

Chapter #

Multi-Start Methods

Rafael Martí

University of Valencia. Spain

Key words: Optimization, Heuristic Search, Re-Starting

Abstract: Heuristic search procedures that aspire to find global optimal solutions to hard combinatorial optimization problems usually require some type of diversification to overcome local optimality. One way to achieve diversification is to re-start the procedure from a new solution once a region has been explored. In this chapter we describe the best known multi-start methods for solving optimization problems. We propose a classification of these methods in terms of their use of randomization, memory and degree of rebuild. We also present a computational comparison of these methods on solving the linear ordering problem in terms of solution quality and diversification power.

1. INTRODUCTION

Search methods based on local optimization that aspire to find global optima usually require some type of diversification to overcome local optimality. Without this diversification, such methods can become localized in a small area of the solution space, eliminating the possibility of finding a global optimum. In recent years many techniques have been suggested for the avoidance of local optima. One way to achieve diversification is to re-start the search from a new solution once a region has been extensively explored. Multi-Start strategies can then be used to guide the construction of new solutions in a long term horizon of the search process.

There are some problems in which we can find that it is more effective to construct solutions than to apply a local search procedure. For example, in constrained scheduling problems it is difficult to define neighborhoods to

keep feasibility whereas solutions can be relatively easily constructed. Therefore, Multi-Start methods provide an appropriate framework to develop algorithms for solving these problems.

The re-start mechanism can be super-imposed on many different search methods. Once a new solution has been generated, we can apply a simple greedy routine, slight perturbations or a complex metaheuristic to improve it. This chapter is focused on studying the different ways, strategies and methods of generating solutions to re-start a search for a global optimum.

2. AN OVERVIEW

Multi-Start methods have two phases: the first one in which the solution is generated and the second one in which the solution is typically (but not necessarily) improved. Then, each global iteration produces a solution (usually a local optima) and the best overall is the algorithm's output.

Figure 1 shows a pseudo-code of the multi-start procedure. A solution x_i is constructed in Step 1 at iteration i . This is typically performed with a constructive algorithm. Step 2 is devoted to improving this solution, obtaining solution x_i' . A simple improvement method can be applied. However, this second phase has recently become more elaborate and, in some cases, is performed with a complex metaheuristic that may or may not improve the initial solution x_i (in this latter case we set $x_i'=x_i$).

```

Initialise  $i=1$ 
while(Stopping condition is not satisfied)
{
    Step 1. (Generation)
        Construct solution  $x_i$ 
    Step 2. (Search)
        Apply a search method to improve  $x_i$ .
        Let  $x_i'$  be the solution obtained
    if(  $x_i'$  improves the best )
        Update the best
     $i=i+1$ 
}

```

Figure 1. Multi-Start procedure

In recent years, a great number of heuristic algorithms have been proposed to solve some combinatorial optimization problems following the outline given in Figure 1. Some of them are problem-dependent and the ideas and strategies implemented are difficult to apply to different problems,

while others are based on a framework that can be directly used to design solving methods for other problems. In this section we describe the most relevant procedures in terms of their application to a wide variety of problems.

Tabu Search is by now a well-known metaheuristic for solving hard combinatorial optimization problems. One of the papers that contains a number of its foundation ideas (Glover, 1977), also includes a focus on applying these ideas within a framework of iterated re-starting. Adaptive memory designs can be used to retain and analyze features of selected solutions and thus, provide a basis for improving future executions of the constructive process. Different forms of memory functions, like frequency and recency information, are proposed to design these restarting mechanisms.

Adaptive memory strategies introduced in this seminal paper have had widespread success in solving a variety of practical and difficult combinatorial optimization problems. They have been adapted to many different fields in the combinatorial optimization theory, since the exploitation of such memory has been proved to be very effective in most of the metaheuristic methods. Some of them are explicitly based on this memory structures like tabu search, while others like simulated annealing or re-starting methods, have evolved incorporating these ideas. Some applications of probabilistic forms of re-starting based on memory functions are given in Rochat and Taillard (1995) and Lokketangen and Glover (1996).

