Black Box Scatter Search for General Classes of Binary Optimization Problems

Francisco Gortazar, Abraham Duarte

Departamento de Ciencias de la Computación
Universidad Rey Juan Carlos, Spain
francisco.gortazar@urjc.es , abraham.duarte@urjc.es

Manuel Laguna

Leeds School of Business
University of Colorado at Boulder, USA
laguna@colorado.edu

Rafael Martí

Departamento de Estadística e Investigación Operativa
Universitat de València, Spain
rafael.marti@uv.es

Abstract

The purpose of this paper is to apply the scatter search methodology to general classes of binary problems. We focus on optimization problems for which the solutions are represented as binary vectors and that may or may not include constraints. Binary problems arise in a variety of settings, including engineering design and statistical mechanics (e.g., the spin glass problem). A distinction is made between two sets of general constraint types that are handled directly by the solver and other constraints that are addressed via penalty functions. In both cases, however, the heuristic treats the objective function evaluation as a black box. We perform computational experiments with four well-known binary optimization problems to study the efficiency (speed) and effectiveness (solution quality) of the proposed method. Comparisons are made against both commercial software and specialized procedures on a set of 376 instances. We chose commercial software that is similar in nature to the proposed procedure, namely, it treats the objective function as a black box and the search is based on evolutionary optimization techniques.

Keywords

Optimization, Metaheuristics, Hard optimization problems

1. Introduction

Black box optimizers have a long tradition in the field of operations research. These procedures treat the objective function evaluation as a black box and therefore do not take advantage of its specific structure. Black box optimizers have also been referred to as context-independent procedures. However, the context-independent notion is more difficult to define because no solver is totally independent from the context. In fact, it can be argued that solvers are developed within a spectrum that ranges from almost no dependence on context to total dependence on context. Knowledge about the context may be divided between the objective function and the set constraints. For instance, some solvers may not have information about the structure of the objective function but have information regarding the feasibility region as defined by a set of constraints. In the case of mathematical programming approaches, such as linear programming, the solvers have a very specific structure that they exploit (even though they "don't know", for instance, if they are solving a transportation or an aggregate production planning problem). The structure — and therefore context dependency — is given by the formulation and not by the unknown (to the solver) real world context.

Nonlinear optimization approaches that do not use derivatives (or estimate them) to find search directions — such as Nelder and Mead or Powell — may be considered as being at the high level of the context-independence range because they treat the objective function as a black box. Methods that estimate derivatives — like those based on generalized reduced gradients (GRG) — assume objective function smoothness and that assumption alone moves them closer to the context-dependence end of the spectrum. However, we can't ignore the fact that some of these procedures (for example the standard version of Microsoft Excel's Solver) are routinely used to search for solutions to problems that do not meet the "smoothness" requirement. The Solver acts as a general purpose procedure in the sense that it does not require that the user provides any context. In the experimental section, we compare our proposed procedure with the Evolutionary Premium Solver (SDK callable library), which implements advanced search strategies for black box optimization.

A similar philosophy is used by the general-purpose commercial optimization software known as OptQuest (by OptTek Systems, Inc.). This software operates by treating the objective function evaluation as a black box. However, OptQuest is not totally "ignorant" of the context given that a solution representation must be chosen in order to run this solver. OptQuest allows users to represent solutions as a mixture of continuous, discrete, integer, binary, permutation and other specialized variables (project, categorical or enumeration). Clearly, the solution representation gives OptQuest some information about the problem context and therefore the solver context-independence changes with each particular application because the amount of information that it receives varies. The software chooses solvers based on the characteristics of the optimization model: Pure or mixed, Constrained or unconstrained and Deterministic or stochastic.

The first characteristic refers to the solution representation. That is, if only one type of variables (e.g., continuous or discrete) is used, then the problem is pure, otherwise the problem is mixed. We consider both constrained and unconstrained problems in the realm of deterministic optimization.

Evolver by Palisade Corp. is a popular black-box optimizer based on genetic algorithms. When coupled with @Risk, a module for risk analysis, Evolver can be used to search for solutions to problems for which the objective function is the result of a Monte Carlo simulation. Evolver can be used as a callable library in the Evolver Development Kit or within Microsoft Excel as its modeling environment, where decision variables are declared as "adjusting cells." The use of specialized problem solving methods, such as "recipe", "order" or "budget", is encouraged because they give the solution process additional information that could result in improved outcomes and thus move the solution process closer to the context-dependent end of the spectrum. We compare our proposed procedure with the Evolver Development Kit, Solver and OPtQuest, selecting the appropriate solving method to tackle binary problems.

In addition to general purpose (and commercially available) software for optimizing unknown objective functions with discrete variables, the literature contains some examples of procedures that have been designed with some degree of context-independence. For instance, Rosen and Harmonosky (2005) proposed a simulated annealing that uses an adaptation of the response surface methodology to optimize simulations with a large number of discrete decision variables. Guikema, Davidson and Çağnan (2004) also address the problem of optimizing simulations when the input variables are discrete. Their method is based on using a standard genetic algorithm aided by ridge regression to limit the number of calls to the simulator.

Holland's (1975) genetic algorithms proposal was in fact a black-box optimizer that used an array of bits as the generic representation. The proposed procedure did not include local search and the standard genetic operators (such as single-point crossover) were not linked to the problem context. As GAs became more popular and researchers and practitioner applied them to many hard optimization problems, the context-independent nature of the original proposal began to vanish when improved outcomes were obtained by the addition of problem structure. Hence, most modern GA implementations are hybridized (e.g., coupled with local searches) and incorporate domain-specific knowledge into the search process.

Cross Entropy (Rubinstein and Kroese, 2004) is a more recent method for optimizing black boxes that originated from the need of estimating very small probabilities in rare-event simulation. Respectable performance has been achieved by implementations of the cross entropy method for some difficult combinatorial problems. However, as shown by Laguna, Duarte and Martí (2007) in the context of the max-cut problem, cross entropy also benefits from hybridization and it is uncertain for how long the method will remain true to its origins as a black box procedure.

The remaining of the article discusses our development and testing of a black box procedure for constrained and unconstrained optimization problems whose solutions are represented as binary strings. The scatter search methodology is used as the basic framework for our solution procedure. Although the scatter search philosophy is to take advantage of domain-specific knowledge, this methodology has been used in the past as the basis for black box procedures (e.g., for OptQuest and for the procedure to tackle permutation problems developed by Campos, Laguna and Martí, 2005).

