PESA-II: Region-based Selection in Evolutionary Multiobjective Optimization

David W. Corne, Nick R. Jerram, Joshua D. Knowles, Martin J, Oates
Department of Computer Science, University of Reading, UK
{D.W.Corne, N.R.Jerram, J.D.Knowles}@reading.ac.uk, moates@btinternet.com
FAX: +44(0) 118 975 1994, TEL: +44 (0) 118 931 8983

Abstract

We describe a new selection technique for evolutionary multiobjective optimization algorithms in which the unit of selection is a hyperbox in objective space. In this technique, instead of assigning a selective fitness to an individual, selective fitness is assigned to the hyperboxes in objective space which are currently occupied by at least one individual in the current approximation to the Pareto frontier. A hyperbox is thereby selected, and the resulting selected individual is randomly chosen from this hyperbox. This method of selection is shown to be more sensitive to ensuring a good spread of development along the Pareto frontier than individual-based selection. The method is implemented in a modern multiobjective evolutionary algorithm, and performance is tested by using Deb's test suite of 'T' functions with varying properties. The new selection technique is found to give significantly superior results to the other methods compared, namely PAES, PESA, and SPEA; each is a modern multi-objective optimization algorithm previously found to outperform earlier approaches on various problems.

1 Introduction

Standing on the shoulders of seminal research and development in the area of multiobjective evolutionary algorithms (MOEAs), such as the Niched Pareto Genetic Algorithm (Horn et al., 1994; Horn and Nafpliotis, 1994), and the Non-Dominated Sorting method (Srinivas and Deb, 1994), the last half-decade has seen an explosion of interest and development of more capable MOEAs. The techniques that have recently emerged seem to provide fast and effective approximations to the Pareto frontier for a variety of benchmark problems. These new methods include, among others, SPEA (Strength Pareto Evolutionary Algorithm –Zitzler and Thiele, 1999), PAES (Pareto Archived Evolution Strategy – Knowles and Corne, 2000), M-PAES (Memetic PAES – Knowles and Corne, 2000a), PESA (Pareto Envelope based Selection – Corne et al, 2000), MOMGA (Multi-Objective Messy Genetic Algorithm, Van Veldhuizen and Lamont, 2000), and NSGA-II (Non-Dominated Sorting genetic Algorithm II – Deb et al, 2000).

PAES, PESA, SPEA, and NSGA-II can each be considered to be 'basic' MOEAs in the sense that their flow of control is essentially a pure evolutionary algorithm framework, while the differences between them amount to explorations of various different ways to do selection and population maintenance in multiobjective spaces. Methods such as M-PAES and MOMGA, on the other hand, are more sophisticated algorithm designs in which a pure evolutionary framework is eschewed in favour of a hybrid or multi-stage flow of control. M-PAES, for example, is a memetic algorithm in which population based search is hybridised with local search, while MOMGA is a messy genetic algorithm (Goldberg et al, 1991) adapted for use in multiobjective search.

In this paper we are interested in the 'basic' evolutionary multiobjective framework, and will therefore no longer consider M-PAES, MOMGA, and other such methods, but the technique developed may of course be incorporated in sophisticated MOEAs such as M-PAES and MOMGA in future work. We describe a variation on how to do selection in basic MOEAs, and compare an MOEA which uses this technique to each of PAES, PESA and SPEA on a variety of test problems. We have not yet compared with NSGA-II, which is an omission we hope soon to rectify.

We should also mention that much impressive multiobjective optimisation work is being done in the fields of multiple criteria decision making (MCDM) and operations research. Until recently, there has been little crosstalk between these communities and the MOEA community. Strong-performing algorithms emerging from these areas include a variety of local-search based multiobjective techniques, e.g. Czyzak and Jaszkiewicz (1998), Gandibleux et al. (1996), and Hansen (1996; 1997). Comparison of such methods with modern MOEAs has been done little so far, although recent work by Zitzler and Thiele (1999) and Knowles and Corne (2000) indicate that methods such as PAES and SPEA are at least comparable in quality to these other methods.

