CIRRELT

Centre interuniversitaire de recherche sur les réseaux d'entreprise, la logistique et le transport

Interuniversity Research Centre on Enterprise Networks, Logistics and Transportation

Metaheuristics: A Canadian Perspective

Michel Gendreau
Jean-Yves Potvin

December 2007

CIRRELT-2007-60

Bureaux de Montréal :
Université de Montréal
C.P. 6128, succ. Centre-ville
Montréal (Québec)
Canada H3C 3J7
Téléphone : 514 343-7575
Télécopie : 514 343-7121

Bureaux de Québec :
Université Laval
Pavillon Palasis-Prince, local 2642
Québec (Québec)
Canada G1K 7P4
Téléphone : 418 656-2073
Télécopie : 418 656-2624

www.cirrelt.ca

Metaheuristics: A Canadian Perspective

Michel Gendreau[1] , Jean-Yves Potvin[1,*]

[1.] Interuniversity Research Centre on Enterprise Networks, Logistics and Transportation (CIRRELT), and Department of Computer Science and Operations Research, Université de Montréal, P.O. Box 6128, Station Centre-ville, Montréal, Canada H3C 3J7

Abstract. Metaheuristics are generic search strategies that can be adapted to solve complex problems. This paper describes in simple terms the most popular metaheuristics for combinatorial optimization problems. It also emphasizes the main contributions of the Canadian research community in the development and application of metaheuristics.

Keywords. Metaheuristics, combinatorial optimization, variable neighbourhood search, tabu search, simulated annealing, greedy randomized adaptive search procedure, ant colony optimization, evolutionary algorithms.

Acknowledgements. Financial support for this work was provided by the Natural Sciences and Engineering Research Council of Canada (NSERC) under grants 36662-06 and 38816-05.This support is gratefully acknowledged.

---

* Corresponding author: Jean-Yves.Potvin@cirrelt.ca

This document is also published as Publication #1313 by the Department of Computer Science and Operations Research of the Université de Montréal.

Dépôt légal – Bibliothèque nationale du Québec,
Bibliothèque nationale du Canada, 2007

1 Introduction

A metaheuristic, according to its original definition, is a general search framework that specifies a heuristic strategy to explore the solution space. The name combines the Greek prefix *meta* (*beyond*, here in the sense of *higher level*) and *heuristic* (from *heuriskein*, to find). A metaheuristic thus provides a means to tackle different types of problems with relatively few modifications. It includes, in particular, strategies to escape from local optima in complex solution spaces, like strategies based on transitions from the current solution to some other solution through local modifications or moves that define a neighbourhood structure.

Metaheuristics have emerged in the 1980's (Glover coined the term *metaheuristics* for such methods in 1986 [56]) and have constantly gained popularity among researchers and practitioners since then. Previously, specialized heuristics were developed to solve complex problems. With the advent of these more general problem-solving approaches, the picture drastically changed. Metaheuristics are now, in most cases, the method of choice for solving complex, ill-defined problems, particularly those of a combinatorial nature. The challenge then becomes one of adapting a metaheuristic to the problem at hand, which is often a much easier task than developing a specialized heuristic from scratch. Furthermore, a good metaheuristic implementation is likely to provide near-optimal solutions in very reasonable computation times. With increasingly powerful computers and parallel platforms, metaheuristics have even been applied with success to real-time problems with stringent response time requirements.

The aim of this paper is to review the most popular metaheuristics and to emphasize the main contributions of the Canadian research community. Accordingly, this paper provides a biased and partial view on metaheuristics and is by no means a survey of the field. For a broader view on metaheuristics, the reader is referred to the following books, monographies and edited volumes [1, 3, 11, 23, 30, 33, 34, 59, 60, 63, 74, 77, 83, 90, 96, 97, 102, 104, 105, 106, 117].

In the following, metaheuristics are divided into two categories: single solution based metaheuristics, where a single solution and search trajectory is considered, and population based metaheuristics, where many solutions evolve concurrently. Within each category, it is also possible to distinguish between primarily constructive metaheuristics, where a solution is built from scratch, through the introduction of new elements at each iteration, and improvement metaheuristics that iteratively modify a solution. In Sections 2 and 3, the constructive metaheuristics are the greedy randomized adaptive search procedure and ant colony optimization, respectively. The others can be considered as primarily improvement metaheuristics.

2 Single Solution Based Metaheuristics

In this class of metaheuristics, a single solution is worked on and a single search trajectory is followed in the solution space.

2.1 Variable Neighbourhood Search (VNS)

We start our review with VNS, as it was designed and developed in the mid 1990's by Hansen and Mladenović [92] at the *Groupe d'études et de recherche en analyse des décisions* (GERAD) in Montreal. The basic idea is that an efficient metaheuristic can be obtained by systematically changing the neighbourhood of a local search heuristic. In practice, nested neighbourhoods that are increasingly distant from the current solution are explored until an improvement is found.

Variable neighbourhood search works as follows. A solution is randomly generated in the first neighbourhood of the current solution, from which a local descent is performed. If the local optimum obtained is not better than the incumbent (best known solution), then the procedure is repeated with the next neighbourhood. The search restarts from the first neighbourhood when either a solution that is better than the incumbent has been found or every neighbourhood structure has been explored. A popular variant is the deterministic *Variable Neighbourhood Descent* where the best neighbour of the current solution is considered instead of a random one (and no local descent is performed on this neighbour). The best neighbour automatically becomes the new current solution if an improvement is obtained, and the search is then restarted from the first neighbourhood. Otherwise, the next neighbourhood is considered. The search stops when all neighbourhood structures have been considered without improvement. At this point, the solution is a local optimum over all neighbourhood structures.

An interesting development proposed by Hansen, Mladenović and Perez-Brito is the *Variable Neighbourhood Decomposition Search* where the kth neighbourhood consists of fixing all but k attributes (or variables) in the current solution [72]. The local descent is then performed on the k free variables only. Another variant is the *skewed VNS* where the evaluation of the current local optimum is biased with the distance from the incumbent [65]. The idea is to favor the exploration of distant solutions from the incumbent to diversify the search. More recently, hybrids with tabu search have been reported [67], where either tabu search is used within VNS or conversely. In the first case, the local descent in VNS is replaced by a tabu search while, in the second case, different neighbourhood structures are exploited by the tabu search. Finally, several ways to parallelize VNS have been reported in the literature [25, 67]. One natural approach is to generate many random neighbours from the current solution (either in the same neighbourhood or in different neighbourhoods) and to assign each neighbour to a different processor to perform the local descent. A somewhat different approach consists in applying *cooperative search* principles: several search threads perform VNS simultaneously and exchange information when they encounter globally improving solutions.

