

Memetic Algorithms in Discrete Optimization

Jin-Kao Hao

Appeared in F. Neri, C. Cotta, P. Moscato (Eds.) Handbook of Memetic Algorithms. Studies in Computational Intelligence 379, Chapter 6, pages 73-94, 2011

Abstract This chapter is dedicated to Memetic Algorithms for Discrete Optimization. It begins with a general survey, and then explains in depth the key ingredients of a successful MA. A particular attention is given to the following issues: Design of semantic combination operators, development of dedicated local search procedures and management of population diversity. Several other important issues are also discussed such as design of rich evaluation function and constraint handling techniques. This chapter includes two case studies with the purpose of showing how these issues can be effectively implemented in practice.

1 Introduction

Discrete optimization concerns in essence the search for a “best” configuration (optimal solution) among a set of *finite* candidate configurations according to a particular criterion. There are several ways to describe a discrete optimization problem. In its most general form, it can be defined as a collection of problem instances, each being specified by a pair (S, f) [53], where S is the set of finite candidate configurations, defining the *search space*; f is the *cost* or *objective function*, given by a mapping $f: S \rightarrow R^+$.

Solving the instance (S, f) is to find an $s^* \in S$ such that $f(s^*) \leq f(s)$ for all $s \in S$ (this minimization formulation can easily be transformed into a maximization problem). Such a configuration s^* is a globally optimal solution (or simply an optimal solution) to the given instance.

Given its generality, discrete optimization allows many problems of practical and theoretical importance to be conveniently formulated. Examples are the classical problems of general integer programming, permutation problems (e.g., trav-

Jin-Kao Hao
LERIA, Université d'Angers, 2 Boulevard Lavoisier, 49045 Angers Cedex 01, France
e-mail: jin-kao.hao@univ-angers.fr

eling salesman problem, bandwidth minimization, linear arrangement), and constraint satisfaction and optimization problems (satisfiability problems in propositional logic, graph partitioning, k -coloring). Discrete optimization naturally covers practical problems of the environment, renewable energy, distribution, infrastructure design, communications and productivity in the manufacturing and service sectors.

However, discrete optimization problems are known to be difficult to solve in general. Most of them, in particular those of practical interest, belong to the class of NP-hard problems, and thus cannot be efficiently solved to optimality. Over the past decades, important efforts have been made to improve the solution methods and important progresses have been achieved in both exact and heuristic strategies in pursuit of optimal or near optimal solutions.

This chapter concerns the design of Memetic Algorithms (MAs) [48, 49] for finding optimal or high quality near optimal solutions to hard discrete optimization problems.

progacionaort257(i12(r)-316.542(ct257(7252(m)0.193594)-0.345529(o)0.433749(c)0.442326(a)0.4240155(i)
Tarogonian a

mentary search strategies. Indeed, it is generally believed that the population-based search framework offers more facilities for exploration while neighborhood search provides more capabilities for exploitation. If they are combined in a suitable way, the resulting hybrid method can then offer a good balance between exploitation and exploration, assuring a high search performance.

Like other metaheuristics, MAs are a general optimization framework that can potentially be applied to various discrete search or optimization problems. Nevertheless, it should be clear that a blind application of MAs (or any other metaheuristics) to a particular problem will not be able to lead to satisfactory solutions. To be effective, the MA framework must be carefully adapted to the given problem and integrate problem-specific knowledge within its search operators and strategies. This is the very key point of a successful MA application in practice.

2.2 Memetic algorithms in overview

Memetic Algorithms [48, 49] are a population-based computational framework and share a number of features with methods like Evolutionary Algorithms [2] and Scatter Search [25]. MAs operate on a set of candidate solutions and use these solutions to create new solutions by applying variation operators such as combinations and local improvements.

From a general perspective, a MA is composed of a number of basic components: a pool of candidate solutions (also called population of individuals) to sample the search space, a combination operator (crossover) to create new candidate solutions (offspring) by blending two or more existing solutions, an improvement operator to ameliorate offspring solutions, and a population management strategy. In addition to these elements, the MA also needs an evaluation or fitness function to assess the quality of each candidate solution as well as a selection mechanism to determine the candidate solutions that will survive and undergo variations.

From an operational perspective, a typical MA starts with an initial population (see §3.4) and then repeats cycles of evolution. Each cycle, also called a generation, consists of four sequential steps.

1. *Selection of parents*: Selection aims to determine the candidate solutions that will survive in the following generations and be used to create new solutions. Selection for reproduction often operates in relation with the fitness (quality) of the candidate solutions; high quality solutions have thus more chances to be chosen. Well-known examples of selection strategies include roulette-wheel and tournament. Selection can also be done according to other criteria such as diversity. In such a case, only “distanced” individuals are allowed to survive and reproduce. If the solutions of the population are sufficiently diversified, selection can also be carried out randomly. The selection strategy influences the diversity of the population (see also §3.3).

2. *Combination of parents for offspring generation*: Combination aims to create new *promising* candidate solutions by blending (suitably) existing solutions (parents), a solution being promising if it can potentially lead the optimization process to new search areas where better solutions may be found. To achieve this, the combination operator is often designed such that it captures the semantics of the targeted problem to ensure the heritage of good properties from parents to offspring. Additionally, the design of the combination operator should ideally take care of creating *diversified* offspring. From a perspective of exploration and exploitation, such a combination is intended to play a role of strategic diversification with a long term goal of reinforcing the intensification. A carefully designed combination operator constitutes a driving force of a successful MA.
3. *Local improvement of offspring*: The goal of local improvement is to improve the quality of an offspring as far as possible. For this purpose, local improvement takes an offspring as its input (current solution) and then iteratively replaces the current solution by another solution taken from a given neighborhood. This process stops and returns the best solution found when a user-defined stop condition is met. Compared with the combination operator, local improvement plays essentially the role of intensifying the search by exploiting search paths delimited by the underlying neighborhood. Like combination, local improvement is another key component and driving force of a MA.
4. *Update of the population*: This step decides whether a new solution should become a member of the population and which existing solution of the population should be replaced. Often, these decisions are made according to criteria related to both quality and diversity. Such a strategy is commonly employed in methods like Scatter Search and many Evolutionary Algorithms. For instance, a basic quality-based updating rule would replace the worst solution of the population while a diversity-based rule would substitute for a similar solution according to a distance metric. Other criteria like recency (age) can also be considered. The policies employed for managing the population are essential to maintain an appropriate diversity of the population, to prevent the search process from premature convergence, and to help the algorithm to continually discover new promising search areas.

The general MA template is described in Algorithm 1 where special attention must be paid to the design of particular components. The stop condition can be a maximum number of cycles (generations), a maximum number of evaluations, a maximum number of cycles without improving the best solution, a solution quality to be reached or a lower-bounded threshold for the population diversity.

We deliberately leave out the mutation operator within this MA template. In some sense, local search can be viewed as a guided macro-mutation operator. However, mutation can also be applied to reinforce population diversity. As a lean design principle, only necessary components are included in a MA, any unjustified and superficial elements must be excluded.