The main contribution of our work is the development or adaptation of scatter search methods that are helpful in the solution of binary problems for which the objective

function is treated as a black-box evaluator. In particular, we developed three diversification generation methods for binary solutions (Section 2.2), an improvement method based on two neighborhoods (Section 2.3) and seven combination methods (Section 2.4). We also adapted to our binary problems a reactive strategy within the combination methods, which was originally proposed by Campos, et al. (2005) for permutation problems. Finally, our procedure includes a strategy for selectively applying the improvement method in order to reduce the computational effort associated with the black-box evaluator.

2. Scatter Search Procedure

Scatter Search (SS) is a metaheuristic that explores solution spaces by evolving a set of reference points. It consists of five methods and their associated strategies. Three of them, the Diversification Generation, the Improvement and the Combination Methods, are typically problem-dependent and are designed for the problem or the class of problems being solved. The other two, the Reference Set Update and the Subset Generation Methods are context independent by design, and standard implementations are available (Laguna and Martí 2003).

2.1 Classes of Problems

We have adapted SS to develop a black box solver for optimization problems with binary variables. We assume that a problem instance consists of finding a set of values for $x = (x_1, x_2, \ldots, x_n)$ — where $x_i = 0$ or 1 — in order to maximize an unknown objective function. The user has the choice of specifying whether the problem is unconstrained or constrained. For *unconstrained* problems, any binary vector of n elements is a feasible solution. If the problem contains constraints the user may choose to transform it into an unconstrained problem by constructing a penalty function whose values are returned by the black box linked to the optimizer. The solver, however, is capable of directly dealing with two general classes of constraints: *multiple choice* and *budget.* Defining these constraints as part of the input to the optimizer adds a level of context-dependency that may result in improved outcomes (just as OptQuest or Evolver benefit from additional specificity in the choice of variable types or input constraints).

Multiple-choice problems are such that k items must be chosen from a total of n. This translates into formulating a multiple-choice constraint that forces that exactly k variables take on the value 1. In mathematical terms:

$$\sum_{i=1}^{n} x_i = k \tag{1}$$

The user specifies the value of k and the solver limits the search to solutions with exactly k variables set to 1.

Budget problems are those containing constraints that limit the amount of resources used by a given solution. In this case, resource utilization increases with the number of variables that are set to the value of one. Infeasible solutions occur when the available resources are exceeded. The mathematical form of this constraint may vary because resource utilization may be linear or nonlinear with respect to the variables that take on

the value of 1. In the case of the knapsack problem, for instance, the constraint has the following mathematical form:

$$\sum_{i=1}^{n} a_i x_i \leq b \tag{2}$$

The budget is given by the value of b and the a_i coefficients indicate the amount of resources needed if the i^{th} option is selected. In our implementation, the solution evaluator returns a Boolean value indicating whether the solution is feasible or not. In other words, the solution method does not know the level of infeasibility or what specific variables need to be given a value of zero to make the solution feasible. The solution method, however, does know that for this type of problems, feasibility may be eventually achieved by switching variable values from 1 to 0.

As mentioned above, if the problem contains constraints that cannot be casted as multiple-choice or budget, these constraints must be handled by the black-box by way of a penalty function. The penalty function should be such that the objective function value of any infeasible solution must be worse than the objective function value of any feasible solution. The specific form of the penalty function is the responsibility of the user and a number of alternatives have been studied in connection with evolutionary algorithms (Yeniay, 2005). In the remainder of the paper, we refer to constrained problems as those containing multiple choice or budget constraints (as specified by the user). Problems with general constraints are considered unconstrained and the solution method focuses on optimizing the penalized objective function.

2.2 Diversification Generation Method

The search starts with the application of a diversification generation method, as shown in Figure 1, that results in a population P of *PSize* points from which a subset is selected as the initial reference set (*RefSet*). To obtain solutions with different structures we apply three different generators of binary vectors and create *PSize*/3 solutions with each one of them.

The first one, G1, proposed by Glover (1998) uses a systematic approach to create a diverse set of binary vectors. It generates a collection of solutions associated with an integer $h = 2, 3,..., h_{max}$, where $h_{max} \leq n - 1$. From a seed solution x, we generate solutions, x', for each value of h, by the following rule: $x'_{1+kh} = 1 - x_{1+kh}$ for $k = 0, 1, 2, 3, ..., \lfloor n/h \rfloor$, where $\lfloor k \rfloor$ is the largest integer satisfying $k \leq n/h$. All other components of x' are equal to x.

We apply this generator directly in the case of unconstrained problems and slightly modified in the case of constrained problems. For multiple-choice problems, the modification consists of an early termination of the method when the number of 1s in the solution matches k. For budget problems, the process stops when setting the value of a variable to one makes the solution infeasible.

---

1. Start with *GlobalIter*=0 and $P = \emptyset$. Use the <u>diversification generation</u> method to construct a solution. Let *x* be the resulting solution. If $x \notin P$ then add *x* to *P* (i.e., $P = P \cup \{x\}$), otherwise, discard *x*. Repeat this step until $|P| = PSize$.
2. Use the <u>reference set update</u> method to build *RefSet* = { x^1, …, x^b } with *b* solutions from *P*. Order the solutions in *RefSet* according to their objective function value such that x^1 is the best solution and x^b the worst. Make *NewSolutions* = TRUE.
3. Apply the <u>improvement</u> method to the *b*/2 best solutions in *RefSet*.

while (*NewSolutions*) **do**
    4. *Pool* = $\emptyset$
    5. Generate *NewSubsets* with the <u>subset generation</u> method. Make *NewSolutions* = FALSE.
    while (*NewSubsets* $\neq \emptyset$) **do**
        6. Select the next subset *s* in *NewSubsets*.
        7. Apply the <u>solution combination</u> method to *s* to obtain one new trial solution *x*.
        8   If $x \notin Pool$ then add *x* to *Pool* (i.e., $Pool = Pool \cup \{x\}$)
        9. Delete *s* from *NewSubsets*.
    end while
    10. Apply the <u>improvement</u> method to the *b*/2 best solutions in *Pool*. Replace the original solutions in *Pool* with the improved ones.
    11. Apply the <u>reference set update</u> method. Update the *RefSet* with the *b* best solutions in *RefSet*$\cup$*Pool*
    if (*RefSet* has changed) **then**
        12. Make *NewSolutions* = TRUE.
    else
        13. Apply the <u>reference set update</u> method. Rebuild the *RefSet* replacing the worst *b*/2 solutions with new diverse solutions from *P*.
        14. Make *NewSolutions* = TRUE.
    end if
    if (*ExecutionTime* > MAXTIME) **then**
        15. STOP. Return the best solution in *RefSet* as the output of the method.
    end if
end while

