explicit	Implicit
$i \neq d_i$ and $s_i \neq d_i$	a, b, d	е	a, b, e	f
$i = d_i$ and $s_i \neq d_i$	b, c, d		a, e	d
$i \neq d_i$ and $s_i = d_i$	a	e, f	b, c	f
$i = d_i$ and $s_i = d_i$	с	f	с	d

Table 2: Arising cases conditioned on the known deleted arc (i, s_i) .

Finally, the symbol ρ is reserved for the threshold given by the standard cases displayed in Figure 1a–1c. The latter appears as a common term (and therefore the lower bound) in all our thresholds. Each case 'y' indeed gives rise to a local threshold $\rho + \tau_y$ where τ_y is a correction term. We make the multi-depot threshold clear by expressing it as $\rho_{MD} = \rho + \tau$, where τ is a case-dependent expression contributing to the correction term. We now distinguish between the 2-opt* and the 2-opt neighborhoods

4.1. 2-Opt* Moves

Recall that Figures 1c and 2a-2f display the seven cases to handle for the 2-opt* neighborhood. Following the same line of arguments as explained in Section 3.2.3 for 2-opt, the radius for the standard intra-depot 2-opt* case 1c is given by

а

$$c_{j,s_i} < \rho \quad \text{with} \quad \rho = c_{i,s_i} - \gamma/2$$

$$\tag{4a}$$

nd
$$c_{i,s_j} < \rho$$
 with $\rho = c_{j,s_j} - \gamma/2.$ (4b)

Since dynamic-radius search exploits that the 2-opt^{*} is completely symmetric with respect to i and j, there is a single inner loop say on $j \in \mathbf{N}(s_i)$ for the known deleted arc (i, s_i) . Observe that in cases 2e and 2f, arc (j, s_i) is absent from the final move. This implies that its cost is irrelevant and cannot be subjected to a threshold. We sidestep this by observing that case 2c is symmetric in the affected routes to case 2e and likewise 2d to 2f. The idea is then that we must define a threshold in such a way that if an improving move with $g > \gamma$ exists in cases 2e or 2f, it is found by their symmetric counterpart. Our case-by-case analysis works as follows:

In the standard source case 2a, the arcs affected by the 2-opt* move are distinct from those needed to fix the depots. A new best gain is established by such a move if $g = [c_{i,s_i} - c_{j,s_i} + c_{d_i,f_i} - c_{d_j,f_i}] + [c_{j,s_j} - c_{i,s_j} + c_{d_j,f_j} - c_{d_i,f_j}] > \gamma$. As before, we break down the gain into two expressions according to the bracketed parts. This implies that $c_{j,s_i} < c_{i,s_i} + c_{d_j,f_i} - \gamma/2$. The point here is that the term c_{d_j,f_i} (depending on j) is unknown at the moment when (i, s_i) is deleted and the threshold ρ must be computed. We can however replace c_{d_j,f_i} by a lower bound over any depot reconnection yielding a radius large enough:

$$c_{j,s_i} < \rho + \tau_{\mathbf{a}}$$
 with $\tau_{\mathbf{a}} = c_{d_i,f_i} - \min_{d \in D} c_{d,f_i}$. (5a)

We find the interpretation of this correction term quite elegant because it goes in line with intuitive expectations of the multi-depot environment: If customer f_i is already attached to the nearest depot, then the right side simplifies to the original ρ value given by (4a). Otherwise, the radius takes into account the potential for swapping source depots as the cost difference between the current depot assignment and one of a nearest depot.

In the standard sink case 2b, the gain can be decomposed into $g = [c_{i,s_i} - c_{j,s_i} + c_{l_i,d_i} - c_{l_i,d_j}] + [c_{j,s_j} - c_{i,s_i} + c_{l_i,d_j} - c_{l_i,d_i}]$ so that the resulting radius is given by

$$c_{j,s_i} < \rho + \tau_{\rm b}$$
 with $\tau_{\rm b} = c_{l_i,d_i} - \min_{d \in D} c_{l_i,d}$, (5b)

where we note the similarity with the standard source case 2a that comes with exchanged roles of source and sink depot swap in the gain formula and the nearest-depot interpretation

In the exception case 2c, the gain is computed as $g = c_{i,s_i} - c_{j,s_i} + [c_{d_j,f_j} - c_{d_i,f_j}] + [c_{j,s_j} - c_{d_j,s_j}]$. Observe that arc (i, s_j) is omitted from the final move (see Table 1a), since it would otherwise be inserted and removed upon swapping source depots. It turns out that we cannot reasonably decompose the gain's components so that the gain criterion can be applied. Moreover, remember that we must ensure a move from case 2e can be found anyway. The following threshold bounds from above the gain seen in case 2c thus fulfilling the latter requirement:

$$c_{j,s_i} < \rho + \tau_c$$
 with $\tau_c = \max_{r \in R} \left[c_{d_r,f_r} - c_{d_i,f_r} + \max_{j \in N(r)} (c_{j,s_j} - c_{d_r,s_j}) \right] - \gamma/2,$ (5c)

where N(r) is the set of customers in route r. Moreover, we have a subtraction of $\gamma/2$. As we did not decompose the gain (into two independent parts) to test condition $g > \gamma$, the whole γ can be subtracted from the computed radius. With $-\gamma = -2 \cdot \gamma/2$ and the first half being already present in ρ (4), the validity of (5c) becomes clear. Fortunately, since the terms are consciously organized, intuition still answers the call. The first bracket describes some route r for which the first customer f_j would be closer to the source depot d_i than to the currently assigned source depot d_j , whereas the second bracket describes some customer j for which short-cutting from d_j to s_j is favorable.