Algorithm 1 Memetic Algorithm Template

```

1: Input:  $|P|$  (size of population  $P$ )
2: Output:  $s^*$  (records best solution found)
3:  $P = \text{PopGeneration}(|P|)$ 
4:  $\text{PopEvaluation}(P)$  (fitness evaluation of each individual)
5:  $s^* = \text{best}(P)$ 
6:  $f^* = f(s^*)$  (records the fitness of the best solution)
7: while Stop Condition is not verified do
8:    $(p_1 \dots p_k) = \text{ParentsSelection}(P)$  ( $k \geq 2$ , two or more parents are selected)
9:    $s' = \text{Recombination}(p_1 \dots p_k)$  (one or more offspring can be generated)
10:   $s = \text{OffspringImprovement}(s')$  (improvement of offspring solution(s) typically by local search)
11:   $P = \text{PopulationUpdate}(s, P)$  (population is updated according to a quality-diversity rule)
12:   $(s^*, f^*) = \text{BestSolutionUpdate}(s^*, f^*, P)$  (the best solution and its fitness are always recorded)
13: end while

```

2.3 Performance of memetic algorithms for discrete optimization

The computational performance of a MA depends first on the representation of the solution space (solution encoding) which should preferably be problem dependent and ease the design of efficient search operators.

The performance of a MA depends then on the design of its two key search components: Combination and local improvement operators. Their design should integrate useful problem-specific knowledge of the given problem in order to ensure aggressive exploitation and guided exploration.

The performance of a MA is also conditioned by the way the population is managed to promote and maintain a fertile diversity during the search process. Indeed, much like conventional Evolutionary Algorithms, premature convergence can easily occur if the population loses its diversity. Diversity management is particularly important with MAs because of the specific nature of their aggressive and intensified search strategies. Consequently, it is crucial for a MA to maintain with rigor a “good” population diversity as long as possible.

The interaction between the components of a MA can directly influence the behavior and the performance of the MA. A long or short local search phase after each combination could change the search trajectories. Similarly, a very effective local search procedure may weaken the role of the combination operator while a very strong combination operator may make it less critical to have a highly efficient local improvement procedure.

Finally, the runtime efficiency of a MA depends for a large part on the choice of the data structures employed to implement the different components of the MA. A typical example concerns local improvement procedures that explore the candidate solutions of a neighborhood and represent the most time-consuming part of a MA. In such a situation, it is critical to devise appropriate data structures to enable and streamline a fast neighborhood evaluation (see §3.1.3). Otherwise, the computational overheads will jeopardize the search power of the method.

3 Special design considerations

3.1 Design of dedicated local search

Local improvement is one of the most important components of a MA and ensures essentially the role of intensive exploitation of the search space. This is typically achieved either by dedicated local search heuristics (see examples in [39, 38, 40]) or by tailored general neighborhood search methods. In this part, we focus our discussion on adaptation of local search metaheuristics [35], but a large part of the discussion applies to the design of local improvement procedures based on specific heuristics.

3.1.1 Local search template

Let (S, f) be our search problem where S and f are respectively the search space and optimization objective. A neighborhood N over S is any function that associates to each solution $s \in S$ some other solutions $N(s) \subset S$. Any solution $s' \in N(s)$ is called a neighboring solution or simply a neighbor of s . For a given neighborhood N , a solution s is a *local optimum* with respect to N if s is the best in terms of f among the solutions in $N(s)$.

The notion of neighborhood can be explained in terms of the *move* operator. Typically applying a move mv to a solution s changes slightly s and leads to a neighboring solution s' . This transition from a solution to a neighbor is denoted by $s' = s \oplus mv$. Let $\Gamma(s)$ be the set of all possible moves which can be applied to s , then the neighborhood $N(s)$ of s can be defined by: $N(s) = \{s \oplus mv | mv \in \Gamma(s)\}$.

A typical local search algorithm begins with an initial configuration s in S and proceeds iteratively to visit a series of configurations following the neighborhood. At each iteration, a particular neighbor $s' \in N(s)$ is sought to replace the current configuration and the choice of s' is determined by the underlying metaheuristic and by referring to the quality of the neighboring solution. For instance, a strict Descent algorithm always replaces the current solution s by a *better* neighbor s' while tabu search replaces the current solution by a *best* neighbor s' even if the latter is of inferior quality. Still with simulated annealing, the transition from s to a randomly selected neighbor s' is conditioned by a changing probability.

3.1.2 Neighborhood design

The success of a LS algorithm depends strongly on its neighborhood. The neighborhood defines the subspace of the search problem to be explored by the method. For a given problem, the definition of the neighborhood should structure the search space such that it helps the search process to find its ways to good solutions.

The choice of neighborhood is conditioned by the representation (genotype) used to encode the candidate solutions of the search space (phenotype). It may further depend on the structure and constraints of the problem on hand. Here we briefly review some neighborhoods associated to three conventional representations, which have a variety of applications.

- *Binary representation*: With this representation, each solution of the search space is coded by a binary string. Binary representation is very popular in discrete optimization due to the fact that many problems are naturally formulated with binary variables. Typical examples include SAT/Max-SAT, Knapsack, Unconstrained Quadratic Optimization, graph bi-partitioning etc. For these binary problems, two basic neighborhoods are defined by the *k-flip* and *Swap* move operators. The *k-flip* move changes the values of k ($k \geq 1$) variables. So any neighbor $s' \in N(s)$ has a Hamming distance of k to solution s . A larger k induces a larger (and stronger) neighborhood. Nevertheless, whether a larger neighborhood should be preferred in practice depends on the computational cost to evaluate the neighborhood. *Swap* exchanges the values of two variables that have different values. Note that *Swap* can be simulated by two 1-*flip* moves.
- *Permutation representation*: Here, each solution of the search space corresponds to a permutation $\pi : \{1..n\} \rightarrow \{1..n\}$. Permutation representation has a large range of applications in discrete optimization. Prominent examples include Traveling Salesman Problem, Flow-Shop/Job-Shop scheduling, Linear Arrangement, Bandwidth Minimization etc. Two basic neighborhoods for this representation are available using *Swap* and *Rotation* moves. Given a permutation (solution) π , The *Swap* move exchanges $\pi(i)$ and $\pi(j)$ for some i and j ($i \neq j$). If π' is a neighbor of π by swapping i and j , then $\pi'(k) = \pi(k)$ for $k \neq i, j$, $\pi'(i) = \pi(j)$ and $\pi'(j) = \pi(i)$. The *Rotation* move rotates all the values between $\pi(i)$ and $\pi(j)$ for some $i < j$. Thus, if π' is a rotation neighbor of π obtained with $i < j$, then $\pi'(k) = \pi(k) + 1$ for $i \leq k < j$, $\pi'(j) = \pi(i)$, and $\pi'(k) = \pi(k)$ for all other k . Note that *Rotation*(i, j) can be simulated by $j - i$ successive *Swap* moves starting with *Swap*($i, i + 1$).
- *Integer representation*: With this representation, each solution of the search space corresponds to an integer vector whose values are taken from some discrete domains. Integer representation is very useful and convenient for many constraint satisfaction and optimization problems. A common neighborhood is defined by a “one-change” move that consists in replacing the current value of a single variable by a new domain value. The set of candidate variables under consideration for a value change can be identified with a number of rules specific to the problem at hand. For instance, if the search algorithm deals with unfeasible solutions, i.e. some variables are receiving conflicting values relative to some constraints, the set of candidate variables can be constituted of the subset of *conflicting* variables [13, 52, 15]. Such a neighborhood is typically employed in local search algorithms for solving Constraint Satisfaction Problems. More generally, candi-

date variables for a value change can be identified as those that are critical for improving the objective function or for reaching the feasibility.

These neighborhoods can be applied directly to a given problem if the problem fits well the required representation. A common practice is to adapt a conventional neighborhood with problem-specific knowledges of the problem. Moreover, in some situations, it is useful to investigate the possibility of multiple neighborhoods that can be applied at different stages of the search process (see §3.1.4 below).