Figure 1. Outline of the scatter search procedure

Once we generate *PSize*/3 solutions with the method above, which focuses on diversification and not on the quality of the resulting solutions, we compute *score*(*i*) for each variable x_i to estimate its contribution to the objective function value in order to generate the remaining 2*PSize*/3 solutions considering both quality and diversity. The score calculation has the following mathematical form, where $\bar{f}_A$ indicates the average of the objective function *f* over the solutions in *A*:

$$score(i) = \frac{\bar{f}_{P_i^1}}{\bar{f}_{P_i^1} + \bar{f}_{P_i^0}} \quad \text{where} \quad P_i^1 = \left\{ x \in P : x_i = 1 \right\}, \ P_i^0 = \left\{ x \in P : x_i = 0 \right\} \quad (3)$$

The second generator, G2, constructs a solution step by step, starting with all variables set to 0, and switching in each step one variable from 0 to 1. The variables are randomly selected and the probability of changing a selected variable from 0 to 1 is given by

$$Prob(x_i = 1) = \min\ (0.1 + score(i)\ ,\ 1) \quad (4)$$

The addition of 0.1 to the computation of the probability reflects a bias to change the variable value from 0 to 1 (this is an arbitrary value and any other "small" value would

work as well). For unconstrained problems, G2 performs steps as long as the solution under construction improves. For multiple-choice problems, G2 stops after k variables change their values from 0 to 1. For budget problems, it stops when the solution becomes infeasible (i.e., G2 changes variables from 0 to 1 as long as the solution remains feasible).

The third generator, G3, can be viewed as a destructive method. It was suggested by Glover (1998) and adapted by Duarte and Martí (2007). It starts with all variables set to 1 and at each step it switches the value of one variable from 1 to 0. Variables are probabilistically selected according to their *score* values. The complement of the scores (i.e., $1 - Prob(x_i = 1)$) is used to calculate the probability of changing the value from 1 to 0. The stopping criterion is customized for each type of problem in a similar way as in G2.

After each construction the *score* values are updated. We have empirically found that better results are obtained when the contribution of the last constructed solution to the *score* value is smoothed. The smoothed score, *sscore*, is computed as a function of the score in the previous constructions and the current score. Let $score_t(i)$ be the score of variable i after the t-th construction, and let $score_{t-1}(i)$ be its score in step t-1 (before the construction), both computed with expression (3). Then, we compute the smoothed score, $sscore(i)$, as:

$$sscore(i) = \alpha \, score_{t-1}(i) + (1-\alpha) \, score_t(i) \tag{5}$$

where the parameter α controls the contribution of the current score relative to its previous value. Then, we use *sscore* instead of *score* to generate the values of the variables in the next constructions. We point out that the smoothing of the scores is similar to the way in which probability values are updated in the cross entropy method (Rubinstein and Kroese, 2004).

The initial *RefSet* must balance solution quality and diversity, and we follow the standard reference set update method, selecting the best $b/2$ solutions (where $b = |RefSet|$) from P and then the $b/2$ solutions in P that are most diverse with respect to those already in the *RefSet* (computing the maximum of the minimum distances between each solution and the solutions already in *RefSet* as it is customary in scatter search). Diversity is measured according to the Hamming distance between solutions.

In the standard SS design (Laguna and Martí 2003) the improvement method is applied to all the solutions in P. However, in context-independent solvers, local search methods are usually extremely time-consuming since every trial move must be evaluated by invoking the black-box evaluator. Therefore, the improvement method in our procedure is applied selectively instead of across the board.

The improvement method consists of a local search procedure, which we describe in Subsection 2.3. It is applied to the best $b/2$ solutions in the *RefSet* (see step 3 in the pseudo-code of Figure 1). The reference set is a collection of b solutions that are used to generate new solutions by way of applying a combination method. In order to design a context-independent methodology that performs well across a wide collection of different binary problems, we propose — in Subsection 2.4 — a set of seven combination methods from which one is probabilistically selected according to its

performance in previous iterations. The selection process is reactive, as described below.

We follow the standard subset generation method and, in step 5 of Figure 1, the procedure generates all pairs of reference solutions that have not been combined before. Here again, we limit the application of the improvement method to promising solutions. Specifically, we store the trial solutions resulting from the application of the combination method in a temporary *Pool*, and we apply the improvement method to the $b/2$ best solutions in *Pool*. The reference set is then updated (see steps 10 and 11 in Figure 1) by selecting the best b solutions from the union of the *RefSet* and *Pool* (where the improved solutions replaced the original ones in *Pool*). If the *RefSet* changes (i.e. a new trial solution in *Pool* improves upon the worst solution in the *RefSet*), then step 12 sets the Boolean variable *NewSolutions* equal to TRUE and performs a new iteration in the outer while loop, applying again the subset generation method. Alternatively, if the combination method is incapable of creating solutions that can be admitted to the *RefSet*, the reference set is rebuilt in step 14. The rebuilding consists of keeping the best $b/2$ solutions intact and replacing the worst $b/2$ solutions in the *RefSet* with new diverse solutions from *P*. The method stops after a pre-established *MaxTime* CPU seconds.

2.3. Improvement Method

A global iteration of the improvement method consists of three steps. First, we construct the candidate list CL of elements (variables) to be changed, which simply consists of all the variables in the problem. In the second step we order the elements (variables) in CL according to their *score* value (where those with the largest score are located first) and then, in the third step, we scan them in this order in search for improving moves. We perform several iterations in the third step alternating *flip* (changing the value of a single variable from 1 to 0 or from 0 to 1) and *swap* (exchanging the values of two different variables) moves, where a swap consists of exchanging the values of two different variables. In the first iteration we try *flip* moves with all the elements in CL (examined in order), changing their value if it improves the objective function. The CL is reconstructed at this point. Then, in the second iteration we consider *swap* moves with all the elements in CL (examined in order), in which we try to exchange the value of each variable with the value of another variable. To do this we implement a *first* strategy, which scans the list of variables in search for the first one whose movement results in a strictly positive change of the objective function value. After this process, the CL is reconstructed again. Further iterations are performed alternating between *flip* and *swap* moves. The method stops after *MaxImpIter* iterations or before, if no improvements are achieved after trying all *flips* and *swaps*.

It must be noted that in constrained problems these moves may lead to an infeasible solution. However, we filter the moves and only consider those that lead to feasible solutions. Specifically, when a move produces an infeasible solution, we discard the move and examine the next move in the exploration.

2.4. Reactive Combination Method

As mentioned earlier, we propose seven different combination methods and a reactive mechanism that probabilistically selects among them according to their ability to

produce high quality solutions during the current search. This method was successfully applied by Campos et al. (2005) in the context of permutation problems.

