

Metaheuristics in Combinatorial Optimization: Overview and Conceptual Comparison

CHRISTIAN BLUM

Université Libre de Bruxelles

AND

ANDREA ROLI

Università degli Studi di Bologna

The field of metaheuristics for the application to combinatorial optimization problems is a rapidly growing field of research. This is due to the importance of combinatorial optimization problems for the scientific as well as the industrial world. We give a survey of the nowadays most important metaheuristics from a conceptual point of view. We outline the different components and concepts that are used in the different metaheuristics in order to analyze their similarities and differences. Two very important concepts in metaheuristics are intensification and diversification. These are the two forces that largely determine the behavior of a metaheuristic. They are in some way contrary but also complementary to each other. We introduce a framework, that we call the *I&D* frame, in order to put different intensification and diversification components into relation with each other. Outlining the advantages and disadvantages of different metaheuristic approaches we conclude by pointing out the importance of hybridization of metaheuristics as well as the integration of metaheuristics and other methods for optimization.

Categories and Subject Descriptors: G.2.1 [**Discrete Mathematics**]: Combinatorics—*combinatorial algorithms*; I.2.8 [**Artificial Intelligence**]: Problem Solving, Control Methods, and Search—*heuristic methods*

General Terms: Algorithms

Additional Key Words and Phrases: Metaheuristics, combinatorial optimization, intensification, diversification.

C. Blum acknowledges support by the “Metaheuristics Network,” a Research Training Network funded by the Improving Human Potential program of the CEC, contract HPRN-CT-1999-00106.

A. Roli acknowledges support by the CEC through a “Marie Curie Training Site” fellowship, contract HPMT-CT-2000-00032.

The information provided is the sole responsibility of the authors and does not reflect the Community’s opinion. The Community is not responsible for any use that might be made of data appearing in this publication. Authors’ addresses: C. Blum, Université Libre de Bruxelles, IRIDIA, Avenue Franklin Roosevelt 50, CP 194/6, 1050 Brussels, Belgium; email: cblum@ulb.ac.be; A. Roli, DEIA—Università degli Studi di Bologna, Viale Risorgimento, 2-Bologna, Italy; email: aroli@deis.unibo.it.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or direct commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 1515 Broadway, New York, NY 10036 USA, fax: +1 (212) 869-0481, or permissions@acm.org.

©2003 ACM 0360-0300/03/0900-0268 \$5.00

1. INTRODUCTION

Many optimization problems of practical as well as theoretical importance consist of the search for a “best” configuration of a set of variables to achieve some goals. They seem to divide naturally into two categories: those where solutions are encoded with *real-valued* variables, and those where solutions are encoded with *discrete* variables. Among the latter ones we find a class of problems called Combinatorial Optimization (CO) problems. According to Papadimitriou and Steiglitz [1982], in CO problems, we are looking for an object from a finite—or possibly countably infinite—set. This object is typically an integer number, a subset, a permutation, or a graph structure.

Definition 1.1. A *Combinatorial Optimization* problem $P = (S, f)$ can be defined by:

- a set of variables $X = \{x_1, \dots, x_n\}$;
- variable domains D_1, \dots, D_n ;
- constraints among variables;
- an *objective function* f to be minimized,¹ where $f : D_1 \times \dots \times D_n \rightarrow \mathbb{R}^+$;

The set of all possible feasible assignments is

$$S = \{s = \{(x_1, v_1), \dots, (x_n, v_n)\} \mid v_i \in D_i, s \text{ satisfies all the constraints}\}.$$

S is usually called a *search (or solution) space*, as each element of the set can be seen as a candidate solution. To solve a combinatorial optimization problem one has to find a solution $s^* \in S$ with minimum objective function value, that is, $f(s^*) \leq f(s) \forall s \in S$. s^* is called a globally optimal solution of (S, f) and the set $S^* \subseteq S$ is called the set of globally optimal solutions.

Examples for CO problems are the Travelling Salesman problem (TSP), the Quadratic Assignment problem (QAP), Timetabling and Scheduling problems. Due to the practical importance of CO

problems, many algorithms to tackle them have been developed. These algorithms can be classified as either *complete* or *approximate* algorithms. Complete algorithms are guaranteed to find for every finite size instance of a CO problem an optimal solution in bounded time (see Papadimitriou and Steiglitz [1982] and Nemhauser and Wolsey [1988]). Yet, for CO problems that are \mathcal{NP} -hard [Garey and Johnson 1979], no polynomial time algorithm exists, assuming that $\mathcal{P} \neq \mathcal{NP}$. Therefore, complete methods might need exponential computation time in the worst-case. This often leads to computation times too high for practical purposes. Thus, the use of approximate methods to solve CO problems has received more and more attention in the last 30 years. In approximate methods we sacrifice the guarantee of finding optimal solutions for the sake of getting good solutions in a significantly reduced amount of time.

Among the basic approximate methods we usually distinguish between *constructive* methods and *local search* methods. Constructive algorithms generate solutions from scratch by adding—to an initially empty partial solution—components, until a solution is complete. They are typically the fastest approximate methods, yet they often return solutions of inferior quality when compared to local search algorithms. Local search algorithms start from some initial solution and iteratively try to replace the current solution by a better solution in an appropriately defined neighborhood of the current solution, where the neighborhood is formally defined as follows:

Definition 1.2. A *neighborhood structure* is a function $\mathcal{N} : S \rightarrow 2^S$ that assigns to every $s \in S$ a set of neighbors $\mathcal{N}(s) \subseteq S$. $\mathcal{N}(s)$ is called the neighborhood of s .

The introduction of a neighborhood structure enables us to define the concept of *locally* minimal solutions.

Definition 1.3. A *locally minimal solution (or local minimum)* with respect to a neighborhood structure \mathcal{N} is a solution \hat{s} such that $\forall s \in \mathcal{N}(\hat{s}) : f(\hat{s}) \leq f(s)$. We

¹As maximizing an objective function f is the same as minimizing $-f$, in this work we will deal, without loss of generality, with minimization problems.

call \hat{s} a strict locally minimal solution if $f(\hat{s}) < f(s) \forall s \in \mathcal{N}(\hat{s})$.

In the last 20 years, a new kind of approximate algorithm has emerged which basically tries to combine basic heuristic methods in higher level frameworks aimed at efficiently and effectively exploring a search space. These methods are nowadays commonly called *metaheuristics*.² The term *metaheuristic*, first introduced in Glover [1986], derives from the composition of two Greek words. *Heuristic* derives from the verb *heuriskein* ($\epsilon\upsilon\pi\iota\sigma\kappa\epsilon\iota\nu$) which means “to find”, while the suffix *meta* means “beyond, in an upper level”. Before this term was widely adopted, metaheuristics were often called *modern heuristics* [Reeves 1993].

This class of algorithms includes³—but is not restricted to—Ant Colony Optimization (ACO), Evolutionary Computation (EC) including Genetic Algorithms (GA), Iterated Local Search (ILS), Simulated Annealing (SA), and Tabu Search (TS). Up to now there is no commonly accepted definition for the term *metaheuristic*. It is just in the last few years that some researchers in the field tried to propose a definition. In the following we quote some of them:

“A metaheuristic is formally defined as an iterative generation process which guides a subordinate heuristic by combining intelligently different concepts for exploring and exploiting the search space, learning strategies are used to structure information in order to find efficiently near-optimal solutions.” [Osman and Laporte 1996].

“A metaheuristic is an iterative master process that guides and modifies the operations of subordinate heuristics to efficiently produce high-quality solutions. It may manipulate a complete (or incomplete) single solution or a collection of solutions at each iteration. The subordinate

heuristics may be high (or low) level procedures, or a simple local search, or just a construction method.” [Voß et al. 1999].

“Metaheuristics are typically high-level strategies which guide an underlying, more problem specific heuristic, to increase their performance. The main goal is to avoid the disadvantages of iterative improvement and, in particular, multiple descent by allowing the local search to escape from local optima. This is achieved by either allowing worsening moves or generating new starting solutions for the local search in a more “intelligent” way than just providing random initial solutions. Many of the methods can be interpreted as introducing a bias such that high quality solutions are produced quickly. This bias can be of various forms and can be cast as descent bias (based on the objective function), memory bias (based on previously made decisions) or experience bias (based on prior performance). Many of the metaheuristic approaches rely on probabilistic decisions made during the search. But, the main difference to pure random search is that in metaheuristic algorithms randomness is not used blindly but in an intelligent, biased form.” [Stützle 1999b].

“A metaheuristic is a set of concepts that can be used to define heuristic methods that can be applied to a wide set of different problems. In other words, a metaheuristic can be seen as a general algorithmic framework which can be applied to different optimization problems with relatively few modifications to make them adapted to a specific problem.” [Metaheuristics Network Website 2000].

Summarizing, we outline fundamental properties which characterize metaheuristics:

- Metaheuristics are strategies that “guide” the search process.
- The goal is to efficiently explore the search space in order to find (near-) optimal solutions.
- Techniques which constitute metaheuristic algorithms range from simple local search procedures to complex learning processes.

²The increasing importance of metaheuristics is underlined by the biannual Metaheuristics International Conference (MIC). The 5th is being held in Kyoto in August 2003 (<http://www-or.amp.i.kyoto-u.ac.jp/mic2003/>).

³In alphabetical order.

- Metaheuristic algorithms are approximate and usually non-deterministic.
- They may incorporate mechanisms to avoid getting trapped in confined areas of the search space.
- The basic concepts of metaheuristics permit an abstract level description.
- Metaheuristics are not problem-specific.
- Metaheuristics may make use of domain-specific knowledge in the form of heuristics that are controlled by the upper level strategy.
- Today's more advanced metaheuristics use search experience (embodied in some form of memory) to guide the search.

In short, we could say that metaheuristics are high level strategies for exploring search spaces by using different methods. Of great importance hereby is that a dynamic balance is given between *diversification* and *intensification*. The term diversification generally refers to the exploration of the search space, whereas the term intensification refers to the exploitation of the accumulated search experience. These terms stem from the Tabu Search field [Glover and Laguna 1997] and it is important to clarify that the terms *exploration* and *exploitation* are sometimes used instead, for example in the Evolutionary Computation field [Eiben and Schippers 1998], with a more restricted meaning. In fact, the notions of exploitation and exploration often refer to rather short-term strategies tied to randomness, whereas intensification and diversification also refer to medium- and long-term strategies based on the usage of memory. The use of the terms diversification and intensification in their initial meaning becomes more and more accepted by the whole field of metaheuristics. Therefore, we use them throughout the article. The balance between diversification and intensification as mentioned above is important, on one side to quickly identify regions in the search space with high quality solutions and on the other side not to waste too much time in regions of the search space which are either already ex-

plored or which do not provide high quality solutions.

The search strategies of different metaheuristics are highly dependent on the philosophy of the metaheuristic itself. Comparing the strategies used in different metaheuristics is one of the goals of Section 5. There are several different philosophies apparent in the existing metaheuristics. Some of them can be seen as “intelligent” extensions of local search algorithms. The goal of this kind of metaheuristic is to escape from local minima in order to proceed in the exploration of the search space and to move on to find other hopefully better local minima. This is for example the case in Tabu Search, Iterated Local Search, Variable Neighborhood Search, GRASP and Simulated Annealing. These metaheuristics (also called trajectory methods) work on one or several neighborhood structure(s) imposed on the members (the solutions) of the search space.

We can find a different philosophy in algorithms like Ant Colony Optimization and Evolutionary Computation. They incorporate a learning component in the sense that they implicitly or explicitly try to learn correlations between decision variables to identify high quality areas in the search space. This kind of metaheuristic performs, in a sense, a biased sampling of the search space. For instance, in Evolutionary Computation this is achieved by recombination of solutions and in Ant Colony Optimization by sampling the search space in every iteration according to a probability distribution.

The structure of this work is as follows: There are several approaches to classify metaheuristics according to their properties. In Section 2, we briefly list and summarize different classification approaches. Section 3 and Section 4 are devoted to a description of the most important metaheuristics nowadays. Section 3 describes the most relevant trajectory methods and, in Section 4, we outline population-based methods. Section 5 aims at giving a unifying view on metaheuristics with respect to the way they achieve intensification and diversification. This is done by the

introduction of a unifying framework, the *I&D frame*. Finally, Section 6 offers some conclusions and an outlook to the future.

We believe that it is hardly possible to produce a completely accurate survey of metaheuristics that is doing justice to every viewpoint. Moreover, a survey of an immense area such as metaheuristics has to focus on certain aspects and therefore has unfortunately to neglect other aspects. Therefore, we want to clarify at this point that this survey is done from the conceptual point of view. We want to outline the different concepts that are used in different metaheuristics in order to analyze the similarities and the differences between them. We do not go into the implementation of metaheuristics, which is certainly an important aspect of metaheuristics research with respect to the increasing importance of efficiency and software reusability. We refer the interested reader to Whitley [1989], Grefenstette [1990], Fink and Voß [1999], Schaerf et al. [2000], and Voß and Woodruff [2002].

2. CLASSIFICATION OF METAHEURISTICS

There are different ways to classify and describe metaheuristic algorithms. Depending on the characteristics selected to differentiate among them, several classifications are possible, each of them being the result of a specific viewpoint. We briefly summarize the most important ways of classifying metaheuristics.

Nature-inspired vs. non-nature inspired. Perhaps, the most intuitive way of classifying metaheuristics is based on the origins of the algorithm. There are nature-inspired algorithms, like Genetic Algorithms and Ant Algorithms, and non nature-inspired ones such as Tabu Search and Iterated Local Search. In our opinion this classification is not very meaningful for the following two reasons. First, many recent hybrid algorithms do not fit either class (or, in a sense, they fit both at the same time). Second, it is sometimes difficult to clearly attribute an algorithm to one of the two classes. So, for example, one might ask the question if

the use of memory in Tabu Search is not nature-inspired as well.

Population-based vs. single point search. Another characteristic that can be used for the classification of metaheuristics is the number of solutions used at the same time: Does the algorithm work on a population or on a single solution at any time? Algorithms working on single solutions are called *trajectory methods* and encompass local search-based metaheuristics, like Tabu Search, Iterated Local Search and Variable Neighborhood Search. They all share the property of describing a trajectory in the search space during the search process. Population-based metaheuristics, on the contrary, perform search processes which describe the evolution of a set of points in the search space.

Dynamic vs. static objective function. Metaheuristics can also be classified according to the way they make use of the objective function. While some algorithms keep the objective function given in the problem representation “as it is”, some others, like Guided Local Search (GLS), modify it during the search. The idea behind this approach is to escape from local minima by modifying the search landscape. Accordingly, during the search the objective function is altered by trying to incorporate information collected during the search process.

One vs. various neighborhood structures. Most metaheuristic algorithms work on one single neighborhood structure. In other words, the fitness landscape topology does not change in the course of the algorithm. Other metaheuristics, such as Variable Neighborhood Search (VNS), use a set of neighborhood structures which gives the possibility to diversify the search by swapping between different fitness landscapes.

Memory usage vs. memory-less methods. A very important feature to classify metaheuristics is the use they make of the search history, that is, whether they use memory or not.⁴ Memory-less algorithms

⁴Here we refer to the use of adaptive memory, in contrast to rather rigid memory, as used for instance in Branch & Bound.

perform a Markov process, as the information they exclusively use to determine the next action is the current state of the search process. There are several different ways of making use of memory. Usually we differentiate between the use of short term and long term memory. The first usually keeps track of recently performed moves, visited solutions or, in general, decisions taken. The second is usually an accumulation of synthetic parameters about the search. The use of memory is nowadays recognized as one of the fundamental elements of a powerful metaheuristic.

In the following, we describe the most important metaheuristics according to the single point vs. population-based search classification, which divides metaheuristics into trajectory methods and population-based methods. This choice is motivated by the fact that this categorization permits a clearer description of the algorithms. Moreover, a current trend is the hybridization of methods in the direction of the integration of single point search algorithms in population-based ones. In the following two sections, we give a detailed description of nowadays most important metaheuristics.

3. TRAJECTORY METHODS

In this section we outline metaheuristics called *trajectory methods*. The term trajectory methods is used because the search process performed by these methods is characterized by a trajectory in the search space. Hereby, a successor solution may or may not belong to the neighborhood of the current solution.

The search process of trajectory methods can be seen as the evolution in (discrete) time of a discrete dynamical system [Bar-Yam 1997; Devaney 1989]. The algorithm starts from an initial state (the initial solution) and describes a trajectory in the state space. The system dynamics depends on the strategy used; simple algorithms generate a trajectory composed of two parts: a *transient* phase followed by an *attractor* (a fixed point, a cycle or a complex attractor). Algorithms with advanced strategies generate more complex

```

s ← GenerateInitialSolution()
repeat
  s ← Improve(N(s))
until no improvement is possible

```

Fig. 1. Algorithm: Iterative Improvement.

trajectories which can not be subdivided in those two phases. The characteristics of the trajectory provide information about the behavior of the algorithm and its effectiveness with respect to the instance that is tackled. It is worth underlining that the dynamics is the result of the combination of algorithm, problem representation and problem instance. In fact, the problem representation together with the neighborhood structures define the search landscape; the algorithm describes the strategy used to explore the landscape and, finally, the actual search space characteristics are defined by the problem instance to be solved.

We will first describe basic local search algorithms, before we proceed with the survey of more complex strategies. Finally, we deal with algorithms that are general explorative strategies which may incorporate other trajectory methods as components.

3.1. Basic Local Search: Iterative Improvement

The basic local search is usually called *iterative improvement*, since each move⁵ is only performed if the resulting solution is better than the current solution. The algorithm stops as soon as it finds a local minimum. The high level algorithm is sketched in Figure 1.