The remainder of this paper is set out as follows. In section 2 we briefly review selection schemes in modern evolutionary multiobjective algorithms, and introduce the simple concept of region-based selection. Some simple analysis is done to suggest why region-based selection may be favoured over other methods, in terms of its maintained strong bias towards developing isolated regions of the Pareto front. In section 3 we note the algorithms and describe the test functions used in later experiments. These experiments are described in section 4 and their results are presented in section 5, and we have a concluding discussion in section 6.

2 Region-Based Selection in Evolutionary Multiobjective Algorithms

2.1 Individual-Based Selection

We will use Figure 1 to illustrate the main selection schemes used in current multiobjective evolutionary algorithms. In the figure, a number of points are plotted in objective space for a supposed two-objective problem, and we imagine that the goal is to minimize along both axes (in the directions shown by the arrows). Objective space is divided into squares (generally, 'hyperboxes' in higher dimensional objective spaces). In both PAES and PESA, the algorithms incorporate a subdivision of the objective space into hyperboxes as shown in the figure. In PESA, information concerning the occupation of hyperboxes is used for selection as follows. An archive is maintained containing only non-dominated solutions, and as such represents the algorithm's current approximation to the Pareto frontier. Selection is only from this archive. The selective fitness of an individual is simply the number of other solutions which occupy the same hyperbox as that individual. This is called the 'squeeze fac-

tor'. Tournament selection (or any other basic selection scheme) can then be used to select parents with a bias towards small squeeze factors.

In PAES, selection is rather a different affair since PAES is essentially a local search method. There is just one current solution at any time, and this is therefore always selected to be the parent of a mutant. However, when the mutant and current solution are non-dominated, a decision has to be made as regards which will become the new current solution (which can be seen indirectly as selecting the parent for the next iteration). The full details of this decision are in Knowles and Corne (2000) but for present purposes we note that, like PESA, it makes use of hyperbox occupancy.

Selection in SPEA is done via a 'Strength Pareto' scheme developed by Zitzler and Thiele (1999). This is a way of assigning selective fitness to an individual based on the number of individuals in the population which it covers – an individual *covers* another if it dominates it, or is equal to it. This method therefore relies on having population members around which are not in the current approximation to the Pareto front. In SPEA, this is organised by having two populations, an *internal* and *external* population. The external population only contains non-dominated individuals, while the internal population contains the latest crop of children produced via genetic operators, and as such may contain individuals which are dominated by members of the external population. Figure 1, may represent the combined populations at a snapshot in a run of SPEA. The point labelled X is nondominated, and hence in SPEA's external population, and it dominates two members of the external population (those contained in the region enclosed by the lines emanating from X. The Strength measure for a nondominated individual is just the number of individuals in the internal population which it covers. Strength measures for members of the internal population are derived by summing the strengths of the external population individuals which cover them. Selection is biased towards minimising this strength figure, thus preferring the exploration of less populated regions of the objective space. So, in Figure 1, Y will have a better selective strength than X.

Finally, NSGA-II uses a rather different selection technique which has been found to be both highly efficient and to perform very favourably in comparison to others. In NSGA-II, a selective fitness measure is derived for an individual by first finding the distance to the closest other individual to it for each objective in turn. The product of these distances gives a hypervolume which in turn estimates the isolation of this individ-
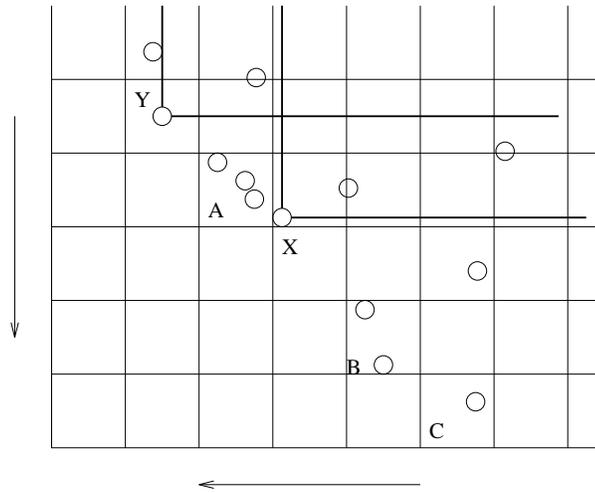
Figure 1: Illustration of Selection Methods in Modern MOEAs.

