
SSPMO: A Scatter Search Procedure for Non-Linear Multiobjective Optimization

JULIÁN MOLINA¹, MANUEL LAGUNA², RAFAEL MARTÍ³ AND RAFAEL CABALLERO¹

¹ Universidad de Málaga, Spain, {julian.molina, rafael.caballero}@uma.es

² University of Colorado at Boulder, USA, laguna@colorado.edu

³ Universitat de València, Spain, Rafael.marti@uv.es

Latest revision: September 3, 2004

Abstract — We describe the development and testing of a metaheuristic procedure, based on the scatter search methodology, for the problem of approximating the efficient frontier of nonlinear multiobjective optimization problems with continuous variables. Recent applications of scatter search have shown its merit as a global optimization technique for single-objective problems. However, the application of scatter search to multiobjective optimization problems has not been explored fully in the literature. We test the proposed procedure on a suite of problems that have been used extensively in multiobjective optimization. Additional tests are performed on instances that are an extension of those considered classic. The tests indicate that our extension of the basic scatter search framework is a viable alternative for multiobjective optimization.

1. Introduction

Most decision problems in the real world require the simultaneous optimization of more than one criterion. That is, decision makers are not able to base their decisions on a single criterion in order to identify one or more courses of action that they may be willing to take in any given situation. Moreover, real decision problems are such that often conflicting criteria are used to evaluate alternative solutions. For example, a decision problem in a service operation may involve the maximization of personnel utilization as well as the increase in quality of service. An easy way of rising utilization is to reduce the workforce. However, this may be at odds with increasing service quality. Therefore, in most multiobjective optimization problems, there is no single solution that yields an optimal performance for every objective function used for evaluation.

Multiobjective Programming is the branch of Mathematical Programming that deals with problems for which more than one objective function is required to evaluate the merit of alternative decisions. Formally, the problem is stated as follows:

$$\begin{array}{ll} \text{Max} & (f_1(x), f_1(x), \dots, f_p(x)) \\ \text{s.t.} & x \in X \end{array}$$

where

$x = (x_1, x_2, \dots, x_n)$ are the decision variables

X is the set of feasible solutions

$y_i = f_i(x)$ is the i^{th} objective function (or decision criterion)

$Y = f(X)$ is the objective function space

Comparing alternative solutions to a given problem is one of the first difficulties encountered when moving from single-objective optimization to multiobjective optimization. In other words, how do we know that one solution is better than another when performance is measured by more than one objective function value? Given that the merit of a solutions is represented by a vector of objective function values, it is not always possible to determine when a vector is strictly larger or smaller than another. Generally, there is no complete order in Y . Multiobjective optimization introduces the concept of efficiency, developed by Vilfredo Pareto in 1896. Essentially, efficiency means that a solution to a multiobjective function is such that no single objective can

be improved without deteriorating another one. Assuming that we want to maximize all objective functions, the Pareto order can be defined mathematically as:

Given two points $y, y' \in Y$, we say that y is preferred to y' if

$$y_i \geq y'_i \quad \forall i = 1, \dots, p \quad \text{and there exists at least one } j \in \{1, \dots, p\} \text{ such that}$$

$$y_j > y'_j.$$

Clearly, the Pareto order is a partial one and therefore it is not possible to select a solution that is preferred over all other solutions, that is, it is not possible to find “an optimal solution”. Hence, the concept of optimality must be generalized in such a way that more than one solution can be considered ideal (optimal) in problems with multiple objectives. This generalization follows from the definition of the Pareto order and is known as Pareto efficiency:

A solution $x^* \in X$ is efficient if there is no other solution $x \in X$ such that $f(x)$ is preferred to $f(x^*)$ according to the Pareto order. That is,

$$x^* \in X \quad \text{is efficient if there is no solution } x \in X \text{ such that}$$

$$f_i(x) \geq f_i(x^*) \quad \forall i = 1, \dots, p \quad \text{and at least one } j \in \{1, \dots, p\} \text{ such that}$$

$$f_j(x) > f_j(x^*).$$

Most multiobjective programming techniques focus on finding the set of efficient points (E) for a given problem or, in the case of heuristic procedures, an approximation of the efficient set (\hat{E}). In this paper, we describe the development and testing of a metaheuristic procedure for multiobjective optimization problems for which $f_i(x)$, for $i = 1, \dots, p$, are nonlinear functions and x are continuous and bounded variables. Since our approach is not exact, our goal is to search for the best \hat{E} . Metaheuristics have been applied to this problem, so before discussing our proposed procedure, we review the approaches that are relevant to our current investigations.

2. Metaheuristics for Multiobjective Optimization

Most exact techniques for multiobjective optimization deal with continuous and discrete linear problems. The application of these techniques to complex multiobjective optimization problems has been considered impractical. We refer to complex problems to those with nonlinearities in the objective function and/or constraints, with a large

number of discrete variables or uncertainty in key data. To deal with these difficulties, which are not unusual in real settings, researchers and practitioners have relied on metaheuristic procedures, as indicated in the survey articles by Ehrgott and Gandibleux (2000) and Jones, Mirrazavi and Tamiz (2002).

