



Université Libre de Bruxelles

*Institut de Recherches Interdisciplinaires
et de Développements en Intelligence Artificielle*

Iterated Greedy

T. STÜTZLE and R. RUIZ

IRIDIA – Technical Report Series

Technical Report No.
TR/IRIDIA/2018-006

April 2018

IRIDIA – Technical Report Series
ISSN 1781-3794

Published by:

IRIDIA, *Institut de Recherches Interdisciplinaires
et de Développements en Intelligence Artificielle*
UNIVERSITÉ LIBRE DE BRUXELLES
Av F. D. Roosevelt 50, CP 194/6
1050 Bruxelles, Belgium

Technical report number TR/IRIDIA/2018-006

The information provided is the sole responsibility of the authors and does not necessarily reflect the opinion of the members of IRIDIA. The authors take full responsibility for any copyright breaches that may result from publication of this paper in the IRIDIA – Technical Report Series. IRIDIA is not responsible for any use that might be made of data appearing in this publication.

Chapter 1

Iterated Greedy*

Thomas Stützle¹ and Rubén Ruiz²

Abstract

Iterated greedy is a search method that iterates through applications of construction heuristics using the repeated execution of two main phases, the partial destruction of a complete candidate solution and a subsequent re-construction of a complete candidate solution. Iterated greedy is based on a simple principle and methods based on this principle have been proposed and published several times in the literature under different names such as simulated annealing, iterative flattening, ruin-and-recreate, large neighborhood search and others. Despite its simplicity, iterated greedy has led to rather high-performing algorithms. In combination with other heuristic optimization techniques such as a local search, it has given place to state-of-the-art algorithms for various problems. This paper reviews the main principles of iterated greedy algorithms, relates the basic technique to the various proposals based on this principle, discusses its relationship with other optimization techniques, and gives an overview of problems to which iterated greedy has been successfully applied.

Key words: Stochastic local search; metaheuristics; iterated greedy; greedy methods; local search, constructive search

* Preprint of an article that appears in the Handbook of Heuristics, edited by R. Martí, P. Pardalos, and M.G.C. Resende, Springer International Publishing, 2018.

IRIDIA, Université Libre de Bruxelles (ULB), CP 194/6.
Av. F. Roosevelt 50, B-1050 Brussels, Belgium.
E-mail: stuetzle@ulb.ac.be · Grupo de Sistemas de Optimización Aplicada,
Instituto Tecnológico de Informática,
Ciudad Politécnica de la Innovación, Edificio 8G, Acc. B.
Universitat Politècnica de València,
Camino de Vera s/n, 46021, València, Spain.
E-mail: rruiz@eio.upv.es

1.1 Introduction

Many effective algorithms for \mathcal{NP} -hard combinatorial optimization problems rely on the efficient, repeated execution of some simple, underlying mechanisms. A common example are perturbative search methods that iterate over neighborhood searches or iteratively apply underlying iterative improvement procedures. Examples of such methods include simulated annealing, tabu search, iterated local search, memetic algorithms, or dynamic local search (Glover and Kochenberger, 2002; Hoos and Stützle, 2005). In fact, simulated annealing, tabu search or dynamic local search at each step explore the neighborhood of a current solution, while iterated local search and memetic algorithms can be seen as iterating across repeated applications of improvement algorithms. Less frequently, stochastic local search (SLS) methods make use of the iterative application of solution construction algorithms. Examples for these latter methods are ant colony optimization (ACO) (Dorigo and Stützle, 2004; Dorigo et al, 2006), greedy randomized adaptive search procedures (GRASP) (Feo and Resende, 1989, 1995), or the pilot and roll-out method (Duin and Voß, 1999; Bertsekas et al, 1997).

In this chapter, we review another SLS method that relies on the iterative application of solution construction procedures: iterated greedy. The method builds a sequence of solutions by iterating through phases of (partial) solution destruction and subsequent re-construction of a complete candidate solution. A first complete solution in this sequence is generated by some constructive method. Then the following three steps are iteratively executed. First, components of a complete candidate solution s are removed, resulting in a partial solution s_p . Second, starting from s_p , a complete candidate solution s' is rebuilt. Third, an acceptance criterion decides whether to continue this process from s or s' .

Iterated greedy has a clear underlying principle and it is generally applicable to any problem for which constructive methods can be conceived. As such, iterated greedy is clearly a general-purpose method. Iterated greedy is a rather simple method that needs typically only short development times, especially if already a constructive heuristic is available. Iterated greedy provides also a rather simple way of improving over the single application of a constructive method and for various problems very high quality solutions are generated. Additionally, basic versions of iterated greedy do only incur few main parameters and their impact on the search process is rather intuitive to understand. All these reasons make iterated greedy a desirable technique for developers of heuristic algorithms.

Given the simple underlying principle, it is maybe also not surprising that the method that we here call iterated greedy has been (re-)discovered and applied a large number of times under different names by different authors (including ourselves). Algorithms that rely to a significant extent on the same underlying principle have been given names such as simulated annealing (Jacobs and Brusco, 1995), evolutionary heuristic (Marchiori and Steenbeek,

2000), iterative flattening (Cesta et al, 2000), ruin-and-recreate (Schrimpf et al, 2000), iterative construction heuristic (Richmond and Beasley, 2004), large neighborhood search (Shaw, 1998), or, as here, iterated greedy (Hoos and Stützle, 2005; Ruiz and Stützle, 2007). We will review these different developments and other related procedures in Section 1.5.

Despite the possible confusion that may arise for the reader due to the different names and the in part different views on the method, we want to stress that the really important aspect is the principle that underlies all these algorithms. In fact, in all these proposals a repeated usage of constructive methods is made that start from some intermediate, partial candidate solutions. This principle is a generic one of a potential large utility and it should be understood as one of the basic principles that can be used to develop optimization algorithms. This basic principle may also be only one of the principles that is used in the development of a hybrid optimization algorithm that combines elements from different techniques. For example, several algorithms that make use of the iterated greedy principle, include also a local search phase that may improve the solutions generated by the constructive mechanisms (Ruiz and Stützle, 2007).

The structure of the chapter is as follows. In Section 1.2 we review the basic principles of iterated greedy algorithms. Next, in Section 1.3, we give some concrete examples of iterated greedy algorithms. In Section 1.4 we give some results of an experimental study that discusses the main trade-offs in the design of an iterated greedy algorithm for the permutation flow shop problem. We then present other algorithms that make use of the same principle as iterated greedy in Section 1.5. The relationship of iterated greedy to some other methods is discussed in Section 1.6. References to some noteworthy applications of iterated greedy are given in Section 1.7 and we conclude in Section 1.8.

1.2 Iterated Greedy

1.2.1 Greedy construction heuristics

Constructive algorithms build candidate solutions to optimization or decision problems step by step, starting from an empty solution. At each step, they add a solution component to the current partial solution and repeat these steps until a complete candidate solution is obtained. Commonly, constructive algorithms use a heuristic function that estimates for each solution component the benefit of including it into a partial candidate solution. A baseline construction algorithm is formed by so called greedy (constructive) algorithms that at each step add a solution component for which the value of the heuristic function is the best (see also Figure 1.1). If more than one solution component has

```

procedure Greedy Constructive Heuristic:
   $s = \emptyset$ 
  while  $s$  is not a complete solution do
    choose a best rated solution component  $c$ 
     $s = s + c$ 
  end while
end

```

Fig. 1.1 Algorithmic outline of a greedy heuristic.

the same best heuristic value, a tie-breaking criterion is used to decide which solution component is actually added; in the simplest case, this tie-breaking is done uniformly at random but it also may be done by a secondary heuristic function.

Greedy construction heuristics are frequently used when tackling combinatorial optimization problems due to a number of reasons. First, greedy construction heuristics are rather fast and at the same time they generate solutions that are usually much better than those generated uniformly at random or by a randomized but heuristically biased construction. Second, these algorithms are often used to seed (perturbative) local search methods such as iterative improvement algorithms; more sophisticated SLS methods such as tabu search, simulated annealing; or population-based methods such as memetic algorithms. In the latter case, typically some members of the population are generated by greedy constructive methods while others may be randomly generated. Seeding perturbative local search methods with solutions from greedy construction algorithms can incur advantages such as improved quality of local optima, faster identification of local optima, and a better trade-off between computation times and solution quality, that is, better anytime behavior (Zilberstein, 1996). Third, sometimes one can prove guarantees on the quality of the solutions that are generated in the worst case, leading to so-called approximation algorithms. Often, the best provable guarantees that can be obtained even for more complex SLS algorithms are the guarantees that directly stem from those of the initial greedy construction. Fourth, for various polynomially solvable problems, greedy algorithms are also guaranteed to generate optimal solutions, the Kruskal algorithm for minimum spanning trees being a well-known example. However, for \mathcal{NP} -hard problems this is not the case. Finally, they build the basis for a number of other methods such as GRASP (Feo and Resende, 1989, 1995), ACO (Dorigo and Stützle, 2004) or squeaky wheel optimization (Joslin and Clements, 1999).

One straightforward way to improve over the generation of a single greedy solution is in some cases the repeated application of a greedy heuristic to generate a variety of different candidate solutions and then to choose the best one. Obviously, repetition in this sense is only reasonable if in the construction process different solutions can be generated. For example, for the well-known

nearest neighbor heuristic for the traveling salesperson problem (TSP), n distinct nearest neighbor tours may result (assuming no random tie-breaking is done): for each of the n possible cities that may be chosen (randomly) as the initial city for the solution construction, a different tour may result. However, in other cases where the greedy construction is fully deterministic, additional randomization of the construction process, as proposed in the semi-greedy heuristics (Hart and Shogan, 1987) and in GRASP (Feo and Resende, 1989; Resende and Ribeiro, 2010), may be required to generate different solutions.

Repeated construction of solutions also has inherent disadvantages. Constructing a full solution is relatively time-consuming as especially the initial construction steps require a large amount of computation when compared to later construction steps. Furthermore, no information is taken from one solution construction to another one and, thus, such a repeated construction does not exploit knowledge gained from previous solutions. A method that alleviates these problems and that allows to invest, in principle, arbitrary computing times to generate different solutions by constructive heuristics is iterated greedy.

1.2.2 Iterated greedy framework

The main principle of iterated greedy is to iterate over (greedy) construction methods by first generating a complete candidate solution and then cycling through a main loop that consists of two main steps. In the first step, some solution components are removed from the current complete candidate solution s to result in some intermediate partial candidate solution s_p . We call this the *destruction step*. In the next step, starting from s_p , a construction heuristic is used to generate a new complete solution s' . We call this step the *construction step*. An acceptance test then decides from which of the two solutions, s or s' , the next destruction step applies. While in the simplest case, the acceptance test may accept only improved solutions, other choices may lead to more search diversification and, thus, to possibly better results when many iterations of the iterated greedy algorithm are done.

An algorithmic outline of an IG algorithm is given in Figure 1.2. It starts by first generating an initial candidate solution using a procedure `GenerateInitialSolution` and then iterates through a main loop that consists of the application of the three procedures `Destruction`, `Construction`, and `AcceptanceCriterion`. Note that the construction procedures used in `GenerateInitialSolution` and `Construction` may be different and, hence, may also use different greedy heuristics. In the simplest case, however, they may be the same procedures.