ual. Selection is therefore biased towards individuals with a high isolation value. In Figure 1, for example, the points in box A would have a low isolation value, but that of point Y would be relatively high.

Each of the selection techniques is oriented towards maintaining development of the Pareto front in a well spread manner. That is, by biasing search in the region of relatively lonely regions of the current approximation to the Pareto front, the aim is to promote an even spread of individuals along it. The main difference between the methods is the precise way in which the degree of isolation of an individual is estimated. PAES and PESA use hyperbox counts, NSGA-II uses distance to nearby individuals, and SPEA uses a somewhat indirect method which estimates an individual's isolation based on how many previously generated individuals it covers. An aspect which all of these methods share is that selection is *individual-based*. That is, the unit of selection is an individual. The different variations can therefore be seen as imposing different distributions of selection probability on the individuals, with the goal of achieving higher probabilities for those in isolated regions than those in crowded regions.

2.2   Region-Based Selection

Region-based selection provides an alternative, in which the above goal is achieved more directly. In region-based selection, the unit of selection is now a hyperbox, rather than an individual. A selective fitness is derived for a hyperbox. Using any standard selection method, a hyperbox is therefore selected, and the resulting individual chosen for genetic operations is randomly chosen from the selected hyperbox. In

Figure 1, for example, hyperbox C would have a better selective fitness than hyperbox B, which in turns would have a better selective fitness than hyperbox C.

The following simple analysis suggests why region-based selection may be favoured over an individual based scheme. Assume we are using binary tournament selection without replacement in both cases. That is, binary tournament selection is used to select an individual based on selective fitnesses, whether those selective fitness are individual-based (measures of isolation such as strength or Deb's crowded-comparison measure (Deb et al, 2000)) or hyperbox-based. It is worth first considering a pathological case in which just two hyperboxes are occupied in the current approximation to the Pareto front. One is occupied by 9 individuals, and the other by a single individual. We will also assume, which seems reasonable, that the single individual is the most isolated in respect of the typical individual based selective fitness measures we have considered.

With binary tournament selection, the chance of selecting the best individual (the most isolated one) in an individual based selection scheme will be $1 - (9/10)^2 = 0.19$. The chance of selecting any one of the 9 overcrowded individuals will therefore be 0.81. This does not seem to provide suitably high bias towards development in the less-crowded region. With region-based selection, however, the units of selection are the two occupied hyperboxes. The chance of choosing the least-occupied box (and hence choosing the best individual) is $1 - (1/2)^2 = 0.75$. The chance of choosing any one of the more crowded individuals is therefore 0.25. With individual based selection in this example,

we are actually more likely to choose a highly non-isolated individual than the most isolated one. With region-based selection, we are three times more likely to choose the isolated than any of the non-isolated individuals.

We will now take a slightly more formal look, stepping away from the pathological case to see what may be the more typical situation. We will remain interested in the relative probabilities of choosing a most isolated individual over a most crowded individual, and will continue to assume the use of binary tournament selection. Consider an approximation to the Pareto front which has b occupied hyperboxes, with n_i individuals in box i, and P individuals altogether in occupied hyperboxes, such that $\sum_{i=1}^{b} n_i = P$. Assume now, with a slight loss of generality, that a single hyperbox j has the largest b_i and another single hyperbox has the smallest n_i. The numbers of individuals in these least and most crowded boxes will be l and m respectively.

