



Design and analysis of stochastic local search for the multiobjective traveling salesman problem

Luís Paquete^{a,*}, Thomas Stützle^b

^aCISUC, Department of Informatics Engineering, University of Coimbra, Coimbra, Portugal

^bIRIDIA, CoDE, Université Libre de Bruxelles (ULB), Bruxelles, Belgium

ARTICLE INFO

Available online 7 December 2008

Keywords:

Multiobjective combinatorial optimization
Meta-heuristics

ABSTRACT

Stochastic local search (SLS) algorithms are typically composed of a number of different components, each of which should contribute significantly to the final algorithm's performance. If the goal is to design and engineer effective SLS algorithms, the algorithm developer requires some insight into the importance and the behavior of possible algorithmic components. In this paper, we analyze algorithmic components of SLS algorithms for the multiobjective travelling salesman problem. The analysis is done using a careful experimental design for a generic class of SLS algorithms for multiobjective combinatorial optimization. Based on the insights gained, we engineer SLS algorithms for this problem. Experimental results show that these SLS algorithms, despite their conceptual simplicity, outperform a well-known memetic algorithm for a range of benchmark instances with two and three objectives.

© 2008 Elsevier Ltd. All rights reserved.

1. Introduction

Stochastic local search (SLS) algorithms are among the most successful techniques for tackling computationally hard problems [1]. In recent years, they have become very popular also for tackling multiobjective combinatorial optimization problems (MCOPs) [2,3]. The currently best performing SLS algorithms for MCOPs typically involve a number of different algorithmic components that are combined into a more complex algorithm. An algorithm developer should therefore have some form of insights into the importance of these algorithmic components and know how they interact with problem characteristics with respect to performance. Ideally, such insights are gained first to make the SLS algorithm design more informed and directed.

In this paper, we present an in-depth experimental analysis of SLS algorithms for the multiobjective traveling salesman problem (MTSP), a paradigmatic NP-hard MCOP. Our analysis is based on a sound experimental design that investigates some usual algorithmic components that can be found in a general algorithmic framework for tackling MCOPs, the scalarized acceptance criterion (SAC) search model [3]. The SAC search model mimics local search approaches that are based on the scalarization of the multiple objective functions. Many such scalarizations are then tackled using local search

(or exact algorithms, if the scalarized problems are efficiently solvable by such algorithms) and the resulting approximate solutions are possibly further treated. Essential components of algorithms following the SAC search model are the number of scalarizations, the search strategy followed (for example, whether information between various scalarizations is exchanged or not), the computation time invested for tackling each of the resulting single objective problems and various others. In fact, decomposing an SLS algorithm into its components allows to employ an experimental design perspective for the analysis of its performance: algorithmic components are seen as *factors*, that is, as abstract characteristics of an SLS algorithm that can affect the response variables such as solution quality. Designing the experiments in a careful way and analyzing them by methods from experimental design allows then to arrive at statistically sound conclusions on the importance of these components and their mutual interdependencies.

While there exist few researches where experimental designs have been used to analyze SLS algorithms for optimization problems with a single objective function [4,5], the usage of experimental designs for analyzing SLS algorithms for MCOPs is rather recent [6,7]. One reason for this is certainly that the outcomes of algorithms for MCOPs are difficult to compare. In fact, fundamental criticisms have been raised against the usage of many unary and binary performance measures [8], which also makes it difficult, if not virtually impossible, to apply the classical ANOVA-type analysis for comparing approximations to the efficient set. Instead, we employ a sound methodology that follows three steps. In a first step, the outcomes of

* Corresponding author. Tel.: +351 239 790 000.

E-mail address: paquete@dei.uc.pt (L. Paquete).

algorithms are compared pairwise with respect to outperformance relations [9]; if these comparisons do not yield clear conclusions, we compute in a next step the attainment functions to detect *significant* differences between sets of outcomes [10,11]. If such differences are detected, the usage of graphical illustrations is used in a third step to examine the areas in the objective space where the results of two algorithms differ more strongly [12].

Our experimental analysis allows an identification of the key-success algorithm components. For example, our results indicate that the two-phase search strategy and the component-wise step [13] are two component levels that yield a significant improvement with respect to solution quality. In addition, the experimental analysis gives insights into the behavior of specific components such as the effectiveness of increasing either the number of scalarizations or the search length.

There is yet another aspect that makes the analysis through the lens of experimental design useful: the insights gained can be exploited to define new high-performing algorithms or at least indicate directions into which existing algorithms should be extended. In other words, the insights gained from the experimental analysis can be helpful to direct the design and engineering of successful SLS algorithms. In fact, based on our experimental analysis, we define SLS algorithms that are assembled from the most promising levels of components we have identified; still, these algorithms remain conceptually rather simple. An extensive experimental comparison of these SLS algorithms on the MTSP with two and three objectives to a well-known state-of-the-art algorithm for this problem shows that they are very competitive or often superior.

The article is structured as follows. In Section 2, we introduce basic notions on MCOPs and the MTSP. Section 3 introduces the SAC model and explains the particular components of these SLS algorithms studied in our experiments. Next, in Section 4, we give an overview of the experimental design, the methodology that was used for comparing the performance of the algorithms and we describe the experimental results obtained. Finally, in Section 5, we compare the performance of our SLS algorithms to a well-known state-of-the-art algorithm. We conclude in Section 6.

2. Multiobjective optimization and the MTSP

The main goal of solving MCOPs in terms of Pareto optimality is to find (all) feasible solutions that are not worse than any other solution and strictly better in at least one objective. The objective function vector for a feasible solution $s \in S$ to an MCOP can be defined as a mapping $f : s \rightarrow \mathbb{R}^Q$, where Q is the number of objectives and S is the set of all feasible solutions. The following order holds for objective function vectors in \mathbb{R}^Q . Let \mathbf{u} and \mathbf{v} be vectors in \mathbb{R}^Q ; we define the *component-wise order* as $\mathbf{u} \leq \mathbf{v}$, i.e., $\mathbf{u} \neq \mathbf{v}$ and $u_i \leq v_i$, $i = 1, \dots, Q$. In optimization, we say (i) $f(s)$ *dominates* $f(s')$ if $f(s) \leq f(s')$; (ii) $f(s)$ and $f(s')$ are *non-dominated* if $f(s) \not\leq f(s')$ and $f(s') \not\leq f(s)$. We use the same notation and wording among solutions if these relations hold between their objective function vectors.

A feasible solution s is said to be a *Pareto optimum solution* if and only if there is no feasible solution s' such that $f(s') \leq f(s)$. There may be more than one Pareto optimum solution; a *Pareto optimum set* is the subset $S' \subseteq S$ that contains *only* and *all* Pareto optimum solutions. We call the image of the Pareto optimum set in the objective space the *efficient set*. In most cases, solving an MCOP in terms of Pareto optimality would correspond to finding solutions that are representative of the efficient set.

The optimization problem handled in this study is the MTSP. In the well-known single-objective version of this problem, a traveling salesman has to visit a set of cities without passing more than once through each city and return to the starting one. The goal is to find a tour such that the total distance traveled is minimized. In the MTSP,

the traveling salesman not only has to minimize the total distance but also the overall traveling time, total cost and so forth. Therefore, it is assumed that several quantities, such as distance, time and cost, are assigned to the connection between each pair of cities. More formally, we define the MTSP as follows: Given Q , a set C of n cities, and distance vectors $d(c_i, c_j) \in \mathbb{N}^Q$ for each pair of cities $c_i, c_j \in C$, the goal is to find every tour in C , that is, a permutation $\pi : [1 \dots n] \rightarrow [1 \dots n]$, such that the length (a vector) of the tour, that is,

$$f(\pi) = d(c_{\pi(n)}, c_{\pi(1)}) + \sum_{i=1}^{n-1} d(c_{\pi(i)}, c_{\pi(i+1)})$$

belongs to the efficient set. The MTSP is known to be NP-hard [14]; additionally, it is known that the lower bound on the expected size of the efficient set for the MTSP is an exponential function of the instance size [15].

The MTSP was chosen for three main reasons. Firstly, its single objective counterpart is one of the best studied NP-hard combinatorial optimization problems and it has been intensively used as a test-bed for experimenting new algorithmic ideas [16], including many SLS algorithms. Hence, experimental results obtained for the multiobjective version may also be interpreted in light of the experience on the performance of these techniques for the single objective case. Secondly, despite the fact that the small instances of the single-objective TSP can be solved in a few seconds to optimality by exact algorithms such as *concorde* (<http://www.tsp.gatech.edu/concorde>), there are two facts that limit their use under fixed time constraints: the typically large variability in the computation times and the potentially very large number of solutions in the efficient set [15]. Thirdly, significant research efforts have been targeted towards applying SLS algorithms to this problem and it has been studied from several different perspectives: from an approximation [17,18], local search [19–21], theoretical [15] and experimental [22] point of view; some related problems have also been studied in the literature [23,24].

3. The search model and algorithmic components

A large number of SLS algorithms have been proposed for MCOPs. Many of these algorithms can be classified as following one of two main search models, the SAC and the component-wise acceptance criterion search (CWAC) search model, or some hybrid thereof [3]. In this article, we analyze the influence of generic components that together mimic the underlying search principles of the SAC search model. (An analysis of generic components of the CWAC search model can be found in [25]; the algorithms following the CWAC model were found to be inferior to those of the SAC model and, hence, here we present only the analysis concerning the more successful model.) Essentially, the SAC search model comprises approaches that use the value returned for solving a scalarization of the objectives for deciding upon the quality of solutions. For a reasonable approximation to the efficient set, it is well known that it is necessary to solve a number of such scalarizations.