Evolutionary algorithms are the most visible and common metaheuristic technique in the realm of multiobjective optimization. These procedures, known as MOEA (multiobjective evolutionary algorithms), dominate the multiobjective optimization research whose focal point is not the development of exact procedures¹. VEGA (Vector Evaluated Genetic Algorithm) by Schaffer (1985) has become a classic method in the MOEA literature. Since the development of VEGA, many other alternative procedures have been developed. One of the most complete references for MOEA is the book by Coello, Van Veldhuizen and Lamont (2002). Chapter 4 (MOEA Testing and Analysis) of this book provides a brief description of 11 MOEA procedures, although experimental results for only 4 of these MOEA variants (MOGA, MOMGA, NPGA and NSGA) are presented. The testing was done over a set of 6 unconstrained nonlinear multiobjective problems with bounded continuous variables. These procedures plus one more (SPEA) are recommended as the best MOEA implementations in Van Veldhuizen and Lamont (2000). For completeness, we include brief descriptions of these multiobjective optimization methods (Van Veldhuizen and Lamont, 2000):

MOGA (Multiobjective Genetic Algorithm) — Developed by Fonseca and Fleming (1998) and used to explore the incorporation of decision-maker goals to the priorities in the multiobjective search process. MOGA is based on the classical GA methodology and for each solution (chromosome) x_i , the method calculates the number of solutions n_i that dominates it. Solutions are ranked according to the n_i -value. The fitness of each solution is computed by means of linear interpolation so that the solution with the lowest rank has maximum fitness. A correction on the fitness is computed as the average of the fitness values of chromosomes with same rank value. Finally, niching methods are used to distribute chromosomes over the Pareto-optimal region.

MOMGA (Multiobjective Messy Genetic Algorithm) — Developed by Van Veldhuizen (1999) and initially used to explore the relationship between the multiobjective problem solution building blocks and their use in a MOEA search. MOMGA incorporates Horn

¹ Jones, Mirrazavi and Tamiz (2002) estimate that about 70% of the metaheuristic applications to multiobjective optimization reported in the literature are evolutionary algorithms, while 24% are based on simulated annealing and only 6% are based on tabu search.

et al.'s (1994) tournament selection and is an extension of mGA (Goldberg et al., 1990) to the multiobjective domain. The original mGA consists of three phases: initialization, primordial and juxtapositional. The modifications are related to the fitness function, the fitness evaluation, the parent selection, the operators in the primordial and juxtapositional phases and the sharing to prevent genetic drift.

NPGA (Niche—Pareto Genetic Algorithm) — Developed by Horn et al. (1994) and initially used to explore benefits of providing the set of best known solutions as input to a decision analysis technique. The method operates on a sample of the population to determine the winner between two candidate solutions to be selected and to choose one of them based on nondominance with respect to the sample taken. Since only a portion of the population is used with this technique, its computational burden is less than that of traditional Pareto ranking approaches.

NSGA (Nondominated Sorting Genetic Algorithm) — Developed by Srinivas and Deb (1994) and initially used to explore bias prevention towards certain regions of the Pareto front. The method uses the well-known sharing function approach, which has been found to maintain diversity in a population based on a parameter that measures the largest value of the distance between solutions with the same fitness. It also implements a non-dominated sorting that makes, with the extensive comparison required by the sharing function, NSGA a computationally expensive method. To overcome these and other related problems, Deb et al. (2000) propose NSGA II which alleviates these difficulties with a fast non-dominated approach and a new selection operator.

SPEA (Strength Pareto Genetic Algorithm) — Developed by Zitzler and Thiele (1999) and initially used to explore the active use of the currently known Pareto front in assigning generational fitness. The method maintains an external population at every generation storing all non-dominated solutions found so far. This population is used for combination during the application of the method. Specifically, a combined population is first constructed at each generation merging the external and the current population. All non-dominated solutions in the combined population are assigned a fitness based on the number of solutions they dominate, and dominated solutions are assigned a fitness value worse than the worst of any non-dominated one. The authors propose a clustering technique to preserve diversity among non-dominated solutions.

All these procedures are Pareto-sampling MOEAs that incorporate fitness sharing. Pareto sampling refers to techniques that use the population of solutions generated during the evolutionary process to construct an approximation of the Pareto optimal set in a single run. The Pareto ranking scheme used in each particular implementation varies. Goldberg (1989) was the first one to suggest a Pareto ranking procedure in connection with multiobjective optimization.

Fitness sharing is a specialization method for evolutionary algorithms that allows them to find multiple optima in multimodal search space. It has been used widely in connection with multiobjective function optimization and machine learning. The main idea of fitness sharing is to modify the landscape by reducing payoff in densely-populated regions and thus encouraging the search to move to unexplored regions of the solution space (Darwen and Yao, 1995).

Serafini (1992) was the first one to apply simulated annealing (SA) to a multiobjective optimization problem and this application motivated similar SA implementations. The most relevant work in SA-based procedures for multiobjective optimization can be found in Jones, Mirrazavi and Tamiz (2002).

Most tabu search (Glover and Laguna, 1997) applications to multiobjective optimization employ the so-called *independent sampling technique*. This technique is based on aggregating the objective functions by assigning a weight to each of them. Each sample consists of solving the single-objective optimization problem that results from applying a given set of weights. To obtain an approximation of the efficient frontier (\hat{E}) the procedure must be run as many times as the desired number of points, using different weight values. The performance of implementations based on independent sampling deteriorates as the need for generating more efficient solutions increases, since this is directly proportional to the number of times that the procedure must be executed. Tabu search implementations based on independent sampling are due to Hertz, et al. (1994), Dahl, et al. (1995), Gandibleux, et al. (1997), Hansen (1997), Ben Abdelaziz, et al. (1999), Alves and Climaco (2000), and Gandibleux and Freville (2000). For the most part, these procedures do not use advanced tabu search strategies for diversification and intensification.