VNS has been applied to many classical combinatorial optimization problems, including the traveling salesman, quadratic assignment, p-median, graph coloring and maximum clique problems, and also to industrial problems like oil pipeline and radar code design [66]. Surveys and tutorials about VNS and VND may be found in [67, 68, 69, 70, 71].

2.2 Tabu Search (TS)

The principles of TS come from the work of Fred Glover. The origins of the name *Tabu Search* itself are traced back to the seminal 1986 paper in which Glover also coined the term metaheuristics [56].

Tabu search is basically a deterministic local search strategy where, at each iteration, the best solution in the neighbourhood of the current solution is selected as the new current solution, even if it leads to an increase in solution cost. As opposed to a pure local descent, the method can thus escape from local optima. A short-term memory, known as the tabu list, stores attributes of recently visited solutions to avoid short-term cycling. Typically, the search stops after a fixed number of iterations or after a maximum number of consecutive iterations without any improvement to the incumbent. Detailed descriptions of TS can be found in [43, 44, 50, 57, 58, 60, 61].

Starting from the simple search scheme described above, a number of refinements have been proposed over the years. These include:

- *Intensification and diversification.* Intensification focuses the search in a promising region of the search space. Conversely, diversification drives the search toward unexplored regions. These features are typically implemented via different forms of medium to long-term memories.
  - *Frequency memories* record how often certain solution attributes are found in previously visited solutions. Neighbourhood solutions that contain elements with high frequency counts can then be penalized to allow the search to visit other regions of the search space. This mechanism provides a form of continuous diversification by introducing a bias in the evaluation of neighbourhood solutions at each iteration. Solutions with high frequency counts may also be deliberately excluded from solutions to induce more radical diversification.
  - *Adaptive memories* [109] provide a means to both diversify and intensify the search. Such memories contain a pool of previously generated elite solutions, which can be used to restart the search. This is usually done by taking different fragments of elite solutions and by combining them to generate a new starting solution, similarly to some population based metaheuristics (see section 3). Intensification or diversification is obtained depending if the fragments are taken from solutions that lie in a common region of the search space or not. Adaptive memories provide a generic paradigm for guiding local search and can be coupled with different types of metaheuristics [116].
  - *Path relinking* [62] is a mechanism that generates new solutions by exploring trajectories between elite solutions. Starting from one of these solutions, it generates a path in the neighbourhood space leading to another solution, called the guiding solution. This can be done by selecting modifications that introduce attributes found in the guiding solution. This mechanism can be used to diversify or intensify the search, depending on the path generation mechanism and the choice of the initiating and guiding solutions.
  - *Strategic oscillation* [60] requires defining an oscillation boundary, which is often a feasibility boundary. Then, the search is allowed to go for a specified depth beyond the boundary before turning around. When the boundary is crossed again from the opposite direction, the search goes beyond it for a specified depth before turning around again. By repeating this procedure an oscillatory search pattern is produced. It is possible to vary the amplitude of the oscillation to explore a particular region of the search space. For example, tight oscillations favor a more thorough search around the boundary.
- *Reactive tabu search.* This approach provides a mechanism for dynamically adjusting the search parameters, based on the search history [6]. In particular, the size of the tabu list is automatically increased when some configurations are repeated too often to avoid short-term cycles (and conversely).
- *Probabilistic TS.* This approach is aimed at introducing randomization into the search method [60]. This can be achieved by associating probabilities with moves leading to neighbourhood solutions, based on their evaluation. As a complete neighbourhood evaluation is often too computationally expensive, randomization can also be introduced through the design of a candidate list strategy, where only a subset of neighbourhood solutions is considered.

Tabu search is a very popular metaheuristic in North America. Thus, a substantial number of developments and applications are reported in Canada, in particular at the *Centre for research on transportation* (CRT). We only emphasize a few in the following.

The *TABUROUTE* algorithm of Gendreau, Hertz and Laporte [47] was one of the first truly effective implementations of TS for solving the Vehicle Routing Problem (VRP). In this problem, customers must be served by vehicles of finite capacity, through routes that start and end at a central depot, with the objective of minimizing the total distance traveled by the vehicles. The neighbourhood structure is defined through the removal of a customer from a route followed by its reinsertion into another route that contains at least one of its nearest neighbours, using the generalized insertion heuristic *GENI* [46]. The tabu search implementation also contains features that were quite innovative at the time. As the search is allowed to explore the infeasible domain, two penalty terms are included in the objective to account for the capacity and maximum route length constraints. These penalties are each weighted by a self-adjusting parameter as follows: every 10 iterations, each parameter value is divided by 2 if the 10 previous solutions were all feasible, to allow the search to reach the infeasible domain, or multiplied by 2 if they were all infeasible, to drive the search toward the feasible domain. Also, the diversification strategy penalizes customers that were moved frequently in the past to allow to search to consider alternative moves. To this end, an artificial term is added to the objective, which is proportional to the absolute frequency of movement for the customer being considered. Finally, *TABUROUTE* exploits *false starts*, where several initial solutions are generated and a limited search is carried out from each one. The best solution observed at the end then becomes the true starting point of the tabu search.

The *Unified Tabu Search Algorithm* (UTSA) of Cordeau, Gendreau and Laporte [20] was developed to provide a simple, yet efficient, search framework. The distinctive features of UTSA come from easily adaptable neighbourhood structures based on solution attributes, coupled with a dynamic adjustment of a few parameters. The proposed framework thus provides both generality and flexibility. It was successfully applied to different types of vehicle routing problems [20, 21, 22]. A tabu search heuristic that follows similar principles is also described in [4].

A tabu search heuristic for the vehicle routing problem with time windows (VRPTW) is reported in [115]. Here, a time interval is associated with each customer to constrain the service time. The tabu search exploits a *CROSS* exchange neighbourhood, which is a generalization of the *2-opt** exchanges of Potvin and Rousseau [100]. Both types of exchanges are now widely used to create neighbourhood structures for problems with time windows. The *2-opt** consists in exchanging the end parts of two routes, while preserving their original orientation (as the time windows induce an implicit orientation on each route). The *CROSS* exchanges generalize this idea by exchanging segments of consecutive customers of varying length between the two routes.

Recently, Desaulniers et al. have developed an exact problem-solving approach for the VRPTW based on column generation and tabu search [29]. In the proposed algorithm, TS is used to quickly identify columns (feasible routes) with a negative reduced cost in a first stage, while an exact dynamic programming algorithm is applied later to generate the remaining columns.

Le Bouthillier et al. [84, 85] have proposed a parallel algorithmic scheme where different metaheuristics explore the search space concurrently and communicate through a central warehouse made of elite solutions (i.e., an adaptive memory). When a metaheuristic returns its best solution found, the latter is included in the warehouse if it is better than its worse solution. This warehouse is then used to provide new starting points for the metaheuristics. Under this paradigm, different tabu and genetic searches are run in parallel to solve the

VRPTW. The authors show that their parallel scheme is competitive with state-of-the-art approaches on Solomon's and Gehring and Homberger's data sets [42, 114].