3.1.3 Neighborhood evaluation

Another design issue that arises is the evaluation of a given neighborhood. Indeed, a local search procedure moves iteratively from the current solution to a new solution chosen within the neighborhood. To make this choice, local search needs to know the cost variation (also called the *move value*) between the current solution s and a candidate neighbor $s' \in N(s)$. The move value indicates whether the neighbor s' is of better, worse or equal quality relative to s . Let $\Delta f = f(s') - f(s)$ denote this move value.

- *Incremental evaluation*: Basically, there are two ways to obtain Δf for a neighbor. The trivial way is to calculate $f(s')$ from “scratch” using the objective function¹ f . Doing this way may be expensive if f needs to be evaluated very often or if the evaluation of f itself involves complex calculations. A more efficient alternative aims to derive the value of $f(s')$ from the value $f(s)$ by updating only what is strictly necessary. Indeed, if a neighbor s' is close to its initial solution s , which is true for many neighborhoods, then the evaluation of $f(s')$ can be carried out in this incremental manner. For a number of basic neighborhoods, like those shown previously, such an incremental evaluation is often possible.
- *Full search of neighborhood*: The incremental evaluation can be applied to *all the neighbors* of a given neighborhood relation. In this case, it is generally useful to investigate dedicated data structures (call it Δ -table) to store the move values for all the neighbors of the current solution. Δ -table provides a convenient way to know the quality of each neighbor and enables an efficient search of the full neighborhood. With such a Δ -table, the local search algorithm can decide easily at each iteration which neighbor to take according to its search strategy. For instance, a best-improvement descent algorithm will take the move that is identified by the most negative value in the Δ -table to minimize the objective function. After each move, the Δ -table (often only a portion of it) is updated accordingly using the incremental evaluation technique to propagate the effect of the move. Δ -table is a very useful technique for local search algorithms. This is particularly the case for descent-based methods like Tabu Search where a best neighbor needs

¹ For the reason of simplicity, the term “objective function” is used here. A more precise term is “evaluation function”, see §3.4.

to be identified (see examples in [35]).

- *Approximative evaluation*: The practical usefulness of Δ -table depends on both the complexity and the number of updates needed after each move transition. It may happen that, the move value can not be incrementally calculated or the Δ updates need to change a large portion of Δ -table. In this case, it would be useful to replace the initial evaluation function by a faster, but approximative evaluation function [36]. More generally, approximate evaluation is useful if the the evaluation function is computationally expensive to calculate or if the function is ill-defined.
- *Order of evaluation*: If the neighborhood is not completely searched, one must decide the order in which the neighborhood is explored. For instance, the first-improvement descent technique moves to any improving neighbor. If there are several improving neighbors, the descent search picks the “first” one encountered in the order the neighbors are examined. To allow such a method to increase its search diversity, a random order may be preferred [53].

3.1.4 Combination of neighborhoods

Very often, different neighborhoods may be available, enabling alternative ways to explore the search space. In such a situation, it is interesting to consider combined use of multiple neighborhoods. For illustrative purpose, consider two neighborhoods N_1 and N_2 . Then one can consider at least three ways to use them in a combined way.

First, *neighborhood union* $N_1 \cup N_2$ includes all the neighbors of the two underlying neighborhoods, so that any member of N_1 and N_2 is a member of $N_1 \cup N_2$. A local search algorithm using this combined neighborhood selects the next neighboring solution among all the solutions in both neighborhoods. This combination has no sense if one neighborhood is fully included in the other one.

With *Probabilistic neighborhood union* $N_1 \circ N_2$, a neighbor solution in N_1 (resp. N_2) belongs to $N_1 \circ N_2$ with probability p (resp. probability $1-p$). A local search algorithm using this combined neighborhood selects at each iteration the next neighbor from N_1 with probability p and from N_2 with probability $1-p$.

Token-ring combination $N_1 \rightarrow N_2$ is time-dependent and defined alternatively either by N_1 or N_2 according to some pre-defined conditions [17]. A local search algorithm using this combined neighborhood cycles through these neighborhoods. It typically starts with one neighborhood until the search stagnates, then changes to the other neighborhood until the search stagnates again to switch back to the first neighborhood and so on.

The advantage of combined neighborhood was already demonstrated long time ago in [40] for solving the Traveling Salesman Problem. More generally, the issue of transitioning among alternative neighborhoods was discussed with the Tabu Search framework and strategic oscillation design in [21]. More recent examples of local search methods focusing on multiple neighborhoods include Variable Neighborhood

Search [33], Neighborhood Portfolio Search [17] and Progressive Neighborhood Search [27]. Examples of studies on neighborhood combinations can be found in [32, 41].

3.2 Design of semantic combination operator

3.2.1 Solution combination

Combination is another key component of a MA and constitutes one leading force to explore the search space. The basic idea of combination is very appealing since it provides a very general way of generating new solutions by mixing existing solutions. Contrary to local changes of local improvement, combination can bring into new solutions more useful information, that may be beneficial for a healthy evolution of the search process.

As a first step, it would be tempting to consider the application of a blind (random) crossover operators for solution combinations. Doing this has the advantage of ease of application. However, one question should be asked before this approach is attempted: Is the crossover operator meaningful with respect to the optimization objective? If the answer is negative, the crossover operator is probably not appropriate and the sole role it would play in this case would be to introduce some random diversification in the search process.

In practice, instead of applying blind crossovers, it is often preferable to consider dedicated combination operators that have strong “semantics” with respect to the optimization objective. A semantic combination aims to pass intrinsic good properties from parents to offspring. The design of such a combination operator is far from trivial and in fact represents a challenging issue. Although there are some theoretical guidances, the discovery of such a semantic combination operator in practice relies basically on a deep analysis and understanding of the given problem. Compared with the design of local search procedures, the design of a meaningful combination operator constitutes probably one of the most creative parts of an effective MA.

3.2.2 Theoretical foundations

The schemata theory [34] and the building block hypothesis [29] are often mentioned to explain (partially) the performance of Genetic Algorithms. Intuitively, building blocks are promising patterns of solutions that can be progressively assembled by crossover to get improved solutions. Given that this theory is defined for binary and simple Genetic Algorithm, it is not directly applicable in the context of MAs. Nevertheless, assembling building blocks to generate new solutions remains an appealing idea. In [56, 57], the concept of forma is introduced to generalize the schemata theory. A formal framework is even proposed to try to capture some fundamental aspects of MA in [58]. The forma theory suggests a set of gen-

eral principles for the design of solution representations and recombination operators. According to this theory, a suitable recombination operator is required to fulfill two conditions called *respect* and *proper assortment*. Intuitively, the *respect* condition advocates the heritage of shared characteristics of parents to offspring, while *proper assortment* ensures the heritage of desirable characteristics of each parent by their offspring. This is in accordance with the general principle of conserving good features through inheritance and discarding bad features developed in Grouping Genetic Algorithms [6].

3.2.3 Design of combination operator

These abstract considerations only provide us with very general guidances for designing recombination operators. For a particular problem, it is still necessary to find out what are the building blocks (interesting patterns or characteristics) of solutions that can be assembled and inherited through the recombination process. Unfortunately, there is no short-cut to this quest and a fine analysis and deep understanding of the given problem is indispensable to find useful clues.