The method developed by Caballero, Gandibleux and Molina (2004) is, to the best of our knowledge, the only tabu search implementation that employs Pareto-sampling instead

of independent sampling in the context of multiobjective combinatorial optimization. Also, this procedure, known as MOAMP (Multiobjective Metaheuristic using an Adaptive Memory Procedure), is the only implementation that is capable of include any solution visited during the search into the final approximation of the efficient frontier. MOAMP separates itself from the other tabu search implementations by looking for efficient points with an intensification process (second phase of this procedure) around an initial set of efficient points (first phase of this procedure). To build this initial set of efficient points, MOAMP carries out a series of linked tabu searches (linked means that the last point of one search becomes the initial point of the next search) where each point visited could be included in the final \hat{E} . This is achieved by checking the dominance criteria for each solution around its neighbourhood. Solutions that are not dominated are declared “possibly efficient” and are added to a list used to update \hat{E} .

The second phase of MOAMP exploits the proximate optimality principle (POP), which stipulates that good solutions at one level are likely to be found close to good solutions at an adjacent level. The POP concept may be viewed as a heuristic counterpart of the so called Principle of Optimality in dynamic programming (see section 5.5 of Glover and Laguna, 1997). The interpretation of POP within multiobjective optimization is that efficient points are “connected” by a curve inside the efficient set. This is why the second phase of MOAMP intensifies the search around the initial set of efficient points found in the first phase.

Our current interest and the purpose of this paper is to extend the basic scatter search methodology (Laguna and Martí, 2003) to tackle nonlinear multiobjective optimization problems. Some applications of scatter search to multiobjective optimization problems have been developed in recent years. Gomes da Silva, Clímaco and Figueira (2004) described a scatter search method with surrogate constraints for solving multi-dimensional knapsacks with two criteria. Garcia, et al. (2002) applied scatter search to a multiobjective location problem. Beausoleil (2004) developed MOSS (Multiobjective Scatter Search), a tabu/scatter search hybrid for nonlinear multiobjective optimization problems. We address the differences between MOSS and our proposed extension of scatter search in section 3.3.

3. Scatter Search Procedure for Multiobjective Optimization (SSPMO)

Our solution procedure (SSPMO) extends the scatter search methodology within a hybridized procedure that includes two different phases:

- Generation of an initial set of efficient points through various tabu searches
- Combination of solutions and updating of \hat{E} via a scatter search

We don't include a detailed description of the scatter search methodology because this information has been published widely in journal articles, book chapters, and conference proceedings. For a complete description, we refer the interested reader to the book by Laguna and Martí (2003). Hence, it suffices to indicate that a scatter search consists of constructing and then maintaining a reference set (*RefSet*) of solutions through the application of five methods: diversification generation, subset generation, combination, improvement, and reference set update.

Scatter search orients its explorations systematically relative to a set of reference points that typically consist of good solutions obtained by prior problem solving efforts, where the criteria for "good" are not restricted to objective function values, and may apply to sub-collections of solutions rather than to a single solution, as in the case of solutions that differ from each other according to certain specifications. The reference set is a collection of both high quality solutions and diverse solutions that are used to generate new solutions by way of applying a combination method. The meaning of diverse is especially important in the context of multiobjective optimization. As stated in Deb et al. (2000), the necessary conditions to convert an evolutionary method into a multiobjective method are both assigning fitness to population members based on non-dominated sorting and preserving diversity among solutions of the same non-dominated front. The following sections describe how our implementation deals with both conditions within the scatter search framework.

3.1 Initial Phase

Our procedure starts with the application of the first phase of MOAMP (Caballero, Gandibeux and Molina, 2004). This phase consists of linking $p+1$ tabu searches. The first tabu search starts from an arbitrary point and attempts to find the optimal solution to the problem with the single objective $f_1(x)$. Let x^1 be the last point visited at the end of this search. Then, a tabu search is applied again to find the best solution to the problem with the single-objective $f_2(x)$ using x^1 as the initial solution. This process is repeated until the problem with the single objective $f_1(x)$ is attempted one

more time starting from x^p . This phase yields the p points that approximate the best solutions to the single-objective problems that result from ignoring all but one objective function. Additional efficient solutions may be found during this phase because all visited points are checked for inclusion in \hat{E} .

Still within the MOAMP framework, we launch several tabu searches using a *global criterion method*. In this step, the aim is to minimize a function that measures the distance to the ideal solution. The ideal solution f^{\max} is that for which each criterion i (for $i = 1, \dots, p$) achieves its maximum value f_i^{\max} . Similarly, the anti-ideal solution f^{\min} is that for which each criterion i (for $i = 1, \dots, p$) achieves its minimum value f_i^{\min} . We approximate f^{\max} and f^{\min} with the maximum and minimum values, respectively, that each objective achieves in the current \hat{E} . Note that with all likelihood each objective achieves its maximum in a different point of the solution space. Hence, the purpose of determining f^{\max} is not searching for an efficient solution in X but for an ideal value in the image space in order to measure the quality of solutions generated during this step. This global criterion method follows the notion of *compromise programming* (Duckstein, 1984; Yu, 1973; Zeleny, 1973). Given that the ideal point consists of the optimal (or best known) values of the individual objective functions, compromise programming assumes that it is logical for the decision maker to prefer a point that is closer to the ideal point over one that is farther away. An essential element of compromise programming is therefore the notion of distance between points in \mathfrak{R}^p . The L_q metric ($1 \leq q \leq \infty$), a generalization of the Euclidian distance ($q = 2$), is commonly used in compromise programming. A characteristic of the L_q family is that a larger weight is given to the maximum deviation as q increases. L_1 and L_∞ are the only linear metrics out of the entire family. L_∞ is often preferred because it has been shown to lead to balanced efficient solutions. The L_∞ metric results in a min-max global criterion:

$$L_\infty(x) = \max_{i=1, \dots, p} \left\{ w_i \left(\frac{f_i^{\max} - f_i(x)}{f_i^{\max} - f_i^{\min}} \right) \right\}$$

The motivation for using this metric for the global criterion tabu searches is that a solution x is efficient for a given set of weights w if it minimizes $L_\infty(x)$. In general, a

point that minimizes an L_q distance to f^{\max} is an efficient point. The set of all points obtained in this way is called the *compromise set*. Compromise solutions have the characteristic of providing a good balance among the values of the p objective functions. Romero (1993) has shown that the best balance is achieved when using the L_∞ metric, and hence our choice. MOAMP's first phase generates random weights until *InitPhase* consecutive searches fail to produce a new efficient point in \hat{E} . Through experimentations, we have determined that the final \hat{E} obtained in this initial phase contains a diverse set of efficient points.




