Ann Oper Res (2011) 183: 1–5
DOI 10.1007/s10479-010-0821-z

Preface

César Rego

Published online: 21 December 2010

1 Introduction

Adaptive Memory Programming (AMP) adopts the premise that methods for complex optimization problems, particularly those arising in real world applications, can often function more effectively if they incorporate flexible and responsive memory. Accompanying this premise is the corollary that such memory is employed together with strategies expressly designed for exploiting it. More broadly, AMP embodies the following principle. If a problem has exploitable features, but contains a structure sufficiently complex to prevent these features from being known in advance, then a method can derive advantages by monitoring its behavior in relation to the space in which it operates. The purpose of the monitoring is effectively to generate a map of the regions the method has visited as a foundation for modifying its behavior, where this map can take multiple forms that ultimately become expressed in the decision rules employed to negotiate the solution space. The hallmark of an AMP method is therefore a capacity to guide its progress by reference to its own unfolding history. Such a method evidently is implicitly or explicitly structured to employ learning.

This philosophy, which entered the metaheuristic field by way of tabu search, is becoming increasingly embraced by a variety of other metaheuristics. In several settings this has occurred by forming hybrids of other methods with tabu search, but it has also been brought about entirely independently as other metaheuristics, both well-established methods and more recent ones, are coming to adopt a focus that qualifies them as members of the AMP class. The vitality borne of this mix of ideas from multiple realms promises to expand the application of adaptive memory programming and to produce more effective strategies for complex optimization problems.

The papers of this special volume offer a case in point, offering an array of innovations in AMP methods and their applications. New methods for exploiting memory and learning are complemented by new forms of memory and learning to be exploited. A brief preview of the contents of these contributions follows.

C. Rego (✉)
University of Mississippi, Oxford, MS, USA
e-mail: crego@bus.olemiss.edu

2 The Papers

Campos, Piñana, and Marti develop two new AMP metaheuristics, one based on tabu search and one based on scatter search.

Bandwidth minimization is generally applied in practice as a means to solve non-singular systems of linear algebraic equations. Preprocessing the coefficient matrix to reduce its bandwidth results in substantial savings in computational effort associated with solving the system of equations. The context of these applications includes aircraft structures, liquid nitrogen gas tanks, propeller blades, and submarines.

The authors employ three different types of moves in their approach which are defined in terms of frequency labels assigned to vertices of a graph. A memory structure is used for short term memory that implements tabu restrictions for vertices that differ according to the role taken by the vertices in structuring a move. Long term memory is introduced in the context of a restarting strategy.

The scatter search component of the study uses the customary five point design for such procedures, and combines solutions (labelings) by creating a pseudo-labeling based variously on averaging, convex combinations, MinMax criteria, and multi-structuring. New labels are then produced by assigning labels consecutively according to the ordering of the pseudo-labeling.

Computational tests on a wide range of problems show that the resulting tabu search and scatter search methods can compete with the most recent and best algorithms for matrix bandwidth minimization, in both solution quality and running time.

Hanafi and Yanev address the problem of identifying a hyperplane to separate two groups so as to minimize the number of misclassified points (which fall on the "wrong side" of the hyperplane). The authors exploit a *system infeasibility formulation* by a tabu search approach that operates relative to a neighborhood consisting of pivot operations.

Hanafi and Yanev's method employs a collection of tabu search strategies that go beyond those utilized in prior related studies. To select pivots efficiently they make use of an Aspiration Plus candidate list strategy implemented by a partial pivoting operation. The method uses a form of short term tabu memory based on two interacting tabu lists, thus obtaining greater differentiation than would be provided by the customary single list approach. This short term memory is complemented by a form of long term memory based on transition frequencies and residence frequencies. Aspiration criteria used by the method differentiate among solutions with tied costs, responding to the fact that many different bases reached by pivoting can produce such ties.

Intensification is managed in the proposed algorithm according to a geometric interpretation that allows a search for improved solutions to be carried out by modifying input points and control parameters. Counterbalancing this, diversification is achieved by incorporating a penalty function that modifies the choice rules.