The exception case 2d slightly differs from the former case 2c because of the different arc costs that are unknown in j and known in i. The gain $g = c_{i,s_i} - c_{j,s_i} + [c_{l_i,d_i} - c_{i,d_i}] + [c_{l_j,d_j} - c_{l_i,d_j}]$ still offers no acceptable decomposition, and we also want to cover case 2f. The resulting radius condition becomes

$$c_{j,s_i} < \rho + \tau_{\rm d}$$
 with $\tau_{\rm d} = [c_{l_i,d_i} - c_{i,d_i}] + \max_{r \in R} (c_{l_r,d_r} - c_{l_i,d_r}) - \gamma/2.$ (5d)

The difference compared to (5c) is that now there is no inner max-term over $j \in N(r)$.

In summary, the case-by-case analysis has led to four different case-dependent correction terms given by the equations (5a)-(5d). We still face the complication that the threshold must be computed when deleting arc (i, s_i) and before knowing which correction term to apply. However, we do know whether $i = d_i$ or $i \neq d_i$ as well as whether $s_i = d_i$ or $s_i \neq d_i$. Depending on these four possibilities, we can filter out which of the seven cases may happen (using Table 2). Accordingly, we define the final radius, tailored to the first deleted arc (i, s_i) , as the maximum of the corresponding radii.

A careful examination of the conditioned cases in Table 2 allows the nodes i and s_i to be treated independently in the final formula. For example, the term τ_c occurs only when $i = d_i$ and the term τ_a only when $i \neq d_i$ which is indeed irrespective of s_i . Judiciously collecting all terms results in an elegant convoluted threshold expressed with respect to the various correction terms (5a)–(5d):

$$\rho_{MD} = \rho + \max \left[\begin{cases} \tau_{\rm c} & \text{if } i = d_i \\ \tau_{\rm a} & \text{if } i \neq d_i \end{cases}, \begin{cases} 0 & \text{if } s_i = d_i \\ \max\{ \tau_{\rm b} , \tau_{\rm d} \} & \text{if } s_i \neq d_i \end{cases} \right].$$
(6)

In summary, this radius definition, which depends on the type of the first deleted arc (i, s_i) , covers all cases of 2-opt^{*}. The different correction terms added to ρ were precisely highlighted. The test $c_{s_i,j} < \rho_{MD}$ is clearly a relaxed radius condition compared to the standard case, but it allows us to exactly explore the 2-opt^{*} neighborhood including inter-depot moves.

4.2. 2-Opt Moves

Deriving the correction term for 2-opt is slightly more intricate yet we find very similar expressions. Let us again be supported by the broken down cases as depicted in Figure 3 together with their respective gains in Table 1b. Exception cases 3c and 3f can indeed be merged into a one-sided test because the arc costs are symmetric (same for 3e and 3d). The following correction terms cover the relevant cases 3a, 3b, 3c, and 3e:

$$\tau_{\mathbf{a}} = c_{l_i,d_i} - \min_{d \in D} c_{d,l_i} \tag{7a}$$

$$\tau_{\rm b} = c_{d_i, f_i} - \min_{d \in D} c_{d, f_i} \tag{7b}$$

$$\tau_{\rm c} = \max_{r \in R} \left[(c_{d_r, f_r} - c_{f_r, d_i}) + \max_{j \in N(r)} \left(c_{j, s_j} - c_{d_r, s_j} \right) \right] - \gamma/2 \tag{7c}$$

$$\tau_{\rm e} = c_{l_i, d_i} - c_{i, d_i} + \max_{r \in R} \left(c_{d_r, f_r} - c_{d_r, l_i} \right) - \gamma/2 \,. \tag{7e}$$

Using Table 2, the final radius is obtained by collecting the terms with respect to the deleted arc (i, s_i) and independently treating the vertices i and s_i similarly to 2-opt^{*} as

$$\rho_{MD} = \rho + \max \left[\begin{array}{cc} 0 & \text{if } i = d_i \\ \tau_{\rm b} & \text{if } i \neq d_i \end{array}, \begin{array}{c} \tau_{\rm c} & \text{if } s_i = d_i \\ \max\{ \tau_{\rm a} , \tau_{\rm e} \} & \text{if } s_i \neq d_i \end{array} \right].$$
(8)

5. Iterated Local Search

In this section, we describe the algorithmic details of our metaheuristic. We have designed it with simplicity in mind so that local search is the fundamental building block. A good pick in this respect is *iterated local search* (ILS, Lourenço et al. 2002) as it combines local search with a *perturbation* mechanism. Local optima are perturbed into new solutions so that local search can be applied repeatedly.

In our ILS, capacity constraints and duration constraints (if any) are handled as hard constraints such that feasibility of all routes is maintained starting from the construction heuristic to the perturbation and throughout the local search. The vehicle fleet-size limit however is artificially construed as a soft constraint by tolerating up to δm , $\delta \geq 1$, vehicles per depot in all components of the metaheuristic, i.e., construction, local search, and perturbation. We therefore speak of δ -fleet-size feasibility, where a solution R fulfills the instance requirements if $\delta = 1$. In addition, we introduce an intermediate *fleet-reduction operation* between local search and perturbation in an effort to recover fleet-size feasibility. This operation is allowed to fail in which case we do not obtain a local optimum.