Solutions in the *RefSet* are ordered according to their objective-function value (where the best solution occupies the first place). When a solution obtained with a combination method CM$_i$ qualifies to be the j^{th} member of the current *RefSet*, we add b-j+1 to *success*(CM$_i$). Therefore, combination methods that generate good solutions accumulate higher *success* values and increase proportionally their probability of being selected. To avoid initial biases, this mechanism is activated after the first *InitIter* combinations, and before this, selections are made completely at random. A description of the seven combination methods follows. These methods generate the new trial solution z from the combination of two reference solutions x and y. It should be mentioned that the *score* value, initially computed with the solutions in *P*, is updated during the entire search process (i.e. computed considering not only the solutions in *P* but all the solutions examined during the search so far), thus providing an estimation of the contribution of each variable (when it takes on the value of 1) to the objective function value.

The **CM$_1$ combination method** first computes the partial solution z formed with the "union" of x and y, that is

$$z_i = 1 \ \text{ if } x_i = 1 \text{ or } y_i = 1$$
$$z_i = 0 \ \text{ otherwise}$$

It then performs a series of steps switching some variables from 1 to 0 in a similar way as the G3 generator. Specifically, at each step, the method uses the score values to select probabilistically one variable with the value of 1 to switch it to 0. For unconstrained problems, it continues in this fashion while the solution improves; for multiple-choice problems the method performs steps until the number of variables set to 1 matches k. For budget problems, it performs steps until the solution becomes feasible. Note that for this type of problems, the infeasibility is always reduced by switching variable values from 1 to 0. The **CM$_2$ combination method** is similar to **CM$_1$** with the only difference that the variable selection (to switch the value from 1 to 0 in z) is performed completely at random, that is, without considering the *score* values.

The **CM$_3$ combination method** is an adaptation of the one proposed by Laguna and Martí (2003) in the context of the knapsack problem. The method calculates a *weight* for each variable, based on the objective function value of the two reference solutions being combined. The *weight* for variable i that corresponds to the combination of reference solutions x and y is calculated with the following formula:

$$weight(i) = \frac{f(x)x_i + f(y)y_i}{f(x) + f(y)} \tag{6}$$

where $f(x)$ is the objective function value of solution x and x_i is the value of the i^{th} variable. Then, the trial solution z is constructed by using the *weight* as the probability for setting each variable to one, i.e., Prob($z_i = 1$) = $weight(i)$. Note that this combination method assumes that the objective function is being maximized.

In constrained problems a test is included to ensure the feasibility of the resulting solution. In the case of *multiple-choice problems*, the method stops when k variables in z have been set to 1. If less than k variables are set to 1, the resulting infeasible solution is discarded. In the case of the *budget problems* the method stops when the resource limit b is reached (and the selection of an extra variable would violate it).

The **CM₄ combination method** first computes the partial solution z formed with the "intersection" of x and y:

$$z_i = 1 \quad \text{if } x_i = 1 \text{ and } y_i = 1$$
$$z_i = 0 \quad \text{otherwise}$$

A number of steps are then performed in which a variable with the value of zero is chosen. The probability of selecting variable i is proportional to *weight*(i). (Note that the larger the weight the more attractive it is, at least from the history of the search, to set the value of the variable to one.) The difference between **CM₃** and **CM₄** is that **CM₃** starts the process of switching variable values from 0 to 1 with all the variables set to 0 while **CM₄** starts with the partial solution z that represents the intersection of the reference solutions being combined. Also, the value of r in **CM₃** is a random number between 0 and 1 while for **CM₄** it is set to 0.5. That is, **CM₄** uses simple rounding to set the value of the selected variable. **CM₄** stops in the same way as **CM₁**, depending on the type of problem.

The **CM₅ combination method** is very similar to **CM₄**. Starting from the partial solution constructed with the intersection of the reference solutions, variables that are set to 0 are chosen to change their value to 1. The only difference with **CM₄** is that variables with a value of 0 are selected completely at random, that is, ignoring their *weight* values. This method is the "fixed crossover" developed by Dolezal et al. (1999) for their genetic algorithm designed to tackle max-cut problems.

The **CM₆ combination method** starts with all the variable values set to 0. Then, it applies the G2 constructive method with the restriction that the z_i variables that are candidates to be switched to 1 are those with a value of 1 in x or y (i.e., those for which $x_i + y_i \geq 1$).

The **CM₇ combination method** is based on the path relinking methodology adapted for GRASP (Laguna and Martí 1999). It generates new solutions by exploring trajectories that connect high-quality solutions — by starting from one of these solutions, called an *initiating solution*, and generating a path in the neighborhood space that leads toward the other solutions, called *guiding solutions*. The **CM₇** combination method explores the path between two solutions x and y in the *RefSet*. It starts with the first solution x, and gradually transforms it into the guiding solution y, by changing the value of the variables in x with their value in y. If the value is the same in both solutions, then no change is made and the procedure moves to the next variable. Our procedure examines the variables in lexicographical order and in at most n steps it reaches y. The procedure uses a first-improving strategy, meaning that if during the relinking process it finds an intermediate solution that is better than either x or y, then the procedure stops. If no better solution is found, the solution that is most distant from x and y is the combined solution resulting from the application of this method. In addition to exploring the path

from x to y, the procedure also constructs the path from y to x and chooses the best solution found during both processes to be the outcome of **CM$_7$**.

3. Optimization Problems used for Testing

We have used four combinatorial optimization problems to test our procedure. Solutions to these problems are naturally represented as binary vectors:

1. the max-cut problem
2. the maximum diversity problem
3. the knapsack problem
4. the multi-demand multi-dimensional knapsack problem

We target these problems because they are well known, they are different in nature, and optimal (or high-quality) solutions to several problem instances are readily available. Existing methods to solve these problems range from construction heuristics and metaheuristics to exact procedures. We now provide a brief description of each problem class.

The **max-cut** problem consists of finding a partition of the nodes of a weighted graph into two subsets such that the sum of the weights on the edges connecting the two subsets is maximized. A solution to this problem can be represented as a binary vector with cardinality equal to the number of nodes in the graph (where the value 0 or 1 indicates that the associated node belongs to one or other subset). The max-cut problem falls within the class of binary unconstrained problems.

Beginning with the simple approach introduced by Sahni and Gonzales (1976) different heuristics and metaheuristics have been proposed for this problem. Recently Festa et al. (2002) developed six different algorithms based on the variable neighborhood search, GRASP and Path Relinking (PR) methodologies. We will compare our method with the SS algorithm (Martí et al. 2009) that was shown to outperform existing methods.