Figure 1. Initial tabu searches

A graphical representation of this initial phase is shown in Figure 1. The figure assumes that there are three objective functions and x^i is the best solution found by the tabu searches either during the initial single-objective problems (x^1 , x^2 and x^3) or during the executions associated with the global criterion method (x^4 , x^5 and x^6). Arcs 1, 2 and 3 represent the initial search for the optimal solution of the single-objective problems that ignore all objective problems except $f_i(x)$, for $i = 1, 2$, and 3, respectively. Arc 5, 6, and 7 show the optimization trajectory of three global criterion tabu searches that result in the compromise solutions x^4 , x^5 and x^6 .

The result of applying tabu search in the way described above is that the \hat{E} at the end of this phase contains both efficient points dominated by single objectives and points where none of the objectives dominate (i.e., they are near the compromise points). This

is graphically shown in Figure 2, where the crosses indicate the efficient solutions found during the tabu searches and the circles are the best solution found at the end of these searches.




Figure 2. Efficient solutions found during the initial phase

3.2 Scatter Search Phase

The core of our scatter search implementation consists of combining solutions that currently are considered efficient and therefore belong to \hat{E} . The solutions being combined are selected from the reference set, $RefSet \subset \hat{E}$. $RefSet$ consists of b solutions ($b > p$) and initially is constructed as follows:

1. Select the best solution in \hat{E} for each of the p objective functions and add them to $RefSet$. (Note that it is possible, but unlikely, to select fewer than p solutions in this step if a solution happens to be best for more than one objective function.)
2. Select $b-p$ solutions from $\hat{E} \setminus RefSet$ that maximize the distance between them and those solutions already in $RefSet$. Since the solutions are selected sequentially, the distance measure is updated after each selection. Distance is measured with a normalized L_∞ metric.

In the basic scatter search design, the initial reference set is built according to the reference set update method. This method starts with the selection of the best $b/2$ solutions from a set P of previously generated solutions. These $b/2$ solutions are added

to *RefSet* and deleted from P . In our multiobjective implementation of SS, we use \hat{E} as the set P and p instead of $b/2$. The role of P is different than in the basic design for single-objective optimization since we maintain and update P during the search to record the efficient solutions while in the basic SS design P is discarded after the *RefSet* is built.

A list of solutions that have been selected as reference points is kept to prevent the selection of those solutions in future iterations. Therefore, every solution that is added to *RefSet* is also added to *TabuRefSet*. The size of *TabuRefSet* increases as the search progresses because this memory function is an explicit record of past reference solutions. The motivation for creating and maintaining *TabuRefSet* is that we would like to obtain a final \hat{E} of adequate density. That is, we would like to encourage a uniform generation of points in the efficient frontier and avoid gaps that may be the result of generating too many points in one region while neglecting other regions.

A linear combination method is used to combine reference solutions. All pairs of solutions in *RefSet* are combined and each combination yields four new trial solutions. Let x^i and x^j be the reference solutions being combined, then four trial solutions x^k are obtained with the following line search:

$$x^k = \lambda x^i + (1 - \lambda)x^j \quad \lambda = -1/3, 1/3, 2/3 \text{ and } 4/3$$

The selection of this combination method follows the suggestions made by Glover (1994) in connection with non-linear optimization of single-objective problems. This method does not rely in randomization and includes both convex and non-convex combinations of the reference solutions. Both of these characteristics make the method appropriate within the philosophy of the scatter search framework.

Each of the new trial solutions is subjected to an improvement method. Tabu search is the mechanism used to improve new trial solutions. The objective function that guides the search for an improved solution is the L_∞ metric with $w_i = 1$ for $i = 1, \dots, p$. The f_i^{\max} and f_i^{\min} are calculated considering only the reference points x^i and x^j whose combination resulted in the trial point x^k currently being improved. Solutions generated during this improvement phase are tested for possible inclusion in \hat{E} . Note that any addition to \hat{E} may cause some previously efficient points to become dominated

and therefore expelled from the set. Figure 3 is a graphical representation of the improvement phase within the scatter search. By creating the ideal point from the reference points that originated the four new trial solutions, the procedure attempts to find efficient points that “fill the gap” between x^i and x^j .




Figure 3. Graphical representation of the improvement method

Once all the solution pairs in $RefSet$ are combined and the new trial solutions are improved, the procedure updates the reference set in preparation for the next scatter search iteration. The first step in the updating process is to choose the best solutions according to each of the objective functions taken separately. This is the same first step as in the construction of the initial reference set. The step does not consider whether these efficient points belong to $TabuRefSet$. The remaining $b-p$ reference solutions are chosen as follows:

1. For each solution $x \in \hat{E} \setminus TabuRefSet$, a normalized (using the range of each function) L_∞ is calculated. The normalization is such that all distances are between 0 and 1. The distance calculations involve x and all the solutions in $TabuRefSet$. Let the minimum of these normalized distances be $L_\infty^{\min}(x)$. This minimum normalized distance is used as the probability that x is declared eligible as a reference solution. That is, the larger the minimum distance between the candidate solution x and all the solutions in $TabuRefSet$ the better

- the chance for x of becoming eligible as a reference solution. A uniform random number is generated and if it is less than $L_{\infty}^{\min}(x)$ then x is declared eligible.
2. From the list of all eligible solutions, we choose sequentially the $b-p$ solutions with largest minimum distance to *TabuRefSet*. The distance is measured against *TabuRefSet* instead of *RefSet* to move away from areas that have been explored in the past by virtue of combining former reference points. Also, *TabuRefSet* is updated after each selection in order to avoid choosing points that are too close to each other .