We then perform a perturbation operation to induce relatively significant routing changes in the current solution. The main complication is to maintain δ -fleet-size feasibility which we handle with a fail-safe savings heuristic. This explains why we had to design the ILS with a somewhat more involved perturbation mechanism.

We describe the construction heuristic in Section 5.1, the local search in Section 5.2, the perturbation and fleet-reduction operations in Section 5.3, and we provide an overview and pseudo-code of the entire ILS in Section 5.4.

5.1. Construction Heuristic

Our construction procedure is based on the savings heuristic of Clarke and Wright (1964). The core idea remains to process an arc list sorted decreasingly by their potential saving, but the multi-depot aspect and the additional duration constraints are accounted for as follows. First, we draw uniformly distributed parameters $\zeta \in [0,2]$ and $\lambda \in [\zeta - 0.25, \zeta + 1.75]$. Then, we compute the saving of every arc (i,j) with respect to each depot $d \in D$ as $\sigma_{ij}^d = -\lambda c_{ij} + c_{di} + c_{jd} + \zeta(c_{di} - c_{jd})$, where λ and ζ influence the comparison between the cost of arc (i, j) and of connecting i and j to d. We reject those combinations where (d, i, j, d)is an infeasible route. To further randomize the procedure, for each arc (i, j), we take from these |D| depotspecific savings values $(\sigma_{ij}^d)_{d\in D}$ an arbitrary one that is non-negative, denoted σ_{ij} . Next, we sort these savings $(\sigma_{ij})_{(i,j)}$ decreasingly.

At the start, all customers $i \in N$ form separate segments (i). After computing and sorting the savings values $(\sigma_{ij})_{(i,j)}$, the main loop considers the associated arcs (i, j) one by one. If the vertices i and j are the last/first of their segments and the concatenation of their segments $(v \rightsquigarrow i)$ and $(j \rightsquigarrow w)$ gives a feasible

route $(d, v \rightsquigarrow i, j \rightsquigarrow w, d)$ for some depot $d \in D$, we join the segments together. At the end, when no more segments can be joined, each segment is finally assigned to the depot d that leads to the cheapest feasible route. Note that this type of depot assignment may lead to a solution R that is infeasible regarding the fleet-size constraints. We accept slightly infeasible solutions R if $|R \cap R_d| \leq \delta m$, where δ is the parameter described in the introduction of Section 5.

The above procedure is repeated with $\delta = 1.25$ and new random parameters ζ and λ until the constructed solution R is δ -fleet feasible.

5.2. Local Search

A reasonable local search-based ILS must use additional neighborhoods besides 2-opt and 2-opt^{*}. For the purpose of this study, we complement them with six other neighborhoods relocation, swap, string exchange (ordered and inverted), and Or-opt (ordered and inverted) as commonly defined (see, e.g., Aarts and Lenstra 1997; Funke et al. 2005). Figures 6 and 7 describe the general composition of Or-opt and string-exchange moves in ordered and inverted variants. Both neighborhoods restrict the length of their relocated chains to a length parameter L. Note that a relocation move is an Or-opt move with L = 1, and likewise a swap move is a string exchange with L = 1. We nevertheless implemented independent relocation- and swap-neighborhood exploration algorithms to benefit from the specialization, because for relocation and swap the distinction between ordered and inverted chains is irrelevant. In the following, we use L = 5 for Or-opt and string exchange unless stated otherwise.



Figure 6: Inter-depot **Or-opt**, $|s_i \rightsquigarrow j| \le L$.



Figure 7: Inter-depot string exchange, $|s_i \rightsquigarrow j| \le L$ and $|s_k \rightsquigarrow l| \le L$.

Relocation, swap, string exchange, and Or-opt are naturally compliant with the multi-depot environment because all these inter-depot moves result in routes that have matching source and sink depots.

Our local search is kept as simple as possible: All eight neighborhoods (we consider inverted and ordered Or-opt and string exchange as different neighborhoods) are explored in a cyclic fashion. The exploration is done with a best-improvement strategy. Whenever an improving move is returned from the neighborhood exploration, it is performed and we move to the next neighborhood. Local search terminates when all neighborhoods are explored without success, so that the solution R returned is always a local optimum with regards to all eight neighborhoods.

Note finally that δ -fleet feasibility can easily be maintained in the local search if the starting solution fulfills it. We must only use a feasible number of copies of the depots: Exactly $2\lfloor \delta m \rfloor$ copies are needed per depot $d \in D$ (two for each route, see Section 2).

Algorithm 4: Iterated Local Search (ILS)					
1	/ Initialization				
1 S	$\operatorname{avings}(R)$	// Section 5.1			
2 c	$2 count \leftarrow 0$				
1	// Main Loop				
3 fo	or n_{ILS} iterations do				
	// Local Search Phase				
4	$\text{LocalSearch}(R, \delta = 1.25)$	// Section 5.2			
5	if not fleet-size feasible then				
6	Fleet-reduction(R)	// Section 5.3			
7	if fleet-size feasible then				
8	LocalSearch $(R, \delta = 1.0)$	// Section 5.2			
	// Perturbation				
9	9 if not fleet-size feasible then				
10	$count \leftarrow count + 1$				
11	if $random() < 0.7^{count}$ then				
12	Perturbation (R)	// Section 5.3			
13	else				
14	Savings (R)	// Section 5.1			
15 else					
16	$count \leftarrow 0$				
17	$\operatorname{Perturbation}(R)$	// Section 5.3			

5.3. Perturbation and Fleet-Reduction

The perturbation consists of a multi-phase re-clustering on a δ -fleet-size feasible solution R. We first permute the routes randomly in the giant tour and then apply a circular shift on a random position. From this new customer sequence, routes are filled in order while satisfying resource consumption. This process is repeated 2 to 5 times (uniformly random). If the re-clustering fails to produce a suitable customer assignment, i.e., a δ -fleet feasible solution, a new solution R is constructed with the construction heuristic.