As simple default version of an iterated greedy algorithm could use the following choices. As constructive heuristic one may take one that is either already available or implement a known state-of-the-art constructive heuristic. The solution destruction may delete some randomly chosen solution com-

```

procedure Iterated Greedy
   $s_0 = \text{GenerateInitialSolution}$ 
  repeat
     $s_p = \text{Destruction}(s^*)$ 
     $s = \text{Construction}(s_p)$ 
     $s^* = \text{AcceptanceCriterion}(s^*, s')$ 
  until termination condition met
end

```

Fig. 1.2 Algorithmic outline of Iterated Greedy (IG).

ponents, where the number d of solution components to be removed is a parameter of the algorithm. As acceptance criterion, one may force the cost to decrease by only accepting improved or equal quality candidate solutions.

The iterative process around construction heuristics gives IG some specific advantages when compared to the repeated construction of complete candidate solution from scratch. In fact, by starting from a partial solution, one may reduce significantly the time necessary to generate a new candidate solution as less constructive steps need to be done and, the time per construction decision is also reduced as the number of available solution components to choose from is smaller the larger the partial solutions. In addition, through the potential bias exerted by the acceptance criterion, the search process can more easily intensify around the best solutions found in the search process.

When trying to develop a more performing version of an iterated greedy algorithm, many different options for each of the specific operators may be taken into account. Some of the main relevant issues to be considered we discuss in the following. Other ideas will be discussed also later when considering the relationship of iterated greedy to other methods or applications of iterated greedy algorithms.

Destruction There are a number of different possible choices for the solution destruction. A first consideration concerns the number of solution components that should be removed, as defined by a parameter d . The extreme settings would correspond to removing only a single component that is, $d = 1$ or all of them. Even if one may argue that the resulting algorithms should be considered as iterated greedy algorithms, these parameter settings would not correspond to the main ideas underlying iterated greedy. The first case would be akin to a randomized local search, while the latter be akin to a repeated application of a construction heuristic. (We discuss these relationships in more depth in Section 1.6.) Intermediate values of d result in a trade-off between search intensification and diversification: removing a large number of solution components allows to jump to rather distant solutions in the construction phase, while removing few components leads to a more localized search.

Another choice is whether to leave the number of solution components to be removed fixed or variable during the algorithm run. In the case d is left

variable, a scheme of how to adapt its value is required. If the value of d is left fixed, it requires proper tuning. Possibilities for varying the value of d could be to modify this value randomly within some interval every few iterations, to choose the value according to a scheme similar as those introduced for variable neighbourhood search (Hansen and Mladenović, 2001) or to adapt the value of d at computation time, exploiting ideas of reactive search (Battiti et al, 2008).

Once it is known how many components are to be removed, which ones should be chosen? There are a number of possible answers to this question. Intuitively, a randomized destruction is preferable over deterministic choices to reduce the danger of cycling. If a stochastic destruction is used, the simplest case is to choose components uniformly at random. More involved could be choices that take into account cost measures on the components or the partial solutions that remain, thus introducing a bias in the choice. In that case, solution components that have a strong contribution to the cost of a solution would then have a larger probability of being removed than low cost components. Alternatively, one may use lower bounds on the partial solutions resulting after having removed a component: the smaller the lower bound the higher the probability.

Construction One of the crucial ideas of IG is that restoring a fully specified solution is done by some form of (greedily biased) constructive mechanism. Hence, the usage of a constructive mechanism is essential. Using a random perturbation (corresponding to a move to a random solution in a large neighbourhood), is more in the spirit of reduced variable neighbourhood search than in the spirit of IG.

The naming iterated greedy suggests that the construction heuristics used in the algorithm are greedy construction heuristics and typically deterministic (modulo random tie-breaking). In fact, in many implementations of iterated greedy algorithms this is also the case. However, we want to emphasize here that this need not necessarily be the case as, in principle, any suitable constructive mechanism that starting from some partial candidate solution s_p can generate a complete candidate solution—be it deterministically greedy, probabilistically greedy or else—may be used as the underlying construction mechanism in an iterated greedy algorithm.

Among the construction rules, one may distinguish between *adaptive* heuristics and *static* ones. In the first case, the heuristic value assigned to a particular choice in the solution construction depends on the partial solution. Typically, adaptive construction heuristics will result in better quality solutions than static ones; however, this improved solution quality is often reached at the cost of higher computation times.

In the simplest case, the solution construction is done following a deterministic construction heuristic—deterministic except of maybe random tie breaking. Depending on the construction mechanism, it is possible to apply a deterministic rule: by applying stochastic destruction, either the order in

which solution components are added is modified or, if adaptive construction heuristics are used, their heuristic evaluation.

Such basic considerations about solution construction in iterated greedy algorithms can be extended in many natural ways by adopting techniques that have been proposed in other methods. One may take inspiration from other constructive methods and use randomized selection in candidate lists built at each construction step as in GRASP algorithms, use biased constructive decisions exploiting past experience such as the pheromone trails in ACO, or use prohibitions by avoiding adding again solution components that have been removed in the solution destruction, taking ideas from tabu search. Another option may be to choose among different possible construction rules and use ideas from the hyperheuristics community for this task (Burke et al, 2013). In fact, all kind of methods and techniques that are compatible with the idea of constructing solutions may be adopted within iterated greedy algorithms if this seems promising, making iterated greedy in this sense also a very flexible technique.

Acceptance criterion. The acceptance criterion has a strong influence on the diversification / intensification behaviour of an IG algorithm. On the extreme cases are the possibilities of accepting any new solution independent of its solution quality or to only accept solutions that improve over the previous one. There are many intermediate choices such as occasionally accepting worse solutions or allowing backtracks to previously seen solutions. A popular choice when accepting worse solutions is to use acceptance criteria from simulated annealing such as the Metropolis criterion: if a new solution is better or equal, it is accepted; otherwise it is accepted with a probability $\exp\{(f(s) - f(s'))/T\}$, where T is a parameter called temperature. More elaborated approaches might probably make use of short-term memory as in tabu search. The acceptance criterion, whatever it might be, does not need to be applied at every iteration, i.e., a given incumbent solution can be destructed and reconstructed a number of times before deciding on its value. For the choice of the acceptance criterion very much the same issues arise as in iterated local search (Ramalhinho Lourenço et al, 2002) and an appropriate choice may be crucial to an IG algorithm's performance.

Hybridization with local search. Iterated greedy algorithms can form directly the basis of hybrid algorithms that combine various search mechanisms. In fact, for constructive heuristics a natural extension is to improve the generated solutions by the application of a (perturbative) local search method, in the simplest case this being an iterative improvement algorithm. Such an extension is also straightforward to be adopted in an iterated greedy algorithm and actually has been done in a number of such approaches. This extension results in an outline of iterated greedy as given in Figure 1.3. With this additional local search phase, the IG algorithm also strongly resembles iterated local search (ILS) algorithms (Ramalhinho Lourenço et al, 2002). In fact, the destruction and construction phases implement a solution perturbation

```

procedure Iterated Greedy with local search
   $s_0 = \text{GenerateInitialSolution}$ 
   $s^* = \text{LocalSearch}(s_0)$ 
  repeat
     $s'_p = \text{Destruction}(s^*)$ 
     $s = \text{Construction}(s'_p)$ 
     $s' = \text{LocalSearch}(s)$ 
     $s^* = \text{AcceptanceCriterion}(s^*, s')$ 
  until termination condition met
end

```

Fig. 1.3 Algorithmic outline of an IG with an additional local search step.

in the ILS sense. However, a minor difference is that ILS often makes use of perturbations that are randomly chosen from some large neighbourhood, while in IG algorithms an underlying constructive method is exploited. More importantly, IG can also reach very high performance when used without the local search phase, which is not necessarily true for ILS algorithms. More on relations of IG to other methods is given in Section 1.6.

In any case, for an algorithm to qualify as an iterated greedy algorithm, it is necessary that one can distinguish between a clearly available construction mechanism and a clear destruction mechanism that are repeatedly applied in an alternating order. An acceptance criterion may be used in an explicit way or in an implicit way; the latter is the case, for example, if every new solution is accepted and no mentioning of an acceptance criterion is done. In that case, the acceptance criterion would correspond to a “random walk” type acceptance criterion used sometimes in ILS algorithms (Ramalhinho Lourenço et al, 2002).

1.3 Some simple examples of iterated greedy algorithms

Let us consider a few examples of constructive heuristics for basic combinatorial problems that will serve throughout the chapter for illustrating various details of IG algorithms. We consider here three examples for the traveling salesman problem (TSP), the set covering problem (SCP), and the permutation flow shop problem (PFSP). We first shortly introduce the problem and then describe a basic constructive heuristic and basic IG algorithms.

1.3.1 TSP example

Traveling salesman problem (TSP). In the TSP is given a graph $G = (N, E)$ where N is the set of $n = |N|$ nodes and E is the set of edges that fully connects the nodes. To each edge (i, j) is associated a distance d_{ij} . Here we assume that the distance matrix is symmetric, that is, we have $d_{ij} = d_{ji}$ for all $(i, j) \in E$; this type of TSP instances are called symmetric and are among the most widely studied types of TSP instances. The objective in the TSP is to determine a Hamiltonian cycle of minimal length. Such a cycle can be represented by a permutation $\pi = \langle \pi(1), \dots, \pi(n) \rangle$, where $\pi(i)$ is the node index at position i . The objective function to be minimized is

$$\min_{\pi \in \mathcal{S}} d_{\pi(n)\pi(1)} + \sum_{i=1}^{n-1} d_{\pi(i)\pi(i+1)} \quad (1.1)$$

where \mathcal{S} is the search space consisting of the set of all permutations.

Probably the best known constructive heuristic for the TSP is the nearest neighbor heuristic. It chooses randomly a first node to start from to obtain a partial tour $\langle \pi(1) \rangle$. Then at each step it appends to the partial tour $\langle \pi(1) \dots \pi(l) \rangle$ a still unvisited node that has minimum distance to $\pi(l)$. Once all nodes are visited the tour is completed by closing the tour and going back to $\pi(1)$. The nearest neighbor heuristic usually contains subtours that may be close to optimal ones but also contains long edges that are added often towards the end of the construction. Empirically, for Euclidean two-dimensional TSP instances, the nearest neighbor heuristic generates tours that are about 20 to 40% above optimal (see Johnson and McGeoch (2002)).

One possibility to build an iterated greedy algorithm on top of the nearest neighbor heuristic would be to remove a subtour of d consecutive cities of a tour, resulting in one subtour of $n - d$ cities and to restart the nearest neighbour heuristic from it. Another possibility would be to remove randomly cities, reconnect the remaining subtours and then restart the nearest neighbor heuristic. However, it is unclear what the performance of such an algorithm would be.