The scatter search continues until the mean value of $L_{\infty}^{\min}(x)$ for the set of eligible solutions in step 2 above falls below a pre-specified threshold *MeanDist*.

3.3 Differences between MOSS and SSPMO

MOSS (Beausoleil, 2004) is related to SSPMO in the sense that it applies similar techniques to the same class of problems. The tabu search elements included in MOSS create restrictions that are used to prevent moves toward solutions that are “too close” to previously visited solutions. The definition of distance is adjusted during the search to control the number of solutions that are classified tabu during given iterations. These solutions are not allowed to be combined within the scatter search iterations. A sequential fan candidate list strategy is used to explore solution neighborhoods. A weighted linear function is used to aggregate the objective function values and provide a way of choosing the best move in a neighborhood. Long term memory is used to encourage diversification by giving incentives to sampling neglected subranges of values within the feasible range of each decision variable.

Although both MOSS and SSPMO are optimization procedures based on a hybridization of scatter and tabu search, there are significant differences between them:

1. MOSS does not utilize a guiding function to direct the tabu searches as SSPMO does. Instead, MOSS employs an aspiration level criterion to avoid moving the search to a dominated point. The guiding functions in SSPMO are both the objective functions in the problem and the L_{∞} compromise functions.
2. The update of the *RefSet* in MOSS mimics the updating procedure that is traditional in SS implementations for single-objective problems. In particular, a reference point may stay in the *RefSet* for several iterations. In contrast,

SSPMO updates the *RefSet* around the best solutions found for each objective function. This is achieved with the inclusion of all former reference points in to *TabuRefSet*.

3. The subset generation method in MOSS considers the Kramer choice function to create the solution subsets to be combined. SSPMO, we take into consideration the characteristics of a point according to its $L_{\infty}^{\min}(x)$ value. This produces a uniform sampling of the efficient frontier and avoids the selection of points for a combination subset that are concentrated in the same area.

4. Computational Experiments

The performance of our solution method depends on the value of three search parameters: *InitPhase*, b and *MeanDist*. The *InitPhase* parameter controls the termination of the initial phase of the procedure. In particular, the initial phase terminates after *InitPhase* consecutive searches fail to produce a new efficient point. The most effective value of 3 for *InitPhase* was determined with a series of tuning experiments. Our experiments also showed that the procedure is almost insensitive to specific values of b in the $(2p, 3p)$ range. We selected $b = 2p$ for our computational tests. The *MeanDist* value controls the termination of the scatter search phase. The parameter value is a threshold of required average distance at which potential reference points must be from the current *RefSet*. Since the distance is normalized to be between 0 and 1, the *MeanDist* value also should be within such a range. *MeanDist* values close to zero make the scatter search phase longer. We employ a value of 0.1 for this parameter because this value provides a good balance between computational effort and solution quality.

We perform tests on a Pentium 4 at 2.4 GHz with three different sets of problems from the literature. The ZDT and DTLZ instances respectively from Zitzler, Deb and Thiele (2000) and Deb, et al. (2002) are well-known, standard problems in the multiobjective optimization literature. We also use instances from Deb (1999) that include non-convex and disjoint efficient frontiers as well as non-uniform distribution of efficient points. A summary of the characteristics of the 27 problems in our test set is shown in Table 1.

Table 1. Characteristics of test problems

<i>Set</i>	<i>Number of problems</i>	<i>Number of variables</i>	<i>Number of objectives</i>
Deb	18	2	2
DTLZ	4	12	3
ZDT	5	10 and 30	2

We don't include a detailed description of the instances because this information is available in the published articles associated with each problem set. However, computer code for all the problems in our test set is available from julian.molina@uma.es.

The main goal of our experiments is to show that the proposed procedure is capable of creating better approximations of the efficient frontiers than exiting methods for nonlinear multiobjective optimization. Specifically, we are interested in comparing the performance of SSPMO with MOAMP, SPEA2 and MOSS. The performance measures that we employ are:

1. *Number of points*: This refers to the ability of finding efficient points. We assume that the decision maker prefers more rather than fewer efficient points.
2. *SSC*: This metric suggested by Zitzler and Thiele (1999) measures the size of the space covered (SSC). In other words, SSC measures the volume of the dominated points. Hence, the larger the SSC value the better.
3. *k-distance*: This density estimation technique used by Zitzler, Laumanns and Thiele (2001) in connection with the computational testing of SPEA2 is based on the k^{th} nearest neighbor method of Silverman (1986). The metric is simply the distance to the k^{th} nearest efficient point. We use $k = 5$ and calculate both the mean and the max of k -distance values. The k -distance value is such that the smaller the better in terms of frontier density.
4. M_1^* : This metric suggested by Zitzler (1999) consists of calculating an average Euclidean distance (L_2) between the estimated efficient frontier and the Pareto-optimal set. Clearly, smaller values of M_1^* indicate better approximations of the Pareto-optimal set.

5. $C(A,B)$: This is known as the coverage of two sets measure (Zizler and Thiele, 1999). $C(A,B)$ represents the proportion of points in the estimated efficient frontier B that are dominated by the efficient points in the estimated frontier A .