The function $\text{Improve}(\mathcal{N}(s))$ can be in the extremes either a *first improvement*, or a *best improvement* function, or any intermediate option. The former scans the neighborhood $\mathcal{N}(s)$ and chooses the first solution that is better than s , the latter exhaustively explores the neighborhood and returns one of the solutions with the lowest objective function value. Both methods

⁵A move is the choice of a solution s' from the neighborhood $\mathcal{N}(s)$ of a solution s .

```

s ← GenerateInitialSolution()
T ← T0
while termination conditions not met do
  s' ← PickAtRandom(N(s))
  if (f(s') < f(s)) then
    s ← s'           % s' replaces s
  else
    Accept s' as new solution with probability p(T, s', s)
  endif
  Update(T)
endwhile

```

Fig. 2. Algorithm: Simulated Annealing (SA).

stop at local minima. Therefore, their performance strongly depends on the definition of S , f and \mathcal{N} . The performance of iterative improvement procedures on CO problems is usually quite unsatisfactory. Therefore, several techniques have been developed to prevent algorithms from getting trapped in local minima, which is done by adding mechanisms that allow them to escape from local minima. This also implies that the termination conditions of metaheuristic algorithms are more complex than simply reaching a local minimum. Indeed, possible termination conditions include: maximum CPU time, a maximum number of iterations, a solution s with $f(s)$ less than a predefined threshold value is found, or the maximum number of iterations without improvements is reached.

3.2. Simulated Annealing

Simulated Annealing (SA) is commonly said to be the oldest among the metaheuristics and surely one of the first algorithms that had an explicit strategy to escape from local minima. The origins of the algorithm are in statistical mechanics (Metropolis algorithm) and it was first presented as a search algorithm for CO problems in Kirkpatrick et al. [1983] and Cerny [1985]. The fundamental idea is to allow moves resulting in solutions of worse quality than the current solution (uphill moves) in order to escape from local minima. The probability of doing such a move is decreased during the search. The high level algorithm is described in Figure 2.

The algorithm starts by generating an initial solution (either randomly or heuris-

tically constructed) and by initializing the so-called temperature parameter T . Then, at each iteration a solution $s' \in \mathcal{N}(s)$ is randomly sampled and it is accepted as new current solution depending on $f(s)$, $f(s')$ and T . s' replaces s if $f(s') < f(s)$ or, in case $f(s') \geq f(s)$, with a probability which is a function of T and $f(s') - f(s)$. The probability is generally computed following the Boltzmann distribution $\exp(-\frac{f(s')-f(s)}{T})$.

The temperature T is decreased⁶ during the search process, thus at the beginning of the search the probability of accepting uphill moves is high and it gradually decreases, converging to a simple iterative improvement algorithm. This process is analogous to the annealing process of metals and glass, which assume a low energy configuration when cooled with an appropriate cooling schedule. Regarding the search process, this means that the algorithm is the result of two combined strategies: random walk and iterative improvement. In the first phase of the search, the bias toward improvements is low and it permits the exploration of the search space; this erratic component is slowly decreased thus leading the search to converge to a (local) minimum. The probability of accepting uphill moves is controlled by two factors: the difference of the objective functions and the temperature. On the one hand, at fixed temperature, the higher the difference $f(s') - f(s)$, the lower the probability to accept a move from s to s' . On the other hand, the higher T , the higher the probability of uphill moves.

The choice of an appropriate cooling schedule is crucial for the performance of the algorithm. The cooling schedule defines the value of T at each iteration k , $T_{k+1} = Q(T_k, k)$, where $Q(T_k, k)$ is a function of the temperature and of the iteration number. Theoretical results on non-homogeneous Markov chains [Aarts et al. 1997] state that under particular conditions on the cooling schedule, the algorithm converges in probability to a global

⁶ T is not necessarily decreased in a monotonic fashion. Elaborate cooling schemes also incorporate an occasional increase of the temperature.

minimum for $k \rightarrow \infty$. More precisely:

$$\begin{aligned} &\exists \Gamma \in \mathbb{R} \text{ such that} \\ &\lim_{k \rightarrow \infty} \mathbf{p}(\text{global minimum found after} \\ &\quad k \text{ steps}) = 1 \\ &\text{iff } \sum_{k=1}^{\infty} \exp\left(\frac{\Gamma}{T_k}\right) = \infty \end{aligned}$$

A particular cooling schedule that fulfils the hypothesis for the convergence is the one that follows a logarithmic law: $T_{k+1} = \frac{\Gamma}{\log(k+k_0)}$ (where k_0 is a constant). Unfortunately, cooling schedules which guarantee the convergence to a global optimum are not feasible in applications, because they are too slow for practical purposes. Therefore, faster cooling schedules are adopted in applications. One of the most used follows a geometric law: $T_{k+1} = \alpha T_k$, where $\alpha \in (0, 1)$, which corresponds to an exponential decay of the temperature.

The cooling rule may vary during the search, with the aim of tuning the balance between diversification and intensification. For example, at the beginning of the search, T might be constant or linearly decreasing, in order to sample the search space; then, T might follow a rule such as the geometric one, to converge to a local minimum at the end of the search. More successful variants are *nonmonotonic* cooling schedules (e.g., see Osman [1993] and Lundy and Mees [1986]). Nonmonotonic cooling schedules are characterized by alternating phases of cooling and reheating, thus providing an oscillating balance between diversification and intensification.

The cooling schedule and the initial temperature should be adapted to the particular problem instance, since the cost of escaping from local minima depends on the structure of the search landscape. A simple way of empirically determining the starting temperature T_0 is to initially sample the search space with a random walk to roughly evaluate the average and the variance of objective function values. But also more elaborate schemes can be implemented [Ingber 1996].

The dynamic process described by SA is a *Markov chain* [Feller 1968], as it follows

```
s ← GenerateInitialSolution()
TabuList ← ∅
while termination conditions not met do
  s ← ChooseBestOf(N(s) \ TabuList)
  Update(TabuList)
endwhile
```

Fig. 3. Algorithm: Simple Tabu Search (TS).

a trajectory in the state space in which the successor state is chosen depending only on the incumbent one. This means that basic SA is memory-less. However, the use of memory can be beneficial for SA approaches (see, e.g., Chardaire et al. [1995]).

SA has been applied to several CO problems, such as the Quadratic Assignment Problem (QAP) [Connolly 1990] and the Job Shop Scheduling (JSS) problem [Van Laarhoven et al. 1992]. References to other applications can be found in Aarts and Lenstra [1997], Ingber [1996] and Fleischer [1995]. SA is nowadays used as a component in metaheuristics, rather than applied as stand-alone search algorithm. Variants of SA called Threshold Accepting and The Great Deluge Algorithm were presented by Dueck and Scheuer [1990] and Dueck [1993].

3.3. Tabu Search

Tabu Search (TS) is among the most cited and used metaheuristics for CO problems. TS basic ideas were first introduced in Glover [1986], based on earlier ideas formulated in Glover [1977].⁷ A description of the method and its concepts can be found in Glover and Laguna [1997]. TS explicitly uses the history of the search, both to escape from local minima and to implement an explorative strategy. We will first describe a simple version of TS, to introduce the basic concepts. Then, we will explain a more applicable algorithm and finally we will discuss some improvements.

The simple TS algorithm (see Figure 3) applies a best improvement local search as basic ingredient and uses a *short term memory* to escape from local minima and

⁷Related ideas were labelled steepest ascent/mildest descent method in Hansen [1986].

to avoid cycles. The short term memory is implemented as a *tabu list* that keeps track of the most recently visited solutions and forbids moves toward them. The neighborhood of the current solution is thus restricted to the solutions that do not belong to the tabu list. In the following we will refer to this set as *allowed set*. At each iteration the best solution from the allowed set is chosen as the new current solution. Additionally, this solution is added to the tabu list and one of the solutions that were already in the tabu list is removed (usually in a FIFO order). Due to this dynamic restriction of allowed solutions in a neighborhood, TS can be considered as a dynamic neighborhood search technique [Stützle 1999b]. The algorithm stops when a termination condition is met. It might also terminate if the allowed set is empty, that is, if all the solutions in $\mathcal{N}(s)$ are forbidden by the tabu list.⁸

The use of a tabu list prevents from returning to recently visited solutions, therefore it prevents from endless cycling⁹ and forces the search to accept even uphill moves. The length l of the tabu list (i.e., the *tabu tenure*) controls the memory of the search process. With small tabu tenures the search will concentrate on small areas of the search space. On the opposite, a large tabu tenure forces the search process to explore larger regions, because it forbids revisiting a higher number of solutions. The tabu tenure can be varied during the search, leading to more robust algorithms. An example can be found in Taillard [1991], where the tabu tenure is periodically reinitialized at random from the interval $[l_{\min}, l_{\max}]$. A more advanced use of a dynamic tabu tenure is presented in Battiti and Tecchiolli [1994] and Battiti and Protasi [1997], where the tabu tenure is increased if there is evidence for repetitions of solutions (thus a higher diversification is needed),

⁸Strategies for avoiding to stop the search when the allowed set is empty include the choice of the least recently visited solution, even if it is tabu.

⁹Cycles of higher period are possible, since the tabu list has a finite length l which is smaller than the cardinality of the search space.

while it is decreased if there are no improvements (thus intensification should be boosted). More advanced ways to create dynamic tabu tenure are described in Glover [1990].

However, the implementation of short term memory as a list that contains complete solutions is not practical, because managing a list of solutions is highly inefficient. Therefore, instead of the solutions themselves, solution *attributes* are stored.¹⁰ Attributes are usually components of solutions, moves, or differences between two solutions. Since more than one attribute can be considered, a tabu list is introduced for each of them. The set of attributes and the corresponding tabu lists define the *tabu conditions* which are used to filter the neighborhood of a solution and generate the allowed set. Storing attributes instead of complete solutions is much more efficient, but it introduces a loss of information, as forbidding an attribute means assigning the tabu status to probably more than one solution. Thus, it is possible that unvisited solutions of good quality are excluded from the allowed set. To overcome this problem, *aspiration criteria* are defined which allow to include a solution in the allowed set even if it is forbidden by tabu conditions. Aspiration criteria define the aspiration conditions that are used to construct the allowed set. The most commonly used aspiration criterion selects solutions which are better than the current best one. The complete algorithm, as described above, is reported in Figure 4.

Tabu lists are only one of the possible ways of taking advantage of the history of the search. They are usually identified with the usage of short term memory. Information collected during the whole search process can also be very useful, especially for a strategic guidance of the algorithm. This kind of long-term memory is usually added to TS by referring to four principles: *recency*, *frequency*, *quality* and *influence*. Recency-based memory records for each solution (or attribute)

¹⁰In addition to storing attributes, some longer term TS strategies also keep complete solutions (e.g., elite solutions) in the memory.

```

s ← GenerateInitialSolution()
InitializeTabuLists(TL1, . . . , TLr)
k ← 0
while termination conditions not met do
    AllowedSet(s, k) ← {s' ∈ N(s) | s does not violate a tabu condition,
                        or it satisfies at least one aspiration condition}
    s ← ChooseBestOf(AllowedSet(s, k))
    UpdateTabuListsAndAspirationConditions()
    k ← k + 1
endwhile

```

Fig. 4. Algorithm: Tabu Search (TS).

```

while termination conditions not met do
    s ← ConstructGreedyRandomizedSolution()           % see Figure 6
    ApplyLocalSearch(s)
    MemorizeBestFoundSolution()
endwhile

```

Fig. 5. Algorithm: Greedy Randomized Adaptive Search Procedure (GRASP).

the most recent iteration it was involved in. Orthogonally, frequency-based memory keeps track of how many times each solution (attribute) has been visited. This information identifies the regions (or the subsets) of the solution space where the search was confined, or where it stayed for a high number of iterations. This kind of information about the past is usually exploited to diversify the search. The third principle (i.e., quality) refers to the accumulation and extraction of information from the search history in order to identify good solution components. This information can be usefully integrated in the solution construction. Other metaheuristics (e.g., Ant Colony Optimization) explicitly use this principle to learn about good combinations of solution components. Finally, influence is a property regarding choices made during the search and can be used to indicate which choices have shown to be the most critical. In general, the TS field is a rich source of ideas. Many of these ideas and strategies have been and are currently adopted by other metaheuristics.

TS has been applied to most CO problems; examples for successful applications are the Robust Tabu Search to the QAP [Taillard 1991], the Reactive Tabu Search to the MAXSAT problem [Battiti and Protasi 1997], and to assignment problems [Dell'Amico et al. 1999]. TS ap-

proaches dominate the Job Shop Scheduling (JSS) problem area (see, e.g., Nowicki and Smutnicki [1996]) and the Vehicle Routing (VR) area [Gendreau et al. 2001]. Further current applications can be found at [Tabu Search website 2003].

3.4. Explorative Local Search Methods

In this section, we present more recently proposed trajectory methods. These are the Greedy Randomized Adaptive Search Procedure (GRASP), Variable Neighborhood Search (VNS), Guided Local Search (GLS) and Iterated Local Search (ILS).

3.4.1. GRASP. The Greedy Randomized Adaptive Search Procedure (GRASP), see Feo and Resende [1995] and Pitsoulis and Resende [2002], is a simple metaheuristic that combines constructive heuristics and local search. Its structure is sketched in Figure 5. GRASP is an iterative procedure, composed of two phases: solution construction and solution improvement. The best found solution is returned upon termination of the search process.

The solution construction mechanism (see Figure 6) is characterized by two main ingredients: a dynamic constructive heuristic and randomization. Assuming that a solution s consists of a subset of a set of elements (solution components),

```

s ← ∅           % s denotes a partial solution in this case
α ← DetermineCandidateListLength()   % definition of the RCL length
while solution not complete do
  RCLα ← GenerateRestrictedCandidateList(s)
  x ← SelectElementAtRandom(RCLα)
  s ← s ∪ {x}
  UpdateGreedyFunction(s)           % update of the heuristic values (see text)
endwhile

```

Fig. 6. Greedy randomized solution construction.

the solution is constructed step-by-step by adding one new element at a time. The choice of the next element is done by picking it uniformly at random from a candidate list. The elements are ranked by means of a heuristic criterion that gives them a score as a function of the (myopic) benefit if inserted in the current partial solution. The candidate list, called *restricted candidate list* (RCL), is composed of the best α elements. The heuristic values are updated at each step, thus the scores of elements change during the construction phase, depending on the possible choices. This constructive heuristic is called *dynamic*, in contrast to the *static* one which assigns a score to elements only before starting the construction. For instance, one of the static heuristics for the TSP is based on arc costs: the lower the cost of an arc, the higher its score. An example of a dynamic heuristic is the *cheapest insertion* heuristic, where the score of an element is evaluated depending on the current partial solution.

The length α of the restricted candidate list determines the strength of the heuristic bias. In the extreme case of $\alpha = 1$ the best element would be added, thus the construction would be equivalent to a deterministic Greedy Heuristic. On the opposite, in case $\alpha = n$ the construction would be completely random (indeed, the choice of an element from the candidate list is done at random). Therefore, α is a critical parameter which influences the sampling of the search space. In Pitsoulis and Resende [2002] the most important schemes to define α are listed. The simplest scheme is, trivially, to keep α constant; it can also be changed at each iteration, either randomly or by means of an adaptive scheme.

The second phase of the algorithm is a local search process, which may be a basic local search algorithm such as iterative improvement, or a more advanced technique such as SA or TS. GRASP can be effective if two conditions are satisfied:

- the solution construction mechanism samples the most promising regions of the search space;
- the solutions constructed by the constructive heuristic belong to basins of attraction of different locally minimal solutions;

The first condition can be met by the choice of an effective constructive heuristic and an appropriate length of the candidate list, whereas the second condition can be met by choosing the constructive heuristic and the local search in a way such that they fit well.

The description of GRASP as given above indicates that a basic GRASP does not use the history of the search process.¹¹ The only memory requirement is for storing the problem instance and for keeping the best so-far solution. This is one of the reasons why GRASP is often outperformed by other metaheuristics. However, due to its simplicity, it is generally very fast and it is able to produce quite good solutions in a very short amount of computation time. Furthermore, it can be successfully integrated into other search techniques. Among the applications of GRASP, we mention the JSS problem

¹¹ However, some extensions in this direction are cited in Pitsoulis and Resende [2002], and an example for a metaheuristic method using an adaptive greedy procedure depending on search history is Squeaky Wheel Optimization (SWO) [Joslin and Clements 1999].

```

Select a set of neighborhood structures  $\mathcal{N}_k, k = 1, \dots, k_{max}$ 
 $s \leftarrow \text{GenerateInitialSolution}()$ 
while termination conditions not met do
   $k \leftarrow 1$ 
  while  $k < k_{max}$  do           % Inner loop
     $s' \leftarrow \text{PickAtRandom}(\mathcal{N}_k(s))$            % Shaking phase
     $s'' \leftarrow \text{LocalSearch}(s')$ 
    if  $(f(s'') < f(s))$  then
       $s \leftarrow s''$ 
       $k \leftarrow 1$ 
    else
       $k \leftarrow k + 1$ 
    endif
  endwhile
endwhile

```

Fig. 7. Algorithm: Variable Neighborhood Search (VNS).

[Binato et al. 2001], the graph planarization problem [Resende and Ribeiro 1997] and assignment problems [Prais and Ribeiro 2000]. A detailed and annotated bibliography references many more applications [Festa and Resende 2002].

3.4.2. Variable Neighborhood Search. Variable Neighborhood Search (VNS) is a metaheuristic proposed in Hansen and Mladenović [1999, 2001], which explicitly applies a strategy based on dynamically changing neighborhood structures. The algorithm is very general and many degrees of freedom exist for designing variants and particular instantiations.¹²

At the initialization step, a set of neighborhood structures has to be defined. These neighborhoods can be arbitrarily chosen, but often a sequence $|\mathcal{N}_1| < |\mathcal{N}_2| < \dots < |\mathcal{N}_{k_{max}}|$ of neighborhoods with increasing cardinality is defined.¹³ Then an initial solution is generated, the neighborhood index is initialized and the algorithm iterates until a stopping condition is met (see Figure 7). VNS' main cycle is composed of three phases: *shaking*, *local search* and *move*. In the *shaking* phase a solution s' in the k th neighborhood of the current solution s is randomly selected. Then, s'

becomes the local search starting point. The local search can use any neighborhood structure and is not restricted to the set of neighborhood structures $\mathcal{N}_k, k = 1, \dots, k_{max}$. At the end of the local search process (terminated as soon as a predefined termination condition is verified) the new solution s'' is compared with s and, if it is better, it replaces s and the algorithm starts again with $k = 1$. Otherwise, k is incremented and a new shaking phase starts using a different neighborhood.