The **maximum diversity** problem (MDP) consists of selecting a subset of k elements from a set of n elements in such a way that the sum of the distances between the chosen elements is maximized. Clearly, a solution to this problem can be represented as a binary string x, where variable x_i takes on the value of 1 if element i is selected and 0 otherwise, $i = 1, ..,n$. There is only one constraint in this problem that forces that exactly k variables in the string are assigned the value of 1. This problem belongs to the class of multiple choice problems. Many different methods have been proposed to search for solutions to the MDP. Recently, Silva et al. (2004) presented two GRASP approaches, Duarte and Martí (2007) introduced a tabu search and Gallego et al. (2008) a scatter search method. We will use the latter method in our comparison since it has been shown to outperform previous approaches.

Knapsack problems are well known in the operations research literature. The problem consists of choosing, from a set of items, the subset that maximizes the value of the objective function subject to a capacity constraint. Mathematically, the problem can be expressed as follows:

$$\text{Maximize} \qquad \sum_i c_i x_i$$

$$\text{Subject to} \qquad \sum_i a_i x_i \leq b$$

$$x_i \in \{0,1\} \qquad \forall i$$

This is a budget-constrained problem that attempts to maximize the benefit associated with selecting a subset of n available objects. For a comprehensive examination of the knapsack and other related problems we refer the reader to Martello and Toth (1990). Pisinger (1995) considered a three step method based on the well known ratios c_i/a_i and forward and backward computations. We will use this specialized method in our experiments and compare results against optimal solutions.

The 0/1 **multi-demand multi-dimensional knapsack** problem (MDMKP) represents a class of practical problems, including portfolio selection, capital budgeting and some facility location problems. Mathematically, the problem can be formulated as follows, where the m knapsack constraints are followed by the q demand constraints:

$$\text{Maximize} \qquad \sum_i c_i x_i$$

$$\text{Subject to} \qquad \sum_j a_{ij} x_j \leq b_i \quad \forall i \in \{1,...,m\}$$

$$\sum_j a_{ij} x_j \geq b_i \quad \forall i \in \{m+1,...,m+q\}$$

$$x_i \in \{0,1\} \qquad \forall i$$

It must be noted that the MDMKP represents a class of binary problems with general constraints for which finding feasible solutions may not be trivial. Arntzen, et al. (2006) proposed an adaptive memory programming algorithm for this problem. Their method is based on a short-term tabu search coupled with strategic oscillation around the feasibility boundaries. This method outperforms previous heuristics due to Cappanera and Trubian (2005) and therefore we will include it in our computational testing. Cappanera and Trubian (2005) noted that the intriguing combinatorial structure of the MDMKP makes it a challenging problem for commercial integer linear programming solvers.

4. Experimental Results

This section describes the computational experiments that we have performed to test the efficiency of our black box scatter search procedure as well as comparing it with various methods from the literature. We have implemented the methods in Java SE 6 and all the experiments were conducted on a Pentium 4 computer at 3 GHz with 2 GB of RAM. We have employed four sets of instances in our experimentation (available at http://www.uv.es/rmarti):

Max-cut: This data set consists of 94 instances from two sources. The *Hartmann* problems are 10 instances with 125 nodes and 375 edges all with weight values equal to -1 or 1. The second set consists of 84 instances (n=50 to 300) generated with *rudy,* a machine independent

graph generator by Giovanni Rinaldi. The toroidal, planar and random graphs have weights taking the values of -1, 0, or 1.

MDP: This data set consists of 92 instances from two sources. The first one with 40 instances, referred to as *Glover*, for which the values are calculated as the Euclidean distances from randomly generated points with coordinates in the 0 to 100 range. The second one with 52 instances, referred to as *Silva*, for which the values are integer numbers randomly generated between 50 and 100. The value of *n* ranges from 50 to 300 and the value of *k* ranges from 0.1*n* to 0.3*n*.

Knapsack: This data set consists of 96 instances from four types, as documented in Pisinger (1994): *strongly correlated, subset sum, uncorrelated* and *weakly correlated* instances. We consider 24 instances in each group with *n*= 100, 300, 1000 and 3000 (6 instance for each *n* value).

MDMKP: This data set consists of 94 instances from Cappanera and Trubian (2005) in which basic multi-knapsack instances were modified by adding demand constraints and also allowing for negative cost coefficients. We consider 47 instances with positive coefficients and 47 with both positive and negative coefficient values. The value of *n* ranges from 50 to 250, the value of *m* ranges from 5 to 15 and the value of *q* ranges from *m*/2 to *m*.

We first compute, for each instance, the overall best solution value, *BestValue*. The *BestValue* may be the optimal value if, as in the case of the knapsack problem instances, it is known. Otherwise, *BestValue* represents the best-known as found by the execution of all methods under consideration. Then, in each experiment, we compute for each method the relative percent deviation (RPD) between the best solution value (*Value*) obtained with a particualr method and *BestValue* for that instance. That is,

$$RPD = \frac{\left(BestValue - Value\right)}{BestValue} \cdot 100\%$$

We report the average *RPD* across all the instances considered in each particular experiment. We also report, for each method, the number of instances (*#Best*) for which the value of the best solution obtained with this method matches *BestValue*. In the preliminary experiments, however, *#Best* refers to the solutions found within the experiment since simplified versions of the method are not expected to match best-known solutions. In addition, we calculate the *Rank* statistic — proposed by Ribeiro, et al. (2002) — associated with each method. For each instance, the *n_rank* of a method M is defined as the number of methods that found a better solution than the one found by M. In case of ties, all the methods receive the same *n_rank*, equal to the number of methods strictly better than all of them. The value of *Rank* is the sum of the *n_rank* values for all the instances in the experiment, thus, the lower the *Rank* the better the method.

In our preliminary experimentation we employed a set of 138 representative instances that included 42 *Max-cut* (*rudy* instances with $n \geq 200$), 48 *MaxDiv* (*Glover* and *Silva* instances with $n \geq 125$) and 48 *Knapsack* instances (from the four sources with

$n \geq 1000$). We did not include *MDMKP* instances in this preliminary experimentation because we report the results obtained with our method when solving this problem without parameter tuning, which is relevant information regarding expected performance of a general purpose solver.

In the first preliminary experiment, we study the impact of changes in the value of the α parameter, which is used to update the *score* in the construction procedures within the diversification generation method. We also test the contribution of the *score* and the smoothed score, *sscore*, to the final solution. We run three versions of our SS algorithm (as it appears in Figure 1) for 5 global iterations with only one combination method, CM_1, instead of the reactive mechanism. The first version, SS-random, does not incorporate the score and therefore probabilistic selections in the diversification and improvement methods are replaced here with completely random selections. The second version, SS-score, incorporates the score strategy, without the smoothness factor. The third version, SS-sscore(α), adds the smoothness parameter α for the updating of the score values. Note that by virtue of equation (5), SS-score is equivalent to SS-sscore(0). Table 1 reports the *RPD*, *#Best* and *Rank* values for the set of 138 instances and each version (with two values of α for SS-sscore(α)) as well as CPU time in seconds (Time).

