Two Heuristic Operations to Improve the Diversity of Two-objective Pareto Solutions

Rinku Dewri and Darrell Whitley
Computer Science, Colorado State University
Fort Collins, CO 80524
{rinku,whitley}@cs.colostate.edu

ABSTRACT

Multi-objective evolutionary algorithms using niching strategies fail to provide any information about the global density of the non-dominated front. This paper shows how the diversity of solutions can also be affected without adopting any explicit niching strategy, and by aiding the evolutionary algorithm with simple information pertaining to the global density of the non-dominated front. The concepts here are formulated for two objective problems and embedded into a real parameter elitist evolutionary algorithm.

Categories and Subject Descriptors

I.2.8 [**Problem Solving, Control Methods, and Search**]: Heuristic methods

General Terms

Algorithms, Experimentation

Keywords

Multi-objective optimization, Diversity

1. INTRODUCTION

Multi-objective Evolutionary Algorithms (MOEAs) have gained importance for their ability to navigate the entire decision space in a highly parallel fashion. This trait coupled with a competent convergence strategy can enable MOEAs to effectively determine multiple points on the Pareto front.

Under situations where no prior information about the Pareto-optimal front is available, converging to the true Pareto-optimal front and maintaining diversity among the solutions at the same time becomes important issues. Niching operators [1]-[4] are applied to help the MOEAs distinguish between densely and rarely neighbored solutions, or, clustering techniques are employed to shrink an overgrown population while preserving the diversity. Although these methods are able to substantially improve the performance of an evolutionary algorithm, the bottleneck lies in their computational complexity and the need to specify the parameters associated with these methods. Algorithms to simplify computationally intensive tasks like non-dominated sorting and population clustering, as well as self-adaptive mechanisms to fix niching parameters [5] have been proposed, but no focus was directed on obtaining an uniform distribution of the solutions along the Pareto-optimal front.

Most existing evolutionary algorithms rely on genetic operators to accomplish the task of generating new solution vectors. Fitness assignment and selection strategies are modified to direct the EA towards better solutions, but no explicit operators are incorporated to make the EA look for underexposed portions of the Pareto-optimal front. Intuitively, finding uniformly distributed solution vectors along the Pareto-optimal front becomes easier for an EA if the population members carry information about underdeveloped parts of the front.

This paper suggests the ideas of infusion and mirroring to address these issues in a two objective problem domain, and incorporates them into a real parameter elitist evolutionary algorithm: the Extended-MOEA (EMOEA). These operations are neither based on any probability measure nor manipulate randomly selected individuals from a population.

2. INFUSION AND MIRRORING

The non-uniformity of the Pareto front is hard to avoid even after integrating a diversity operator into the MOEA. Diversity operators helps the EA evolve solution vectors with a good spread along the Pareto-optimal front, but no control is possible over the distribution of these points. This is primarily because niching operators only help adjusting the local concentration of the population members and do not carry any information about the global density of the population. They mostly rely on the selection strategy to take care of this global density. Moreover, crossover and mutation operators are not always effective in quickly discovering unknown parts of the Pareto-optimal front.

The whole process of a MOEA can be viewed as the displacement of the current non-dominated front towards the Pareto-optimal one. EMOEA tries to uniformize this non-dominated front at every generation by inserting objective vector points at intermediate positions on the Pareto front. With reference to Figure 1, it would be better if the search could be directed more along ab, rather than fa. We try to achieve this by explicitly inserting more points along ab. However, since no direct mapping is available from the objective space to the decision space, finding solution vectors that correspond to these objective vector points is also not possible.

The current non-dominated front can be uniformized using a simple geometrical concept. The length of a curve $(A_1, A_2, \ldots, A_m)$ can be approximated by the line segments $A_1A_2, A_2A_3,$, and $A_{m-1}A_m$. A set of n uniform points can then be approximated along this curve, each at a euclidean distance of,

$$d_{avg} = \frac{\sum_{i=1}^{m-1} distance(A_i, A_{i+1})}{n-1} \quad (1)$$

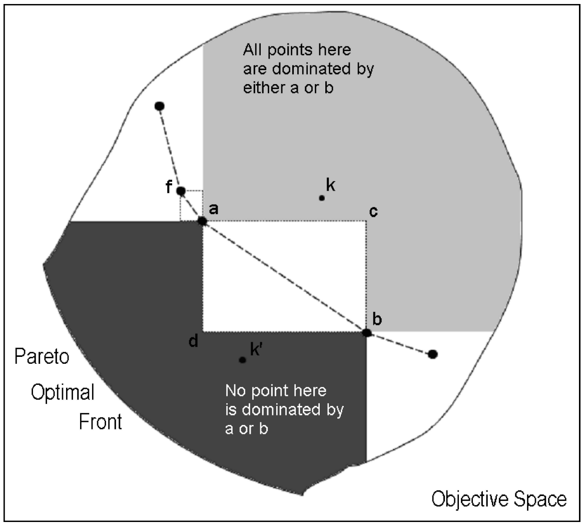

Figure 1: Infusion tries to fill in the region *acbd*; mirroring tries to convert dominated point k to non-dominated point $k\prime$.