As one representative of the SAC search model, we examine the straight-forward approach that solves several scalarizations of the objective function vector and tackles each of these problems with an underlying algorithm for the corresponding single objective version, an approach which underlies many earlier proposed algorithms [20,21,26–29]. We follow the well-known principle of defining scalarizations of the objective function vector with respect to a weighted sum. Hence, the scalarized (single) objective function is defined as

$$f_j(s) = \sum_{q=1}^Q \lambda_q f_q(s), \quad (1)$$

where $\lambda = (\lambda_1, \dots, \lambda_Q)$ is a weight vector. Typically, λ is normalized such that its components sum to one. We follow this convention and each λ we use is an element from the set of normalized weight vectors A given by

$$A = \left\{ \lambda \in \mathbb{R}^Q : \lambda_q > 0, \sum_{q=1}^Q \lambda_q = 1, q = 1, \dots, Q \right\}. \quad (2)$$

Algorithms following the SAC model then solve a number of scalarized problems that are obtained by different weight vectors. The resulting scalarized problems could be solved by any algorithm for the resulting single-objective version; in our case, we apply an effective SLS algorithm for the TSP, which is described later in more detail. In the following subsections, we describe the algorithmic components for the SAC model that are analyzed in this article. For each component at least two and at most three levels are studied in the experimental design; although for several components more options would be possible and interesting, the number of levels was kept restricted to limit the exponential increase of the number of experiments.

3.1. Component: search strategy

The search strategy determines the series of scalarized problems that are defined and tackled. In particular, this concerns the strategy for defining the sequence of weight vectors and how information is transferred from one scalarized problem to another one. We consider two different search strategies.

Restart strategy: The probably most straightforward strategy is to use different weight vectors and to not transfer the results from one scalarized problem to another one. Such a strategy is obtained, for example, by starting the search process for each scalarization from a random initial solution (or, alternatively, by some known construction heuristic that does not use information from previous runs). We call this approach the `Restart` strategy and its pseudo-code is given in Algorithm 1; it results in multiple, independent runs of the single objective SLS algorithm. The procedure `SLS` at the third line is the underlying SLS algorithm that tackles the problems obtained by weight vector λ_i . Since this process could result in dominated solutions, at step 5 only the non-dominated solutions are kept, which is implemented by the procedure `Filter`.

Algorithm 1. Restart search strategy

```

1: for all weight vectors  $\lambda$  do
2:    $s$  is a randomly generated solution
3:    $s' = \text{SLS}(s, \lambda)$ 
4:   Add  $s'$  to Archive
5:   Filter Archive
6: return Archive

```

A set of m weight vectors is used by the `Restart` strategy. Here, we assume that each component of the weight vector has a value $i/z, i=0, \dots, z$, where z is a parameter, and the sum of the components is equal to one, as required by Eq. (2). Since this set of weight vectors can be seen as the set of all compositions of z in Q parts, we have that $m = \binom{z+Q-1}{Q-1}$.

2phase strategy: A different possibility is to transfer results from one scalarization to another one. Here we adopt the `2phase` strategy [13]. In a first phase, a high quality solution for one objective is generated. This solution is the starting solution for the second phase that solves a sequence of scalarizations of the objective

function vector. In this sequence, the initial solution for a scalarization i is the one that is returned from the previous scalarization $i-1$; the first scalarization of the second phase is initialized with the solution returned from the first phase. A pseudo-code of the `2phase` search strategy is given in Algorithm 2. Note that the first and the second phase may make use of two distinct SLS algorithms, which is indicated in Algorithm 2 by `SLS1` and `SLS2`.

Algorithm 2. 2phase search strategy

```

1:  $s$  is a randomly generated solution
2:  $s' = \text{SLS}_1(s)$  /* First phase */
3: for all weight vectors  $\lambda$  do
4:    $s = s'$ 
5:    $s' = \text{SLS}_2(s, \lambda)$  /* Second phase */
6:   Add  $s'$  to Archive
7:   Filter Archive
8: return Archive

```

Concerning the definition of the sequence of weight vectors that define each scalarization, several strategies may be followed. Here, we adopt a *minimal change* strategy between two successive weight vectors to define this sequence. It can be built for two objectives by generating $m=z$ weight vectors such that $\lambda_i = (1-i/z, i/z)$, $i=1, \dots, m$, if the first objective is the one that is optimized in the first phase. For more than two objectives, we define this sequence such that two successive weight vectors differ only by $\pm 1/z$ in any two components. Thus, a minimal change is incurred between components of two successive weight vectors generalizing the biobjective case. To generate such a sequence, an algorithm for generating compositions of z into Q parts can be used. In our particular case, we need a *combinatorial Gray code* for this task, which can be generated by the *Gray code for compositions* [30].

Finally, one may run the `2phase` strategy considering different orders of the objectives. One possibility would be to run it once for each permutation of the Q objectives, that is, $Q!$ times. Computationally less expensive is to consider only the combinations of each pair of objectives, totalizing $\binom{Q}{2}$ runs, or to apply just one run for each objective.

3.2. Component: number of scalarizations

The number of scalarizations is defined by the parameter z . It is expected that an increase of the number of scalarizations would increase also the number of different (non-dominated) solutions returned. However, how much the number of solutions grows when increasing the number of scalarizations is not clear in advance. Therefore, we consider the number of scalarizations as a numerical parameter, that is, also as an algorithmic component, whose influence is studied in the experimental part.

3.3. Component: neighborhood structure

In our analysis, we study two main components of the underlying SLS algorithm. Both are related to the quality of the solutions it returns. The first component is the neighborhood structure that is used; it identifies which solutions are neighbored and it has a significant influence on the performance of local search algorithms. The typical trade-off is that the larger the neighborhood, the better the quality of the solutions that are found by, for example, iterative improvement algorithms. However, an increase of neighborhood size also corresponds to an increase of computation time to find improving neighbors.

3.4. Component: search length

The second component of the underlying SLS algorithm is the number of iterations. The reason for studying this component is that by increasing the number of iterations, the final solution quality returned by the SLS algorithm tends to increase, but at the same time does also the computation time. In other words, there is a trade-off between these two criteria and a good compromise needs to be found for the design of an algorithm following the SAC model.

3.5. Component: component-wise step

The number of solutions returned is bounded by the number of compositions of z in Q parts. One possibility for increasing this number is by accepting, for each scalarization, non-dominated solutions in the neighborhood of the solution returned by the underlying single-objective SLS algorithm. We call this additional component *component-wise step* [13]. In our particular case, this step uses the neighborhood that is defined by the neighborhood component.

4. Experimental analysis

4.1. Experimental design

The experimental design considered all the five algorithm factors that were described in the previous section plus two factors concerning the MTSP instances, namely the instance type and the instance size.

4.1.1. MTSP instances

For the experimental part of the study, we have generated MTSP instances of different sizes and using different ways to generate the distances between the cities. We used the random instance generator available from the 8th DIMACS implementation challenge site¹ and, for each objective, we generated one distance matrix. We only considered two objectives for the experimental analysis of the algorithm components, while the comparison with a state-of-the-art SLS algorithm for the MTSP also included instances with three objectives. Three types of biobjective instances were generated²:

- *Random uniform Euclidean* (RUE) instances, where each component of the distance vector is generated as usual in RUE instances: each distance value corresponds to the Euclidean distance between two points in a two-dimensional plane rounded to the next integer; the coordinates of each point are integers that are uniformly and independently generated in the range [0,3163].
- *Random distance matrix* (RDM) instances, where each component of the distance vector assigned to an edge is chosen as an integer value taken from a uniform distribution in the range [0,4473].
- *Mixed* instances, where one objective assigns distances to the edges as in RUE instances while the other assigns distances as in RDM instances.

The range of edge lengths for the RUE instances was chosen in order to meet the range of values of the Krolak/Felts/Nelson instances available in TSPLIB (files with prefix `kro`). The range of the edge lengths for the RDM instances and the RDM objective of the mixed instances were chosen in order to have a range similar to the one of the RUE instances (note that $\lfloor \sqrt{2 \times 3163^2 + 0.5} \rfloor = 4473$). The instance sizes considered were $n=100, 300$ and 500 , which includes instances

Table 1

List of components and the corresponding levels that we considered for our experimental setup.

Components	Levels
Search strategy	{Restart, 2phase}
Number of scalarizations	{ $n, 5n, 10n$ }
Neighborhood structure	{2-exchange, 3-exchange}
Search length	{0, 50, 100}
Component-wise step	{True, False}

that are larger than those considered in most of the literature on the MTSP. A wide range of values allow us to test whether instance size plays an important role on the algorithm performance.

For each instance size and type of instance, three instances were generated, resulting in a total of 27 instances.

4.1.2. Algorithmic component levels

For each of the factors concerning algorithmic components, the main effects of two or three levels were studied. A summary of the components and their associated levels is given in Table 1. Necessary details on the components are explained as follows.

Search strategy: If the search strategy `2phase` is chosen, the first phase will optimize the first objective and the solution returned is also the starting solution for the second phase. For RUE and RDM instances, the first phase of `2phase` consisted in running an iterated local search (ILS) algorithm [31] (which is described below in some more detail) for 50 iterations only for the first objective; this usually gives a high quality solution for the single objective case. For the mixed instances, we considered two variants of the `2phase` strategy: `2phaseE` starts the first phase optimizing only the objective defined by the Euclidean distance matrix; `2phaseR` starts optimizing the objective defined by the RDM distance matrix. Using these two versions of the `2phase` strategy, we can thus also analyze the dependence of the final performance on the structure of the objective that is used for generating the initial solution for the second phase. Since the solutions found in preliminary experiments for the mixed instances had a lower range of objective function values for the RUE objective than for the RDM objective, we equalized the ranges of both objectives by multiplying the i th component of the objective function value vector by a range equalization factor F_i [32], which for each objective i is

$$F_i = \frac{R_i}{\sum_{j=1}^Q R_j},$$

where Q is the number of objectives and R_i is the range of objective function values for the objective i . We computed an approximation to the range of the efficient set as follows. First, we run an ILS algorithm [31] 10 times for each objective; then, given the best solutions to the first and the second objective, s_1 and s_2 , respectively, the ranges are computed as $R_1 = f_1(s_2) - f_1(s_1)$ for the first objective and as $R_2 = f_2(s_1) - f_2(s_2)$ for the second objective.

Number of scalarizations: For the number of scalarizations, we analyze three levels, $n, 5n$ and $10n$, where n is the number of cities in the MTSP instance.