Method	RPD	#Best	Rank	Time (sec.)
SS-random	4.43%	37	368	38
SS-score	3.09%	40	311	54
SS-sscore(0.3)	2.72%	52	295	55
SS-sscore(0.5)	2.87%	51	296	55

Table 1. Analysis of the *score* parameter

Given the results summarized in Table 1, we determined that the best choice for the value of α is 0.3, which yields the lowest relative deviation and the best *Rank* value. The table also reveals the contribution of probabilistic selection based on smoothed scores over totally random choices.

In this preliminary experiment we also examined the contribution of each diversification generation method (G1, G2 and G3) to the solution quality. Specifically, we employed G1 to generate the *PSize* solutions in *P* and obtained a RPD value of 3.23%, instead of the 2.72% shown in Table 1 in which the three methods are applied (each one generating *PSize*/3 solutions). Similarly, we applied G2 or G3 to populate *P* and obtained RPD values of 3.09% and 2.88%, respectively. This reinforced our conjecture that the combined use of three different generation methods is more effective than the application of a single method in isolation.

In the second preliminary experiment we consider the contribution of the improvement method to the overall scatter search procedure. Table 2 reports the average of the RPD values of two procedures, the BinarySS without the improvement method and the BinarySS with the improvement method. These were both run with the *PSize* parameter set to 100. The results in Table 2 are quite conclusive about the contribution of the improvement method to the quality of the best solutions found during the search.

	Max-cut	MaxDiv	Knapsack
BinarySS without Imp.	46.40%	8.98%	10.43%
BinarySS with Imp.	7.16%	1.00%	1.42%

Table 2. *RPD* values of 10 runs of the improvement method

In our third preliminary experiment we search for an effective value for the *MaxImpIter* parameter in the improvement method. This parameter determines the amount of search effort invested in improving solutions. A large value emphasizes search intensification while a small value favors diversification (by allowing the SS to spend more time combining solutions than improving them). Clearly, intensification and diversification are induced by a number of factors in the scatter search method and here we are only controlling one of them as it relates to the number of iterations in the improvement method. Our experimentation showed that a value of 30 for *MaxImpIter* achieves a good balance between exploration and exploitation and therefore we chose this value to complete the setting of the search parameters.

One of the key elements in the way solutions are combined within our implementation of scatter search is the self-adapting nature of the combination methods. Therefore, as part of the performance analysis of our SS implementation, we wanted to know whether some combination methods are more effective than others for particular problems. With this goal in mind, we performed two additional preliminary experiments. In the first one we recorded the number of times that each combination method was used throughout the search when the procedure was applied to three problem classes. Table 3 reports the relative frequency of the use of each combination method by problem class. For this experiment, the full SS procedure was used with a time limit of 5 seconds. The full SS includes the improvement method and the seven combination methods.

Problem	CM_1	CM_2	CM_3	CM_4	CM_5	CM_6	CM_7
Max-cut	170	244	170	185	171	172	169
MDP	49824	5209	2221	4092	2272	23800	49461
Knapsack	89	93	96	102	99	70	161

Table 3. Frequency of use of the combination methods

In the second experiment relative to the contribution of the combination methods, we consider seven different scatter search versions, SS-CM$_i$ for i = 1 to 7, in which only one combination method, CMi, is applied. In the last row of Table 4, we include the complete SS version, BinarySS, in which the reactive mechanism with the seven combination methods is activated.

Method	RPD	#Best	Rank	Time (sec.)
SS-CM1	3.52%	31	611	14s
SS-CM2	3.50%	26	593	15s
SS-CM3	3.07%	42	404	12s
SS-CM4	3.44%	23	639	15s
SS-CM5	3.76%	29	593	12s
SS-CM6	3.22%	30	533	12s
SS-CM7	3.31%	38	475	12s
BinarySS	2.77%	48	390	11s

Table 4. Contribution of combination methods

Tables 3 and 4 reveal the importance of embedding several mechanisms for combining reference solutions within the scatter search framework. Consider, for instance, **CM₅** in Table 3. This method is rarely used for *MDP* instances (only on 2% of the cases) while it is the fourth most used for *Max-cut* instances (on 13% of the cases). Similar examples of varying performance are found in Table 3, with **CM₁** and **CM₇** as the possible exceptions. Both methods seem to perform well across all three problem types. Table 4 shows the advantage (both in terms of solution quality and speed) of employing several combination methods and the reactive mechanism over versions with a single combination method.

We now compare our scatter search implementation (BinarySS) — employing the parameter values that resulted from the experiments above — against three well known commercial solvers:

- the Evolutionary Premium Solver by Frontline Systems (http://www.frontsys.com) in the Solver Software Development Kit (version 7.2). This solver manages the constraints with the *fcnconstraint* function, defining an *upper bound* in the case of budget problems and *lower and upper bounds* for multiple-choice problems.

- OptQuest (version 6.2) by OptTek Systems (http://www.opttek.com). This engine allows us to define *upper requirements* to handle constraints in Budget problems and *dual requeriments* for multiple-choice problems.

- Evolver by Palisade Corporation (http://www.palisade.com) in the Evolver Development Kit (version 4.1.2). This method manages constraints with the *evconstraintadd* function for both budget and multiple-choice problems.

We also consider state-of-the-art specialized methods for each problem class, all of which are expected to outperform general context-independent solvers. We are only using them as a baseline for comparison. Specifically, for the *Max-cut* instances, we consider the SS method by Martí et al. (2009), for the *MDP* instances we use the SS method by Gallego et al. (2008), and for the *Knapsack* instances we use Expknap by Pisinger (1995). In this experiment, we add the *MDMKP* instances and the Almha procedure by Arntzen et al. (2006). None of the *MDMKP* instances were used for fine tuning and therefore they will be a good test for the generalization attributes of the solver that we have created. For *Max-cut* and *MDP* we compare against the best known

solutions. For the *Knapsack* and *MDMKP* instances we compare against the optimal solutions. The termination limit was set to 30 seconds, however, some methods finished earlier, as triggered by their internal logic. Tables 5 to 8 show the RPD, #*Best*, *Rank*, and the CPU time for each problem class, respectively. We also report in tables 6 to 8 the number of instances for which each method is able to find a feasible solution, given that it is not unusual that black-box solvers fail to find feasible solutions for highly constrained instances. A RPD value of 100% is assigned when a feasible solution is not found. In this way, the average RPD value reflects the method's ability to find feasible solutions.