The objective of the shaking phase is to perturb the solution so as to provide a good starting point for the local search. The starting point should belong to the basin of attraction of a different local minimum than the current one, but should not be “too far” from s , otherwise the algorithm would degenerate into a simple random multi-start. Moreover, choosing s' in the neighborhood of the current best solution is likely to produce a solution that maintains some good features of the current one.

The process of changing neighborhoods in case of no improvements corresponds to a diversification of the search. In particular the choice of neighborhoods of increasing cardinality yields a progressive diversification. The effectiveness of this dynamic neighborhood strategy can be explained by the fact that a “bad” place on the search landscape given by one neighborhood could be a “good” place on the search landscape given by another neighborhood.¹⁴ Moreover, a solution that is locally optimal with respect to a neighborhood is probably not locally optimal with respect to another neighborhood. These concepts are known as “*One Operator, One Landscape*” and explained in Jones [1995a, 1995b]. The core idea is that the neighborhood structure determines the topological properties of the search landscape, that is, each neighborhood defines one landscape. The properties of a landscape are in general different from those of other landscapes, therefore

¹²The variants described in the following are also described in Hansen and Mladenović [1999, 2001].

¹³In principle they could be one included in the other, $\mathcal{N}_1 \subset \mathcal{N}_2 \subset \dots \subset \mathcal{N}_{k_{max}}$. Nevertheless, such a sequence might produce an inefficient search, because a large number of solutions could be revisited.

¹⁴A “good” place in the search space is an area from which a good local minimum can be reached.




Fig. 8. Two search landscapes defined by two different neighborhoods. On the landscape that is shown in the graphic on the left, the best improvement local search stops at s_1 , while it proceeds till a better local minimum s_2 on the landscape that is shown in the graphic on the right.

a search strategy performs differently on them (see an example in Figure 8).

This property is directly exploited by a local search called Variable Neighborhood Descent (VND). In VND a *best improvement* local search (see Section 3.1) is applied, and, in case a local minimum is found, the search proceeds with another neighborhood structure. The VND algorithm can be obtained by substituting the inner loop of the VNS algorithm (see Figure 7) with the following pseudo-code:

```

s' ← ChooseBestOf( $\mathcal{N}_k(s)$ )
if ( $f(s') < f(s)$ )
then % i.e., if a better solution is found in  $\mathcal{N}_k(s)$ 
    s ← s'
else % i.e., s is a local minimum
    k ← k + 1
endif

```

As can be observed from the description as given above, the choice of the neighborhood structures is the critical point of VNS and VND. The neighborhoods chosen should exploit different properties and characteristics of the search space, that is, the neighborhood structures should provide different *abstractions* of the search space. A variant of VNS is obtained by selecting the neighborhoods in such a way as to produce a problem decomposition (the algorithm is called Variable Neighborhood Decomposition Search—VNDS). VNDS follows the usual VNS scheme, but the neighborhood structures and the local search are defined on sub-problems. For each solution, all attributes (usually variables) are kept fixed except for k of them. For each, k , a neighborhood struc-

ture \mathcal{N}_k is defined. Local search only regards changes on the variables belonging to the sub-problem it is applied to. The inner loop of VNDS is the following:

```

s' ← PickAtRandom( $\mathcal{N}_k(s)$ ) % s and s' differ in k
                           attributes
s'' ← LocalSearch(s', Attributes) % only moves
                                   involving the k
                                   attributes are
                                   allowed

if ( $f(s'') < f(s)$ ) then
    s ← s''
    k ← 1
else
    k ← k + 1
endif

```

The decision whether to perform a move can be varied as well. The acceptance criterion based on improvements is strongly steepest descent-oriented and it might not be suited to effectively explore the search space. For example, when local minima are clustered, VNS can quickly find the best optimum in a cluster, but it has no guidance to leave that cluster and find another one. Skewed VNS (SVNS) extends VNS by providing a more flexible acceptance criterion that takes also into account the distance from the current solution.¹⁵ The new acceptance criterion is the following: besides always accepting improvements, worse solutions can be accepted if the distance from the current one is less than a value $\alpha\rho(s, s'')$. The function $\rho(s, s'')$ measures the distance between s and s''

¹⁵A distance measure between solutions has thus to be formally defined.




Fig. 9. Basic GLS idea: Escaping from a valley in the landscape by increasing the objective function value of its solutions.

and α is a parameter that weights the importance of the distance between the two solutions in the acceptance criterion. The inner loop of SVNS can be sketched as follows:

```

if ( $f(s'') - \alpha \rho(s, s'') < f(s)$ ) then
     $s \leftarrow s''$ 
     $k \leftarrow 1$ 
else
     $k \leftarrow k + 1$ 
endif
    
```

VNS and its variants have been successfully applied to graph based CO problems such as the p -Median problem [Hansen and Mladenović 1997], the degree constrained minimum spanning tree problem [Ribeiro and Souza 2002], the Steiner tree problem [Wade and Rayward-Smith 1997] and the k -Cardinality Tree (KCT) problem [Mladenović and Urošević 2001]. References to more applications can be found in Hansen and Mladenović [2001].

3.4.3. Guided Local Search. Tabu Search and Variable Neighborhood Search explicitly deal with dynamic neighborhoods with the aim of efficiently and effectively exploring the search space. A different approach for guiding the search is to dynamically change the objective function. Among the most general methods that use this approach is Guided Local Search (GLS) [Voudouris and Tsang 1999; Voudouris 1997].

The basic GLS principle is to help the search to gradually move away from local minima by changing the search landscape. In GLS, the set of solutions and the neighborhood structure are kept fixed, while the objective function f is dynamically changed with the aim of making the current local optimum “less desirable”. A pictorial description of this idea is given in Figure 9.

The mechanism used by GLS is based on *solution features*, which may be any kind of properties or characteristics that can be used to discriminate between solutions. For example, solution features in the TSP could be arcs between pairs of cities, while in the MAXSAT problem they could be the number of unsatisfied clauses. An indicator function $I_i(s)$ indicates whether the feature i is present in solution s :

$$I_i(s) = \begin{cases} 1 & : \text{if feature } i \text{ is present in} \\ & : \text{solution } s \\ 0 & : \text{otherwise.} \end{cases}$$

The objective function f is modified to yield a new objective function f' by adding a term that depends on the m features:

$$f'(s) = f(s) + \lambda \sum_{i=1}^m p_i \cdot I_i(s),$$

where p_i are called *penalty parameters* and λ is called the *regularization*

```

s ← GenerateInitialSolution()
while termination conditions not met do
  s ← LocalSearch(s, f')
  for all feature i with maximum utility Util(s, i) do
    pi ← pi + 1
  endfor
  Update(f', p)      % p is the penalty vector
endwhile

```

Fig. 10. Algorithm: Guided Local Search (GLS).

parameter. The penalty parameters weight the importance of the features: the higher p_i , the higher the importance of feature i , thus the higher the cost of having that feature in the solution. The regularization parameter balances the relevance of features with respect to the original objective function.

The algorithm (see Figure 10) works as follows: It starts from an initial solution and applies a local search method until a local minimum is reached. Then the array $\mathbf{p} = (p_1, \dots, p_m)$ of penalties is updated by incrementing some of the penalties and the local search is started again. The penalized features are those that have the maximum *utility*:

$$Util(s, i) = I_i(s) \cdot \frac{c_i}{1 + p_i},$$

where c_i are *costs* assigned to every feature i giving a heuristic evaluation of the relative importance of features with respect to others. The higher the cost, the higher the utility of features. Nevertheless, the cost is scaled by the penalty parameter to prevent the algorithm from being totally biased toward the cost and to make it sensitive to the search history.

The penalties update procedure can be modified by adding a multiplicative rule to the simple incrementing rule (that is applied at each iteration). The multiplicative rule has the form: $p_i \leftarrow p_i \cdot \alpha$, where $\alpha \in (0, 1)$. This rule is applied with a lower frequency than the incrementing one (for example every few hundreds of iterations) with the aim of smoothing the weights of penalized features so as to prevent the landscape from becoming too rugged. It is important to note that the penalties up-

date rules are often very sensitive to the problem instance.

GLS has been successfully applied to the weighted MAXSAT [Mills and Tsang 2000], the VR problem [Kilby et al. 1999], the TSP and the QAP [Voudouris and Tsang 1999].

3.4.4. Iterated Local Search. We conclude this presentation of explorative strategies with Iterated Local Search (ILS), the most general scheme among the explorative strategies. On the one hand, its generality makes it a framework for other metaheuristics (such as VNS); on the other hand, other metaheuristics can be easily incorporated as subcomponents. ILS is a simple but powerful metaheuristic algorithm [Stützle 1999a,1999b; Lourenço et al. 2001,2002; Martin et al. 1991]. It applies local search to an initial solution until it finds a local optimum; then it perturbs the solution and it restarts local search. The importance of the perturbation is obvious: too small a perturbation might not enable the system to escape from the basin of attraction of the local optimum just found. On the other side, too strong a perturbation would make the algorithm similar to a random restart local search.

A local search is effective if it is able to find good local minima, that is, if it can find the basin of attraction of those states. When the search space is wide and/or when the basins of attraction of good local optima are small,¹⁶ a simple multi-start algorithm is almost useless. An effective search could be designed as a trajectory only in the set of local optima \hat{S} , instead of in the set S of all the states. Unfortunately, in most cases there is no feasible way of introducing a neighborhood structure for \hat{S} . Therefore, a trajectory along local optima $\hat{s}_1, \hat{s}_2, \dots, \hat{s}_t$ is performed, without explicitly introducing a neighborhood structure, by applying the following scheme:

¹⁶The basin of attraction size of a point s (in a finite space), is defined as the fraction of initial states of trajectories which converge to point s .

```

s0 ← GenerateInitialSolution()
ŝ ← LocalSearch(s0)
while termination conditions not met do
  s' ← Perturbation(ŝ, history)
  ŝ' ← LocalSearch(s')
  ŝ ← ApplyAcceptanceCriterion(ŝ, ŝ', history)
endwhile

```

Fig. 11. Algorithm: Iterated Local Search (ILS).

- (1) Execute local search (LS) from an initial state s until a local minimum \hat{s} is found.
- (2) *Perturb* \hat{s} and obtain s' .
- (3) Execute LS from s' until a local minimum \hat{s}' is reached.
- (4) On the basis of an *acceptance criterion* decide whether to set $\hat{s} \leftarrow \hat{s}'$.
- (5) Goto step 2.

The requirement on the *perturbation* of \hat{s} is to produce a starting point for local search such that a local minimum different from \hat{s} is reached. However, this new local minimum should be *closer* to \hat{s} than a local minimum produced by a random restart. The *acceptance criterion* acts as a counterbalance, as it filters and gives feedback to the perturbation action, depending on the characteristics of the new local minimum. A high level description of ILS as it is described in Lourenço et al. [2002] is given in Figure 11. Figure 12 shows a possible (lucky) ILS step.

The design of ILS algorithms has several degrees of freedom in the choice of the initial solution, perturbation and acceptance criteria. A key role is played by the *history* of the search which can be exploited both in form of short and long term memory.

The construction of initial solutions should be fast (computationally not expensive), and initial solutions should be a good starting point for local search. The fastest way of producing an initial solution is to generate it at random; however, this is the easiest way for problems that are constrained, whilst in other cases the construction of a feasible solution requires also constraint checking. Constructive methods, guided by heuristics, can also be adopted. It is worth underlining

that an initial solution is considered a good starting point depending on the particular LS applied and on the problem structure, thus the algorithm designer's goal is to find a trade-off between speed and quality of solutions.

The perturbation is usually non-deterministic in order to avoid cycling. Its most important characteristic is the *strength*, roughly defined as the amount of changes made on the current solution. The strength can be either fixed or variable. In the first case, the distance between \hat{s} and s' is kept constant, independently of the problem size. However, a variable strength is in general more effective, since it has been experimentally found that, in most of the problems, the bigger the problem size, the larger should be the strength. More sophisticated schemes are possible; for example, the strength can be adaptive: it increases when more diversification is needed and it decreases when intensification seems preferable. VNS and its variants belong to this category. A second choice is the mechanism to perform perturbations. This may be a random mechanism, or the perturbation may be produced by a (semi-)deterministic method (e.g., a LS different from the one used in the main algorithm).

The third important component is the acceptance criterion. Two extreme examples consist in (1) accepting the new local optimum only in case of improvement and (2) in always accepting the new solution. In-between, there are several possibilities. For example, it is possible to adopt a kind of annealing schedule: accept all the improving new local optima and accept also the nonimproving ones with a probability that is a function of the temperature T and the difference of objective function values. In formulas:

$$\begin{aligned}
 & \mathbf{p}(\text{Accept}(\hat{s}, \hat{s}', \text{history})) \\
 &= \begin{cases} 1 & \text{if } f(\hat{s}') < f(\hat{s}) \\ \exp\left(-\frac{f(\hat{s}') - f(\hat{s})}{T}\right) & \text{otherwise} \end{cases}
 \end{aligned}$$

The cooling schedule can be either monotonic (non-increasing in time) or




Fig. 12. A desirable ILS step: the local minimum \hat{s} is perturbed, then LS is applied and a new local minimum is found.

non-monotonic (adapted to tune the balance between diversification and intensification). The nonmonotonic schedule is particularly effective if it exploits the history of the search, in a way similar to the Reactive Tabu Search [Taillard 1991] mentioned at the end of the section about Tabu Search. When intensification seems no longer effective, a diversification phase is needed and the temperature is increased.

Examples for successful applications of ILS are to the TSP [Martin and Otto 1996; Johnson and McGeoch 1997], to the QAP [Lourenço et al. 2002], and to the Single Machine Total Weighted Tardiness (SMTWT) problem [den Besten et al. 2001]. References to other applications can be found in Lourenço et al. [2002].

4. POPULATION-BASED METHODS

Population-based methods deal in every iteration of the algorithm with a set (i.e., a population) of solutions¹⁷ rather than with a single solution. As they deal with a population of solutions, population-based algorithms provide a natural, intrinsic way for the exploration of the search space. Yet, the final performance depends strongly on the way the population is manipulated. The most studied population-

¹⁷In general, especially in EC algorithms, we talk about a population of individuals rather than solutions.

based methods in combinatorial optimization are Evolutionary Computation (EC) and Ant Colony Optimization (ACO). In EC algorithms, a population of individuals is modified by recombination and mutation operators, and in ACO a colony of artificial ants is used to construct solutions guided by the pheromone trails and heuristic information.

4.1. Evolutionary Computation

Evolutionary Computation (EC) algorithms are inspired by nature's capability to evolve living beings well adapted to their environment. EC algorithms can be succinctly characterized as computational models of evolutionary processes. At each iteration a number of operators is applied to the individuals of the current population to generate the individuals of the population of the next generation (iteration). Usually, EC algorithms use operators called *recombination* or *crossover* to recombine two or more individuals to produce new individuals. They also use *mutation* or *modification* operators which cause a self-adaptation of individuals. The driving force in evolutionary algorithms is the *selection* of individuals based on their *fitness* (this can be the value of an objective function or the result of a simulation experiment, or some other kind of quality measure). Individuals with a higher

fitness have a higher probability to be chosen as members of the population of the next iteration (or as parents for the generation of new individuals). This corresponds to the principle of *survival of the fittest* in natural evolution. It is the capability of nature to adapt itself to a changing environment, which gave the inspiration for EC algorithms.

There has been a variety of slightly different EC algorithms proposed over the years. Basically they fall into three different categories which have been developed independently from each other. These are Evolutionary Programming (EP) developed by Fogel [1962] and Fogel et al. [1966], Evolutionary Strategies (ES) proposed by Rechenberg [1973] and Genetic Algorithms initiated by Holland [1975] (see Goldberg [1989], Mitchell [1998], Reeves and Rowe [2002], and Vose [1999] for further references). EP arose from the desire to generate machine intelligence. While EP originally was proposed to operate on discrete representations of finite state machines, most of the present variants are used for continuous optimization problems. The latter also holds for most present variants of ES, whereas GAs are mainly applied to solve combinatorial optimization problems. Over the years, there have been quite a few overviews and surveys about EC methods. Among those are the ones by Bäck [1996], Fogel [1994], Spears et al. [1993] and Michalewicz and Michalewicz [1997]. Calégary et al. [1999] propose a taxonomy of EC algorithms.

In the following we provide a “combinatorial optimization”-oriented introduction to EC algorithms. For doing this, we follow an overview work by Hertz and Kobler [2000], which gives, in our opinion, a good overview of the different components of EC algorithms and of the possibilities to define them.

Figure 13 shows the basic structure of every EC algorithm. In this algorithm, P denotes the population of individuals. A population of offspring is generated by the application of *recombination* and *mutation* operators and the individuals for the next population are *selected* from the union of the old population and the

```

 $P \leftarrow \text{GenerateInitialPopulation}()$ 
Evaluate( $P$ )
while termination conditions not met do
   $P' \leftarrow \text{Recombine}(P)$ 
   $P'' \leftarrow \text{Mutate}(P')$ 
  Evaluate( $P''$ )
   $P \leftarrow \text{Select}(P'' \cup P)$ 
endwhile

```

Fig. 13. Algorithm: Evolutionary Computation (EC).

offspring population. The main features of an EC algorithm are outlined in the following.