Neighborhood structure: We consider two standard TSP neighborhood structures, the 2- and 3-exchange neighborhoods. (In general, two solutions are neighbored in the k -exchange neighborhood if they differ by at most k edges.) The iterative improvement algorithms we use for each scalarization, a weighted sum of the objectives, make use of well-known TSP speed-up techniques such as *don't look bits* and fixed radius search [33] within nearest neighbor candidate lists [16]. Clearly, one could also use more complex neighborhood structures like the ones used in the Lin-Kernighan local search

¹ The generator is available at <http://www.research.att.com/~dsj/chtsp/download>.

² The instances are available at <http://eden.dei.uc.pt/~paquete/tspl/>.

algorithms [34]; however, we restricted to the 2- and 3-exchange neighborhoods for the sake of limiting the number of configurations.

Search length: On top of the resulting two iterative improvement algorithms, we used a general-purpose SLS method called ILS [35]. This allows larger search lengths and in this way to intensify the search for each scalarization. (In fact, ILS forms the basis for many high-performing SLS algorithms for the TSP [1].) An algorithmic outline of ILS is given in Algorithm 3. The perturbation used in the ILS algorithm is a random double-bridge move and the acceptance criterion accepts a new solution only if it improves over the previous one. For the ILS algorithm, we have used the code provided at <http://www.sls-book.net/>. Since the ILS algorithm uses the iterative improvement algorithm as its subsidiary local search procedure, the application of the iterative improvement algorithm corresponds to “zero iterations” of the ILS algorithm. Hence, the three different levels of the *search length* component can be indicated as 0, 50 and 100 iterations of ILS, respectively, which we denote in the following as ILS(0), ILS(50) and ILS(100).

Algorithm 3. Algorithmic outline of an iterated local search algorithm as used in the experimental setting. `no_iterations` is a parameter that defines the number of times the main loop of the algorithm is executed. If `no_iterations=0`, the algorithm corresponds to a single application of iterative improvement.

```

1: input:  $s_0$  (an initial solution), no_iterations (number
   of iterations of main loop)
2:  $s^* := \text{IterativeImprovement}(s_0)$ 
3: for  $i := 1$  to no_iterations do
4:    $s' := \text{Perturbation}(s^*)$ 
5:    $s^{*'} := \text{IterativeImprovement}(s')$ 
6:    $s^* := \text{AcceptanceCriterion}(s^*, s^{*'})$ 
7: return  $s^*$ 

```

Component-wise step: The effect of the component-wise step was only explicitly analyzed for instances of size 100; in fact, on all instance sizes the component-wise step has a (strongly) positive effect and, hence, we decided to use it always on the large instances. This has the side-advantage of reducing the number of experiments on the instances of size 300 and 500 by a factor of two and, thus, saving significant computational effort.

4.2. Performance assessment methodology

Assessing the performance of algorithms for MCOPs is by far more complex than in the single-objective case and a number of serious problems, in particular of unary performance indicators, have been described [8]. Our experimental analysis is based on a three step evaluation that avoids these known drawbacks. In a first step we use the *better* relations, which provide the most basic assertion of performance; the second step computes attainment functions and tests the equality of the attainment functions [36]; the third step consists in detecting the largest differences of performance in the objective space between pairs of algorithms. Most aspects of this three-step experimental analysis were already described in the literature [6,12] and are summarized here for the sake of comprehensibility of the remainder of the paper.

Step 1: Better relations. A set of points A is *better* than a set of points B if every point of B is dominated or equal to any point of A and A is different from B . This relation was introduced in Hansen and Jaszkiewicz [9] as one of the outperformance relations that can be established between pairs of outcomes of SLS algorithms for MCOPs. Thus, as a first step, we count how many times each outcome associated with each level of a component is *better* than the ones from another level of the same component. However, we restrict the

comparison of outcomes to those that were produced *within* the same levels of other components in order to reduce variability. This allows us to detect if some level is *clearly* responsible for a good or bad performance. If no clear answers are obtained from this first step, we can conclude that the outcomes are mostly *incomparable*, that is, neither A is *better* than B nor *vice versa*. Since then we do not know to what extent they really differ, we test the equality of their attainment functions.

Step 2: Attainment functions. In Fonseca and Fleming [10], the performance of an SLS algorithm for multiobjective problems is associated with the probability of attaining (dominating or being equal to) an *arbitrary point* in the objective space in one single run. This function is called *attainment function* [36] and it can be seen as a generalization of the distribution function of solution cost [1] to the multiobjective case. These probabilities can be estimated empirically from the outcomes obtained in several runs of an SLS algorithm by the *empirical attainment function* (EAF). Then, we can formulate statistical hypotheses and test them based on the EAFs of several algorithms for a certain problem instance. A suitable test statistic for the comparison of two algorithms is the maximum absolute distance between their corresponding EAFs, analogous to the Kolmogorov–Smirnov statistic [37]. For the case of $k > 2$ algorithms, we choose the maximum absolute distance between the k EAFs, analogous to the Birnbaum–Hall test [37]; if the global null hypothesis of equality is rejected, we test the equality between each pair of EAFs, where the p -values are corrected by Holm’s procedure [38]. Since the distribution of these test statistics is not known, permutation tests [39] based on the above test statistics have to be performed [11]. The permutation procedure has to be changed according to the experimental design chosen. For instance, in the presence of several factors, restricted randomizations [39], as done in Paquete and Fonseca [6] in a similar context, can be applied; for testing the main effects of each component, we allow permutations of the outcomes *between* different levels of the component of interest, but *within* the same levels of the other components.

Step 3: Location of differences. If the previous analysis indicates that the null hypothesis of equality of the attainment functions should be rejected, the largest performance differences can be visualized by plotting the points in the objective space with a large absolute difference of the EAFs. In fact, large has a subjective meaning; here, we plot the points whose absolute differences were above or equal to 20%, assuming that lower values are negligible.³ Since the sign of the difference at each point gives information about which algorithm performed better at that point, we may plot positive and negative differences separately, if differences in both directions exist.

Fig. 2 illustrates the main idea. Each of the two plots give the differences of the EAFs associated with two algorithms that were run several times on one instance. The lower line on each plot is a lower bound on the efficient set,⁴ while the upper line connects the set of points attained by all runs of both algorithms. On both plots are shown the regions where the EAF of Algorithm 1 (using 2phase strategy) takes larger values by at least 20% than that of

³ The minimum number of outcomes that were used for statistical tests in this thesis were 10 (5 runs associated with each level of a factor); thus, a difference of 20% corresponds also to the minimum difference that can be observed in these comparisons.

⁴ The lower bound is used simply as a visual reference when plotting the differences with respect to the EAFs. We use a lower bound based on the solution of the 2-matching problem; it yields a lower bound approximately at 14% of the optimum for the TSP [40]. (Note that better quality lower bounds exist, but this is for our purposes not important.) For our particular case, we solved the 2-matching problem using as input a matrix resulting from the weighted sum of distances assigned to each edge of an MTSP instance. This procedure is repeated for 5000 maximally dispersed weight vectors and the lower bounds have been computed using CPLEX.

Algorithm 2 (using Restart strategy); the observed differences are encoded using a grey scale—the darker the stronger are the differences. In this case, no point in the EAF of Algorithm 2 was larger than the corresponding one of Algorithm 1. If differences in favor of each of the algorithms occur, the positive differences in favor of each algorithm can be given in one plot, as it is done in Fig. 1.

4.3. Experimental results

Each configuration resulting from any of the possible combinations of levels of the factors, as described in Section 4.1, was run five times on an AMD Athlon (TM) 1.2 GHz CPU, 512 MB of RAM under Suse Linux 7.3. We have permuted randomly the original order of the runs to remove a possible bias. In the following, we discuss the results of the analysis for each SLS component under study. Each permutation test for testing hypotheses on the equality of EAFs used 10000 permutations and the significance level was set to $\alpha = 0.05$. Due to space restrictions, we only show the most relevant plots of the locations of the differences as explained in Section 4.2; a full collection of the results comprising all the data on the comparisons and all plots is available at <http://eden.dei.uc.pt/~paquete/mtsp>.⁵

4.3.1. Component: search strategy

The results in Table 2 with respect to the better relations indicate that the 2phase strategy performs slightly better than the Restart strategy for larger instances and that the difference is more relevant on RDM instances. In addition, the null hypothesis of equality of the EAFs was always rejected. Hence, the search strategies behave statistically different with respect to the corresponding EAFs in the instances tested.

Figs. 1 and 2 indicate the location of differences above 20% between the search strategies. The two plots of Fig. 1 indicate that the Restart strategy covers a wider part of the trade-off than the 2phase strategy on the RUE instances of size 100, though with only a small difference (which is reflected by the fact that the differences are almost imperceptible); the latter performs better only towards the first objective, where the first phase terminated. However, as instance size increases in RUE and RDM instances, the 2phase strategy performs clearly better than the Restart strategy, as shown in the plots of Fig. 2.

Finally, the comparisons on the mixed instances have shown interesting tendencies. When comparing 2phase_E and 2phase_R , each strategy has advantages towards the objective that is optimized in the first phase and the observed differences between the two are roughly the same across the various instance sizes. Differences in favor of the Restart search strategy for instance size 100 are located in the center of the trade-off; however, on larger instances, no differences above 20% in favor of the Restart strategy were found.