Early papers in multi-start methods are devoted to the Monte Carlo random re-start in the context of nonlinear unconstrained optimization, where the method simply evaluates the objective function at randomly generated points. The probability of success approaches one as the sample size tends to infinity under very mild assumptions about objective function. Many algorithms have been proposed that combine the Monte Carlo method with local search procedures (Rinnooy Kan and Timmer, 1989). Solis and Wets (1981) study convergence for random re-starts methods in which the probability distribution for choosing the next starting point can depend on the evolution of the search. Some extensions of these methods seek to reduce the number of complete local searches that are performed and increase the probability that they are started from points close to the global optimum (Mayne and Meewella, 1988).

Ulder et al. (1990) combines genetic algorithms with local search strategies improving previous genetic approaches for the travelling salesman

problem. The method allows an iterative algorithm to improve each individual, either before or while being combined with other individuals to form a new solution “offspring”. The combination of these three elements: *Generation*, *Combination* and *Local Search*, extends the paradigm of Re-Start and links with other areas of the metaheuristics such as Scatter Search (Glover et al., 2000) or Memetic Algorithms (Moscato, 1999).

Hu et al. (1994) study from a theoretical point of view, the combination of the “gradient” algorithm with random initializations for finding a global optimum. Efficacy of parallel processing, choice of the restart probability distribution and number of restarts are studied for both discrete and continuous models. The authors show that the uniform probability is a good measure for restarting procedures.

Boese et al. (1994) analyze relationships among local minima “from the perspective of the best local minimum”, finding convex structures in the cost surfaces. From that study they propose a multi-start method where starting points for greedy descent are adaptively derived from the best previously found local minima. In the first step, Adaptive Multi-Start heuristics (AMS) generate r random starting solutions and run a greedy descent method from each one to determine a set of corresponding random local minima. In the second step, *adaptive starting solutions* are constructed based on the local minima obtained so far and improved with a greedy descent method. This improvement is applied several times from each adaptive starting solution to yield corresponding *adaptive local minima*. The authors test this method for the traveling salesman problem and obtain significant speedups over previous multi-start implementations. Hagen and Kahng (1997) apply this method for the iterative partitioning problem.

Moreno et al. (1995) proposed a stopping rule for the multi-start method based on a statistical study of the number of iterations needed to find the global optimum. The authors introduce two random variables that provide together a way to estimate the number of global iterations needed to find the global optima: the number of initial solutions generated and the number of objective function evaluations performed on finding the global optima. From these measures, the probability that the incumbent solution is the global optimum is evaluated via a normal approximation. Thus, at each global iteration, this value is computed and if it is greater than a prefixed threshold, the algorithm stops, otherwise a new solution is generated. The authors illustrate these method in the median p-hub problem.

Simple forms of multi-start methods are often used as a way to compare other methods and measure their relative contribution. Baluja (1995) compare different genetic algorithms for six sets of benchmark problems commonly found in the GA literature: Traveling salesman problem, job-shop scheduling, knapsack, bin packing, neural network weight optimization, and numerical function optimization. The author uses the multi-start method (Multiple restart stochastic hill-climbing, MRSH) as a baseline in the computational testing. Since solutions are represented with strings, the improvement step consist of a local search based on random flip of bits. The results indicate that using genetic algorithms for the optimization of static functions does not yield a benefit, in terms of the final answer obtained, over simpler optimization heuristics. Other comparisons between MRSH and GAs can be found, for example, in Ackley (1987) or Wattenberg and Juels (1994).

One of the most well known Multi-start methods is the greedy adaptive search procedures (GRASP). The GRASP methodology was introduced by Feo and Resende (1995). It was first used to solve computationally difficult set covering problems (Feo and Resende, 1989). Each GRASP iteration consists of constructing a trial solution and then applying a local search procedure to find a local optimum (i.e., the final solution for that iteration). The construction step is an adaptive and iterative process guided by a greedy evaluation function. It is iterative because the initial solution is built considering one element at a time. It is greedy because the addition of each element is guided by a greedy function. It is adaptive because the element chosen at any iteration in a construction is a function of those previously chosen. (That is, the method is adaptive in the sense of updating relevant information from one construction step to the next.). At each stage, the next element to be added to the solution is randomly selected from a candidate list of high quality elements according to the evaluation function. Once a solution has been obtained, it is typically improved by a local search procedure. The improvement phase performs a sequence of moves towards a local optimum solution, which becomes the output of a complete GRASP iteration. Some examples of successful applications are given in Laguna et al. (1994), Resende (1998) and Laguna and Martí (1999).