When using individual based selection, the chance of choosing an individual from the least crowded box will be $1 - ((P - l)/P)^2$. The corresponding term for the most crowded box is simply $(m/P)^2$; the ratio of these probabilities $(2Pl - l^2)/m^2$. When m is high with respect to l, the relative chance of choosing an isolated individual rather than a crowded one reduces fairly sharply, this would seem to unreasonably draw selective attention towards the crowded regions. In contrast, the corresponding ratio for region-based selection turns out to be $2b - 1$. It is unaffected by the relative numbers of individuals in the different boxes, and never less than 1 (in fact, always at least 3 when more than 1 hyperbox are occupied).

It might be thought that the same effect – that is, duly high attention to isolated regions rather than crowded ones, could be achieved by individual-based selection with a higher tournament size. However, notice that the chance of choosing an individual from the most crowded box in this case will be $(m/P)^k$, where k is the tournament size. When the tournament size is large, this will drop very sharply with a large population and a fairly even distribution of individuals among them. In these conditions, the chance of choosing an individual from the least crowded box would become unacceptably low, affecting the exploratory capabilities of the algorithm.

2.3 Complexity Issues

Here we briefly reflect on the complexity issues inherent in individual-based versus region-based selection schemes. In the context of multi-objective search, the issue of main interest to us here is the complexity of calculating selective fitness based on crowding in phenotype space. For simplicity, we will assume generational approaches in which a new population of size n is in every generation.

Individual-based selection requires estimates of the degree of 'isolation' of each individual. Accurate estimation of the relative isolation of the individuals in a population of size n would of course require n^2 comparison operations, where the distances between all distinct pairs are calculated. However, it has been found, in both NSGA-II and PAES, for example, that approximate estimates of isolation can be achieved more quickly, with quite adequate results. For example, the metric used to approximate isolation in NSGA-II (Deb et al., 2000) requires $\mathcal{O}(k \cdot n \log n)$ time, where k is the number of objectives.

In region-based selection using hyperboxes, the key computational concern is to caclulate a hyperbox ID for each individual. As indicated in Knowles & Corne (2000), in a k-objective problem using a grid of g^k hyperboxes, only $\mathcal{O}k \cdot n$ comparison operations need be made per generation. Efficiency is improved if g is a power of 2, but the broad order of complexity is just linear in n. A single pass through the hyperbox IDs then easily yields the selective fitness information required by either region-based or individual-based selection.

3 Algorithms and Test Functions

The algorithms we test in this paper are PAES, SPEA, PESA, and PESA-II. PAES is described in full in Knowles and Corne (2000), SPEA is described in Zitzler and Thiele (1999), and PESA is described in Corne et al (2000). PESA-II is a version of PESA which uses region-based selection. the parameter settings used are detailed in section 4.

Deb (1998) gives a procedure for designing tunable test functions for multiobjective optimisation. This technique enables the incorporation into objective space of a range of characteristics to varying degrees. These include discontinuity, concavity, non-uniformity of individuals along the Pareto front, and deception, each of which are considered by many to be the key characteristics which capable evolutionary multiobjective optimisers need to cope with.

In this paper we use six test functions designed using Deb's scheme. These are the functions $\mathcal{T}_1$-$\mathcal{T}_6$ which were used in a comparison of the performance of eight different MOEAs by Zitzler et al (1999), and in a comparison of three different algorithms by Corne et al (2000). The important characteristics of these func-

tions are as follows. $\mathcal{T}_1$ plays the role of a baseline, simple test; it has a convex Pareto front, and no characteristics which should lead to particular difficulty; $\mathcal{T}_2$ has a non-convex Pareto front – this causes difficulties, for example, for several techniques from the operations research and MCDM communities, which attempt to iteratively optimise weighted sums of the objectives for different sets of weights, since solutions in the concave region are not optima of any such scalarisation; $\mathcal{T}_3$ has many discontinuities in the Pareto front; $\mathcal{T}_4$ is highly multimodal and has 21^9 Pareto fronts; $\mathcal{T}_5$ is a deceptive problem, and $\mathcal{T}_6$ has a non-uniformly distributed search space with solutions non-uniformly distributed along the Pareto front.