Table 2 shows means (first row) and standard deviations (second row) of the first 4 measures described above for the methods that we have tested. The last column of this table shows the average and standard deviation of the CPU seconds associated with each procedure.

Table 2. Average and standard deviation values of six performance measures over the 27 problems.

<i>Methods</i>	<i>N. of points</i>	<i>k-distance (mean)</i>	<i>k-distance (max)</i>	<i>SSC</i>	M_1^*	<i>Time</i>
SSPMO	1622.185	0.006	0.149	0.534	0.061	22.926
	1708.762	0.007	0.188	0.287	0.160	25.859
MOAMP	764.963	0.032	0.144	0.482	0.095	6.111
	1011.204	0.036	0.119	0.291	0.236	10.467
SPEA2	68.889	0.232	0.450	0.507	0.073	12.519
	92.696	0.274	0.327	0.273	0.187	13.791
MOSS	83.111	0.055	0.527	0.498	0.066	237.630
	28.860	0.031	0.347	0.303	0.168	446.714

The results in Table 2 indicate that SSPMO is capable of finding efficient frontiers with a large number of points and high density, as indicated by the small k -distance values. The number of points is not an input parameter of SSPMO. The termination criteria for SSPMO allow it to find a large number of points while improving on the density of \hat{E} . This is not the case of MOAMP, which can construct efficient frontiers with a fair amount of points but at the same time is not capable of identifying gaps that ultimately result in sparse areas of the frontiers. The SSC and M_1^* values also show a superior performance of SSPMO over the competing approaches. Regarding execution time, it is clear that the large number of efficient points that SSPMO is able to generate makes the procedure slower than MOAMP and SPEA2. The tradeoff, however, still favors SSPMO if we consider the time per point in the final efficient set.

Another way of looking at the results of the same experiment is to count the number of times that each procedure performs best according to each performance measure. Table 3 shows the summary of such computation.

Table 3. Number of “wins” for each method over the 27 problems

<i>Methods</i>	<i>N. of points</i>	<i>k-distance (mean)</i>	<i>k-distance (max)</i>	<i>SSC</i>	M_1^*	<i>Time</i>
SSPMO	25	22	15	15	5	1
MOAMP	2	5	11	9	15	21
SPEA2	0	0	0	1	5	3
MOSS	0	0	1	2	2	2

The values in Table 3 confirm the high performance of SSPMO according to most of the standard performance measures. SSPMO is inferior to MOAMP only in execution time and number of best M_1^* values. It is interesting to note that although MOAMP achieves the best M_1^* values 15 times, the average M_1^* value still favors SSPMO (as shown in Table 2).

The $C(A,B)$ measure allows us to make a comparison according to the dominance of one efficient frontier over another. Table 4 shows the average $C(A,B)$ values over the entire set of test problems.

Table 4. Coverage of two sets

A/B	SSPMO	MOAMP	SPEA2	MOSS
SSPMO	0.000	0.160	0.261	0.197
MOAMP	0.230	0.000	0.196	0.181
SPEA2	0.106	0.140	0.000	0.113
MOSS	0.115	0.111	0.093	0.000

The values in Table 4 show that the efficient points generated by SSPMO tend to dominate those generated by other methods. That is, $C(\text{SSPMO}, -) > C(-, \text{SSPMO})$ except in the case of $C(\text{MOAMP}, \text{SSPMO})$. It is interesting to point out that even though $C(\text{MOAMP}, \text{SSPMO}) > C(\text{SSPMO}, \text{MOAMP})$, SSPMO performs better than MOAMP according to the *SSC* and *k-distance* measures. The explanation is that while SSPMO approximations are better distributed along the frontier MOAMP dominates SSPMO in some specific areas. This is not surprising, given the tendency of MOAMP to over-intensify in some areas of the efficient frontier while neglecting others.

To conclude this section on computational experiments, we examine the results obtained by SSPMO and SPEA2 on two problem instances in the Deb set (see Table 1). The reason for choosing these two problems is to gain additional insight of the strengths and weakness of these approaches. In particular, we would like to address the issue of finding efficient frontiers that are dense and whose points are uniformly

distributed. This is possible only by including mechanisms that can detect gaps in the approximation and dynamically adapt the search. First, we examine the second problem in section 4.1 of Deb (1999). The approximations of the efficient frontier obtained by SSPMO and SPEA2 are shown in Figure 4.




Figure 4. Approximations of the efficient frontier of the second problem in section 4.1 of Deb (1999)

SPEA2 includes a mechanism to “fine-tune” the approximation of the efficient frontier once a pre-specified number of efficient points have been found. The mechanism is capable of identifying zones of the efficient frontier where there is an excessive concentration of points with the purpose of eliminating some of them. Once these points have been deleted, however, SPEA2 is incapable of identifying zones that have been under-explored and this produces the gaps in the efficient frontier that are evident in Figure 4. SSPMO, on the other hand, does not limit a priori the total number of efficient points and instead it focuses on achieving a desired density along the entire efficient frontier. In the specific case depicted in Figure 4, SSPMO finds 1,660 efficient points that result in an average k -distance value of 0.0009 while SPEA2 finds only 60 efficient points with an average k -distance value of 0.0553.

The lack of density of the approximations found by SPEA2 is more evident in problems with a disjoint efficient frontier. Figure 5 shows the approximations found by SPEA2 and SSPMO for the problem described in section 5.1.3 of Deb (1999) with parameters $q = 8$ and $\alpha = 2$.