Hickernell and Yuan (1997) present a multi-start algorithm for unconstrained global optimization based on *quasirandom samples*. Quasirandom samples are sets of deterministic points, as opposed to random, that are evenly distributed over a set. The algorithm applies an inexpensive local search (steepest descent) on a set of quasirandom points to concentrate the sample. The sample is reduced replacing worse points by new

quasirandom points. Any point that is retained for a certain number of iterations is used to start an efficient complete local search. The algorithm terminates when no new local minimum is found after several iterations. An experimental comparison shows that the method performs favorably with respect to other global optimization procedures.

Hagen and Kang (1997) used an adaptive multi start method for the partitioning optimization VLSI problem where the objective is to minimize the number of signals which pass between components. The method consists of two phases: (1) Generate a set of random starting points and call the iterative (local search) algorithm, thus determining a set of local minimum solutions; and (2) construct adaptive starting points that are central from the best local minimum solutions found so far. The authors add a preprocessing cluster module to reduce the size of the problem. The resulting Clustering Adaptive Multi-Start method (CAMS) is fast and stable and improves upon previous partitioning results in the literature.

Fleurent and Glover (1999) propose some adaptive memory search principles to enhance multi-start approaches. The authors introduce a template of a constructive version of Tabu Search based on both, a set of elite solutions and the intensification strategies that rely on the identification of *strongly determined* and *consistent variables*. Strongly determined variables are those whose values cannot be changed without significantly eroding the objective function value or disrupting the values of other variables. A consistent variable is defined as one that receives a particular value in a significant portion of good solutions. The authors propose the inclusion of memory structures within the multi-start framework as those used in tabu search: *recency*, *frequency* and *attractiveness*. Computational experiments for the quadratic assignment problem disclose that these methods improve significantly over previous multi-start methods like GRASP and Random Restart that do not incorporate memory based strategies.

Multi-Start procedures usually follow the global scheme given in Figure 1; but there are some applications in which Step 2 can be applied several times within a global iteration. In the *incomplete construction methods*, the improvement phase was periodically invoked during the construction process to the partial solution rather than the standard implementation after the complete construction. See Russell (1995) and Chiang and Russell (1995) for successful applications to this approach in vehicle routing.

Patterson et al (1999) introduce a multi-start framework (Adaptive Reasoning Techniques, ART) based on memory structures. The authors implement the short term and long term memory functions, proposed in the tabu search framework, to solve the Capacitated Minimum Spanning Tree Problem. ART is an iterative, constructive solution procedure that implements learning methodologies on top of memory structures. ART derives its success from being able to learn about, and modify the behavior of a primary greedy heuristic. The greedy heuristic is executed repeatedly, and for each new execution we probabilistically introduce constraints that prohibit certain solution elements from being considered by the greedy heuristic. The active constraints are held in a short term memory. A long term memory holds information regarding which constraints were in the active memory for the best set of solutions.

Glover (2000) proposes approaches for creating improved forms of constructive multi-start and strategic oscillation methods, based on new search principles: *persistent attractiveness* and *marginal conditional validity*. These concepts play a key role in deriving appropriate measures to capture information during prior search. Applied to constructive neighborhoods, strategic oscillation operates by alternating constructive and destructive phases, where each solution generated by a constructive phase is dismantled (to a variable degree) by the destructive phase, after which a new phase builds the solution anew. The conjunction of both phases and their associated memory structures provides the basis of an improved multi-start method.

Prais and Ribeiro (2000) propose an improved GRASP implementation, called reactive GRASP, for a matrix decomposition problem arising in the context of traffic assignment in communication satellites. The method incorporates a memory structure to record information about the solutions previously found. In Reactive GRASP, the basic parameter which defines the restrictive-ness of the candidate list during the construction phase is self-adjusted, according to the quality of the solutions previously found. The proposed method matches most of the optimal solutions known.

An open question in order to design a good search procedure is whether it is better to implement a simple improving method that allows a great number of global iterations or, alternatively, to apply a complex routine that significantly improves a few generated solutions. A simple procedure depends heavily on the initial solution but a more elaborate method takes much more running time and therefore can only be applied a few times, thus reducing the sampling of the solution space. Some metaheuristics, such as

GRASP, launch limited local searches from numerous constructions (i.e., starting points). In most of the tabu search implementations, the search starts from one initial point and if a restarting procedure is also part of the method, it is invoked only a limited number of times. However, the inclusion of re-starting strategies within the tabu search framework has been well documented in several papers (Glover, 1977, Glover and Laguna, 1997).