Description of the Individuals. EC algorithms handle populations of individuals. These individuals are not necessarily solutions of the considered problem. They may be partial solutions, or sets of solutions, or any object which can be transformed into one or more solutions in a structured way. Most commonly used in combinatorial optimization is the representation of solutions as bit-strings or as permutations of n integer numbers. Tree-structures or other complex structures are also possible. In the context of Genetic Algorithms, individuals are called *genotypes*, whereas the solutions that are encoded by individuals are called *phenotypes*. This is to differentiate between the representation of solutions and solutions themselves. The choice of an appropriate representation is crucial for the success of an EC algorithm. Holland’s [1975] schema analysis and Radcliffe’s [1991] generalization to formae are examples of how theory can help to guide representation choices.

Evolution Process. In each iteration, it has to be decided which individuals will enter the population of the next iteration. This is done by a selection scheme. To choose, the individuals for the next population exclusively from the offspring is called *generational replacement*. If it is possible to transfer individuals of the current population into the next population, then we deal with a so-called *steady state* evolution process.

Most EC algorithms work with populations of fixed size keeping at least the best individual always in the current

population. It is also possible to have a variable population size. In case of a continuously shrinking population size, the situation where only one individual is left in the population (or no crossover partners can be found for any member of the population) might be one of the stopping conditions of the algorithm.

Neighborhood Structure. A neighborhood function $\mathcal{N}_{EC} : \mathcal{I} \rightarrow 2^{\mathcal{I}}$ on the set of individuals \mathcal{I} assigns to every individual $i \in \mathcal{I}$ a set of individuals $\mathcal{N}_{EC}(i) \subseteq \mathcal{I}$ which are permitted to act as recombination partners for i to create offspring. If an individual can be recombined with any other individual (as for example in the simple GA) we talk about *unstructured* populations, otherwise we talk about *structured* populations. An example for an EC algorithm that works on structured populations is the Parallel Genetic Algorithm proposed by Mühlenbein [1991].

Information Sources. The most common form of information sources to create offspring (i.e., new individuals) is a couple of parents (two-parent crossover). But there are also recombination operators that operate on more than two individuals to create a new individual (multi-parent crossover), see Eiben et al. [1994]. More recent developments even use population statistics for generating the individuals of the next population. Examples are the recombination operators called Gene Pool Recombination [Mühlenbein and Voigt 1995] and Bit-Simulated Crossover [Syswerda 1993] which make use of a distribution over the search space given by the current population to generate the next population.

Infeasibility. An important characteristic of an EC algorithm is the way it deals with infeasible individuals. When recombining individuals, the offspring might be potentially infeasible. There are basically three different ways to handle such a situation. The most simple action is to *reject* infeasible individuals. Nevertheless, for many problems (e.g., for timetabling problems) it might be very difficult to find feasible individuals. Therefore, the strategy of *penalizing* infeasible individuals in the function that measures the quality of

an individual is sometimes more appropriate (or even unavoidable). The third possibility consists in trying to *repair* an infeasible solution (see Eiben and Ruttkay [1997] for an example).

Intensification Strategy. In many applications it proved to be quite beneficial to use improvement mechanisms to improve the fitness of individuals. EC algorithms that apply a local search algorithm to every individual of a population are often called Memetic Algorithms [Moscato 1989, 1999]. While the use of a population ensures an exploration of the search space, the use of local search techniques helps to quickly identify “good” areas in the search space.

Another intensification strategy is the use of recombination operators that explicitly try to combine “good” parts of individuals (rather than, e.g., a simple one-point crossover for bit-strings). This may guide the search performed by EC algorithms to areas of individuals with certain “good” properties. Techniques of this kind are sometimes called *linkage learning* or *building block learning* (see Goldberg et al. [1991], van Kemenade [1996], Watson et al. [1998], and Harik [1999] as examples). Moreover, generalized recombination operators have been proposed in the literature, which incorporate the notion of “neighborhood search” into EC. An example can be found in Rayward-Smith [1994].

Diversification Strategy. One of the major difficulties of EC algorithms (especially when applying local search) is the premature convergence toward sub-optimal solutions. The most simple mechanism to diversify the search process is the use of a mutation operator. The simple form of a mutation operator just performs a small random perturbation of an individual, introducing a kind of *noise*. In order to avoid premature convergence there are ways of maintaining the population diversity. Probably the oldest strategies are *crowding* [DeJong 1975] and its close relative, *preselection*. Newer strategies are *fitness sharing* [Goldberg and Richardson 1987], respectively *niching*, whereby the reproductive fitness allocated to an individual

```

Initial Phase:
SeedGeneration()
repeat
  DiversificationGenerator()
  Improvement()
  ReferenceSetUpdate()
until the reference set is of cardinality  $n$ 

Scatter Search/Path Relinking Phase:
repeat
  SubsetGeneration()
  SolutionCombination()
  Improvement()
  ReferenceSetUpdate()
until termination criteria met

```

Fig. 14. Algorithm: Scatter Search and Path Relinking.

in a population is reduced proportionally to the number of other individuals that share the same region of the search space.

This concludes the list of the main features of EC algorithms. EC algorithms have been applied to most CO problems and optimization problems in general. Recent successes were obtained in the rapidly growing bioinformatics area (see, e.g., Fogel et al. [2002]), but also in multi-objective optimization [Coello Coello 2000], and in evolvable hardware [Sipper et al. 1997]. For an extensive collection of references to EC applications, we refer to Bäck et al. [1997]. In the following two sections, we are going to introduce two other populations-based methods that are sometimes also regarded as being EC algorithms.

4.1.1. Scatter Search and Path Relinking. Scatter Search and its generalized form called Path Relinking [Glover 1999; Glover et al. 2000] differ from EC algorithms mainly by providing unifying principles for joining (or recombining) solutions based on generalized path constructions in Euclidean or neighborhood spaces. They also incorporate some ideas originating from Tabu Search methods, as, for example, the use of adaptive memory and associated memory-exploiting mechanisms. The template for Scatter Search (respectively, Path Relinking) is shown in Figure 14.

Scatter Search (respectively, Path Relinking) is a search strategy that generates a set of solutions from a chosen set of *reference solutions* corresponding to feasible solutions to the problem under consideration. This is done by making combinations of subsets of the current set of reference solutions. The resulting solutions are called *trial solutions*. These trial solutions may be infeasible solutions and are therefore usually modified by means of a repair procedure that transforms them into feasible solutions. An improvement mechanism is then applied in order to try to improve the set of trial solutions (usually this improvement procedure is a local search). These improved solutions form the set of *dispersed solutions*. The new set of reference solutions that will be used in the next iteration is selected from the current set of reference solutions and the newly created set of dispersed solutions. The components of the pseudo-code, which is shown in Figure 14, are explained in the following:

SeedGeneration(): One or more seed solutions, which are arbitrary trial solutions, are created and used to initiate the remainder of the method.

DiversificationGenerator(): This is a procedure to generate a collection of diverse trial solutions from an arbitrary trial solution (or seed solution).

Improvement(): In this procedure, an improvement mechanism—usually a local search—is used to transform a trial solution into one or more enhanced trial solutions. Neither the input nor the output solutions are required to be feasible, though the output solutions will more usually be expected to be so. It might be necessary to apply repair methods to infeasible solutions.

ReferenceSetUpdate(): The procedure for updating the reference set is responsible for building and maintaining a reference set consisting of a number of “best” solutions found in the course of the algorithm. The attribute “best” covers features such as quality of solutions and diversity of solutions (the solutions in the reference set should be of good quality and they should be diverse).

SubsetGeneration(): This method operates on the reference set, to produce a subset of its solutions as a basis for creating combined solutions.

SolutionCombination(): A procedure to transform a given subset of solutions produced by the subset generation method into one or more combined solutions. In Scatter Search, which was introduced for solutions encoded as points in the Euclidean space, new solutions are created by building linear combinations of reference solutions using both positive and negative weights. This means that trial solutions can be both, inside and outside the convex region spanned by the reference solutions. In Path Relinking, the concept of combining solutions by making linear combinations of reference points is generalized to neighborhood spaces. Linear combinations of points in the Euclidean space can be re-interpreted as paths between and beyond solutions in a neighborhood space. To generate the desired paths, it is only necessary to select moves that satisfy the following condition: upon starting from an *initiating solution*, the moves must progressively introduce attributes contributed by a *guiding solution*. Multiparent path generation possibilities emerge in Path Relinking by considering the combined attributes provided by a set of guiding solutions, where these attributes are weighted to determine which moves are given higher priority.

Scatter Search enjoys increasing interest in recent years. Among other problems, it has been applied to multi-objective assignment problems [Laguna et al. 2000] and the Linear Ordering Problem (LOP) [Campos et al. 2001]. For further references, we refer to [Glover et al. 2002]. Path relinking is often used as a component in metaheuristics such as Tabu Search [Laguna et al. 1999] and GRASP [Aiex et al. 2003; Laguna and Martí 1999].

4.1.2. Estimation of Distribution Algorithms. In the last decade, more and more researchers tried to overcome the drawbacks of usual recombination operators of EC algorithms, which are likely to

```

P ← InitializePopulation()
while termination criteria not met do
  Psel ← Select(P)           % Psel ⊆ P
  p(x) = p(x | Psel) ← EstimateProbabilityDistribution()
  P ← SampleProbabilityDistribution()
endwhile

```

Fig. 15. Algorithm: Estimation of Distribution Algorithms (EDAs).

break good building blocks. So, a number of algorithms—sometimes called Estimation of Distribution Algorithms (EDA) [Mühlenbein and Paaß 1996]—have been developed (see Figure 15 for the algorithmic framework). These algorithms, which have a theoretical foundation in probability theory, are also based on populations that evolve as the search progresses. EDAs use probabilistic modelling of promising solutions to estimate a distribution over the search space, which is then used to produce the next generation by sampling the search space according to the estimated distribution. After every iteration, the distribution is re-estimated. For a survey of EDAs, see [Pelikan et al. 1999b].

One of the first EDAs that was proposed for Combinatorial Optimization is called Population-based Incremental Learning (PBIL) [Baluja 1994; Baluja and Caruana 1995]. The objective of this method is to create a real valued probability vector (each position corresponds to a binary decision variable) which—when used to sample the search space—generates high quality solutions with high probability. Initially, the values of the probability vector are initialized to 0.5 (for each variable there is equal probability to be set to 0 or 1). The goal of shifting the values of this probability vector in order to generate high quality solutions is accomplished as follows: a number of solution vectors are generated according to the probability vector. Then, the probability vector is shifted toward the generated solution vector(s) with highest quality. The distance that the probability vector is shifted depends on the learning rate parameter. Then, a mutation operator is applied to the probability vector. After that, the cycle is repeated. The probability vector can be regarded as a prototype vector for generating high-quality solution vectors

with respect to the available knowledge about the search space. The drawback of this method is the fact that it does not automatically provide a way to deal with constrained problems. In contrast to PBIL, which estimates a distribution of promising solutions assuming that the decision variables are independent, various other approaches try to estimate distributions taking into account dependencies between decision variables. An example for EDAs regarding pairwise dependencies between decision variables is MIMIC [de Bonet et al. 1997] and an example for multivariate dependencies is the Bayesian Optimization Algorithm (BOA) [Pelikan et al. 1999a].

The field of EDAs is still quite young and much of the research effort is focused on methodology rather than high-performance applications. Applications to Knapsack problems, the Job Shop Scheduling (JSS) problem, and other CO problems can be found in Larrañaga and Lozano [2002].

4.2. Ant Colony Optimization

Ant Colony Optimization (ACO) is a metaheuristic approach proposed in Dorigo 1992, 1996, 1999. In the course of this section, we keep close to the description of ACO as given in Dorigo and Di Caro [1999]. The inspiring source of ACO is the foraging behavior of real ants. This behavior—as described by Deneubourg et al. [1990]—enables ants to find shortest paths between food sources and their nest. While walking from food sources to the nest and vice versa, ants deposit a substance called *pheromone* on the ground. When they decide about a direction to go, they choose with higher probability paths that are marked by stronger pheromone concentrations. This basic behavior is the basis for a cooperative interaction which leads to the emergence of shortest paths.

ACO algorithms are based on a parametrized probabilistic model—the *pheromone model*—that is used to model the chemical pheromone trails. Artificial ants incrementally construct solutions by adding opportunely defined solution

components to a partial solution under consideration.¹⁸ For doing that, artificial ants perform randomized walks on a completely connected graph $\mathcal{G} = (\mathcal{C}, \mathcal{L})$ whose vertices are the solution components \mathcal{C} and the set \mathcal{L} are the connections. This graph is commonly called *construction graph*. When a constrained CO problem is considered, the problem constraints Ω are built into the ants' constructive procedure in such a way that in every step of the construction process only feasible solution components can be added to the current partial solution. In most applications, ants are implemented to build feasible solutions, but sometimes it is unavoidable to also let them construct infeasible solutions. Components $c_i \in \mathcal{C}$ can have associated a *pheromone trail parameter* T_i , and connections $l_{ij} \in \mathcal{L}$ can have associated a *pheromone trail parameter* T_{ij} . The set of all pheromone trail parameters is denoted by \mathcal{T} . The values of these parameters—the *pheromone values*—are denoted by τ_i , respectively τ_{ij} . Furthermore, components and connections can have associated a *heuristic value* η_i , respectively η_{ij} , representing *a priori* or run time heuristic information about the problem instance. We henceforth denote the set of all heuristic values by \mathcal{H} . These values are used by the ants to make probabilistic decisions on how to move on the construction graph. The probabilities involved in moving on the construction graph are commonly called *transition probabilities*. The first ACO algorithm proposed in the literature is called Ant System (AS) [Dorigo et al. 1996]. The pseudo-code for this algorithm is shown in Figure 16. For the sake of simplicity, we restrict the following description of AS to pheromone trail parameters and heuristic information on solution components.

In this algorithm, \mathcal{A} denotes the set of ants and s_a denotes the solution constructed by ant $a \in \mathcal{A}$. After the initialization of the pheromone values, at each step of the algorithm each ant constructs a

¹⁸Therefore, the ACO metaheuristic can be applied to any CO problem for which a constructive procedure can be defined.

```

InitializePheromoneValues( $\mathcal{T}$ )
while termination conditions not met do
  for all ants  $a \in \mathcal{A}$  do
     $s_a \leftarrow$  ConstructSolution( $\mathcal{T}, \mathcal{H}$ )
  endfor
  ApplyOnlineDelayedPheromoneUpdate( $\mathcal{T}, \{s_a \mid a \in \mathcal{A}\}$ )
endwhile

```

Fig. 16. Algorithm: Ant System (AS).

solution. These solutions are then used to update the pheromone values. The components of this algorithm are explained in more detail in the following.

InitializePheromoneValues(\mathcal{T}): At the beginning of the algorithm the pheromone values are initialized to the same small value $ph > 0$.

ConstructSolution(\mathcal{T}, \mathcal{H}): In the construction phase an ant incrementally builds a solution by adding solution components to the partial solution constructed so far. The probabilistic choice of the next solution component to be added is done by means of transition probabilities, which in AS are determined by the following *state transition rule*:

$$\mathbf{p}(c_r | s_a[c_l]) = \begin{cases} \frac{[\eta_r]^\alpha [\tau_r]^\beta}{\sum_{c_u \in \mathcal{J}(s_a[c_l])} [\eta_u]^\alpha [\tau_u]^\beta} & \text{if } c_r \in \mathcal{J}(s_a[c_l]) \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

In this formula, α and β are parameters to adjust the relative importance of heuristic information and pheromone values and $\mathcal{J}(s_a[c_l])$ denotes the set of solution components that are allowed to be added to the partial solution $s_a[c_l]$, where c_l is the last component that was added.

ApplyOnlineDelayedPheromoneUpdate($\mathcal{T}, \{s_a \mid a \in \mathcal{A}\}$): Once all ants have constructed a solution, the *online delayed pheromone update rule* is applied:

$$\tau_j \leftarrow (1 - \rho) \cdot \tau_j + \sum_{a \in \mathcal{A}} \Delta \tau_j^{s_a} \quad (2)$$

$\forall T_j \in \mathcal{T}$, where

$$\Delta \tau_j^{s_a} = \begin{cases} F(s_a) & \text{if } c_j \text{ is a component of } s_a \\ 0 & \text{otherwise,} \end{cases} \quad (3)$$

```

while termination conditions not met do
  ScheduleActivities
    AntBasedSolutionConstruction()
    PheromoneUpdate()
    DaemonActions() % optional
  end ScheduleActivities
endwhile

```

Fig. 17. Algorithm: Ant Colony Optimization (ACO).

where $F: \mathcal{S} \mapsto \mathbb{R}^+$ is a function that satisfies $f(s) < f(s') \Rightarrow F(s) \geq F(s')$, $\forall s \neq s' \in \mathcal{S}$. $F(\cdot)$ is commonly called the quality function. Furthermore, $0 < \rho \leq 1$ is the pheromone evaporation rate. This pheromone update rule aims at an increase of pheromone on solution components that have been found in high-quality solutions.

In the following, we describe the more general ACO metaheuristic, which is based on the same basic principles as AS. The ACO metaheuristic framework that is shown in Figure 17 covers all the improvements and extensions of AS which have been developed over the years. It consists of three parts gathered in the **ScheduleActivities** construct. The **ScheduleActivities** construct does not specify how these three activities are scheduled and synchronized. This is up to the algorithm designer.

AntBasedSolutionConstruction(): An ant constructively builds a solution to the problem by moving through nodes of the construction graph \mathcal{G} . Ants move by applying a stochastic local decision policy that makes use of the pheromone values and the heuristic values on components and/or connections of the construction graph (e.g., see the state transition rule of AS). While moving, the ant keeps in memory the partial solution it has built in terms of the path it was walking on the construction graph.