First, one can analyze the samples of optimal or high quality solutions to possibly identify regular patterns shared by these solutions. Indeed, if such a pattern exists, then the recombination operator can be constrained to conserve the pattern from the parent solutions and to avoid breaking the pattern. Alternatively, the recombination operator can also be encouraged to promote the emergence of favorable building blocks. For instance, such an analysis applied to the Traveling Salesman Problem shows that high quality local optima share sub-tours [39, 40]. This property has been used by several highly successful crossover operators which conserve common edges or sub-tours in offspring solutions [50, 68, 12, 51, 46]. Similarly, for the graph k -coloring problem, an analysis of coloring solutions discloses that some nodes are always grouped to the same color class (i.e. colored with the same color). This characteristic has helped to devise powerful combination operators, as shown in [5, 14].

3.2.4 Multi-parent combination

Combination may operate with more than two parents. Multiple parent combination is even a general rule for the Scatter Search metaheuristic which uses, in its original form, linear combinations of several solutions to create new solutions [19]. Although there is no theoretical justifications, the practical advantage of multiple parent recombinations was demonstrated in several occasions for discrete optimization. For instance, for the graph k -coloring problem, several recent and top-performing algorithms integrate multiple parent combination [16, 45, 43, 54], where color classes from different solutions are assembled to build offspring colorings. More generally, when multiple solutions are used for creating a new solution, one can define spe-

cial rules to score the solution components of each parent solution and use strategic voting rules to combine components from different parents solutions.

A question that arises for multi-parent combination is how to determine the number of the parents. By using two parents, the offspring is expected to inherit 50% material from each parent. The contribution of each parent to the new solution decreases with an increasing number of parents. If the building blocks from different parents are independent from one another, taking more parents into account would be interesting to build good and diversified offspring. Otherwise, if a building block from a parent is epistatic with respect to the building blocks of other parents, blending more parents means more disruption, and thus should be avoided.

3.3 Population diversity management

Population diversity is another important issue that should be considered in the design of an effective MA [14, 63]. If the population diversity is not properly managed, the population will converge prematurely and the search process stops with poor local optima. This is particularly true when a small population is used by the MA. In what follows, we first provide some precisions about the nature of diversity and explain how fertile diversity can be promoted and maintained within a population. Note however that diversity is not interesting per se within a MA. The ultimate goal of population diversity is to help the search process not only to avoid premature converge, but also to continually discover interesting new solutions in order to explore non-visited promising search areas.

3.3.1 Diversity

Population diversity can be measured by a similarity (or distance) metric applied to the members of the population. The metric can be defined either on the solution representation level (genotype metric) or solution level (phenotype metric) [30]. For instance, pair-wise *Hamming distance* can be used as a genotype metric to measure population diversity. Diversity can also be measured in terms of *entropy* [18] or by the so-called *moment of inertia* [47]. Genotype metric is usually problem independent, and thus may or may not reflect the intrinsic diversity of a population with respect to the given optimization objective.

Population diversity can also be measured at the phenotype level over the solution space. For instance, for partition problems like graph k -coloring, the distance between two partitions can be measured by the so-called *transfer distance* which is the minimum number of elements that need to be moved between classes of one partition so that the resulting partition becomes the other partition [59, 3]. A phenotype metric is defined over the solution space and thus is more likely to measure the real diversity of a population.

In order to observe suitably the population diversity, it is useful to first determine the most appropriate distance or similarity metric with respect to the optimization objective of the given problem. Moreover, if the population diversity needs to be continually monitored, it becomes important to pay attention to the cost of computing the underlying metric.

3.3.2 Promoting and maintaining useful diversity

Population diversity can be promoted and managed at several levels of a MA. One evident possibility is to define specific selection rules to favor the selection of *distanced parents* for mating. Another possibility concerns the variation operators which can be designed in such a way that they favors the generation of divers and varied offspring. For instance, the “Distance Preserving Crossover” introduced in [12, 46] is constrained to generate an offspring which is at the same distance from both parents. More generally, the path-relinking type of combinations typically construct offspring solutions by considering both the solution quality and its distance to its parent solutions [25] (see also [42] for an example).

Population diversity can also be controlled by the offsprin

by controlling the distance between each new solution and the existing solutions of the population. Only distant new solutions are allowed to join the population.

- *Distance*: At several places, MAs may need to measure the distance between two solutions or between a solution and a group of solutions. For instance, parents selection may operate in such a way that the selected parents are sufficiently distant. Similarly, a population management strategy may decide the acceptance or rejection of an offspring by considering its distance to the members of the population. When an operation refers to the notion of distance, it is preferable to employ an appropriate distance metric which is meaningful with respect to the given problem. For instance, for partition problems like graph coloring (see §4.1), Hamming distance is not a suitable metric to characterize the difference of two partitions. Instead, transfer distance between partitions should be preferred. Once again, the choice of the distance metric should ideally be correlated with the semantics of the problem on hand.
- *Rich evaluation function*: Evaluation function assesses the quality of a candidate solution with respect to the optimization objective and orients the search method to “navigate” through the search space. A good evaluation function is expected to be able to distinguish each solution from the other solutions and thus to effectively guide the search method to make the most appropriate choice at each iteration. Very often, the initial optimization objective f is directly used as evaluation function. However, such a function may not be sufficiently discriminant to distinguish different solutions. To improve the discriminating power, it is useful to incorporate in the evaluation function additional information, e.g. relative to the structure of the problem instance to be solved. Examples can be found in [37, 6, 60]. Moreover, when constrained optimization problems are considered, some constraints may be hard to satisfy, and thus are relaxed. Among various constraint relaxation techniques, a common practice is to integrate the relaxed constraints into the evaluation function as a (weighted) component or as a part of a multi-component evaluation function (see examples in [22, 65, 66]).
- *Constraints*: The constraints in the considered problem may influence the design of some MA components. For instance, suppose that the MA algorithm is expected to explore only feasible solutions. Then one must decide whether a combination operator is constrained to create only feasible solutions. If infeasible offspring is allowed, it is necessary to consider a dedicated mechanism to repair the broken constraints. Similarly, neighborhood design can take into consideration the constraints to identify eligible moves. For instance, in feasibility search problems, this is often done by identifying problem variables involving violated constraints and restricting the set of authorized moves to those defined on these conflicting variables. Finally, as previously stated, constraints that are difficult to solve can be used in the design of the evaluation function.

- *Connections with Scatter Search and Path Relinking*: As discussed in [20] and [23] (Chapter 9), the MA framework shares ideas with Scatter Search and Path Relinking [24, 25]. These latter methods provide unifying principles for joining solutions based on generalized path constructions (in both Euclidean and neighborhood spaces) and by using strategic design. Solution combination in Scatter Search originated historically from strategies for combining decision rules and combining constraints. In Scatter Search, dispersed new solutions are created from a set of reference solutions by weighted combinations of subsets of the reference solutions that are selected as elite solutions. With Path Relinking, offspring solutions are generated by exploring, within a neighborhood space, trajectories that connect two or more reference solutions. One notices that the reference solutions or subsets of them can be considered as parent solutions for combination while combination resorts to diverse strategies such as attribute voting and weighting.

4 Case Studies

In this section, we show two case studies of quite different nature with the purpose of showing how these issues can be effectively implemented in practice. We particularly focus on the design of combination and local search operators.

4.1 Graph coloring problems

4.1.1 Problem description

Given an integer k and a undirected graph $G = (V, E)$ with a set V of vertices and a set E of edges, a legal k -coloring of G is a partition of V into k distinct color classes such that each color class is composed of pairwise non-adjacent vertices. The graph k -coloring problem (k -COLOR) aims at finding a legal k -coloring for a fixed k while the graph coloring problem (COLOR) determines the smallest k for a given graph G (its *chromatic number* χ_G) such that G has a legal k -coloring. Since COLOR can be handled by solving a series of k -COLOR with decreasing k values, we only consider here k -COLOR.