Martí et al. (2000) study the balance between restarting and search-depth (i.e., the time spent searching from a single starting point) in the context of the bandwidth matrix problem. They tested both alternatives and concluded that it was better to invest the time searching from a few starting points than restarting the search more often. Although we cannot draw a general conclusion from these experiments, the experience in the current context and in previous projects indicates that some metaheuristics like tabu search need to reach a critical search depth to be effective. If this search depth is not reached, the effectiveness of the method is severely compromised.

3. A CLASSIFICATION

We have found three key elements in multi-start methods that can be used for classification purposes: memory, randomization and degree of rebuild. The choices for each one of these elements are not restricted to the extreme case of “present” or “not present”, but they represent the whole range between both extremes that can be labeled as *Memory/Memory-less*, *Systematic/Randomized* and *Rebuild/Build from scratch*, respectively.

The first element is the **Memory** and it is used to identify elements that are common to good solutions previously generated. As in the Tabu Search framework (Glover and Laguna, 1997), it provides a foundation for incentive-based learning, where inducements are provided to reinforce actions that lead to good solutions. Thus, instead of simply resorting to randomized re-starting processes, in which current decisions derive no benefit from knowledge accumulated during prior search, specific types of information are identified for exploiting history. On the other hand, the avoidance of memory (*Memory-less*) is not as unreasonable as might be imagined since the construction of “unconnected” solutions may be interpreted as a means of strategically sampling the solution space. It should be noted that the meaning of good is not restricted to the objective function, but also includes the notion of diversity as described later.

Starting solutions can be randomly generated or, on the contrary, they can be generated in a systematic way. **Randomization** is a very simple way of achieving diversification, but with no control over the diversity achieved since in some cases we can obtain very similar solutions. We can add some mechanisms to control the similarities in order to discard some solutions or generate the solutions in a deterministic way that guarantees a certain degree of difference. The extremes of this element can be described as *Randomized* where solutions are generated in a random way and *Systematic* where solutions are generated in a deterministic way. Between both extremes there are a great number of possibilities for combining random elements with deterministic rules. GRASP constructions is an example of a combined method.

The **Degree of Rebuild** indicates the elements that remain fixed from one generation to another. Most applications *build from scratch* the solution at each generation, but recent implementations have fixed, for a certain number of iterations, some elements in the construction process that have appeared in previously generated solutions. Such an approach was proposed in the context of identifying and then iteratively exploiting ‘strongly determined and consistent variables’ in Glover (1977). This selective fixing of elements by reference to their previous impact and frequency of occurrence in various solution classes is a memory-based strategy of the type commonly used in tabu search. It can also be considered as an instance of Path Relinking (Glover and Laguna, 1993) that generates new solutions by exploring trajectories that connect high-quality solutions. This approach seeks to incorporate the attributes of elite solutions previously generated by creating inducements to favor these attributes in the solutions. In an extreme case all the elements in the new solution will be fixed by the information generated from the set of elite solutions considered. This is labeled as *Rebuild*.

The constructive algorithm depicted in Figure 2 has no memory structures, a combination between randomization and systematic construction rules and, at each iteration, the solution is completely built from scratch.

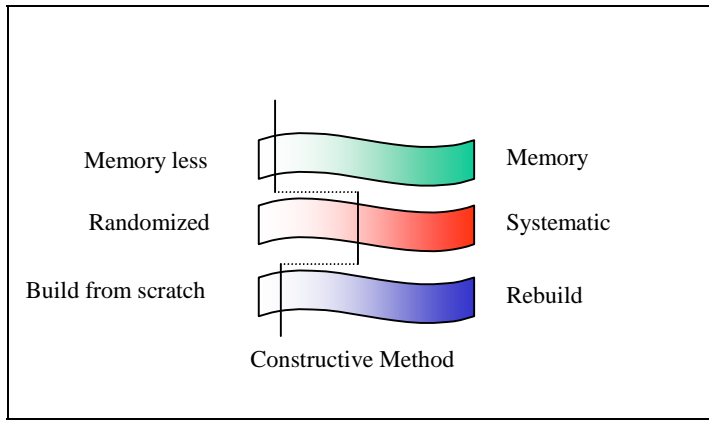


Figure 2. Multi-Start classification

Given different re-starting methods for a problem, one way of comparing them is to generate a set of solutions with each one and compare their quality and diversity. Since the quality is trivially measured by the objective function, we now propose two measures of diversity. We restrict our attention to solutions represented by permutations.