Computational tests comparing the method to others yield highly favorable outcomes, demonstrating both the method's robustness and its ability to solve large problems efficiently, independent of the number of misclassifications.

Glover et al. deal with the problem of protecting confidentiality which poses a major challenge within government agencies and commercial organizations that report data. The paper approaches the problem within the context of controlled tabular adjustment (CTA), a protocol for confidentiality protection that has been heralded for its value from a theoretical standpoint, but that presents a significant challenge in practice due to requiring the solution of a difficult zero-one mixed integer programming problem. State-of-the-art integer programming software fails to perform effectively for CTA problems of modest size, taking excessive run time to locate even relatively poor solutions, and failing completely for problems of sizes routinely encountered in practice.

The authors develop a succession of approaches to improving the effectiveness of solving the integer programming formulation, ranging from problem partitioning to a special heuristic ordering method. While gains are obtained, the problem continues to resist solution. A more dramatic step is undertaken that replaces the commercial software by a method based on scatter search (SS). The SS approach succeeds in handling significantly larger problems, but still does not obtain solutions of high quality according to estimates derived from theoretical optimality bounds.

Finally, a new metaheuristic learning method related to statistical experimental design is examined, with the goal of selecting subsets of 0–1 variables estimated to be important candidates for composing optimal solutions (by receiving a value of 1 in these solutions). The resulting method greatly reduced the deviation from optimality bounds, cutting the gap for the largest problem from 117% to 6%.

Such a learning process is conjectured to be useful in other applications of zero-one optimization as a strategy for pinpointing candidate sets of 0–1 variables to belong to optimal solutions.

Blazewicz et al. tackle a long standing challenge from the field of biological analysis by addressing the problem of reconstructing evolutionary relationships of groups of organisms.

Adopting a framework that expresses the evolutionary relationships in the form of phylogenetic trees, the authors use local search algorithms from the class of methods known as character-based parsimony algorithms. Trees are transformed by rules that correspond to defining three neighborhoods of moves: (1) Nearest Neighbor Exchanges, (2) Subtree Pruning and Regrafting, and (3) Tree Bisection and Reconnection. Adaptive memory is added to guide the process, producing three algorithms that integrate the neighborhoods in different ways, embodied in parallel search procedures using a master-slave structure.

The key conclusion of the study is that the use of adaptive memory gives better results than a method based on randomized choices of starting points. The parallelization of the method provides a superlinear speedup compared to the sequential algorithm, disclosing the efficacy of their design, and the methods were additionally verified to yield advantages over existing software packages.

Duarte et al. produce a method for global nonlinear function optimization that combines scatter search and tabu search. The scatter search component starts by generating a diverse set which samples solutions randomly but only accepts a solution if it lies a specified minimum threshold distance from other solutions already in the set. However, instead of selecting solutions one at a time, they select them simultaneously by solving a maximum diversity problem. The resulting set of solutions is then integrated with an elite reference set, so that the full reference set consists of both an elite component and a diversified component. Improvement steps are carried out by line searches and solutions are combined to create points lying both inside and outside the line segments spanned by the chosen solutions.

The tabu search portion of the method uses a short term TS procedure to produce modified line searches, coupled with an implementation of the Nelder-Mead simplex method that is likewise modified by short term memory. Alternative ways of integrating the scatter search and tabu search procedure produce five different algorithms, which are tested on two data sets. The best algorithm, STS (for "Scatter Tabu Search") is then compared to leading algorithms for the global function optimization problem. The performance of the STS method displays a distinctive pattern. As the number of iterations grows, the performance of the STS method improves relative to other methods, as measured by the quality of the average of the best solutions found over multiple trials of each method. Compared to 12 other methods, STS moves from 8th place to 4th place and then to 1st place when 100,000 iterations are executed. Also, when considering the average of all the average best solutions over repeated tests, the STS ranks 3rd over a long term horizon. This behavior accords with the expectation that an effective diversification procedure may not improve the average quality of solutions generated over time, but can be valuable by improving the quality of best solutions found.