PheromoneUpdate(): When adding a component c_j to the current partial solution, an ant can update the pheromone trail(s) τ_i and/or τ_{ij} (in case the ant was walking on connection l_{ij} in order to reach component c_j). This kind of pheromone update is called *online step-by-step pheromone update*. Once an ant has built a solution, it can retrace the same path backward (by using its memory) and update

the pheromone trails of the used components and/or connections according to the quality of the solution it has built. This is called *online delayed pheromone update*. Pheromone evaporation is the process by means of which the pheromone trail intensity on the components decreases over time. From a practical point of view, pheromone evaporation is needed to avoid a too rapid convergence of the algorithm toward a sub-optimal region. It implements a useful form of *forgetting*, favoring the exploration of new areas in the search space.

DaemonActions(): Daemon actions can be used to implement centralized actions which cannot be performed by single ants. Examples are the use of a local search procedure applied to the solutions built by the ants, or the collection of global information that can be used to decide whether it is useful or not to deposit additional pheromone to bias the search process from a nonlocal perspective. As a practical example, the daemon can observe the path found by each ant in the colony and choose to deposit extra pheromone on the components used by the ant that built the best solution. Pheromone updates performed by the daemon are called *offline pheromone updates*.

Within the ACO metaheuristic framework, as shortly described above, the currently best performing versions in practice are Ant Colony System (ACS) [Dorigo and Gambardella 1997] and *MAX-MIN* Ant System (*MMAS*) [Stützle and Hoos 2000]. In the following, we are going to briefly outline the peculiarities of these algorithms.

Ant Colony System (ACS). The ACS algorithm has been introduced to improve the performance of AS. ACS is based on AS but presents some important differences. First, the daemon updates pheromone trails offline: At the end of an iteration of the algorithm—once all the ants have built a solution—pheromone is added to the arcs used by the ant that found the best solution from the start of the algorithm. Second, ants use a different decision rule to decide to which component to move next in the construction graph. The rule is called *pseudo-random-proportional*

rule. With this rule, some moves are chosen deterministically (in a greedy manner), others are chosen probabilistically with the usual decision rule. Third, in ACS, ants perform only online step-by-step pheromone updates. These updates are performed to favor the emergence of other solutions than the best so far.

MAX-MIN Ant System (MMAS). *MMAS* is also an extension of AS. First, the pheromone trails are only updated offline by the daemon (the arcs that were used by the iteration best ant or the best ant since the start of the algorithm receive additional pheromone). Second, the pheromone values are restricted to an interval $[\tau_{\min}, \tau_{\max}]$ and the pheromone trails are initialized to their maximum value τ_{\max} . Explicit bounds on the pheromone trails prevent that the probability to construct a solution falls below a certain value greater than 0. This means that the chance of finding a global optimum never vanishes during the course of the algorithm.

Recently, researchers have been dealing with finding similarities between ACO algorithms and probabilistic learning algorithms such as EDAs. An important step into this direction was the development of the Hyper-Cube Framework for Ant Colony Optimization (HC-ACO) [Blum et al. 2001]. An extensive study on this subject has been presented in Zlochin et al. [2004], where the authors present a unifying framework for so-called Model-Based Search (MBS) algorithms. Also, the close relation of algorithms like Population-Based Incremental Learning (PBIL) [Baluja and Caruana 1995] and the Univariate Marginal Distribution Algorithm (UMDA) [Mühlenbein and Paaß 1996] to ACO algorithms in the Hyper-Cube Framework has been shown. We refer the interested reader to Zlochin et al. [2004] for more information on this subject. Furthermore, connections of ACO algorithms to Stochastic Gradient Descent (SGD) algorithms are shown in Meuleau and Dorigo [2002].

Successful applications of ACO include the application to routing in communication networks [Di Caro and Dorigo

1998], the application to the Sequential Ordering Problem (SOP) [Gambardella and Dorigo 2000], and the application to Resource Constraint Project Scheduling (RCPS) [Merkle et al. 2002]. Further references to applications of ACO can be found in Dorigo and Stützle [2002, 2003].

5. A UNIFYING VIEW ON INTENSIFICATION AND DIVERSIFICATION

In this section, we take a closer look at the concepts of *intensification* and *diversification* as the two powerful forces driving metaheuristic applications to high performance. We give a view on metaheuristics that is characterized by the way intensification and diversification are implemented. Although the relevance of these two concepts is commonly agreed, so far there is no unifying description to be found in the literature. Descriptions are very generic and metaheuristic specific. Therefore most of them can be considered incomplete and sometimes they are even opposing. Depending on the paradigm behind a particular metaheuristic, intensification and diversification are achieved in different ways. Even so, we propose a unifying view on intensification and diversification. Furthermore, this discussion could lead to the goal-directed development of hybrid algorithms combining concepts originating from different metaheuristics.

5.1. Intensification and Diversification

Every metaheuristic approach should be designed with the aim of effectively and efficiently exploring a search space. The search performed by a metaheuristic approach should be “clever” enough to both intensively explore areas of the search space with high quality solutions, and to move to unexplored areas of the search space when necessary. The concepts for reaching these goals are nowadays called intensification and diversification. These terms stem from the TS field [Glover and Laguna 1997]. In other fields—such as the EC field—related concepts are often denoted by exploitation (related to intensification) and exploration (related to diversi-

fication). However, the terms exploitation and exploration have a somewhat more restricted meaning. In fact, the notions of exploitation and exploration often refer to rather short term strategies tied to randomness, whereas intensification and diversification refer to rather medium and long term strategies based on the usage of memory. As the various different ways of using memory become increasingly important in the whole field of metaheuristics, the terms intensification and diversification are more and more adopted and understood in their original meaning.

An implicit reference to the concept of “locality” is often introduced when intensification and diversification are involved. The notion of “area” (or “region”) of the search space and of “locality” can only be expressed in a fuzzy way, as they always depend on the characteristics of the search space as well as on the definition of metrics on the search space (distances between solutions).

The literature provides several high level descriptions of intensification and diversification. In the following, we cite some of them.

“Two highly important components of Tabu Search are intensification and diversification strategies. Intensification strategies are based on modifying choice rules to encourage move combinations and solution features historically found good. They may also initiate a return to attractive regions to search them more thoroughly. Since elite solutions must be recorded in order to examine their immediate neighborhoods, explicit memory is closely related to the implementation of intensification strategies. **The main difference between intensification and diversification is that during an intensification stage the search focuses on examining neighbors of elite solutions. [...]** **The diversification stage on the other hand encourages the search process to examine unvisited regions and to generate solutions that differ in various significant ways from those seen before.**” [Glover and Laguna 1997]

Later in the same book, Glover and Laguna write: “In some instances we may conceive of intensification as having the function of an intermediate term strategy, while diversification applies to considerations that emerge in the longer run.”

Furthermore, they write: “Strategic oscillation is closely linked to the origins of tabu search, and provides a means **to achieve an effective interplay between intensification and diversification.**”

“After a local minimizer is encountered, all points in its attraction basin lose any interest for optimization. The search should avoid wasting excessive computing time in a single basin and diversification should be activated. On the other hand, in the assumptions that neighbors have correlated cost function values, some effort should be spent in searching for better points located close to the most recently found local minimum point (intensification). **The two requirements are conflicting and finding a proper balance of diversification and intensification is a crucial issue in heuristics.**” [Battiti 1996].

“A metaheuristic will be successful on a given optimization problem if it can provide a balance between the exploitation of the accumulated search experience and the exploration of the search space to identify regions with high quality solutions in a problem specific, near optimal way.” [Stützle 1999b].

“Intensification is to search carefully and intensively around good solutions found in the past search. Diversification, on the contrary, is to guide the search to unvisited regions. These terminologies are usually used to explain the basic elements of Tabu Search, but these are essential to all the metaheuristic algorithms. In other words, various metaheuristic ideas should be understood from the viewpoint of these two concepts, and **metaheuristic algorithms should be designed so that intensification and diversifica-**

tion play balanced roles.” Yagiura and Ibaraki [2001].

“Holland frames adaption as a tension between exploration (the search for new, useful adaptations) and exploitation (the use and propagation of these adaptations). The tension comes about since any move toward exploration—testing previously unseen schemas or schemas whose instances seen so far have low fitness—takes away from the exploitation of tried and true schemas. In any system (e.g., a population of organisms) required to face environments with some degree of unpredictability, **an optimal balance between exploration and exploitation must be found.** The system has to keep trying out new possibilities (or else it could “over-adapt” and be inflexible in the face of novelty), but it also has to continually incorporate and use past experience as a guide for future behavior.”—M. Mitchell citing J.H. Holland in Mitchell [1998].

All these descriptions share the common view that there are two forces for which an appropriate balance has to be found. Sometimes these two forces were described as opposing forces. However, lately some researchers raised the question on how opposing intensification and diversification really are.

In 1998, Eiben and Schippers [1998] started a discussion about that in the field of Evolutionary Computation. They question the common opinion about EC algorithms, that they explore the search space by the genetic operators, while exploitation is achieved by selection. In their paper they give examples of operators that one cannot unambiguously label as being either intensification or diversification. So, for example, an operator using a local search component to improve individuals is not merely a mechanism of diversification because it also comprises a strong element of intensification (e.g., in Memetic Algorithms). Another example is the heuristically guided recombination of good quality solutions. If the use of the accumulated search experience is identified with intensification, then a recombination operator is not merely a means of




Fig. 18. The *I&D frame* provides a unified view on intensification and diversification in metaheuristics (**OG** = I&D components solely guided by the objective function, **NOG** = I&D components solely guided by one or more function other than the objective function, **R** = I&D components solely guided by randomness).

diversification, it also—as in the example above—has a strong intensification component.

Especially the TS literature advocates the view that intensification and diversification cannot be characterized as opposing forces. For example, in Glover and Laguna [1997], the authors write: “Similarly, as we have noted, intensification and diversification are not opposed notions, for the best form of each contains aspects of the other, along a spectrum of alternatives.”

Intensification and diversification can be considered as effects of algorithm components. In order to understand similarities and differences among metaheuristics, a framework may be helpful in providing a unified view on intensification and diversification components. We define an *I&D component* as any algorithmic or functional component that has an intensification and/or a diversification effect on the search process. Accordingly, examples of I&D components are genetic operators, perturbations of probability distributions, the use of tabu lists, or changes in the objective function. Thus, I&D components are operators, actions, or strategies of metaheuristic algorithms.

In contrast to the still widely spread view that there are components that have either an intensification or a diversification effect, there are many I&D components that have both. In I&D components that are commonly labelled as intensification, the intensification component is stronger than the diversification com-

ponent, and vice versa. To clarify this, we developed a framework to put I&D components of different metaheuristics into relation with each other. We called this framework—shown in Figure 18—the *I&D frame*.

We depict the space of all I&D components as a triangle with the three corners corresponding to three extreme examples of I&D components. The corner denoted by **OG** corresponds to I&D components solely guided by the objective function of the problem under consideration. An example of an I&D component which is located very close to the corner **OG** is the steepest descent choice rule in local search. The corner denoted by **NOG** covers all I&D components guided by one or more functions other than the objective function, again without using any random component. An example for such a component is a deterministic restart mechanism based on global frequency counts of solution components. The third corner, which is denoted by **R**, comprises all I&D components that are completely random. This means that they are not guided by anything. For example, a restart of an EC approach with random individuals is located in that corner. From the description of the corners, it becomes clear that corner **OG** corresponds to I&D components with a maximum intensification effect and a minimum diversification effect. On the other hand, corners **NOG**, **R** and the segment between the two corners correspond to I&D components with a maximum

diversification effect and a minimum intensification effect.¹⁹ All I&D components can be located somewhere on or inbetween the three corners, where the intensification effect is becoming smaller the further away a mechanism is located from **OG**. At the same time the diversification effect is growing. In step with this gradient is the use I&D components make of the objective function. The less an I&D component is using the objective function, the further away from corner **OG** it has to be located. There is also a second gradient to be found in this frame (which is shown in the second graphic of Figure 18). Corner **R** stands for complete randomness. The less randomness is involved in an I&D component, the further away from corner **R** it has to be located. Finally, a third gradient describes the influence of criteria different from the objective function, which generally stem from the exploitation of the search history that is in some form kept in the memory. In the following, we analyze some basic I&D components intrinsic to the basic versions of the metaheuristics with respect to the I&D frame.

5.2. Basic I&D Components of Metaheuristics

The I&D components occurring in metaheuristics can be divided in basic (or intrinsic) ones and strategic ones. The basic I&D components are the ones that are defined by the basic ideas of a metaheuristic. On the other side, strategic I&D components are composed of techniques and strategies the algorithm designer adds to the basic metaheuristic in order to improve the performance by incorporating medium- and long-term strategies. Many of these strategies were originally developed in the context of a specific metaheuristic. However, it becomes more and more apparent that many of these strategies can also be very useful when applied in other metaheuristics. In the following, we exemplarily choose some basic I&D com-

ponents that are inherent to a metaheuristic and explain them in the context of the I&D frame. With that, we show that most of the basic I&D components have an intensification character as well as a diversification character.

For many components and strategies of metaheuristics, it is obvious that they involve an intensification as well as a diversification component, because they make an explicit use of the objective function. For example, the basic idea of TS is a neighbor choice rule using one or more tabu lists. This I&D component has two effects on the search process. The restriction of the set of possible neighbors in every step has a diversifying effect on the search, whereas the choice of the best neighbor in the restricted set of neighbors (the best non-tabu move) has an intensifying effect on the search. The balance between these two effects can be varied by the length of the tabu list. Shorter tabu lists result in a lower influence of the diversifying effect, whereas longer tabu lists result in an overall higher influence of the diversifying effect. The location of this component in Figure 18 is on the segment between corner **OG** and **NOG**. The shorter the tabu lists, the closer is the location to corner **OG**, and vice versa.

Another example for such an I&D component is the probabilistic acceptance criterion in conjunction with the cooling schedule in SA. The acceptance criterion is guided by the objective function and it also involves a changing amount of randomness. The decrease of the temperature parameter drives the system from diversification to intensification eventually leading to a convergence²⁰ of the system. Therefore, this I&D component is located in the interior of the I&D space between corners **OG**, **NOG** and **R**.

A third example is the following one. Ant Colony Optimization provides an I&D component that manages the update of the pheromone values. This component has

¹⁹There is no quantitative difference between corners **NOG** and **R**. The difference is rather qualitative.

²⁰Here, we use the term *convergence* in the sense of getting stuck in the basin of attraction of a local minimum.

the effect of changing the probability distribution that is used to sample the search space. It is guided by the objective function (solution components found in better solutions than others are updated with a higher amount of pheromone) and it is also influenced by a function applying the pheromone evaporation. Therefore, this component is located on the line between corners **OG** and **NOG**. The effect of this mechanism is basically the intensification of the search, but there is also a diversifying component that depends on the greediness of the pheromone update (the less greedy or deterministic, the higher is the diversifying effect).

For other strategies and components of metaheuristics, it is not immediately obvious that they have both, an intensification and a diversification effect. An example is the random selection of a neighbor from the neighborhood of a current solution, as it is done for example in the kick-move mechanism of ILS. Initially one might think that there is no intensification involved and that this mechanism has a pure diversification effect caused by the use of randomness. However, for the following reason, this is not the case. Many strategies (such as the kick-move operator mentioned above) involve the explicit or implicit use of a neighborhood. A neighborhood structures the search space in the sense that it defines the topology of the so-called *fitness landscape* [Stadler 1995, 1996; Jones 1995a; Kauffman 1993], which can be visualized as a labelled graph. In this graph, nodes are solutions (labels indicate their objective function value) and arcs represent the neighborhood relation between states.²¹ A fitness landscape can be analyzed by means of statistical measures. One of the common measures is the *auto-correlation*, that provides information about how much the fitness will change when a move is made from one point to a neighboring

one. Different landscapes differ in their *ruggedness*. A landscape with small (average) fitness differences between neighboring points is called *smooth* and it will usually have just a few local optima. In contrast, a landscape with large (average) fitness differences is called *rugged* and it will be usually characterized by many local optima. Most of the neighborhoods used in metaheuristics provide some degree of smoothness that is higher than the one of a fitness landscape defined by a random neighborhood. This means that such a neighborhood is, in a sense, preselecting for every solution a set of neighbors for which the average fitness is not too different. Therefore, even when a solution is randomly selected from a set of neighbors, the objective function guidance is implicitly present. The consequence is that even for a random kick-move there is some degree of intensification involved, as far as a nonrandom neighborhood is considered.

For a mutation operator of an EC method that is doing a random change of a solution, it is neither immediately clear that it can have both, an intensification as well as a diversification effect. In the following, we assume a bit-string representation and a mutation operator that is characterized by flipping every bit of a solution with a certain probability. The implicit neighborhood used by this operator is the completely connected neighborhood. However, the neighbors have different probabilities to be selected. The ones that are (with respect to the Hamming distance) closer to the solution to which the operator is applied to, have a higher probability to be generated by the operator. With this observation, we can use the same argument as above in order to show an implicit use of objective function guidance. The balance between intensification and diversification is determined by the probability to flip each bit. The higher this probability, the higher the diversification effect of the operator. In contrast, the lower this probability, the higher the intensification effect of this operator.

On the other side, there are some strategies that are often labelled as intensification supposedly without having any

²¹The discussion of definitions and analysis of fitness landscapes is beyond the scope of this article. We forward the interested reader to Stadler [1995,1996]; Jones [1995a, 1995b]; Fonlupt et al. [1999], Hordijk [1996], Kauffman 1993], and Reeved [1999].

diversifying effect. One example is the selection operator in EC algorithms. However, nearly all selection operators involve some degree of randomness (e.g., proportionate selection, tournament selection) and are therefore located somewhere between corners **OG**, **NOG** and **R** of the I&D frame. This means that they also have a diversifying effect. The balance between intensification and diversification depends on the function that assigns the selection probabilities. If the differences between the selection probabilities are quite high, the intensification effect is higher, and similarly for the other extreme of having only small differences between the selection probabilities.