3.1 Diversity Measures

The first measure consists of computing the distances between each solution and a “center” of the set of solutions P . The sum (or alternatively the average) of these $|P|$ distances provides a measure of the diversity of P . The second one consists of computing the distance between each pair of solutions in P . The sum of these $|P| \times |P|$ distances provides another way of measuring the diversity of P .

The first diversity measure is calculated as follows:

1. Calculate the median position of each element i in the solutions in P .
2. Calculate the dissimilarity of each solution in the population with respect to the median solution. The dissimilarity is calculated as the sum of the absolute difference between the position of the elements in the solution under consideration and the median solution.
3. Calculate d as the sum of all the individual dissimilarities.

To illustrate, suppose that P consists of the following three orderings: (A,B,C,D), (B,D,C,A), (C,B,A,D). The median position of element A is therefore 3, since it occupies positions 1, 3 and 4 in the given orderings. In the same way, the median positions of B, C and D are 2, 3 and 4,

respectively. Note that the median positions might not induce an ordering, as in the case of this example. The dissimilarity of the first solution is then calculated as follows:

$$d_1 = |1-3| + |2-2| + |3-3| + |4-4| = 2.$$

In the same way, the dissimilarities of the other two solutions are $d_2 = 4$ and $d_3 = 2$. The diversity measure of P is then given by $d = 2+4+2 = 8$.

The second measure is calculated, for each pair of solutions in P , as the sum of the absolute differences between the positions of each element in both solutions. The sum of these $|P \times P|$ values provides the measure of the diversity of the set P . The value with solutions (A,B,C,D) and (B,D,C,A) in the previous example is computed as follows:

$$d'_{12} = |1-4| + |2-1| + |3-3| + |4-2| = 6.$$

In the same way, the values of the other three pairs of solutions are $d'_{13} = 4$ and $d'_{23} = 6$. The diversity measure of P is then given by $d' = 6+4+6 = 16$.

We have found computationally that both measures are strongly correlated and provide the same information. Since the second measure is computationally more expensive than the first, we will use the first one in the following experiments.

It should be noted that a third measure could be added to evaluate a set of solutions. The notion of *influence* introduced by Glover (1990) in the context of Tabu Search, can be adapted to our study. The influence considers the potential and the structure of a solution in the search process. The authors propose memory functions that classify moves relative to their attractiveness within “distance classes” and other measures of their impact. Consider, for example, two solutions a and b with the same objective and diversity values, but a is close to a local optimum with a better objective function value than a and b , while b is itself a local optimum. Consequently, we probably obtain a better solution with a search from a rather than from b . Therefore it is more valuable to have a than b in the set of solutions since it has more influence in the search for the global optimum. Obviously we do not know *a priori* if a given solution is closer to a local optimum than another, but if we identify some properties of good solutions we will be able to define evaluators and measures to reflect the “importance” or influence of the solutions. Good starting points for this study are given by the solution structure, landscape and neighborhood induced by the local search method.

4. COMPUTATIONAL EXPERIMENTS

The linear ordering problem (LOP) has generated a considerable amount of research interest over the years, as documented in Grotschel, et al. (1984) and Campos et al. (1999). Because of its practical and theoretical relevance, we use this problem as a test case for re-start mechanisms.

Given a matrix of weights $E = \{ e_{ij} \}_{m \times m}$, the LOP consists of finding a permutation p of the columns (and rows) in order to maximize the sum of the weights in the upper triangle. In mathematical terms, we seek to maximize:

$$C_E(p) = \sum_{i=1}^{m-1} \sum_{j=i+1}^m e_{p(i)p(j)}$$

where $p(i)$ is the index of the column (and row) in position i in the permutation. Note that in the LOP, the permutation p provides the ordering of both the columns and the rows. The equivalent problem in graphs consists of finding, in a complete weighted graph, an acyclic tournament with a maximal sum of arc weights (Reinelt, 1985).