Hanafi and Wilbaut provide a new modification of the Soyster, Lev and Slivka (SLS) algorithm for 0–1 integer programming problems. The SLS method works by solving linear programming relaxations and generating restricted subproblems in which a subset of nonbasic variables having largest reduced costs are held fixed at their associated bounds. An integer programming optimization is then performed relative to the remaining free variables. The authors develop a series of propositions that identify useful relationships between constrained versions of successively generated subproblems which are augmented to include additional constraints.

The addition and removal of constraints, although capable of being managed using TS recency and frequency memory, are handled by using memory in a different way. The oldest constraints in this approach, instead of being dropped, are replaced by surrogate constraints, thus continuing to preserve a useful portion of the information contained in the constraints that are discarded. Recency and frequency memory are used to implement intensification and diversification processes, by rules that govern which and how many variables are to be freed.

The resulting algorithm is shown by computational testing to be highly effective. Depending on parameters of testing, the method finds up to 52 new best known solutions to a set of standard benchmark test problems.

Hashimoto et al. address a problem proposed by France Telecom for the 2007 ROADEF challenge. The problem consists of identifying a daily schedule for technicians by assigning them to work groups and selecting which tasks (interventions) the groups will carry out. Considerations include priorities and durations of tasks, together with requirements for specified numbers of technicians with specified skills, and precedence constraints.

The "Greedy Randomized Adaptive Search Procedure" (GRASP) is employed to execute a succession of steps each consisting of a construction phase, an improvement phase and an update phase (to update the current best solution). The authors undertake to extend GRASP to produce an AMP procedure by introducing memory to record and exploit information about choices and their outcomes from previous steps as a means to direct the search toward potentially improving solutions. This is handled by using the recorded information to modify the weights used in evaluating the candidate elements to be selected during the construction phase.

The improving phase of the approach uses a critical path method and a packing method, employing swap and insertion moves and restricting attention to moves that are feasible. The critical path phase undertakes to decrease the ending times of each priority and of the entire schedule simultaneously, while the packing phase undertakes to schedule interventions more efficiently without increasing the ending time of any priority element.

The authors generate optimality bounds that confirm the quality of the solutions produced. The bounds were also instrumental in an analysis to suggest where possible improvements may be achieved in future versions of the method.

Hard problems for metaheuristic search can be a source of insight for developing better methods. **Glover and Hao** examine a challenging instance of such a problem that possesses just two local optima, one the global optimum and the other a somewhat inferior solution that lies far away from the global optimum, separated by a maximum possible number of 0–1 flip (complementation) moves in the feasible space. Nevertheless, starting from any feasible solution except the worst one of all, every possible sequence of descent moves leads to the inferior local optimum. Once reaching this local optimum, all sequences of feasible flips to reach the global optimum must ultimately pass through solutions that are progressively worse until reaching the worst solution of all, which is adjacent to the global optimum.

This raises the question of what kind of method can prove successful in solving the problem efficiently. Moves that flip multiple variables simultaneously are likewise stymied by the problem structure, offering negligible improvement in the number of steps required for any reasonable number of simultaneous flips, and entailing greatly increased effort. Alternative neighborhoods are identified that can locate the global more readily, but each of these approaches encounters serious difficulties upon slightly changing the problem formulation. Other possible approaches that seem at first to be promising turn out to have deficiencies.

These difficulties are shown to be capable of resolution by a strategic oscillation approach which takes the form of inducing transitions between feasible and infeasible space. This approach provides a useful connection to AMP strategies by constituting one of the initial forms of oscillation proposed in tabu search. The analysis concludes by noting connections with recent results on the structure of feasible and infeasible trajectories, and by identifying more general ways for implementing such a strategic oscillation that invite further study in the context of adaptive memory programming.

3 Conclusion

In sum, the contributions appearing this special volume disclose some of the useful advances provided by AMP methods along with identifying considerations that may form the foundation for future advances. The successes of these approaches motivate a fuller exploration of the potential for applying AMP methods in additional contexts, giving special consideration to processes for inferring and exploiting structure made possible by the use of adaptive memory.