Another simple constructive heuristic for the TSP is the random insertion heuristic (RIH). It starts by choosing randomly two nodes that are ordered arbitrarily. At each construction step, a next node is chosen randomly and inserted into a position in the current path such that the cost increase is minimum. The RIH is one of simplest heuristics for the TSP; however, it is among the best performing constructive algorithms for the TSP (Johnson and McGeoch, 2002).

IG can be used on top of the RIH as follows. The first solution is constructed by RIH. The procedure `Destruct` would remove l nodes that are chosen uniformly at random. Next, `Construct` adds the removed nodes using again the same rules as in the RIH. Finally, `AcceptanceCriterion` could be chosen to only accept improved solutions. This means that the solution destruction would be applied to the best tour found so far. Different choices for iterated greedy

algorithms based on such steps have also been experimentally examined by Schrimpf et al (2000).

1.3.2 SCP example

Set covering problem (SCP). In the set covering problem (SCP) is given a set $A = \{a_1, \dots, a_n\}$ of items and a set $B = \{B_1, \dots, B_m\}$ of subsets of elements of A that covers A ; in other words, we have that for each $B_i \subseteq A$ and $\bigcup_{i=1}^m B_i = A$. A set B_i covers an item a_j if $a_j \in B_i$. Each set B_i has a cost of c_i . The objective in the SCP is to find a subset C of the sets in B that covers each element in A and that is of minimal total cost.

For the SCP, constructive heuristics differ in the choice of the heuristic function that is used. A standard heuristic function is to compute the cover value of the sets in B , which is defined as $\gamma_i = c_i/b_i$, where b_i is the number of items that would be covered when adding subset B_i to a current partial cover. Note that when starting with an empty solution, that is, none of the subsets B_i is chosen to be in the cover C , we have that $b_i = |B_i|$. However, each time a subset is chosen, b_i needs to be updated taking into account the items already covered in a partial cover C_p . The cover value is an example of an adaptive heuristic, where the heuristic values depend on the partial solution already generated. Adaptive heuristics require higher computation times than their corresponding static heuristics, which do not update heuristic values in dependence of the partial solution, but typically lead also to better quality solutions. Once a complete cover is obtained, some of the subsets may have become redundant if all the items they cover are also covered by other subsets; removing such redundant subsets then improves the solution cost of the generated cover. Note that the removal of redundant subsets is not yet a destruction step as through the removal of redundant subsets the candidate solution remains a complete cover. The destruction step, then may remove a number of subsets from a complete solution, resulting in a partial cover C_p , from which again the construction heuristic starts.

A first IG algorithm for the SCP has been proposed by Jacobs and Brusco (1995). They construct the first solution using a simple greedy construction heuristic (Balas and Ho, 1980) that at each step first selects a random, not covered item and adds the least cost subset that covers that item. Once a complete cover is found, redundant columns are removed. The solution destruction removes $k_1 \cdot |C|$ subsets that are chosen uniformly at random; here $0 < k_1 < 1$ is a parameter and $|C|$ is the number of subsets on the cover, the cover size. The solution construction uses the heuristic based on cover values explained above. (The iterated greedy algorithm by Jacobs and Brusco (1995) is, hence, a first example where the constructive heuristics for generating the initial solution and for extending partial solutions in the main loop differ.) However, the subsets to be considered are limited to a candidate set that comprises all those subsets that have a cost less than $k_2 \cdot \max\{c_i | i \in C\}$,

where $k_2 > 0$ is a parameter that influences the size of the candidate set and C is the current cover. At each step of the construction, a subset with a minimum cover value is added to the current partial solution, breaking ties uniformly at random. Once a complete cover is obtained, redundant columns are removed as explained above. The acceptance criterion of a new cover C' is based on the Metropolis condition that is frequently used in simulated annealing algorithms. This simple IG algorithm for SCP obtained very high performance improving over several earlier proposed SCP heuristics.

1.3.3 PFSP example

Permutation flow shop scheduling problem (PFSP). In the PFSP, n jobs have to be scheduled on m machines. All jobs visit the machines in the same order, each job having an operation at each machine. Time p_{ij} denotes the non-negative, known and deterministic processing time that job j needs on machine i . In the PFSP the same processing sequence of the jobs is maintained throughout all machines and, hence, the processing sequence is obtained as a permutation of the jobs. The standard objective is to minimize the completion time of the last job in this order, which is also known as makespan (C_{\max}). The PFSP arises in many practical situations as it is common to have production lines where machines are disposed in series. The PFSP is a thoroughly studied problem in the scheduling literature with literally hundreds of papers published each year. The minimization of the makespan is the most used criterion in the literature, but not so in practice (Framiñán et al, 2014). Some reviews on the PFSP are Framiñán et al (2004); Hejazi and Saghafian (2005); Ruiz and Maroto (2005).

A popular and high performing constructive heuristic for the PFSP is the NEH heuristic (Nawaz et al, 1983), named after the initials of the last names of the paper's authors. NEH is an insertion heuristic, a kind of heuristics that iteratively extends a permutation by inserting at each step one new element into the current partial solution. The NEH heuristic first computes for each job j its sum of the processing time $P_j = \sum_{i=1}^m p_{ij}$ and then orders them according to non-increasing P_j values, resulting in a sequence $\phi(1) \dots \phi(n)$. The first two jobs according to that order are taken, their two possible sequences $\langle \phi(1)\phi(2) \rangle$ and $\langle \phi(2)\phi(1) \rangle$ are evaluated and the better of the two is adopted. Then at each step l , the job $\phi(l)$ is considered and tentatively inserted in all possible l positions in the current partial solution $\pi(1) \dots \pi(l-1)$. Among these tentative insertions, the one resulting in the least increase of the makespan is taken. These steps are repeated until all jobs are inserted. Note that there might be two levels of ties, both at the ordering of the P_j values as well as at the insertion phase. The NEH heuristic has a computational complexity of $O(n^3m)$ which is lowered to $O(n^2m)$ using the efficient implementation of Taillard (1990). There is a rich literature of methods that propose variants of the NEH heuristic, mainly proposing mechanisms to break ties or reinsertions. As shown by Ruiz and Maroto (2005) and more recently by Fernandez-Viagas

and Framiñán (2014), the NEH is a state-of-the-art constructive heuristic for the PFSP.

A simple IG algorithm for the PFSP has been proposed by Ruiz and Stützle (2007). It is based on insertion heuristics such as NEH. In fact, this iterated greedy algorithm uses NEH to construct an initial solution. At each destruction step it removes a number of jobs that are chosen uniformly at random from the current permutation. In the construction step, these jobs are then re-inserted in the same order in which they have been removed. For the insertion of the jobs, the same procedure as described for the NEH heuristic is followed. The acceptance criterion accepts a new candidate solution using the Metropolis condition known from simulated annealing algorithms but with the temperature T set to a constant value. The performance of this simple IG algorithm was very good, outperforming many metaheuristic algorithms for the PFSP. When combined with a local search phase, the proposed IG algorithm was shown to be a new state-of-the-art algorithm for the PFSP (Ruiz and Stützle, 2007). The IG implementation of Fernandez-Viagas and Framiñán (2014) has given some further improvements of the above presented iterated greedy algorithm. Currently, the top-performing iterated greedy algorithm for the PFSP is the variant of Dubois-Lacoste et al (2017), which additionally adds a local search on the partial solution that is obtained after the destruction step. As a matter of fact, this IG method has shown to outperform other more recent, and arguably much more complex, approaches.

1.4 Case study: iterated greedy for flow shop scheduling

The development of an effective iterated greedy algorithm requires the algorithm designer to choose which of the various possible implementation choices to take. In this section, we exemplify the development of an iterated greedy algorithm for the PFSP and study the impact specific alternative choices have on iterated greedy performance. Here we examine the impact of various alternative choices for the iterated greedy algorithm by Ruiz and Stützle (2007) that we presented in the previous section.

In particular, we consider the following design choices in our experimental analysis.

Initial solution. It may seem advisable to start from an as good initial solution as possible. For the PFSP, this would be the NEH heuristic (Ruiz and Maroto, 2005; Fernandez-Viagas and Framiñán, 2014) using the accelerations of Taillard (1990). However, at least for some instances it is not clear whether a random or heuristic initialization of iterated greedy provides the best results. In fact, Rad et al (2009) study cases in which only for hard instances and in short CPU times it is advisable to use NEH (or extensions of the NEH) to obtain competitive results.

Destruction strength. A first and foremost decision to be taken is how many elements of a complete candidate solution should be removed. Here, we consider different fixed values for this parameter d , although it may also be interesting to consider schemes of how to vary the value of d at run-time.

Destruction type. The type of destruction determines which components of the incumbent solution are removed. In the simplest case, one may remove components of a solution uniformly at random. Alternatively, one may remove blocks of consecutive elements. In this case, a random block of d jobs could be removed from the sequence. Note that for a same solution π and a block-based destruction, there are only $n - d + 1$ possible choices for the destruction move. These two latter possibilities will be examined here. However, one may consider a biased destruction, where components of the incumbent solution may be chosen randomly but in a biased way. For the PFSP, for example, jobs generating a large idle time at machines before or after their processing might be disrupting the sequence and are therefore likely to have been misplaced. We leave the study of such a destruction operator for future work.

Construction type. Similarly to the type of destruction move, different ways of how to construct a solution could be examined. The default choice in the iterated greedy algorithm is to insert jobs using the NEH heuristic. As an alternative, we consider a random insertion of the removed jobs, which could be used if no efficient greedy heuristic is available. However, in this case, the difference between iterated greedy and ILS becomes somewhat blurry, especially when additional local search is used.

Acceptance criterion. The acceptance criterion has a direct impact on the balance between intensification and diversification of the search. A simple idea is to accept only better quality solutions, while alternatively one may use the Metropolis condition that occasionally also accepts worse candidate solutions.

Acceptance iterations. Instead of applying the acceptance criterion at each iteration, one may apply an acceptance criterion only each l iterations. This would correspond to accepting for a few iterations every new candidate solution that is generated and only applying the acceptance test after each sequence of l steps.

Local search. Finally, it is well known that local search can have a tremendous impact on the quality of the results achieved, even though iterated greedy algorithms may reach high quality solutions even without local search. Hence, it may be worthwhile to test the impact of an additional local search phase.