Each is a two-objective problem defined on m parameters, in which both objectives are to be minimized. In five of the problems the parameters x_i were coded as a binary string decoded such that $x_i \in [0, 1]$. The remaining function ($\mathcal{T}_5$) also employed a binary chromosome but this time unitation was used to evaluate each of the parameters. We encode the functions here in precisely the same way as done in Zitzler et al (1999) and Corne et al (2000).

To briefly summarize Zitzler's study, SPEA seemed to be the best algorithm overall of the eight tested. Those compared included several of the classic methods such as the Niched Pareto Genetic Algorithm (Horn et al., 1994; Horn and Nafpliotis, 1994), the Non-Dominated Sorting method (Srinivas and Deb, 1994), and various versions of SPEA. Later, in Corne et al's study (2000), SPEA was compared with PESA and PAES. PESA was found to be best overall, although on $\mathcal{T}_5$, the deceptive problem, SPEA was slightly, but certainly, the best of the algorithms compared.

4   Experimental Design

4.1   Experiments

Our experiments sought to determine the relative quality of PESA, SPEA, PAES and PESA-II, a version of PESA which incorporates region-based selection, on the Deb test functions. Parameter settings are given in Table 1.

In the next section we summarise the statistical comparison method used to analyse the results within a set of experiments.

4.2   Statistics

Given the results of several trial runs for each algorithm, we compare the performance of two or more

Crossover rate	0.7 in PESA, PESA-II and SPEA; not used in PAES
Crossover method	uniform in PESA, PESA-II and SPEA; not used in PAES
Mutation rate	bit-flip rate set to $1/L$ where L is chromosome length
Populations	archive 100 in all algorithms, IP size 10 in PESA and SPEA
Chromosome lengths	900 in $\mathcal{T}_1$, $\mathcal{T}_2$ and $\mathcal{T}_3$, 300 in $\mathcal{T}_4$ and $\mathcal{T}_6$, 80 in $\mathcal{T}_5$
Hyper-grid size	32×32 grid in PESA, PESA-II and PAES, not used in SPEA

Table 1: Parameter settings

multiobjective optimisers using a method proposed originally by Fonseca and Fleming (1995a) which we have implemented with certain extensions. When comparing two multiobjective algorithms (A and B), this method returns two numbers: the percentage of the Pareto frontier on which A conclusively beats B (based on a Mann-Whitney U test at the 95% confidence level), and the percentage of the Pareto frontier on which B beats algorithm A. For example, two well-matched MOEAs might yield a result like [3.7, 4.1], indicating that each algorithm was definitely better than the other in small regions of the space, but they performed similarly well on the majority of the Pareto frontier. A clear indication that one algorithm is superior to another, however, is given by a comparison result such as [68.3, 2.2], or [100, 0.0].

In a comparison of more than two algorithms, the comparison code performs pairwise statistical comparisons, as before, for each distinct pair of algorithms. The results then show, for each algorithm, on what percentage of the discovered Pareto frontier we can be confident that it was unbeaten by any of the others, and on what percentage of the space it beat *all* of the others. For example, in Table 2, we can see that, on problem $\mathcal{T}_2$, PESA-II was unbeaten by any of the other algorithms individually on the entire Pareto surface, and conclusively superior to *all* of the others on 27.4% of the Pareto Tradeoff surface.