Even an operator like the neighbor choice rule of a steepest descent local search, which might be regarded as pure intensification, has a diversifying component in the sense that the search is “moving” in the search space with respect to a neighborhood. A neighborhood can be regarded as a function other than the objective function, making implicit use of the objective function. Therefore, a steepest descent local search is located between corners **OG** and **NOG**, and has both, a strong intensification effect but also a weak diversification character.

Based on these observations we conclude that probably most of the basic I&D components used in metaheuristics have both, an intensification and a diversification effect. However, the balance between intensification and diversification might be quite different for different I&D components. Table 1 attempts to summarize the basic I&D components that are inherent to the different metaheuristics.

5.3. Strategic Control of Intensification and Diversification

The right balance between intensification and diversification is needed to obtain an effective metaheuristic. Moreover, this balance should not be fixed or only changing into one direction (e.g., continuously increasing intensification). This balance should rather be dynamical. This issue is often treated in the literature, both implic-

Table 1. I&D-components intrinsic to the basic metaheuristics

| Metaheuristic | I&D component |
|---------------|--|
| SA | acceptance criterion + cooling schedule |
| TS | neighbor choice (tabu lists) aspiration criterion |
| EC | recombination mutation selection |
| ACO | pheromone update probabilistic construction |
| ILS | black box local search kick-move |
| VNS | acceptance criterion black box local search neighborhood choice shaking phase |
| GRASP | acceptance criterion black box local search restricted candidate list |
| GLS | penalty function |

itly and explicitly, when strategies to guide search algorithms are discussed.

The distinction between intensification and diversification is often interpreted with respect to the temporal horizon of the search. Short-term search strategies can be seen as the iterative application of tactics with a strong intensification character (for instance, the repeated application of greedy moves). When the horizon is enlarged, usually strategies referring to some sort of diversification come into play. Indeed, a general strategy usually proves its effectiveness especially in the long term.

The simplest strategy that coordinates the interplay of intensification and diversification and can achieve an oscillating balance between them is the restart mechanism: under certain circumstances (e.g., local optimum is reached, no improvements after a specific number of algorithm cycles, stagnation, no diversity) the algorithm is restarted. The goal is to achieve a sufficient coverage of the search space in the long run, thus the already visited regions should not be explored again. The computationally least expensive attempt to address this issue is a random restart. Every algorithm applying this naive diversification mechanism therefore incorporates an I&D component located in corner **R** of the I&D frame.

Usually, the most effective restart approaches make use of the search history. Examples for such restart strategies are the ones based on concepts such as *global frequency* and *global desirability*. The concept of *global frequency* is well known from TS applications. In this concept, the number of occurrences of solution components is counted during the run of the algorithm. These numbers, called the global frequency numbers, are then used for changing the heuristic constructive method, for example to generate a new population for restarting an EC method or the initial solution for restarting a trajectory method. Similarly, the concept of *global desirability* (which keeps for every solution component the objective function value of the best solution it had been a member of) can be used to restart algorithms with a bias toward good quality solutions. I&D components based on global frequency can be located in corner **NOG**, while global desirability-based components are located along the segment **NOG-OG**. Examples of the use of nonrandom restarts can be found also in population-based methods. In EC algorithms, the new population can be generated by applying constructive heuristics²² (line **R-OG**). In ACO, this goal is addressed by smoothing or resetting pheromone values [Stützle and Hoos 2000]. In the latter case, if the pheromone reset is also based on the search history, the action is located inside the I&D frame.

There are also strategies explicitly aimed at dynamically changing the balance between intensification and diversification during the search. A fairly simple strategy is used in SA, where an increase in diversification and simultaneous decrease in intensification can be achieved by “reheating” the system and then cooling it down again (which corresponds to increasing parameter T and decreasing it again according to some scheme). Such a cooling scheme is called *nonmonotonic cooling scheme* (e.g., see Lundy and Mees [1986] or Osman [1993]). Another example can be found in Ant Colony System

(ACS). This ACO algorithm uses an additional I&D component aimed at introducing diversification during the solution construction phase. While an ant is walking on the construction graph to construct a solution it reduces the pheromone values on the nodes/arcs of the construction graph that it visits. This has the effect to reduce for the other ants the probability of taking the same path. This additional pheromone update mechanism is called *step-by-step online pheromone update rule*. The interplay between this component and the other pheromone update rules (*online delayed pheromone update rule* and *online pheromone update rule*) leads to an oscillating balance between intensification and diversification.

Some more advanced strategies can be found in the literature. Often, they are described with respect to the particular metaheuristic in which they are applied. However, many of them are very general and can be easily adapted and reused also in a different context. A very effective example is *Strategic Oscillation* [Glover and Laguna 1997].²³ This strategy can be applied both to constructive methods and improvement algorithms. Actions are invoked with respect to a critical level (*oscillation boundary*), which usually corresponds to a steady state of the algorithm. Examples for steady states of an algorithm are local minima, completion of solution constructions, or the situation where no components can be added to a partial solution such that it can be completed to a feasible solution. The oscillation strategy is defined by a pattern indicating the way to approach the critical level, to cross it and to cross it again from the other side. This pattern defines the distance of moves from the boundary and the duration of phases (of intensification and diversification). Different patterns generate different strategies; moreover, they can also be adaptive and change depending on the current state and history of the search process. Other representative examples of general

²²See, for example, Freisleben and Merz [1996] and Grefenstette [1987].

²³Indeed, in Glover and Laguna [1997] and in the literature related to TS, many strategies are described and discussed.

strategies that dynamically coordinate intensification and diversification can be found in Battiti and Protasi [1997] and Blum [2002a, 2002b].

Furthermore, strategies are not restricted to single actions (e.g., variable assignments, moves), but may also guide the application of coordinated sequences of moves. Examples of such a strategy are given by so-called *ejection chain* procedures [Glover and Laguna 1997; Rego 1998, 2001]. These procedures provide a mechanism to perform compound moves, that is, compositions of different types of moves. For instance, in a problem defined over a graph (e.g., the VRP), it is possible to define two different moves: insertion and exchange of nodes; a compound move can thus be defined as the combination of an insertion and an exchange move. These procedures describe general strategies to combine the application of different neighborhood structures, thus they provide an example of a general diversification/intensification interplay. Further examples of strategies that can be interpreted as mechanisms to produce compositions of interlinked moves can also be found in the literature concerning the integration of metaheuristics and complete techniques [Caseau and Laburthe 1999; Shaw 1998].

In conclusion, we would like to stress again that most metaheuristic components have both an intensification and a diversification effect. The higher the objective function bias, the higher the intensification effect. In contrast, diversification is achieved by following guiding criteria other than the objective function and also by the use of randomness. With the introduction of the I&D frame, metaheuristics can be analyzed by their signature in the I&D frame. This can be a first step toward the systematic design of metaheuristics, combining I&D components of different origin.

5.4. Hybridization of Metaheuristics

We conclude our work by discussing a very promising research issue: the hybridization of metaheuristics. In fact, many of

the successful applications that we have cited in previous sections are hybridizations. In the following, we distinguish different forms of hybridization. The first one consists of including components from one metaheuristic into another one. The second form concerns systems that are sometimes labelled as cooperative search. They consist of various algorithms exchanging information in some way. The third form is the integration of approximate and systematic (or complete) methods. For a taxonomy of hybrid metaheuristics, see Talbi [2002].

Component Exchange Among Metaheuristics. One of the most popular ways of hybridization concerns the use of trajectory methods in population-based methods. Most of the successful applications of EC and ACO make use of local search procedures. The reason for that becomes apparent when analyzing the respective strengths of trajectory methods and population-based methods.

The power of population-based methods is certainly based on the concept of recombining solutions to obtain new ones. In EC algorithms and Scatter Search, explicit recombinations are implemented by one or more recombination operators. In ACO and EDAs, recombination is implicit, because new solutions are generated by using a distribution over the search space which is a function of earlier populations. This allows to make guided steps in the search space, which are usually “larger” than the steps done by trajectory methods. In other words, a solution resulting from a recombination in population-based methods is usually more “different” from the parents than, say, a predecessor solution to a successor solution (obtained by applying a move) in TS. We also have “big” steps in trajectory methods like ILS and VNS, but in these methods the steps are usually not guided (these steps are rather called “kick move” or “perturbation” indicating the lack of guidance). It is interesting to note, that in all population-based methods there are mechanisms in which good solutions found during the search influence the search process in the hope to find better solutions in-between those

solutions and current solutions. In Path Relinking, this idea is implemented in the most explicit way. The basic elements are guiding solutions (which are the good solutions found) and initiating solutions. New solutions are produced by applying moves to decrease the distance between the resulting solution and the guiding solution. In EC algorithms, this is often obtained by keeping the best (or a number of best) solution(s) found since the beginning of the respective run of the algorithm in the population. This is called a *steady state* evolution process. Scatter Search performs a steady state process by definition. In some ACO implementations (see, e.g., Stützle and Hoos [2000] and Blum [2002b]) a pheromone updating schedule is applied such that in a situation where the algorithm has nearly converged to a solution, only the best found solution since the start of the algorithm is used for updating the pheromone trails. This corresponds to “changing direction” and directing the search process toward a very good solution in the hope to find better ones on the way.

The strength of trajectory methods is rather to be found in the way they explore a promising region in the search space. As in those methods local search is the driving component, a promising area in the search space is searched in a more structured way than in population-based methods. In this way the danger of being close to good solutions but “missing” them is not as high as in population-based methods.

In summary, population-based methods are better in identifying promising areas in the search space, whereas trajectory methods are better in exploring promising areas in the search space. Thus, metaheuristic hybrids that in some way manage to combine the advantage of population-based methods with the strength of trajectory methods are often very successful.

Cooperative Search. A loose form of hybridization is provided by *cooperative search* [Hogg and Williams 1993; Hogg and Huberman 1993; Bachelet and Talbi 2000; Denzinger and Offerman 1999; Toulouse et al. 1999a, 1999b; Sondergeld and Voß 1999] which consists of a search

performed by possibly different algorithms that exchange information about states, models, entire sub-problems, solutions or other search space characteristics. Typically, cooperative search algorithms consist of the parallel execution of search algorithms with a varying level of communication. The algorithms can be different or they can be instances of the same algorithm working on different models or running with different parameters setting. The algorithms composing a cooperative search system can be all approximate, all complete, or a mix of approximate and complete approaches.

Cooperative search nowadays receives more attention, which is among other reasons due to the increasing research on parallel implementations of metaheuristics. The aim of research on parallelization of metaheuristics is twofold. First, metaheuristics should be redesigned to make them suitable for parallel implementation in order to exploit intrinsic parallelism. Second, an effective combination of metaheuristics has to be found, both to combine different characteristics and strengths and to design efficient communication mechanisms. Since the aim of this article is to provide an overview of the core ideas and strategies of metaheuristics, we refer to Crainic and Toulouse [2002a, 2002b] for a survey on the state-of-the-art in parallel metaheuristics.

Integrating Metaheuristics and Systematic Methods. Concluding this discussion on hybrid metaheuristics, we briefly overview the integration of metaheuristics and systematic search techniques. This approach has recently produced very effective algorithms especially when applied to real-world problems. Discussions on similarities, differences and possible integration of metaheuristics and systematic search can be found in Freuder et al. [1995], Ginsberg [1993], Harvey [1995], and Glover and Laguna [1997]. A very successful example of such an integration is the combination of metaheuristics and Constraint Programming (CP) [Focacci et al. 2002; Pesant and Gendreau 1996, 1999; De Backer et al. 2000]. CP enables to model a CO problem by means of

variables, domains²⁴ and constraints, which can be mathematical or symbolic (*global*). The latter ones involve a set of variables and describe subproblems, thus reducing the modelling complexity by encapsulating well defined parts of the problem into single constraints. Every constraint is associated to a *filtering* algorithm that deletes those values from a variable domain that do not contribute to feasible solutions. A CP system can be seen as the interaction of components (constraints) which communicate through shared variables. Constraints are activated as soon as the domain of any variable involved has been changed. Then, they perform a propagation phase, that is, they apply the filtering algorithm. This behavior stops as soon as there are no more values that can be removed from the domains or at least one domain is empty (i.e., no feasible solution exists). Since the complexity of the full constraint propagation is often exponential, the filtering is usually not complete. Therefore, at the end of the propagation phase, some domains may still contain unfeasible values. Hence, a search phase is started, such as Branch & Bound. A survey on the integration of metaheuristics and CP is provided by Focacci et al. [2002].

There are three main approaches for the integration of metaheuristics (especially trajectory methods) and systematic techniques (CP and tree search):

- A metaheuristic and a systematic method are sequentially applied (their execution can be also interleaved). For instance, the metaheuristic algorithm is run to produce some solutions that are then used as heuristic information by the systematic search. Vice-versa, the systematic algorithm can be run to generate a partial solution which will then be completed by the metaheuristic.
- Metaheuristics use CP and/or tree search to efficiently explore the neighborhood, instead of simply enumerating the neighbors or randomly sampling the neighborhood.

—The third possibility consists of introducing concepts or strategies from either class of algorithms into the other. For example, the concepts of tabu list and aspiration criteria—defined in Tabu Search—can be used to manage the list of open nodes (i.e., the ones whose child nodes are not yet explored) in a tree search algorithm.

The first approach can be seen as an instance of cooperative search and it represents a rather loose integration.

The second approach combines the advantages of a fast search space exploration by means of a metaheuristic with the efficient neighborhood exploration performed by a systematic method. A prominent example of such a kind of integration is *Large Neighborhood Search* and related approaches [Shaw 1998; Caseau and Laburthe 1999]. These approaches are effective mainly when the neighborhood to explore is very large. Moreover, many real-world problems have additional constraints (called *side constraints*) which might make them unsuitable for usual neighborhood exploration performed by metaheuristics, since they usually just sample the neighborhood or enumerate its solutions. For instance, time window constraints often reduce the number of feasible solutions in a neighborhood, which might make a local search inefficient. Thus, domain filtering techniques can effectively support neighborhood exploration. In fact, for such a kind of neighborhoods, both sampling and enumeration are usually inefficient. More examples can be found in Pesant and Gendreau [1996, 1999] and Focacci et al. [2002].

The third approach preserves the search space exploration based on a systematic search (such as tree search), but sacrifices the exhaustive nature of the search [Ginsberg 1993; Harvey 1995; Harvey and Ginsberg 1995; Milano and Roli 2002]. The hybridization is usually achieved by integrating concepts developed for metaheuristics (e.g., probabilistic choices, aspiration criteria, heuristic construction) into tree search methods. A typical application of this integration is the use

²⁴We restrict the discussion to finite domains.

of a probabilistic backtracking, instead of a—deterministic—chronological backtracking. For instance, an extreme case is the random choice of a backtracking move. The list of possible backtracking moves can also be sorted by means of a dynamic heuristic or a *sampling* of the leaves of the search tree. This sampling can be performed by a metaheuristic: the result of each possible backtracking move is chosen as the starting point for producing a complete solution by a metaheuristic (more than one solution can be generated from each partial solution). Then, the quality of these complete solutions is used as a score for—probabilistically—selecting a backtracking move. Another prominent example is the introduction of randomization in systematic techniques, as described in Gomes et al. [2000]. Many examples of this approach can be found in Focacci et al. [2002], Jussien and Lhomme [2002], Schaerf [1997], Dell’Amico and Lodi [2002], Prestwich [2002], and Della Croce and T’kindt [2003].

6. CONCLUSIONS

In this work, we have presented and compared the current most important metaheuristic methods. In Sections 3 and 4, we have outlined the basic metaheuristics as they are described in the literature. In Section 5, we then proposed a conceptual comparison of the different metaheuristics based on the way they implement the two main concepts for guiding the search process: Intensification and diversification. This comparison is founded on the *I&D frame*, where algorithmic components can be characterized by the criteria they depend upon (objective function, guiding functions and randomization) and their effect on the search process. Although metaheuristics are different in the sense that some of them are population-based (EC, ACO), and others are trajectory methods (SA, TS, ILS, VNS, GRASP), and although they are based on different philosophies, the mechanisms to efficiently explore a search space are all based on intensification and diversification. Nevertheless, it is possible to identify “sub-

tasks” in the search process where some metaheuristics perform better than others. This has to be examined more closely in the future in order to be able to produce hybrid metaheuristics performing considerably better than their “pure” parents. In fact we can find this phenomenon in many facets of life, not just in the world of algorithms. Mixing and hybridizing is often better than purity.

ACKNOWLEDGMENTS

We would like to thank Marco Dorigo, Joshua Knowles, Andrea Lodi, Michela Milano, Michael Sampels, and Thomas Stützle for suggestions and useful discussions. We also would like to thank the anonymous referees for many useful suggestions and comments.