In order to study the different design and implementation alternatives, we have carried out a Design of Experiments (DOE) approach (Montgomery, 2012) where the previous factors are analyzed. Seven factors are coded and tested at two levels each, as summarized in Table 1.1. As a result, a total of $2^7 = 128$ combinations are to be tested in a full factorial experimental design. However, we use a half fractional design, which has most of the power of a

Factor	abbreviation	level one	level two
Initial solution	<i>Initialization</i>	NEH	Random
Destruction strength	<i>Destruct</i>	4	6
Destruction type	<i>Destruction_T</i>	Random	Block
Construction	<i>Reconstruction_T</i>	NEH	Random
Acceptance criterion	<i>Acceptance_C</i>	SA	Descent
Acceptance iterations	<i>Iterations_Acc</i>	1	5
Local search	<i>LS</i>	no	yes

Table 1.1 Summary of the factors and the levels studied in the experimental analysis.

full factorial one but requires only half of the runs. We use a 2^{7-1}_{VII} design, which has a high resolution *VII*: interactions between four factors are aliased (confounded) with interactions between three factors, interactions between five factors aliased with interactions between two factors and so on. As a result, we can safely study the experimental data since it is highly unlikely that two high level interactions might be significant. More elaborated techniques for the analysis and calibration of algorithms than the simple exploratory analysis we perform here have been published elsewhere (Bartz-Beielstein et al, 2010). Each studied combination has been tested on 30 of the hardest instances from Taillard’s benchmark (Taillard, 1993). This benchmark is composed of 12 groups of 10 instances, each ranging from 20 jobs and 5 machines to 500 jobs and 20 machines. The instances of one particular group are denoted as $n \times m$. For the tests we pick three of the hardest groups, namely 50×20 , 100×20 and 200×20 . We are interested in these instances since for most of the remaining problems, the optimum solution is already known and today’s state-of-the-art methods are capable of obtaining near optimum solutions.

Each combination is run 5 independent times (replicates) with each instance. We also control the number of jobs (n) as a blocking factor with three levels in the experiment. Therefore, the experimental design contains $2^{7-1} \cdot 3 = 192$ treatments, $192 \cdot 10 = 1920$ experimental units (for ten instances) and $1920 \cdot 5 = 9600$ experiments (for five replicates).

The basic iterated greedy algorithm along with all studied alternatives has been implemented in Delphi language and run during a predefined CPU time which is fixed to 30, 60 and 120 seconds for the instances of 50, 100 and 200 jobs, respectively. The machine used for the tests is a Pentium IV PC/AT computer running at 3.2 GHz with 2 GBytes of RAM memory. The performance measure and response variable of the experimental design is the so called relative percentage deviation (*RPD*) over the optimum or best known solution (upper bound) for each tested instance:

$$\text{Relative Percentage Deviation (RPD)} = \frac{Heu_{sol} - Best_{sol}}{Best_{sol}} \times 100 \quad (1.2)$$

where Heu_{sol} is the solution given by any of the replicates of any iterated greedy variation for a given instance and $Best_{sol}$ is the optimum solution

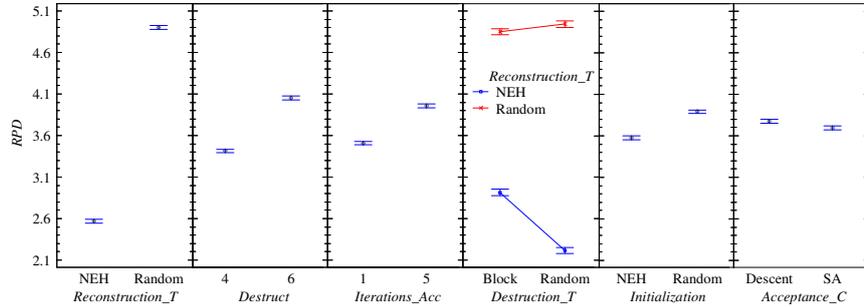


Fig. 1.4 Means plot of the average Relative Percentage Deviation (RPD) and 99% Tukey HSD confidence intervals for the studied factors. Simple iterated greedy without local search.

or the lowest known upper bound for that specific Taillard instance as of November 2015. The results are analyzed by means of the parametric ANOVA technique. Notice that in this case, the three assumptions of this parametric test (normality, homoscedasticity and independence of the residual) have to be satisfied. In our experiment, the factor LS results to be extremely significant and it creates large differences in the response variable. As expected, applying the local search results in a statistically significant difference in the average RPD of 1.46 considering all results, compared with the average RPD of 3.73 without local search. Such a large difference generates normality problems. In order to avoid this situation, we study two separate ANOVAs, one for each level of the LS factor.

Results of the simple iterated greedy without local search

All remaining factors after fixing the local search have p-values very close to zero in the resulting ANOVA table. As a result, we focus on the F-Ratio, which is the ratio between the variance generated by a given factor and the residual variance in the studied 2-level interaction linear model. The higher this ratio, the more significant the factor or interaction is. Figure 1.4 shows the means plots for the remaining six factors in order of importance.

From Figure 1.4 we can observe how all six studied factors have levels and variants that result in statistically significant differences. Observed differences in the average performance of the iterated greedy when a factor has been set to a given level or variant are depicted with non-overlapping confidence intervals in the plots. The importance of the F-Ratio is shown in Figure 1.4 with the most significant factors to the left of the figure and the least significant to the right. As can be seen, without local search, it is clearly preferable to greedily reconstruct the solution as the difference in performance is extremely large.

This is one of the main keys behind the iterated greedy algorithm. A random reconstruction of jobs means that first the jobs are randomly extracted and then randomly inserted. This is more or less a sequence of insertion moves. Too many moves (and no local search) result in an algorithm that is not performing well.

The second factor in importance is *Destruct* or d , the number of jobs to be removed in the destruction phase. Following the previous discussion, six jobs impose a rather large overall disruption and removing only four jobs gives much better results. Notice that a similar conclusion was reached by Ruiz and Stützle (2007) in their calibration of the iterated greedy for the PFSP. Although not shown here, removing four jobs is better for all the studied instances with 50, 100 and 200 jobs.

The third most important factor results to be *Iterations_Acc*. Among the two tested levels, applying the acceptance criterion at every iteration gives substantially better results. The analysis is continued until all factors have been fixed through the most important effects in the response variable. For example, the third most significant effect in the ANOVA is not a single factor but the interaction between factors *Destruction_T* and *Reconstruction_T*. It is shown in the fourth plot from the left in Figure 1.4. We see that both factors interact greatly. When *Reconstruction_T* is set to Random, the Block destruction appears to be slightly better. However, when *Reconstruction_T* is set to NEH, as in the original NEH method, the destruction has to be done randomly as much better results are obtained. *Initialization* is best set at NEH but we observe that the difference in performance with the Random initialization is rather small. Finally, the acceptance criterion factor (*Acceptance_C*) is also slightly better at SA but the difference with Descent is very small, even though it is statistically significant. More specifically, the average *RPD* given by the simple iterated greedy across all factors with a SA-like acceptance criterion is 3.69% whereas for the Descent criterion it is 3.77%. It is interesting to note that most results confirm the choices made in previous studies as in Ruiz and Stützle (2007).

Results of the simple iterated greedy with a local search step

Now we proceed with the analysis of the experimental data after focusing on the results of the simple iterated greedy with the local search enabled. In this experiment, the first interesting result is that the observed differences as regards *RPD* among the levels of the factors are much smaller than in the previous one. It is safe to say that the local search is capturing much of the variability of the experiment. Figure 1.5 contains the means plots for the six factors, in order of importance according to their F-Ratios.

As can be seen, the relative importance of the studied factor is not the same when compared to the previous experiment without local search. The most important factor is now *Iterations_Acc*. Similar to the previous exper-

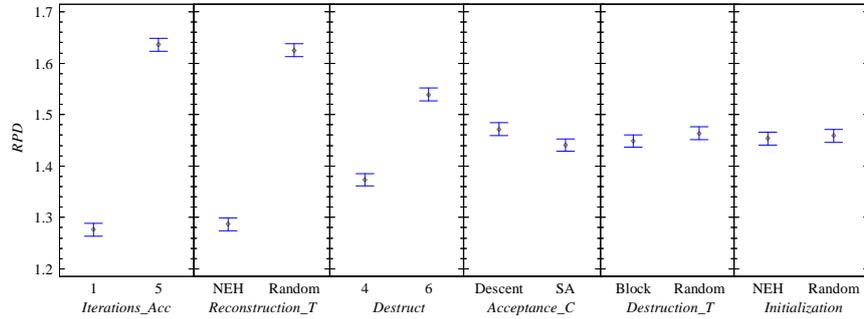


Fig. 1.5 Means plot of the average Relative Percentage Deviation (RPD) and 99% Tukey HSD confidence intervals for the studied factors. Simple iterated greedy with local search enabled.

iment, applying the acceptance criterion at each iteration improves results significantly. The next factor in importance is *Reconstruction_T*. The plot is similar to the previous experiment and represents an important result. Considering that the algorithm already carries out a local search step after reconstruction, it could be argued that a random reconstruction might suffice. However, this is not the case. Carrying out a local search step from a randomly re-constructed solution yields, on average, much worse results than if the local search is carried out on a greedily (and therefore, of higher quality) re-constructed solution. This supports the point that iterating over greedy heuristics gives very good results. The remaining factors per importance are *Destruct* and *Acceptance_C*. The last two remaining factors, *Destruction_T* and *Initialization*, are not statistically significant at a 99% confidence level as it can be seen by the fact that their corresponding means plots overlap. Hence, among the two tested types of destruction, it is not important how the jobs are removed. This is probably due to the fact that even though jobs are removed in blocks, they are re-inserted one by one. As for the initialization, the advantage gained by an NEH initialization is nullified during the run of the iterated greedy algorithm, confirming the observation of Rad et al (2009).

As a conclusion from this study, the most significant factor is the local search and the NEH reconstruction. Most other factors have less importance. iterated greedy can even do without a full NEH initialization for the hardest instances of Taillard in the PFSP problem.

1.5 IG applications: Historical development

Among the examples of iterated greedy algorithms in Section 1.3, we had described an algorithm for the SCP by Jacobs and Brusco (1995) who have

actually called their algorithm a *simulated annealing* algorithm. This naming is probably due to the use of the Metropolis condition as an acceptance criterion, which was also used in the first proposals of simulated annealing (Kirkpatrick et al, 1983). This algorithm has later been extended by the same authors (Brusco et al, 1999). Somehow related is the more complex algorithm by Marchiori and Steenbeek (2000) for the SCP; in this algorithm a construction starts from a partial solution following the steps of the iterated greedy method; however, instead of obtaining the partial solution by choosing the subsets to be removed, they use a mechanism to choose the subsets of the best solution found so far that are maintained in the partial solution. Maybe surprisingly, in that article the proposed algorithm is viewed as an evolutionary algorithm. This algorithm is preceded by another algorithm for the unweighted SCP (Marchiori and Steenbeek, 1998), which more directly follows the ideas of the iterated greedy method.

Interestingly, the above mentioned algorithms for the SCP are not the only ones that are directly based on the same destruction–construction process that is the basis of iterated greedy algorithms. In fact, a large number of other algorithms are based on the same principle. However, different names have been used for the underlying method or no specific name at all, which complicates the transfer of knowledge to other researchers.

Probably the earliest techniques that use some mechanism akin to the destructive–constructive moves of iterated greedy are found in the VLSI design automation community and, in particular, in the routing step, where components are connected by wires trying to obey design rules for integrated circuits. In the rip-up and reroute approach (Rubin, 1974; Dees, Jr. and Karger, 1982) some routes related to bottlenecks are removed and components are reconnected in a different order. The rip-up and reroute technique is commonly used in the design automation community but the principle underlying this method seems not to have been generalized to the level of metaheuristic for tackling a wide variety of other optimization problems.