For a given k -COLOR instance, i.e. an integer k and graph $G = (V, E)$, let $s = \{C_1, C_2, \dots, C_k\}$ denote a partition of V into k distinct color classes such that each C_i ($i \in \{1, 2, \dots, k\}$) contains all the vertices that are colored with color i . Let S denote all such partitions. For any $s \in S$, define its conflict number $f(s)$ to be the number of pairs of adjacent vertices x and y ($\{x, y\} \in E$) belonging to a same color class of s . Then k -COLOR can be solved by minimizing $f(s)$; $f(s)=0$ implies that s is a legal k -coloring, i.e. all its color classes C_i are conflict-free.

Notice that among the large number of existing heuristic algorithms for k -COLOR, Memetic Algorithms are certainly among the most powerful ones and provide the best results on the well-known DIMACS benchmark instances of this well-known NP-complete problem.

4.1.2 Partition crossovers

In order to design a semantic combination operator, let us try to get an idea about the possible “building blocks” for our problem. The goal of k -COLOR is to determine a set of k *distinct* conflict-free color classes. In this context, color classes can be considered our basic “building blocks”. If there are several “good” color classes among some candidate solutions, then these color classes can favorably be recombined to obtain new candidate solutions. This idea was first explored by the Greedy Partition Crossover (GPX) described in [14] and the Union of Independent Sets crossover in [5], which are also related to the design of grouping crossovers described in [6].

Operating with two parent k -colorings s_1 and s_2 , GPX builds step by step the k classes C_1^0, \dots, C_k^0 of the offspring s_0 . At the first step, GPX creates C_1^0 by choosing a *largest* class from one parent and remove its vertices from both parents s_1 and s_2 . GPX repeats then the same operations for the next $k-1$ steps, but alternates each time the parent considered. If some vertices remain unassigned at the end of these k steps, they are randomly assigned to one of the k color classes. The alternation between the parents aims at a balanced mixture of information from both parents and avoiding the dominance of one parent over the other one during the recombination.

Table 1 shows an example with 3 color classes ($k = 3$) and 10 vertices represented by capital letters A,B, ..., J.

parent s_1 →	<table border="1"><tr><td>A B C</td><td>D E F G</td><td>H I J</td></tr></table>	A B C	D E F G	H I J	$C_1^0 := \{D, E, F, G\}$ remove D, E, F and G	<table border="1"><tr><td>A B C</td><td></td><td>H I J</td></tr></table>	A B C		H I J
A B C	D E F G	H I J							
A B C		H I J							
parent s_2 →	<table border="1"><tr><td>C D E G</td><td>A F I</td><td>B H J</td></tr></table>	C D E G	A F I	B H J	<table border="1"><tr><td>C</td><td>A I</td><td>B H J</td></tr></table>	C	A I	B H J	
C D E G	A F I	B H J							
C	A I	B H J							
offspring s	<table border="1"><tr><td></td><td></td><td></td></tr></table>				<table border="1"><tr><td>D E F G</td><td></td><td></td></tr></table>	D E F G			
D E F G									
parent s_1 →	<table border="1"><tr><td>A B C</td><td></td><td>H I J</td></tr></table>	A B C		H I J	$C_2^0 := \{B, H, J\}$ remove B, H and J	<table border="1"><tr><td>A C</td><td></td><td>I</td></tr></table>	A C		I
A B C		H I J							
A C		I							
parent s_2 →	<table border="1"><tr><td>C</td><td>A I</td><td>B H J</td></tr></table>	C	A I	B H J	<table border="1"><tr><td>C</td><td>A I</td><td></td></tr></table>	C	A I		
C	A I	B H J							
C	A I								
offspring s	<table border="1"><tr><td>D E F G</td><td></td><td></td></tr></table>	D E F G			<table border="1"><tr><td>D E F G</td><td>B H J</td><td></td></tr></table>	D E F G	B H J		
D E F G									
D E F G	B H J								
parent s_1 →	<table border="1"><tr><td>A C</td><td></td><td>I</td></tr></table>	A C		I	$C_3^0 := \{A, C\}$ remove A and C	<table border="1"><tr><td></td><td></td><td>I</td></tr></table>			I
A C		I							
		I							
parent s_2 →	<table border="1"><tr><td>C</td><td>A I</td><td></td></tr></table>	C	A I		<table border="1"><tr><td></td><td>I</td><td></td></tr></table>		I		
C	A I								
	I								
offspring s	<table border="1"><tr><td>D E F G</td><td>B H J</td><td></td></tr></table>	D E F G	B H J		<table border="1"><tr><td>D E F G</td><td>B H J</td><td>A C</td></tr></table>	D E F G	B H J	A C	
D E F G	B H J								
D E F G	B H J	A C							

Table 1 The Greedy Partition Crossover: An example from [14]

The basic idea underlying GPX was also explored with multiple parent combination operators [31, 16, 45, 43, 54]. Using multiple parents for combination is fertile for k -COLOR since this offers more possibilities to obtain good (large) color classes for each step of the recombination operation. By generalizing two parents to multiple parents, refined and additional strategies were also introduced to make the combination process as most effective as possible. For instance with the AMaPX

operator of [43], in order to favor the creation of *diversified offspring*, each time a color class from a parent is transmitted to the offspring, this parent k -coloring will not be considered for the next few steps of offspring building. In [54, 55], in order to measure the goodness of the color classes of the parent colorings, the combination operator takes into account the size of each color class, the number of conflicting vertices as well as the degrees of the vertices in the color class.

A question that arises when multiple parents are used is how to determine the number of parents. It is clear that by using more parents, fewer classes will be transmitted from each parent to the offspring and this also implies that the class blending from each parent is also more disrupted. An analysis of the relations between the number of vertices, the number of color classes and the number of parents permits to identify a heuristic rule to fix the right number of parents [55].

In [16], the combination operation is performed within a slightly different context. The algorithm maintains a pool of conflict-free color classes obtained during the search process. From time to time, these color classes are used to generate new k -colorings. Other combination operators using similar ideas are investigated in [5, 31, 45]

4.1.3 Local improvement by Tabu Search

In memetic coloring algorithms, Tabu Search is frequently used for local improvement to ameliorate a new offspring created by the combination operator. For illustration purpose, we use the TS algorithm described in [14] as an example. It uses the constrained “one-change” move described in §3.1.2 such that a *neighbor* of a given configuration s is obtained by moving a single *conflicting* vertex v from a color class C_i to another color class C_j . When such a move $\langle v, i \rangle$ is performed, the couple $\langle v, i \rangle$ is classified tabu for the next tl iterations. Therefore, v cannot be reassigned to the class i during this period, unless moving v back to the color class i leads to a configuration better than the best configuration found so far (*aspiration criterion*). The tabu tenure tl for a move $\langle v, i \rangle$ depends on the number of iterations since the move was first performed. The tabu tenure tl for a move $\langle v, i \rangle$ is defined as $tl = \lfloor \alpha \cdot (it - it_0) \rfloor + 1$, where it is the current iteration number, it_0 is the iteration number when the move was first performed, and α is a constant between 0 and 1.

4.2 Maximum parsimony phylogeny

4.2.1 Problem description

Phylogenetics is the study of evolutionary relationships among various groups of organisms (for example, species or populations). These connections are represented graphically through phylogenetic trees. Computational phylogenetics aims to infer phylogenetic trees from molecular data such as protein or DNA sequences [10]. The main phylogenetic approaches include methods using a distance-matrix, the maximum likelihood or maximum parsimony criterion.