from the previous one. In the context of EMOEA, the non-dominated population is sorted with respect to an objective function. The objective function chosen is varied iteratively to remove bias towards any particular function. Also, since the drive here is to uniformize the solution vectors in the objective space, all distance calculations are done in the same space.

2.1 Infusion

Line search methods [6]-[10] are well known in the classical optimization community for their ability to exploit the local convexities present in the search space. These methods generates iterates based on a step size and a search direction, obtained suitably to satisfy the *sufficient decrease* and the *curvature* conditions. The idea used for infusion is similar in nature, with differences arising in the non-derivative computation of the search direction. A heuristic approach is adopted to obtain this direction, and step sizes are chosen to facilitate the creation of new points in the non-dominated front.

Given two vectors F_A and F_B, the parametric equation,

$$F_Z(t) = F_A + t(F_B - F_A), \;\; 0 \leq t \leq 1 \tag{2}$$

represents the translation of the vector point F_A along the direction F_AF_B. In a multi-objective optimization scenario, if F_A and F_B are mutually non-dominated objective vectors, then no vector point in the set $V_{non-dom} = F_Z(t)$ is dominated by any other vector belonging to the set. This is an easy way of creating non-dominated objective vectors as long as we know the mapping from the objective space to the decision space. Otherwise, even though we are able to find more non-dominated objective vectors, the corresponding solution vectors will still be unknown. However, this mapping is not always evident.

The heuristic method suggested here is to extract the direction information from the objective space and apply this in the decision space. The non-dominated solution vectors are sorted with respect to one objective function and then as we move from one solution vector to the successive one, the corresponding objective vector are also assumed to move from one to the next. The movement of the solution vectors can be represented by an equation similar to (2),

$$X_{jk}(t) = X_{ik} + t(X_{(i+1)k} - X_{ik}), \;\; 0 < t < 1; \quad i = 1, \cdots, \eta - 1; \quad k = 1, 2 \tag{3}$$

where X_{ik} is the solution vector corresponding to the objective vector F_{ik}, all vectors being sorted with respect to the k^{th} objective function. η is the size of the non-dominated population of the current generation. Since all sorting is done in objective space, while transformations are applied in the decision space, direction information gets transferred from one to the other. The parameter t can be varied depending on how many intermediate points are required for uniformization. Once these solution vectors are obtained, they can be mapped directly to the objective space by using the specified objective functions. We refer to this process as *infusion*.

2.2 Mirroring

One easily perceivable flaw in the method of infusion is the non-availability of a dominance check on the newly generated vectors. Since all transformations are applied in the decision space, there is no assurance that the resulting objective vector points will also be non-dominated. Figure 1 sketches the possible areas where new objective vectors may be formed during infusion in a two dimensional space. a and b are two non-dominated objective vector points between whom infusion is carried out. Any point inside the region $acbd$ is not dominated by either a or b. The shaded region to the right of the line segment ab constitutes all points in the space that are dominated by either a or b. Noticeably, the bi-lateral reflection of this region about the line segment ab is an area which is greatly desired. The points in this region dominate either one of a or b, or both.

It is clear that transforming a dominated point to a non-dominated one requires the flipping of the objective function values about that of points a and b. If $a = (a_1, a_2)$ and $b = (b_1, b_2)$, then a point $k = (k_1, k_2)$ in the dominated region can be converted to a non-dominated point $k' = (k'_1, k'_2)$ using,

$$k'_i = a_i + b_i - k_i, \;\; i = 1, 2 \tag{4}$$

Once again, these mirroring transformations are applied to the solution vectors in the variable space. A solution vector can be mirrored with respect to the variable bounds,

$$X_{j'k} = X_{lower} + X_{upper} - X_{jk} \tag{5}$$

or the variable values corresponding to the two objective vector points subjected to infusion,

$$X_{j'k} = X_{ik} + X_{(i+1)k} - X_{jk} \tag{6}$$

Here X_{lower} and X_{upper} are the lower and upper bound vectors for the variables in the problem. The success of a mirroring transformation depends on the nature of interaction present between the decision space and the objective space. In case of a success, the mirroring operation manages to push the non-dominated curve towards the Pareto-optimal front, thus improving the convergence considerably. Mirroring is applied only when infused points fail to satisfy the dominance criteria, but it can as well be an effective transformation, in terms of convergence, even when the points are within the non-dominated region *abc* (in Figure 1).