5   Results and Discussion

Table 2 summarises all results for the set of experiments in which each trial run was allowed 5,000 fitness evaluations. The best performing algorithm for each problem has its table entries highlighted in bold; when there is little difference between the best two (or all three), each such entry are highlighted in bold.

There are two rows for each problem; the first give the *unbeaten* statistic for each algorithm, and the second gives the *beats all* statistic. For example, on problem $\mathcal{T}_7$ PESA-II was unbeaten by any of the other algorithms on 100% of the combined Pareto frontier discovered over all trials, and on this frontier it was significantly superior to *all* of the others on 12.5% of it. In the case of $\mathcal{T}_7$ PESA-II was so much better than the other methods, we did additional trials to 20,000 evaluations to see if the other methods could 'catch up'.

Problem	PAES	SPEA	PESA	PESA-II
$\mathcal{T}_1$	66.1	1.1	**99.8**	**100**
	0	0	0	0
$\mathcal{T}_2$	0	0	72.3	**100**
	0	0	0	27.4
$\mathcal{T}_3$	65.4	22.3	78.4	**100**
	0	0	0	0
$\mathcal{T}_4$	64.4	**100**	**100**	**99.8**
	0	0.1	0	0
$\mathcal{T}_5$	0	**100**	98.6	**99.7**
	0	0	0	0
$\mathcal{T}_6$	16.7	74.5	18.8	**100**
	0	0	0	12.5
$\mathcal{T}_6$-long	2.8	0.5	1.1	**99.8**
	0.2	0	0	96.1

Table 2: Comparison of PAES, SPEA, PESA and PESA-II at 5,000 evaluations (plus an extra comparison at 20,000 evaluations for $\mathcal{T}_6$

As Table 2 shows, PESA was clearly the best method on three of the functions, and joint best with SPEA on a further two. On the one remaining function it achieved the second-best performance. SPEA is clearly best on just one function, and joint best with PESA on two. PAES is the worst performer here, being clearly worst on three of the test functions, and second or joint second best on the remaining three.

The results are summarised in Table 3, in which we show the rank for each algorithm on each problem. The rank is simply one plus the number of algorithms which clearly did better. For example, PAES has rank 3 for $\mathcal{T}_\infty$ since two algorithms (PESA and PESA-II) performed better than it on this function.

With reference to both tables 2 and 3, PESA-II clearly outperforms the other methods on the test functions examined overall. The performance on the $\mathcal{T}_7$ is especially marked.

We will now briefly consider the differential performance in terms of the $\mathcal{T}$ problem characteristics. $\mathcal{T}_\infty$ is a straightforward problem, and we find that both

Problem	PAES	SPEA	PESA	PESA-II
$\mathcal{T}_1$	3	4	1	1
$\mathcal{T}_2$	3	3	2	1
$\mathcal{T}_3$	3	4	2	1
$\mathcal{T}_4$	4	1	1	1
$\mathcal{T}_5$	4	1	1	1
$\mathcal{T}_6$	4	2	3	1
$\mathcal{T}_6$-long	2	4	3	1

Table 3: Summary of PAES, SPEA, PESA and PESA-II comparisons on Functions $\mathcal{T}_1$-$\mathcal{T}_6$

PESA and PESA-II perform excellently on it, with PAES doing well too, but SPEA doing rather badly. Since the problem lacks deception, and PAES is essentially a local search procedure, the good performance of PAES, especially in comparison to SPEA, is understandable. SPEA, as hinted at in Section 2, and unlike any of the other algorithms tested here, spends significant algorithmic effort in considering non-elitist solutions. This strategy seems to be unnecessary for $\mathcal{T}_\infty$, and seems to have prevented SPEA from performing well on it in the available time. In contrast, SPEA's non-elitism is likely to be responsible for it maintaining overall best performance on the deceptive problem, $\mathcal{T}_\bigtriangledown$, and also the highly multimodal problem, $\mathcal{T}_\bigtriangleup$. $\mathcal{T}_\in$ has a concave front; this makes it non-trivial for a local search based method and also seems to have confounded the strength pareto approach; PESA and PESA-II, especially the latter, perform very well on it. PESA-II also particularly shines on the remaining two problems, $\mathcal{T}_\ni$ and $\mathcal{T}_7$ which, respectively, have highly discontinuous and highly non-uniform Pareto fronts.