Instances of input-output tables from sectors in the European Economy can be found in the public-domain library LOLIB (1997). We employ these problem instances to compare different restarting methods.

We have tested 10 re-starting generation methods. Six of these methods are based on GRASP (Feo and Resende, 1995) constructions with a greedy function that selects sectors based on a measure of attractiveness.

G1. A GRASP construction where the attractiveness of a row is the sum of the elements in its corresponding row. The method randomly selects from a short list of the most attractive sectors and constructs the solution starting from the first position of the permutation.

G2. A GRASP construction where the attractiveness of a sector is the sum of the elements in its corresponding column. The method randomly selects from a short list of the most attractive sectors and constructs the solution starting from the first position of the permutation.

G3. A GRASP construction where the attractiveness of a sector is the sum of the elements in its corresponding row divided by the sum of the elements in its corresponding column. The method randomly selects from a

short list of the most attractive sectors and constructs the solution starting from the first position of the permutation.

G4, G5 and G6. These methods are identical to the first three, except that the selection of sectors is from a short list of the least attractive and the solution is constructed starting from the last position of the permutation.

MIX. A mixed procedure derived from the previous six. The procedure generates an even number of solutions from each of the previous six methods and combines these solutions into a single set. That is, if n solutions are required, then each method G_i (for $i = 1, \dots, 6$) contributes with $n/6$ solutions.

RND. A random generator. This method simply generates random permutations.

DG. A general purpose diversification generator suggested in Glover (1998) which generates diversified permutations in a systematic way without reference to the objective function.

FQ. A method using frequency-based memory, as proposed in Tabu Search (Glover and Laguna, 1997). This method is based on modifying a measure of attractiveness with a frequency measure that discourages sectors from occupying positions that they have frequently occupied in previous solution generations. See Campos et al. (1999) for a description of the method.

In our first experiment we use the instance *stabu3* from the LOLIB. We have generated a set of $N=100$ solutions with each of the 10 generation methods. Figures 3 and 4 show the box and whiskers plot of the objective function value and dissimilarity, respectively, of the solution set obtained with each method.




Figure 3. Objective function value box plot for each method




Figure 4. Dissimilarity box plot for each method

With both diagrams together (Figures 3 and 4) we can observe the performance of the 10 generators on the problem *stabu3*. We have repeated the same experiments on 10 other problems from the LOLIB, obtaining similar diagrams.

A good re-starting method must produce a set of solutions with high quality and high diversity. If we compare, for example, generators MIX and

G3 we observe in Figure 3 that G3 produces slightly better solutions in terms of solution quality, but Figure 4 shows that MIX outperforms G3 in terms of diversity. Therefore, we will probably select MIX as a better method than G3.

In order to rank the methods we have computed the average of both measures across each set. Figure 5 shows in the x-axis the average of the dissimilarity and in the y-axis the average of the quality. A point is plotted for each method.