A number of TS implementations for transportation network design problems have also been developed at the CRT. As such problems often involve both integer location and continuous flow variables, they offer a variety of opportunities for the application of TS. First, from any feasible set of values for the location variables, one can retrieve optimal values for the flow variables by solving the associated transportation problem. The tabu search can thus be used to explore the set of feasible vectors of location variables, by alternatively closing and opening locations, with any solution in that search space being completed by computing the associated optimal flow variables, see [26, 51, 53]. One could also decide to search the extreme points of the set of feasible flow vectors, retrieving the associated location variables by simply noting that a location must be opened whenever some flow is allocated to it. This kind of approach is used in [24] for the fixed-charge capacitated multicommodity network design problem. In [54], cycle-based neighbourhoods are exploited for solving the same problem. Here, two different paths associated with the same origin-destination pair are selected and the flow on one of these paths is transferred to the other. A number of arcs are then opened or closed depending on the resulting flow. It is worth noting that telecommunications network design problems have also been addressed by Canadian researchers and the interested reader is referred to [5, 7, 14, 15, 55] for details.

An interesting tabu search implementation for a real-world problem in the steel industry has been developed at the University of Toronto [88]. The goal here is to produce coils by scheduling steel slabs through a hot strip mill under multiple and conflicting objectives and constraints. A solution is made up of blocks, where each block is a sequence of slabs, that are processed in a particular order. The authors introduce a so-called *cannibalization* procedure within the TS, where a bad subsequence of slabs in a given block is exchanged with a good subsequence in another block, using *CROSS* exchanges [115]. An application of TS for a real-word forest management problem is also reported by Richards and Gunn [107, 108]. In this work, TS is used to explore different harvest period assignments.

Metaheuristics have been perceived, for quite a long time, as too slow to handle real-time contexts where quick response times are required. But this has proved to be false. In fact, the first real-time applications of tabu search come, for a large part, from Canada. Real-time vehicle routing applications motivated from courier services are found in [45, 48, 78, 79], while a real-time ambulance redeployment problem is reported in [49]. Also, an interesting military application with very stringent response time requirements is described in [10]. In this case, the tabu search manipulates operating scenarios encoded as tree structures.

2.3 Simulated Annealing (SA)

Simulated annealing was developed in the early 1980's [13, 81]. It is a randomized local search procedure where a modification to the current solution leading to an increase in solution cost can be accepted with some probability. This algorithm is motivated from an analogy with the physical annealing process used to find low-energy states of solids. In condensed matter physics, annealing denotes a process in which a solid is first melted by increasing its temperature; this is followed by a progressive temperature reduction aimed at recovering a solid state of lower energy. If the cooling is done too fast, widespread irregularities emerge in the structure of the solid, thus leading to relatively high energy states. Conversely, a careful annealing through a series of levels, where the temperature is held long enough at each level to reach equilibrium, leads to more regular structures associated with low-energy states. Basically, the process is less likely to get trapped in a high-energy state when the temperature is prevented from getting too far from the current energy level.

In a combinatorial optimization context, a solution corresponds to a state of the physical system and the solution cost to the energy of the system. At each iteration, the current solution is modified by randomly selecting a move from a particular class of transformations (which defines a neighbourhood of solutions). If the new solution provides an improvement, it is automatically accepted and becomes the new current solution. Otherwise, the new solution is accepted according to the Metropolis criterion, where the probability of acceptance is related to the magnitude of the cost increase and a parameter called the *temperature*. Basically, a move is more likely to be accepted if the temperature is high and the cost increase is low. The temperature parameter is progressively lowered, according to some predefined cooling schedule, and a certain number of iterations are performed at each temperature level. When the temperature is sufficiently low, only improving moves are accepted and the method stops at a local optimum. As opposed to most metaheuristics, this method asymptotically converges to a global optimum (assuming an infinite number of iterations). Finite-time implementations, however, do not provide such a guarantee.

Developments around this basic search scheme have focused on a number of issues [2, 73]. For example, deterministic variants like threshold accepting [35] and record-to-record travel [36] have been studied. In these variants, a transition is accepted if it does not increase the cost by more than some predefined value; the transition is rejected otherwise. The predefined value is progressively reduced as the search progresses. Research work has also been devoted to the design of static and dynamic cooling schedules aimed at increasing the speed of convergence, without compromising solution quality. Finally, substantial efforts have taken place on hybridization with other metaheuristics and on parallel implementations.

It should be observed that SA is popular in Europe but has not induced a sustained research activity in North-America, as opposed to VNS and TS.

2.4 Greedy Randomized Adaptive Search Procedure (GRASP)

Multi-start local search methods repeatedly apply a local search from different initial solutions. The use of a quick greedy heuristic to generate starting solutions looks attractive, if the greedy solutions are different enough to allow for a good sampling of local optima. Semi-greedy or randomized greedy heuristics have thus been proposed in the late 1980's to add variability to greedy heuristics, which led to the search scheme known as GRASP. The latter is a multi-start procedure designed by Feo and Resende [37, 38], where each restart applies a randomized greedy construction heuristic to generate an initial solution, which is then improved through local search. This is repeated for a given number of restarts and the best overall solution is returned. At each step of the construction heuristic, the elements not yet incorporated into the partial solution are evaluated with a greedy function, and the best elements are kept in a so-called *restricted candidate list* (RCL). One element is then randomly chosen from this list and incorporated into the solution. Through randomization, the best current element is not necessarily chosen, thus leading to a diversity of solutions.

One shortcoming of GRASP comes from the fact that each restart is independent of the previous ones, thus preventing the exploitation of information from previously obtained solutions to guide the search. Some recent developments are thus aimed at providing such capabilities. One example is the reactive GRASP [101], where the size of the RCL is dynamically adjusted, depending on the quality of recently generated solutions. Another example is the use of memories to guide the search. In Fleurent and Glover [39], a pool of elite solutions is maintained to bias the probability distribution associated with the elements in the RCL. Intensification or diversification can be obtained by either rewarding or penalizing elements that are often found in the pool of elite solutions. Such a pool can also be used to implement path relinking [62], by generating a search trajectory between a randomly chosen elite solution and the current local optimum [82].

Other recent developments propose hybrids where the construction step is followed by more sophisticated local search procedures, in particular variable neighbourhood search [89]. Parallel implementations are also reported, mostly based on the distribution of the restarts over multiple processors.

Quite surprisingly, the research community in Canada has mostly ignored this metaheuristic for reasons that remain unclear. This is particularly surprising given that GRASP was designed and developed in the United States. One exception is found in the work of Gendron, Potvin and Soriano [52], where a GRASP-like approach iteratively constructs new solutions for a network design problem, based on information gathered from previously generated solutions.