Maximum parsimony phylogeny generally takes as input a multiple sequence alignment which is a matrix M of characters composed of n lines (related to a set S of species, where $|S| = n$) and k columns which represent the characters of the sequences [7]. Each sequence is also called a taxon. Each character of the matrix belongs to an alphabet Σ . A phylogenetic tree T of the given input is a binary tree such that (1) the leaves of T are the set of n species, and (2) each internal node is induced by the sequence of parsimony of its two descendant sequences. Given two sequences $S_1 = \langle x_1, \dots, x_k \rangle$ and $S_2 = \langle y_1, \dots, y_k \rangle$ with $\forall i \in \{1..k\}, x_i, y_i$ belonging to the power set $\mathcal{P}(\Sigma = \{-, A, C, G, T\})$, the sequence of parsimony $P(S_1, S_2) = \langle z_1, \dots, z_k \rangle$ of S_1 and S_2 is given by ([8]) :

$$\forall i, 1 \leq i \leq k, z_i = \begin{cases} x_i \cup y_i, & \text{if } x_i \cap y_i = \emptyset \\ x_i \cap y_i, & \text{otherwise} \end{cases}$$

The score of the sequence of parsimony defines the “distance” separating its two descent sequences:

$$f_{P(S_1, S_2)} = \sum_{i=1}^k c_i \quad \text{where} \quad c_i = \begin{cases} 1, & \text{if } x_i \cap y_i = \emptyset \\ 0, & \text{otherwise} \end{cases}$$

Let T be a binary parsimony tree with n leafs or species. T has then $n - 1$ sequences of parsimony (internal nodes). Let I denote the set of these internal nodes. The Fitch parsimony score $f(T)$ of T is defined as follows:

$$f(T) = \sum_{i \in I} f_i(T)$$

The aim of the Maximum Parsimony problem (MP) is then to find a most parsimonious phylogenetic tree T^* such that T^* minimizes the parsimony score. Since there are $\prod_{i=3}^n (2i - 3)$ possible binary trees with n leafs, this problem is a highly combinatorial search problem. The MP problem is computationally difficult since its associated decision problem is equivalent to the NP-complete Steiner problem in a hypercube [11]. MP has been subject of many studies for many years. Among them, neighborhood-based local search and various hybrid algorithms are certainly the most popular solution methods. In what follows, we show a Memetic Algorithm

called HYDRA, which combines a dedicated tree crossover called DiBIP [26] and a progressive neighborhood local search method [27].

4.2.2 Distance-based information preservation crossover

First, let us notice that conventional tree crossovers known in genetic programming are not suitable here. The Distance-Based Information Preservation crossover (DiBIP) is specifically designed for the MP problem. DiBIP is based on a topological distance between species (leaves) and aims to preserve common properties of parents in terms of this distance between species. For instance, two species that are close (resp. far) in both parents should stay close (resp. distant) in the offspring. Given two parents trees, the DiBIP crossover is realized in three steps: Calculate a distance matrix for each parent tree, then combine the two resulting matrices to get a third matrix and finally create a child tree from this last matrix.

The general DiBIP crossover scheme is described in Algorithm 2 where T_1 and T_2 denote two parents trees. δ is a distance metric to measure the distance of each pair of species of a tree T , Δ a tree-to-distance operator to obtain a distance matrix of a tree, \oplus a matrix operator to combine 2 distance matrices to produce a new distance matrix, Λ a distance-to-tree operator to construct a tree from a given distance matrix.

Algorithm 2 The general DiBIP crossover scheme

Input: $T_1, T_2, \delta, \Delta, \oplus, \Lambda$

Output: A child tree T^*

1. Apply the tree-to-distance operator Δ to each parent tree T_i ($i=1,2$) to obtain the corresponding distance matrix $D_i = \Delta(T_i)$;
 2. Apply the matrix operator \oplus to D_1 and D_2 to obtain D^* : $D^* = D_1 \oplus D_2$;
 3. Apply the distance-to-tree operator Λ to D^* to obtain a child tree: $T^* = \Lambda(D^*)$.
-

A specific DiBIP crossover operator is obtained once δ , Δ , \oplus , and Λ are provided. The distance measure δ should be ideally correlated to the evolutionary changes between species. For instance, 2 species separated in the tree by a small number of evolutionary changes should have a smaller distance than 2 species separated by a large number of changes. The distance measure should additionally be tree-topology dependent. In this sense, the length of the elementary path between 2 species is a possible option while Hamming distance is not suitable here because this metric is totally independent of tree topologies.

Moreover, since we want to preserve representative features of the parents during the crossover operation, a valid matrix operator \oplus should favor such an inheritance from parents to offspring and meet some relation preservation property. For instance, if a pair of species (a,b) is closer than another pair (c,d) in both parents, then this relation should be conserved. Consider the operation \oplus such that for a pair of species (i,j) , $(D_1 \oplus D_2)(i,j) = \alpha \cdot \min\{D_1(i,j), D_2(i,j)\} + (1 - \alpha) \cdot \max\{D_1(i,j), D_2(i,j)\}$ with $\alpha \in [0, 1]$. This defines indeed a valid \oplus operator. Furthermore, this definition

offers in fact many possibilities and seems particularly relevant to MP. For instance, the arithmetic average ($\alpha = 0.5$) and the max operator \max ($\alpha = 0$) are 2 special cases. At last, let us mention that the arithmetic addition is another simple valid \oplus operator.

We now show a concrete example. Given two species i and j , define their distance δ_{ij} to be the *topological distance*, i.e. the length of the elementary path between the respective ascendants of i and j , (minus 1 if the path contains the root of the tree T). The matrix operator \oplus is the addition $+$ such that $D(i, j) = D_1(i, j) + D_2(i, j)$, which satisfies the relation preservation property previously mentioned. The distance-to-tree operator Λ is a non-deterministic variant of the well-known UPGMA method [62]. Fig. 1 shows an application of this crossover operator. One observes that the closeness of species in both parents is conserved in the child. This observation applies equally to distant species.

4.2.3 Progressive neighborhood search

For local improvement, HYDRA uses *Progressive Neighborhood Search* (PNS) which operates with a variable-size neighborhoods [27]. Given a parsimony tree T , a neighboring tree T' is typically obtained by a move that consists in cutting a subtree from T and reinserting the sub-tree elsewhere in the initial tree. If a meaningful metric can be defined to measure the distance between the cutting and inserting points, then it would be possible to define neighborhoods of variable sizes. In [27], the topological distance δ shown in Section 4.2.2 is used for this purpose. A distance parameter d is introduced to constraint the distance between the pruned edge i and the edge j receiving the insertion such that $\delta_{ij} \leq d$.

So, setting $d = \infty$ leads to a large neighborhood where the pruned edge (with its subtree) can be reinserted anywhere in the tree. Consequently, the topological change can be important. This case corresponds in fact to the well-known *Subtree Pruning Regrafting* neighborhood [64] whose size equals $2(n-3)(2n-7)$ [1]. Reversely, setting $d = 1$ gives a small neighborhood where neighboring trees are close to the current tree. This case corresponds to another well-known neighborhood called *Nearest Neighbor Interchange* [67] which swaps two adjacent branches of the tree leading to $(2n-6)$ neighbors [61]. By varying the parameter d , one gets neighborhoods of intermediate sizes.