6   Conclusion

We have described region-based selection as an alternative selection scheme for use in evolutionary multiobjective optimisation. We have implemented it within PESA, although it may of course be employed in most evolutionary multiobjective frameworks. Experiments on functions from Deb's test suite seem to confirm that region-based selection is a very promising technique. PESA-II, which employed the new selection method, was only beaten (and then slightly) on two of the test problems. These were the highly multimodal problem, and the deceptive problem; hence, the relatively low profile for region-based selection in these cases can potentially be explained by the fact that SPEA is non-elitist (a highly helpful feature in problems with such characteristics), while region-based selection was implemented in an entirely elitist method

(PESA).

One promising avenue for further work would seem to be the deployment of region-based selection in a non-elitist framework. However this is not trivial; considering occupied hyperboxes in dominated regions of the space requires us to have a way of preferring, for example, a hyperbox on the Pareto frontier over a dominated hyperbox which has the same number of occupants. Another complicating factor is that a hyperbox may contain both dominated and nondominated individuals. We are thinking along the lines of using the Strength Pareto approach (Zitzler and thiele, 1999) to deal with these issues. A simple alternative might be to simply do region-based selection on the nondominated frontier most of the time, but for a portion of the time select from dominated individuals based on their strength Pareto fitnesses. A further alternative would be to only use region-based selection, but apportion algorithm effort between selecting in this way from different Pareto frontiers, akin to the nondominated sorting approach (Srinivas and Deb, 1994).

There are runtime complexity issues which we have not dealt with in this paper. For example, finding occupied hyperboxes can be done quickly, though is not trivial. Depending on the enclosing algorithm framework, the hyperbox and related datastructures may or may not need constant updating. Also, region-based selection (and any hypergrid method) requires the choice of a parameter to define the individual hyperbox dimensions. NSGA-II (Deb et al, 2000), for example, requires no such parameter. Preliminary investigations suggest that results are not overly sensitive to the hyperbox dimension parameter, although much more investigation needs to be done to determine if this is generally the case.

Acknowledgments

The authors are grateful to British Telecommunications Plc for financial support of the authors. We also acknowledge the helpful criticism of the anonymous reviewiers.

References

Corne, D.W., Knowles, J.D. and Oates, M.J. (2000). The Pareto-Envelope based Selection Algorithm for Multiobjective Optimisation, in Schoenauer, M., Deb, K., Rudolph, G., Yao, X., Lutton, E., Merelo, J.J. and Schwefel, H-P. (eds.), *Parallel Problem Solving from Nature - PPSN VI*, Springer Lecture Notes in Computer Science, pp. 869–878.

Czyzak, P. and Jaszkiewicz, A. (1998). Pareto simulated annealing - a metaheuristic technique for multiple-objective combinatorial optimization. *Journal of Multi-Criteria Decision Analysis 7*, 34–47.

Deb, K. (1998). Multi-Objective Genetic Algorithms: Problem Difficulties and Construction of Test Problems. Technical Report CI-49/98, Department of Computer Science, University of Dortmund.

Deb, K., Agrawal, S., Pratap, A. and Meyarivan, T. (2000). A Fast Elitist Non-dominated Sorting Genetic Algorithm for Multi-Objective Optimization: NSGA-II, in Schoenauer, M., Deb, K., Rudolph, G., Yao, X., Lutton, E., Merelo, J.J. and Schwefel, H-P. (eds.), *Parallel Problem Solving from Nature - PPSN VI*, Springer Lecture Notes in Computer Science, pp. 849–858.