Mirroring is analogous to the *reflection* step of Nelder-Mead's simplex algorithm [11], used mostly for nonlinear unconstrained optimization. Nelder-Mead generates a new test position by extrapolating the behavior of the objective function measured at each test point arranged as a simplex. Similar to Nelder-Mead's algorithm, mirroring replaces the worst point in the simplex with a point reflected through the remaining points considered as a plane, provided an improvement is obtained by the operation.

2.3 Extended-MOEA

The algorithm starts by creating a random population P_0 of size N, which is referred to as the general population. An elite population E_g is also maintained, where $E_0 = \phi$. The non-dominance check,

$$E_g = \Gamma(E_g \cup \Gamma(P_g)) \quad (7)$$

ensures that only the non-dominated solutions are maintained in the elite population. Here, $\Gamma(P)$ represents the non-dominated solutions in the population P, and g the generation counter starting at zero. The elite population is then sorted with respect to an objective function. An elite population size η is taken as user input and the sorted elite population is subjected to uniformization. Fitness is assigned to members from both the populations and recombination operators are applied to members selected from a union of the two populations. Child members replace the old general population.

The uniformization process first calculates the average distance d_{avg} between the points. We then move along the front, from one population member to the next, until a distance of d_{avg} is covered. If a population member is found at that location, it is included; otherwise infusion is carried out to fill in the distance between the two adjacent population members. Upon failure of the infused point to be a non-dominated point, mirroring is carried out on it; the outcome is discarded if no improvement is obtained. The extreme points are always included in the uniformized population.

Once the task of uniformization is complete, the elite set E_g contains approximately η members. A general population member is assigned a fitness value equal to its minimum distance from the elite set. For the elite population, fitness is assigned based on its distance from the two neighboring population members. The neighborhood is defined in terms of the value of one of the objective function. The extreme points of E_g are assigned a fitness value less than all the others (ideally $-\infty$). Fitness comparisons assume we are minimizing. All elite member fitness values are negated to give them more preference than the general population members.

3. SIMULATION RESULTS

A set of six test problems is taken from past studies. Two performance metrics introduced by Deb in [12] are used to quantitatively estimate the convergence and diversity measures. It should be noted that better values on such performance measures can at best indicate that an algorithm is not worse than another, but does not help infer the superiority of the algorithm [13]. Simulation results are compared with that of NSGA-II [12] and SPEA [14] using these two metrics. These methods are chosen for their relative simplicity and effectiveness in obtaining a diverse set of solutions in the Pareto optimal front. The effect of infusion and mirroring on the performance of EMOEA are also reported.

The EMOEA simulation for these problems is run for 250 generations with a population size of 80. The elite population size, η, is set at 20. Probability of crossover is set at $p_{cross} = 0.9$ and that of mutation at $p_{mut} = \frac{1}{n}$, where n is the number of variables in the problem. The results in Table 1 and Table 2 are averaged over 20 runs for each problem.

3.1 Comparative results

Out of the six test problems used, FON in [15], KUR in [16] and POL in [17] are two objective minimization problems. Zitzler, Deb and Thiele suggested a systematic procedure to construct two objective minimization problems in [18], out of which we chose ZDT3, ZDT4 and ZDT6 for our study. These problems pose difficulties to an MOEA in terms of diversity maintenance, convergence and irregular density mappings between the objective space and the decision variable space.

Table 1: Convergence metric values on test problems. Values shown as mean over variance.

Problem	NSGA-II	SPEA	EMOEA
FON	$\mathbf{\frac{0.001931}{0.000000}}$	$\frac{0.010611}{0.000005}$	$\frac{0.004051}{0.000000}$
KUR	$\mathbf{\frac{0.028964}{0.000018}}$	$\frac{0.049077}{0.000081}$	$\frac{0.076553}{0.001450}$
POL	$\mathbf{\frac{0.015553}{0.000001}}$	$\frac{0.054531}{0.000179}$	$\frac{0.075542}{0.000779}$
ZDT3	$\frac{0.114500}{0.007940}$	$\frac{0.044212}{0.000019}$	$\mathbf{\frac{0.005317}{0.000015}}$
ZDT4	$\frac{0.513053}{0.118460}$	$\frac{9.513615}{11.321067}$	$\mathbf{\frac{0.001002}{0.000000}}$
ZDT6	$\frac{0.296564}{0.013135}$	$\frac{0.020166}{0.000923}$	$\mathbf{\frac{0.019356}{0.000082}}$

Table 1 shows the convergence metric values for the six two objective problems, obtained using NSGA-II, SPEA and EMOEA. The diversity metric for the same is listed in Table 2. EMOEA performs well in almost all the six problems, although NSGA-II performs better on KUR and POL. For the ZDT series of problems, EMOEA's performance is much better than both NSGA-II and SPEA. We show typical simulation runs of the algorithm on a few problems (Figure 2-7).