4.3.2. Component: component-wise step

The comparison based on the better relation with respect to the use or not of the component-wise step always resulted in incomparable cases, but the null hypothesis with respect to the equality of EAFs was always rejected. Hence, there are always significant differences between using or not the component-wise step. The location of differences above 20% clearly indicates that the use of this step yields a significant advantage, as shown in the top plot of Fig. 3. However, this advantage is not constant over all types of instances tested: the differences are stronger for RUE instances than for RDM

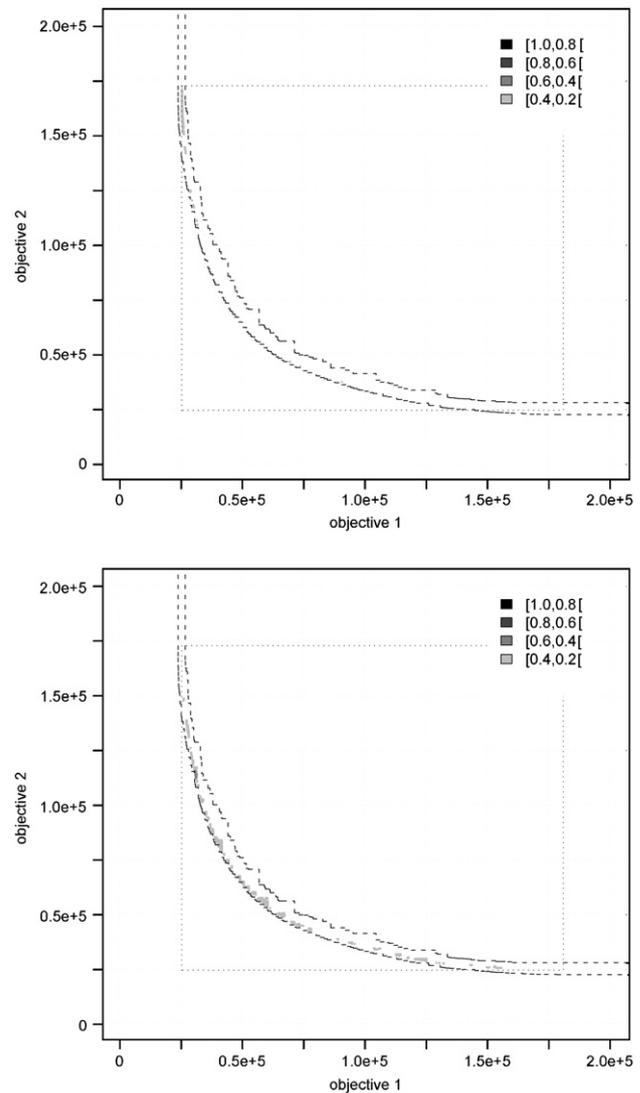


Fig. 1. Location of differences between the 2phase and the Restart search strategy in favor of the former (top) and in favor of the latter (bottom), for RUE instances of size 100. Note that the differences between the two strategies in this case are all minor, that is, in the range of [0.2,0.4], and are therefore almost imperceptible.

Table 2

Given is the percentage of the pairwise comparisons in which the 2phase search strategy (for mixed instances, 2phase_E and 2phase_R , respectively) was better than the Restart strategy.

Size	RUE	RDM	Mixed	
	2phase (%)	2phase (%)	2phase_E (%)	2phase_R (%)
100	0.0	6.4	0.0	0.0
300	4.1	31.6	5.2	2.1
500	10.1	31.7	15.8	6.9

In no comparison, Restart was found to be better than 2phase (or 2phase_E and 2phase_R , respectively).

instances. In addition, the differences for mixed instances lie more towards the Euclidean objective.

Concerning the computation time, it is remarkable that the addition of the component-wise step increases the computation time by only about 1%, which is negligible in most applications. Furthermore, the number of non-dominated solutions is increased by a factor of

⁵ Note that the computation of all the experimental results including the execution of the hypothesis tests took more than 6 months of CPU-time; in fact, the exponential increase of the number of experiments was one reason for limiting the experiments to a small number of levels for each factor.

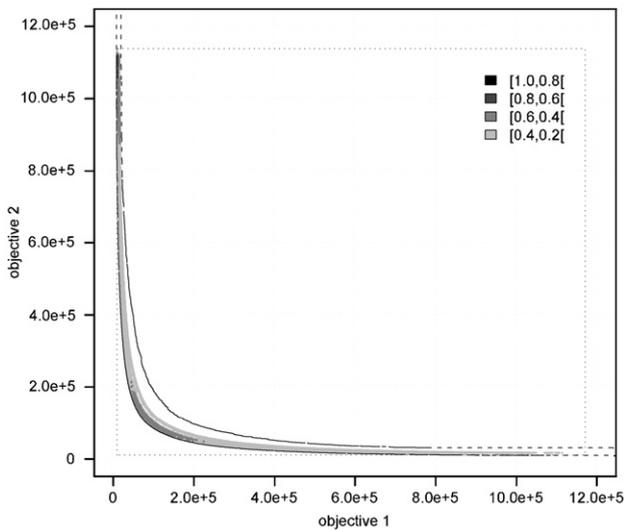
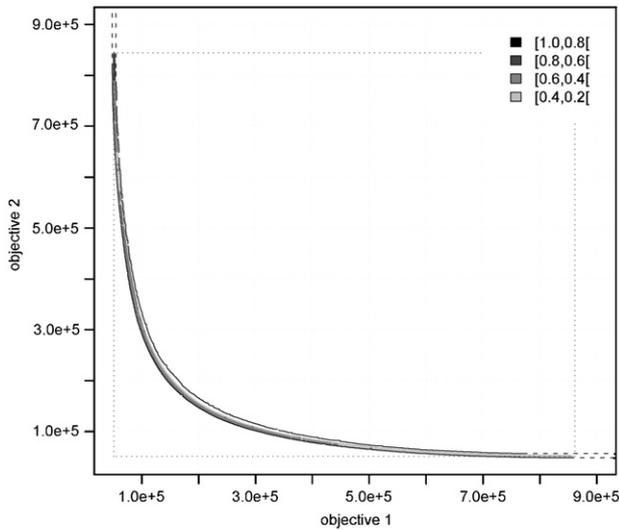


Fig. 2. Location of differences between search strategies for an RUE (top) and an RDM instance (bottom) of size 500. All differences are in favor of the 2phase strategy.

about six for RUE instances and by a factor of about three for RDM instances.

4.3.3. Component: neighborhood structure

The results based on the better relation indicated that using a 3-exchange neighborhood results in significantly better performance than 2-exchange with the advantage of 3-exchange over 2-exchange increasing strongly with instance size; the advantage of 3-exchange is most evident for the RDM instances. Table 3 gives a summary of the observed percentages of 3-exchange being better than 2-exchange for the different instance sizes and the different instance types. Given these strong differences, clearly also the null hypothesis with respect to the equality of the EAFs was always rejected. The differences above 20% were always in favor of 3-exchange and the differences were more pronounced for larger instances (see bottom plot of Fig. 3). Hence, this result is analogous to the relative behavior between these neighborhoods in iterative improvement algorithms for the single-objective TSP [1,16,40].

4.3.4. Component: search length

As said before, the search length defines the number of iterations for a single execution of the ILS algorithm, denoted by ILS(*i*). According to the use of the better relation, we observed that ILS(50)

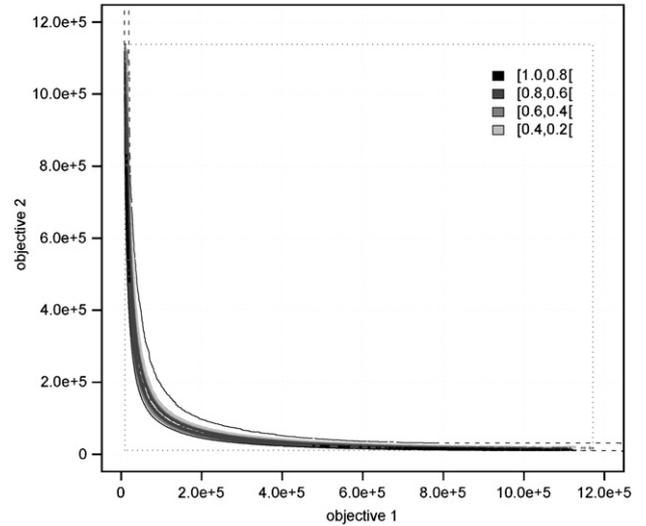
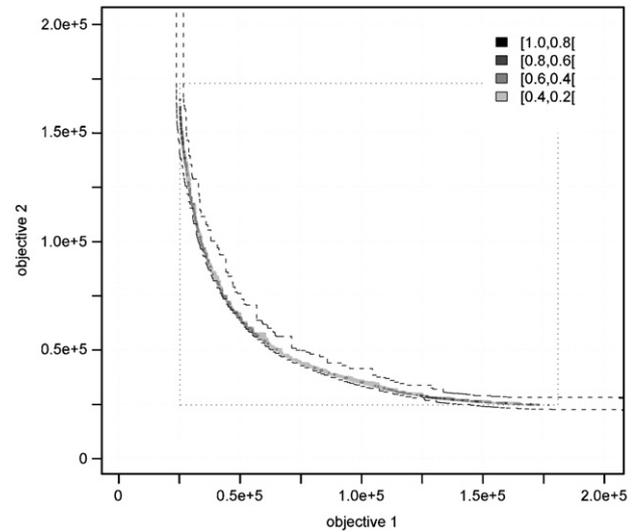


Fig. 3. Location of differences between using or not the component-wise step in favor of the latter for an RUE instance of size 100 (top) and between the 2-exchange and the 3-exchange neighborhood in favor of the latter for an RDM instance of size 500 (bottom).

Table 3

Given is the percentage of the pairwise comparisons in which 3-exchange is better than 2-exchange.

Size	RUE (%)	RDM (%)	Mixed (%)
100	1.2	50.3	6.5
300	43.0	67.4	22.3
500	53.1	75.0	26.2

In no comparison 2-exchange was found to be better than 3-exchange.

and ILS(100) perform better than ILS(0), that is, iterative improvement, and that the frequency of better performance is higher for instances of size 300 and 500 than for those of size 100. The relative performance seems similar across all types of instances, though less emphasized for mixed instances. See Table 4 for more details. The results also indicate that the comparisons between the outcomes obtained by ILS(50) and ILS(100) are incomparable.

The statistical tests on the equality of the EAFs clearly indicate the rejection of the null hypothesis for all instances. Thus, any increase of the number of iterations from 50 to 100 results still in statistically

Table 4

Given is the percentage of the pairwise comparisons in which a search length of 50 or 100 was better than a search length of 0.