The purpose of the fleet-reduction operation is to transform a given solution R that is δ -fleet feasible into one that is 1-fleet-size feasible. As this is an NP-hard and sometimes practically difficult task, the fleet-reduction operation may terminate with a solution that is only partly improved regarding fleet-size feasibility. Note that in any case such an improvement likely comes at the cost of worsening the objective value.

The fleet-reduction operation tries to patch up the given solution by moving chains of customers from overused to underused depots and their routes. We reuse the exploration of 2-opt^{*} and Or-opt neighborhoods to find a chain inside a route belonging to an overused depot that can be moved to another depot at minimal cost. 2-opt^{*} and Or-opt moves are repeated until a fleet-size feasible solution is constructed or the search for such a feasible 2-opt^{*} or Or-opt move fails. This modified solution is then perturbed irrespective of its feasibility status, but we do save it as a local optimum if the fleet-reduction operation is successful.

5.4. Iterated Local Search

The general design is an ILS with a limit of n_{ILS} local search iterations as summarized by the pseudocode in Algorithm 4. The current solution is denoted by R and it is initialized by the result of the savings heuristic in Step 1. The counter *count* (initialized at zero in Step 2) keeps track of the number of consecutive iterations for which the fleet-reduction operation fails to produce a fleet-size feasible solution.

We perform up to two passes in each local descent (Steps 4 and 8). In the first pass (Step 4), the limit on the number of vehicles is relaxed. When a local optimum is reached, if said limit is satisfied, the algorithm moves on to the perturbation operation. Otherwise, the fleet-reduction operation tries to make the solution fleet-size feasible (Step 6) and, if so, the second local-search pass is performed, for which the strict fleet-size limit is imposed (Step 8). The perturbation mechanism (Steps 9 to 17) uses the savings heuristic as a fallback whenever the actual perturbation procedure described in Section 5.3 fails to produce a δ -fleet feasible solution.

6. Computational Results

The implementation of the ILS algorithm is written in C++ and compiled in 64-bit release mode under Microsoft Visual Studio 2015. The experiments are conducted on a Microsoft Windows 10 standard personal computer equipped with an Intel i7-6700 CPU clocked at 3.40 GHz and 16 GB of RAM. A single thread is allocated to each run.

Section 6.1 describes the benchmark instances used in this study. A comparison with the previous dynamic-radius search implementation follows in Section 6.2. We then take a look in Section 6.3 at an alternative way to recover inter-depot 2-opt and 2-opt* moves that foregoes both the challenging implementation and the correction term. The impact of the correction term is analyzed in Section 6.4 by evaluating its contribution under various usage scenarios.

6.1. Instances

We start our analysis of dynamic-radius neighborhood exploration techniques by reproducing a comparative assessment with lexicographic search using

• 560 CVRP instances [10 (seed) × 4 (load factor) × 14 (size)] from Irnich et al. (2006), where the load factor f is the quotient of the capacity Q and the average demand q_i of the customers $i \in N$. As a result, we see on average approximately f customers per route in good solutions.

We then look at how dynamic-radius search behaves with respect to multi-depots and the presence or not of the correction terms using the proposed ILS on commonly used MDVRP instances from the literature: • 33 MDVRP instances [p01-pr10] from Cordeau et al. (1997);

The CVRP instances from the previous work (Irnich et al. 2006) are available at https://logistik.bwl. uni-mainz.de/research/benchmarks/.

6.2. Improved Implementation

Dynamic-radius search has been re-implemented with a greater focus on low-level efficiency but we also gave due attention to lexicographic search. The reader may compare the previous results (Irnich et al. 2006; Figures 7 and 9) with Figures 8 and 9 referring to the same 560 CVRP instances. These figures compare lexicographic and radius search by analyzing two indicators: the acceleration factor describing the average ratio of computing times needed with lexicographic compared to dynamic-radius search (ratios are dimensionless) and average neighborhood-exploration times (in milliseconds [ms]). To understand what is depicted, one must know that the effectiveness of (dynamic-)radius search for the CVRP strongly depends on the *load factor f*. Therefore, the results presented in Figures 8 and 9 are exactly grouped by both number |N| of customers and load factor *f*.

Figure 8 shows the average acceleration factors. Comparing relative performance to lexicographic search with the 2-opt neighborhood (old factor 35 versus new factor 20) and the string-exchange neighborhood (900 versus 600) for the largest instances with |N| = 2,500 customers, it appears that the benefit of dynamic-radius search has slightly decreased. This is due to significant improvements in the lexicographic search implementation but also the fact that the increased maximum string length from L = 3 to 5 favors the lexicographic paradigm. Contrarily to the previous computational study in (Irnich et al. 2006), the trend lines also do not show any significant negative slope when increasing the number of customers. This is particularly striking on the string-exchange neighborhood and can be explained by the fact that we have incorporated feasibility pruning inside the string examinations of dynamic-radius search. Some neighborhoods such as relocation now also exploit the threshold whenever a new best gain is identified rather than only in the inner loop head.

Figure 9 depicts the average computation time for a single neighborhood exploration. We achieve a reduction in absolute computation time by at least a factor of 10 for every neighborhood including those



Figure 8: [CVRP, Irnich et al. (2006)] Average acceleration factor of dynamic-radius search over lexicographic search for various neighborhood operators and instances ranging from 300 to 2,500 in customer size.