Fonseca, C. M. and Fleming, P. J. (1995). An Overview of Evolutionary Algorithms in Multiobjective Optimization. *Evolutionary Computation 3*, 1–16.

Fonseca, C. M. and Fleming, P. J. (1995a). On the Performance Assessment and Comparison of Stochastic Multiobjective Optimizers. In Voigt, H-M., Ebeling, W., Rechenberg, I. and Schwefel, H-P., editors, *Parallel Problem Solving From Nature - PPSN IV*, pages 584–593, Springer.

Gandibleux, X., Mezdaoui N. and Freville, A. (1996). A tabu search procedure to solve multiobjective combinatorial optimization problems. In Caballero, R. and Steuer, R., editors, in *Proceedings volume of Multiple Objective Programming and Goal Programming '96*, Springer-Verlag.

Golberg, D.E., Deb, K., Korb, B. (1991) Don't Worry, Be Messy. in *Proceedings of the Fourth International Conference on Genetic Algorithms, pp. 24-30.*

Hansen, M. P. (1996). Tabu Search for Multiobjective Optimization : MOTS. Presented at MCDM '97, *Cape Town, South Africa, Jan 6-10.*

Hansen, M. P. (1997). Generating a Diversity of Good Solutions to a Practical Combinatorial Problem using Vectorized Simulated Annealing. Submitted to Control and Cybernetics, *August 1997.*

Horn, J., Nafpliotis, N., Goldberg, D.E. (1994). A niched Pareto genetic algorithm for multiobjective optimization, in Proceedings of the First IEEE Conference on Evolutionary Computation IEEE World Congress on Computational Intelligence, Volume 1, *pages 67-72. Piscataway, NJ: IEEE Service Centre.*

Horn, J. and Nafpliotis, N. (1994). Multiobjective Optimization Using The Niched Pareto Genetic Al-

gorithm. *IlliGAL Report 93005, Illinois Genetic Algorithms Laboratory, University of Illinois, Urbana, Champaign.*

Knowles, J. D. and Corne, D. W. (1999). *Local Search, Multiobjective Optimisation and the Pareto Archived Evolution Strategy, in* Proceedings of the Third Australia-Japan Joint Workshop on Intelligent and Evolutionary Systems, *pp. 209–216.*

Knowles, J. D. and Corne, D. W. (2000). *Approximating the Nondominated Front Using the Pareto Archived Evolution Strategy.* Evolutionary Computation, *8(2):149–172.*

Knowles, J.D. and Corne, D. (2000a). *M-PAES: A Memetic Algorithm for Multiobjective Optimisation, in* Proceedings of the 2000 Congress on Evolutionary Computation, *IEEE Neural Networks Council, pp. 325-332.*

Parks, G. T., Miller, I. (1998). *Selective Breeding in a Multiobjective Genetic Algorithm. In* Fifth International Conference on Parallel Problem Solving from Nature (PPSN-V), *pages 250–259. Springer.*

Srinivas, N., Deb, K. (1994). *Multiobjective optimization using nondominated sorting in genetic algorithms.* Evolutionary Computation, 2(3), *221–248.*

Van Veldhuizen, D.A. and Lamont, G.B. (2000). *Multiobjective Optimization with Messy Genetic Algorithms, in* Proceedings of the 2000 ACM Symposium on Applied Computing, *ACM Press, pages 470-476.*

Zitzler, E. and Thiele, L. (1999). *Multiobjective Evolutionary Algorithms: A Comparative Case Study and the Strength Pareto Approach.* IEEE Transactions on Evolutionary Computation, 2(4), *257–272.*

Zitzler, E., Deb, K. and Thiele, L. (1999) *Comparison of multiobjective evolutionary algorithms: empirical results. Technical report 70, Computer Engineering and Networks Laboratory (TIK), Swiss Federal Institute of Technology (ETH) Zurich.*