The *Progressive Neighborhood Search* is based on this parametric neighborhood and its neighborhood changes during the search process by varying the value of d . In the particular MP context, PNS carries out a descent search starting with a large neighborhood (i.e. with large d) and reduces progressively the neighborhood. Indeed, at the beginning of the search, it is possible to obtain strong quality improvement by important topological modifications of the tree with large d . When the search progresses, the quality of the trees becomes better and better, only small improvements can be expected with small tree modifications. It is thus more judicious to switch to smaller and small neighborhoods to accelerate the search.




Fig. 1 Application of the DiBIP Tree Crossover [26]

One notices that PNS shares some features with Variable Neighborhood Search (VNS) [33]. However, contrary to VNS, the neighborhoods explored by PNS are not systematically of increasing sizes. Within the context of our Maximum Parsimony problem, PNS even reduces progressively its neighborhood.

5 Conclusions

In this chapter we have presented the basic concepts of Memetic Algorithms for Discrete Optimization. Focus is given to the key design issues of an effective MA algorithm. We have explained the usefulness of a deep study and understanding of

the optimization problem on hand. We have insisted on the importance of a careful adaptation of the general search strategies offered by the MA framework, a suitable incorporation of problem specific knowledge in different components of the MA as well as a logical integration of these components. The pursuit goal is clearly to build an effective MA algorithm that is able to ensure a balanced exploitation and exploration of the search space.

It should be clear that a blind MA application would have little chance to deliver good results for difficult optimization problems. High performance can only be possible by a disciplined and careful specialization of the general MA framework to the targeted problem. It is equally important to apply the “lean design” principle in order to avoid redundant or superficial algorithmic components.

The framework of Memetic Algorithms constitutes an interesting enrichment to the arsenal of existing discrete optimization methods and offers a valuable alternative for tackling hard discrete optimization problems.

Acknowledgement

Acknowledgements This work was partially supported by “Angers Loire Métropole” and the Region of “Pays de la Loire” within the following projects: MILES, BIL, Radapop and LigeRO. I’m grateful to my research collaborators within and outside the Group “Metaheuristics, Optimization and Applications” of the LERIA Laboratory.

References

1. B. L. Allen and M. Steel. Subtree transfer operations and their induced metrics on evolutionary trees. *Annals of Combinatorics* 5(1):1-15, 2000.
2. T. Bäck, D. B. Fogel, and Z. Michalewicz (Eds). *Handbook of evolutionary computation*. Springer, 1997.
3. W.H.E. Day. The complexity of computing metric distances between partitions. *Mathematical Social Sciences* 1:269–287, 1981.
4. K.A. De Jong. An analysis of the behavior of a class of genetic adaptive systems. PhD thesis, University of Michigan Ann Arbor, MI, USA, 1975.
5. R. Dorne and J.K. Hao. A new genetic local search algorithm for graph coloring. In *PPSN 98*, volume 1498 of *LNCS*, pages 745–754. Springer, 1998.
6. E. Falkenauer. *Genetic algorithms and grouping problems*. John Wiley & Sons, Inc. New York, USA, 1998.
7. J. Felsenstein. Evolutionary trees from DNA sequences: a maximum likelihood approach. *Journal of Molecular Evolution*, 17:368–376, 1981.
8. W. Fitch. Towards defining course of evolution: minimum change for a specified tree topology. *Systematic Zoology*, 20:406–416, 1971.
9. W.M. Fitch and E. Margoliash. Construction of phylogenetic trees. *Science*, 155(3760):279–284, 1967.
10. J. Felsenstein. *Inferring phylogenies*. Sinauer Associates, 2003.
11. L.R. Foulds and R.L. Graham. The steiner problem in phylogeny is NP-complete. *Advances in Applied Mathematics*, 3:43–49, 1982.

12. B. Freisleben and P. Merz. New genetic local search operators for the traveling salesman problem. in H.M. Voigt *et al.* (Eds.): PPSN 1996, *Lecture Notes in Computer Science*, 1141:890–899, Springer, 1996.
13. P. Galinier and J.K. Hao. Tabu search for maximal constraint satisfaction problems. G. Smolka (Ed.): CP 1997, *Lecture Notes in Computer Science*, 1330:196–208, Springer 1997.
14. P. Galinier and J.K. Hao. Hybrid evolutionary algorithms for graph coloring. *Journal of Combinatorial Optimization*, 3(4):379–397, 1999.
15. P. Galinier and J.K. Hao. A general approach for constraint solving by local search. *Journal of Mathematical Modelling and Algorithms*, 3(1):73–88, 2004.
16. P. Galinier, A. Hertz, N. Zufferey. An adaptive memory algorithm for the k-coloring problem. *Discrete Applied Mathematics*, 156(2):267–279, 2008.
17. L. Di Gaspero and A. Schaerf. Neighborhood portfolio approach for local search applied to timetabling problems. *Journal of Mathematical Modeling and Algorithms*, 5(1):65–89, 2006.
18. C. Fleurent and J.A. Ferland. Object-oriented implementation of heuristic search methods for graph coloring, maximum clique, and satisfiability. In *Cliques, Coloring, and Satisfiability: Second DIMACS Implementation Challenge*, volume 26 of DIMACS Series in Discrete Mathematics and Theoretical Computer Science, 619–652, 1996.
19. F. Glover. Heuristics for integer programming using surrogate constraints. *Decision Sciences*, 8(1):156–166, 1977.
20. F. Glover. Tabu search for nonlinear and parametric optimization (with links to genetic algorithms). *Discrete Applied Mathematics*, 49: 231–255, 1994.
21. F. Glover. Tabu Search and adaptive memory programming advances, applications and challenges, pp 1-75. In R.S. Barr, R.V. Helgason, J.L. Kennington (Eds.) *Interfaces in Computer Science and Operations Research*, 7: 1-75, 1997.
22. F. Glover and G.A. Kochenberger. Critical event tabu search for multidimensional knapsack problems. In I.H. Osman, J.P. Kelly (Eds): *Metaheuristics: The Theory and Applications*, pages 407-427. Kluwer Academic Publishers, 1996
23. F. Glover and M. Laguna. Tabu Search. Kluwer Academic Publishers, 1997.
24. F. Glover. A template for scatter search and path relinking. in J.K. Hao, E. Lutton, E. Ronald, M. Schoenauer, D. Snyers (Eds): AE 1997, *Lecture Notes in Computer Science*, 1363:1–51, 1998.
25. F. Glover, M. Laguna, R. Marti. Fundamentals of scatter search and path-relinking. *Control and Cybernetics* 39:654–684, 2000.
26. A. Goëffon, J.M. Richer, J.K. Hao. A distance-based information preservation tree crossover for the maximum parsimony problem. In T.P. Runarsson, H. G. Beyer, E.K. Burke, J.J. Merelo Guervs, L.D. Whitley, X. Yao (Eds.): PPSN 2006, *Lecture Notes in Computer Science*, 4193:761–770, 2006.
27. A. Goëffon, J.M. Richer, J.K. Hao. Progressive tree neighborhood applied to the maximum parsimony problem. *IEEE/ACM Transactions on Computational Biology and Bioinformatics* 5(1): 136-145, 2008.
28. A. Goëffon, J.M. Richer, J.K. Hao. A memetic algorithm for phylogenetic reconstruction with maximum parsimony. In C. Pizzuti, M.D. Ritchie, and M. Giacobini (Eds.): *EvoBIO 2009, Lecture Notes in Computer Science* 5483: 164–175, 2009.
29. D.E. Goldberg. Genetic algorithms in search, optimization and machine learning. Addison-Wesley, 1989.
30. D. E. Goldberg and J. Richardson. Genetic algorithms with sharing for multimodal function optimization. In J. J. Grefensette (Ed.), Proc. 2nd Int. Conf. Genetic Algorithms, Hillsdale, NJ: Lawrence Erlbaum, pp. 41-49, 1987.
31. J. P. Hamiez and J. K. Hao. Scatter search for graph coloring. P. Collet, C. Fonlupt, J.K. Hao, E. Lutton, M. Schoenauer (Eds.): AE 2001, *Lecture Notes in Computer Science* 2310:168–179. Springer, 2002.
32. J. P. Hamiez, J. Robet, J.K. Hao. A tabu search algorithm with direct representation for strip packing. C. Cotta and P. Cowling (Eds.): *EvoCOP 2009, Lecture Notes in Computer Science* 5482:61-72, Springer, 2009.