Figure 5. Approximations of the efficient frontier of the problem in section 5.1.3 of Deb (1999) with parameters $q = 8$ and $\alpha = 2$

The approximations depicted in Figure 5 correspond to 128 SSPMO points and 22 SPEA2 points. The corresponding average k -distance values are 0.0267 and 0.2075 for SSPMO and SPEA2, respectively. This figure shows that SSPMO visits and fills all the disjoint parts of the efficient frontiers because it is able to detect the presence of gaps and react accordingly. On the other hand, SPEA2 is able to achieve good approximations only on the extremes of the efficient frontier because it lacks a mechanism to detect areas with low density and redirect the search.

4.1 Additional Analysis of Results

The scatter search phase of SSPMO is designed to address the main problem found in the approximation methods for multiobjective optimization that can be found in the literature. Namely, we attempt to find dense efficient frontiers with a sufficient number of points that are well-distributed. As our experiments show, some of the methods in

the literature either produce approximations with an insufficient number of points (SPEA2) or the approximations contain gaps (MOAMP) or the points are not well-distributed (MOSS). These deficiencies are caused by the following design problems:

1. The size of \hat{E} is pre-specified as an input parameter (SPEA2 and MOSS). This entails that the analyst must estimate the number of points needed for a good approximation. In the case of SPEA2, this is not an easy question to answer. A small number results in sparse approximation. A large number prevents the launching of a diversification phase and results in an inferior distribution of points in the final approximation of the efficient frontier.
2. While diversifying criteria are included in some methods (SPEA2 and MOSS), none include “filling criteria”. That is, these methods include mechanisms to delete or penalize solutions too close to other solutions, but they don’t complement these mechanisms with others that encourage the search to intensify around regions with low density of points.
3. The stopping criteria of the existing methods are not related with density estimations or the presence of gaps. This causes either to stop the search before neglected areas are visited or to continue searching even after a well-distributed approximation has been found.

SSPMO addresses these deficiencies by including three strategies: 1) the search focuses on filling the gaps between carefully selected reference points, 2) low density areas are visited by moving the search away from previous reference points, and 3) the search terminates only if a measure designed to identify the largest gap falls below a specified value.

5. Conclusions

We have described the development and implementation of a metaheuristic procedure for the optimization of multiobjective non-linear functions. Our procedure extends the application of scatter search in an innovative way by assigning new roles to the reference set and strategically including tabu search elements. One of our main goals has been to test our procedure employing all the problem instances that we could find in the literature. In order to make a valid comparison against competing procedures, we have used several metrics as well as graphical output (when applicable). Our computational experiments show that SSPMO has merit when compared to approximation procedures that are also metaheuristic in nature.

Throughout the development of our procedure and the ensuing experimentation we have been able to identify three issues that limit the performance of exiting procedures. The identification these issues (summarized in section 4.1) have allowed us to not only design a more robust solution procedure but also explain why the procedure performs at a higher level than the ones we have used for comparison.

References

Alves, M. J. and J. Climaco (2000) “An Interactive Method for 0-1 Multiobjective Problems Using Simulated Annealing and Tabu Search,” *Journal of Heuristics*, vol. 6, no. 3, pp. 385-403.

Beausoleil, R. P. (2004) “MOSS: Multiobjective Scatter Search Applied to Non-linear Multiple Criteria Optimization,” to appear in the *European Journal of Operational Research*.

Ben Abdelaziz, F., S. Krichen and J. Chaouachi (1999) “A Hybrid Metaheuristic for the Multiobjective Knapsack Problem,” *Meta-Heuristics: Advances and Trends in Local Search Paradigms for Optimization*, in S. Voss, S. Martello, I. Osman and C. Roucairol (eds.), Kluwer Academic Publishers: Boston, pp. 205-212.

Caballero, R., X. Gandibleux, and J. Molina (2003) “MOAMP- A Generic Multiobjective Metaheuristic using an Adaptive Memory, Technical Report, University of Valenciennes, France.

Coello, C., D. A. Van Veldhuizen and G. B. Lamont (2002) *Evolutionary Algorithms for Solving Multi-Objective Problems*, Kluwer Academic Publishers: Boston, ISBN 0-306-46762-3.

Dahl, G., K. Jörnsten and A. Lokketangen (1995) “A Tabu Search Approach to the Channel Minimization Problem,” Paper presented at ICOTA'95, Chengdu, China.

Darwen P. and X. Yao (1995) “A Dilemma for Fitness Sharing with a Scaling Function,” in *Proc. of the 1995 IEEE International Conference on Evolutionary Computation (ICEC'95)*, Perth, Australia, IEEE Press: New York, NY 10017-2394, pp. 166-171.

Deb, K. (1999) “Multi-Objective Genetic Algorithms: Problem Difficulties and Construction of Test Problems,” *Evolutionary Computation*, vol. 7, no. 3, pp. 205-230.

Deb, K., A. Pratap, S. Agarwal and T. Meyarivan (2000) “A Fast and Elitist Multi-Objective Genetic Algorithm-NSGA-II,” KanGAL Report Number 2000001.

Deb, K., L. Thiele, M. Laumanns and E. Zitzler (2002) “Scalable Multi-Objective Optimization Test Problems,” in Congress on Evolutionary Computation (CEC'2002) Piscataway, New Jersey, IEEE Service Center, vol. 1, pp. 825-830.

Erhgott, M. and X. Gandibleux (2000) “A Survey and Annotated Bibliography on Multiobjective Combinatorial Optimization,” *OR Spektrum*, vol. 22, pp. 425-460.

Fonseca, C. M. and P. J. Fleming (1998) “Multiobjective Optimization and Multiple Constraint Handling with Evolutionary Algorithms – Part I: A Unified Formulation,” *IEEE Transactions on Systems, Man and Cybernetics – Part A: Systems and Humans*, vol. 28, no. 1, pp. 26–37.