Size	RUE		RDM		Mixed	
	0 vs. 50 (%)	0 vs. 100 (%)	0 vs. 50 (%)	0 vs. 100 (%)	0 vs. 50 (%)	0 vs. 100 (%)
100	15.1	20.4	42.6	54.9	22.8	26.6
300	65.3	80.0	67.0	77.7	34.6	52.8
500	71.8	79.9	69.0	74.1	40.0	49.7

In no comparison, a search length of 0 was better than 50 or 100. The pairwise comparisons between search lengths 50 and 100 resulted always in incomparable cases.

Table 5

Given is the percentage of the pairwise comparisons in which a number of scalarizations j is better than a number of scalarizations i (indicated by i vs. j).

Size	RUE		RDM		Mixed	
	n vs. $5n$ (%)	$5n$ vs. $10n$ (%)	n vs. $5n$ (%)	$5n$ vs. $10n$ (%)	n vs. $5n$ (%)	$5n$ vs. $10n$ (%)
100	0.0	0.0	0.0	0.0	0.0	0.0
300	0.0	0.0	8.3	11.9	0.9	7.3
500	0.0	0.0	7.0	10.5	1.1	7.6

In no comparison, a number of scalarizations $i < j$ was better than j .

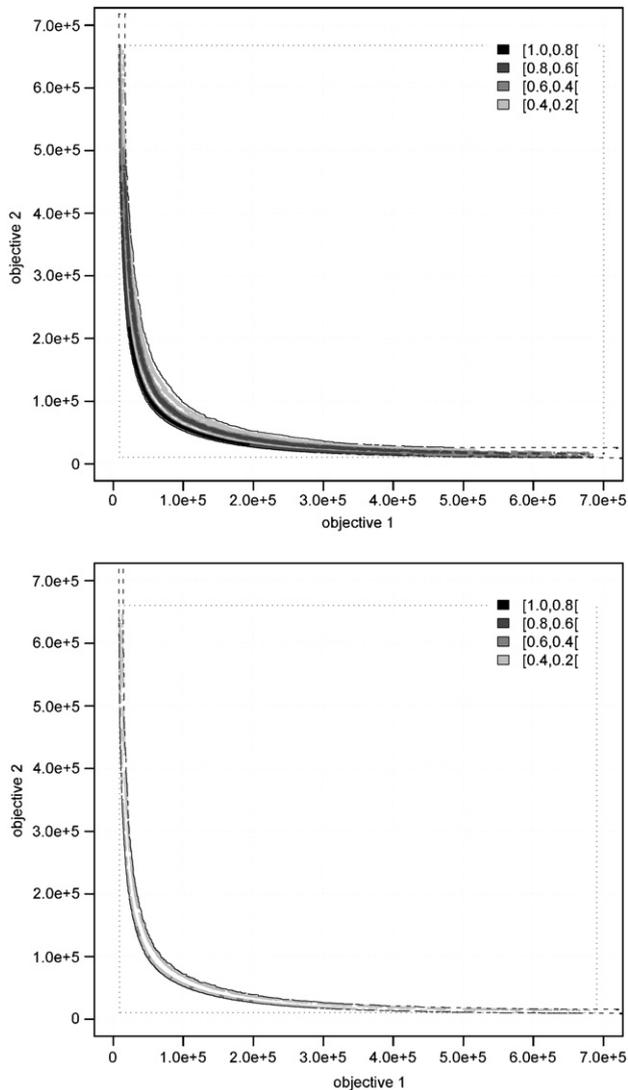


Fig. 4. Location of differences between iterative improvement and ILS(50) in favor of the latter (top) and between ILS(50) and ILS(100) in favor of the latter (bottom) for an RDM instance of size 300.

significantly better performance. However, the examination of the location of the differences above 20% indicates that, for all the instances tested, the major leap in performance is given by moving from ILS(0) to ILS(50), while moving from ILS(50) to ILS(100) yields somewhat less pronounced (but still significant) differences. For an illustration of this behavior, we refer to Fig. 4.

4.3.5. Component: number of scalarizations

Differently from an increase of the search length, an increase of the number of scalarizations does not correspond to an evidently better performance with respect to the better relation; as can be seen in Table 5, some minor evidence for improved performance is only found for large RDM and mixed instances. However, the null hypothesis of equality with respect to the EAFs is always rejected for any instance, which means that an increase of the number of scalarizations results in a significant effect. The location of differences above 20% indicates that the performance differences are not very pronounced. While the differences between n against $5n$ scalarizations are still rather clearly visible, as shown in Fig. 5 on the top plot, the differences between $5n$ and $10n$ scalarizations are almost imperceptible (see bottom plot).

4.4. Summary

The main insights from the experimental analysis for the SAC search model applied to the MTSP are the following. A substantial gain in solution quality can be obtained by choosing an underlying high performing SLS algorithm. Two ways of improving the performance of the underlying SLS algorithm have been studied: the underlying neighborhood (solution quality is known to improve considerably also for the single objective case when moving from the 2-exchange to 3-exchange neighborhood in an iterative improvement algorithm) and the search length, here defined by the number of iterations of the underlying ILS algorithm. In fact, these insights would suggest that a further improvement of the performance might be expected when moving to effective implementations of the Lin-Kernighan algorithms as provided by Helsgaun [41] or the concorde library [42]. (This is the case because for the single-objective TSP, the Lin-Kernighan algorithm is known to reach better quality solutions when compared to the iterative improvement algorithms we use; in fact, this same ranking transfers to the case once the iterative improvement algorithms are included into an ILS algorithm [1,16,43].)

The component-wise step was also shown to have a significant impact on the final solution quality, as we observed on all instances studied. In addition, the computational overhead caused by its introduction seems to be minor at least in the biobjective case studied here: on average it leads to an increase of the computation time by approximately one percent. Hence, concerning computation time the impact of adding that component-wise step is much less than when moving from 2-exchange to 3-exchange neighborhood, which actually increases the computation time significantly.

Concerning the choice of the search strategy, there is a clear interaction between the search strategy and the type of instance. For small RUE instances with 100 cities, the Restart strategy has slight advantages over 2phase. However, for larger mixed and Euclidean instances, and for all RDM instances tested, the 2phase strategy is clearly preferable, often by a large margin. In fact, the experimental analysis of the SAC search model indicated that instance features

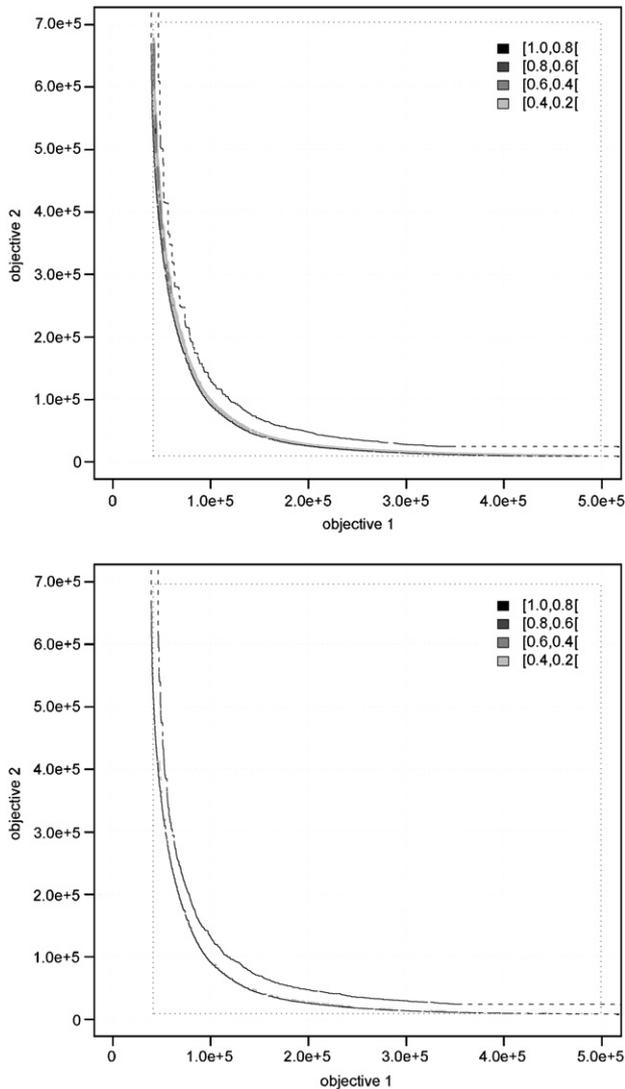


Fig. 5. Location of differences between n and $5n$ scalarizations in favor of the latter (top) and between $5n$ and $10n$ scalarizations in favor of the latter (bottom) for a mixed instance of size 300.

play a strong role in the performance of the algorithms under study; it is therefore expected that those features are also relevant in the performance of many other SLS algorithms.

The impact of the number of scalarizations seems to be the smallest, at least when judging from the better relations. In the plots of the differences, the step from n to $5n$ scalarizations was most noticeable, while further increasing it to $10n$ gave no further strong advantages.

5. Comparison with a state-of-the-art algorithm

The insights gained from an extensive study of a class of algorithms through experimental designs may, beyond the scientific insights gained, also be useful to define new high-performing algorithms. Here we show that, indeed, this step can effectively be done by deriving such algorithms. These algorithms are then compared to a multiobjective memetic algorithm called MOGLS. This algorithm was proposed by Jaszkiwicz [21], who kindly provided us the source code of this algorithm, and in earlier studies it was shown to outperform other SLS algorithms for the MTSP for two and three objectives [21]: the algorithms to which MOGLS was compared include MOGA [44], MOSA [29] and Ishibushi and Murata's memetic algorithm [45].

MOGLS works as follows. It initializes two archives CS and A with l solutions obtained from runs of an iterative improvement algorithm based on the 2-exchange neighborhood with respect to a weighted sum scalarization with randomly generated weights. Then, it iterates r times over the following two steps. First, it chooses two among the best m solutions in A with respect to a weighted sum based on randomly generated weights; these two solutions are then *recombined* by the *distance-preserving crossover* [46]. Next, it applies a different iterative improvement algorithm based on the 2-exchange neighborhood to the new recombined solution using the same weighted sum scalarization; the resulting local optimum s^* is added to archive CS if it is better than the worst solution among the m best solutions according to the weight vector considered; finally s^* is added to archive A, if no solution dominates it and the archive A is updated. The performance of this approach depends on parameters m , l and r . These two steps (recombination and local improvement) are repeated for $r \cdot l$ iterations, thus generating a total of $(r + 1) \cdot l$ local optima between the initialization and the following iterations.