3 Population Based Metaheuristics

This class of metaheuristics explicitly works with a population of solutions (rather than a single solution). In such a case, different solutions are combined, either implicitly or explicitly, to create new solutions.

3.1 Ant Colony Optimization (ACO)

It is well known that ants can manage to find shortest paths from their nest to food sources. The medium used for communication among the ants is a chemical compound, known as *pheromone*, which is laid down in various quantities on the ground. While an isolated ant would more or less wander randomly, an ant detecting a pheromone trail will follow it (with some probability) and will strengthen it with its own pheromone. Through this positive reinforcement loop, the probability that ants will follow that path in the future increases with the number of ants that previously followed it. It has been shown experimentally that this mechanism can give rise to the emergence of shortest paths (as the pheromone tends to accumulate faster on these paths). In ACO, at each cycle, a number of artificial ants sequentially construct solutions in a randomized and greedy way. Each ant chooses the next element to be incorporated into its current partial solution on the basis of some heuristic evaluation and the amount of pheromone associated with that element. The pheromone represents the memory of the system, and is related to the presence of that element in previously constructed solutions (in the same way as the strength of a pheromone trail is related to how many ants previously chose to follow that path). Randomization is used to allow the construction of a variety of different solutions. Basically, a probability distribution is defined over all elements that can be incorporated into the current partial solution, with a bias in favor of the best elements. In particular, an element with a good heuristic evaluation and a high level of pheromone is more likely to be selected. Each time an element is selected by an ant, its pheromone level is updated by first removing a fraction of it, to mimic pheromone evaporation, and then by adding some new pheromone. When all ants have constructed a complete solution, the procedure is restarted with the updated pheromone levels. This is repeated for a fixed number of cycles or until the search stagnates.

Since the original ant system framework described in [31] for solving the Traveling Salesman Problem, refinements have been integrated into this general iterative scheme to lead to the ACO. These refinements provide mechanisms to either intensify or diversify the search. For example, more pheromone may be associated with elements found in the incumbent, to favor a more intense search around that solution. Conversely, the pheromone levels may be reduced for some elements, to force the construction of diversified solutions. For a precise description and historical perspective on ACO, see [33].

Important developments with regard to this metaheuristic are the *Hybrid Ant System* [41], where the pheromone trails are used to guide a local search heuristic rather than a construction heuristic, the use of multiple ant colonies that interact by exchanging information about good solutions [40, 91], the coupling of ACO with local search to further improve the solutions constructed by the ants and the design of parallel implementations where subcolonies evolve on different processors.

In Canada, an interesting development to the ant system [32] was proposed by Hoos at the University of British Columbia in 2000 [75]. Together with Stützle, he developed the so-called *Max-Min Ant System*. In this variant, premature convergence is avoided through a simple mechanism that forces the pheromone strength to lie within a minimum and maximum value. Also, some form of intensification is provided by allowing only the best ants to add pheromone during the pheromone update. Hoos has also worked for the last fifteen years on the design and empirical analysis of stochastic local search algorithms, from rather simple constructive and iterative improvement procedures to more complex metaheuristics [76]. Another interesting development in Canada comes from the work of Le Louarn, Gendreau and Potvin [87] who proposed more powerful ants that construct tours based on the generalized insertion heuristic GENI [46]. With regard to real-world applications, Gravel et al. describe an ant-based system for a production problem in the aluminum industry [64], while an application in computational molecular biology by Shmygelska and Hoos is reported in [113].

3.1.1 Evolutionary Algorithms (EAs)

Evolutionary algorithms represent a large class of problem-solving methodologies, with genetic algorithms (GAs) being the most widely known [74]. These algorithms are motivated by the way species evolve and adapt to their environment, based on the Darwinian principle of natural selection. Under this paradigm, a population of solutions (often encoded as a bit or integer string, referred to as a *chromosome*) evolves from one generation to the next through the application of operators that mimic those found in nature, namely, selection of the fittest, crossover and mutation. Through the selection process, which is probabilistically biased toward the best elements in the population, only the best solutions are allowed to become *parents* and to generate *offspring*. The mating process, called crossover, then takes two selected parent solutions and combine some of their features to create one or two offspring solutions. This is repeated until a new population of offspring solutions is created. Before replacing the old population, each member of the new population undergoes small random perturbations via the mutation operator. Starting from a randomly or heuristically generated initial population, this renewal cycle is repeated for a number of iterations, and the best solution found is returned. Details about this metaheuristic can be found in [63, 90, 94].

Genetic algorithms have been successfully applied to many combinatorial optimization problems, although this success has often been achieved by departing from the classical scheme outlined above. In particular, the encoding of solutions into chromosomes is either completely ignored (by working directly on the solutions) or specifically designed for specialized crossover and mutation operators. The distinctive feature of GAs remains the exploitation of a population of solutions and the creation of new solutions through the recombination of good attributes of two parent solutions. Many single solution metaheuristics now integrate this feature (e.g., via adaptive memories). There is also a clear connection between recombination, where an *intermediate* solution is generated from two parent solutions, and path relinking [103].

One of the main improvements comes from the observation that, although the population improves on average over the generations, the algorithm often fails to generate near-optimal solutions. Thus, some form of intensification is needed, and this is usually achieved in modern GAs by considering powerful local search operators as a form of mutation. That is, rather than introducing small random perturbations into the offspring solution, a local search is applied to improve the solution until a local optimum is reached (this combination of GA and local search is often called a memetic algorithm [93]).

Fined-grained and coarse-grained parallel GAs are also widely used [12]. In the former case, each solution is implemented on a different processor and is allowed to mate only with solutions that are not too distant in the underlying connection topology. This approach may be implemented in a totally asynchronous way, with each solution evolving independently of the others. Restricting the pool of mating candidates to a subset of the population has beneficial effects of its own, by allowing good attributes to diffuse more slowly into the population, thus preventing premature convergence to grossly suboptimal solutions. In coarse-grained implementations, subpopulations of solutions evolve in parallel until migration of the best individuals from one subpopulation to another takes place to refresh the genetic pool. This so-called *island model* however requires some form of synchronization.

Canadian researchers have been active in the development of GA applications, in particular in the field of vehicle routing. Potvin and Bengio [99] report one of the first application of genetic algorithms to the VRPTW. In this work, the objective is to minimize the number of vehicles first and, for the same number of vehicles, to minimize the total distance. The authors apply the crossover and mutation operators directly on the solutions, thus avoiding the problematic issue of encoding a solution into a chromosome (an idea later adopted by Berger et al. [9], see below). The crossover operator, where two routes from two different solutions are exchanged, followed by a repair operator to transform the offspring solution into a valid and feasible solution, was adopted by many researchers in the following years. The authors also recognized the need to minimize the number of vehicles through special mutation operators that try to empty routes with only a few customers. The authors reported at the time the best cumulative number of vehicles and found many new best known solutions on Solomon's test set [114].