Figure 9: [CVRP, Irnich et al. (2006)] Average neighborhood-exploration time (in milliseconds) of dynamic-radius search for various neighborhood operators and instances ranging from 300 to 2,500 in customer size.

with larger maximal string lengths L. We believe that technological progress since 2006 cannot singlehandedly explain this improvement as evidenced by the average search times of the swap neighborhood (70 milliseconds [ms] versus 4 ms) and string-exchange neighborhood (350 ms versus 9 ms) for instances with |N| = 2,500 customers. We point here to the aforementioned transition to an explicit double inner-loop design which allows finer treatment of depot arcs, see Section 3.2.2.

As a last note, we rectify a mistake in the pseudo-code (Irnich et al. 2006; Algorithm 8, Line 6) of the swap neighborhood. Making abstraction of the different nomenclature, it should read as "LET $B_1 = (c_{v_1,t_1} + c_{t_1,w_1})/2 - G^*/4$ " instead of **2** in the last denominator.

Figure 10 shows how this behavior remains consistent on the benchmark MDVRP instances of Cordeau et al. (1997). In particular, we see that despite being considerably smaller than 2,500, the largest instance (#23 with 360 customers and 9 depots) benefits comparatively more from dynamic-radius search. We also use the same plot symbol as for load factor f = 25, since it is usually does not surpass 10 in these instances.



Figure 10: [MDVRP, Cordeau et al. (1997)] Average acceleration factor of dynamic-radius search over lexicographic search for various neighborhood operators and instances ranging from 48 to 360 in customer size.

6.3. Inter-Depot Moves via String Exchange and Or-Opt

In this section, we exploit that all inter-depot 2-opt^{*} and 2-opt moves are specific string exchange and Or-opt moves on the giant tour. For example, 2-opt^{*} moves (Figure 2) can be reproduced by ordered variants of the string exchange (standard cases 2a, 2b) and Or-opt (exception cases 2c, 2d, 2e, 2f) moves. In particular, one can see that the string $f_i \sim i$ in Figure 2e (2-opt^{*}) can have any length and corresponds to the string $s_i \sim j$ in Figure 6a (ordered Or-opt). With respect to 2-opt moves (Figure 3), inverted variants of the string exchange and Or-opt respectively reproduce standard and exception cases.

Accounting for inter-depot 2-opt^{*} and 2-opt moves in dynamic-radius search implies that two algorithmic tasks have to be fulfilled: the computation of correction terms and obviously the implementation of the actual inter-depot moves. The exception correction terms are likely to be larger than their standard counterparts to account for worst-case scenarios in unknown customer j. Fortunately, exceptions occur only sporadically, see Table 2. With this in mind, we discuss three alternatives to cope with exhaustively testing for inter-depot moves.

First, it is possible to duplicate the 2-opt^{*} and 2-opt methods and specialize these copies to account for specific inter-depot cases. In this fashion, we inevitably face redundant move tests and therefore have an overall slower method. Moreover, it is a cumbersome implementation for which one must indeed implement inter-depot cases.

Second, specializing the string exchange and Or-opt neighborhoods to test for specific depot cases is even more cumbersome (number and complexity of loop blocks) and slower, since we must also pay for the overhead of these richer neighborhoods.

Third, we can forbid inter-depot 2-opt and 2-opt^{*} moves and herewith get rid of all the correction terms as well as the repair operations. Instead, we allow arbitrary string lengths, i.e., $L = \infty$ in the string exchange and Or-opt neighborhood exploration.

We have implemented this third alternative and tested how well string exchange and Or-opt scale with this length increase. Obviously, the time of neighborhood exploration increases with an unbounded length $L = \infty$ but it is limited by the longest route in the candidate solution. The results of the comparison between the maximum string length $L = \infty$ and L = 5 are shown in Figure 11. We ultimately observe a factor of at least 1.5 which reaches 4 on instances with load factor f = 100.





Figure 11: [CVRP, Irnich et al. (2006)] Dynamic-radius search computing time comparison of allowed length $L = \infty$ over L = 5 in string exchange and Or-opt neighborhood explorations for instances ranging from 300 to 2,500 in customer size.

The reader may expect now that we present a direct comparison of the implementations of 2-opt and 2-opt^{*} using the correction terms of Section 4 versus Or-opt and string exchange with unlimited string length $L = \infty$. Such a comparison would reveal that already longer neighborhood exploration times of Or-opt and string exchange (see Figure 9) must be compounded with the observed factors of Figure 11. However, by allowing arbitrary string lengths, we do not only recover all inter-depot 2-opt^{*} and 2-opt moves, but we also enrich the local optima space: additional improving Or-opt and string exchange moves that do not represent 2-opt or 2-opt^{*} moves are found. Hence, such a direct comparison considering relative computation times is an oversimplification. We therefore omit further analyses.

Finally, we arrive at the conclusion that the simplest possibility is to rely on string exchange and Oropt neighborhoods to produce inter-depot 2-opt and 2-opt* moves. However, even if this implementation shortcut is functional, it does not compete with a full-fledged inter-depot adaptation for the 2-opt* and 2-opt neighborhoods.

6.4. Inter-Depot Moves and the Correction Term

In this section, we investigate the central hypotheses that we already discussed in Section 1 as a main motivation for our work:

Hyp1: Allowing inter-depot moves helps finding better local optima on average.

Hyp2: Allowing inter-depot moves helps finding better *best solutions* on average.