The Pareto-optimal set in KUR is non-convex and has three distinct disconnected regions. These regions are constituted with a disconnected set of decision variable vectors. Such a discontinuity in the mapping between the decision space and objective space can create difficulties for an MOEA in finding Pareto-optimal solutions in all the regions. Although EMOEA's convergence is not poor in this problem,

Table 2: Diversity metric values on test problems. Values shown as mean over variance.

Problem	NSGA-II	SPEA	EMOEA
FON	0.378065 / 0.000639	0.804113 / 0.002961	**0.054982 / 0.001239**
KUR	**0.411477 / 0.000992**	0.880424 / 0.009066	0.418119 / 0.005973
POL	0.452150 / 0.002862	0.954327 / 0.013170	**0.361416 / 0.002507**
ZDT3	0.738540 / 0.019706	0.732097 / 0.011284	**0.575132 / 0.002975**
ZDT4	0.702612 / 0.064648	0.732097 / 0.011284	**0.081626 / 0.008399**
ZDT6	0.668025 / 0.009923	0.900793 / 0.004124	**0.242728 / 0.013266**

both NSGA-II and SPEA perform better than EMOEA. Figure 2 shows the non-dominated solutions obtained by SPEA. Given the constraint of finding only 20 points along the Pareto-optimal region, EMOEA performs better in finding a diverse set of solution vectors (Figure 3) in all the regions of the Pareto-optimal front. EMOEA's ability to find more uniformly distributed non-dominated vectors is better demonstrated when we ask for a higher number of points along the Pareto-optimal front. Figure 4 shows the non-dominated solutions obtained by EMOEA with the same general population size and η=80.

Problem ZDT4 has 21^9 local Pareto-optimal solutions, each corresponding to $0 \leq x_1 \leq 1$ and $x_i = 0.5m$, where m is any integer in [-10,10] and $i = 2, \cdots, 10$. The global Pareto-optimal front corresponds to $m = 0$. The presence of multiple local Pareto-optimal fronts creates many obstacles for an MOEA. This is verified by the convergence metric values of both NSGA-II and SPEA. EMOEA's performance is much better in this problem, and the solution vectors obtained are almost on the global Pareto-optimal front (Figure 5). The diversity of the solutions obtained is near uniform and better than both NSGA-II and SPEA. This is achieved by the mirroring operation which locally forces the local non-dominated front to move towards the global optima.

Figure 6 shows the non-dominated solutions in ZDT6 obtained by NSGA-II and SPEA. ZDT6 has a non-convex Pareto-optimal set with a variable density of solutions across the Pareto-optimal region. Such adverse density mappings coupled with the non-convex nature of the Pareto-optimal front can also cause difficulties to an MOEA. NSGA-II and SPEAs performance suffers in terms of convergence and diversity. However, EMOEA is able to maintain an almost uniform set of non-dominated solution points along the true Pareto-optimal front (Figure 7).

Noticeably, the diversity metric values for EMOEA are much better than both NSGA-II and SPEA in almost all the problems. The diversity can be improved by taking a

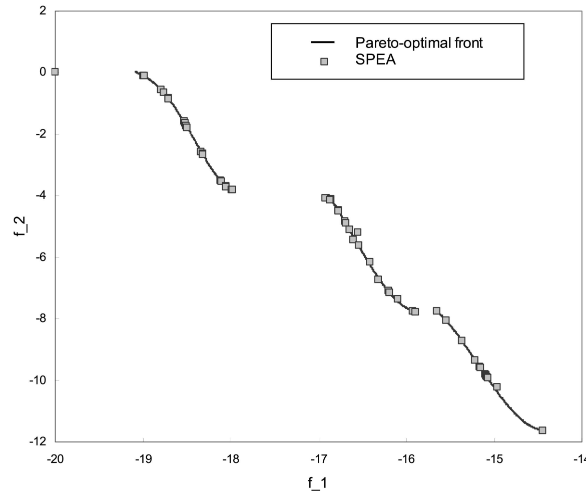

Figure 2: Non-dominated solutions with SPEA on KUR.