The possible use of destructive–constructive moves has also been mentioned earlier in the context of the type of techniques that have become known as strategic oscillation (Glover, 1977, 1986, 1989). These type of approaches are often embedded into algorithms that make use of tabu search features to provide additional guidance to the search. Some examples of such implementations that have also a clear component related to the iterated greedy methodology are Glover and Kochenberger (1996); Glover et al (1998); Lozano et al (2014).

Among the first researchers to identify the potential of the principles underlying iterated greedy and to formulate these as a general-purpose SLS method are Schrimpf et al (2000). They called their method *ruin-and-recreate*, where the *ruin* stands for the solution destruction and the *recreate* for the re-construction of a complete solution.³ In their seminal article they applied

³ Ruin-and-recreate is protected by US patent *Optimization with ruin recreate* No. 6418398; see <http://www.patentstorm.us/patents/6418398-fulltext.html>.

ruin-and-recreate to the symmetric TSP, the vehicle routing problem with time windows and to a network design problem, reporting good overall performance of the method. Following this work, a number of other papers have been published using the name ruin-and-recreate, however, not all papers that use the name ruin-and-recreate should be considered as iterated greedy algorithms; for example, Misevičius (2003a,b) essentially uses a random mutation of a current solution instead of a clearly separable destruction phase and greedy solution re-construction.

Ahmadi and Osman (2004) study the Capacitated Clustering p -Median Problem. Among different proposed methods and heuristics, the authors introduce a *periodic construction-deconstruction* procedure, which is basically a simple form of iterated greedy. After a given number of clusters have been built in the constructive method, some of them are randomly deconstructed and the constructive method is re-applied.

Richmond and Beasley (2004) have presented an *iterative construction heuristic* for a problem arising in ore selection. In this problem, processing options need to be chosen for a number of mining blocks. The destruction operator deletes the processing option for some mining blocks; the size of the destruction is chosen randomly following a uniform distribution in some interval. After the re-construction of a complete solution, the algorithm accepts only better quality solutions as new incumbent ones. It should be clear from the description that the algorithm follows directly the very same principles of iterated greedy.

Cesta et al (2000) describe an algorithm for a capacitated scheduling problem that makes usage of the ideas underlying the concept of iterated greedy algorithms. They have developed a method that uses a greedy constructive heuristic to generate feasible schedules by introducing precedence constraints. In the destruction phase (called retraction in the article), they remove some of the introduced precedence relations to obtain a partial solution and then reconstruct a complete feasible one. They call their method *iterative flattening*, where flattening refers to the reduction of resource conflicts, which is an effect of the introduced precedence constraints. The term iterative flattening has gained some popularity and several follow-up articles have been published on improvements of the algorithm (Michel and van Hentenryck, 2004), studies of combinations of basic algorithm components (Oddi et al, 2008) or applications to other problems (Oddi et al, 2010, 2011).

The authors of iterative flattening characterize the method in their initial paper as a local search method. Few years before, Shaw (1998) has proposed the *large neighborhood search* method. The initial proposal consists in the removal of solution components and the re-insertion of solution components by using constraint programming techniques that exploit a tree search and constraint propagation techniques. Even though the method was proposed in a constraint programming framework, where a tree search is used to restore a complete solution instead of simple constructive heuristic, the latter is mentioned as one possible option. In fact, in many later articles on large

neighborhood search, simple constructive heuristics are used (Pisinger and Ropke, 2010), implementing, hence, directly an iterated greedy method.

Finally, the name *iterated greedy* has been used by Ruiz and Stützle (2007) and also in Hoos and Stützle (2005) to denominate the type of SLS method we describe in this chapter. However, these are not the first publications that are using the name *iterated greedy*. A much earlier mentioning of the term iterated greedy is by Culberson (1992). His iterated greedy algorithm for graph colouring uses a greedy algorithm for generating a colouring, removing the colour of all nodes and then reapplying the greedy algorithm using a specific ordering of the nodes that guarantees the next generated colouring to be not worse than the previous one but potentially better. The algorithm by Culberson (1992) may be seen as an extreme case of iterated greedy where the destruction operator destroys the complete solution. However, our view is rather that it is a specific algorithm that works in the way it is proposed mainly for graph colouring.

1.6 Relationship to other approaches

Iterated greedy is a general-purpose SLS method that has many links to other SLS methods. In the following, we discuss similarities and differences to a number of SLS methods considering other constructive methods, local search methods, and tree search algorithms. These relationships also open up many possibilities for combining techniques proposed for different methods.

1.6.1 Repeated (greedy) construction algorithms

Iterated Greedy has natural connections to other methods that repeatedly construct complete candidate solutions. A key difference between Iterated Greedy and few other repeated construction methods is that the solution construction in Iterated Greedy usually starts from a (non-empty) partial solution while the other methods repeatedly generate new candidate solutions starting from *empty* candidate solutions.

A well known simple such method is the greedy randomized adaptive search procedures (GRASP) that combines a greedily biased but randomized solution construction with a subsequent improvement of the generated candidate solutions by a local search procedure (Feo and Resende, 1989; Resende and Ribeiro, 2010). GRASP in turn extends on some initial ideas of how to generate different solutions by a randomized greedy heuristic (Hart and Shogan, 1987). While GRASP does not start the construction from partial solutions, it is an example of a (simple) randomization scheme during the solution construction, which could be included directly in iterated greedy algorithms. An appealing

feature of GRASP is the usage of adaptive constructive procedures and the large number of articles on GRASP (Resende and Ribeiro, 2010) may be a source for randomized greedy heuristics to be used in iterated greedy algorithms.

Many useful ideas that underlie other constructive methods may be adopted into iterated greedy algorithms. One such possibility are ways of how to improve the constructive mechanisms, for example, through look-ahead methods or, when taking look-ahead to an extreme, as advocated in the roll-out (Bertsekas et al, 1997) and pilot method (Duin and Voß, 1999). In the roll-out and pilot methods for each constructive decision to be taken, a full solution is constructed (or at least well approximated) and the construction decision that results in the best complete solution is taken as the next one. This process is repeated for each construction decision, resulting in a relatively time-consuming constructive approach. However, starting the roll-out/pilot method from partial solutions with the resulting reduction in computation time, may make this approach promising inside an iterated greedy algorithm.

Squeaky wheel optimization (SWO) (Joslin and Clements, 1999) uses the idea of biasing the solution construction by information that was gained by analyzing complete candidate solutions. This is done by assigning priorities to specific solution components. The solution re-construction in SWO is done, differently from iterated greedy, from empty initial solutions corresponding to a complete destruction of the current candidate solution. Within the SWO framework, Aickelin et al (2006) proposed the idea of seeding the re-construction of a full candidate solution by some partial solution that was obtained by removing some of the solution components of a complete candidate solution, making it a variant of iterated greedy.

A rather different way of biasing the solution construction is underlying the ant colony optimization (ACO) metaheuristic (Dorigo and Stützle, 2004; Dorigo et al, 2006). In ACO, (artificial) ants implement stochastic solution construction heuristics that make their constructive decisions based on so-called pheromone information and heuristic information associated to specific solution components. The pheromone information tries to bias solution construction towards the best solutions that have been found so far. Again, differently from iterated greedy algorithms, the ants start their solution construction from empty candidate solutions. However, there have been a number of proposals in the ACO area that suggest starting the solution construction from partial solutions that are either stored (Acan, 2004, 2005) or obtained by deconstructing complete solutions (Tsutsui, 2007, 2006; Wiesemann and Stützle, 2006) in a way akin to iterated greedy algorithms. These approaches generally try to transfer the advantages of iterated greedy to ACO algorithms and for several of these approaches positive results are reported.

1.6.2 (*Perturbative*) Local search techniques

Iterated greedy has tight links to local search algorithms and, in particular, to very large-scale neighborhood searches (Ahuja et al, 2002). In fact, as also mentioned in Section 1.5, one of the proposals of iterated greedy style algorithms called the method large neighborhood search (Pisinger and Ropke, 2010). The analogy stems from the fact that the removal of k solution components in the destruction step and the subsequent re-construction of a complete candidate solution can be seen as a move in a large neighborhood that is implicitly defined by the number of components to be removed and/or added.⁴ The (greedy) re-construction of a complete candidate solution typical for iterated greedy algorithms can then be seen as a heuristic examination of the neighborhood either in a fully greedy way or using randomization steps—details that simply depend on the particular design and implementation of the construction step in an iterated greedy algorithm. By varying the value of k , the size of the implicitly defined neighborhood is varied. Such changes may be done either (i) randomly within some limited range $[k_{\min}, k_{\max}]$ resembling known strategies in local search methods such as robust tabu search (Taillard, 1991), where one tries to avoid overcommitment to a single value for a parameter, (ii) in a more systematic way such as advocated in the strategies defined for variable neighborhood search (Hansen and Mladenović, 2001), or (iii) by using feedback from the search performance such as advocated in reactive search methods (Battiti et al, 2008). An interesting possibility for exploring the neighborhoods is used in the adaptive large neighborhood algorithms by Pisinger and Ropke (2007); Ropke and Pisinger (2006), who propose to consider different heuristics to be used in the solution re-construction. In their method, they additionally adapt the probability with which specific heuristics are chosen over the run-time of the algorithm based on the feedback of the search process.

An alternative to the heuristic exploration of neighborhoods provide methods that try to determine the best possible solution that can be reached from the current one, akin to best-improvement neighborhoods. In the case of large neighborhood searches, this corresponds to an exact examination of all the possible solutions that can be reached from a current partial candidate solution (Ahuja et al, 2002; Dumitrescu and Stützle, 2009). The original proposal of large neighborhood search by Shaw (1998) actually considered such an exact exploration of the resulting neighborhood by means of a constraint programming approach embedded within a branch-and-bound scheme. Due to the large variation in computation time that such a scheme however incurs, alternatives have been tested such as pure insertion-based reconstruction or limited enumeration schemes.

⁴ Note that the number of solution components removed in a destruction step may be different from the number of solution components added in the construction step and so we refrain from talking of k -exchange neighborhoods here. A common example where this happens are subset problems such as the SCP we discussed earlier.

If we consider other local search based SLS methods, the closest related one is certainly iterated local search (Ramalhinho Lourenço et al, 2010), especially when iterated greedy algorithms exploit additional (perturbative) local search methods to improve candidate solutions, as outlined in Figure 1.3. In that case, the destruction-construction cycle corresponds to what in ILS terms would be a solution perturbation. In general, it seems to be commendable in ILS to apply problem specific perturbations whenever possible and the destruction-construction cycle in iterated greedy introduces such problem specific information through the use of heuristic information. Basic variable neighborhood search is an ILS type algorithm that systematically varies the strength of the perturbation and, thus, the link to iterated greedy is also immediate.

Besides the fact that complete solutions may be improved by an additional local search phase in iterated greedy algorithms, it could be worthwhile to also consider the local re-optimization of the partial solutions that are obtained during the destruction and re-construction process. It is noteworthy that such occasional local re-optimization of partial candidate solutions has shown to be useful in a number of other contexts such as vehicle routing (Caseau and Laburthe, 1999).