REFERENCES

- AARTS, E. H. L., KORST, J. H. M., AND LAARHOVEN, P. J. M. v. 1997. Simulated annealing. In *Local Search in Combinatorial Optimization*, E. H. L. Aarts and J. K. Lenstra, Eds. Wiley-Interscience, Chichester, England, 91–120.
- AARTS, E. H. L. AND LENSTRA, J. K., Eds. 1997. *Local Search in Combinatorial Optimization*. Wiley, Chichester, UK.
- AIEX, R. M., BINATO, S., AND RESENDE, M. G. C. 2003. Parallel GRASP with path-relinking for job shop scheduling. *Paral. Comput.* To appear.
- BACHELET, V. AND TALBI, E. 2000. Cosearch: A co-evolutionary metaheuristics. In *Proceedings of Congress on Evolutionary Computation—CEC2000*. 1550–1557.
- BÄCK, T. 1996. *Evolutionary Algorithms in Theory and Practice*. Oxford University Press, New York.
- BÄCK, T., FOGEL, D. B., AND MACHALEWICZ, Z., Eds. 1997. *Handbook of Evolutionary Computation*. Institute of Physics Publishing Ltd, Bristol, UK.
- BALUJA, S. 1994. Population-based incremental learning: A method for integrating genetic search based function optimization and competitive learning. Tech. Rep. No. CMU-CS-94-163, Carnegie Mellon University, Pittsburgh, Pa.
- BALUJA, S. AND CARUANA, R. 1995. Removing the genetics from the standard genetic algorithm. In *The International Conference on Machine Learning 1995*, A. Prieditis and S. Russel, Eds. Morgan-Kaufmann Publishers, San Mateo, Calif., 38–46.
- BAR-YAM, Y. 1997. *Dynamics of Complex Systems*. Studies in nonlinearity. Addison-Wesley, Reading, Mass.
- BATTITI, R. 1996. Reactive search: Toward self-tuning heuristics. In *Modern Heuristic Search*

- Methods*, V. J. Rayward-Smith, I. H. Osman, C. R. Reeves, and G. D. Smith, Eds. Wiley, Chichester, UK, 61–83.
- BATTITI, R. AND PROTASI, M. 1997. Reactive Search, A history-base heuristic for MAX-SAT. *ACM J. Exper. Algor.* 2, Article 2.
- BATTITI, R. AND TECCHIOLLI, G. 1994. The reactive tabu search. *ORSA J. Comput.* 6, 2, 126–140.
- BINATO, S., HERY, W. J., LOEWENSTERN, D., AND RESENDE, M. G. C. 2001. A greedy randomized adaptive search procedure for job shop scheduling. In *Essays and Surveys on Metaheuristics*, P. Hansen and C. C. Ribeiro, Eds. Kluwer Academic Publishers.
- BLUM, C. 2002a. ACO applied to group shop scheduling: A case study on intensification and diversification. In *Proceedings of ANTS 2002—Third International Workshop on Ant Algorithms*, M. Dorigo, G. Di Caro, and M. Sampels, Eds. Lecture Notes in Computer Science, vol. 2463. Springer Verlag, Berlin, Germany, 14–27.
- BLUM, C. 2002b. Ant colony optimization for the edge-weighted k -cardinality tree problem. In *GECCO 2002: Proceedings of the Genetic and Evolutionary Computation Conference*, W. B. Langdon, E. Cantú-Paz, K. Mathias, R. Roy, D. Davis, R. Poli, K. Balakrishnan, V. Honavar, G. Rudolph, J. Wegener, L. Bull, M. A. Potter, A. C. Schultz, J. F. Miller, E. Burke, and N. Jonoska, Eds. Morgan-Kaufmann, New York, 27–34.
- BLUM, C., ROLI, A., AND DORIGO, M. 2001. HC-ACO: The hyper-cube framework for ant colony optimization. In *Proceedings of MIC'2001—Meta-heuristics International Conference*. Vol. 2. Porto, Portugal, 399–403.
- CALÉGARY, P., CORAY, G., HERTZ, A., KOBLER, D., AND KUONEN, P. 1999. A taxonomy of evolutionary algorithms in combinatorial optimization. *J. Heuristics* 5, 145–158.
- CAMPOS, V., GLOVER, F., LAGUNA, M., AND MARTÍ, R. 2001. An experimental evaluation of a scatter search for the linear ordering problem. *J. Global Opt.* 21, 397–414.
- CASEAU, Y. AND LABURTHER, F. 1999. Effective forget-and-extend heuristics for scheduling problems. In *Proceedings of CP-AI-OR'02—Fourth Int. Workshop on Integration of AI and OR techniques in Constraint Programming for Combinatorial Optimization Problems*. Ferrara (Italy). Also available at: www.deis.unibo.it/Events/Deis/Workshops/Proceedings.html.
- CERNY, V. 1985. A thermodynamical approach to the travelling salesman problem: An efficient simulation algorithm. *J. Optim. Theory Appl.* 45, 41–51.
- CHARDAIRE, P., LUTTON, J. L., AND SUTTER, A. 1995. Thermostatistical persistency: A powerful improving concept for simulated annealing algorithms. *Europ. J. Oper. Res.* 86, 565–579.
- COELLO COELLO, C. A. 2000. An updated survey of GA-based multiobjective optimization techniques. *ACM Comput. Surv.* 32, 2, 109–143.
- CONNOLLY, D. T. 1990. An improved annealing scheme for the QAP. *Europ. J. Oper. Res.* 46, 93–100.
- CRAINIC, T. G. AND TOULOUSE, M. 2002a. Introduction to the special issue on parallel metaheuristics. *J. Heuristics* 8, 3, 247–249.
- CRAINIC, T. G. AND TOULOUSE, M. 2002b. Parallel strategies for meta-heuristics. In *Handbook of Metaheuristics*, F. Glover and G. Kochenberger, Eds. International Series in Operations Research & Management Science, vol. 57. Kluwer Academic Publishers, Norwell, MA.
- DE BACKER, B., FURNON, V., AND SHAW, P. 2000. Solving Vehicle Routing Problems Using Constraint Programming and Metaheuristics. *J. Heuristics* 6, 501–523.
- DE BONET, J. S., ISBELL JR., C. L., AND VIOLA, P. 1997. MIMIC: Finding optima by estimating probability densities. In *Proceedings of the 1997 Conference on Advances in Neural Information Processing Systems (NIPS'97)*, M. C. Mozer, M. I. Jordan, and T. Petsche, Eds. MIT Press, Cambridge, MA, 424–431.
- DEJONG, K. A. 1975. An analysis of the behavior of a class of genetic adaptive systems. Ph.D. thesis, University of Michigan, Ann Arbor, MI. Dissertation Abstracts International 36(10), 5140B, University Microfilms Number 76-9381.
- DELLA CROCE, F. AND T'KINDT, V. 2003. A Recovering Beam Search algorithm for the one machine dynamic total completion time scheduling problem. *J. Oper. Res. Soc.* To appear.
- DELL'AMICO, M. AND LODI, A. 2002. On the integration of metaheuristic strategies in constraint programming. In *Adaptive Memory and Evolution: Tabu Search and Scatter Search*, C. Rego and B. Alidaee, Eds. Kluwer Academic Publishers, Boston, MA.
- DELL'AMICO, M., LODI, A., AND MAFFIOLI, F. 1999. Solution of the cumulative assignment problem with a well-structured tabu search method. *J. Heuristics* 5, 123–143.
- DEN BESTEN, M. L., STÜTZLE, T., AND DORIGO, M. 2001. Design of iterated local search algorithms: An example application to the single machine total weighted tardiness problem. In *Proceedings of EvoStim'01*. Lecture Notes in Computer Science. Springer, 441–452.
- DENEUBOURG, J.-L., ARON, S., GOSS, S., AND PASTEELS, J.-M. 1990. The self-organizing exploratory pattern of the argentine ant. *J. Insect Behav.* 3, 159–168.
- DENZINGER, J. AND OFFERMAN, T. 1999. On cooperation between evolutionary algorithms and other search paradigms. In *Proceedings of Congress on Evolutionary Computation—CEC'1999*. 2317–2324.
- DEVANEY, R. L. 1989. *An introduction to chaotic dynamical systems*, second ed. Addison-Wesley, Reading, Mass.

- DI CARO, G. AND DORIGO, M. 1998. AntNet: Distributed stigmergetic control for communication networks. *J. Artif. Int. Res.* 9, 317–365.
- DORIGO, M. 1992. Optimization, learning and natural algorithms (in italian). Ph.D. thesis, DEI, Politecnico di Milano, Italy. pp. 140.
- DORIGO, M. AND DI CARO, G. 1999. The ant colony optimization meta-heuristic. In *New Ideas in Optimization*, D. Corne, M. Dorigo, and F. Glover, Eds. McGraw-Hill, 11–32.
- DORIGO, M., DI CARO, G., AND GAMBARDELLA, L. M. 1999. Ant algorithms for discrete optimization. *Art. Life* 5, 2, 137–172.
- DORIGO, M. AND GAMBARDELLA, L. M. 1997. Ant colony system: A cooperative learning approach to the travelling salesman problem. *IEEE Trans. Evolution. Comput.* 1, 1 (Apr.), 53–66.
- DORIGO, M., MANIEZZO, V., AND COLORNI, A. 1996. Ant system: Optimization by a colony of cooperating agents. *IEEE Trans. Syst. Man Cybernet.—Part B* 26, 1, 29–41.
- DORIGO, M. AND STÜTZLE, T. 2002. The ant colony optimization metaheuristic: Algorithms, applications and advances. In *Handbook of Metaheuristics*, F. Glover and G. Kochenberger, Eds. International Series in Operations Research & Management Science, vol. 57. Kluwer Academic Publishers, Norwell, MA, 251–285.
- DORIGO, M. AND STÜTZLE, T. 2003. *Ant Colony Optimization*. MIT Press, Boston, MA. To appear.
- DUECK, G. 1993. New Optimization Heuristics. *J. Comput. Phy.* 104, 86–92.
- DUECK, G. AND SCHEUER, T. 1990. Threshold accepting: A general purpose optimization algorithm appearing superior to simulated annealing. *J. Comput. Phy.* 90, 161–175.
- EIBEN, A. E., RAUÉ, P.-E., AND RUTTKAY, Z. 1994. Genetic algorithms with multi-parent recombination. In *Proceedings of the 3rd Conference on Parallel Problem Solving from Nature*, Y. Davidor, H.-P. Schwefel, and R. Manner, Eds. Lecture Notes in Computer Science, vol. 866. Springer, Berlin, 78–87.
- EIBEN, A. E. AND RUTTKAY, Z. 1997. Constraint satisfaction problems. In *Handbook of Evolutionary Computation*, T. Bäck, D. Fogel, and M. Michalewicz, Eds. Institute of Physics Publishing Ltd, Bristol, UK.
- EIBEN, A. E. AND SCHIPPERS, C. A. 1998. On evolutionary exploration and exploitation. *Fund. Inf.* 35, 1–16.
- FELLER, W. 1968. *An Introduction to Probability Theory and Its Applications*. Wiley, New York.
- FEO, T. A. AND RESENDE, M. G. C. 1995. Greedy randomized adaptive search procedures. *J. Global Optim.* 6, 109–133.
- FESTA, P. AND RESENDE, M. G. C. 2002. GRASP: An annotated bibliography. In *Essays and Surveys on Metaheuristics*, C. C. Ribeiro and P. Hansen, Eds. Kluwer Academic Publishers, 325–367.
- FINK, A. AND VOß, S. 1999. Generic metaheuristics application to industrial engineering problems. *Comput. Indust. Eng.* 37, 281–284.
- FLEISCHER, M. 1995. Simulated Annealing: past, present and future. In *Proceedings of the 1995 Winter Simulation Conference*, C. Alexopoulos, K. Kang, W. Lilegdon, and G. Goldsman, Eds. 155–161.
- FOCACCI, F., LABURTHE, F., AND LODI, A. 2002. Local search and constraint programming. In *Handbook of Metaheuristics*, F. Glover and G. Kochenberger, Eds. International Series in Operations Research & Management Science, vol. 57. Kluwer Academic Publishers, Norwell, MA.
- FOGEL, D. B. 1994. An introduction to simulated evolutionary optimization. *IEEE Trans. Neural Netw.* 5, 1 (Jan.), 3–14.
- FOGEL, G. B., PORTO, V. W., WEEKES, D. G., FOGEL, D. B., GRIFFEY, R. H., MCNEIL, J. A., LESNIK, E., ECKER, D. J., AND SAMPATH, R. 2002. Discovery of RNA structural elements using evolutionary computation. *Nucleic Acids Res.* 30, 23, 5310–5317.
- FOGEL, L. J. 1962. Toward inductive inference automata. In *Proceedings of the International Federation for Information Processing Congress*. Munich, 395–399.
- FOGEL, L. J., OWENS, A. J., AND WALSH, M. J. 1966. *Artificial Intelligence through Simulated Evolution*. Wiley, New York.
- FONLUPT, C., ROBILLIARD, D., PREUX, P., AND TALBI, E. 1999. Fitness landscapes and performance of meta-heuristics. In *Meta-heuristics: advances and trends in local search paradigms for optimization*, S. Voß, S. Martello, I. Osman, and C. Roucairol, Eds. Kluwer Academic.
- FREISLEBEN, B. AND MERZ, P. 1996. A genetic local search algorithm for solving symmetric and asymmetric traveling salesman problems. In *International Conference on Evolutionary Computation*. 616–621.
- FREUDER, E. C., DECHTER, R., GINSBERG, M. L., SELMAN, B., AND TSANG, E. P. K. 1995. Systematic versus stochastic constraint satisfaction. In *Proceedings of the 14th International Joint Conference on Artificial Intelligence, IJCAI 1995*. Vol. 2. Morgan-Kaufmann, 2027–2032.
- GAMBARDELLA, L. M. AND DORIGO, M. 2000. Ant colony system hybridized with a new local search for the sequential ordering problem. *INFORMS J. Comput.* 12, 3, 237–255.
- GAREY, M. R. AND JOHNSON, D. S. 1979. *Computers and Intractability; A Guide to the Theory of NP-Completeness*. W.H. Freeman.
- GENDREAU, M., LAPORTE, G., AND POTVIN, J.-Y. 2001. Metaheuristics for the vehicle routing problem. In *The Vehicle Routing Problem*, P. Toth and D. Vigo, Eds. SIAM Series on Discrete Mathematics and Applications, vol. 9. 129–154.
- GINSBERG, M. L. 1993. Dynamic backtracking. *J. Artif. Int. Res.* 1, 25–46.
- GLOVER, F. 1977. Heuristics for integer programming using surrogate constraints. *Dec. Sci.* 8, 156–166.

- GLOVER, F. 1986. Future paths for integer programming and links to artificial intelligence. *Comput. Oper. Res.* 13, 533–549.
- GLOVER, F. 1990. Tabu search Part II. *ORSA J. Comput.* 2, 1, 4–32.
- GLOVER, F. 1999. Scatter search and path relinking. In *New Ideas in Optimization*, D. Corne, M. Dorigo, and F. Glover, Eds. Advanced topics in computer science series. McGraw-Hill.
- GLOVER, F. AND LAGUNA, M. 1997. *Tabu Search*. Kluwer Academic Publishers.
- GLOVER, F., LAGUNA, M., AND MARTÍ, R. 2000. Fundamentals of scatter search and path relinking. *Control*, 29, 3, 653–684.
- GLOVER, F., LAGUNA, M., AND MARTÍ, R. 2002. Scatter search and path relinking: Advances and applications. In *Handbook of Metaheuristics*, F. Glover and G. Kochenberger, Eds. International Series in Operations Research & Management Science, vol. 57. Kluwer Academic Publishers, Norwell, MA.
- GOLDBERG, D. E. 1989. *Genetic Algorithms in Search, Optimization and Machine Learning*. Addison Wesley, Reading, MA.
- GOLDBERG, D. E., DEB, K., AND KORB, B. 1991. Don't worry, be messy. In *Proceedings of the 4th International Conference on Genetic Algorithms*. Morgan-Kaufmann, La Jolla, CA.
- GOLDBERG, D. E. AND RICHARDSON, J. 1987. Genetic algorithms with sharing for multimodal function optimization. In *Genetic Algorithms and their Applications*, J. J. Grefenstette, Ed. Lawrence Erlbaum Associates, Hillsdale, NJ, 41–49.
- GOMES, C. P., SELMAN, B., CRATO, N., AND KAUTZ, H. 2000. Heavy-Tailed phenomena in Satisfiability and Constraint Satisfaction Problems. *J. Automat. Reason.* 24, 67–100.
- GREFENSTETTE, J. J. 1987. Incorporating problem specific knowledge into genetic algorithms. In *Genetic Algorithms and Simulated Annealing*, L. Davis, Ed. Morgan-Kaufmann, 42–60.
- GREFENSTETTE, J. J. 1990. A user's guide to GENESIS 5.0. Tech. rep., Navy Centre for Applied Research in Artificial Intelligence, Washington, D.C.
- HANSEN, P. 1986. The steepest ascent mildest descent heuristic for combinatorial programming. In *Congress on Numerical Methods in Combinatorial Optimization*. Capri, Italy.
- HANSEN, P. AND MLADENOVIĆ, N. 1997. Variable neighborhood search for the p -median. *Loc. Sci.* 5, 207–226.
- HANSEN, P. AND MLADENOVIĆ, N. 1999. An introduction to variable neighborhood search. In *Metaheuristics: Advances and trends in local search paradigms for optimization*, S. Voß, S. Martello, I. Osman, and C. Roucairol, Eds. Kluwer Academic Publishers, Chapter 30, 433–458.
- HANSEN, P. AND MLADENOVIĆ, N. 2001. Variable neighborhood search: Principles and applications. *Europ. J. Oper. Res.* 130, 449–467.
- HARIK, G. 1999. Linkage learning via probabilistic modeling in the ECGA. Tech. Rep. No. 99010, IlliGAL, University of Illinois.
- HARVEY, W. D. 1995. Nonsystematic backtrack-search. Ph.D. thesis, CIRL, University of Oregon.
- HARVEY, W. D. AND GINSBERG, M. L. 1995. Limited discrepancy search. In *Proceedings of the 14th International Joint Conference on Artificial Intelligence, IJCAI 1995* (Montréal, Qué, Canada). C. S. Mellish, Ed. Vol. 1. Morgan-Kaufmann, 607–615.
- HERTZ, A. AND KOBLER, D. 2000. A framework for the description of evolutionary algorithms. *Europ. J. Oper. Res.* 126, 1–12.
- HOGG, T. AND HUBERMAN, A. 1993. Better than the best: The power of cooperation. In *SFI 1992 Lectures in Complex Systems*. Addison-Wesley, 163–184.
- HOGG, T. AND WILLIAMS, C. 1993. Solving the really hard problems with cooperative search. In *Proceedings of AAAI93*. AAAI Press, 213–235.
- HOLLAND, J. H. 1975. *Adaption in natural and artificial systems*. The University of Michigan Press, Ann Harbor, MI.
- HORDIJK, W. 1996. A measure of landscapes. *Evolut. Comput.* 4, 4, 335–360.
- INGBER, L. 1996. Adaptive simulated annealing (ASA): Lessons learned. *Cont. Cybernet.—Special Issue on Simulated Annealing Applied to Combinatorial Optimization* 25, 1, 33–54.
- JOHNSON, D. S. AND MCGEOCH, L. A. 1997. The traveling salesman problem: a case study. In *Local Search in Combinatorial Optimization*, E. Aarts and J. Lenstra, Eds. Wiley, New York, 215–310.
- JONES, T. 1995a. Evolutionary algorithms, fitness landscapes and search. Ph.D. thesis, Univ. of New Mexico, Albuquerque, NM.
- JONES, T. 1995b. One operator, one landscape. Santa Fe Institute Tech. Rep. 95-02-025, Santa Fe Institute.
- JOSLIN, D. E. AND CLEMENTS, D. P. 1999. “Squeaky Wheel” Optimization. *J. Artif. Int. Res.* 10, 353–373.
- JUSSIEN, N. AND LHOMME, O. 2002. Local search with constraint propagation and conflict-based heuristics. *Artif. Int.* 139, 21–45.
- KAUFFMAN, S. A. 1993. *The Origins of Order: Self-Organization and Selection in Evolution*. Oxford University Press.
- KILBY, P., PROSSER, P., AND SHAW, P. 1999. Guided Local Search for the Vehicle Routing Problem with time windows. In *Meta-heuristics: Advances and trends in local search paradigms for optimization*, S. Voß, S. Martello, I. Osman, and C. Roucairol, Eds. Kluwer Academic, 473–486.
- KIRKPATRICK, S., GELATT, C. D., AND VECCHI, M. P. 1983. Optimization by simulated annealing. *Science*, 13 May 1983 220, 4598, 671–680.
- LAGUNA, M., LOURENÇO, H., AND MARTÍ, R. 2000. Assigning Proctors to Exams with Scatter Search.