For our experiments, we follow the parameter settings proposed by Jaszkiwicz [21] as far as possible: we set $m = 16$, the number of iterations $r = 50$, which was the maximum value tested earlier; for l we used the value 142 for instances of size 100 with two objectives, but for instances of size 300 we extrapolate it linearly, which resulted in $l = 278$.

The configurations of our SLS algorithms were chosen according to the main insights from the experimental design. The only exception is that we use only local search based on the 2-exchange neighborhood, as MOGLS. (Using the 3-exchange neighborhood would clearly bias the results in our favor.) For the comparison we tested the two approaches on the RUE and RDM instances with 100 and 300 cities. We used the following configurations: all SLS algorithms use (i) the component-wise step, (ii) 100 iterations of the ILS algorithm, (iii) $10n$ scalarizations and (iv) the 2-exchange neighborhood. Regarding the search strategy, we use always the 2phase strategy, with the only exception being the RUE instances with 100 cities: for these instances, our experimental analysis indicated slightly better performance with the Restart strategy and, hence, we follow our conclusions of the experimental analysis. Each of the algorithms was run 10 times on a single CPU of a computer with two AMD Athlon (TM) 1.2 GHz CPUs with 512 MB of RAM running under Suse Linux 7.3. For measuring the computational effort spent by the algorithms, we use the number of times a neighboring solution is evaluated (recombinations in case of MOGLS, and perturbations in the ILS were not counted). This was done, because the implementations were not coded using the same programming languages and the very same data structures and, hence, measuring CPU time would have been unfair. (We verified that our code was actually much faster concerning CPU time than MOGLS and, hence, in this way we also avoid the bias in favor of our code that would occur if we stopped both algorithms at a same CPU time.)

Given that most of the outcomes returned by the two algorithms were incomparable in some preliminary experiments, we decided to directly apply the statistical tests at the 5% significance level for checking the null hypothesis of equality between attainment functions: for all experiments, the null hypothesis was always rejected. The location of the differences above 20% showed a clearly better performance of our SLS algorithms over MOGLS, except in the RUE instance with 100 cities and two objectives.⁶ On this instance, the differences were rather small, except towards the improvement of the second objective where our configuration performs better (see

⁶ The complete EAF plots are available online at <http://eden.dei.uc.pt/~paquete/mtsp>.

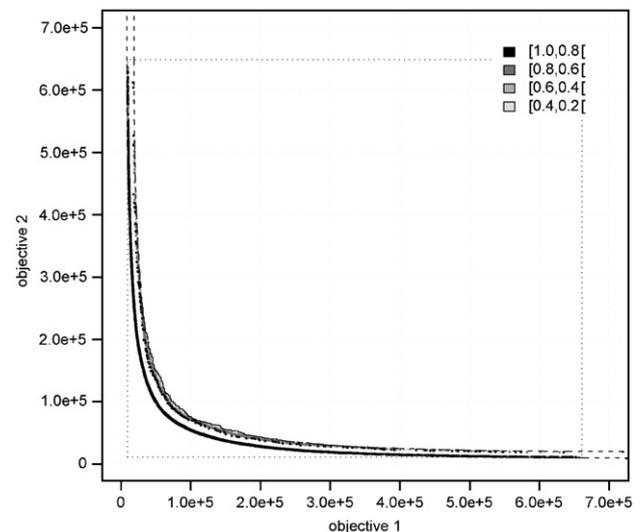
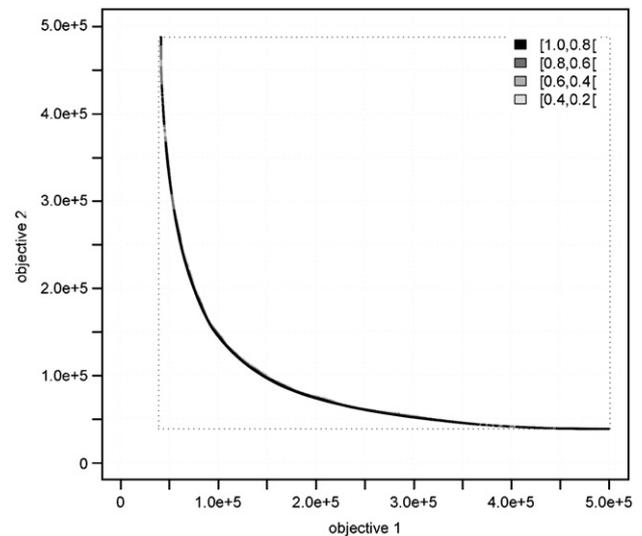
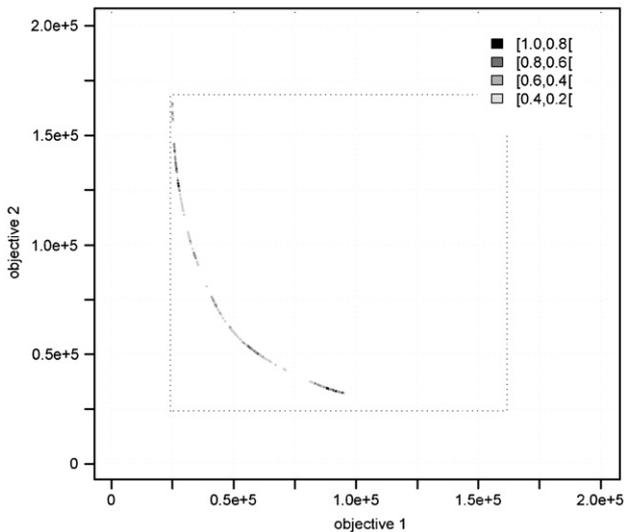
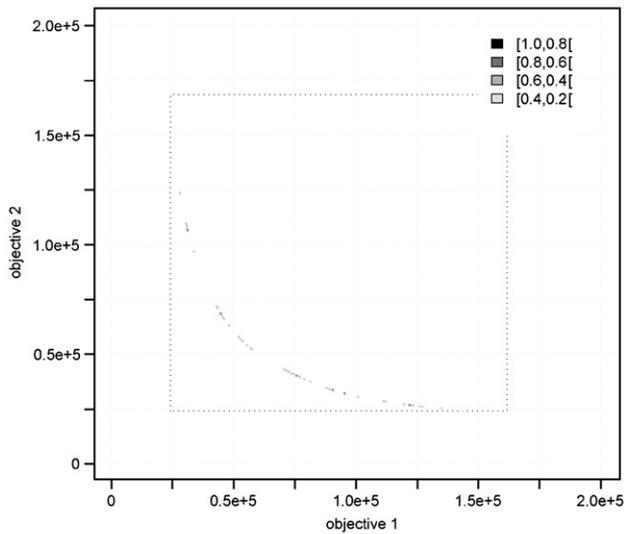


Fig. 6. Location of differences between MOGLS and our SLS algorithms in favor of MOGLS (top) and in favor of ours (bottom) for an RUE instance with 100 cities. The differences are not very marked, but slightly stronger in favor of our algorithm. Note that, in order to improve visibility, the gap between the worst case and the lower bound was removed.

Fig. 6). These two plots show that the differences are larger and more spread in favor of our SLS algorithm, though we can notice that there are still regions of the objective space where MOGLS performs slightly better. However, for all RDM instances, and for all RUE instances with 300 cities, we found only differences in favor of our approach; for an example, see Fig. 7.

Given the high performance advantage of our SLS algorithms, we decided to run some experiments comparing them to MOGLS for MTSP instances with three objectives. Since the experimental study was done only for the two objectives case, we first performed some exploratory experiments concerning the configuration of the SLS algorithms. Based on these, we decided to increase strongly the number of scalarizations to 5151 scalarizations (that is, $z = 100$), for all instance sizes. We did not apply the component-wise step for three objectives; the main reason was that the component-wise step for more than two objectives has the drawback of returning many clusters of solutions in the objective space. We therefore compensate the lack of the component-wise step by the increase of the number of scalarizations. All other components remained the same (that is, the `Restart` search strategy is used for RUE instances of 100 cities,

Fig. 7. Location of differences between MOGLS and our SLS algorithms in favor of the latter on an RUE (top) and an RDM instance (bottom) with 300 cities. No differences in favor of MOGLS have been observed. Note that, in order to improve visibility, the gap between the worst case and the lower bound was removed.

while `2phase` is applied in all other cases). For MOGLS, we increased the values of the parameter l , the initial size of the archive, to 662 for instances with 100 cities, as suggested by Jaskiewicz [21], and to 1786 for instances of size 300 by linear extrapolation.