1.6.3 Tree search algorithms

Constructive mechanisms can be extended to an exhaustive search method by adding a simple backtracking mechanism; adding further bounding schemes, one quickly comes to methods such as branch-and-bound or, more in general, tree search techniques (Hoos and Stützle, 2005). Hence, it is clear that iterated greedy algorithms also share some relationships to such methods.

Some links have already been mentioned in the previous section, when exact methods are exploited to generate the best possible completion of a partial candidate solution. Such a completion may be generated using tree-search methods and various of these ideas have been explored (Shaw, 1998). Even if tree search is used, the search need not be necessarily complete. For example, Shaw (1998) has considered the usage of limited discrepancy search (Harvey and Ginsberg, 1995), which consists in examining a limited number of alternative choices in the decision points during the tree generation. Any other variants such as depth-bounded discrepancy search (Walsh, 1997) could also be interesting. Another alternative that does not seem to have been applied so far would be to apply beam search (Ow and Morton, 1988).

Yet another connection to tree search techniques could rely on the exploitation of lower bound information to generate heuristic information or to prune choices that are guaranteed not to lead to improved solutions. So far, we are not aware of an exploitation of such ideas inside iterated greedy algorithms,

although in other constructive SLS algorithms such uses have been explored (Maniezzo, 1999; Blum, 2005).

1.7 Applications

In this section, we give a short overview of the applications of iterated greedy algorithms. Few applications have already been mentioned in Section 1.5 when discussing other methods that use the same principle as iterated greedy. Here, we focus mainly on applications that identified their algorithm as an iterated greedy algorithm. In fact, depending on the class of problems for which the respective methods have been proposed, the usage of names to refer to iterated greedy type methods differs. For example, in the scheduling area, the excellent results of the iterated greedy algorithms for the permutation flow shop scheduling problem has spawned a lot of follow-up work on similar problems, while the name large neighborhood search is frequently found in application to vehicle routing problems, given the prominent role these algorithms have played in that domain.

1.7.1 *Iterated greedy for scheduling problems*

As mentioned, iterated greedy, as defined in this chapter, was initially applied to the permutation flow shop scheduling problem by Ruiz and Stützle (2007), so many applications to other scheduling problems were published after that.⁵ Ruiz and Stützle (2008) extended their iterated greedy algorithm to tackle the permutation flow shop scheduling problem when considering additional sequence-dependent setup times and other objectives than makespan, namely the total tardiness. Other variants of flow shop problems have been studied by Ribas et al (2011); Tasgetiren et al (2017) (blocking), Pan et al (2008) and Ding et al (2015) (no-wait), Ruiz et al (2009), Tasgetiren et al (2013) and Pan and Ruiz (2014) who studied no-idle and mixed no-idle problems. Non-permutation flow shops were approached by Ying (2008b) or Benavides and Ritt (2015). Distributed flow shop scheduling problems have been solved with iterated greedy algorithms by Lin et al (2013) and Fernandez-Viagas and Framinan (2015). Iterated greedy has also proven valuable in other optimization criteria apart from makespan. Tardiness is studied by Framinan and Leisten (2008) or total flowtime in Pan and Ruiz (2012) to name just a few. Multi-objective extensions of iterated greedy have proven effective for Pareto flow shops without and with setups in Minella et al (2011) and in Ciavotta et al (2013), respectively; Dubois-Lacoste et al (2011) have embedded

⁵ Various applications of iterative flattening to scheduling problems have been referenced in Section 1.5.

an iterated greedy algorithm into the two-phase local search framework to tackle various bi-objective flow shop problems. Parallel machines problems were successfully tackled with iterated greedy algorithms by Fanjul-Peyro and Ruiz (2010), Rodriguez et al (2013), or Arroyo and Leung (2017). Iterated greedy algorithms have also shown to be effective for various single machine scheduling problems as shown by Deng and Gu (2014) and Ying et al (2009). Some other more complex problems have been studied, including job-shop scheduling with blocking constraints (Pranzo and Pacciarelli, 2016), job-shop scheduling with sequence-dependent set-up times and job families (Kim et al, 2017), hybrid flow shops (Ying, 2008a), and real life problems (Urlings and Ruiz, 2007; Urlings et al, 2010a,b).

1.7.2 Iterated greedy for routing problems

As mentioned above, in the context of vehicle routing problems several algorithms that follow the main structure of an iterated greedy algorithm have been applied but usually under the name large neighborhood search. The article making these approaches popular is the already mentioned one by Shaw (1998). A major impact in that line of applications had the adaptive large neighborhood search approaches by Pisinger and Ropke (2007); Ropke and Pisinger (2006), which led to a large number of follow-up work mainly in the vehicle routing domain. Some overview is also given in Pisinger and Ropke (2010). In fact, the usage of such iterated greedy type methods as local searches inside SLS algorithms for vehicle routing problems is increasingly wide-spread and can by now be considered a standard in this area. Apart from vehicle routing, also iterated greedy approaches have been proposed to few other routing type problems such as scheduling and routing problems of freight trains (Yuan et al, 2008) or in the context of TSP variants (Karabulut and Tasgetiren, 2014).

1.7.3 Iterated greedy for other problems

There have been a number of other applications where explicitly iterated greedy algorithms have been devised as the main solution techniques. Lozano et al (2011) presented an iterated greedy approach to the maximum diversity problem, where from a set of elements a subset with maximum diversity has to be chosen. After proper tuning, the algorithm was shown to perform better than various competing algorithms and was established as a new state-of-the-art algorithm. García-Martínez et al (2014) have developed an iterated greedy algorithm enhanced by a short tabu list in the destructive phase for the quadratic multiple knapsack problem. Lozano et al (2014) have developed also

another iterated greedy algorithm for the quadratic minimum spanning tree problem obtained excellent results on large instances; further embedding the algorithm into a strategic oscillation approach by essentially extending it with tabu criteria led to further improvements on some problem instance classes. Early implementations of similar ideas for binary quadratic programming have been presented by Glover et al (1998); more recently, Toyama et al (2012) apply iterated greedy to the same problem tackling also very large scale instances successfully. Kang et al (2013) study the problem of allocating parallel tasks to processors in computing systems that are distributed and heterogeneous. Population-based iterated greedy algorithms and iterated greedy algorithms that exploit further exact solutions to large neighborhood searches have been proposed by Rodríguez et al (2012) for the maximal covering location problem, the goal of which is to cover clients such that the largest amount of demand possible is satisfied. A problem in market segmentation, where a company asks to partition a set of customers subject to some specific requirements related to homogeneity of customers and compactness of the areas is tackled by Huerta-Muñoz et al (2017). A population-based iterated greedy algorithm has also been applied by Porta et al (2013) for delimiting and zoning rural settlements and by Bouamama et al (2012) to the minimum weight vertex cover problem. First applications of iterated greedy algorithms for machine learning tasks, here in particular the generation of classification rules has been explored by Pedraza et al (2014).

Given the wide applicability, the flexibility, and the often high performance of iterated greedy, we would expect this list of applications to further grow significantly in the future.

1.8 Conclusions

The main principle of iterated greedy is to build a sequence of solutions by iterating over constructive algorithms through a loop of solution destructions and (re-)constructions. Deconstruction removes solution components resulting in partial solutions from which again full solutions are re-constructed. This loop may be extended by an additional local search phase, where the generated complete candidate solutions are further improved, and by many other techniques from other heuristic and exact search techniques, making iterated greedy a very flexible and malleable method.

We prefer to call this principle iterated greedy because (i) this name directly refers to what the principle implies, namely making iterative use of construction methods, (ii) it does not obfuscate the name with natural or unnatural analogies, (iii) it is a short and punchy name describing the essence of the method, and (iv) and it has by now been used in a large number of publications. A bit unfortunate is that in the literature there have been proposed several methods under different names that make use of

the same (or a very similar) principle including large neighborhood search (Shaw, 1998; Pisinger and Ropke, 2010), simulated annealing (Jacobs and Brusco, 1995), evolutionary heuristic (Marchiori and Steenbeek, 2000), ruin-and-recreate (Schrimpf et al, 2000), iterative construction search (Richmond and Beasley, 2004), or iterative flattening (Cesta et al, 2000). The multitude of different names for the same kind of approach can probably be explained in two ways. First, different researchers have a different perspective on the same method and, thus, the iterated greedy principle can be viewed and proposed, for example, from the perspective of a (perturbative) neighborhood search leading to the heuristic or exact exploration of large neighborhoods or from the perspective of a (constructive) algorithm background where the method is simply iterating through applications of constructive algorithms. Second, the method has been proposed in somewhat different communities such as ruin-and-recreate in a physics journal (Schrimpf et al, 2000), large neighborhood search at a constraint programming conference (Shaw, 1998), iterative flattening at an artificial intelligence conference (Cesta et al, 2000), and publications using other names have appeared in operations research journals (Jacobs and Brusco, 1995; Richmond and Beasley, 2004; Ruiz and Stützle, 2007). It apparently has taken a significant amount of time that these ideas transpired from one community to another.

Even if this leads to some confusion and we maybe contribute to this confusion, we think that what is really important is (i) the principle underlying these methods, (ii) the fact that the iterated greedy principle can lead to very powerful algorithms and (iii) that iterated greedy is a very flexible method that can easily be combined with other techniques. By making the relationship among the different proposals clear, we hope to contribute also to a transfer of experience between these different algorithms. In any case, we hope that this review article will be useful for other researchers in stochastic local search methods by clearly identifying the potential of iterated greedy and in this way also contribute to its further development.

Acknowledgements This work received support from the COMEX project (P7/36) within the Interuniversity Attraction Poles Programme of the Belgian Science Policy Office. Thomas Stützle acknowledges support from the Belgian F.R.S.-FNRS, of which he is a research director. Rubén Ruiz is partially supported by the Spanish Ministry of Economy and Competitiveness, under the project “SCHEYARD – Optimization of Scheduling Problems in Container Yards” (No. DPI2015-65895-R) financed by FEDER funds.