At the Defence Research Establishment Valcartier, Berger et al. [9] developed a few years later a GA with a crossover and a mutation operator inspired by the ruin and recreate principle [111] and the large neighbourhood search [112]. Their approach consists of removing a number of customers from a solution and to create a new solution by reinserting these customers in some (hopefully) better way. An innovative approach is also proposed to handle the hierarchical objective of the VRPTW. Basically, the authors relax the capacity and time window constraints and evolve two populations of solutions with a fixed number of vehicles. Each solution in the first population contains m vehicles, among which at least one solution is feasible, with the objective of minimizing the total distance. Each solution in the second population contains $m-1$ vehicles with the objective of minimizing constraint violations. As soon as a feasible solution with $m-1$ vehicles is found in the second population, this population is duplicated and replaces the first population. Then, each solution in the second population is transformed into a solution with one fewer vehicle. This is repeated until no feasible solution can be found in the second population with the given number of vehicles. A parallel master-slave implementation of this algorithm is also reported in [8]. Ombuki et al. [95] from Brock University in Ontario, address the issue of minimizing the number of vehicles versus minimizing the total distance by handling the VRPTW as a true bi-objective problem. Namely, a solution dominates another if it is better for one objective and better or equal for the other objective. Then, they exploit ideas from multi-objective GAs [27] to approximate the Pareto front of the VRPTW.

In the transportation domain at large, B. Abdulhai, who holds the Canada Research Chair in Intelligent Transportation Systems (ITS) at the University of Toronto, used GAs and other adaptive methods to address ITS-related topics, like estimating origin-destination traffic matrices, detecting freeway incidents and optimizing traffic signal control at street intersections [80, 86, 110].

A memetic algorithm where TS is used to perform the local search is reported in [19] for the Generalized Quadratic Assignment Problem (GQAP). Here, an assignment of weighted facilities to sites must be determined so that the sum of assignment costs and traffic costs between facilities is minimized and the total weight of all facilities assigned to the same site does not exceed the site capacity. Subsets of *cohesive* (close) and *non-cohesive* (far) sites are first identified for every site. Then, given two parent solutions and a so-called *pivot* site, the assignment of facilities to the corresponding cohesive and non cohesive sites are taken from the parent with the best evaluation and the other parent, respectively. The solution is completed by randomly assigning the remaining facilities. This procedure is repeated using every site as the pivot site and the best solution is returned. The TS then improves this solution using a neighbourhood structure where facilities are moved from one site to another.

A facility layout problem is addressed with a GA by Delmaire, Langevin and Riopel [28]. In this application, the position and dimension of departments along a corridor (spine) must be determined. A GA first explores the possible orderings of departments along the spine, using order-based crossover and mutation operators, like those developed for the TSP [98]. Each ordering is then given to a linear program that produces a complete solution by locating and determining the dimension of each department. Using a similar approach, Cheung, Langevin and Delmaire develop a GA that explores possible values for integer variables in mixed integer non-linear programming problems [18]. For a given set of values for these integer variables, a grid search algorithm [17] is used to determine the values of the real variables. Finally, a multi-buyer Joint Replenishment Problem (JRP) is solved with a GA in [16]. The JRP is a multi-item inventory problem where policies for the replenishment of a group of items that are jointly ordered from a single supplier are designed to minimize the inventory and setup costs over the planning horizon. When a multi-branch firm has all its branches ordering the same group of items from the supplier, the multi-buyer variant is obtained. In this work, the GA evolves the replenishement cycle times of every item and branch.

4 Conclusion

In spite of the recent advances in the performance of mathematical programming software, realistically-sized instances of many important operations research problems, in particular, combinatorial ones, remain intractable for exact solution approaches and must thus be tackled using approximate solution techniques. Among these, metaheuristics stand out as the family of methods that is the most effective and the most susceptible to deliver near-optimal solutions. Metaheuristics are thus of prime importance for the operations research community in Canada who has to face routinely large instances of difficult problems in application areas such as logistics, transportation, telecommunications networks, manufacturing, forestry, mining, health care and other services. Not surprisingly, the importance of metaheuristics for the operations research community has led to a large number of very significant contributions to this area from Canadian researchers.

Acknowledgments. Financial support for this work was provided by the Canadian Natural Sciences and Engineering Research Council (NSERC) under grants 36662-06 and 38816-05. This support is gratefully acknowledged.

References

[1] Aarts E.H.L. and J.K. Lenstra (eds.), Local Search in Combinatorial Optimization, Wiley, Chichester, 1997.

[2] Aarts E.H.L and H.M.M. Ten Eikelder, "Simulated Annealing", in *Handbook of Applied Optimization*, P.M. Pardalos and M.G.C. Resende (eds.), Oxford University Press, New York, 209–220, 2002.

[3] Alba E. (ed.), Parallel Metaheuristics: A New Class of Algorithms, Wiley, Hoboken, 2005.

[4] Archetti C., A. Hertz and M.G. Speranza, "A tabu search algorithm for the split delivery vehicle routing problem", *Transportation Science 40*, 64–73, 2006.

[5] Bahri A. and S. Chamberland, "On the wireless local area network design problem with performance guarantees", *Computer Networks 48*, 856–866, 2005

[6] Battiti R. and G. Tecchiolli, "The Reactive Tabu Search", *ORSA Journal on Computing 6*, 126–140, 1994.

[7] Berger D., B. Gendron, J.-Y. Potvin, S. Raghavan and P. Soriano, "Tabu Search for a Network Loading Problem with Multiple Facilities", *Journal of Heuristics 6*, 253–268, 2000.

[8] Berger J. and M. Barkaoui, "A Parallel Hybrid Genetic Algorithm for the Vehicle Routing Problem with Time Windows", *Computers & Operations Research 31*, 2037–2053, 2004.

[9] Berger J., M. Barkaoui and O. Bräysy, "A Route-Directed Hybrid Genetic Approach for the Vehicle Routing Problem with Time Windows", *INFOR 41*, 179–194, 2003.

[10] Blodgett D.E., M. Gendreau, F. Guertin, J.-Y. Potvin and R. Séguin, "A Tabu Search Heuristic for Resource Management in Naval Warfare", *Journal of Heuristics 9*, 145–169, 2003.

[11] Burke E.K. and G. Kendall (eds.), Search Methodologies: Introductory Tutorials in Optimization and Decision Support Techniques, Springer, New York, 2005.

[12] Cantú-Paz E., Efficient and Accurate Parallel Genetic Algorithms, Kluwer, Boston, 2001.

[13] Cerny V., "Thermodynamical Approach to the Traveling Salesman Problem: An Efficient Simulation Algorithm", *Journal of Optimization Theory and Applications 45*, 41–51, 1985.