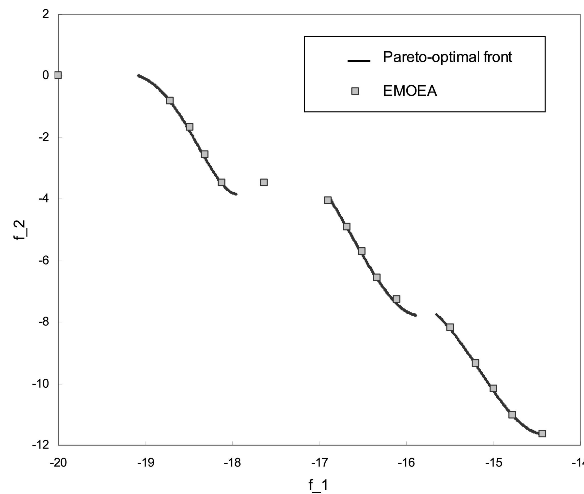

Figure 3: Non-dominated solutions with EMOEA on KUR.

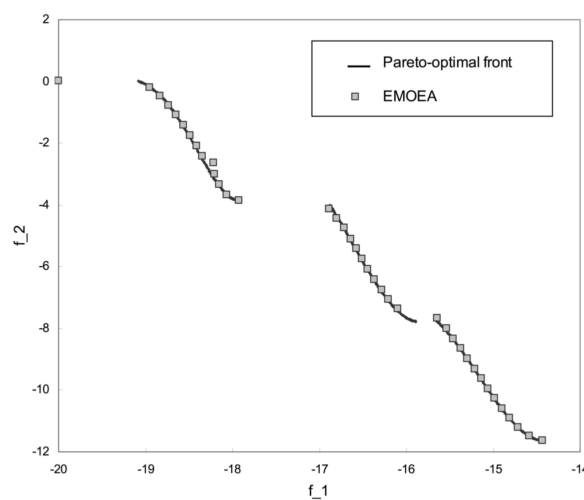

Figure 4: Non-dominated solutions with EMOEA on KUR (η=80).

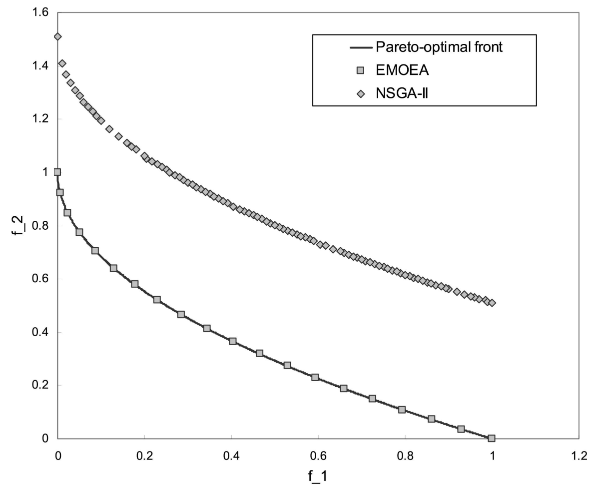

Figure 5: Non-dominated solutions with NSGA-II and EMOEA on ZDT4.

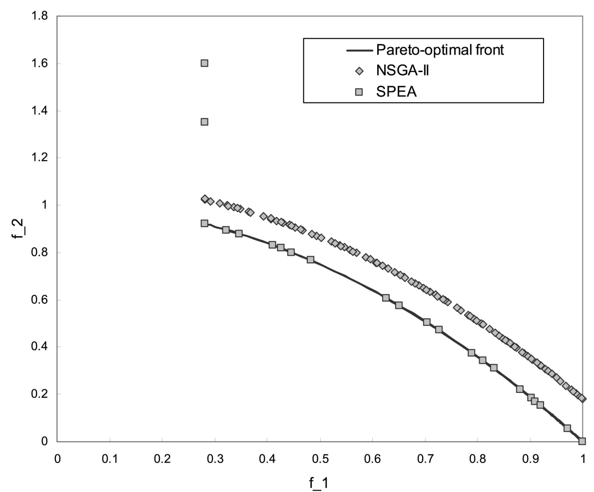

Figure 6: Non-dominated solutions with NSGA-II and SPEA on ZDT6.