For the instances with three objectives, we did not apply the statistical test because the computation of the EAFs, which consisted in more than one million points, took already about one week for each experiment. However, the maximum absolute difference of one between the EAFs was always detected and, hence, we suspect that the null hypothesis of equal EAFs would anyway be rejected. (Recall that the statistical test is based on the maximum difference between EAFs associated with two different algorithms.) For an illustration of the results, we used the parallel coordinates graphical technique [47], where each line corresponds to one point in the objective space for detecting the location of the differences. Examples of the resulting plots are given in Fig. 8, where points with differences in the range of (0.8,1.0] are plotted on top and points in the range of (0.6,0.8] are plotted on bottom in favor of MOGLS (left side) and in favor of our SLS algorithms (right side). Since, for each point a line is drawn, the much darker plots on the right side indicate a strong advantage of

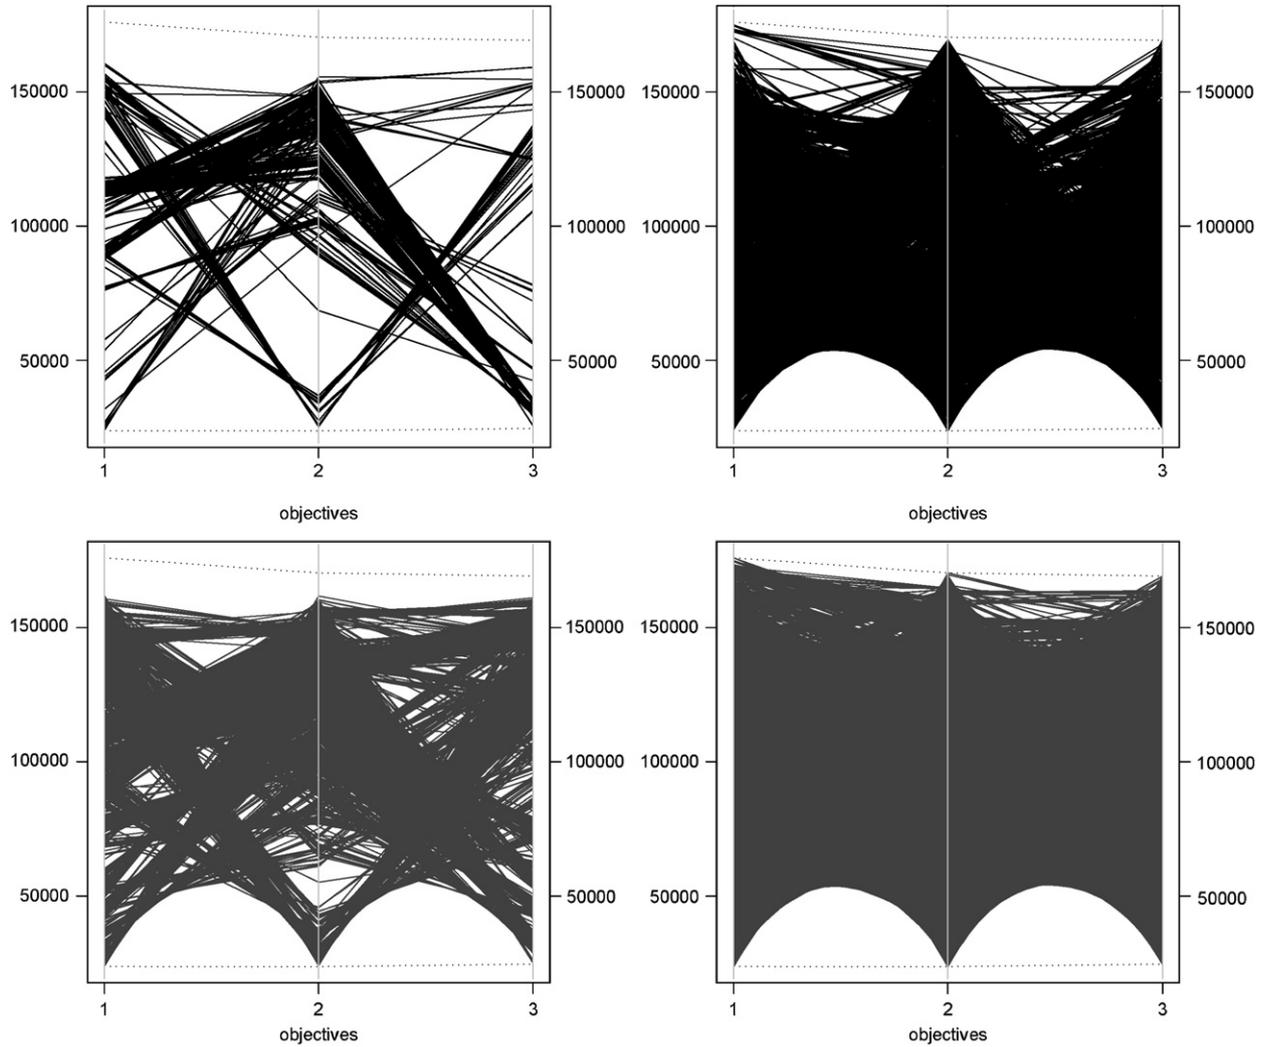


Fig. 8. Location of differences between MOGLS and our configuration on an RUE instance of size 100 with three objectives in favor of the former (left) and in favor of the latter (right) in the range (0.8,1.0] (top) and (0.6,0.8] (bottom). See the text for more details.

our SLS algorithms over MOGLS. In fact, when counting the number of points, for the RUE instance with 100 cities, there are 669 points in favor of MOGLS against 16 484 in favor of our approach in the range of (0.8,1.0] and 8938 in favor of MOGLS against 200 618 in favor of ours in the range of (0.6,0.8]. For the RDM instance of size 100, no differences above 60% were found in favor of MOGLS, and only one point was found whose difference was above 40%. Finally, for the RDM instance of size 300 only differences in favor of our approach were found.

Table 6 presents the average number and standard deviation of evaluations performed by our approach and by MOGLS for each instance. It is possible to observe that our approach performs much less evaluations than MOGLS. Therefore, these results indicate that our approach is highly competitive, both in terms of solution quality and time.

6. Discussion and conclusions

The main goal of this paper is to make a step towards the understanding of the working mechanisms of SLS algorithms applied to MCOPs from a component-based point of view. In fact, SLS algorithms are usually assembled from several components that can, or not, be instantiated for tackling a problem at hand. Hence,

Table 6

The average number and standard deviation of the number of evaluations performed by our approach and by MOGLS for each instance.

Type	Objectives	Size	Our approach	MOGLS
RUE	2	100	$0.04 \times 10^9 \pm 0.03 \times 10^6$	$0.10 \times 10^9 \pm 3.96 \times 10^6$
		300	$0.75 \times 10^9 \pm 1.29 \times 10^6$	$8.00 \times 10^9 \pm 0.85 \times 10^9$
	3	100	$0.41 \times 10^9 \pm 0.21 \times 10^6$	$0.69 \times 10^9 \pm 31.13 \times 10^6$
		300	$0.75 \times 10^9 \pm 3.52 \times 10^6$	$80.91 \times 10^9 \pm 1.36 \times 10^9$
RDM	2	100	$0.07 \times 10^9 \pm 0.26 \times 10^6$	$0.45 \times 10^9 \pm 25.33 \times 10^6$
		300	$0.85 \times 10^9 \pm 3.19 \times 10^6$	$57.66 \times 10^9 \pm 2.00 \times 10^9$
	3	100	$0.40 \times 10^9 \pm 0.26 \times 10^6$	$1.93 \times 10^9 \pm 60.92 \times 10^6$
		300	$1.39 \times 10^9 \pm 7.31 \times 10^6$	$287.21 \times 10^9 \pm 5.63 \times 10^9$

immediate questions for an algorithm designer are: how relevant are these components for the overall algorithm performance? Is there a component that can be removed in order to reduce the fine-tuning effort?

In this article, we focused on SLS algorithms for MCOPs that follow a general algorithmic template, the so-called SAC model. We studied the importance of their components by means of a systematic experimental design using their example application to the biobjective TSP. The experimental results were analyzed using a sound

methodology for the evaluation of the outcomes of SLS algorithms for multiobjective problems.

Our analysis gave clear hints on the effectiveness of each algorithmic component for the MTSP. First, strong intensification for each scalarized problem provides better performance, as shown by the results on the components neighborhood structure and search length. This gives a clear indication that further improvements could be obtained by using iterative improvement algorithms based on more advanced neighborhood structures such as used by the Lin–Kernighan heuristic [34] and its iterated versions [16,41,42]. In fact, initial results by other researchers [48,49] indicate that this conjecture may be true. The component-wise step is always recommendable, at least, for two objectives. For more objectives it may, however, induce an undesirable clustering, which may be circumvented by using more scalarizations; nevertheless, more research in this direction is certainly necessary to give a more detailed answer. Finally, the 2phase strategy seems to be a better option than Restart. This fact is certainly connected to recent results on the closeness of approximate solutions for this problem [50].

As a proof-of-concept, the insights we gained from this analysis were used to assemble SLS algorithms for the MTSP. Despite their simplicity, they showed to be highly competitive with other well-established algorithms for this problem.

There are a number of possibilities for further investigations. More experimental research for this and other MCOPs is certainly required to further increase the understanding of the importance of SLS algorithm components. One methodological aspect that should be treated is to explore other measures for comparing the efficient sets returned by the SLS algorithms such as the hypervolume indicator, the R measure and the ε -indicator. Such measures would certainly speed-up computations in the analysis of the experimental results; at the same time they may incur some loss of relevant information, which is avoided by our methodology. Another direction is to analyze the importance of the algorithmic components in dependence of search space characteristics of MCOPs and the connectedness of the solutions in the efficient set.

Ultimately, we hope that systematic experimental designs and their rigorous analysis will help to make the development of effective SLS algorithms for MCOPs less an art but more a well-established algorithm engineering process.

Acknowledgments

The authors gratefully thank Dr. Andrzej Jaszkiwicz for the source code of MOGLS and Dr. Carlos Fonseca for the source code for computing the EAFs.