[14] Chamberland S. and B. Sansò, "On the Design Problem of Multitechnology Networks", *INFORMS Journal on Computing 13*, 245–256, 2001.

[15] Chamberland S., B. Sansò and O. Marcotte, "Topological Design of Two-Level Telecommunication Networks with Modular Switches", *Operations Research 48*, 745–760, 2000.

[16] Chan C.K., B.K.S. Cheung, K.S. Bernard and A. Langevin, "Solving Multi-Buyer Joint Replenishment Problems with a Modified Genetic Algorithm", *Transportation Research B37*, 291–299, 2003.

[17] Cheung B.K.S. and A.C.L. Ng, "An Efficient and Reliable Algorithm for Nonsmooth Nonlinear Optimization", *Neural, Parallel & Scientific Computations 3*, 115–128, 1995.

[18] Cheung B.K.S., A. Langevin and H. Delmaire, "Coupling Genetic Algorithm with a Grid Search Method to Solve Mixed Integer Nonlinear Programming Problems ", *Computers and Mathematics with Applications 14*, 13–23, 1997.

[19] Cordeau J.-F., M. Gaudioso, G. Laporte and L. Moccia. "A Memetic Heuristic for the Generalized Quadratic Assignment Problem", *INFORMS Journal on Computing 18*, 433–443, 2006.

[20] Cordeau J.-F., M. Gendreau and G. Laporte, "A Tabu Search Heuristic for Periodic and Multi-Depot Vehicle Routing Problems", *Networks 30*, 105–119, 1997.

[21] Cordeau J.-F. and G. Laporte. "A Tabu Search Algorithm for the Site Dependent Vehicle Routing Problem with Time Windows", *INFOR 39*, 292–298, 2001.

[22] Cordeau J.-F., G. Laporte and A. Mercier. "A Unified Tabu Search Heuristic for Vehicle Routing Problems with Time Windows", *Journal of the Operational Research Society 52*, 928–936, 2001.

[23] Corne D., M. Dorigo and F. Glover (eds.), New Ideas in Optimization, McGraw-Hill, Maidenhead, 1999.

[24] Crainic T.G., M. Gendreau and J.M. Farvolden, "A Simplex-Based Tabu Search Method for Capacitated Network Design", *INFORMS Journal on Computing 12*, 223–236, 2000.

[25] Crainic T.G., M. Gendreau, P. Hansen and N. Mladenović, "Cooperative Parallel Variable neighborhood Search for the p-Median", *Journal of Heuristics 10*, 289–310, 2004.

[26] Crainic T.G., M. Gendreau, P. Soriano and M. Toulouse, "A Tabu Search Procedure for Multicommodity Location/Allocation with Balancing Requirements", *Annals of Operations Research 41*, 359–383, 1993.

[27] Deb K., Multi-Objective Optimization Using Evolutionary Algorithms, Wiley, Chichester, 2001.

[28] Delmaire H., A. Langevin and D. Riopel, "Skeleton-Based Facility Layout Design Using Genetic Algorithms", *Annals of Operations Research 69*, 85–104, 1997.

[29] Desaulniers G., F. Lessard and A. Hadjar, "Tabu Search, Partial Elementary, and Generalized k-Path Inequalities for the Vehicle Routing Problem with Time Windows", Technical Report G-2006-45, GERAD, Montreal, 2006.

[30] Doerner K.F., M. Gendreau, P. Greistorfer, W.J. Gutjahr, R.F. Hartl and M. Reimann (eds.), Metaheuristics - Progress in Complex Systems Optimization, Springer Science+Business Media, New York, 2007.

[31] Dorigo M., Optimization, Learning, and Natural Algorithms, Ph.D. Thesis, Politecnico di Milano, Italy, 1992.

[32] Dorigo M., V. Maniezzo and A. Colorni, "The Ant System: Optimization by a Colony of Cooperating Agents", *IEEE Transactions on Systems, Man and Cybernetics B26*, 29–42, 1996.

[33] Dorigo M. and T. Stützle, Ant Colony Optimization, The MIT Press, Cambridge, MA, 2004.

[34] Dréo J., A. Pétrowski, P. Siarry and É. Taillard, Metaheuristics for Hard Optimization: Methods and Case Studies, Springer, Berlin, 2006.

[35] Dueck G. and T. Scheurer, "Threshold Accepting: A General Purpose Optimization Algorithm", *Journal of Computational Physics 90*, 161-175, 1990.

[36] Dueck G., "New Optimization Heuristics: The Great Deluge Algorithm and the Record-To-Record Travel", *Journal of Computational Physics 104*, 86-92, 1993.

[37] Feo T.A. and M.G.C. Resende, "A probabilistic heuristic for a computationally difficult set covering problem", *Operations Research Letters 8*, 67–71, 1989.

[38] Feo T.A. and M.G.C. Resende, "Greedy Randomized Adaptive Search Procedures", *Journal of Global Optimization 6*, 109–133, 1995.

[39] Fleurent C. and F. Glover, "Improved Constructive Multistart Strategies for the Quadratic Assignment Problem using Adaptive Memory", *INFORMS Journal on Computing 11*, 198–204, 1999.

[40] Gambardella L.-M., É.D. Taillard and G. Agazzi, "MACS-VRPTW: A Multiple Ant Colony System for Vehicle Routing Problems with Time Windows", in *New Ideas in Optimization*, D. Corne, M. Dorigo and F. Glover (eds.), McGraw-Hill, Maidenhead, 63–76, 1999.

[41] Gambardella L.-M., É.D. Taillard and M. Dorigo, "Ant Colonies for the Quadratic Assignment Problem", *Journal of the Operational Research Society 50*, 167–176, 1999.

[42] Gehring H. and J. Homberger, "A Parallel Hybrid Evolutionary Metaheuristic for the Vehicle Routing Problem with Time Windows", In *Proceedings of EUROGEN99*, K. Miettinen, M.M. Mäkelä and J. Toivanen (eds.), Reports of the Department of Mathematical Information Technology, No. A 2/1999, University of Jyväskylä, Finland, 57–64, 1999.

[43] Gendreau M., "Recent Advances in Tabu Search", in *Essays and Surveys in Metaheuristics*, C.C. Ribeiro and P. Hansen (eds.), 369–378, Kluwer, Boston, 2002.

[44] Gendreau M., "An Introduction to Tabu Search", in *Handbook of Metaheuristics*, F. Glover and G. Kochenberger (eds.), Kluwer, Boston, 37–54, 2003.

[45] Gendreau M., F. Guertin, J.-Y. Potvin and É.D. Taillard, "Parallel Tabu Search for Real-Time Vehicle Routing and Dispatching", *Transportation Science 33*, 381–390, 1999.

[46] Gendreau M., A. Hertz and G. Laporte, "New Insertion and Postoptimization Procedures for the Traveling Salesman Problem", *Operations Research 40*, 1086–1094, 1992.