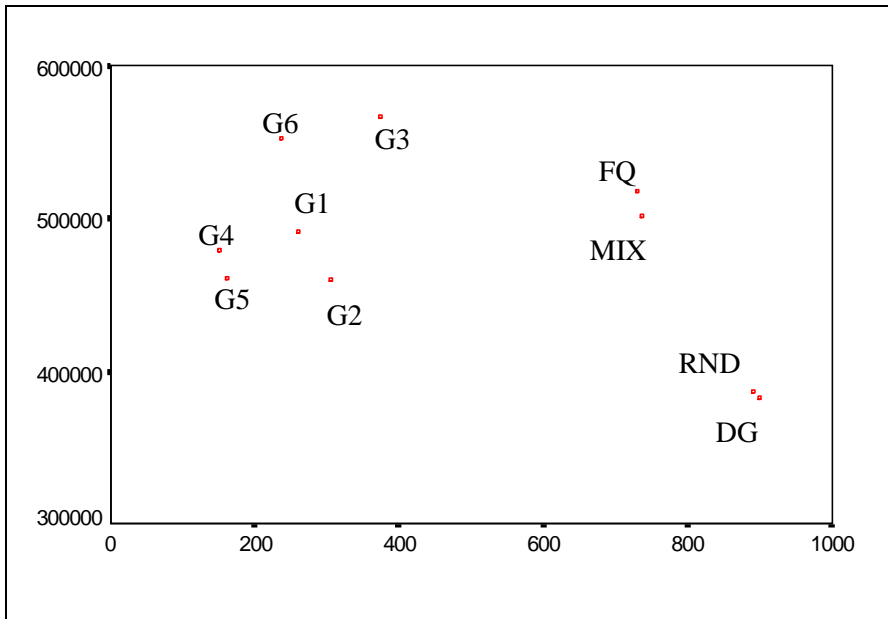


Figure 5. Quality and Dissimilarity for each method

As expected, the random generator (RND) produces the maximum diversity (as measured by the dissimilarity value). DG matches the diversity of RND using a systematic approach instead of randomness. The mixed method MIX provides a good balance between dissimilarity and quality, by the union of solutions generated with methods G1 to G6.

We have standardized both averages in order to directly compare them. We think that quality and diversity have the same importance, so we have added both relative averages, obtaining the following ranking where the overall best is the FQ generator:

G5, G4, G2, G1, DG, RND, G6, G3, MIX and FQ.

This results are in line with previous works where the inclusion of memory structures has been shown to be effective within the multi-start framework. However, it should be noted that this ranking of methods has been obtained considering both measures, quality and diversity, with equal weight. If we vary this criterion, the ranking would also change.

5. CONCLUSIONS

The objective of this study has been to expand and advance the knowledge associated with the implementation of multi-start procedures. Unlike other well-known methods, it has not yet become widely implemented and tested as a metaheuristic itself for solving complex optimization problems. We have shown new ideas that have recently emerged within the multi-start area that add a clear potential to this framework which has yet to be fully explored.

ACKNOWLEDGEMENTS

The author wishes to thank professor Vicente Campos for his valuable help in the computational testing.

REFERENCES

- Ackley, D.H. (1987) "An empirical study of bit vector function optimization" Genetic Algorithms and simulated annealing, Ed. Davis, Morgan Kaufmann Publishers.
- Baluja, S. (1995) "An Empirical Comparison of 7 iterative evolutionary function optimization heuristics", School of computer science, Carnegie Mellon University.
- Boese, K.D., Kahng, A.B. and Muddu S. (1994) " A new adaptive multi-start technique for combinatorial global optimisation" *Operations Research Letters*, 16, pp 103-113
- Campos, V., Laguna, M. and Martí R. (1999) "Scatter Search for the Linear Ordering Problem" *New Ideas in Optimization*, Eds Corne, Dorigo and Glover, Mc Graw Hill, pp 331-341.
- Chiang, W.C. and Russell, R.A. (1995) "Simulated annealing metaheuristics for the vehicle routing problems with time windows" *Annals of Operations Research*, Vol. 60.

- Feo, T. and M. G. C. Resende (1989) "A probabilistic Heuristic for a Computationally Difficult Set Covering Problem," *Operations Research Letters*, Vol. 8, pp. 67-71
- Feo, T. and M. G. C. Resende (1995) "Greedy Randomized Adaptive Search Procedures," *Journal of Global Optimization*, vol. 2, pp. 1-27.
- Fleurent C. and Glover F. (1999) "Improved constructive multi-start strategies for the Quadratic assignment problem using adaptive memory" *INFORMS Journal on Computing*.
- Glover, F. (1977) "Heuristics for Integer Programming Using Surrogate Constraints," *Decision Sciences*, Vol. 8, pp. 156-166.
- Glover, F. (1990) "Tabu Search: A tutorial", *Interfaces*, Vol. 20, pp. 74-94.
- Glover, F. (2000) "Multi-Start and strategic oscillation methods – Principles to exploit adaptive memory" *Computing tools for modeling optimization and simulation*, Eds Laguna and Gonzalez-Velarde, Kluwer, pp 1-25.
- Glover, F. and M. Laguna (1993) "Tabu Search", *Modern Heuristic Techniques for Combinatorial Optimization Problems*, C. Reeves (Ed.), Blackwell Scientific Publishing, Oxford, pp. 70-141.
- Glover, F. and M. Laguna (1997) *Tabu Search*, Kluwer Academic Publishers, Boston.
- Glover, F. , M . Laguna and R. Martí (2000) "Scatter Search" to appear in *Theory and Applications of Evolutionary Computation: Recent Trends*, A. Ghosh and S. Tsutsui (Eds.), Springer-Verlag
- Grotschel, M., M. Junger and G. Reinelt (1984), "A Cutting Plane Algorithm for the Linear Ordering Problem," *Operations Research*, vol. 32, no. 6, pp. 1195-1220.
- Hagen, L.W. and Kahng A.B. (1997) "Combining problem reduction and adaptive multi-start: a new technique for superior iterative partitioning", *IEEE Trans. on CAD*, Vol. 16 (7), pp. 709-717.
- Hickernell, F.J. and Yuan, Y. (1997) "A simple multistart algorithm for global optimization", *OR Transactions*, Vol. 1 (2).
- Hu, X., Shonkwiler, R. and Spruill, M.C. (1994) "Random restarts in global optimization" Georgia Institute of technology, Atlanta.
- Laguna, M., T. Feo and H. Elrod (1994) "A Greedy Randomized Adaptive Search Procedure for the 2-Partition Problem", *Operations Research*, vol. 42, no. 4, pp. 677-687.
- Laguna, M. and Martí, R. (1999) "GRASP and Path relinking for 2-layer straight line crossing minimization" *INFORMS journal on computing*, vol. 11 (1), pp. 44-52.