References

- Acan A (2004) An external memory implementation in ant colony optimization. In: Dorigo M, et al (eds) *Ant Colony Optimization and Swarm Intelligence*, 4th International Workshop, ANTS 2004, Lecture Notes in Computer Science, vol 3172, Springer,

- Heidelberg, Germany, pp 73–84
- Acan A (2005) An external partial permutations memory for ant colony optimization. In: Raidl GR, Gottlieb J (eds) *Proceedings of EvoCOP 2005 – 5th European Conference on Evolutionary Computation in Combinatorial Optimization*, Lecture Notes in Computer Science, vol 3448, Springer, Heidelberg, Germany, pp 1–11
- Ahmadi S, Osman IH (2004) Density based problem space search for the capacitated clustering p -median problem. *Annals of Operations Research* 131:21–43
- Ahuja RK, Ergun O, Orlin JB, Punnen AP (2002) A survey of very large-scale neighborhood search techniques. *Discrete Applied Mathematics* 123(1–3):75–102
- Aickelin U, Burke EK, Li J (2006) Improved squeaky wheel optimisation for driver scheduling. In: Runarsson et al (2006)
- Arroyo J, Leung JT (2017) An effective iterated greedy algorithm for scheduling unrelated parallel batch machines with non-identical capacities and unequal ready times. *Computers and Industrial Engineering* 105:84–100
- Balas E, Ho A (1980) Set covering algorithms using cutting planes, heuristics, and subgradient optimization: A computational study. *Mathematical Programming Study* 12:37–60
- Bartz-Beielstein T, Chiarandini M, Paquete L, Preuss M (eds) (2010) *Experimental Methods for the Analysis of Optimization Algorithms*. Springer, Berlin, Germany
- Battiti R, Brunato M, Mascia F (2008) *Reactive Search and Intelligent Optimization*, Operations Research/Computer Science Interfaces, vol 45. Springer, New York, NY, DOI 10.1007/978-0-387-09624-7
- Benavides AJ, Ritt M (2015) Two simple and effective heuristics for minimizing the makespan in non-permutation flow shops. *Computers & Operations Research* 66:160–169
- Bertsekas DP, Tsitsiklis JN, Wu C (1997) Rollout algorithms for combinatorial optimization. *Journal of Heuristics* 3(3):245–262
- Blum C (2005) Beam-ACO—Hybridizing ant colony optimization with beam search: An application to open shop scheduling. *Computers & Operations Research* 32(6):1565–1591
- Bouamama S, Blum C, Boukerram A (2012) A population-based iterated greedy algorithm for the minimum weight vertex cover problem. *Applied Soft Computing* 12(6):1632–1639
- Brusco MJ, Jacobs LW, Thompson GM (1999) A morphing procedure to supplement a simulated annealing heuristic for cost- and coverage-correlated set covering problems. *Annals of Operations Research* 86:611–627
- Burke EK, Gendreau M, Hyde MR, Kendall G, Ochoa G, Özcan E, Qu R (2013) Hyperheuristics: A survey of the state of the art. *Journal of the Operational Research Society* 64(12):1695–1724
- Caseau Y, Laburthe F (1999) Heuristics for large constrained vehicle routing problems. *Journal of Heuristics* 5(3):281–303
- Cesta A, Oddi A, Smith SF (2000) Iterative flattening: A scalable method for solving multi-capacity scheduling problems. In: *Proceedings of AAAI 2000 – Seventeenth National Conference on Artificial Intelligence*, AAAI Press/MIT Press, Menlo Park, CA, pp 742–747
- Ciavotta M, Minella G, Ruiz R (2013) Multi-objective sequence dependent setup times flowshop scheduling: a new algorithm and a comprehensive study. *European Journal of Operational Research* 227(2):301–313
- Culberson JC (1992) Iterated greedy graph coloring and the difficulty landscape. Tech. Rep. 92-07, Department of Computing Science, The University of Alberta, Edmonton, Alberta, Canada
- Dees, Jr WA, Karger PG (1982) Automated rip-up and reroute techniques. In: *DAC’82, Proceedings of the 19th Design Automation Workshop*, IEEE Press, pp 432–439

- Deng G, Gu X (2014) An iterated greedy algorithm for the single-machine total weighted tardiness problem with sequence-dependent setup times. *International Journal of Systems Science* 45(3):351–362
- Ding JY, Song S, Gupta JND, Zhang R, Chiong R, Wu C (2015) An improved iterated greedy algorithm with a tabu-based reconstruction strategy for the no-wait flowshop scheduling problem. *Applied Soft Computing* 30:604–613
- Dorigo M, Stützle T (2004) *Ant Colony Optimization*. MIT Press, Cambridge, MA
- Dorigo M, Birattari M, Stützle T (2006) Ant colony optimization: Artificial ants as a computational intelligence technique. *IEEE Computational Intelligence Magazine* 1(4):28–39
- Dubois-Lacoste J, López-Ibáñez M, Stützle T (2011) A hybrid TP+PLS algorithm for bi-objective flow-shop scheduling problems. *Computers & Operations Research* 38(8):1219–1236, DOI 10.1016/j.cor.2010.10.008
- Dubois-Lacoste J, Pagnozzi F, Stützle T (2017) An iterated greedy algorithm with optimization of partial solutions for the makespan permutation flowshop problem. *Computers & Operations Research* 81:160–166
- Duin C, Voß S (1999) The pilot method: A strategy for heuristic repetition with application to the Steiner problem in graphs. *Networks* 34(3):181–191
- Dumitrescu I, Stützle T (2009) Usage of exact algorithms to enhance stochastic local search algorithms. In: Maniezzo V, Stützle T, Voß S (eds) *Matheuristics—Hybridizing Metaheuristics and Mathematical Programming*, *Annals of Information Systems*, vol 10, Springer, New York, NY, pp 103–134
- Fanjul-Peyro L, Ruiz R (2010) Iterated greedy local search methods for unrelated parallel machine scheduling. *European Journal of Operational Research* 207(1):55–69
- Feo TA, Resende MGC (1989) A probabilistic heuristic for a computationally difficult set covering problem. *Operations Research Letters* 8(2):67–71
- Feo TA, Resende MGC (1995) Greedy randomized adaptive search procedures. *Journal of Global Optimization* 6:109–113
- Fernandez-Viagas V, Framiñán JM (2014) On insertion tie-breaking rules in heuristics for the permutation flowshop scheduling problem. *Computers & Operations Research* 45:60–67
- Fernandez-Viagas V, Framinan JM (2015) A bounded-search iterated greedy algorithm for the distributed permutation flowshop scheduling problem. *International Journal of Production Research* 53(4):1111–1123
- Framinan JM, Leisten R (2008) Total tardiness minimization in permutation flow shops: a simple approach based on a variable greedy algorithm. *International Journal of Production Research* 46(22):6479–6498
- Framiñán JM, Gupta JN, Leisten R (2004) A review and classification of heuristics for permutation flow-shop scheduling with makespan objective. *Journal of the Operational Research Society* 55(12):1243–1255
- Framiñán JM, Leisten R, Ruiz R (2014) *Manufacturing Scheduling Systems: An Integrated View on Models, Methods, and Tools*. Springer, New York, NY
- García-Martínez C, Rodríguez FJ, Lozano M (2014) Tabu-enhanced iterated greedy algorithm: A case study in the quadratic multiple knapsack problem. *European Journal of Operational Research* 232(3):454–463
- Gendreau M, Potvin JY (eds) (2010) *Handbook of Metaheuristics*, *International Series in Operations Research & Management Science*, vol 146, 2nd edn. Springer, New York, NY
- Glover F (1977) Heuristics for integer programming using surrogate constraints. *Decision Sciences* 8:156–166
- Glover F (1986) Future paths for integer programming and links to artificial intelligence. *Computers & Operations Research* 13(5):533–549
- Glover F (1989) Tabu search – Part I. *INFORMS Journal on Computing* 1(3):190–206, DOI 10.1287/ijoc.1.3.190

- Glover F, Kochenberger G (eds) (2002) Handbook of Metaheuristics. Kluwer Academic Publishers, Norwell, MA
- Glover F, Kochenberger GA (1996) Critical even tabu search for multidimensional knapsack problems. In: Osman IH, Kelly JP (eds) Metaheuristics: Theory & Applications, Kluwer Academic Publishers, Norwell, MA, pp 407–427
- Glover F, Kochenberger GA, Alidaee B (1998) Adaptive memory tabu search for binary quadratic programs. *Management Science* 44(3):336–345
- Hansen P, Mladenović N (2001) Variable neighborhood search: Principles and applications. *European Journal of Operational Research* 130(3):449–467
- Hart JP, Shogan AW (1987) Semi-greedy heuristics: An empirical study. *Operations Research Letters* 6(3):107–114
- Harvey WD, Ginsberg ML (1995) Limited discrepancy search. In: Mellish CS (ed) Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence (IJCAI-95), Morgan Kaufmann Publishers, pp 607–615
- Hejazi SR, Saghafian S (2005) Flowshop-scheduling problems with makespan criterion: A review. *International Journal of Production Research* 43(14):2895–2929
- Hoos HH, Stützle T (2005) Stochastic Local Search—Foundations and Applications. Morgan Kaufmann Publishers, San Francisco, CA
- Huerta-Muñoz DL, Ríos-Mercado RZ, Ruiz R (2017) An iterated greedy heuristic for a market segmentation problem with multiple attributes. *European Journal of Operational Research* 261(1):75–87
- Jacobs LW, Brusco MJ (1995) A local search heuristic for large set-covering problems. *Naval Research Logistics* 42(7):1129–1140
- Johnson DS, McGeoch LA (2002) Experimental analysis of heuristics for the STSP. In: Gutin G, Punnen A (eds) The Traveling Salesman Problem and its Variations, Kluwer Academic Publishers, Dordrecht, The Netherlands, pp 369–443
- Joslin DE, Clements DP (1999) Squeaky wheel optimization. *Journal of Artificial Intelligence Research* 10:353–373
- Kang Q, He H, Wei J (2013) An effective iterated greedy algorithm for reliability-oriented task allocation in distributed computing systems. *Journal of Parallel and Distributed Computing* 73(8):1106–1115
- Karabulut K, Tasgetiren FM (2014) A variable iterated greedy algorithm for the traveling salesman problem with time windows. *Information Sciences* 279:383–395
- Kim JS, Park JH, Lee DH (2017) Iterated greedy algorithms to minimize the total family flow time for job-shop scheduling with job families and sequence-dependent set-ups. *Engineering Optimization* 49(10):1719–1732
- Kirkpatrick S, Gelatt CD, Vecchi MP (1983) Optimization by simulated annealing. *Science* 220:671–680
- Lin SW, Ying KC, Huang CY (2013) Minimising makespan in distributed permutation flowshops using a modified iterated greedy algorithm. *International Journal of Production Research* 51(16):5029–5038
- Lozano M, Molina D, García-Martínez C (2011) Iterated greedy for the maximum diversity problem. *European Journal of Operational Research* 214(1):31–38
- Lozano M, Glover F, García-Martínez C, Rodríguez FJ, Martí R (2014) Tabu search with strategic oscillation for the quadratic minimum spanning tree. *IIE Transactions* 46(4):414–428
- Maniezzo V (1999) Exact and approximate nondeterministic tree-search procedures for the quadratic assignment problem. *INFORMS Journal on Computing* 11(4):358–369
- Marchiori E, Steenbeek AG (1998) An iterated heuristic algorithm for the set covering problem. In: Mehlhorn K (ed) Algorithm Engineering, 2nd International Workshop, WAE'92, Max-Planck-Institut für Informatik, Saarbrücken, Germany, pp 155–166
- Marchiori E, Steenbeek AG (2000) An evolutionary algorithm for large scale set covering problems with application to airline crew scheduling. In: Cagnoni S, et al (eds)