References

- [1] Hoos H, Stützle T. Stochastic local search—foundations and applications. San Francisco, CA: Morgan Kaufmann; 2004.
- [2] Ehrgott M, Gandibleux X. Approximate solution methods for combinatorial multicriteria optimization. *TOP* 2004;12(1):1–90.
- [3] Paquete L, Stützle T. Stochastic local search algorithms for multiobjective combinatorial optimization: a review. In: Gonzalez TF, editor. Handbook of approximation algorithms and metaheuristics. Computer and information science series. Boca Raton, FL, USA: Chapman & Hall/CRC; 2007. p. 29.1–29.15.
- [4] Ruiz R, Stützle T. A simple and effective iterated greedy algorithm for the permutation flowshop scheduling problem. *European Journal of Operational Research* 2007;177(3):2033–49.
- [5] Xu J, Chiu SY, Glover F. Fine-tuning a tabu search algorithm with statistical tests. *International Transactions in Operational Research* 1998;5(4):233–44.
- [6] Paquete L, Fonseca CM. A study of examination timetabling with multiobjective evolutionary algorithms. In: Proceedings of the fourth metaheuristics international conference, Porto, 2001. p. 149–54.
- [7] Paquete L, Stützle T, López-Ibáñez M. Using experimental design to analyze stochastic local search algorithms for multiobjective problems. In: Doerner KF, Gendreau M, Greistorfer P, Gutjahr WJ, Hartl RF, Reimann M, editors. Metaheuristics—progress in complex systems optimization. Operations research/computer science interface series, vol.39. Berlin: Springer Verlag; 2007. p. 325–44.
- [8] Zitzler E, Thiele L, Laumanns M, Fonseca CM, Grunert da Fonseca V. Performance assessment of multiobjective optimizers: an analysis and review. *IEEE Transactions on Evolutionary Computation* 2003;7(2):117–32.
- [9] Hansen MP, Jaszkiwicz A. Evaluating the quality of approximations to the non-dominated set. Technical report IMM-REP-1998-7, Institute of Mathematical Modelling, Technical University of Denmark, Lyngby, Denmark, 1998.
- [10] Fonseca CM, Fleming P. On the performance assessment and comparison of stochastic multiobjective optimizers. In: Voigt HM, Ebeling W, Rechenberg I, Schwebel HP, editors. Proceedings of PPSN-IV, fourth international conference on parallel problem solving from nature. Lecture notes in computer science, vol. 1141. Berlin, Germany: Springer; 1996. p. 584–93.
- [11] Shaw KJ, Fonseca CM, Nortcliffe AL, Thompson M, Love J, Fleming PJ. Assessing the performance of multiobjective genetic algorithms for optimization of a batch process scheduling problem. Proceedings of the 1999 congress on evolutionary computation (CEC'99), vol. 1, 1999. p. 34–75.
- [12] López-Ibáñez M, Paquete L, Stützle T. Hybrid population-based algorithms for the bi-objective quadratic assignment problem. *Journal of Mathematical Modelling and Algorithms* 2006;5(1):111–37.
- [13] Paquete L, Stützle T. A two-phase local search for the biobjective traveling salesman problem. In: Fonseca CM, Fleming P, Zitzler E, Deb K, Thiele L, editors. Evolutionary multi-criterion optimization, second international conference, EMO 2003. Lecture notes in computer science, vol. 2632. Berlin, Germany: Springer; 2003. p. 479–93.
- [14] Serafini P. Some considerations about computational complexity for multiobjective combinatorial problems. In: Jahn J, Krabs W, editors. Recent advances and historical development of vector optimization. Lecture notes in economics and mathematical systems, vol. 294. Berlin, Germany: Springer; 1986. p. 222–31.
- [15] Emelichev VA, Perepelitsa VA. On the cardinality of the set of alternatives in discrete many-criterion problems. *Discrete Mathematics and Applications* 1992;2(5):461–71.
- [16] Johnson DS, McGeoch LA. The travelling salesman problem: a case study in local optimization. In: Aarts EHL, Lenstra JK, editors. Local search in combinatorial optimization. Chichester, UK: Wiley; 1997. p. 215–310.
- [17] Angel E, Bampis E, Gourvés L. Approximating the Pareto curve with local search for the bicriteria TSP(1,2) problem. *Theoretical Computer Science* 2004;310:135–46.
- [18] Ehrgott M. Approximation algorithms for combinatorial multicriteria optimization problems. *International Transactions in Operational Research* 2000;7:5–31.
- [19] Angel E, Bampis E, Gourvés L. A dynasearch neighborhood for the bicriteria traveling salesman problem. In: Gandibleux X, Sevaux M, Sörensen K, T'kindt V, editors. Metaheuristics for multiobjective optimisation. Lecture notes in computer science, vol. 535. Berlin, Germany: Springer; 2004. p. 153–76.
- [20] Hansen MP. Use of substitute scalarizing functions to guide a local search base heuristics: the case of moTSP. *Journal of Heuristics* 2000;6:419–31.
- [21] Jaszkiwicz A. Genetic local search for multiple objective combinatorial optimization. *European Journal of Operational Research* 2002;1(137):50–71.
- [22] Borges P. CHESS—changing horizon efficient set search: a simple principle for multiobjective optimization. *Journal of Heuristics* 2000;6(3):405–18.
- [23] Keller CP. Algorithms to solve the orienteering problem: a comparison. *European Journal of Operational Research* 1989;41:224–31.
- [24] Sigal IK. Algorithms for solving the two-criterion large-scale travelling salesman problem. *Computational Mathematics and Mathematical Physics* 1994;34(1):33–43.
- [25] Paquete L. Stochastic local search algorithms for multiobjective combinatorial optimization: methodology and analysis. PhD thesis, Fachbereich Informatik, Technische Universität Darmstadt, 2005.
- [26] Czyzak P, Jaszkiwicz A. Pareto simulated annealing—a metaheuristic technique for multiple objective combinatorial optimization. *Journal of Multi-Criteria Decision Analysis* 1998;7:34–47.
- [27] Gandibleux X, Mezdaoui N, Fréville A. A tabu search procedure to solve multiobjective combinatorial optimization problems. In: Caballero R, Ruiz F, Steuer R, editors. Advances in multiple objective and goal programming. Lecture notes in economics and mathematics systems, vol. 455. Berlin: Springer; 1997.
- [28] Hamacher HW, Ruhe G. On spanning tree problems with multiple objectives. *Annals of Operations Research* 1994;52:209–30.
- [29] Ulungu EL. Optimisation combinatoire multicritère: détermination de l'ensemble des solutions efficaces et méthodes interactives. PhD thesis, Université de Mons-Hainaut, Mons, Belgium, 1993.
- [30] Klingsberg P. A gray code for compositions. *Journal of Algorithms* 1982;3:41–4.
- [31] Stützle T, Hoos HH. Analysing the run-time behaviour of iterated local search for the travelling salesman problem. In: Hansen P, Ribeiro C, editors. Essays and surveys on metaheuristics. Operations research/computer science interfaces series. Boston, MA: Kluwer Academic Publishers; 2001. p. 589–611.
- [32] Steuer RE. Multiple criteria optimization: theory, computation and application. Wiley series in probability and mathematical statistics, New York, NY: Wiley; 1986.
- [33] Bentley JL. Fast algorithms for geometric traveling salesman problems. *ORSA Journal on Computing* 1992;4(4):387–411.
- [34] Lin S, Kernighan BW. An effective heuristic algorithm for the traveling salesman problem. *Operations Research* 1973;21(2):498–516.
- [35] Lourenço HR, Martin O, Stützle T. Iterated local search. In: Glover F, Kochenberger G, editors. Handbook of metaheuristics. International series in operations research & management science, vol. 57. Norwell, MA: Kluwer Academic Publishers; 2002. p. 321–53.

- [36] Grunert da Fonseca V, Fonseca CM, Hall A. Inferential performance assessment of stochastic optimizers and the attainment function. In: Zitzler E, Deb K, Thiele L, Coello CC, Corne D, editors. *Evolutionary multi-criterion optimization (EMO 2001)*. Lecture notes in computer science, vol. 1993. Berlin, Germany: Springer; 2001. p. 213–25.
- [37] Conover J. *Practical nonparametric statistics*. New York, NY: Wiley; 1980.
- [38] Hsu J. *Multiple comparisons—theory and methods*. Boca Raton, FL, USA: Chapman & Hall/CRC; 1996.
- [39] Good PI. *Permutation tests: a practical guide to resampling methods for testing hypothesis*. Springer series in statistics. 2nd ed., New York, USA: Springer; 2000.
- [40] Reinelt G. *The traveling salesman: computational solutions for TSP Applications*. Lecture notes in computer science Berlin, Germany: Springer; vol. 840, 1994.
- [41] Helsgaun K. An effective implementation of the Lin-Kernighan traveling salesman heuristic. *European Journal of Operational Research* 2000;126(1): 106–30.
- [42] Applegate DL, Bixby RE, Chvátal V, Cook WJ. *The traveling salesman problem: a computational study*. Princeton, NJ, USA: Princeton University Press; 2006.
- [43] Johnson DS, McGeoch LA. Experimental analysis of heuristics for the STSP. In: Gutin G, Punnen A, editors. *The traveling salesman problem and its variations*. Dordrecht, The Netherlands: Kluwer Academic Publishers; 2002. p. 369–443.
- [44] Fonseca CM, Fleming P. Genetic algorithms for multiobjective optimization: Formulation, discussion and generalization. In: Forrester S, editor. *Proceedings of the fifth international conference on genetic algorithms*. San Mateo, CA: Morgan Kaufmann; 1993. p. 416–23.
- [45] Ishibuchi H, Murata T. A multi-objective genetic local search algorithm and its application to flow-shop scheduling. *IEEE Transactions on Systems, Man, and Cybernetics—Part C* 1998;28(3):392–403.
- [46] Freisleben B, Merz P. A genetic local search algorithm for solving symmetric and asymmetric traveling salesman problem. In: Bäck T, Fukuda T, Michalewicz Z, editors. *Proceedings of the 1996 IEEE international conference on evolutionary computation (ICEC 1996)*. 1996. p. 616–21.
- [47] Inselberg A. The plane with parallel coordinates. *Visual Computer* 1985;1(4): 69–91.
- [48] Jaszkiwicz A, Zielniewicz P. Pareto memetic algorithm with path relinking for bi-objective traveling salesperson problem. *European Journal of Operational Research* 2009; 193(3):885–90.
- [49] Lust T, Teghem J. Two phase stochastic local search algorithms for the biobjective traveling salesman problem. In: Ridge E, Stützle T, Birattari M, Hoos HH, editors. *Proceedings of SLS-DS 2007, doctoral symposium on engineering stochastic local search algorithms*, Brussels, Belgium, 2007. p. 21–5.
- [50] Paquete L, Stützle T. Clusters of non-dominated solutions in multiobjective combinatorial optimization. In: Barichard V, Ehrgott M, Gandibleux X, T'Kindt V, editors. *Multiobjective programming and goal programming: theoretical results and practical applications*. Lecture notes in economics and mathematical systems, vol. 618. Berlin: Springer, Germany; 2009. p. 69–77.