To this end, we computationally compare five implementations of ILS that differ in the way in which the 2opt and 2-opt* neighborhoods are explored, i.e., the only neighborhoods that implement inter-depot moves. More precisely, we examine how the following five options *With*, *Without*, *Last Resort*, *Random*, and *Forbid* perform against each other by comparing the respective relative gap measures:

- (i): Option *With* (W) uses the correction term (6) and thus guarantees that a move with provably maximum gain is identified.
- (ii): Option *Without* (WO) uses the simple radii of the intra-depot cases, i.e., (3) and (4). It loses the guarantee but still requires the whole machinery of the depot repair operation.
- (iii): Option Last Resort (LR) uses option Without until no improving move is found at which point it switches to option With.
- (iv): Option *Forbid* (F) literally forbids inter-depot moves from happening thus eliminating the need for the repair implementation.
- (v): Option Random (R) uniformly picks among those four options at every local search iteration.

It is clear that, for a fixed number n_{ILS} of iterations, the five options differ in the computation time consumed by the respective ILS implementation. Formally, we have two vectors of result data $[z_{no}^{I}]$ and $[t_{no}^{I}]$, for $n \in \{1, \ldots, n_{\text{ILS}}\}$ and o denotes the option. The vectors respectively give the local optima we have found at each iteration and the total time it takes since the beginning. Figure 12 illustrates how we normalize the computation log of the different options on a specific instance I in order to be able to summarize the benchmark results. On the left, we have iterations on the abscissa and objective values on the ordinate. Normalizing the ordinate is fairly intuitive since we simply compute relative gaps with respect to the instance's best known solution \bar{z}^{I} :

$$\gamma^{I}(z_{no}^{I}) = 100(z_{no}^{I} - \bar{z}^{I})/\bar{z}^{I}.$$

Normalizing the abscissa in the range [0, 1] allows to compare where we stand for each option with respect to a computation time budget. Since every option impacts the search time of neighborhood explorations, the cumulative time difference across the options can become significant. The circle markers on the tail of the best local optima plots for each option indicate the iteration for which the computation time is the same as that of the fastest option, say τ^{I} (in Figure 12a, the ILS with option WO is the fastest). One can then find the iterations with matching points in time for every option. We can see the normalized time axis as pulling each plot line by its circle marker rightwards thus more or less keeping their shape. Next, we define for each option o the function $h_o^I : [0,1] \mapsto \mathbb{R}_+$ giving the relative gaps in percent of the average local optima up to a percent time, given by

$$h_o^I(t\,\%) = \frac{1}{\nu} \sum_{n=1}^{\nu} \gamma^I(z_{no}^I) \tag{9}$$

where $\nu = \arg \max_{n \in \{1, \dots, n_{ILS}\}} t_{no}^I \leq t \% \cdot \tau^I$.



Figure 12: [MDVRP, Cordeau et al. (1997)] Normalizing times and average relative gaps of five options for a specific MDVRP instance *I*.

From Figure 12b we see that the function h_o^I is generally not monotone. The trend of the average local optima is decreasing, which can be attributed to the design of the metaheuristic, e.g., the acceptance criterion and the perturbation that reuses solution parts. If randomized initial solutions were used at every iteration, we would by the law of large numbers expect rather horizontal lines in Figure 12 after sufficiently many iterations, i.e., larger values of t %.

Overall Average Quality of Local Optima and Hypothesis Hyp1. The normalization allows us to aggregate the functions h_o^I over several MDVRP instances for a given option o. Recall that hypothesis Hyp1 refers to the average quality of local minima. Formally, let \mathcal{I} be a finite set of instances $I \in \mathcal{I}$. The aggregated function is

$$h_o(t\,\%) = \frac{1}{|\mathcal{I}|} \sum_{I \in \mathcal{I}} h_o^I(t\,\%),\tag{10}$$

and the upper part of Figure 13 shows the result for the benchmark instances described in Section 6.1. We can see that overall averages are worst in ILS with options *Forbid* and *Random*, while with options *With*, *Without*, and *Last Resort* the overall averages are significantly better. Moreover, differences between the latter three ILS are minor. The point is that options *With*, *Without*, and *Last Resort* are those that systematically allow

inter-depot moves. This confirms our leading hypothesis Hyp1, i.e., the ability to identify better gains by allowing inter-depot moves pays off with regard to the average quality of local minima.

Overall Average Quality of Best Solutions and Hypothesis Hyp2. To analyze the average quality of best solutions found, we introduce functions f_o^I and f_o in a similar manner as (9) and (10):

$$f_o^I(t\,\%) = \min_{n \in \{1,...,\nu\}} \gamma(z_{no}^I) \quad \text{and} \quad f_o(t\,\%) = \frac{1}{|\mathcal{I}|} \sum_{I \in \mathcal{I}} f_o^I(t\,\%), \tag{11}$$

where again $\nu = \arg \max_{n \in \{1, \dots, n_{ILS}\}} t_{no}^{I} \leq t \% \cdot \tau^{I}$. The functions f_{o}^{I} and herewith also f_{o} for the incumbent relative gap are monotonically non-increasing for all options o by definition.

We can see in the lower part of Figure 13 that the best local optima plots are tightly bunched together fairly early on. For the sake of better visibility, we provide explicit coordinates in Figures 12 and 13, even though functions h_o^I and f_o^I are defined for continuous values (see (9) and (11)). These explicit coordinates are found by the same time percent within some small tolerance across all options. Furthermore, we discard all coordinates at t% if $f_o^I(t\%)$ is equal (again within some tolerance) to the value we found in the previous time percent for all options. This means that the changes in average relative gap values do not appear where the minimum functions f_o^I are plateaus. The rest of each function is filled by interpolation. Deriving explicit coordinates for the aggregated functions follows the same treatment.