33. P. Hansen and N. Mladenovic. An introduction to variable neighborhood search. In S. Voss et al. (Eds): *Metaheuristics, Advances and Trends in Local Search Paradigms for Optimization*, pages 433–458, Kluwer Academic Publishers, Dordrecht, 1999.
34. J.H. Holland. *Adaptation and artificial systems*. Ann Arbor, University of Michigan Press, 1975.
35. H. Hoos and T. Stutzle. *Stochastic local search: foundations and applications*. Morgan Kaufmann Publishers Inc. San Francisco, CA, USA, 2004.
36. Y. Jin. A comprehensive survey of fitness approximation in evolutionary computation *Soft Computing - A Fusion of Foundations, Methodologies and Applications* 9(1):3–12, 2005.
37. D. S. Johnson, C. R. Aragon, L. A. McGeoch, C. Schevon. Optimization by simulated annealing: An experimental evaluation; part II, graph coloring and number partitioning. *Operations Research*, 39(3), 378–406, 1991.
38. B.W. Kernighan and S. Lin. An efficient heuristic procedure for partitioning graphs. *The Bell System Tech J.*, 49(2):291–307, 1970.
39. S. Lin. Computer solutions of the traveling salesman problem. *The Bell System Tech Journal*, 44(10):2245–2269, 1965.
40. S. Lin and B.W. Kernighan. An effective heuristic algorithm for the traveling salesman problem. *Operations Research*, 21:498–516, 1973.
41. Z. Lü, F. Glover, J.K. Hao. Neighborhood combination for unconstrained binary quadratic problems. *MIC-09 Post-Conference Volume*, Springer, 2010.
42. Z. Lü, F. Glover, J.K. Hao. A hybrid metaheuristic approach to solving the UBQP problem. Working paper, 2010.
43. Z. Lü and J.K. Hao. A memetic algorithm for graph coloring. *European Journal of Operational Research*, 200(1):235–244, 2010.
44. S.W. Mahfoud1992. Crowding and preselection revisited. R. Manner and B. Manderick (Eds), *PPSN 1992: pages27–36*, Amsterdam, The Netherlands: Elsevier Science, 1992.
45. E. Malaguti, M. Monaci, P. Toth. A metaheuristic approach for the vertex coloring problem. *INFORMS Journal on Computing*, 20(2):302–316, 2008.
46. P. Merz and B. Freisleben. Memetic algorithms for the traveling salesman problem. *Complex Systems*, 13(4):297–345, 2001.
47. R.W. Morrison and K.A. De Jong. Measurement of population diversity. P. Collet, C. Fonlupt, J.K. Hao, E. Lutton, M. Schoenauer (Eds.): *AE 2001, Lecture Notes in Computer Science* 2310:31–41, Springer, 2002.
48. P. Moscato. On evolution, search, optimization, genetic algorithms and martial arts: towards memetic algorithms. Caltech Concurrent Computation Program, C3P Report 826, 1989.
49. P. Moscato. Memetic algorithms: a short introduction. In D. Corne, M. Dorigo, F. Glover (Eds), *New Ideas in Optimization*, McGraw-Hill Ltd., Maidenhead, UK. 219–234, 1999.
50. H.Müehlenbein, M. Gorges-Schleuter, O. Krämer. Evolution algorithms in combinatorial optimization. *Parallel Computing*, 7:65–88, 1988.
51. Y. Nagata, S. Kobayashi. Edge assembly crossover: a high-power genetic algorithm for the travelling salesman problem. Thomas Bäck (Ed.): *Proceedings of the 7th International Conference on Genetic Algorithms*, pages 450–457, Morgan Kaufmann, 1997
52. K. Nonobe and T. Ibaraki. A tabu search approach for the constraint satisfaction problem as a general problem solver. *European Journal of Operational Research*, 106(2–3):599–623, 1998.
53. C. Papadimitriou and K. Steiglitz. *Combinatorial optimization: algorithms and complexity*. Prentice-Hall, 1982.
54. D.C. Porumbel, J.K. Hao, P. Kuntz. Diversity control and multi-parent recombination for evolutionary graph coloring algorithms. C. Cotta and P. Cowling (Eds.): *EvoCOP 2009, Lecture Notes in Computer Science* 5482:121–132, Springer, 2009.
55. D.C. Porumbel, J.K. Hao, P. Kuntz. An evolutionary approach with diversity guarantee and well-Informed grouping recombination for graph coloring. *Working paper*, 2009.
56. N.J. Radcliffe. Forma analysis and random respectful recombination In R.K. Belew, L.B. Booker (Eds.): *Proceedings of the 4th International Conference on Genetic Algorithms*, pages 222–229, Morgan Kaufmann, 1991.

57. N.J. Radcliffe. The algebra of genetic algorithms. *Annals of Mathematics and Artificial Intelligence*, 10(4): 339–384, 1993.
58. N.J. Radcliffe and P.D. Surry. Formal memetic algorithms. In Terence C. Fogarty (Ed.) *Lecture Notes in Computer Science* 865: 1–16, Springer 1994.
59. S. Régnier. Sur quelques aspects mathématiques des problèmes de classification automatique. *Mathématiques et Sciences Humaines*, 82:20, 1983. reprint of *ICC Bulletin*, 4:175–191, Rome, 1965.
60. E. Rodriguez-Tello, J.K. Hao, J. Torres-Jimenez. An effective two-stage simulated annealing algorithm for the Minimum Linear Arrangement problem. *Computers and Operations Research*, 35(10):3331–3346, 2008.
61. D.F. Robinson. Comparison of labeled trees with valency three. *J. Combin. Theory*, 11:105–119, 1971.
62. P.H.A. Sneath and R.R. Sokal. Numerical taxonomy. W.H. Freeman and Co., San Francisco, California, 1973.
63. K. Sörensen and M. Sevaux. MA|PM: memetic algorithms with population management. *Computers and Operations Research*, 33(5):1214–1225, 2006
64. D. L. Swofford and G. J. Olsen. Phylogeny Reconstruction. D.M. Hillis and C. Moritz (Ed.) *Molecular Systematics*, chapter 11:411–501, 1990.
65. M. Vasquez and J.K. Hao. A heuristic approach for antenna positioning in cellular networks. *Journal of Heuristics*, 7(5):443–472, 2001.
66. M. Vasquez and J.K. Hao. A "logic-constrained" knapsack formulation and a Tabu algorithm for the daily photograph scheduling of an earth observation satellite. *Computational Optimization and Applications*, 20(2): 137–157, 2001.
67. M.S. Waterman and T.F. Smith. On the similarity of dendograms. *Journal of Theoretical Biology*, 73:789–800, 1978.
68. D. Whiteley, T. Starkweather, D. Fuquay. Scheduling problems and traveling salesmen: the genetic edge recombination operator. *Proceedings of the 3rd International Conference of Genetic Algorithms*, pages 133–140, 1989.