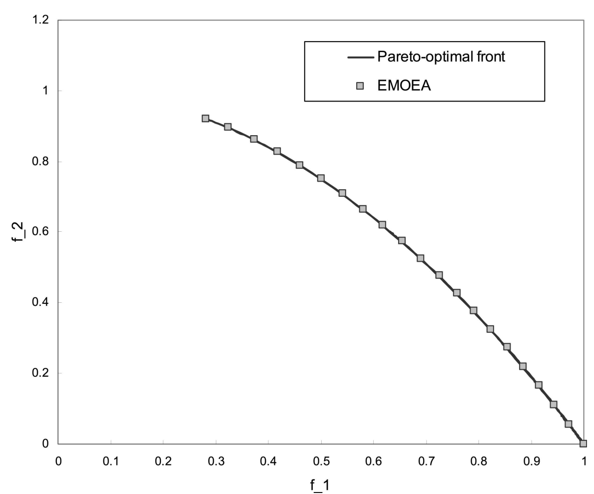

Figure 7: Non-dominated solutions with EMOEA on ZDT6.

bigger elite population. Also, it may be noted that ZDT3 and ZDT4 are problems where one of the objective functions is directly mapped to one of the decision variables, and so they are somewhat biased towards getting a good spread of solutions. However, we have also considered other problems where such linearity is not present and the simulation results corroborates that the better diversity obtained is not solely due to this factor.

3.2 Effect of Infusion

In order to study the effect of infusion, separate EMOEA runs were taken with and without the uniformization process. Mirroring is disabled to see if infusion alone can take care of the convergence requirement. Single simulation runs are shown for ZDT4 and ZDT3 to help visualize the effects.

As seen in Figure 8, EMOEA (without uniformization) has found a number of non-dominated points for ZDT4 but these points are mostly accumulated in a section of the non-dominated front. The diversity obtained after enabling uniformization is quite promising. Apart from giving a better spread, uniformization also helped attain further convergence of the non-dominated points. However, infusion alone was insufficient to find the Pareto-optimal front for ZDT4. ZDT4 has multiple local Pareto-optimal fronts and whenever the non-dominated solutions lie on one of these local fronts, infusion, by definition, is not able to create new values for the variables $x_i; i = 2, \cdots, 10$ and hence the new points generated are points on the old front itself. Mirroring, however, helped avoid such a situation and the non-dominated curve easily converged to the Pareto-optimal front.

Uniformization also helped discover the five discontinuous regions of ZDT3. Even though the convergence at the end of 50 generations (Figure 9) is poorer in this case, a wider stretch of the front has been discovered in a relatively shorter number of generations. Figure 10, which shows the simulation results after 250 generations, corroborates the slow rate of progress of the non-dominated curve if not supplemented by the uniformization process. On the other hand, EMOEA (with uniformization) has converged to a set of optimal solutions well dispersed along all the five regions of the Pareto-optimal front.

3.3 Effect of Mirroring

The effect of mirroring on the performance of EMOEA is well evident from the simulations of ZDT4. Figure 8 clearly shows that mirroring considerably helped the convergence of ZDT4. ZDT4 is a problem having a number of parallel local optima; any success in a mirroring operation creates a dent in the non-dominated front and locally forces the front to move towards optimality. The mirroring operation could be successful in only a fraction of all instances, but can result in much better convergence.

3.4 Computational complexity

EMOEA starts by finding the best non-domination front which requires at most $O(NlogN)$ time [19]. The non-dominated solutions thus obtained are compared for dominance with the existing elite set - an $O(N\eta)$ task. The sorting operation performed on the elite population can be done in $O(\eta log\eta)$ computations. The infusion process iterates over the elite population requiring $O(\eta)$ operations. Fitness assignment can be accomplished in $O(N\eta)$ time for the general population and in $O(\eta^2)$ time for the elite pop-

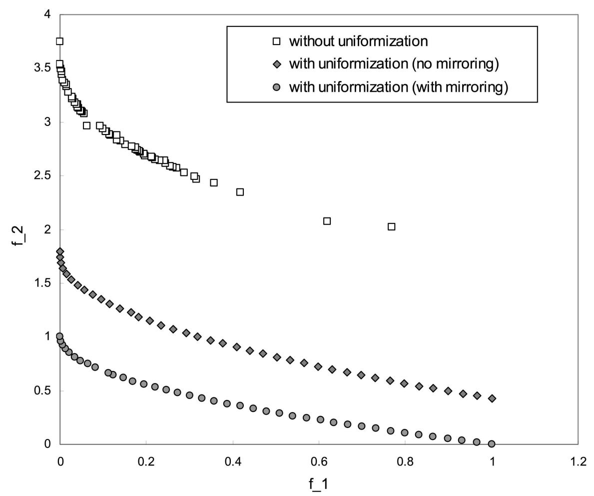

Figure 8: Non-dominated front for ZDT4 obtained after 250 generations with population size 80 and η=40.