- Lokketangen, A. and Glover, F. (1996) "Probabilistic move selection in tabu search for 0/1 mixed integer programming problems" *Meta-Heuristics: Theory and Practice*, Kluwer, pp. 467-488.
- LOLIB (1997) <http://www.iwr.uni-heidelberg.de/iwr/comopt/soft/LOLIB/LOLIB.html>.
- Martí, R., Laguna, M., Glover, F. and Campos, V. (2000) "Reducing the bandwidth of a sparse matrix with tabu search," *European Journal of Operational Research*, Forthcoming.
- Mayne, D.Q. and Meewella, C.C. (1988) "A non-clustering multistart algorithm for global optimization" *Analysis and Optimization of Systems*, Ed. Bensoussan and Lions, Lecture Notes in control and information sciences, vol. 111, springer verlag.
- Moreno J. A., Mladenovic N. and Moreno-Vega J.M. (1995) "An Statistical Analysis of Strategies for Multistart Heuristic Searches for p-Facility Location-Allocation Problems." Eighth Meeting of the EWG on Locational Analysis Lambrecht, Germany.
- Moscato, P. (1999) "Memetic Algorithms" *New ideas in optimization*" Eds. Corne, Dorigo and Glover, Mc Graw Hill, pp. 219-235.
- Osman, I.H. and Kelly, J.P. (1996) "Meta-Heuristics: An Overview", *Meta-Heuristics: Theory and Practice*, Kluwer, pp 1-23.
- Prais, M. and Ribeiro C.C. (2000) "Reactive GRASP: An application to a matrix decomposition problem in TDMA traffic assignment" vol 12, 164-176.
- Rochat and Taillard (1995) "Probabilistic diversification and intensification in local search for vehicle routing" *Journal of heuristics*, Vol. 1, 1, pp 147-167.
- Resende M. G. C. (1998) "Computing approximate solutions for the maximum covering problem using GRASP" *Journal of Heuristics*, Vol. 4, pp. 161 – 171.
- Patterson R., Pirkul H. and Rolland E (1999) "Adaptive Reasoning Technique for the Capacitated Minimum Spanning Tree Problem", *Journal of Heuristics*, 5, pp. 159-180.
- Rinnooy Kan, A.H.G. and Timmer, G.T. (1989) "Global Optimization" *Handbooks in operations research and management science*, Vol. 1, Ed. Rinnooy Kan and Todds, North Holland, pp. 631-662.
- Russell, R.A. (1995) "Hybrid heuristics for the vehicle routing problem with time windows", *Transportation Science*.
- Solis, F. and Wets, R. (1981) "Minimization by random search techniques", *Math. of Operations Research*, Vol. 6, pp. 19-30.
- Ulder, N.L.J., Aarts, E.H.L., Bandelt, H.J., Van Laarhoven, P.J.M. and Pesch, E. (1990) "Genetic Local Search algorithms for the traveling salesman problem", *Parallel problem solving from nature*, Eds. Schwefel and Männer, Springer Verlag, pp. 109-116.

Wattenberg, M. and Juels, A. (1994) "Stochastic Hillclimbing as a baseline method for evaluating genetic algorithms" University of California – Berkeley, CSD94-834.