To handle multiple constraints in the MDMKP problem we consider the static penalty approach described in Yeniay (2005) in which the penalization of an infeasible solution x is computed as $K(1-s/m)$ where K is a large positive constant set to 10^9, m is the number of constraints and s is the number of satisfied constraints. We used this penalty function when executing all the methods in the comparison set.

Method	RPD	#Best	Rank	Time (sec.)
OptQuest	19.09%	6	302	25.14
Evolver	10.34%	10	204	16.16
Solver	8.99%	3	213	14.84
BinarySS	5.20%	13	114	12.33
SS (Marti et al 2009)	0.00%	91	0	13.17

Table 5. Max-cut problem (94 instances)

Method	RPD	#Best	Rank	Time (sec.)	#Feasible
OptQuest	2.48%	7	169	30.51	92
Evolver	88.93%	1	305	23.84	9
Solver	2.40%	2	207	10.85	92
BinarySS	0.22%	31	59	1.97	92
SS (Gallego et al 2008)	0.07%	71	16	30.02	92

Table 6. Maximum diversity problem (92 instances)

Method	RPD	#Best	Rank	Time (sec.)	#Feasible
OptQuest	2.64%	62	36	5.24	96
Evolver	59.03%	3	320	7.95	40
Solver	51.76%	7	290	6.56	47
BinarySS	1.68%	30	124	45.02	96
Expknap (Pisinger 1995)	0.75%	41	76	0.11	96

Table 7. Knapsack problem (96 instances)

Method	RPD	#Best	Rank	Time (sec.)	#Feasible
OptQuest	31.85%	0	189	24.40	65
Evolver	90.63%	0	367	10.95	11
Solver	39.87%	0	289	12.75	65
BinarySS	29.46%	0	207	20.47	71
Almha (Arntzen et al., 2006)	1.86%	28	11	59.97	94

Table 8. Multi-demand multi-dimensional knapsack problem (94 instances)

The behavior of BinarySS is fairly consistent across problem classes. The method delivers low relative percent deviations with respect to the best-known (or optimal) solutions and always ranks best among the black-box solvers, which are also based on evolutionary strategies. As mentioned before, the RPD values in Tables 6 to 8 are computed considering a value of 100% for infeasible solutions. Anyway, one must take into consideration the success rate in finding feasible solutions when assessing RPD values.

The performance of BinarySS on the MDMKP set is more than acceptable. This is a set of instances to which BinarySS has not been exposed before (e.g., during the tuning process) and, as shown in Table 8, the procedure is able to find 71 out of 94 feasible solutions. This compares well with OptQuest, Evolver, and Solver, which found 65, 11, and 65 feasible solutions, respectively. The RPD values in Table 8 also support the conclusion that the BinarySS is able to perform at a higher level than the competing black-box approaches.

In order to test the robustness of the results shown in Tables 5 to 8, we replicate 10 times each method on the entire set of 376 instances. Table 9 shows the standard deviation of the RPD values for each method and each problem type. We can observe that all the methods exhibit very small standard deviations of the RPD values in the four problems (smaller than 0.01 in all the cases), and thus we may conclude that the four black-box solvers exhibit a high level of robustness.

Method	Max-cut	MDP	Knapsack	MDMKP
OptQuest	0.00011	0.00007	0.00008	0.00128
Evolver	0.00302	0.00195	0.00033	0.00687
Solver	0.00327	0.00123	0.00056	0.00448
BinarySS	0.00082	0.00007	0.00017	0.00517

Table 9. Standard deviations of RPD across 10 replications (376 instances)

Our final experiment is designed to show how the average value of the best solution found improves over time. We use the reduced set of problems that were employed for the fine-tuning experiments and run the four context-independent procedures for 180 seconds while recording the average deviation values every second. Figure 2 shows the results of this experiment which does not include the MDMKP instances because the deviation is only informative when the method is able to obtain a feasible solution, and as shown in Table 8, this happens only in a small fraction of the instances for two of the methods.
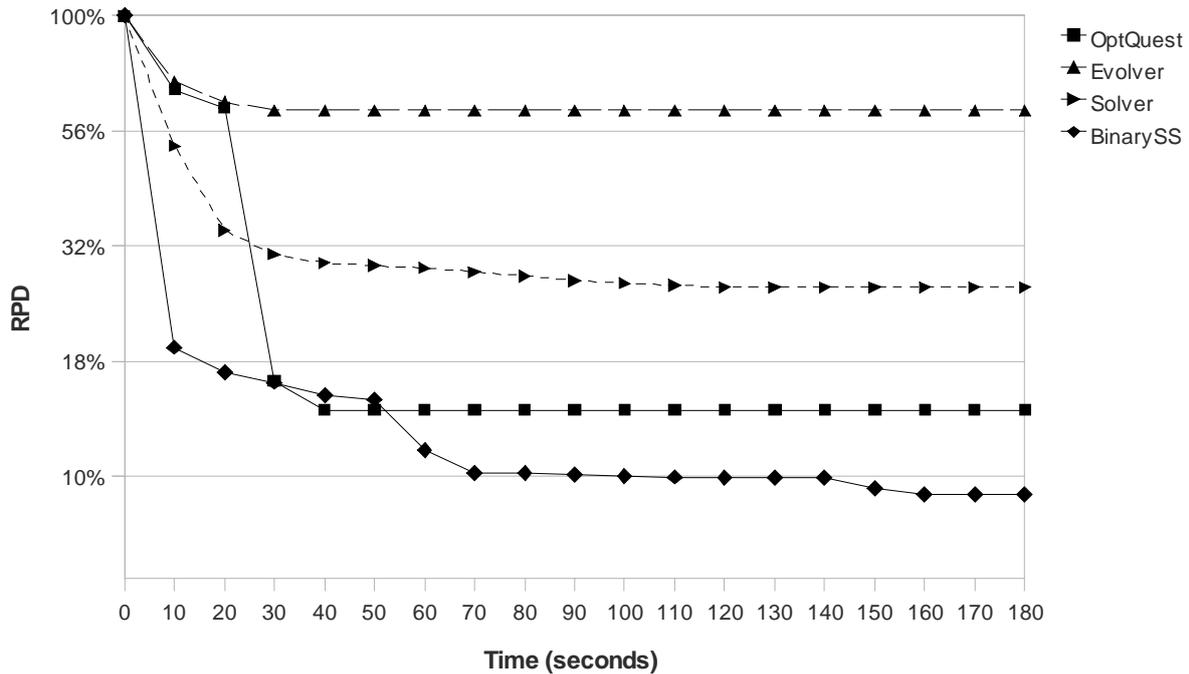
Figure 2. Performance profile for a 180-second run

The profile in Figure 2 indicates that our method is capable of finding high-quality solutions to all the instances in the three classes from the early stages of the search. At no time during the span of the search, the BinarySS produces a RPD value that is worse than any of the competing methods. In this figure, the RPD values are calculated with the best solution found up to that point in the search. Finding high quality solutions early in the search is particularly important in the context of simulation-optimization, where the search horizon (defined as the number of calls to the objective function evaluator) may be particularly limited.