- Real-World Applications of Evolutionary Computing, EvoWorkshops 2000, Springer, Heidelberg, Germany, Lecture Notes in Computer Science, vol 1803, pp 367–381
- Michel LD, van Hentenryck P (2004) Iterative relaxations for iterative flattening in cumulative scheduling. In: Zilberstein S, Koehler J, Koenig S (eds) Proceedings of the Fourteenth International Conference on Automated Planning and Scheduling (ICAPS 2004), AAAI Press/MIT Press, Menlo Park, CA, pp 200–208
- Minella G, Ruiz R, Ciavotta M (2011) Restarted iterated pareto greedy algorithm for multi-objective flowshop scheduling problems. *Computers & Operations Research* 38(11):1521–1533
- Misevičius A (2003a) Genetic algorithm hybridized with ruin and recreate procedure: Application to the quadratic assignment problem. *Knowledge Based Systems* 16(5–6):261–268
- Misevičius A (2003b) Ruin and recreate principle based approach for the quadratic assignment problem. In: Cantú-Paz E, et al (eds) Genetic and Evolutionary Computation - GECCO 2003, Part I, Springer, Heidelberg, Germany, Lecture Notes in Computer Science, vol 2723, pp 598–609
- Montgomery DC (2012) Design and Analysis of Experiments, eighth edn. John Wiley & Sons, New York, NY
- Nawaz M, Enscore E Jr, Ham I (1983) A heuristic algorithm for the m -machine, n -job flow-shop sequencing problem. *Omega* 11(1):91–95
- Oddi A, Cesta A, Policella N, Smith SF (2008) Combining variants of iterative flattening search. *Engineering Applications of Artificial Intelligence* 21(5):683–690
- Oddi A, Cesta A, Policella N, Smith SF (2010) Iterative flattening search for resource constrained scheduling. *Journal of Intelligent Manufacturing* 21(1):17–30
- Oddi A, Rasconi R, Cesta A, Smith SF (2011) Iterative flattening search for the flexible job shop scheduling problem. In: Walsh T (ed) Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence (IJCAI-11), IJCAI/AAAI Press, Menlo Park, CA, pp 1991–1996
- Ow PS, Morton TE (1988) Filtered beam search in scheduling. *International Journal of Production Research* 26:297–307
- Pan QK, Ruiz R (2012) Local search methods for the flowshop scheduling problem with flowtime minimization. *European Journal of Operational Research* 222(1):31–43
- Pan QK, Ruiz R (2014) An effective iterated greedy algorithm for the mixed no-idle flowshop scheduling problem. *Omega* 44(1):41–50
- Pan QK, Wang L, Zhao BH (2008) An improved iterated greedy algorithm for the no-wait flow shop scheduling problem with makespan criterion. *International Journal of Advanced Manufacturing Technology* 38(7-8):778–786
- Pedraza JA, García-Martínez C, Cano A, Ventura S (2014) Classification rule mining with iterated greedy. In: Polycarpou MM, de Carvalho ACPLF, Pan J, Wozniak M, Quintián H, Corchado E (eds) Hybrid Artificial Intelligence Systems - 9th International Conference, HAIS 2014, Salamanca, Spain, June 11-13, 2014. Proceedings, Lecture Notes in Computer Science, vol 8480, Springer, Heidelberg, Germany, pp 585–596
- Pisinger D, Ropke S (2007) A general heuristic for vehicle routing problems. *Computers & Operations Research* 34(8):2403–2435
- Pisinger D, Ropke S (2010) Large neighborhood search. In: Gendreau and Potvin (2010), pp 399–419
- Porta J, Parapar J, Doallo R, Barbosa V, Santé I, Crecente R, Díaz C (2013) A population-based iterated greedy algorithm for the delimitation and zoning of rural settlements. *Computers, Environment and Urban Systems* 39:12–26
- Pranzo M, Pacciarelli D (2016) An iterated greedy metaheuristic for the blocking job shop scheduling problem. *Journal of Heuristics* 22(4):587–611, 10.1007/s10732-014-9279-5
- Rad SF, Ruiz R, Boroojerdian N (2009) New high performing heuristics for minimizing makespan in permutation flowshops. *Omega* 37(2):331–345

- Ramalhinho Lourenço H, Martin O, Stützle T (2002) Iterated local search. In: Glover and Kochenberger (2002), pp 321–353
- Ramalhinho Lourenço H, Martin O, Stützle T (2010) Iterated local search: Framework and applications. In: Gendreau and Potvin (2010), chap 9, pp 363–397
- Resende MGC, Ribeiro CC (2010) Greedy randomized adaptive search procedures: Advances, hybridizations, and applications. In: Gendreau and Potvin (2010), pp 283–319
- Ribas I, Companys R, Tort-Martorell X (2011) An iterated greedy algorithm for the flowshop scheduling problem with blocking. *Omega* 39(3):293 – 301
- Richmond AJ, Beasley JE (2004) An iterative construction heuristic for the ore selection problem. *Journal of Heuristics* 10(2):153–167
- Rodríguez FJ, Blum C, Lozano M, García-Martínez C (2012) Iterated greedy algorithms for the maximal covering location problem. In: Hao JK, Middendorf M (eds) Proceedings of EvoCOP 2012 – 12th European Conference on Evolutionary Computation in Combinatorial Optimization, Lecture Notes in Computer Science, vol 7245, Springer, Heidelberg, Germany, pp 172–181
- Rodríguez FJ, Lozano M, Blum C, García-Martínez C (2013) An iterated greedy algorithm for the large-scale unrelated parallel machines scheduling problem. *Computers & Operations Research* 40(7):1829–1841
- Ropke S, Pisinger D (2006) An adaptive large neighborhood search heuristic for the pickup and delivery problem with time windows. *Transportation Science* 40(4):455–472
- Rubin F (1974) An iterative technique for printed wire routing. In: DAC’74, Proceedings of the 11th Design Automation Workshop, IEEE Press, pp 308–313
- Ruiz R, Maroto C (2005) A comprehensive review and evaluation of permutation flowshop heuristics. *European Journal of Operational Research* 165(2):479–494
- Ruiz R, Stützle T (2007) A simple and effective iterated greedy algorithm for the permutation flowshop scheduling problem. *European Journal of Operational Research* 177(3):2033–2049
- Ruiz R, Stützle T (2008) An Iterated Greedy heuristic for the sequence dependent setup times flowshop problem with makespan and weighted tardiness objectives. *European Journal of Operational Research* 187(3):1143 – 1159
- Ruiz R, Vallada E, Fernández-Martínez C (2009) Scheduling in flowshops with no-idle machines. In: Chakraborty UK (ed) Computational Intelligence in Flow Shop and Job Shop Scheduling, Springer-Verlag, Berlin, Studies in Computational Intelligence, vol 230, pp 21–51
- Runarsson TP, Beyer HG, Burke EK, Merelo JJ, Whitley LD, Yao X (eds) (2006) Proceedings of PPSN-IX, Ninth International Conference on Parallel Problem Solving from Nature, Lecture Notes in Computer Science, vol 4193. Springer, Heidelberg, Germany
- Schrimpf G, Schneider J, Stamm-Wilbrandt H, Dueck G (2000) Record breaking optimization results using the ruin and recreate principle. *Journal of Computational Physics* 159(2):139 – 171
- Shaw P (1998) Using constraint programming and local search methods to solve vehicle routing problems. In: Maher M, Puget JF (eds) Principles and Practice of Constraint Programming, CP98, Lecture Notes in Computer Science, vol 1520, Springer, Heidelberg, Germany, pp 417–431
- Taillard ÉD (1990) Some efficient heuristic methods for the flow shop sequencing problem. *European Journal of Operational Research* 47(1):65–74
- Taillard ÉD (1991) Robust taboo search for the quadratic assignment problem. *Parallel Computing* 17(4-5):443–455
- Taillard ÉD (1993) Benchmarks for basic scheduling problems. *European Journal of Operational Research* 64(2):278–285

- Tasgetiren FM, Pan QK, Suganthan PN, Buyukdagli O (2013) A variable iterated greedy algorithm with differential evolution for the no-idle permutation flowshop scheduling problem. *Computers & Operations Research* 40(7):1729–1743
- Tasgetiren MF, Kizilay D, Pan QK, Suganthan PN (2017) Iterated greedy algorithms for the blocking flowshop scheduling problem with makespan criterion. *Computers & Operations Research* 77:111–126
- Toyama F, Shoji K, Mori H, Miyamichi J (2012) An iterated greedy algorithm for the binary quadratic programming problem. In: Joint 6th International Conference on Soft Computing and Intelligent Systems (SCIS) and 13th International Symposium on Advanced Intelligent Systems (ISIS), 2012, IEEE Press, pp 2183–2188
- Tsutsui S (2006) cAS: Ant colony optimization with cunning ants. In: Runarsson et al (2006), pp 162–171
- Tsutsui S (2007) Ant colony optimization with cunning ants. *Transactions of the Japanese Society for Artificial Intelligence* 22:29–36, DOI 10.1527/tjsai.22.29
- Urlings T, Ruiz R (2007) Local search in complex scheduling problems. In: Stützle T, Birattari M, Hoos HH (eds) *Engineering Stochastic Local Search Algorithms. Designing, Implementing and Analyzing Effective Heuristics*, Lecture Notes in Computer Science, vol 4638, Springer-Verlag, Brussels, Belgium, pp 202–206
- Urlings T, Ruiz R, Sivrikaya-Şerifoğlu F (2010a) Genetic algorithms for complex hybrid flexible flow line problems. *International Journal of Metaheuristics* 1(1):30–54
- Urlings T, Ruiz R, Stützle T (2010b) Shifting representation search for hybrid flexible flowline problems. *European Journal of Operational Research* 207(2):1086–1095, DOI 10.1016/j.ejor.2010.05.041
- Walsh T (1997) Depth-bounded discrepancy search. In: Pollack ME (ed) *Proceedings of the Fifteenth International Joint Conference on Artificial Intelligence (IJCAI-97)*, Morgan Kaufmann Publishers, pp 1388–1395
- Wiesemann W, Stützle T (2006) Iterated ants: An experimental study for the quadratic assignment problem. In: Dorigo M, et al (eds) *Ant Colony Optimization and Swarm Intelligence, 5th International Workshop, ANTS 2006*, Lecture Notes in Computer Science, vol 4150, Springer, Heidelberg, Germany, pp 179–190
- Ying KC (2008a) An iterated greedy heuristic for multistage hybrid flowshop scheduling problems with multiprocessor tasks. *IEEE Transactions on Evolutionary Computation* 60(6):810–817
- Ying KC (2008b) Solving non-permutation flowshop scheduling problems by an effective iterated greedy heuristic. *International Journal of Advanced Manufacturing Technology* 38(3-4):348–354
- Ying KC, Lin SW, Huang CY (2009) Sequencing single-machine tardiness problems with sequence dependent setup times using an iterated greedy heuristic. *Expert Systems with Applications* 36(3):7087–7092
- Yuan Z, Fügenschuh A, Homfeld H, Balaprakash P, Stützle T, Schoch M (2008) Iterated greedy algorithms for a real-world cyclic train scheduling problem. In: Blesa MJ, Blum C, Cotta C, Fernández AJ, Gallardo JE, Roli A, Sampels M (eds) *Hybrid Metaheuristics*, Lecture Notes in Computer Science, vol 5296, Springer, Heidelberg, Germany, pp 102–116
- Zilberstein S (1996) Using anytime algorithms in intelligent systems. *AI Magazine* 17(3):73–83