[47] Gendreau M., A. Hertz and G. Laporte, "A Tabu Search Heuristic for the Vehicle Routing Problem", *Management Science 40*, 1276–1290, 1994.

[48] Gendreau M., F. Guertin, J.-Y. Potvin and R. Séguin, "Neighborhood Search Heuristics for a Dynamic Vehicle Dispatching Problem with Pick-ups and Deliveries", *Transportation Research C14*, 157–174, 2006.

[49] Gendreau M., G. Laporte and F. Semet, "A Dynamic Model and Parallel Tabu Search Heuristic for Real-Time Ambulance Relocation", *Parallel Computing 27*, 1641–1653, 2001.

[50] Gendreau M. and J.-Y. Potvin, "Tabu Search", in *Search Methodologies : Introductory Tutorials in Optimization and Decision Support Techniques*, E. Burke and G. Kendall (eds.), Springer, New York, 165–186, 2005.

[51] Gendron B., J.-Y. Potvin and P. Soriano, "Tabu Search with an Exact Neighbor Evaluation for Multicommodity Location with Balancing Requirements", *INFOR 37*, 255–270, 1999.

[52] Gendron B., J.-Y. Potvin and P. Soriano, "Diversification Strategies in Local Search for a Nonbifurcated Network Loading Problem", *European Journal of Operational Research 142*, 231–241, 2002.

[53] Gendron B., J.-Y. Potvin and P. Soriano, "A Tabu Search with Slope Scaling for the Multicommodity Capacitated Location Problem with Balancing Requirements", *Annals of Operations Research 122*, 193–217, 2003.

[54] Ghamlouche I., T.G. Crainic and M. Gendreau, "Cycle-Based Neighborhoods for Fixed-Charge Capacitated Multicommodity Network Design", *Operations Research 51*, 655–667, 2003.

[55] Girard A., B. Sansò and L. Dadjo, "A Tabu Search Algorithm for Access Network Design", *Annals of Operations Research 106*, 229–262, 2001

[56] Glover F., "Future Paths for Integer Programming and Links to Artificial Intelligence", *Computers & Operations Research 13*, 533–549, 1986.

[57] Glover F., "Tabu Search - Part I", *ORSA Journal on Computing 1*, 190–206, 1989.

[58] Glover F., "Tabu Search - Part II", *ORSA Journal on Computing 2*, 4–32, 1990.

[59] Glover F. and G.A. Kochenberger (eds.), Handbook of Metaheuristics, Kluwer, Boston, 2003.

[60] Glover F. and M. Laguna, Tabu Search, Kluwer, Boston, 1997.

[61] Glover F. and M. Laguna, "Tabu Search", in *Handbook of Applied Optimization*, P.M. Pardalos and M.G.C. Resende (eds.), Oxford University Press, New York, 194–208, 2002.

[62] Glover F., M. Laguna and R. Marti, "Fundamentals of Scatter Search and Path Relinking", *Control and Cybernetics 39*, 653–684, 2000.

[63] Goldberg D.E., Genetic Algorithms in Search, Optimization & Machine Learning, Addison-Wesley, Reading, 1989.

[64] Gravel M., W. Price and C. Gagné, "Scheduling Continuous Casting of Aluminum using a Multiple Objective Ant Colony Optimization Metaheuristic", *European Journal of Operational Research 143*, 218–229, 2002.

[65] Hansen P., B. Jaumard, N. Mladenović and A. Parreira, "Variable Neighborhood Search for Weighted Maximum Satisfiability Problem", Les Cahiers du GERAD, G-2000-62, Montreal, 2000.

[66] Hansen P. and N. Mladenović, "Industrial Applications of the Variable Neighborhood Search Metaheuristic", In *Decision and Control in Management Science*, G. Zaccour (ed.), Kluwer, Boston, 261–274, 2001.

[67] Hansen P. and N. Mladenović, "Variable Neighborhood Search", in *Handbook of Metaheuristics*, F. Glover and G. Kochenberger (eds), Kluwer, Boston, 145–184, 2003.

[68] Hansen P. and N. Mladenović, "Variable Neighborhood Search: Principles and Applications", *European Journal of Operational Research 130*, 449–467, 2001.

[69] Hansen P. and N. Mladenović, "Developments of Variable Neighborhood Search", in *Essays and Surveys in Metaheuristics*, C.C. Ribeiro and P. Hansen (eds.), Kluwer, Boston, 415–440, 2002.

[70] Hansen P. and N. Mladenović, "Variable Neighborhood Search", in *Handbook of Applied Optimization*, P.M. Pardalos and M.G.C. Resende (eds.), Oxford University Press, New York, 221–234, 2002.

[71] Hansen P. and N. Mladenović, "Variable Neighborhood Search", in *Search Methodologies: Introductory Tutorials in Optimization and Decision Support Techniques*, E.K. Burke and G. Kendall (eds.), Springer, New York, 211–238, 2005.

[72] Hansen P., N. Mladenović and D. Perez-Brito, "Variable Neighborhood Decomposition Search", *Journal of Heuristics 7*, 335–350, 2001.

[73] Henderson D., S.H. Jacobson and A.W. Johnson, "The Theory and Practice of Simulated Annealing", in *Handbook of Metaheuristics*, F. Glover and G. Kochenberger (eds.), Kluwer, Boston, 287–319, 2003.

[74] Holland J.H., Adaptation in Natural and Artificial Systems, The University of Michigan Press, Ann Arbor, 1975.

[75] Hoos H. and T. Stützle, "MAX-MIN Ant System", *Future Generation Computer Systems 16*, 889–914, 2000.

[76] Hoos H. and T. Stützle, Stochastic Local Search: Foundations and Applications, Morgan Kaufmann, San Francisco, 2005.

[77] Ibaraki T., K. Nonobe and M. Yagiura (eds.), Metaheuristics: Progress as Real Problem Solvers, Springer, New York, 2005.

[78] Ichoua S., M. Gendreau and J.-Y. Potvin , "Diversion Issues in Real-Time Vehicle Dispatching", *Transportation Science 34*, 426–438, 2000.

[79] Ichoua S., M. Gendreau and J.-Y. Potvin, "Exploiting Knowledge about Future Demands for Real-Time Vehicle Dispatching", *Transportation Science 40*, 211–225, 2006.

[80] Kattan, L. and B. Abdulhai, "Non-Iterative Approach to Dynamic Traffic Origin-Destination Estimation Using Parallel Evolutionary Algorithms, *Journal of the Transportation Research Record 1964*, 201–210, 2006.

[81] Kirkpatrick S., C.D. Gelatt Jr. and M.P. Vecchi, "Optimization by Simulated Annealing", *Science 220*, 671–680, 1983.