There are several results that were completely unexpected for us:

- The difference between the ILS equipped with the five different options is very small. It means that although the average quality of local optima is improved with 2-opt and 2-opt* moves that allow a depot swap, average quality does not translate into better overall quality. Hypothesis Hyp2 is not confirmed but rather proven false.
- The ILS with option With (W) is performing worst, while the ILS with option Without (WO) is performing best. Even though options W and WO allow inter-depot 2-opt and 2-opt* moves, it is hard to make an option recommendation when compared to option Forbidden (F) which is somewhere in between. The relationship shown with f_o is really inconsistent with the relationship shown with h_o .

Retrospectively, we can offer the following possible explanation for the unexpected outcome: For the overall best solution, we only need to find an excellent local optimum one time. It seems to be more important to inspect more local optima (even if not ensured that they are true optima) than proving that the solutions cannot be improved with the given neighborhoods. The explanation goes partly in line with the computation times that we can see in Figure 12a. The ranking by decreasing speed of the options is WO, F, LR, R, and W, while the ranking by increasing best local optima WO, F and LR, R and W (the latter pairs are performing equally well). This outcome is very consistent for values of t% between 0.2 and 1.0.

7. Conclusion and Outlook

In this paper, we have revisited radius search, an effective neighborhood exploration technique, which distinguishes itself from other techniques such as lexicographic search by the way a neighborhood is explored: A lexicographic search prunes an exploration branch whenever a local infeasibility is observed. Dynamic-radius search is closer in spirit to the optimization paradigm, since the pruning is based on coefficients of the objective function, that is, the threshold bound is a function of the best gain found at any given time.

We have extended previous works on radius search to the multi-depot vehicle routing problem including capacity and tour-duration constraints. The focus of our research is on the two fundamental neighborhoods 2-opt and 2-opt*, more specifically, their exploration which includes standard intra-depot moves but also interdepot moves. Case-dependent correction terms to be added to the otherwise incorrect standard search radii have been derived. Dynamic-radius search equipped with this modified pruning criterion allows identifying a provably best-improving move, either intra-depot or inter-depot, with little additional computational effort. Such inter-depot moves can also be found without the repair machinery but we found that it is a computationally expensive neighborhood exploration that results from the need to increase the allowed string length in string exchange and Or-opt neighborhoods.



Figure 13: [MDVRP, Cordeau et al. (1997)] Average and minimum relative gaps evolution over normalized axes averaged across benchmark instances for five options.

In comparison to lexicographic search, speedups of factors of 100 and more are observed for 2-opt, 2-opt^{*}, Or-opt, swap, and string-exchange neighborhoods. Because of the improved implementation, these factors are sometimes much better than in the original study of Irnich et al. (2006) on the capacitated vehicle routing problem. Moreover, the adaptation to the multi-depot vehicle routing problem benefits in the same orders of magnitude.

While there certainly are some fancier metaheuristics out there, we believe our basic iterated-local search implementation is legitimate enough with respect to behavior one would observe in state-of-theart metaheuristics. In particular, we conjectured and have later shown that identifying moves with better gains lead to better local optima on average. We also conjectured that better average local optima translate to better incumbent solutions in local search-based metaheuristics like the iterated-local search that we implemented. This behavior unfortunately does not hold true, proving our second hypothesis as false. Even if false, this is an important insight that has not been reported so explicitly in the literature.

Clearly, the overarching metastrategy does strongly influence the empirical behavior we illustrate in our analysis. This means that the mataheuristic's components like the initial solution construction, the acceptance criterion, and the depth of local search including the choice of neighborhoods all impact the results. However, there is absolutely no reason to believe that switching to another search paradigm like lexicographic would yield any difference. (The metaheuristic would just become slower and therefore loose competitiveness.)

We can think of the following research paths. First, the way the threshold is constructed is particularly interesting because it relies on cost upper bounds rather than the actual cost. Tackling alternative vehicle routing problem variants in which the objective function is not exactly a sum of arc costs such as time-dependent travel costs then becomes possible. Second, for asymmetric problems, the redundancy in the exploration does not mean we get to test asymmetric arc costs for free. Indeed, the worst-case factor is eight rather than four which fortunately is still prone to significant empirical reduction. Finally, we venture that machine learning may help answer the question we raised at the end of Section 3.2.3 concerning the

prediction of the smallest but sufficiently large radius ensuring that a move with maximum gain is identified.

Acknowledgement

This research was funded by the Deutsche Forschungsgemeinschaft (DFG) under grant no. IR 122/7-1 and IR 122/7-2.