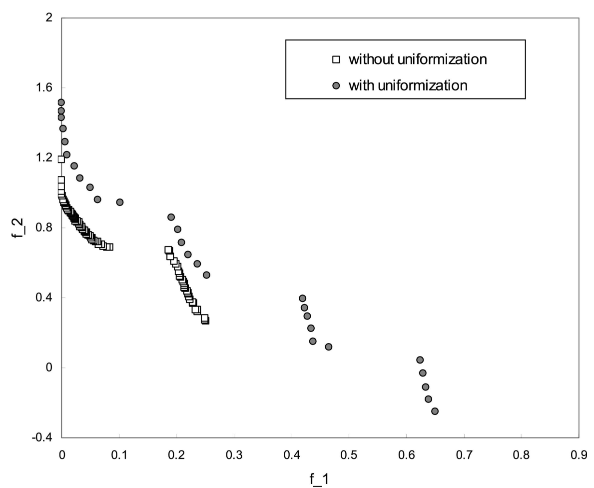

Figure 9: Non-dominated front for ZDT3 obtained after 50 generations with population size 80 and η=40.

ulation. All other steps in the algorithm do not have computational complexities more than the above computations. Assuming that N and η are of the same order, the overall complexity of a single generation of EMOEA is no less than $O(N^2)$. This complexity compares equivalently to the clustering algorithm of SPEA, but is worse than the crowding distance assignment procedure of NSGA-II ($O(NlogN)$).

In contrast to NSGA-II and SPEA, EMOEA actively generates solutions in the least populated areas of the front. This leads to an $O(\eta)$ more function evaluations as an expense to obtaining better diversity. These excess number of function evaluations would be an overhead as the non-dominated set approaches the Pareto optimal front. The mirroring step should be avoided then as it would only generate infeasible solutions.

4. EMOEA ISSUES

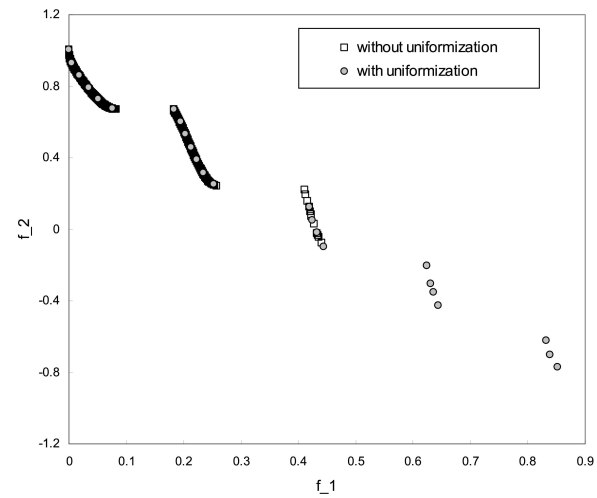

Figure 10: Non-dominated front for ZDT3 obtained after 250 generations with population size 80 and η=40.

EMOEA depends on a regular, or well-ordered, mapping between the objective space and the decision variable space. Transformations applied during the infusion and mirroring operations assume that any sort of ordering of the points in objective space also regulates the decision variable points in accordance to some order. The absence of this mapping may cause EMOEA to not be able to produce a good distribution of solutions.

Problems where the density of the solutions across the Pareto-optimal front is not uniform may pose similar difficulties for the algorithm. As the number of generational passes increase, the adverse density regions gets divided into regularly mapped intervals. When dealing with such problems, a number of generations is required by the algorithm to identify these regions of regular mapping before stabilizing on a uniform distribution of solutions. The number of generational runs should thus be sufficiently large to ensure that the algorithm has effectively dealt with such situations.

EMOEA does not use a sharing parameter to determine the crowdedness of a solution vector. It is always assumed that the uniformization process will take care to see that the solution vectors are evenly distributed across the non-dominated front. However, the credibility of this assumption is highly dependent on the success of the infusion and mirroring operations performed, and no assurance can be given as to what fraction of these operations will be successful. In order to compensate for this, the fitness assignment scheme in EMOEA incorporates the requisite measures to give preference to isolated or rarely populated members of the non-dominated set. The fitness of the elite members being based on its distance from the neighbors, population members spaced far apart from nearby solution vectors gain a higher selection pressure in an evolutionary run.

Beyer and Deb argued in [20] that the selection operator reduces the population diversity because it creates duplicates of a few population members and eliminates a few others. Crossover and mutation operators are employed to enhance this population diversity. An adaptive search using an evolutionary algorithm is viable because of this characteristic balance between exploitation and exploration of the population.

However, these operators are generally probabilistic in nature. Apart from being applied to randomly selected population members, the applicability of these operators is also probabilistically determined. They boost the search power by introducing heterogeneity with respect to the randomly selected parent members. Such a random selection process is of significance in discovering new solution vectors, but not necessarily solutions that maintain diversity with the ones generated during a previous instance of the operation. Here we present infusion and mirroring as aiding operations to genetic operators which follow some heuristics to transform a dominated solution into a non-dominated one. These operations are not applied to random members of the population. Instead, the entire set of population members to which they are to be applied is prearranged in some order to suit the objectives of the algorithm, better diversity in this case. In EMOEA, the uniformization process decides the relative proportion of heterogeneity between two parent vectors and this knowledge is then utilized by the infusion operation to fill in the gap between the two vectors. Mirroring tries to rectify any faults that might have occurred during infusion. It has to be noted that these operators are designed to aid the usual evolutionary framework and do not act as a replacement for the crossover and mutation operators.