[82] Laguna M. and R. Martí, "GRASP and path relinking for 2-layer straight line crossing minimization", *INFORMS Journal on Computing 11*, 44–52, 1999.

[83] Laporte G. and I.H. Osman (eds.), Metaheuristics in Combinatorial Optimization, *Annals of Operations Research 63*, 1996.

[84] Le Bouthillier A. and T.G. Crainic, "A Cooperative Parallel Meta-Heuristic for the Vehicle Routing Problem with Time Windows", *Computers & Operations Research 32*, 1685–1708, 2005.

[85] Le Bouthillier A., T.G. Crainic and P. Kropf, "A Guided Cooperative Search for the Vehicle Routing Problem with Time Windows", *IEEE Intelligent Systems 20*, 36–42, 2005.

[86] Lee J., B. Abdulhai, A. Shalaby and E. Chung, "Real-Time Optimization for Adaptive Traffic Signal Control using Genetic Algorithms", *Journal of Intelligent Transportation Systems 9*, 111–122, 2005.

[87] Le Louarn F.-X., M. Gendreau and J.-Y. Potvin, "GENI Ants for the Traveling Salesman Problem", *Annals of Operations Research 131*, 187–201, 2004.

[88] Lopez L., M.W. Carter and M. Gendreau, "The Hot Strip Mill Production Scheduling Problem: A Tabu Search Approach", *European Journal of Operational Research 106*, 317–335, 1998.

[89] Martins S.L., M.G.C. Resende, C.C. Ribeiro and P. Pardalos, "A Parallel GRASP for the Steiner Tree Problem in Graphs using a Hybrid Local Search Strategy", *Journal of Global Optimization 17*, 267–283, 2000.

[90] Michalewicz Z., Genetic Algorithms + Data Structures = Evolution Programs, Third Edition, Springer, Berlin, 1996.

[91] Middendorf M., F. Reischle and H. Schmeck, "Multi Colony Ant Algorithms", *Journal of Heuristics 8*, 305–320, 2002.

[92] Mladenović N. and P. Hansen, "Variable Neighborhood Search", *Computers & Operations Research 24*, 1097–1100, 1997.

[93] Moscato P., "Memetic Algorithms: A short introduction", in *New Ideas in Optimization*, D. Corne, M. Dorigo and F. Glover (eds.), McGraw-Hill, Maidenhead, 219–234, 1999.

[94] Mühlenbein H., "Genetic Algorithms" in *Local Search in Combinatorial Optimization*, E.H.L. Aarts and J.K. Lenstra (eds.), Wiley, Chichester, 137–171, 1997.

[95] Ombuki B., B.J. Ross and F. Hanshar, "Multi-Objective Genetic Algorithms for Vehicle Routing Problem with Time Windows", *Applied Intelligence 24*, 17–30, 2006.

[96] Osman I.H. and J.P. Kelly (eds.), Meta-Heuristics: Theory & Applications, Kluwer, Boston, 1996.

[97] Pardalos P.M. and M.G.C. Resende (eds.), Handbook of Applied Optimization, Chapter 3.6 on Metaheuristics, Oxford University Press, New York, 123–234, 2002.

[98] Potvin J.-Y., "Genetic Algorithms for the Traveling Salesman Problem", *Annals of Operations Research 63*, 339–370, 1996.

[99] Potvin J.-Y. and S. Bengio, "The Vehicle Routing Problem With Time Windows - Part II: Genetic search", *INFORMS Journal on Computing 8*, 165–172, 1996.

[100] Potvin J.-Y. and J.-M. Rousseau, "An Exchange Heuristic for Routing Problems with Time Windows", *Journal of the Operational Research Society 46*, 1433–1446, 1995.

[101] Prais M., C.C. Ribeiro, "Reactive GRASP: An Application to a Matrix Decomposition Problem", *INFORMS Journal on Computing 12*, 164–176, 2000.

[102] Reeves C.R. (ed.), Modern Heuristic Techniques for Combinatorial Problems, Blackwell, Oxford, 1993.

[103] Reeves C.R. and T. Yamada, "Genetic Algorithms, Path Relinking and the Flowshop Sequencing Problem", *Evolutionary Computation 6*, 45–60, 1998.

[104] Rego C. and B. Alidaee, Metaheuristic Optimization via Memory and Evolution: Tabu Search and Scatter Search, Springer, New York, 2005.

[105] Resende M.G.C. and J.P. de Sousa (eds.), Metaheuristics: Computer Decision-Making, Kluwer, Boston, 2004.

[106] Ribeiro C.C. and P. Hansen (eds.), Essays and Surveys in Metaheuristics, Kluwer, Boston, 2002.

[107] Richards E.W. and E.A. Gunn, "A Model and Tabu Search Method to Optimize Stand Harvest and Road Construction Schedules", *Forest Science 46*, 188–203, 2000.

[108] Richards E.W. and E.A. Gunn, "Tabu Search Design for Difficult Forest Management Optimization Problems", *Canadian Journal of Forest Research 33*, 1126–1133, 2003.

[109] Rochat Y. and É.D. Taillard, "Probabilistic Diversification and Intensification in Local Search for Vehicle Routing", *Journal of Heuristics 1*, 147–167, 1995.

[110] Roy P. and B. Abdulhai, "GAID: Genetic Adaptive Incident Detection for Freeways", *Journal of the Transportation Research Record 1856*, 96–106, 2004.

[111] Schrimpf G., J. Schneider, H. Stamm-Wilmbrandt and G. Dueck, "Record Breaking Optimization Results using the Ruin and Recreate Principle", *Journal of Computational Physics 159*, 139–171, 2000.

[112] Shaw P., "Using Constraint Programming and Local Search Methods to Solve Vehicle Routing Problems", *Lecture Notes in Computer Science 1520*, M. Maher and J.-F. Puget (eds.), Springer, New-York, 417–431, 1998.

[113] Shmygelska A. and H.H. Hoos, "An Ant Colony Optimisation Algorithm for the 2D and 3D Hydrophobic Polar Protein Folding Problem", *BMC Bioinformatics 6:30*, 2005.

[114] Solomon M.M., "Time Window Constrained Routing and Scheduling Problems", *Operations Research 35*, 254–265, 1987.

[115] Taillard É.D., P. Badeau, M. Gendreau, F. Guertin and J.-Y. Potvin, "A Tabu Search Heuristic for the Vehicle Routing Problem with Soft Time Windows", *Transportation Science 31*, 170–186, 1997.

[116] Taillard É.D., L.M. Gambardella, M. Gendreau and J.-Y. Potvin, "Adaptive Memory Programming: A Unified View of Metaheuristics", *European Journal of Operational Research 135*, 1–16, 2001.

[117] Voss S., S. Martello, I.H. Osman and C. Roucairol (eds.), Meta-Heuristics: Advances and Trends in Local Search Paradigms for Optimization, Kluwer, Boston, 1999.