Gandibleux, X., N. Mezdaoui and A. Freville (1997) “A Tabu Search Procedure to solve Multiobjective Combinatorial Optimization Problems,” *Advances in Multiple Objective and Goal Programming*, R. Caballero, F. Ruiz and R. Steuer (eds.), Lecture Notes in Economics and Mathematical Systems, Springer: Berlin, vol. 455, pp. 291-300.

Gandibleux, X and A. Freville (2000) “Tabu Search Based Procedure for Solving the 0-1 Multiobjective Knapsack Problem: The Two objectives Case,” *Journal of Heuristics*, vol. 6, pp. 361-383.

Garcia, F., B. Melian, J. A. Moreno and J. M. Moreno-Vega (2002) “Scatter Search for Multiple Objective p-facility Location problems,” Groupe de Travail ROADEF sur la Programmation Mathématique MultiObjectif, 6ème journée de travail.

Goldberg, D. E. (1989) *Genetic Algorithms in Search, Optimization and Machine Learning*, Addison-Wesley Publishing Company: Reading, Massachusetts.

Goldberg, D. E., Deb, K. and Korb, B. (1990) “Messy Genetic Algorithms Revisited: Studies in Mixed Size and Scale,” *Complex Systems*, vol. 4, pp. 415-444.

Glover, F. (1994) "Tabu Search for Non-Linear and Parametric Optimization (with links to Genetic Algorithms)," *Discrete Applied Mathematics*, vol. 49, pp. 231-255.

Glover, F. and M. Laguna (1997) *Tabu Search*, Kluwer Academic Publishers: Boston, ISBN: 0-7923-8187-4, 408 pp.

Gomes da Silva, C., J. Climaco and J. Figueira (2004) "A Scatter Search Method for the Bi-criteria Multi-dimensional {0,1}-Knapsack Problem Using Surrogate Relaxation," *Journal of Mathematical Modeling and Algorithms*, vol. 3, n°3, pp. 183-208.

Hansen, M. P. (1997) "Tabu Search for Multiobjective Optimization: MOTS," 13th International Conference on Multiple Criteria Decision Making, Cape Town, South Africa.

Hertz, A., B. Jaumard, C. Ribeiro and F. W. Formosinho (1994) "A Multicriteria Tabu Search Approach to Cell Formation Problems in Group Technology with Multiple Objectives," *Recherche Operationelle/Operations Research*, vol. 28, no. 3, pp. 303-328.

Horn, J., N. Nafpliotis, and D. E. Goldberg (1994) "A Niche Pareto Genetic Algorithm for Multiobjective Optimization," in *Proceedings of the First IEEE Conference on Evolutionary Computation*, Z. Michalewicz (ed.), IEEE Press: Piscataway, New Jersey, pp. 82-87.

Jones, D. F., S. K. Mirrazavi and M. Tamiz (2002) "Multiobjective Metaheuristics: An Overview of the Current State of the Art," *European Journal of Operational Research*, vol. 137, pp. 1-9.

Laguna, M. and R. Martí (2003) *Scatter Search: Methodology and Implementations in C*, Kluwer Academic Publishers: Boston, ISBN: 1-4020-7376-3, 312 pp.

Serafini, P. (1992) "Simulated Annealing for Multiobjective Optimization Problems," In *Proceedings of the 10th International Conference on Multiple Criteria Decision Making*, Taipei, Taiwan, vol. I, pp. 87-96.

Silvermann, B.W. (1986) *Density Estimation for Statistics and Data Analysis*, Chapman and Hall, London.

Srinivas, N. and K. Deb (1994) “Multiobjective Optimization Using Nondominated Sorting in Genetic Algorithms,” *Evolutionary Computation*, vol. 2, no. 3, pp. 221–248.

Van Veldhuizen, D. A. (1999) *Multiobjective Evolutionary Algorithms: Classifications, Analyses, and New Innovations*, Ph.D. Thesis, AFIT/DS/ENG/99-01, Air Force Institute of Technology, Wright-Patterson AFB, Ohio.

Van Veldhuizen, D. A. and G. B. Lamont (2000) “Multiobjective Evolutionary Algorithms: Analyzing the State-of-the-Art,” *Evolutionary Computation*, vol. 8, no. 2, pp. 125-147.

Yu, P. L. (1973) “A Class of Solutions for Group Decision Problems,” *Management Science*, vol. 19, pp. 936-946.

Zeleny, M. (1973) “Compromise Programming,” in *Multiple Criteria Decision Making*, J. L. Cochrane and M. Zeleny (eds.) University of South Carolina Press, Columbia, pp. 262-301.

Zitzler, E. (1999) *Evolutionary Algorithms for Multiobjective Optimization: Methods and Applications*, Ph D. thesis, Swiss Federal Institute of Technology (ETH) Zurich, Switzerland.

Zitzler, E., K. Deb and L. Thiele (2000) “Comparison of Multiobjective Evolutionary Algorithms: Empirical Results,” *Evolutionary Computation Journal*, vol. 8, no. 2, pp. 125-148.

Zitzler, E., M. Laumanns and L. Thiele (2001) “SPEA2: Improving the Strength Pareto Evolutionary Algorithm,” Technical Report 103, Computer Engineering and Networks Laboratory (TIK), Swiss Federal Institute of Technology (ETH), Zurich, Switzerland.

Zitzler, E. and L. Thiele (1999) “Multiobjective Evolutionary Algorithms: A Comparative Case Study and the Strength Pareto Approach,” *IEEE Transactions on Evolutionary Computation*, vol. 3, no. 4, pp. 257–271.