5. CONCLUSIONS

We presented a method that can positively affect the convergence and diversity of solutions obtained in a two objective problem domain. Infusion and mirroring operations enforce some form of uniformity in the non-dominated population before the usual genetic operators act on it. The results indicate that such operations can help an evolutionary algorithm harness its search power better; in addition, the need for explicit niching strategies has been avoided. The method suggested here cannot be directly applied to higher dimensional problems but brings forth the question whether similar operations could be devised for such problems.

6. REFERENCES

[1] D. J. Cavicchio. *Adaptive search using Simulated Evolution*. Doctoral dissertation, University of Michigan, 1970.

[2] K. A. DeJong. *An analysis of the behavior of a class of Genetic Adaptive Systems*. Doctoral dissertation, University of Michigan, 1975.

[3] S. W. Mahfoud. *Niching Methods for Genetic Algorithm*. Doctoral dissertation, University of Illinois at Urbana-Champaign, 1992.

[4] D. E. Goldberg and J. Richardson. Genetic algorithms with sharing for multimodal function optimization. In *First International Conference on Genetic Algorithms and Their Applications*, pages 41–49, 1987.

[5] K. Deb and D. E. Goldberg. An investigation of niche and species formation in genetic function optimization. In *Third International Conference on Genetic Algorithms*, pages 42–50, 1989.

[6] J. E. Dennis and R. E. Sghnabel. *Numerical Methods for Unconstrained Optimization And Non-linear Equations*. Prentice Hall, Englewood Cliffs, N.J., 1983.

[7] R. Fletcher. *Practical Methods of Optimization*. 2nd ed. Wiley, New York, 1987.

[8] R. H. Byrd, J. Nocedal and Y. Yun. Global convergence of a class of quassi-Newton methods on convex problems. *SIAM Journal of Numerical Analysis*, 24:1171–1190, 1987.

[9] M. Al-Baali. Descent property and global convergence of the Fletcher-Reeves method with inexact line searches. *IMA Journal of Numerical Analysis*, 5:121–124, 1985.

[10] J. C. Gilber and J. Nocedal. Global convergence properties of conjugate gradient. *SIAM Journal of Optimization*, 2:21–42, 1992.

[11] J. A. Nelder and R. Mead. A simple method for function minimization. *Computer Journal*, 7:308–313, 1965.

[12] K. Deb, S. Agrawal, A. Pratap, and T. Meyarivan. A fast and elitist multi-objective genetic algorithm: NSGA II. *IEEE Transactions on Evolutionary Computation*, 6(2):182–197, 2002.

[13] E. Zitzler, L. Thiele, M. Laumanns, C. M. Fonseca and V. G. da Fonseca. Performance Assessment of Multiobjective Optimizers: An Analysis and Review. *IEEE Transactions on Evolutionary Computation*, 7(2):117–132, 2003.

[14] E. Zitzler and L. Thiele. *An evolutionary algorithm for multiobjective optimization: The strength Pareto approach*. Technical Report 43, Computer Engineering and Networks Laboratory (TIK), Swiss Federal Institute of Technology, 1998.

[15] C. M. Fonseca and P. J. Fleming. Multiobjective optimization and multiple constraint handling with evolutionary algorithms - Part II: Application example. *IEEE Transactions on Systems, Man, and Cybernetics: Part A: Systems and Humans*, pages 38–47, 1998.

[16] F. Kursawe. A variant of evolution strategies for vector optimization. *Parallel Problem Solving from Nature*, I:193–197, 1990.

[17] C. Poloni. Hybrid GA for multi-objective aerodynamic shape optimization. In *Genetic algorithms in engineering and computer science*, pages 397–414, 1997.

[18] E. Zitzler, K. Deb, and L. Thiele. Comparison of multiobjective evolutionary algorithms: Empirical results. *Evolutionary Computation Journal*, 8(2):125–148, 2000.

[19] H. Kung, F. Luccio, and F. Preparata. On finding the maxima of a set of vectors. *Journal of the ACM*, 22:469–476, 1975.

[20] H. G. Beyer and K. Deb. On the desired behavior of self-adaptive evolutionary algorithms. *Parallel Problem Solving from Nature VI*, (PPSN - VI):59–68, 2000.