5. Conclusions

We have described the development of a scatter search application to optimization of problems whose solution representation is given by a binary vector. We have discussed the notion of context-independence and how solver designers attempt to use some limited, but important, information to improve the quality of the outcomes. In our particular procedure, we disclose two types of constraints to the solution method. This gives our solver some advantages when constructing and combining solutions because feasibility can be achieved easier than when dealing with general constraints (linear or nonlinear) that are handled with penalty functions.

The results of our experiments with 376 problem instances are quite conclusive regarding the effectiveness of the method that we have developed. In the process of developing this method, we are able to show the advantages of using multiple combination methods. The first-improving strategy when applied to the complete

neighborhood resulted in an effective use of the allotted searching time. We expect that our experience will help software developers to improve upon the commercial solvers that are based on evolutionary processes.

Acknowledgments

References

April, J., F. Glover, J. Kelly and M. Laguna (2003) "Practical Introduction to Simulation Optimization," *Proceedings of the 2003 Winter Simulation Conference*, S. Chick, P. J. Sánchez, D. Ferrin, and D. J. Morrice, eds., New Orleans, pp. 71-78

Arntzen, H., Hvattum, L.M. and Lokketangen, A. (2006) "Adaptive memory search for multidemand multidimensional knapsack problems," *Computers and Operations Research,* vol. 33, pp. 2508-2525.

Campos, V., M. Laguna and R. Martí (2005) "Context-Independent Scatter and Tabu Search for Permutation Problems," *INFORMS Journal on Computing*, vol. 17, no. 1, pp. 111-122

Cappanera, P. and Trubian, M. (2005) "A local search based heuristic for the demand constrained multidimensional knapsack problem," *INFORMS Journal on Computing*, vol. 17, no. 1, pp. 82-98

Dolezal, O., Hofmeister, T. and Hanno Lefmann Y. (1999) "A Comparison of Approximation Algorithms for the Maxcut-Problem, Technical Report ,Universitat Dortmund, Lehrstuhl Informatik.

Duarte, A. and Martí, R. (2007) "Tabu Search for the Maximum Diversity Problem," *European Journal of Operational Research*, vol. 178, pp. 71-84

Festa, P., P.M. Pardalos, M.G.C. Resende and C.C. Ribeiro (2002) "Randomized Heuristics for the Max-cut Problem," *Optimization Methods and Software*, vol. 7, pp. 1033-1058

Fu, M., F. Glover and J. April (2005) "Simulation Optimization: A Review, New Developments and Applications," *Proceedings of the 2005 Winter Simulation Conference*, M. E. Kuhl, N. M. Steiger, F. B. Armstrong and J. A. Joines, eds., pp. 83-95

Fylstra, D., L. Lasdon, J. Watson and A. Waren (1998) "Design and Use of the Microsoft Excel Solver," *Interfaces*, vol. 28, no. 5, pp. 29-55

Gallego, M., A. Duarte, M. Laguna and R. Martí (2008) "Heuristics Algorithm for the Maximum Diverstity Problem," *Computational Optimization and Application*, In Press

Glover, F. (1998) "A Template for Scatter Search and Path Relinking," *Lecture Notes in Compuer Science*, vol. 1363, pp. 13-54

Guikema, S. D., R. A. Davidson and Z Çağnan (2004) "Efficient Simulation-Based Discrete Optimization," *Proceedings of the 2004 Winter Simulation Conference*, R .G. Ingalls, M. D. Rossetti, J. S. Smith, and B. A. Peters, eds., pp. 536-544

Holland, J. (1975) *Adaptation in Natural and Artificial Systems*, University of Michigan Press, Ann Arbor

Laguna, M. and R . Martí (1999) "GRASP and Path Relinking for 2-Layer Straight Line Crossing Minimization," *INFORMS Journal on Computing*, vol. 11, pp. 44–52

Laguna, M. and R. Martí (2003) *Scatter Search: Methodology and Implementations in C*, Kluwer Academic Publishers, Boston

Laguna, M., A. Duarte and R. Martí (2007) "Hybridizing the Cross Entropy Method: An Application to the Max-Cut Problem", *Computers and Operations Research*, In Press

Lippai, I., J. P. Heaney and M. Laguna (1999) "Robust Water System Design with Commercial Intelligent Search Optimizers," *Computing in Civil Engineering*, vol. 13, no. 3, pp. 135-143

Martello, S. and P. Toth (1990) *Knapsack Problems: Algorithms and Computer Implementations*, John Wiley & Sons, Chichester–New York

Martí, R. A. Duarte, M. Laguna (2009) "Advanced Scatter Search for the Max-Cut Problem," *INFORMS Journal on Computing*, vol. 21, no. 1.

Pisinger, D. (1994) "Core problems in Knapsack Algorithms", *Operations Research*.

Pisinger, D. (1995) "An Expanding-Core Algorithm for the Exact 0-1 Knapsack Problem," *European Journal of Operational Research*, vol. 87, pp. 175-187

Ribeiro, C. C., E. Uchoa and R. F. Werneck (2002) "A Hybrid GRASP with Perturbations for the Steiner Problem in Graphs," *INFORMS Journal on Computing* vol. 14, pp. 228 - 246

Rosen, L. R. and C. M. Harmonosky (2005) "An Improved Simulated Annealing Simulation Optimization Method for Discrete Parameter Stochastic Systems", *Computer and Operations Research*, vol. 32, pp. 343-358

Rubinstein, R. Y. and D. P. Kroese (2004) *The Cross-Entropy Method: A Unified Approach to Combinatorial Optimization, Monte-Carlo Simulation, and Machine Learning*, Springer-Verlag, New York

Sahni, S. and T. Gonzales (1976), "P-Complete Approximation Problem", *Journal of the Association for Computing Machinery*, vol. 46, pp. 48-61

Silva, G.C., Ochi, L.S., Martins, S.L. (2004) "Experimental Comparison of Greedy Randomized Adaptive Search Procedures for the Maximum Diversity Problem," *Lecture Notes in Computer Science*, vol. 3059, pp. 498-512

Yeniay, Ö. (2005) "Penalty Function Methods for Constrained Optimization with Genetic Algorithms," *Mathematical and Computational Applications*, vol. 10, no. 1, pp. 45-56