- In *Computing Tools for Modeling, Optimization and Simulation: Interfaces in Computer Science and Operations Research*, M. Laguna and J. L. González-Velarde, Eds. Kluwer Academic Publishers, Boston, MA, 215–227.
- LAGUNA, M. AND MARTÍ, R. 1999. GRASP and path relinking for 2-layer straight line crossing minimization. *INFORMS J. Comput.* 11, 1, 44–52.
- LAGUNA, M., MARTÍ, R., AND CAMPOS, V. 1999. Intensification and diversification with elite tabu search solutions for the linear ordering problem. *Comput. Oper. Res.* 26, 1217–1230.
- LARRAÑAGA, P. AND LOZANO, J. A., Eds. 2002. *Estimation of Distribution Algorithms: A New Tool for Evolutionary Computation*. Kluwer Academic Publishers, Boston, MA.
- LOURENÇO, H. R., MARTIN, O., AND STÜTZLE, T. 2001. A beginner's introduction to Iterated Local Search. In *Proceedings of MIC'2001—Meta-heuristics International Conference*. Vol. 1. Porto—Portugal, 1–6.
- LOURENÇO, H. R., MARTIN, O., AND STÜTZLE, T. 2002. Iterated local search. In *Handbook of Meta-heuristics*, F. Glover and G. Kochenberger, Eds. International Series in Operations Research & Management Science, vol. 57. Kluwer Academic Publishers, Norwell, MA, 321–353.
- LUNDY, M. AND MEES, A. 1986. Convergence of an annealing algorithm. *Math. Prog.* 34, 1, 111–124.
- MARTIN, O. AND OTTO, S. W. 1996. Combining simulated annealing with local search heuristics. *Ann. Oper. Res.* 63, 57–75.
- MARTIN, O., OTTO, S. W., AND FELTEN, E. W. 1991. Large-step Markov chains for the traveling salesman problem. *Complex Syst.* 5, 3, 299–326.
- MERKLE, D., MIDDENDORF, M., AND SCHMECK, H. 2002. Ant colony optimization for resource-constrained project scheduling. *IEEE Trans. Evolut. Comput.* 6, 4, 333–346.
- METAHEURISTICS NETWORK WEBSITE 2000. <http://www.metaheuristics.net/>. Visited in January 2003.
- MEULEAU, N. AND DORIGO, M. 2002. Ant colony optimization and stochastic gradient descent. *Artif. Life* 8, 2, 103–121.
- MICHALEWICZ, Z. AND MICHALEWICZ, M. 1997. Evolutionary computation techniques and their applications. In *Proceedings of the IEEE International Conference on Intelligent Processing Systems*, (Beijing, China). Institute of Electrical & Electronics Engineers, Incorporated, 14–24.
- MILANO, M. AND ROLI, A. 2002. On the relation between complete and incomplete search: An informal discussion. In *Proceedings of CP-AI-OR'02—Fourth Int. Workshop on Integration of AI and OR techniques in Constraint Programming for Combinatorial Optimization Problems* (Le Croisic, France). 237–250.
- MILLS, P. AND TSANG, E. 2000. Guided local search for solving SAT and weighted MAX-SAT Problems. In *SAT2000*, I. Gent, H. van Maaren, and T. Walsh, Eds. IOS Press, 89–106.
- MITCHELL, M. 1998. *An Introduction to Genetic Algorithms*. MIT press, Cambridge, MA.
- MLADENVIĆ, N. AND UROŠEVIĆ, D. 2001. Variable neighborhood search for the k -cardinality tree. In *Proceedings of MIC'2001—Meta-heuristics International Conference*. Vol. 2. Porto, Portugal, 743–747.
- MOSCATO, P. 1989. On evolution, search, optimization, genetic algorithms and martial arts: Toward memetic algorithms. Tech. Rep. Caltech Concurrent Computation Program 826, California Institute of Technology, Pasadena, Calif.
- MOSCATO, P. 1999. Memetic algorithms: A short introduction. In *New Ideas in Optimization*, F. G. D. Corne and M. Dorigo, Eds. McGraw-Hill.
- MÜHLENBEIN, H. 1991. Evolution in time and space—The parallel genetic algorithm. In *Foundations of Genetic Algorithms*, G. J. E. Rawlins, Ed. Morgan-Kaufmann, San Mateo, Calif.
- MÜHLENBEIN, H. AND PAAß, G. 1996. From recombination of genes to the estimation of distributions. In *Proceedings of the 4th Conference on Parallel Problem Solving from Nature—PPSN IV*, H.-M. Voigt, W. Ebeling, I. Rechenberg, and H.-P. Schwefel, Eds. Lecture Notes in Computer Science, vol. 1411. Springer, Berlin, 178–187.
- MÜHLENBEIN, H. AND VOIGT, H.-M. 1995. Gene pool recombination in genetic algorithms. In *Proc. of the Metaheuristics Conference*, I. H. Osman and J. P. Kelly, Eds. Kluwer Academic Publishers, Norwell, USA.
- NEMHAUSER, G. L. AND WOLSEY, A. L. 1988. *Integer and Combinatorial Optimization*. Wiley, New York.
- NOWICKI, E. AND SMUTNICKI, C. 1996. A fast taboo search algorithm for the job-shop problem. *Manage. Sci.* 42, 2, 797–813.
- OSMAN, I. H. 1993. Metastrategy simulated annealing and tabu search algorithms for the vehicle routing problem. *Ann. Oper. Res.* 41, 421–451.
- OSMAN, I. H. AND LAPORTE, G. 1996. Metaheuristics: A bibliography. *Ann. Oper. Res.* 63, 513–623.
- PAPADIMITRIOU, C. H. AND STEIGLITZ, K. 1982. *Combinatorial Optimization—Algorithms and Complexity*. Dover Publications, Inc., New York.
- PELIKAN, M., GOLDBERG, D. E., AND CANTÚ-PAZ, E. 1999a. BOA: The Bayesian optimization algorithm. In *Proceedings of the Genetic and Evolutionary Computation Conference GECCO-99* (Orlando, Fla.). W. Banzhaf, J. Daida, A. E. Eiben, M. H. Garzon, V. Honavar, M. Jakiela, and R. E. Smith, Eds. Vol. 1. Morgan-Kaufmann Publishers, San Francisco, CA, 525–532.
- PELIKAN, M., GOLDBERG, D. E., AND LOBO, F. 1999b. A survey of optimization by building and using probabilistic models. Tech. Rep. No. 99018, ILLIGAL, University of Illinois.

- PESANT, G. AND GENDREAU, M. 1996. A view of local search in Constraint Programming. In *Principles and Practice of Constraint Programming—CP'96*. Lecture Notes in Computer Science, vol. 1118. Springer-Verlag, 353–366.
- PESANT, G. AND GENDREAU, M. 1999. A constraint programming framework for local search methods. *J. Heuristics* 5, 255–279.
- PITSOULIS, L. S. AND RESENDE, M. G. C. 2002. Greedy randomized adaptive search procedure. In *Handbook of Applied Optimization*, P. Pardalos and M. Resende, Eds. Oxford University Press, 168–183.
- PRAIS, M. AND RIBEIRO, C. C. 2000. Reactive GRASP: An application to a matrix decomposition problem in TDMA traffic assignment. *INFORMS J. Comput.* 12, 164–176.
- PRESTWICH, S. 2002. Combining the scalability of local search with the pruning techniques of systematic search. *Ann. Oper. Res.* 115, 51–72.
- RADCLIFFE, N. J. 1991. Forma Analysis and Random Respectful Recombination. In *Proceedings of the Fourth International Conference on Genetic Algorithms, ICGA 1991*. Morgan-Kaufmann, San Mateo, Calif., 222–229.
- RAYWARD-SMITH, V. J. 1994. A unified approach to tabu search, simulated annealing and genetic algorithms. In *Applications of Modern Heuristics*, V. J. Rayward-Smith, Ed. Alfred Waller Limited, Publishers.
- RECHENBERG, I. 1973. *Evolutionsstrategie: Optimierung technischer Systeme nach Prinzipien der biologischen Evolution*. Frommann-Holzboog.
- REEVES, C. R., Ed. 1993. *Modern Heuristic Techniques for Combinatorial Problems*. Blackwell Scientific Publishing, Oxford, England.
- REEVES, C. R. 1999. Landscapes, operators and heuristic search. *Ann. Oper. Res.* 86, 473–490.
- REEVES, C. R. AND ROWE, J. E. 2002. *Genetic Algorithms: Principles and Perspectives. A Guide to GA Theory*. Kluwer Academic Publishers, Boston (USA).
- REGO, C. 1998. Relaxed Tours and Path Ejections for the Traveling Salesman Problem. *Europ. J. Oper. Res.* 106, 522–538.
- REGO, C. 2001. Node-ejection chains for the vehicle routing problem: Sequential and parallel algorithms. *Paral. Comput.* 27, 3, 201–222.
- RESENDE, M. G. C. AND RIBEIRO, C. C. 1997. A GRASP for graph planarization. *Networks* 29, 173–189.
- RIBEIRO, C. C. AND SOUZA, M. C. 2002. Variable neighborhood search for the degree constrained minimum spanning tree problem. *Disc. Appl. Math.* 118, 43–54.
- SCHAERF, A. 1997. Combining local search and look-ahead for scheduling and constraint satisfaction problems. In *Proceedings of the 15th International Joint Conference on Artificial Intelligence, IJCAI 1997*. Morgan-Kaufmann Publishers, San Mateo, CA, 1254–1259.
- SCHAERF, A., CADOLI, M., AND LENZERINI, M. 2000. LOCAL++: a C++ framework for local search algorithms. *Softw. Pract. Exp.* 30, 3, 233–256.
- SHAW, P. 1998. Using constraint programming and local search methods to solve vehicle routing problems. In *Principle and Practice of Constraint Programming—CP98*, M. Maher and J.-F. Puget, Eds. Lecture Notes in Computer Science, vol. 1520. Springer.
- SIPPER, M., SANCHEZ, E., MANGE, D., TOMASSINI, M., PÉREZ-URIBE, A., AND STAUFFER, A. 1997. A phylogenetic, ontogenetic, and epigenetic view of bio-inspired hardware systems. *IEEE Trans. Evolut. Comput.* 1, 1, 83–97.
- SONDERGELD, L. AND VOß, S. 1999. Cooperative intelligent search using adaptive memory techniques. In *Meta-Heuristics: Advances and Trends in Local Search Paradigms for Optimization*, S. Voss, S. Martello, I. Osman, and C. Roucairol, Eds. Kluwer Academic Publishers, Chapter 21, 297–312.
- SPEARS, W. M., JONG, K. A. D., BÄCK, T., FOGEL, D. B., AND DE GARIS, H. 1993. An overview of evolutionary computation. In *Proceedings of the European Conference on Machine Learning (ECML-93)*, P. B. Brazdil, Ed. Vol. 667. Springer Verlag, Vienna, Austria, 442–459.
- STADLER, P. F. 1995. Towards a theory of landscapes. In *Complex Systems and Binary Networks*, R. López-Peña, R. Capovilla, R. García-Pelayo, H. Waelbroeck, and F. Zertuche, Eds. Lecture Notes in Physics, vol. 461. Springer-Verlag, Berlin, New York, 77–163. Also available as SFI preprint 95-03-030.
- STADLER, P. F. 1996. Landscapes and their correlation functions. *J. Math. Chem.* 20, 1–45. Also available as SFI preprint 95-07-067.
- STÜTZLE, T. 1999a. Iterated local search for the quadratic assignment problem. Tech. rep. aida-99-03, FG Intellektik, TU Darmstadt.
- STÜTZLE, T. 1999b. *Local Search Algorithms for Combinatorial Problems—Analysis, Algorithms and New Applications*. DISKI—Dissertationen zur Künstlichen Intelligenz. infix, Sankt Augustin, Germany.
- STÜTZLE, T. AND HOOS, H. H. 2000. MAX-MIN Ant System. *Fut. Gen. Comput. Syst.* 16, 8, 889–914.
- SYSWERDA, G. 1993. Simulated Crossover in Genetic Algorithms. In *Proceedings of the 2nd Workshop on Foundations of Genetic Algorithms*, L. Whitley, Ed. Morgan-Kaufmann Publishers, San Mateo, Calif., 239–255.
- TABU SEARCH WEBSITE. 2003. <http://www.tabusearch.net>. Visited in January 2003.
- TAILLARD, E. 1991. Robust Taboo Search for the Quadratic Assignment Problem. *Paral. Comput.* 17, 443–455.

- TALBI, E.-G. 2002. A Taxonomy of Hybrid Metaheuristics. *Journal of Heuristics* 8, 5, 541–564.
- TOULOUSE, M., CRAINIC, T., AND SANSÒ, B. 1999a. An experimental study of the systemic behavior of cooperative search algorithms. In *Meta-Heuristics: Advances and Trends in Local Search Paradigms for Optimization*, S. Voß, S. Martello, I. Osman, and C. Roucairol, Eds. Kluwer Academic Publishers, Chapter 26, 373–392.
- TOULOUSE, M., THULASIRAMAN, K., AND GLOVER, F. 1999b. Multi-level cooperative search: A new paradigm for combinatorial optimization and application to graph partitioning. In *Proceedings of the 5th International Euro-Par Conference on Parallel Processing*. Lecture Notes in Computer Science. Springer-Verlag, New York, 533–542.
- VAN KEMENADE, C. H. M. 1996. Explicit filtering of building blocks for genetic algorithms. In *Proceedings of the 4th Conference on Parallel Problem Solving from Nature—PPSN IV*, H.-M. Voigt, W. Ebeling, I. Rechenberg, and H.-P. Schwefel, Eds. Lecture Notes in Computer Science, vol. 1141. Springer, Berlin, 494–503.
- VAN LAARHOVEN, P. J. M., AARTS, E. H. L., AND LENSTRA, J. K. 1992. Job Shop Scheduling by Simulated Annealing. *Oper. Res.* 40, 113–125.
- VOSE, M. 1999. *The Simple Genetic Algorithm: Foundations and Theory*. Complex Adaptive Systems. MIT Press.
- VOß, S., MARTELLO, S., OSMAN, I. H., AND ROUCAIROL, C., Eds. 1999. *Meta-Heuristics—Advances and Trends in Local Search Paradigms for Optimization*. Kluwer Academic Publishers, Dordrecht, The Netherlands.
- VOß, S. AND WOODRUFF, D., Eds. 2002. *Optimization Software Class Libraries*. Kluwer Academic Publishers, Dordrecht, The Netherlands.
- VOUDOURIS, C. 1997. Guided local search for combinatorial optimization problems. Ph.D. dissertation, Department of Computer Science, University of Essex. pp. 166.
- VOUDOURIS, C. AND TSANG, E. 1999. Guided local search. *Europ. J. Oper. Res.* 113, 2, 469–499.
- WADE, A. S. AND RAYWARD-SMITH, V. J. 1997. Effective local search for the steiner tree problem. *Studies in Locational Analysis* 11, 219–241. Also in *Advances in Steiner Trees*, ed. by Ding-Zhu Du, J. M. Smith and J.H. Rubinstein, Kluwer, 2000.
- WATSON, R. A., HORNBY, G. S., AND POLLACK, J. B. 1998. Modeling building-block interdependency. In *Late Breaking Papers at the Genetic Programming 1998 Conference*, J. R. Koza, Ed. Stanford University Bookstore, University of Wisconsin, Madison, Wisconsin, USA.
- WHITLEY, D. 1989. The GENITOR algorithm and selective pressure: Why rank-based allocation of reproductive trials is best. In *Proceedings of the 3rd International Conference on Genetic Algorithms, ICGA 1989*. Morgan-Kaufmann Publishers, 116–121.
- YAGHURA, M. AND IBARAKI, T. 2001. On metaheuristic algorithms for combinatorial optimization problems. *Syst. Comput. Japan* 32, 3, 33–55.
- ZLOCHIN, M., BIRATTARI, M., MEULEAU, N., AND DORIGO, M. 2004. Model-based search for combinatorial optimization: A critical survey. *Ann. Oper. Res.* To appear.

Received July 2002; revised February 2003; accepted June 2003