References

- Emile H. L. Aarts and Jan Karel Lenstra, editors. Local Search in Combinatorial Optimization. Wiley-Interscience series in discrete mathematics and optimization. John Wiley & Sons, New York, NY, USA, 1997.
- Mandell Bellmore and Saman Hong. Transformation of multisalesmen problem to the standard traveling salesman problem. Journal of the Association for Computing Machinery, 21:500–504, 1974. doi:10.1145/321832.321847.
- Jon Jouis Bentley. Fast algorithms for geometric traveling salesman problems. ORSA Journal on Computing, 4(4):387–411, 1992. doi:10.1287/ijoc.4.4.387.
- Geoff Clarke and J. W. Wright. Scheduling of vehicles from a central depot to a number of delivery points. *Operations Research*, 12(4):568–581, 1964. doi:10.1287/opre.12.4.568.
- Jean-François Cordeau, Michel Gendreau, and Gilbert Laporte. A tabu search heuristic for periodic and multi-depot vehicle routing problems. Networks, 30(2):105–119, 1997. doi:10.1002/(SICI)1097-0037(199709)30:2<105::AID-NET5>3.0.CO;2-G.
 G. A. Croes. A method for solving traveling-salesman problems. Operations Research, 6(6):791–812, 1958.
- Guy Desaulniers, Oli B. G. Madsen, and Stefan Røpke. The vehicle routing problem with time windows. In Paolo Toth and Daniele Vigo, editors, *Vehicle Routing*, chapter 5, pages 119–159. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2014. doi:10.1137/1.9781611973594.ch5.
- John Willmer Escobar, Rodrigo Linfati, Paolo Toth, and Maria G. Baldoquin. A hybrid granular tabu search algorithm for the multi-depot vehicle routing problem. *Journal of Heuristics*, 20(5):483–509, 2014. doi:10.1007/s10732-014-9247-0.
- Birger Funke, Tore Grünert, and Stefan Irnich. A note on single alternating cycle neighborhoods for the TSP. Journal of Heuristics, 11(2):135–146, 2005. doi:10.1007/s10732-005-0713-6.
- Keld Helsgaun. An effective implementation of the Lin-Kernighan traveling salesman heuristic. European Journal of Operational Research, 126(1):106–130, 2000. doi:10.1016/s0377-2217(99)00284-2.
- Holger H. Hoos and Thomas Stützle. Stochastic Local Search: Foundations and Applications. The Morgan Kaufmann Series in Artifical Intelligence. Elsevier, San Francisco, CA, USA, 2005.
- Stefan Irnich. Resource extension functions: properties, inversion, and generalization to segments. OR Spectrum, 30(1):113–148, 2008a. doi:10.1007/s00291-007-0083-6.
- Stefan Irnich. A unified modeling and solution framework for vehicle routing and local search-based metaheuristics. INFORMS Journal on Computing, 20(2):270–287, 2008b. doi:10.1287/ijoc.1070.0239.
- Stefan Irnich, Birger Funke, and Tore Grünert. Sequential search and its application to vehicle-routing problems. Computers & Operations Research, 33(8):2405–2429, 2006. doi:10.1016/j.cor.2005.02.020.
- David Stifler Johnson and Lyle A. McGeoch. The traveling salesman problem: A case study in local optimization. In Emile H. L. Aarts and Jan Karel Lenstra, editors, *Local Search in Combinatorial Optimization*, chapter 8, pages 215–310. John Wiley & Sons, 1997.
- Shen Lin. Computer solutions of the traveling salesman problem. The Bell System Technical Journal, 44(10):2245–2269, 1965. doi:10.1002/j.1538-7305.1965.tb04146.x.
- Shen Lin and Brian Wilson Kernighan. An effective heuristic algorithm for the traveling-salesman problem. Operations Research, 21(2):498–516, 1973.
- Helena R. Lourenço, Olivier C. Martin, and Thomas Stützle. Iterated local search. In Fred Glover and Gary A. Kochenberger, editors, *Handbook of Metaheuristics*, pages 321–353. Kluwer Academic Publishers, Norwell, MA, USA, 2002. doi:10.1007/0-306-48056-5 11.
- Olivier Martin, Steve W. Otto, and Edward William Felten. Large-step Markov chains for the TSP incorporating local search heuristics. Operations Research Letters, 1(4):219–224, 1992. doi:10.1016/0167-6377(92)90028-2.
- Jean-Yves Potvin and Jean-Marc Rousseau. An exchange heuristic for routeing problems with time windows. The Journal of the Operational Research Society, 46(12):1433-1446, 1995. doi:10.2307/2584063.
- Gerhard Reinelt. The Traveling Salesman: Computational Solutions for TSP Applications, volume 840 of Lecture Notes in Computer Science. Springer, Berlin Heidelberg, Germany, 1994. doi:10.1007/3-540-48661-5.
- Martin W. P. Savelsbergh. An efficient implementation of local search algorithms for constrained routing problems. *European Journal of Operational Research*, 47(1):75–85, 1990. doi:10.1016/0377-2217(90)90091-O.
- Michael Schneider, Fabian Schwahn, and Daniele Vigo. Designing granular solution methods for routing problems with time windows. *European Journal of Operational Research*, 263(2):493–509, 2017. doi:10.1016/j.ejor.2017.04.059.
- Kenneth Steiglitz and Peter Weiner. Some improved algorithms for computer solution of the traveling salesman problem. In *Proceedings of the Sixth Allerton Conference on Circuit and System Theory*, pages 814–821, Urbana, IL, USA, 1968.
- Paolo Toth and Daniele Vigo. The granular tabu search and its application to the vehicle-routing problem. *INFORMS Journal* on Computing, 15(4):333–346, 2003. doi:10.1287/ijoc.15.4.333.24890.
- Thibaut Vidal, Teodor Gabriel Crainic, Michel Gendreau, and Christian Prins. A unified solution framework for multi-attribute vehicle routing problems. *European Journal of Operational Research*, 234(3):658–673, 2014. doi:10.1016/j.ejor.2013.09.045.

Thomas R. Visser and Remy Spliet. Efficient move evaluations for time-dependent vehicle routing problems. *Transportation Science*, 54(4):1091–1112, 2020. doi:10.1287/trsc.2019.0938.